Bayesian Diversification Rate Analysis with TESS

Sebastian Höhna, Michael R. May and Brian R. Moore

June 17, 2015

This tutorial covers various statistical approaches for estimating rates of lineage diversification (speciation – extinction) from phylogentic trees using the R package TESS. TESS provides a flexible Bayesian framework for specifying an effectively infinite array of diversification models—where diversification rates are constant, vary continuously, or change episodically through time—and implements numerical methods to estimate parameters of these models from molecular phylogenies. We provide robust Bayesian methods for assessing the *relative* fit of these models of lineage diversification to a given study tree—*e.g.,* where stepping-stone simulation is used to estimate the marginal likelihoods of competing models, which can then be compared using Bayes factors. We also provide Bayesian methods for evaluating the *absolute* fit of these branching-process models to a given study tree—*i.e.,* where posterior-predictive simulation is used to assess the ability of a candidate model to generate the observed phylogenetic data.

Contents

1 Getting Started

We assume that the reader has some experience using R and has installed the TESS package (including all dependent packages, such as ape). We also assume some familiarity with Bayesian inference and models of lineage diversification. Nevertheless, we intend this guide to be relatively self-contained: we provide brief explanations of the methods and models in the corresponding tutorials, and direct the reader to the relevant primary literature for more detailed descriptions of the corresponding topics.

We originally developed TESS as a tool for efficiently simulating phylogenies in order to test and validate new inference methods and models (Höhna, 2013). However, TESS has since evolved to include several methods for estimating diversification rates from empirical phylogenies (*e.g.,* Höhna, 2014; May et al., 2015). This is a natural extension, as both simulation and inference methods are based on the same equations and inference machinery.

1.1 Scope of research questions

There are three fundamental questions that can be addressed using TESS:

1. *What are the rates of the process that gave rise to my study tree?*
2. *Have diversification rates changed through time in my study tree?*
3. *Is there evidence that my study tree experienced mass extinction?*

Questions regarding diversification rates can be addressed using TESS simply by estimating the parameters of the branching-process model—*i.e.,* rates of speciation (λ), extinction (μ), net-diversification ($\lambda - \mu$), and relative-extinction ($\mu \div \lambda$). We estimate these parameters in a Bayesian statistical framework, which provides a natural means to accommodate our uncertainty in estimates of the parameters—*i.e.,* rather than estimating rate parameters as point values, TESS provides estimates as marginal posterior probability densities. We describe the branching-process models implemented in TESS—and the methods for estimating parameters of these models—in Section 2 of this guide.

Questions regarding temporal variation in diversification rates can be addressed using TESS by comparing the relative fit of the study tree to candidate branching-process models—*i.e.,* by performing Bayes factor comparisons to assess the relative support for models in which diversification rates are either constant or change through time. Note that the models we have implemented in TESS assume that diversification rates are homogeneous across lineages. Accordingly, even though diversification rates may change—gradually or episodically—*through time*, diversification rates are nevertheless identical across all lineages at any *instant in time.*

We describe how to use TESS to compare the fit of candidate diversification models to a given dataset in Section 4 of this guide.

Questions regarding mass-extinction events can be inferred using TESS by performing specific hypothesis tests (see Section 2.4.4) or analyses under the CPP on Mass-Extinction Times (`CoMET`) model (May et al., 2015). These analyses can identify whether your study tree has been impacted by mass extinction, and if so, can identify the number and timing of these events. Note that the `CoMET` model is comprised of three compound Poisson process (CPP) models that describe three corresponding types of events: (1) instantaneous tree-wide shifts in speciation rate; (2) instantaneous tree-wide shifts in extinction rate, and; (3) instantaneous tree-wide mass-extinction events. *In principle*, the `CoMET` model could therefore be used to explore events other than mass extinction—such as the number of tree-wide diversification-rate shifts, the timing of those events, and the rate parameters (*e.g.,* speciation and extinction rates) associated with those events. *In practice*, however, the diversification-rate components of the `CoMET` model are included as nuisance parameters that improve estimation of the focal parameters—*i.e.,* those associated with mass-extinction events. The ability to detect tree-wide diversification-rate shifts is currently being explored, but we caution users against overinterpretation of these nuisance parameters. We describe how to use TESS to explore mass-extinction events in Section 4.3 of this guide.

1.2 Loading empirical data

Rates of lineage diversification are typically estimated from phylogenies that, in turn, have been inferred from molecular sequence data. For example, consider the conifer phylogeny that is included with the TESS distribution:

```
library(TESS)      # load the package
data(conifers)     # load the conifers dataset
```

More information on this phylogeny can be found in Leslie et al. (2012). You will, of course, want to use your own tree for your diversification-rate analyses. You can do this using the `read.nexus` function provided in the `ape` package:

```
myTree <- read.nexus("data/myTree.nex")
```

You can extract the node ages from the tree using the `ape` function `branching.times`. We often use the node ages for estimating parameters of birth-death processes, so we'll extract them and store them in a variable for later use.

```
times <- as.numeric( branching.times(conifers) )
```

You then can view the phylogeny (Figure 1).

```
plot(conifers,show.tip.label=FALSE,no.margin=TRUE)
```

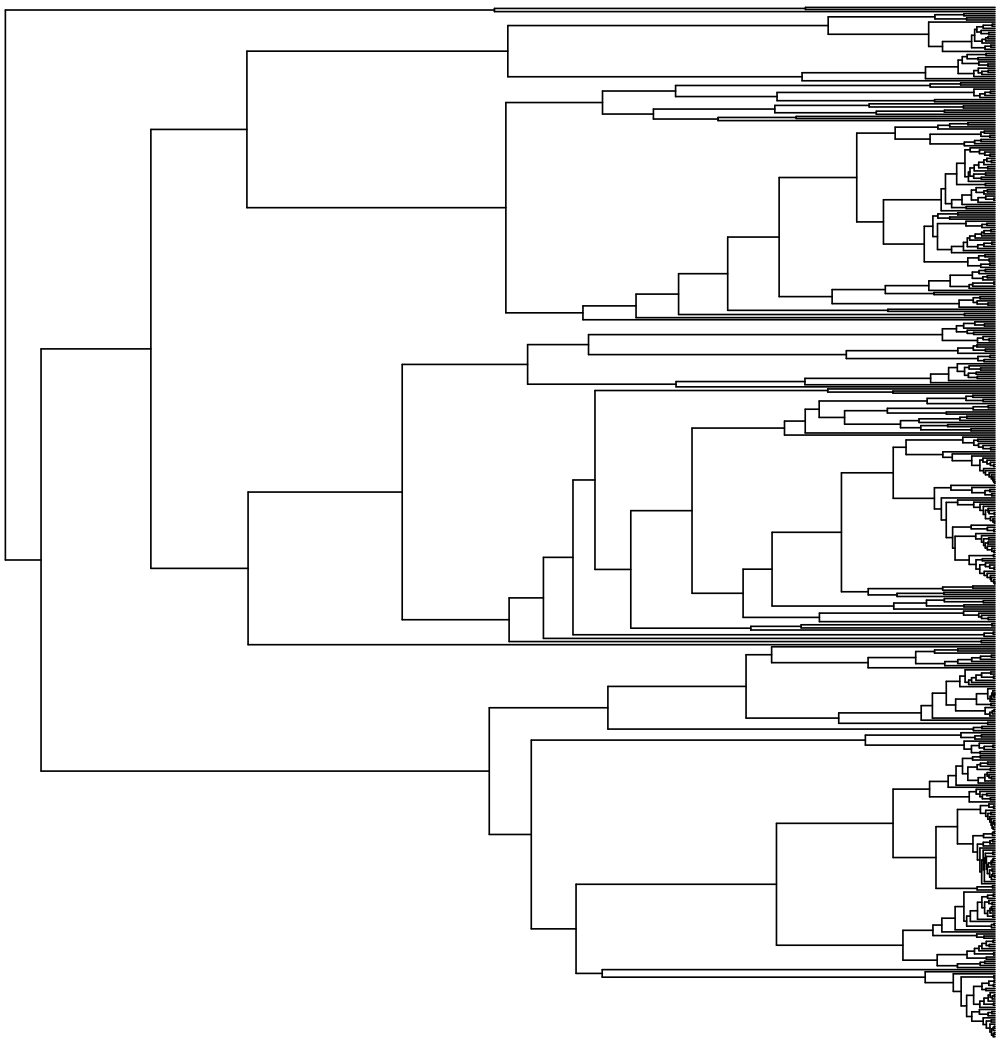

Figure 1: Conifer phylogeny from Leslie et al. (2012) without taxon labels.

Notice that this is an *ultrametric* tree; that is, it is rooted and all of the tips are sampled at the same time horizon (*i.e.*, the present). The models implemented in

TESS are only valid for ultrametric trees. Other trees—*e.g.,* where tips are sampled sequentially through time (Heath et al., 2014)—are currently not supported.

Additionally, you can look at the lineage-through-time (LTT) plot (Figure 2).

```
ltt.plot(conifers,log="y")
```

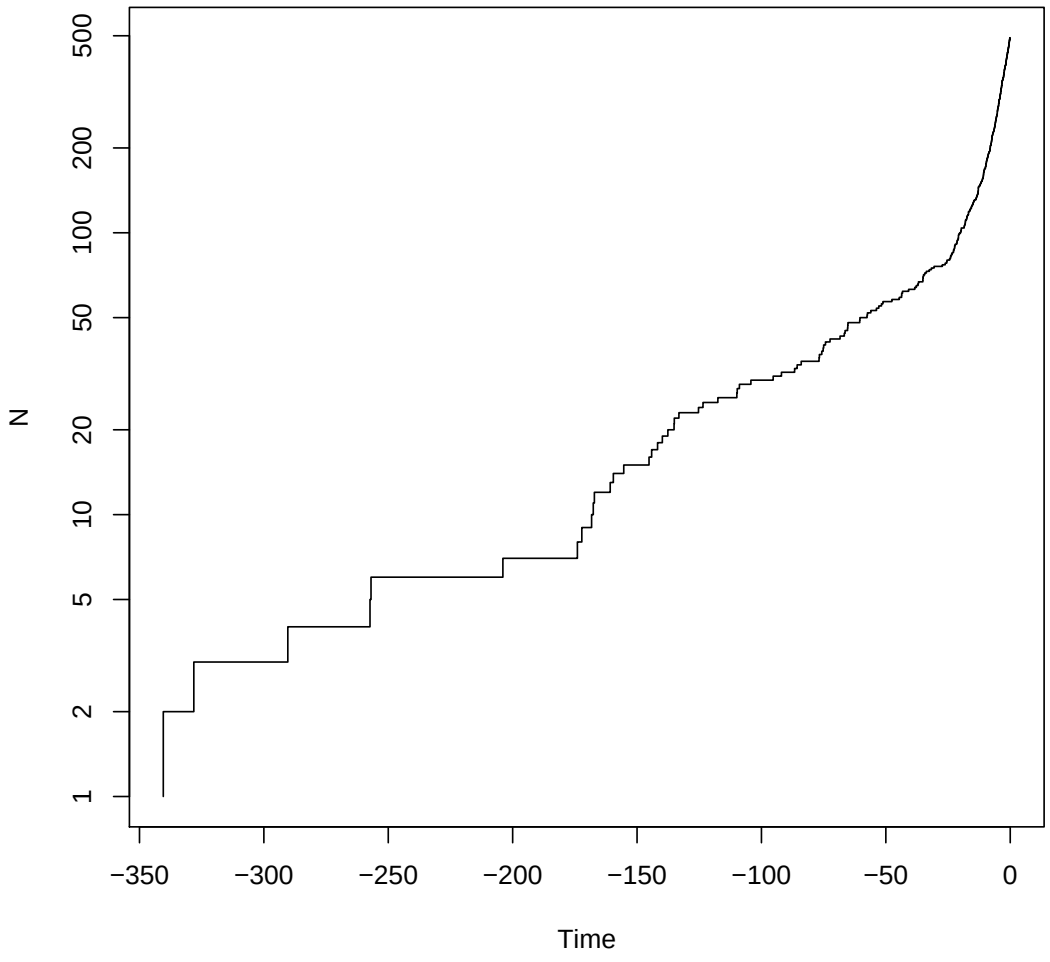

Figure 2: Lineage-through-time plot of the conifer phylogeny.

The LTT plot allows us to visualize the phylogenetic information that is used

for estimating diversification rates. For example, we it appears that the slope of the LTT plot changes slightly at $\approx 175, 70$, and 20 million years ago.

2 Models

We begin this section with a general introduction to the stochastic birth-death branching process that underlies inference of diversification rates in TESS. This primer will provide some details on the relevant theory of stochastic-branching process models. We appreciate that some readers may want to skip this somewhat technical primer; however, we believe that a better understanding of the relevant theory provides a foundation for performing better inferences. We then disscuss a variety of specific birth-death models, but emphasize that these examples represent only a tiny fraction of the possible diversification-rate models that can be specified in TESS.

2.1 The birth-death branching process

Our approach is based on the *reconstructed evolutionary process* described by Nee et al. (1994); a birth-death process in which only sampled, extant lineages are observed. Let $N(t)$ denote the number of species at time t. Assume the process starts at time t_1 (the 'crown' age of the most recent common ancestor of the study group, t_{MRCA}) when there are two species. Thus, the process is initiated with two species, $N(t_1) = 2$. We condition the process on sampling at least one descendant from each of these initial two lineages; otherwise t_1 would not correspond to the t_{MRCA} of our study group. Each lineage evolves independently of all other lineages, giving rise to exactly one new lineage with rate $b(t)$ and losing one existing lineage with rate $d(t)$ (Figure 3 and Figure 4). Note that although each lineage evolves independently, all lineages share both a common (tree-wide) speciation rate $b(t)$ and a common extinction rate $d(t)$ (Nee et al., 1994; Höhna, 2013, 2014, 2015; May et al., 2015). Additionally, at certain times, $t_{\mathbb{M}}$, a mass-extinction event occurs and each species existing at that time has the same probability, ρ, of survival. Finally, all extinct lineages are pruned and only the reconstructed tree remains (Figure 3).

To condition the probability of observing the branching times on the survival of both lineages that descend from the root, we divide by $P(N(T) > 0|N(0) = 1)^2$. Then, the probability density of the branching times, $\mathbb{T}$, becomes

$$P(\mathbb{T}) = \frac{\overbrace{P(N(T) = 1 \mid N(0) = 1)^2}^{\text{both initial lineages have one descendant}}}{\underbrace{P(N(T) > 0 \mid N(0) = 1)^2}_{\text{both initial lineages survive}}} \times \prod_{i=2}^{n-1} \overbrace{i \times b(t_i)}^{\text{speciation rate}} \times \overbrace{P(N(T) = 1 \mid N(t_i) = 1)}^{\text{singleton lineage has one descendant}},$$

and the probability density of the reconstructed tree (topology and branching

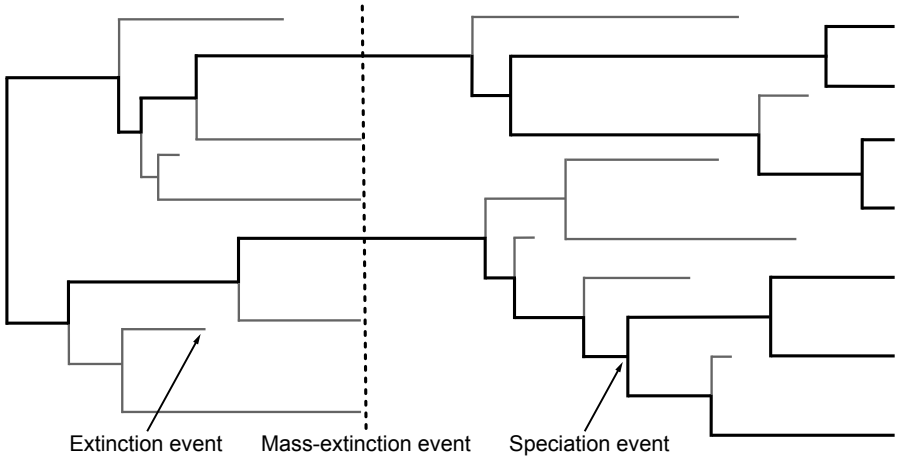

Figure 3: A realization of the birth-death process with mass extinction.

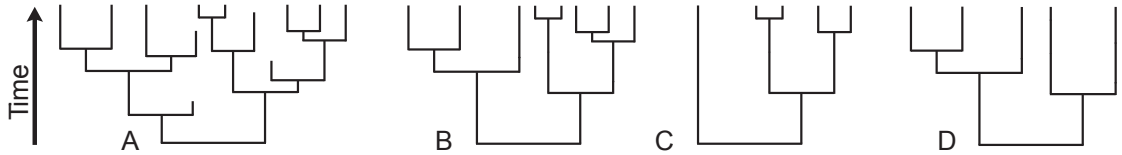

Figure 4: Examples of trees produced under a birth-death process. The process is initiated at the first speciation event (the 'crown-age' of the MRCA) when there are two initial lineages. At each speciation event the ancestral lineage is replaced by two descendant lineages. At an extinction event one lineage simply terminates. (A) A complete tree including extinct lineages. (B) The reconstructed tree of tree from A with extinct lineages pruned away. (C) A *uniform* subsample of the tree from B, where each species was sampled with equal probability, ρ. (D) A *diversified* subsample of the tree from B, where the species were selected so as to maximize diversity.

times) is then

$$P(\Psi) = \frac{2^{n-1}}{n!(n-1)!} \times \left(\frac{P(N(T)=1 \mid N(0)=1)}{P(N(T)>0 \mid N(0)=1)} \right)^2 \times \prod_{i=2}^{n-1} i \times b(t_i) \times P(N(T)=1 \mid N(t_i)=1) \quad (1)$$

We can expand Equation (1) by substituting $P(N(T)>0 \mid N(t)=1)^2 \exp(r(t,T))$ for $P(N(T)=1 \mid N(t)=1)$, where $r(u,v) = \int_u^v d(t) - b(t)dt$; the above equation

becomes

$$
\begin{aligned}
P(\Psi) &= \frac{2^{n-1}}{n!(n-1)!} \times \left(\frac{P(N(T) > 0 \mid N(0) = 1)^2 \exp(r(0,T))}{P(N(T) > 0 \mid N(0) = 1)} \right)^2 \\
&\quad \times \prod_{i=2}^{n-1} i \times b(t_i) \times P(N(T) > 0 \mid N(t_i) = 1)^2 \exp(r(t_i, T)) \\
&= \frac{2^{n-1}}{n!} \times \Big(P(N(T) > 0 \mid N(0) = 1) \exp(r(0,T)) \Big)^2 \\
&\quad \times \prod_{i=2}^{n-1} b(t_i) \times P(N(T) > 0 \mid N(t_i) = 1)^2 \exp(r(t_i, T)). \qquad (2)
\end{aligned}
$$

For a detailed description of this substitution, see Höhna (2015). Additional information regarding the underlying birth-death process can be found in (Thompson, 1975, Equation 3.4.6) and Nee et al. (1994) for constant rates and Lambert (2010); Höhna (2013, 2014, 2015) for arbitrary rate functions.

To compute the equation above we need to know the rate function, $r(t,s) = \int_t^s d(x) - b(x)dx$, and the probability of survival, $P(N(T) > 0|N(t) = 1)$. Yule (1925) and later Kendall (1948) derived the probability that a process survives ($N(T) > 0$) and the probability of obtaining exactly n species at time T ($N(T) = n$) when the process started at time t with one species. Kendall's results were summarized in Equation (3) and Equation (24) in Nee et al. (1994)

$$
\begin{aligned}
P(N(T)>0|N(t)=1) &= \left(1 + \int_t^T \Big(\mu(s) \exp(r(t,s)) \Big) ds \right)^{-1} \qquad (3) \\
P(N(T)=n|N(t)=1) &= (1 - P(N(T)>0|N(t)=1) \exp(r(t,T)))^{n-1} \\
&\quad \times P(N(T)>0|N(t)=1)^2 \exp(r(t,T)) \qquad (4)
\end{aligned}
$$

An overview for different diversification models is given in Höhna (2015).

2.2 The space of birth-death branching-process models

Our preceding discussion of the birth-death process makes it clear that we can define countless birth-death models that specify different speciation- and extinction-rate functions over time. We could assume, for example, that the extinction rate is constant over time, $d(t) = \mu$, or that the speciation rate decreases exponentially, $b(t) = \lambda * \exp(-\alpha * t)$. Furthermore, the constant-rate birth-death process can be parameterized in various ways, for example, by adopting parameters for the rate of speciation, $b(t) = \lambda$, and extinction, $d(t) = \mu$. Alternatively, we could

describe the birth-death process using parameters for the net-diversification rate, $\delta = \lambda - \mu$, and relative-extinction rate, $\epsilon = \mu/\lambda$, such that $b(t) = \delta/(1-\epsilon)$ and $d(t) = \epsilon * (\delta/(1-\epsilon))$. Finally, we could describe the birth-death process using parameters for the net-diversification rate, $\delta = \lambda - \mu$, and turnover rate, $\tau = \mu$, such that $b(t) = \delta + \tau$ and $d(t) = \tau$. Depending on the inference scenario, each of these parameterizations may offer advantages in terms of interpretation.

Below, we list several birth-death process models (*e.g.,* used in Höhna, 2014) to provide a sense of the types of models that can be specified and how they are parametrized in TESS (Table 1).

Table 1: Six different birth-death models with the corresponding parameters.

Model	$b(t)$	$d(t)$
Model 1	λ_0	0
Model 2	$\lambda_1 * \exp(-\alpha * t)$	0
Model 3	λ_0	μ
Model 4	$\lambda_0 + \lambda_1 * \exp(-\alpha * t)$	0
Model 5	$\lambda_1 * \exp(-\alpha * t)$	μ
Model 6	$\lambda_0 + \lambda_1 * \exp(-\alpha * t)$	μ

- Model 1: A constant-rate pure-birth (Yule) process (Yule, 1925). Under this process, the number of species increases monotonically and exponentially.
- Model 2: A decreasing-rate pure-birth process where the speciation rate declines toward zero. This process is equivalent to the decreasing-rate pure-birth process used in Rabosky and Lovette (2008). Under this process, the number of species increases monotonically.
- Model 3: A constant-rate birth-death process, as used in Thompson (1975). Under this process, the expected number of species increases exponentially.
- Model 4: A pure-birth process with a decaying rate of speciation but a constant, non-zero speciation rate the longer the process continues ($\lambda(t) = \lambda_0 + \lambda_1 * \exp(-\alpha * t)$). Thus, the process does not stop producing new species after the initial burst, as in Model 2. As in the other two pure-birth processes, the number of species increases monotonically.
- Model 5: A birth-death process with an initial expansion phase (where the speciation rate exceeds the extinction rate) that subsequently converges to a critical-branching process, *i.e.,* where the speciation and extinction rates

are equal, $\lambda(t) = \mu + \lambda_1 * \exp(-\alpha * t)$ and $\mu(t) = \mu$. Although one might assume that the expected number of species will remain constant for a critical-branching process, this does not hold if the process is conditioned on survival.

- Model 6: A birth-death process where the extinction rate remains constant, but speciation rate has an initially constant phase followed by a decreasing phase. This model corresponds to an early phase of radiation, followed by a phase of steady increase, $\lambda(t) = \lambda_0 + \lambda_1 * \exp(-\alpha * t)$ and $\mu(t) = \mu$.

The parametrizations of these models are listed in Table 1, and the expected number of species, $E[N(T)]$, at time T under each model is depicted in Figure 5. We derive $E[N(T)]$ analytically by using the fact that $N(T)$ is geometrically distributed (see Equation 5 in Höhna, 2013). Note that the process is conditioned on survival to the present, such that $E[N(T)]$ increases even if $\lambda(t) = \mu(t)$.

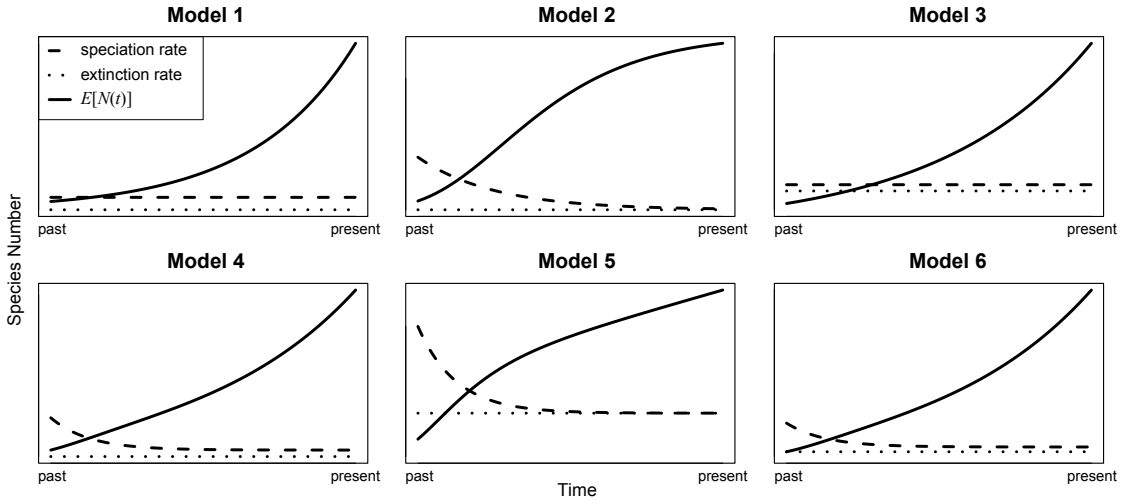

Figure 5: Six possible birth-death models. Each plot shows the speciation and extinction rates over time, and also the expected number of species ($E[N(t)]$). Model 1: A constant-rate pure-birth process. Model 2: A decreasing-rate pure-birth process with speciation rate declining to zero. Model 3: A constant-rate birth-death process. Model 4: A pure-birth process, where the speciation rate passes through a constant phase to a decreasing phase. Model 5: A birth-death process with an initial expansion phase (speciation rate > extinction rate) that later converges to a critical-branching process. Model 6: A birth-death process with a constant extinction rate, where the speciation rate is initially constant and later decreases.

2.3 Simulating data

Simulating phylogenies is critical for validating methods/models of lineage diversification, and is also invaluable for developing our intuition about the behavior of these models. Simulations are also crucial for assessing the adequacy (absolute fit) of a model for a given dataset, which we will describe later. In the previous section we described the expected form of lineage-accumulation curves under different branching-process models. We will now briefly explain how to simulate phylogenies using TESS.

We will explore some common diversification models, including the constant-rate pure-birth process, the constant-rate birth-death process, and the exponentially decaying pure-birth process. Specifically, we will use TESS to simulate 50 trees under these models and look at the corresponding LTT plots. You can experiment with the parameter settings to better understand their impact, *e.g.,* the influence of the extinction rate.

We will first simulate trees under a constant-rate pure-birth process, where we specify a speciation rate of 1.0 and the duration of the process as 3.0 time units.

```
speciation <- 1.0
extinction <- 0.0
tmrca <- 3.0
```

Here, we are explicitly conditioning the simulation on the time of the process. Because it is a stochastic process, this will result in simulated trees of different sizes (number of species), which may be relevant to our question. We might, for example, wish to know whether the observed species diversity in our study tree is improbable under the current model and parameterization. In the next subsection we will show how to simulate trees conditioned on the number of extant species.

We simulate 50 trees under the specified model as follows:

```
trees <- tess.sim.age(n = 50,
                      age = tmrca,
                      lambda = speciation,
                      mu = extinction,
                      MRCA = TRUE)
```

Note that we are initializing the simulation with two species; *i.e.,* from the 'crown age' of the most recent common ancestor (MRCA). Accordingly, the resulting trees will not have 'stem' branches subtending their root nodes; instead, these trees begin at the root node that corresponds to the first speciation event in each tree (*c.f.*, Figure 4). This scenario corresponds well with empirical trees, where (by definition) at least one species from both of these two initial lineages

will survive to the present (otherwise we would not recognize this node as the root of our study tree).

Next, we will generate the lineage-through-time plots for all 50 simulated trees.

```
mltt.plot(trees,
          log = "y",
          dcol = FALSE,
          legend = FALSE,
          backward = FALSE)
```

For a fully specified model, TESS can calculate the expected number of lineages through time. We will overlay a curve describing the the expected number of lineages on the LTT plot.

```
expected <- function(t)
  tess.nTaxa.expected(begin = 0,
                      t = t,
                      end = tmrca,
                      lambda = speciation,
                      mu = extinction,
                      MRCA = TRUE,
                      reconstructed = TRUE)

curve(expected,add=FALSE,col="red",lty=2,lwd=5)
legend("topleft",col="red",lty=2,"Expected Diversity")
```

The results of this simulation are shown in Figure 6A. Here, you can see that the shape of the LTT curve is clearly linear (in log-scale) under a constant-rate pure-birth process. All other curves will be rendered in log-scale for convenience. Notice also that we used the argument `reconstructed = TRUE` which means that we compute the expected number of species (diversity) of a reconstructed phylogeny. This must be a monotonically increasing function. You could plot the expected diversity at any given time and compare it to the diversity of reconstructed phylogeny.

We will now repeat the above simulation under a constant-rate birth-death process. First, we set the parameters of the model.

```
speciation <- 5.0
extinction <- 4.0
tmrca <- 3.0
```

Then simulate 50 trees under these parameters.

```
trees <- tess.sim.age(n = 50,
                      age = tmrca,
                      lambda = speciation,
                      mu = extinction,
                      MRCA = TRUE)
```

Next, we plot the lineage-through-time curves for the simulated trees:

```
mltt.plot(trees,
          log = "y",
          dcol = FALSE,
          legend = FALSE,
          backward = FALSE)
```

Finally, we overlay the expected number of lineages on our LTT plot. In this example you may notice that the expected number of lineages under the birth-death process diverges from the expected number of lineages in the reconstructed tree. This is simply because the expected number of lineages in the reconstructed tree only considers lineages that have at least one descendant sampled at the present time, whereas the expected number of lineages gives the expected diversity at the time without that constraint.

```
expected <- function(t)
  tess.nTaxa.expected(begin = 0,
                      t = t,
                      end = tmrca,
                      lambda = speciation,
                      mu = extinction,
                      MRCA = TRUE,
                      reconstructed = TRUE)

curve(expected,add=TRUE,col="red",lty=2,lwd=5)
legend("topleft",col="red",lty=2,"Expected Diversity")
```

The results of this simulation are shown in Figure 6B. Notice that the slope of the LTT plot increases sharply near the present: this is commonly referred the 'pull-of-the-present' effect. This effect becomes more pronounced as the relative-extinction rate (*i.e.*, extinction ÷ speciation) increases.

Finally, we will consider a pure-birth process with exponentially decreasing speciation rate. In TESS you can either specify a simple numeric value for the speciation and extinction rates or you can specify a function that takes the time t as a parameter. Here, we will use the second option.

```
speciation <- function(t) 0.5 + 2 * exp(-1.0*t)
extinction <- 0.0
tmrca <- 3.0
```

We again simulate 50 trees conditioned on the survival of the two initial lineages.

```
trees <- tess.sim.age(n = 50,
                      age = tmrca,
                      lambda = speciation,
                      mu = extinction,
                      MRCA = TRUE)
```

We generate the LTT plots for the simulated trees.

```
mltt.plot(trees,
          log = "y",
          dcol = FALSE,
          legend = FALSE,
          backward = FALSE)
```

And then add the expected number of lineages in the reconstructed phylogeny.

```
expected <- function(t)
  tess.nTaxa.expected(begin = 0,
                      t = t,
                      end = tmrca,
                      lambda = speciation,
                      mu = extinction,
                      MRCA = TRUE,
                      reconstructed = TRUE)

curve(expected,add=TRUE,col="red",lty=2,lwd=5)
legend("topleft",col="red",lty=2,"Expected Diversity")
```

The results of the three simulations are shown in Figure 6.

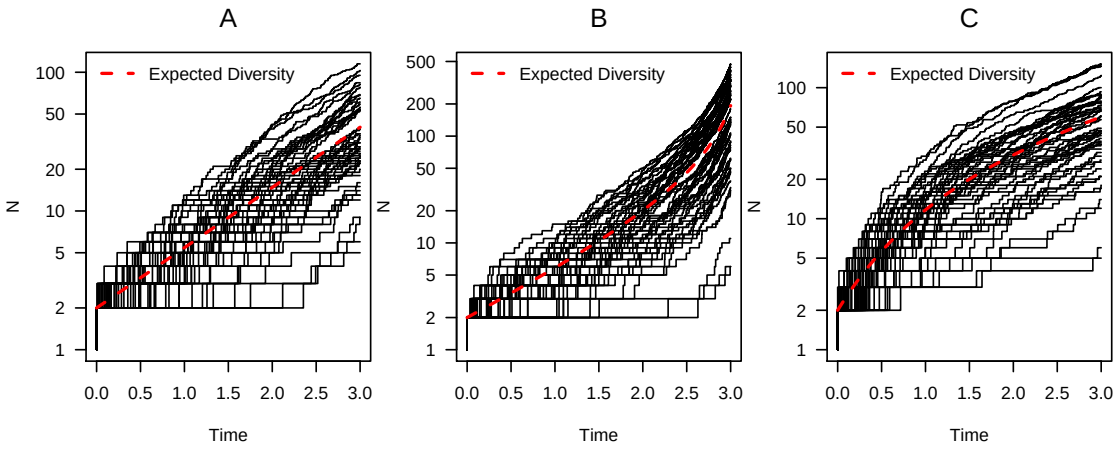

Figure 6: Lineage-through-time curves for pure-birth trees (panel A), birth-death trees (panel B), and pure-birth trees with exponentially decreasing speciation rate (panel C).

2.4 Estimating parameters using Markov chain Monte Carlo (MCMC)

In the previous section we introduced some stochastic-branching process models, and demonstrated how to simulate trees under those models. Here, we turn to the issue of estimating parameters of branching-process models from empirical data. We estimate parameters within a Bayesian statistical framework, which adopts the perspective that parameters are random variables. Accordingly, it is necessary to specify a probability distribution for each parameter that describes the nature of that random variation. These *prior* probability distributions describe our beliefs about the parameter values before evaluating the data at hand. Prior probabilities are updated by the information in the data (via the likelihood function) to provide the corresponding *posterior* probability distributions. These posterior probability distributions reflect our belief about the parameter values after incorporating the new information in our data. We estimate the joint posterior probability density of the model parameters from the data using numerical methods—Markov chain Monte Carlo (MCMC) algorithms.

2.4.1 Birth-death processes with constant rates

We first consider the constant-rate birth-death process. Although we do not explicitly consider the constant-rate pure-birth process, it can easily be specifed by simply setting the extinction rate of the constant-rate birth-death process to zero.

First, we specify prior distributions for our parameters. The constant-rate birth-death process has two parameters; the speciation rate and extinction rate.

There are many possible prior distributions that we might adopt for these two parameters, *e.g.,* the exponential, gamma, lognormal distributions. Here, we will use an exponential distribution, which has a single parameter (the rate parameter) that describes the shape of the distribution. We will specify a value of 0.1 for the rate parameter (such that the mean of the exponential is 1/rate = 10.0). We will use identical priors for both the speciation- and extinction-rate parameters.

In TESS the prior distribution must be functions that can be computed for all values that can be realized by the corresponding parameter (*e.g.,* priors for rates must only include positive-real values). Furthermore, the prior distributions need to return log-transformed probabilities (this is a standard convention adopted to avoid underflow in computer memory).

```
prior_delta <- function(x) { dexp(x,rate=10.0,log=TRUE) }
prior_tau <- function(x) { dexp(x,rate=10.0,log=TRUE) }
priorsConstBD <- c("diversification"=prior_delta,
                   "turnover"=prior_tau)
```

If you provide names for the prior distributions, as we did here, then these names will be used to label that parameters in the MCMC output. Currently, only the names of the priors are used.

Next, we set up the likelihood of the constant-rate birth-death process as an R function. Here, the actual likelihood computation is performed by the function `tess.likelihood`. It is necessary to wrap the TESS likelihood into another R function because you need to specify how the speciation and extinction rates are assembled and which assumptions/conditions are applied. This approach enables maximal flexibility for using TESS.

```
likelihoodConstBD <- function(params) {

  speciation <- params[1] + params[2]
  extinction <- params[2]

  lnl <- tess.likelihood(times,
                         lambda = speciation,
                         mu = extinction,
                         samplingProbability = 1.0,
                         log = TRUE)

  return (lnl)

}
```

It is also possible to specify prior distributions on other parameterizations of the constant-rate birth-death model, *e.g.,* using parameters for the net-diversification rate (speciation−extinction) and the relative-extinction rate (extinction/speciation). This alternative parmaterization of the model would, of course, require modification of the likelihood function.

Next, we use the function `tess.mcmc` to run an MCMC simulation. The function takes in several arguments to describe the MCMC algorithm. Specifically, you must specify the `likelihoodFunction`, `priors`, and initial values for the `parameters`. Additionally, you can specify whether the MCMC proposal mechanisms should operate on the log-transformed parameters, which is advisable for rate parameters but not for location parameters.

We will also specify the value for the `delta` parameter, which defines the (initial) width of the sliding-window proposal mechanism. This `delta` tuning parameter determines the scale (severity) of the proposal mechanism: larger values will specify more severe changes to the current parameter value when that parameter is being updated during the MCMC. We will discuss these issues in more detail in Section 5 of this guide. The remaining parameters specify the number of iterations of the MCMC simulation, the number of iterations for the pre-burnin phase, the thinning schedule, and whether the scale of the poposal mechanisms are to be automatically tuned.

```
set.seed(12345)
samplesConstBD <- tess.mcmc(likelihoodFunction = likelihoodConstBD,
                            priors = priorsConstBD,
                            parameters = runif(2,0,1),
                            logTransforms = c(TRUE,TRUE),
                            delta = c(1,1),
                            iterations = 2000,
                            burnin = 200,
                            thinning = 10,
                            adaptive = TRUE,
                            verbose = TRUE)

## Burning-in the chain ...
## 0--------25--------50--------75--------100
## ========================================
## Finished burnin period!
##
## Running the chain ...
## 0--------25--------50--------75--------100
## ========================================
```

```
## Finished MCMC!
##
## Parameter | delta | Acceptance Probability
## ========================================
## diversification | 0.768 | 0.480
## turnover | 0.244 | 0.372
```

Note that we have specified a starting seed for the random-number generator. We have done this only to ensure that your results will be identical to those in this guide. However, you *should not* specify the starting seed for your analyses, but instead use a random starting seed that is automatically generated from the system clock. This is important, as you will want to perform multiple independent MCMC simulations to assess convergence. The basic idea is to compare parameter estimates from multiple independent analyses: if the chains have converged to the target (joint posterior probability) distribution, then the parameter estimates from the replicate chains should be similar. However, this important diagnostic would be rendered meaningless if the replicate analyses were performed under the same starting seed; in this case, the results are guaranteed to be identical.

Note that we ran a very short MCMC simulation above for covenience. In practice, MCMC simulations are commonly run for 10^5 to 10^8 iterations. We will use the R package coda (which is automatically loaded with TESS) to summarize the samples from our MCMC simulation. TESS saves samples in the coda format, which allows us to easily summarize our samples:

```
summary(samplesConstBD)

##
## Iterations = 1:201
## Thinning interval = 1
## Number of chains = 1
## Sample size per chain = 201
##
## 1. Empirical mean and standard deviation for each variable,
##    plus standard error of the mean:
##
##                   Mean       SD  Naive SE Time-series SE
## diversification 0.0060 0.002262 0.0001596      0.0001873
## turnover        0.1512 0.013622 0.0009608      0.0009608
##
## 2. Quantiles for each variable:
##
```

```
##                      2.5%     25%      50%      75%   97.5%
## diversification 0.001824 0.00449 0.005951 0.007532 0.01051
## turnover        0.124930 0.14165 0.151321 0.160658 0.17833
```

We can also visualize the trace plots and marginal posterior probability densities for these samples (Figure 7).

```
plot(samplesConstBD)
```

2.4.2 Birth-death processes with continuously varying rates

Here we consider a birth-death process with an exponentially decreasing speciation rate. Specifically, we define the speciation rate as $\lambda(t) = \delta + \lambda \exp(-\alpha * t)$ and extinction rate as $\mu(t) = \delta$ (Höhna, 2014). It is not possible to analytically compute the probability density (or likelihood) under this process. Instead, we approximate these quantities using numerical integration techniques. These numerical methods are implemented in TESS and will be performed automatically if you provide functions instead of numerical arguments for the speciation and/or extinction rate. The numerical integration is very convenient but, of course, imposes a higher computational cost that will make these analyses run more slowly.

The decreasing speciation rate birth-death model has three parameters: δ, λ_1, and α. We will use an exponential prior probability distribution with a rate of 0.1 (*i.e.,* with a mean of 10.0) for all three parameters. As before, the prior distributions must be functions that return the log-transformed probability for a given value of the parameter.

```
prior_delta <- function(x) { dexp(x,rate=0.1,log=TRUE) }
prior_lambda <- function(x) { dexp(x,rate=10.0,log=TRUE) }
prior_alpha <- function(x) { dexp(x,rate=0.1,log=TRUE) }
priorsDecrBD <- c("turnover"=prior_delta,
                  "initial speciation"=prior_lambda,
                  "speciation decay"=prior_alpha)
```

We now specify the speciation and extinction rates as functions and pass them into the likelihood, which again must be provided as a function.

```
likelihoodDecrBD <- function(params) {

  speciation <- function(t) params[1] + params[2] * exp(-params[3]*t)
```

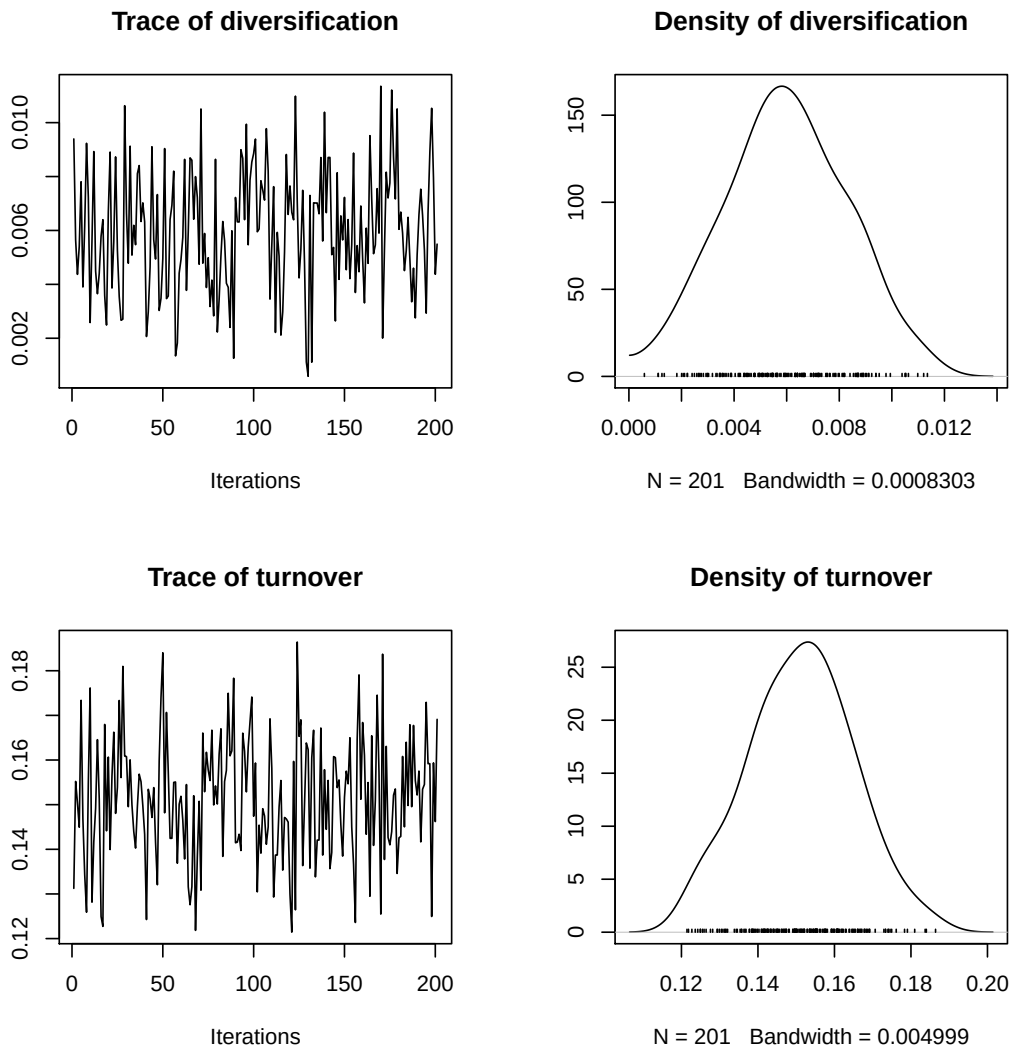

Figure 7: Trace plots (left) and marginal posterior probability densities (right) for the speciation rate (variable 1) and extinction rate (variable 2) from the MCMC simulation under the constant-rate birth-death process.

```
extinction <- function(t) params[1]

lnl <- tess.likelihood(times,
                       lambda = speciation,
                       mu = extinction,
                       samplingProbability = 1.0,
```

```
                          log = TRUE)

  return (lnl)

}
```

Next, we start the analysis by calling the MCMC function in TESS. (The details of this MCMC simulation are similar to those described in the constant-rate birth-death example, above.)

```
set.seed(12345)
samplesDecrBD <- tess.mcmc(likelihoodFunction = likelihoodDecrBD,
                           priors = priorsDecrBD,
                           parameters = runif(3,0,1),
                           logTransforms = c(TRUE,TRUE,TRUE),
                           delta = c(1,1,1),
                           iterations = 2000,
                           burnin = 200,
                           thinning = 10,
                           adaptive = TRUE,
                           verbose = TRUE)

## Burning-in the chain ...
## 0--------25--------50--------75--------100
## ========================================
## Finished burnin period!
##
## Running the chain ...
## 0--------25--------50--------75--------100
## ========================================
## Finished MCMC!
##
## Parameter | delta | Acceptance Probability
## ========================================
## turnover | 0.193 | 0.415
## initial speciation | 2.219 | 0.514
## speciation decay | 2.389 | 0.481
```

We then summarize the parameter estimates from our MCMC samples:

```
summary(samplesDecrBD)

##
## Iterations = 1:201
## Thinning interval = 1
## Number of chains = 1
## Sample size per chain = 201
##
## 1. Empirical mean and standard deviation for each variable,
##    plus standard error of the mean:
##
##                       Mean      SD  Naive SE Time-series SE
## turnover            0.1622 0.01121 0.0007907      0.0009085
## initial speciation 0.1117 0.10594 0.0074726      0.0074726
## speciation decay   9.0703 8.92827 0.6297515      0.5847390
##
## 2. Quantiles for each variable:
##
##                         2.5%     25%     50%     75%   97.5%
## turnover            0.140330 0.15455 0.16246  0.1693  0.1839
## initial speciation 0.001673 0.02818 0.08409  0.1655  0.3817
## speciation decay   0.299961 2.50382 6.91353 12.6620 34.7207
```

We can also visualize the trace plots and marginal posterior probability densities for these samples:

```
plot(samplesDecrBD)
```

2.4.3 Birth-death processes with episodically varying rates

The next model we consider is a birth-death process with piecewise-constant rates. Under this model, rates of speciation and extinction change at some (discrete) number of events; between these rate-shift events, however, the diversification-rate parameters remain constant (Höhna, 2015).

The number of parameters included in the episodic model varies depending on the number of rate-shift events. In general, there are $k_{\mathbb{B}} + 1$ speciation-rate parameters and $k_{\mathbb{D}} + 1$ extinction-rate parameters, where $k_{\mathbb{B}}$ is the number of speciation-rate shifts and $k_{\mathbb{D}}$ is the number of extinction-rate shifts.

In this example, we will assume there is a single speciation-rate shift and a single extinction-rate shift, both occurring at the mid-point of the duration spanned by the conifer tree. First, we specify the time of the rate-shift event.

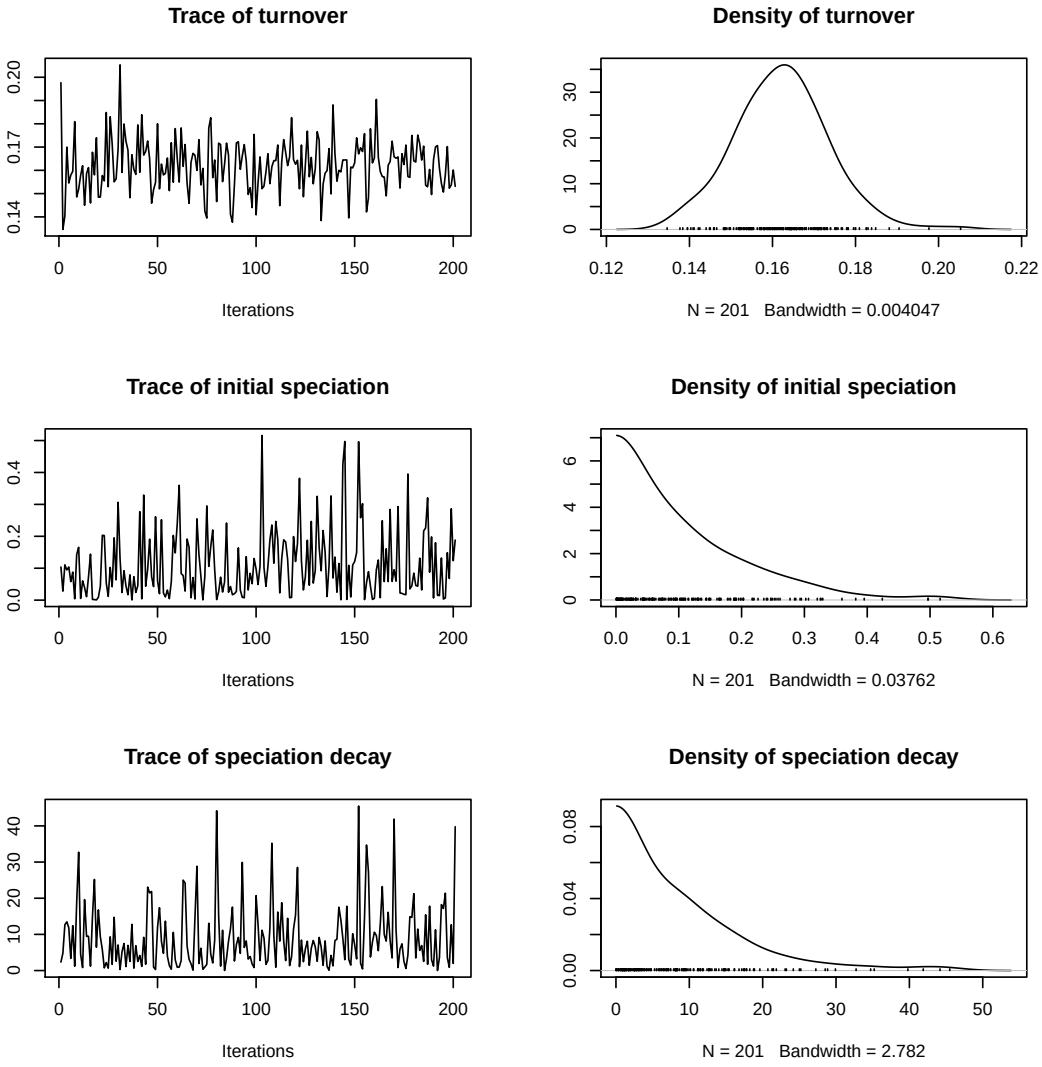

Figure 8: Trace plots and estimated posterior distribution of the parameter under the decreasing speciation rate birth-death model.

```
rateChangeTime <- max( times ) / 2
```

Next, we specify priors for the parameters. There are a total of four parameters (the speciation and extinction rates before and after the rate-shift event). Accordingly, we specify four identical exponential priors for these parameters, all with a rate of 10.0 (and a mean of 0.1).

```
prior_delta_before <- function(x) { dexp(x,rate=10.0,log=TRUE) }
prior_tau_before <- function(x) { dexp(x,rate=10.0,log=TRUE) }
prior_delta_after <- function(x) { dexp(x,rate=10.0,log=TRUE) }
prior_tau_after <- function(x) { dexp(x,rate=10.0,log=TRUE) }
priorsEpisodicBD <- c("diversification before"=prior_delta_before,
                      "turnover before"=prior_tau_before,
                      "diversification after"=prior_delta_after,
                      "turnover after"=prior_tau_after)
```

Next, we specify a likelihood function using the rate-shift model implemented in TESS, `tess.likelihood.rateshift`.

```
likelihoodEpisodicBD <- function(params) {

  speciation <- c(params[1]+params[2],params[3]+params[4])
  extinction <- c(params[2],params[4])

  lnl <- tess.likelihood.rateshift(times,
                                   lambda = speciation,
                                   mu = extinction,
                                   rateChangeTimesLambda = rateChangeTime,
                                   rateChangeTimesMu = rateChangeTime,
                                   samplingProbability = 1.0,
                                   log = TRUE)

  return (lnl)

}
```

Now we can start the analysis by calling the MCMC function in TESS. (The details of this MCMC simulation are similar to those described in the constant-rate birth-death example, above.)

```
set.seed(12345)
samplesEpisodicBD <- tess.mcmc(likelihoodFunction = likelihoodEpisodicBD,
                               priors = priorsEpisodicBD,
                               parameters = runif(4,0,1),
                               logTransforms = c(TRUE,TRUE,TRUE,TRUE),
                               delta = c(1,1,1,1),
                               iterations = 2000,
                               burnin = 200,
                               thinning = 10,
                               adaptive = TRUE,
                               verbose = TRUE)

## Burning-in the chain ...
## 0--------25--------50--------75--------100
## ========================================
## Finished burnin period!
##
## Running the chain ...
## 0--------25--------50--------75--------100
## ========================================
## Finished MCMC!
##
## Parameter | delta | Acceptance Probability
## ========================================
## diversification before | 0.980 | 0.531
## turnover before | 1.592 | 0.430
## diversification after | 1.058 | 0.417
## turnover after | 0.222 | 0.385
```

We then summarize the parameter estimates from our MCMC samples:

```
summary(samplesEpisodicBD)

##
## Iterations = 1:201
## Thinning interval = 1
## Number of chains = 1
## Sample size per chain = 201
##
## 1. Empirical mean and standard deviation for each variable,
##    plus standard error of the mean:
```

```
##
##                             Mean       SD  Naive SE Time-series SE
## diversification before 0.010951 0.006429 0.0004535      0.0005027
## turnover before        0.113643 0.073380 0.0051759      0.0051759
## diversification after  0.006036 0.002500 0.0001764      0.0001764
## turnover after         0.146843 0.011637 0.0008208      0.0008208
##
## 2. Quantiles for each variable:
##
##                             2.5%      25%      50%      75%   97.5%
## diversification before 0.001098 0.006567 0.009774 0.016035 0.02347
## turnover before        0.011698 0.063437 0.099464 0.145658 0.26717
## diversification after  0.001683 0.004280 0.005758 0.007679 0.01105
## turnover after         0.126521 0.138313 0.146133 0.154644 0.16911
```

Finally, we can visualize the trace plots and marginal posterior probability densities for these samples:

```
plot(samplesEpisodicBD)
```

2.4.4 Birth-death processes with explicit mass-extinction events

The final model we consider is one where speciation and extinction rates are constant, but where there is a single mass-extinction event at some unknown time. We'll assume that 10% of the species survive the mass-extinction event.

```
survivalProbability <- 0.1
```

There are three parameters in the model: the speciation rate, the extinction rate, and the mass-extinction time. We must specify priors for each of these parameters. For simplicity, we'll assume *a priori* that the mass-extinction event could happen at any time in the most recent half of the tree with equal probability.

```
prior_delta <- function(x) { dexp(x,rate=10.0,log=TRUE) }
prior_tau <- function(x) { dexp(x,rate=10.0,log=TRUE) }
prior_time <- function(x) { dunif(x,min=max(times)/2,max=max(times),log=TRUE)}
priorsMassExtinctionBD <- c("diversification"=prior_delta,
                            "turnover"=prior_tau,
                            "mass-extinction time"=prior_time)
```

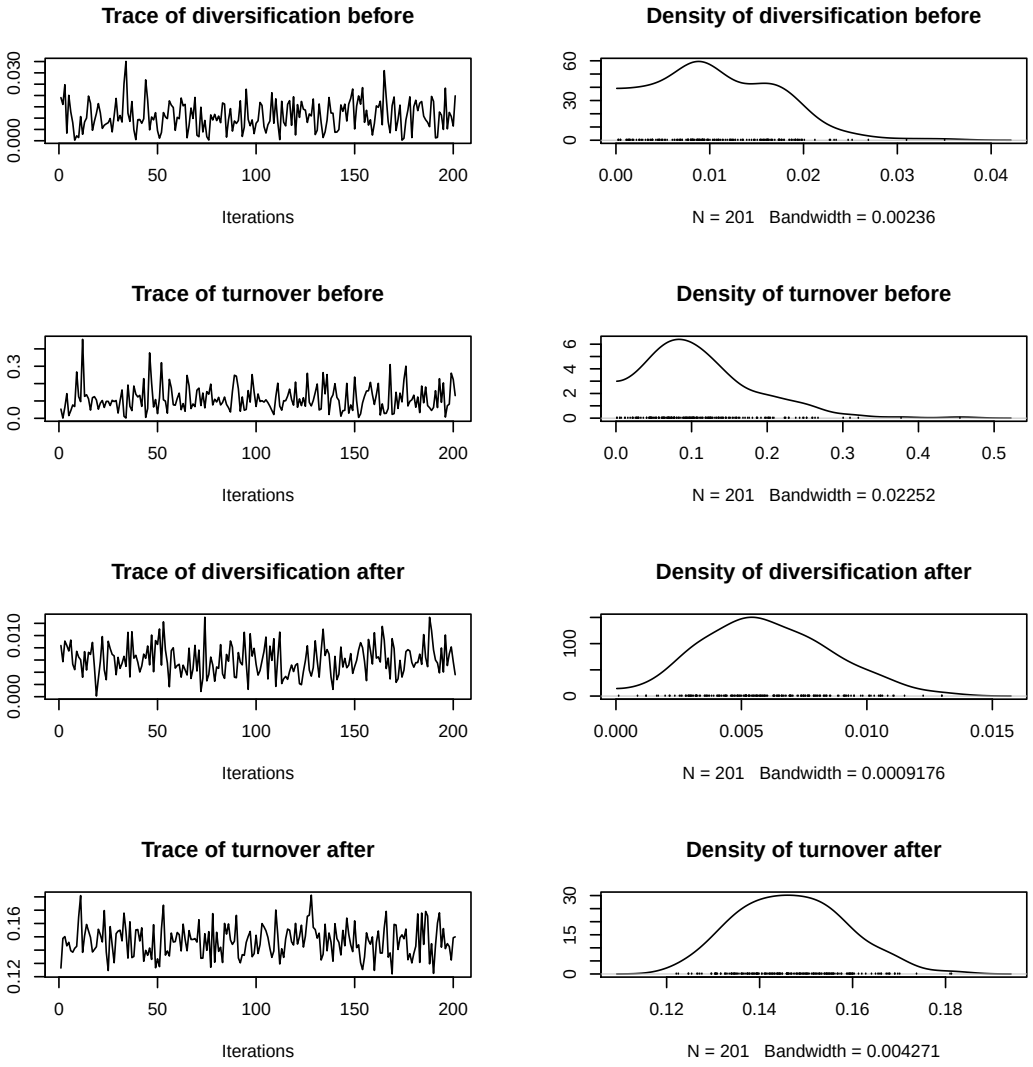

Figure 9: Trace plots and estimated posterior distributions of the parameters under a birth-death-shift model.

Next, we specify a likelihood function. We can use either the standard likelihood function `tess.likelihood` or the likelihood function of the rate-shift model `tess.likelihood.rateshift`.

```
likelihoodMassExtinctionBD <- function(params) {

  speciation <- params[1]+params[2]
  extinction <- params[2]
  time <- params[3]

  lnl <- tess.likelihood(times,
                         lambda = speciation,
                         mu = extinction,
                         massExtinctionTimes = time,
                         massExtinctionSurvivalProbabilities =
                                     survivalProbability,
                         samplingProbability = 1.0,
                         log = TRUE)

  return (lnl)

}
```

Now we can start the analysis by calling the MCMC function in TESS. (The details of this MCMC simulation are similar to those described in the constant-rate birth-death example, above.)

```
set.seed(12345)
samplesMassExtinctionBD <- tess.mcmc(likelihoodFunction =
                                       likelihoodMassExtinctionBD,
                                       priors = priorsMassExtinctionBD,
                                       parameters = c(runif(2,0,1),max(times)*3/4),
                                       logTransforms = c(TRUE,TRUE,FALSE),
                                       delta = c(1,1,1),
                                       iterations = 2000,
                                       burnin = 200,
                                       thinning = 10,
                                       adaptive = TRUE,
                                       verbose = TRUE)

## Burning-in the chain ...
```

```
## 0--------25--------50--------75--------100
## ========================================
## Finished burnin period!
##
## Running the chain ...
## 0--------25--------50--------75--------100
## ========================================
## Finished MCMC!
##
## Parameter | delta | Acceptance Probability
## ========================================
## diversification | 0.370 | 0.386
## turnover | 0.174 | 0.505
## mass-extinction time | 1.668 | 0.431
```

We then summarize the parameter estimates from our MCMC samples:

```
summary(samplesMassExtinctionBD)

##
## Iterations = 1:201
## Thinning interval = 1
## Number of chains = 1
## Sample size per chain = 201
##
## 1. Empirical mean and standard deviation for each variable,
##    plus standard error of the mean:
##
##                           Mean       SD  Naive SE Time-series SE
## diversification        0.01912 0.002312 0.0001631      0.0001631
## turnover               0.12706 0.011475 0.0008094      0.0008094
## mass-extinction time 262.02657 2.242028 0.1581405      0.4675061
##
## 2. Quantiles for each variable:
##
##                            2.5%       25%       50%       75%     97.5%
## diversification        0.01457   0.01778   0.01904   0.02058   0.02417
## turnover               0.10449   0.11963   0.12697   0.13415   0.14789
## mass-extinction time 253.51488 261.81437 262.82410 263.22746 263.70129
```

Finally, we visualize the trace plots and marginal posterior probability densities for these samples:

```
plot(samplesMassExtinctionBD)
```

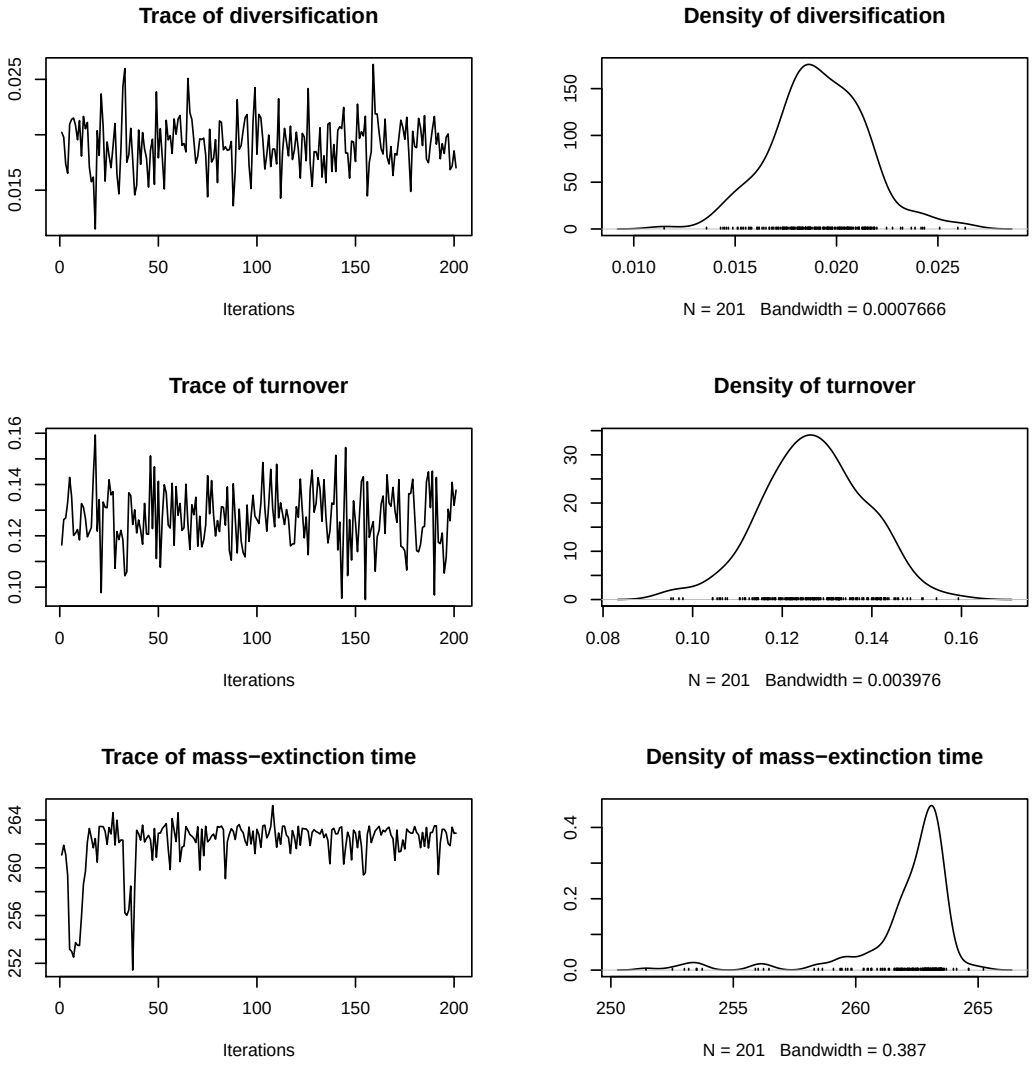

3 Accommodating Incomplete Species Sampling

Most phylogenies do not contain all species of the group under study. Instead, only an *incomplete sample*—or subsample—of all described species are included. Assuming complete species sampling for trees that are actually incomplete is known to bias estimates of diversification rates (Cusimano and Renner, 2010; Höhna et al., 2011). Additionally, the sampling strategy (*e.g.,* whether species are sampled at random or to maximize diversity) also influences the parameter estimates (Höhna et al., 2011; Höhna, 2014). Fortunately, methods for modeling incomplete species sampling exist to correct for the introduced bias.

Here we consider two approach of incomplete species sampling: *uniform* sampling and *diversified* sampling. A sketch of the two sampling methods was provided in Figure 4. As we will demonstrate below, the sampling strategy has a substantial influence on the distribution of branching times in the tree, which results in different patterns in the lineage-through-time curves for the different sampling schemes. Additionally, we note that the patterns induced by incomplete sampling can mimic patterns of decreasing rates of lineage diversification (Pybus and Harvey, 2000; Cusimano and Renner, 2010; Höhna et al., 2011), and thus it is critical to incorporate incomplete sampling in any study of lineage diversification rates.

We will begin by simulating trees under different kinds of sampling schemes, and then demonstrate how we can incorporate these sampling schemes into branching-process models in TESS.

3.1 Patterns of incomplete sampling

To demonstrate the impact of incomplete species sampling, we will simulate trees under each sampling strategy and plot the resulting LTT curves. We simulate $n = 50$ trees, each conditioned on the specified age. First, we simulate trees with complete species sampling to compare against the other sampling schemes. Next, we simulate trees under uniform species sampling with a sampling probability of $\rho = 0.25$ (which means that each species at the present has the same probability of $\rho = 0.25$ being included in the phylogeny; if a species is not sampled then its lineage is removed from the reconstructed tree). Finally, we simulate trees under diversified species sampling with a sampling probability of $\rho = 0.25$ (*i.e.,* only the oldest 25% of divergence events are included in the reconstructed phylogeny, and all later divergence events are excluded).

```
# Birth-death
birthDeathSpeciationSampling <- 2.0
birthDeathExtinctionSampling <- 1.0
```

```
birthDeathTreesComplete <- tess.sim.age(n = 50,
                                age = 3.0,
                                lambda = birthDeathSpeciationSampling,
                                mu = birthDeathExtinctionSampling,
                                MRCA = TRUE)

birthDeathTreesUniform <- tess.sim.age(n = 50,
                                age = 4.0,
                                lambda = birthDeathSpeciationSampling,
                                mu = birthDeathExtinctionSampling,
                                samplingProbability = 0.25,
                                samplingStrategy = "uniform",
                                MRCA = TRUE)

birthDeathTreesDiversified <- tess.sim.age(n = 50,
                                age = 4.0,
                                lambda = birthDeathSpeciationSampling,
                                mu = birthDeathExtinctionSampling,
                                samplingProbability = 0.25,
                                samplingStrategy = "diversified",
                                MRCA = TRUE)



par(mfrow=c(1,3),mar=c(5,4,3,0.1),las=1)

# Plot the trees
mltt.plot(birthDeathTreesComplete,log = "y",dcol = FALSE,
          legend = FALSE,backward = FALSE)
mtext("A", line = 1)

mltt.plot(birthDeathTreesUniform,log = "y",dcol = FALSE,
          legend = FALSE,backward = FALSE)
mtext("B", line = 1)

mltt.plot(birthDeathTreesDiversified,log = "y",dcol = FALSE,
          legend = FALSE,backward = FALSE)
mtext("C", line = 1)
```

For the remainder of this section, we focus on the biases stemming from incomplete species sampling. We will simulate a single incompletely sampled tree

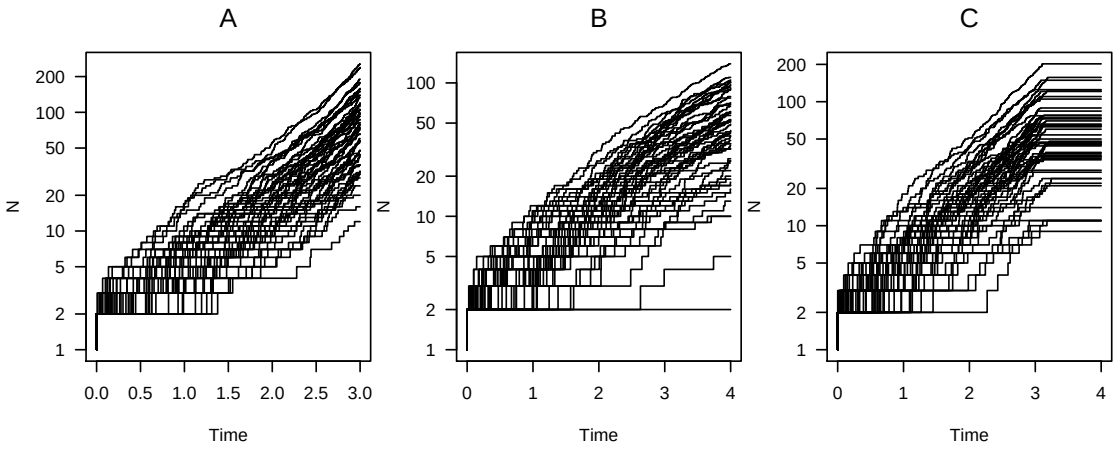

Figure 10: Lineage-through-time plots for completely sampled trees (panel A), incomplete trees with uniform sampling (panel B), and incomplete trees with diversified sampling (panel C).

under a diversified sampling strategy with sampling fraction $\rho = 0.25$, and in the following sections we will estimate parameters under various birth-death processes from this tree. We have simulated the tree under diversified sampling with known speciation and extinction rates, which allows us to compare estimates of parameter value using various approaches to the true parameter values in order to better understand the influence of the sampling strategy on parameter estimates.

```
# simulate a tree under diversified taxon sampling
tree.diversified <- tess.sim.age(n = 1,
                                 age = 4.0,
                                 lambda = 2.0,
                                 mu = 1.0,
                                 samplingProbability = 0.25,
                                 samplingStrategy = "diversified",
                                 MRCA = TRUE)[[1]]

# extract the branching times from the tree
times.diversified <- as.numeric( branching.times(tree.diversified) )
```

First, we will take a look at the simulated tree: it looks similar to many empirical phylogenies.

```
plot(tree.diversified,show.tip.label=FALSE,no.margin=TRUE)
```

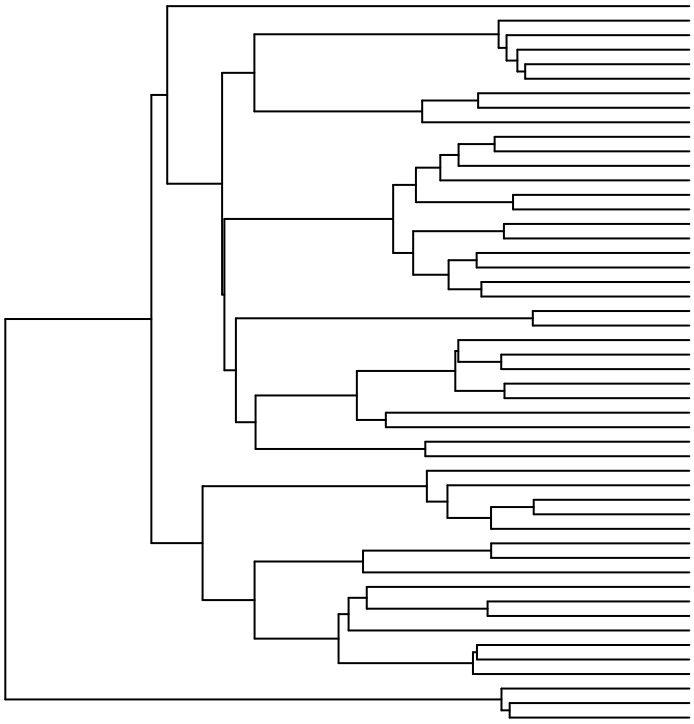

Figure 11: The simulated tree under *diversified* sampling with sampling fraction $\rho = 0.25$.

3.2 Uniform species sampling

Uniform species sampling, which is sometimes also called *random* species sampling, assumes that every species at present has the same probability ρ of being included in the sample. That is, regardless of age or phylogenetic relationship, this method assumes that a researcher flips a coin for each species to decide whether it will be included in the analysis. This sampling scheme may not be realistic, as many factors typically influence the probability that a researcher will include a species in their study. However, the uniform species sampling scheme was initially adopted because it is mathematically convenient. Moroever, the approximation to uniform sampling scheme improves when the sampling fraction is large, *i.e.,* if more than 80% of the species are included.

In principle, we could treat the sampling probablity ρ as a random variable and estimate it from the data. However, estimating all three parameters of the sampled

constant-rate birth-death process model—the speciation rate, the extinction rate and the sampling probability—is not possible because the parameters are nonidentifiable (Stadler, 2009). Therefore, we use the empirical sampling fraction; simply the number of included species divided by the total number of known species (630 for conifers).

```
# There are 630 known conifer species
samplingFraction <- (conifers$Nnode + 1) / 630
```

We can then use this empirical sampling fraction as our sampling probability. All the likelihood functions in TESS, as described in the previous sections, have an argument called `samplingProbability`. You can incorporate incomplete sampling in any of these analyses by setting this argument to the empirical sampling probability.

How well does the *uniform* species sampling scheme perform on our simulated phylogeny? We will assess the performance of this method by estimating the joint posterior density of the diversification-rate parameters by performing an MCMC simulation. As usual, we start by specifying the prior distributions. Here, we use the constant-rate birth-death process with the *net-diversification* rate (speciation − extinction) and *turnover* rate (extinction). We choose exponential prior distributions with a mean of 1.0, which corresponds to the true value. Notice that this is essentially an ideal setting, as the true parameter values are clearly unknown for empirical analyses. Therefore, this represents a *best case scenario* for the impact of species sampling on parameter estimates.

```
prior_delta <- function(x) { dexp(x,rate=1.0,log=TRUE) }
prior_tau <- function(x) { dexp(x,rate=1.0,log=TRUE) }
priorsSampling <- c("diversifiation"=prior_delta,
                    "turnover"=prior_tau)
```

We then define the likelihood function. This is similar to the likelihood function used previously for the constant-rate birth-death process. However, in this case we will specify the sampling probability ($\rho = 0.25$) and the sampling strategy.

```
likelihoodUniform <- function(params) {

  speciation <- params[1] + params[2]
  extinction <- params[2]

  lnl <- tess.likelihood(times.diversified,
```

```
                            lambda = speciation,
                            mu = extinction,
                            samplingProbability = 0.25,
                            samplingStrategy = "uniform",
                            log = TRUE)

  return (lnl)

}
```

Now we are ready to estimate the diversification-rate parameters. The settings for the MCMC simulation are the same as those used previously (*c.f.*, Section 2.4).

```
set.seed(12345)
samplesUniform <- tess.mcmc(likelihoodFunction = likelihoodUniform,
                            priors = priorsSampling,
                            parameters = runif(2,0,1),
                            logTransforms = c(TRUE,TRUE),
                            delta = c(1,1),
                            iterations = 2000,
                            burnin = 200,
                            thinning = 10,
                            adaptive = TRUE,
                            verbose = TRUE)

## Burning-in the chain ...
## 0--------25--------50--------75--------100
## ========================================
## Finished burnin period!
##
## Running the chain ...
## 0--------25--------50--------75--------100
## ========================================
## Finished MCMC!
##
## Parameter | delta | Acceptance Probability
## ========================================
## diversifiation | 0.269 | 0.407
## turnover | 2.321 | 0.454
```

Recall that the true parameter values are $\lambda = 2.0$ and $\mu = 1.0$. This correspods to a *diversification* rate of 1.0 and a *turnover* rate of 1.0. The estimated parameters under *uniform* species sampling are

```
summary(samplesUniform)

##
## Iterations = 1:201
## Thinning interval = 1
## Number of chains = 1
## Sample size per chain = 201
##
## 1. Empirical mean and standard deviation for each variable,
##    plus standard error of the mean:
##
##                   Mean     SD Naive SE Time-series SE
## diversifiation 1.1159 0.1249 0.008812       0.008812
## turnover       0.1568 0.1566 0.011045       0.011045
##
## 2. Quantiles for each variable:
##
##                    2.5%     25%    50%    75%  97.5%
## diversifiation 0.880858 1.03295 1.1167 1.2074 1.3543
## turnover       0.006634 0.04304 0.1073 0.2095 0.5373

plot(samplesUniform)
```

Our estimate of the *diversification* rate is actually quite good. This is because most information about the diversification rate comes from the age of the phylogeny and the number of sampled species, which does not depend on the sampling scheme or the divergence times. However, the estimated *turnover* rate is quite biased. As a result, the speciation- and extinction-rate estimates are both biased. This is caused by the underestimation of the extinction rate (Höhna et al., 2011). The bias is quite severe: the true values are not even contained in the 95% credible interval (see Figure 12).

3.3 Diversified species sampling

Now we will consider diversified sampling in more detail. We assume that our study group contains m species from which we have n sampled species. Under *diversified* sampling, this means that the most recent $m - n$ speciation events

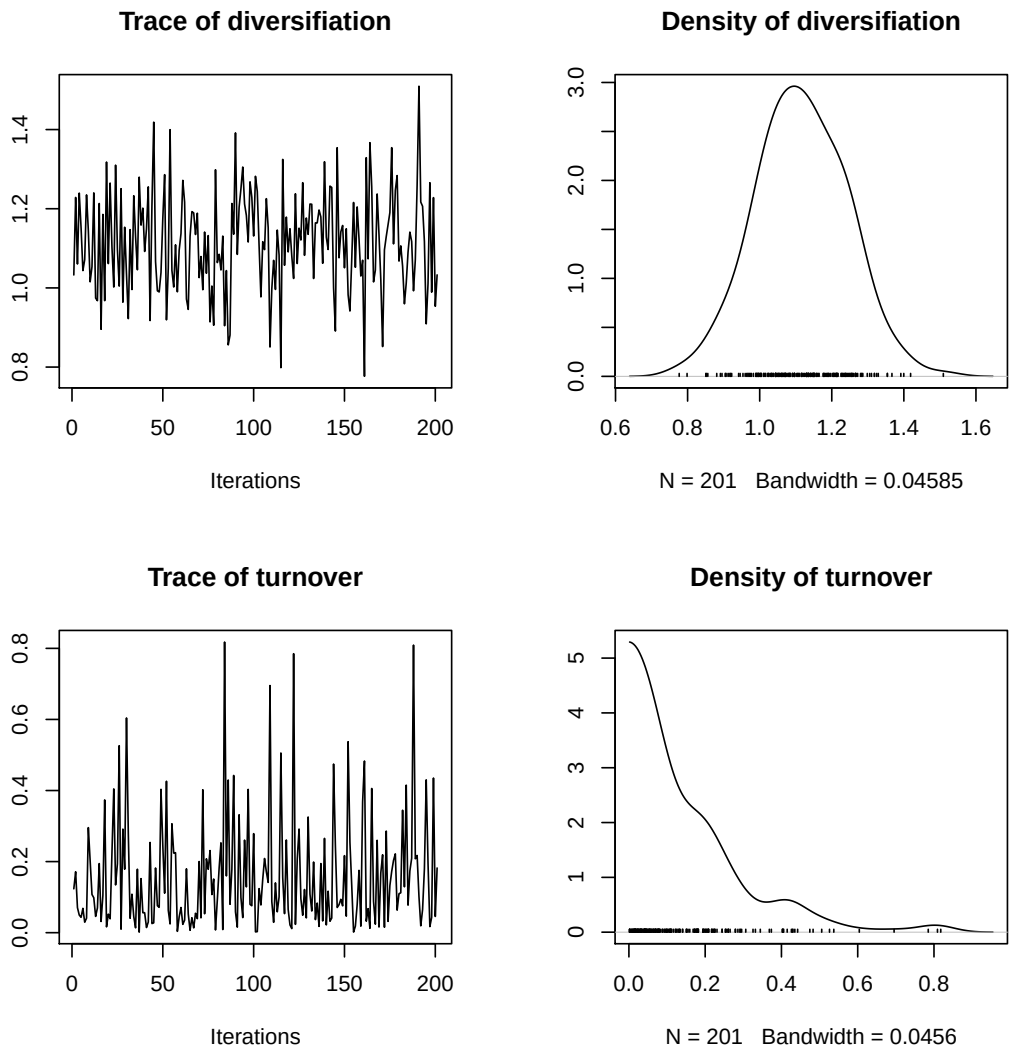

Figure 12: Trace plots (left) and marginal posterior probability densities (right) for the diversification rate (speciation - extinction) and turnover rate (extinction) under *uniform* sampling from the MCMC simulation.

have been discarded. This is the strictly mathematical interpretation of *diversified* sampling (Höhna et al., 2011; Höhna, 2014).

This sampling strategy is intended to mimic empirical datasets where species were selected to include "representatives" from some number of distinct lineages (*e.g.,* all families or all genera). This sampling implicitly maximizes species diversity and comes close to the mathematical description of *diversified* sampling. However, diversified species sampling is not a perfect mathematical description

of this sort of sampling. For example, not all "major lineages" are of the same age and size, and we may therefore sometimes include species that are recently diverged.

As we did for *uniform* species sampling, we want to test how well the method performs in estimating parameters given the simulated phylogeny. We will use the same constant-rate birth-death process with the only difference being the sampling strategy. We therefore use the same prior distributions as in the *uniform* sampling analysis. The likelihood function is adapted by changing the `samplingStrategy` to "diversified".

```
likelihoodDiversified <- function(params) {

  speciation <- params[1] + params[2]
  extinction <- params[2]

  lnl <- tess.likelihood(times.diversified,
                         lambda = speciation,
                         mu = extinction,
                         samplingProbability = 0.25,
                         samplingStrategy = "diversified",
                         log = TRUE)

  return (lnl)

}
```

Then, we perform a very short MCMC simulation to sample from the posterior distribution of the parameters.

```
samplesDiversified <- tess.mcmc(likelihoodFunction = likelihoodDiversified,
                            priors = priorsSampling,
                            parameters = runif(2,0,1),
                            logTransforms = c(TRUE,TRUE),
                            delta = c(1,1),
                            iterations = 2000,
                            burnin = 200,
                            thinning = 10,
                            adaptive = TRUE,
                            verbose = TRUE)

## Burning-in the chain ...
```

```
## 0--------25--------50--------75--------100
## ========================================
## Finished burnin period!
##
## Running the chain ...
## 0--------25--------50--------75--------100
## ========================================
## Finished MCMC!
##
## Parameter | delta | Acceptance Probability
## ========================================
## diversifiation | 0.614 | 0.372
## turnover | 0.830 | 0.407
```

Finally, our estimates of the *diversification* rate and a *turnover* rate under *diversified* species sampling are

```
summary(samplesDiversified)

##
## Iterations = 1:201
## Thinning interval = 1
## Number of chains = 1
## Sample size per chain = 201
##
## 1. Empirical mean and standard deviation for each variable,
##    plus standard error of the mean:
##
##                  Mean     SD Naive SE Time-series SE
## diversifiation 0.7966 0.3178  0.02241        0.03897
## turnover       1.2958 0.6549  0.04619        0.06287
##
## 2. Quantiles for each variable:
##
##                  2.5%    25%    50%   75% 97.5%
## diversifiation 0.2089 0.5602 0.8083 1.044 1.370
## turnover       0.1876 0.7759 1.2273 1.782 2.699

plot(samplesDiversified)
```

Here we see that the true values fall within the 95% credible interval. The mean estimate of the *diversification* rate might be slightly worse, but the transformed

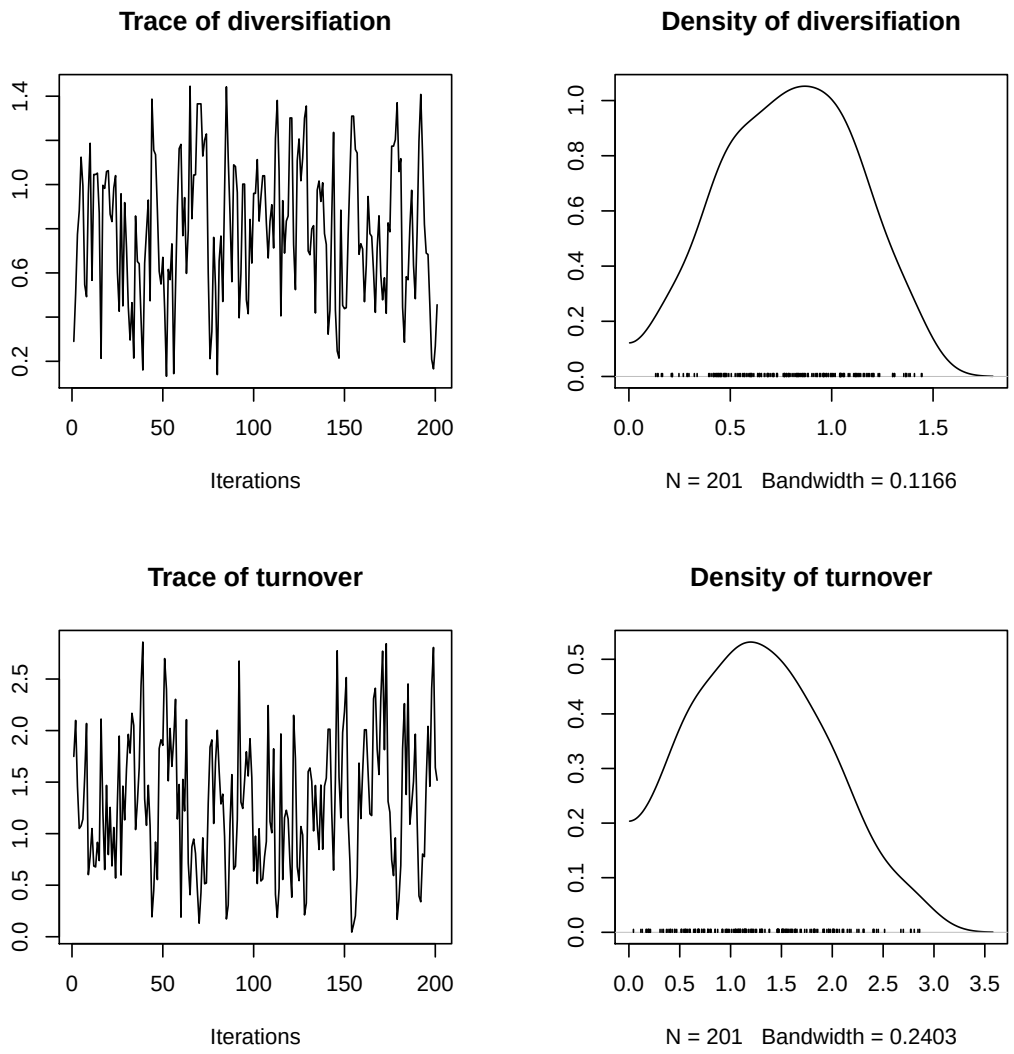

Figure 13: Trace plots (left) and marginal posterior probability densities (right) for the diversification rate (speciation - extinction) and turnover rate (extinction) under *diversified* sampling from the MCMC simulation.

speciation and extinction rate estimates are significantly better. Thus, we can conclude that only if we know the true sampling strategy, and sampling fraction, are we able to make unbiased estimates the speciation and extinction rates.

You may wish to repeat the above experiment for the case when you simulate the tree under *uniform* species sampling, and/or when you assume an incorrect sampling fraction (*i.e.,* that is either too large or too small).

For more information about incomplete species sampling, we refer the reader

to Höhna et al. (2011) and Höhna (2014).

4 Model Evaluation

Model-based inference is, by definition, based on the model. The model describes the process that gave rise to our observed data—in the present case, it describes the stochastic-branching process that gave rise to our study tree. Accordingly, if the model provides a 'poor fit' to the data—*i.e.,* provides a poor description of the process that gave rise to the observed data—then all bets are off.

The model must balance two competing criteria: (1) it must include the relevant parameters to describe important aspects of the diversification process, but; (2) it must exclude any superfluous parameters that only capture stochastic fluctuations in the data. Failure to satisfy criterion (1) will result in biased estimates of parameters; *e.g.,* speciation and extinction rates. Failure to satisfy criterion (2) will inflate the error variance of the parameter estimates; we only have a fixed amount of data at hand, such that the inclusion of additional parameters decreases the data available for estimating each parameter.

Accordingly, we must be vigilant both regarding our choice of model and also with respect to assessing our ability to perform reliable inference under the chosen model. Model evaluation entails three closely related issues. **Model selection** entails assessing the *relative fit* of our dataset to the pool of candidate models. In a Bayesian statistical framework, we compare the relative fit of candidate models based on their marginal likelihood (which measures the average fit of the candidate models to the data). **Model adequacy**—an equally important but relatively neglected issue—entails assessing the *absolute fit* of the dataset to a given model. **Model uncertainty** is related to the common (and commonly ignored) scenario when multiple candidate models provide a similar fit to the data: in this scenario, conditioning on *any single* model (even the best) will lead to biased estimates, and so model averaging is required to accommodate uncertainty in the choice of model. Below, we demonstrate how to address each of these model-evaluation issues using TESS.

4.1 Comparing models with Bayes factors

For most groups of species, several (possibly many) branching-process models of varying complexity are plausible *a priori*. We therefore need a way to objectively identify the model that balances estimation bias and inflated error variance associated with under- and over-parameterized models, respectively. Increasingly, model selection is based on *Bayes factors* (*e.g.*, Kass and Raftery, 1995; Suchard et al., 2001; Lartillot and Philippe, 2006; Xie et al., 2011; Fan et al., 2011; Baele et al., 2012, 2013). This procedure requires that we first calculate the marginal likelihood of each candidate model, and we then compare the ratio of the marginal likelihoods for the set of candidate models.

Note that interpreting Bayes factors (BF) involves a measure of subjectivity. That is, it is up to *you* to decide what BF values appropriately reflect the level of significance in the competing models. Despite the absence of an absolutely objective model-selection threshold, we can refer to the scale (outlined by Jeffreys, 1961) that provides a "rule-of-thumb" for interpreting these measures (Table 2).

Table 2: The scale for interpreting Bayes factors by Harold Jeffreys (1961).

Strength of evidence	$BF(M_0, M_1)$	$\mathbf{log}(BF(M_0, M_1))$	$\mathbf{log_{10}}(BF(M_0, M_1))$
Negative (supports M_1)	< 1	< 0	< 0
Barely worth mentioning	1 to 3.2	0 to 1.16	0 to 0.5
Substantial	3.2 to 10	1.16 to 2.3	0.5 to 1
Strong	10 to 100	2.3 to 4.6	1 to 2
Decisive	> 100	> 4.6	> 2

For a detailed description of Bayes factors see Kass and Raftery (1995)

Given two candidate models, M_0 and M_1, the Bayes-factor comparison to assess the relative fit of each model to the data, $BF(M_0, M_1)$, is:

$$BF(M_0, M_1) = \frac{\mathbb{P}(M_0 \mid \mathbf{X})}{\mathbb{P}(M_1 \mid \mathbf{X})} = \frac{\mathbb{P}(M_0)}{\mathbb{P}(M_1)} \frac{\mathbb{P}(\mathbf{X} \mid M_0)}{\mathbb{P}(\mathbf{X} \mid M_1)},$$

where $\mathbb{P}(\mathbf{X} \mid M_i)$ is the *marginal likelihood* of the data (this may be familiar to you as the denominator of Bayes Theorem, which is variously referred to as the *model evidence* or *integrated likelihood*). Formally, the marginal likelihood is the probability of the observed data ($\mathbf{X}$) under a given model (M_i) that is averaged over all possible values of the parameters of the model (θ_i) with respect to the prior density on θ_i

$$\mathbb{P}(\mathbf{X} \mid M_i) = \int \mathbb{P}(\mathbf{X} \mid \theta_i)\mathbb{P}(\theta_i)d\theta. \tag{5}$$

This makes it clear that more complex (parameter-rich) models are penalized by virtue of the associated prior: each additional parameter entails integration of the likelihood over the corresponding prior density.

Exact solutions for calculating marginal likelihoods are not avaiable for most models, which requires that we resort to numerical integration methods to approximate these values. Below, we first provide a brief description of a robust method for estimating marginal likelihoods—stepping-stone simulation (Xie et al., 2011; Fan et al., 2011)—and then demonstrate how to use the implementation of the stepping-stone algorithm in TESS to estimate the marginal likelihoods for two birth-death branching-process models.

4.1.1 Stepping-stone simulation

Recent developments provide robust methods for estimating marginal likelihoods, including stepping-stone Xie et al. (2011); Fan et al. (2011) and path-sampling estimators Lartillot and Philippe (2006); Baele et al. (2012). These algorithms are similar to the familiar MCMC algorithms, which are intended to sample from (and estimate) the joint posterior probability of the model parameters. Stepping-stone algorithms are like a series of MCMC simulations that iteratively sample from a specified number of discrete steps between the posterior and the prior probability distributions. The basic idea is to estimate the probability of the data for all points between the posterior and the prior—effectively summing the probability of the data over the prior probability of the parameters to estimate the marginal likelihood. Technically, the steps correspond to a series of *power-posteriors*: a series of numbers between 1 and 0 that are iteratively applied to the posterior. When the posterior probability is raised to the power of 1 (typically the first stepping stone), samples are drawn from the (untransformed) posterior. By contrast, when the posterior probability is raised to the power of 0 (typically the last stepping stone), samples are drawn from the prior (Figure 14).

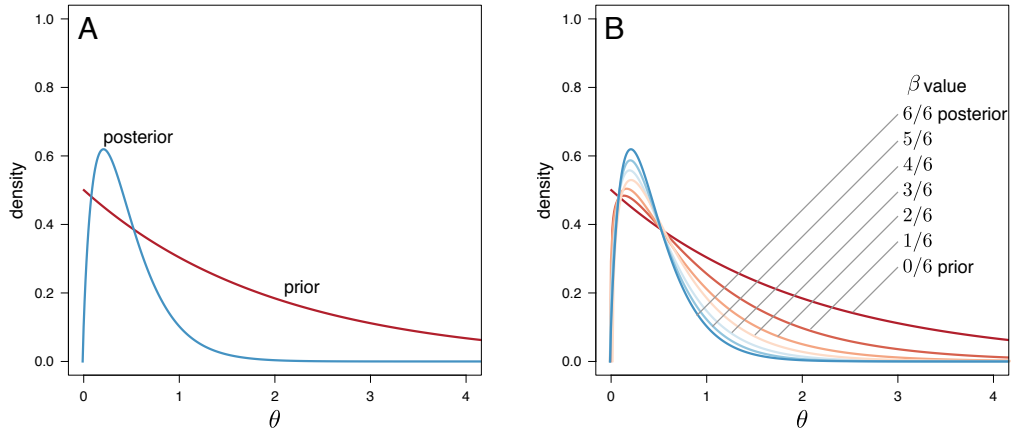

Figure 14: Estimating marginal likelihoods using stepping-stone simulation. Estimating the marginal likelihood involves integrating the likelihood of the data over the entire prior probability density for the model parameters. MCMC algorithms target the posterior probability density, which is typically concentrated in a small region of the prior probability density (A). Accordingly, standard MCMC simulation cannot provide unbiased estimates of the marginal likelihood because it will typically fail to explore most of the prior density. (B) Stepping-stone algorithms estimate the marginal likelihood by means of a series of MCMC-like simulations, where the likelihood is iterativey raised to a series of powers, effectivey forcing the simulation to more fully explore the prior density of the model parameters. Six uniformly spaced stones span the posterior, where the power posterior is $\beta = 6/6 = 1$, to the prior, where the power posterior is $\beta = 0/6 = 0$.

To perform a stepping-stone simulation, we need to specify (1) the *number* of stepping stones (power posteriors) that we will use to traverse the path between the posterior and the prior (*e.g.,* we specify 50 or 100 stones), (2) the *spacing* of the stones between the posterior and prior (*e.g.,* we may specify that the stones are distributed according to a beta distribution), (3) the number (and thinning) of samples to be drawn from each stepping stone, and (4) the direction we will travel (*i.e.,* from the posterior to the prior or *vice versa*).

This method computes a vector of powers from a beta distribution, then executes an MCMC run for each power step while raising the likelihood to that power. As implemented in TESS, the vector of powers starts with 1, initially sampling the likelihood close to the posterior, and incrementally sampling closer and closer to the prior as the simulation progresses across the stepping stones.

4.1.2 Estimating marginal likelihoods of birth-death models

To estimate the marginal likelihoods of the branching-process models, we will again make use of the corresponding likelihood functions that we defined in the previous section. However, rather than using the `tess.mcmc` function to sample from (and so estimate) the posterior distribution of model parameters, we will use the `tess.steppingStoneSampling` function to estimate the marginal likelihood of the data under the various models. The commands to execute this function for each of the three branching-process models are as follows:

```
set.seed(12345)
marginalLikelihoodConstBD <- tess.steppingStoneSampling(
                  likelihoodFunction = likelihoodConstBD,
                  priors = priorsConstBD,
                  parameters = runif(2,0,1),
                  logTransforms = c(TRUE,TRUE),
                  iterations = 200,
                  burnin = 200,
                  K = 10)

marginalLikelihoodDecrBD <- tess.steppingStoneSampling(
                  likelihoodFunction = likelihoodDecrBD,
                  priors = priorsDecrBD,
                  parameters = runif(3,0,1),
                  logTransforms = c(TRUE,TRUE,TRUE),
                  iterations = 200,
                  burnin = 200,
                  K = 10)
```

```
marginalLikelihoodEpisodicBD <- tess.steppingStoneSampling(
                likelihoodFunction = likelihoodEpisodicBD,
                priors = priorsEpisodicBD,
                parameters = runif(4,0,1),
                logTransforms = c(TRUE,TRUE,TRUE,TRUE),
                iterations = 200,
                burnin = 200,
                K = 10)

marginalLikelihoodMassExtinctionBD <- tess.steppingStoneSampling(
                likelihoodFunction = likelihoodMassExtinctionBD,
                priors = priorsMassExtinctionBD,
                parameters = c(runif(2,0,1),max(times)*3/4),
                logTransforms = c(TRUE,TRUE,FALSE),
                iterations = 200,
                burnin = 200,
                K = 10)
```

We can now use the estimated marginal likelihoods to perform Bayes factor comparisons of these three candidate branching-process models.

```
# First, construct a vector of the marginal likelhoods named by the
# model to which they refer.
candidateModels <- c("ConstBD"=marginalLikelihoodConstBD,
                     "DecrBD"=marginalLikelihoodDecrBD,
                     "EpisodicBD"=marginalLikelihoodEpisodicBD,
                     "MassExtinctionBD"=marginalLikelihoodMassExtinctionBD)

# Make all possible combinations of the models.
marginalLikelihoodGrid <- expand.grid(M0=names(candidateModels),
                                      M1=names(candidateModels))

# Add a column that is the 2 ln BF for each pair of models.
marginalLikelihoodGrid$BF <- 2 * (candidateModels[marginalLikelihoodGrid$M0] -
                                  candidateModels[marginalLikelihoodGrid$M1])

# Sort the comparisons by their 2 ln BF in descending order.
marginalLikelihoodGrid <- marginalLikelihoodGrid[order(marginalLikelihoodGrid$BF,
                                                       decreasing=TRUE),]
```

```
marginalLikelihoodGrid

##                   M0                M1         BF
## 13          ConstBD MassExtinctionBD  31.816729
## 15       EpisodicBD MassExtinctionBD  27.020702
## 14           DecrBD MassExtinctionBD  23.956544
## 5           ConstBD           DecrBD   7.860185
## 9           ConstBD       EpisodicBD   4.796027
## 7        EpisodicBD           DecrBD   3.064158
## 1           ConstBD          ConstBD   0.000000
## 6            DecrBD           DecrBD   0.000000
## 11       EpisodicBD       EpisodicBD   0.000000
## 16 MassExtinctionBD MassExtinctionBD   0.000000
## 10           DecrBD       EpisodicBD  -3.064158
## 3        EpisodicBD          ConstBD  -4.796027
## 2            DecrBD          ConstBD  -7.860185
## 8  MassExtinctionBD           DecrBD -23.956544
## 12 MassExtinctionBD       EpisodicBD -27.020702
## 4  MassExtinctionBD          ConstBD -31.816729
```

If we compare these computed Bayes factor values to the thresholds in Table 2, we see that there is decisive support for the constant-rate model (*e.g.,* this model is decisively preferred over either of the variable-rate models; BF $\gg$ 4.6). Furthermore, we see that the decreasing-rate model is decisively preferred over the episodic model.

4.2 Assessing model adequacy with posterior predictive simulation

Bayes factors, dicussed in the previous section, allow us to assess the *relative* fit of two or more competing models to a given dataset. However, even the very best of the competing models may nevertheless be woefully inadequate in an *absolute* sense. Fortunately, we can assess the absolute fit of a candidate model to a given dataset using *posterior-predictive simulation*. The basic premise of this approach is as folows: if the model under consideration provides an adequate description of the process that gave rise to our observed dataset, then we should be able to use that model generate new datasets that are in some sense 'similar' to our dataset.

4.2.1 Posterior-predictive simulation

Posterior-predictive simulation involves six main steps:

1. We first calculate a summary statistic for our observed dataset. This is intended to capture—in a single number—a relevant feature of our dataset. For models of lineage diversification, for example, we might use number of species in the tree or the γ-statistic (Pybus and Harvey, 2000) as our summary statistic.

2. We then estimate parameters of the candidate model from our oberved dataset. This simply involves performing an MCMC simulation to estimate the posterior probability distribution of the candidate model parameters.

3. Next, we specify parameters of the candidate model by drawing values from the inferred joint posterior probability distributions. For example, we would parameterize the diversification model under consideration by drawing rate parameters from the joint posterior densities that we inferred from the study tree.

4. We then use this parameterized model to simulate a tree, and calculate the summary statistic for the resulting tree.

5. We repeat steps $3-4$ many times to generate a distribution of the summary statistic. This is the distribution that is predicted by simulating datasets under the candidate model that has been parameterized using posterior estimates on the observed dataset.

6. Finally, we compare the summary statistic calculated for the observed dataset to the posterior-predictive distribution. If the candidate model provides an adequate description of the process that gave rise to the original dataset, then the statistic for the observed dataset will fall near the center of the simulated distribution. Otherwise, the statistic from the observed data will fall near the tails of the null distribution, indicating that the model cannot be used to predict future data that look like the observed dataset.

We can formalize the relative position of the statistic for the observed data to the posterior-predictive distribution by calculating the *posterior-predictive p-value*. To do so, we simply sum the number of simulated summary statistics that are greater than or equal to the observed value, and divide this by the number of simulated values.

4.2.2 Assessing the adequacy of branching-process models

To assess the adequacy of branching-process models in TESS, we will use the number of species and the γ-statistic (Pybus and Harvey, 2000) as our example test statistics. We will demonstrate how to perform posterior-predictive simulation to assess the absolute fit of the constant-rate birth-death process model to the conifer dataset. Note that we have already estimated the joint posterior probability distribution of this model using MCMC (described in Section 2.4.1, above).

We will condition our simulate trees on the age of conifer phylogeny.

```
tmrca <- max( times )
```

We first define the function that will perform the simulation. This is analogous to the specification of the likelihood function, which means that it is possible to perform posterior-predictive simulation under any birth-death model.

```
# The simulation function
simConstBD <- function(params) {

  # Same model as above.
  speciation <- params[1] + params[2]
  extinction <- params[2]

  # We need trees with at least three tips for the
  # gamma-statistic.
  repeat {
    tree <- tess.sim.age(n = 1,
                         age = tmrca,
                         lambda = speciation,
                         mu = extinction,
                         samplingProbability = 1.0,
                         MRCA = TRUE)[[1]]
    if (tree$Nnode > 1) break
  }
  return (tree)
}
```

Note that the simulation function needs to return a single tree. The next step is to simulate trees by sampling parameter values (for the speciation and extinction rate) from the corresponding posterior probability distributions that we inferred from the conifer tree (these are referred to as the 'posterior-predictive samples').

```
# simulate trees from the posterior-predictive distribution
treesConstBD <- tess.PosteriorPrediction(simConstBD,samplesConstBD)
```

We will specify the number of species as the summary statistic:

```
# compute the number of species in each simulate tree
numTaxaConstBD <- c()
for (i in 1:length(treesConstBD)){
  numTaxaConstBD[i] <- treesConstBD[[i]]$Nnode + 1
}
```

We then compute the posterior-predictive quantiles for the number of species and plot the posterior-predictive distribution and quantiles. Finally, we compare the observed number of species to the posterior-predictive distribution.

```
# Compute the 95% posterior-predictive interval of the
# number of taxa.
numTaxaPPDI <- quantile(numTaxaConstBD,prob=c(0.025,0.975))

# Plot the posterior-predictive distribution with the
# quantiles. Then, compare it to the observed number
# of species, x.
plot(density(numTaxaConstBD),main="Number of taxa",xlab="",
    ylab="Posterior Predictive Density",lwd=2)

abline(v=numTaxaPPDI,lty=2,col="gray",lwd=2)

points(conifers$Nnode+1,0,pch="x")
```

We can plot the posterior-predictive distribution of the lineage-accumulation curves (*i.e.*, the LTT plots for the simulated trees), and compare this predictive distribution to the LTT plot for the observed tree.

```
ltt.plot(treesConstBD[[1]],backward=FALSE,col="gray",log="y",
         ylim=c(1,max(numTaxaConstBD)),main="LTT-plot")

for (i in 2:min(100,length(treesConstBD))) ltt.lines(treesConstBD[[i]],
                           backward=FALSE, col="gray")

ltt.lines(conifers,backward=FALSE,lwd=3)
```

Additional to using species number as the summary statistic, we can also use the gamma statistic. To do so, we simply compute the value of the gamma statistic for the observed tree, and then compare it to the posterior-predictive distribution of the gamma statistic (*i.e.,* the distribution of γ-statistics computed from the simulated trees).

```
# Compute the observed gamma statistic.
observedGamma <- gammaStat(conifers)

# Perform the posterior predictive test, and compute
# the 95% posterior predictive interval.
ppt <- tess.PosteriorPredictiveTest(treesConstBD,conifers,
                                    gammaStat)
gammaPPDI <- quantile(ppt[[1]],prob=c(0.025,0.975))

# Compare the observed statistic to the posterior
# predictive density.
plot(density(ppt[[1]]),main="Gamma Statistic",xlab="",
                       ylab="Posterior Predictive Density",lwd=2)
abline(v=gammaPPDI,lty=2,col="gray",lwd=2)
points(observedGamma,0,pch="x")
```

Encouragingly, the observed values for both summary statistics—the number of species and the gamma statistic—fall near the center of their respective posterior-predictive distributions (see Figure 15 A and C, respectively). This means that the model under consideration—the constant-rate birth-death model—can be used to simulate trees that look like our conifer study tree, indicating that it provides a good absolute fit to our dataset. (This conclusion is consistent with the Bayes factor comparisons that we performed in Section 4.1.2, which indicated that the constant-rate birth-death model provided the best relative fit to the conifer tree.) Conversely, if we had found that the observed values for the summary statistics fell outside the 95% credible intervals of the posterior-predictive distributions, we would conclude that the contant-rate birth-death model cannot be used to predict trees that look like our conifer stuy tree.

We could further quantify the relative position of the observed summary statistic within the posterior-predictive distribution by calculatig the posterior-predictive p-value as follows:

```
mean(ppt[[1]] >= observedGamma)

## [1] 0.1854305
```

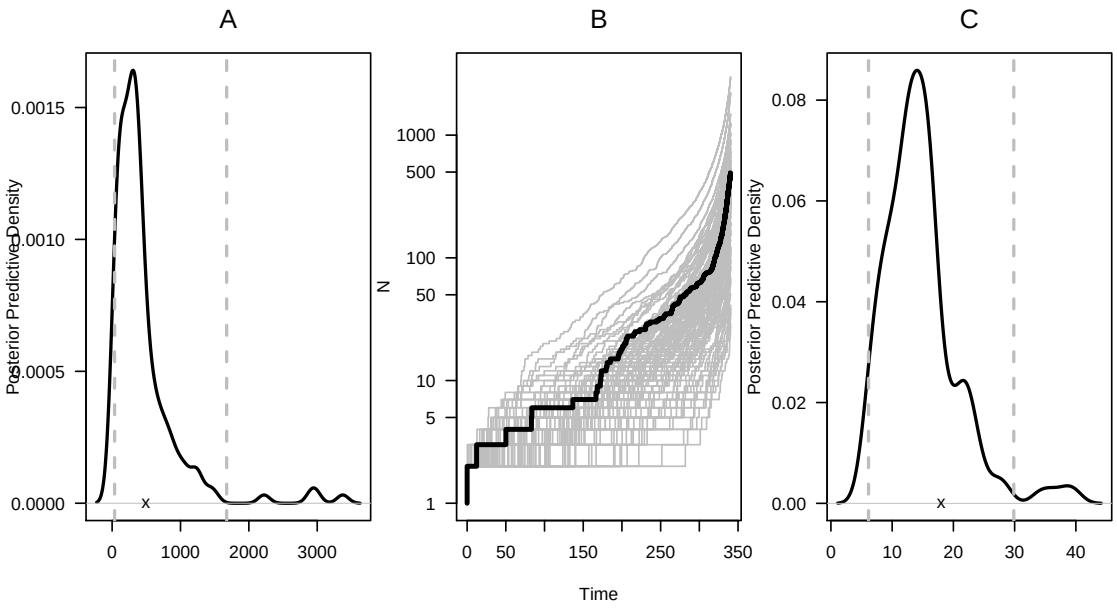

Figure 15: Assessing the absolute fit of the conifer tree to the constant-rate birth-death model using posterior-predictive simulation. (A) The posterior-predictive distribution for the number of species; the dashed gray lines indicate the 95% credible interval, and the 'x' indicates the location of the observed species number. (B) LTT plots for the simulated trees (gray) and for the conifer study tree (black). (C) The posterior-predictive distribution for the gamma statistic; the dashed gray lines indicate the 95% credible interval, and the 'x' indicates the location of the value of the gamma statistic calculated for the conifer tree.

These posterior-predictive p-values can be used to compare the absolute fit of two or more models to a given dataset.

4.3 Model averaging with CoMET

In the preceding sections, we assessed the relative and absolute fit of birth-death process models using Bayes factors and posterior-predictive simulation. These methods are computationally intensive and so are only practical when the set of candidate branching-process models is quite small; however, the number of total birth-death models is vast. Consider, for example, the episodic birth-death process models: there is an infinite array of nested models that differ in the *number* of events, and for a specific number of events, there is an infite number of *times* at which those events could occur. Clearly, the vast space of possible branching-process models precludes their exhaustive pairwise comparison using Bayes factors. This issue may be addressed by means of *Bayesian model-averaging* approaches that treat the model as a random variable, such that inferences are averaged over all birth-death process models.

The `CoMET` method performs Bayesian model averaging over all possible episodically varying birth-death processes with explicitly modeled mass-extinction events. Briefly, this method treats the number of specation-rate shifts, extinction-rate shifts, and mass-extinction events—as well as the parameters associated with these events—as random variables, and estimates their joint posterior distribution using reversible jump Markov chain Monte Carlo. To perform a full `CoMET` analysis, we must therefore specify values for the following quantities: (1) the expected number of speciation-rate shifts, extinction-rate shifts, and mass-extinction events ($\lambda_{\mathbb{B}}, \lambda_{\mathbb{D}}, \lambda_{\mathbb{M}}$, respectively); (2) the hyperpriors describing the speciation and extinction rates ($\mu_{\mathbb{B}}, \sigma_{\mathbb{B}}$ and $\mu_{\mathbb{D}}, \sigma_{\mathbb{D}}$, respectively); and (3) the hyperpriors describing the mass-extinction survival probability (α, β). The settings and accompanying arguments are summarized in Table 3. Note that the method currently assumes that $\lambda_{\mathbb{B}} = \lambda_{\mathbb{D}}$, although the actual number, timing and magnitude of speciation- and extinction-rate shifts are independent of each other. For the full details of this method, and a more complete description of the model parameters, see May et al. (2015).

4.3.1 Specifying hyperpiors *a priori*

Before we can run a `CoMET` analysis, we must specify prior distributions for each of the parameters in the model. We will begin with the most complex model in this section, and then show how to specify special cases of the full model in later sections.

We will start by specifying the prior distributions for the expected number of speciation- and extinction-rate shifts, $\lambda_{\mathbb{B}} = \lambda_{\mathbb{D}}$, and mass-extinction events $\lambda_{\mathbb{M}}$. These values should reflect the number of events you expect to have impacted the study tree based on external/prior information, such as paleontological data. The

Table 3: Settings for the `CoMET` model. The left column lists the priors and hyperpriors used by the `CoMET` model. The middle column lists the associated arguments for use with the `tess.analysis` command. The right column lists the interpretation of the prior or hyperprior.

Prior	Argument	Interpretation
$\lambda_{\mathbb{B}}$	`numExpectedRateChanges`	Expected number of speciation-rate shifts
$\lambda_{\mathbb{D}}$	`numExpectedRateChanges`	Expected number of extinction-rate shifts
$\lambda_{\mathbb{M}}$	`numExpectedMassExtinctions`	Expected number of mass-extinction events
$\mu_{\mathbb{B}}$	`speciationRatePriorMean`	Mean speciation rate
$\sigma_{\mathbb{B}}$	`speciationRatePriorStDev`	Standard deviation of the speciation rate
$\mu_{\mathbb{D}}$	`extinctionRatePriorMean`	Mean extinction rate
$\sigma_{\mathbb{D}}$	`extinctionRatePriorStDev`	Standard deviation of the extinction rate
α	`pMassExtinctionPriorShape1`	Shape parameter of the expected survival probability
β	`pMassExtinctionPriorShape2`	Shape parameter of the expected survival probability

prior expected number of mass-extinction events should reflect information from the fossil record. For example, the conifer tree is approximately 350 million years old, so we suspect that this group may have been exposed to three mass-extinction events: the Permo-Triassic event, the Triassic-Jurassic event, and the Cretaceous-Paleogene event. Accordingly, we specify $\lambda_{\mathbb{M}}$ to reflect the fossil record as follows:

```
numExpectedMassExtinctions <- 3
```

By contrast, it can be difficult to specify an empirically informed prior on the number of diversification-rate shifts; for this reason, we recommend doing many analyses with various values of $\lambda_{\mathbb{B}} = \lambda_{\mathbb{D}}$ to assess the sensitivity of conclusions to this prior. For this example, we only use a single value:

```
numExpectedRateChanges <- log(2)
```

Next, we consider the prior densities for the diversification-rate parameters themselves. Speciation and extinction rates must be greater than 0, so we will

use lognormal prior densities to reflect this fact. We must therefore specify the mean and standard deviation hyperparameters of each lognormal distribution. We start by specifying the hyperparameters in *real space*; *i.e.,* the mean and standard deviation of the actual speciation and extinction rates.

```
# Specify the mean and standard deviation of the lognormal
# prior on the speciation rate in real space
speciationPriorMu <- 0.2
speciationPriorSigma <- 0.5

# Specify the mean and standard deviation of the lognormal
# prior on the extinction rate in real space
extinctionPriorMu <- 0.15
extinctionPriorSigma <- 0.5
```

We then transform the hyperparameters to reflect the mean and standard of the log-transformed speciation and extinction rates. These are the μ and σ parameters of the lognormal priors on speciation and extinction rates.

```
# Transform the priors on the speciation rate into log space.
speciationRatePriorMean <- log((speciationPriorMu^2)
                           /sqrt(speciationPriorSigma^2+
                           speciationPriorMu^2))

speciationRatePriorStDev <- sqrt( log(1+speciationPriorSigma^2
                            /(speciationPriorMu^2)))

# Transform the priors on the extinction rate into log space.
extinctionRatePriorMean <- log((extinctionPriorMu^2)
                           /sqrt(extinctionPriorSigma^2+
                           extinctionPriorMu^2))

extinctionRatePriorStDev <- sqrt( log(1+extinctionPriorSigma^2
                            /(extinctionPriorMu^2)))
```

Finally, we need to specify the prior density on the survival probability of a mass-extinction event. This value reflects the probability that a lineage survives a particular mass-extinction event, and therefore must be between 0 (each lineage will always go extinct) and 1 (each lineage will always survive). A convenient prior density for this parameter is the beta distribution, which has two shape parameters, α and β. We note that the ability to distinguish mass-extinction events from more

prosaic temporal variation in speciation and extinction rates depends critically on this density: high survival probabilities and relatively flat prior densities on the survival probability will greatly decrease our power to disentangle mass-extinction events from temporal variation in diversification rates. Again, we use the fossil record to inform our prior density on this parameter.

We begin parameterizing this distribution by specifying the *expected survival probability*. The fossil record suggests that between 70% and 95% of species diversity was lost during the relevant mass-extinction events, therefore the expected survival probability should be quite low. In this example, we assume *a priori* a survival probability of 10%.

```
expectedSurvivalProbability <- 0.1
```

Using the expected survival probability, we compute the α and β parameters of the beta distribution. We set the value of β to be large, which focuses the prior density more tightly around the expected survival probability. Then, we compute α based on the expected survival probability and the specified β value.

```
pMassExtinctionPriorShape2 <- 18
pMassExtinctionPriorShape1 <- - pMassExtinctionPriorShape2 *
                                expectedSurvivalProbability /
                                (expectedSurvivalProbability - 1)
```

We can inspect this beta distribution to confirm that it accurately reflects our prior belief regarding the survival probability (Figure 16).

```
# Plot the density function of our beta distribution.
curve(dbeta(x,shape1=pMassExtinctionPriorShape1,
            shape2=pMassExtinctionPriorShape2),n=1001,
            xlab='survival probability',ylab='density',las=1)

# Plot the 95% prior interval on the survival probability.
abline(v = qbeta(c(0.025,0.975),shape1=pMassExtinctionPriorShape1,
                 shape2=pMassExtinctionPriorShape2),lty=2)
```

This beta distribution seems to reflect our prior belief that the survival probability is expected to be $\sim 10\%$, but can range from about 1% to about 25%.

Having specified prior distributions for all of the parameters of the `CoMET` model, we can now perform an analysis.

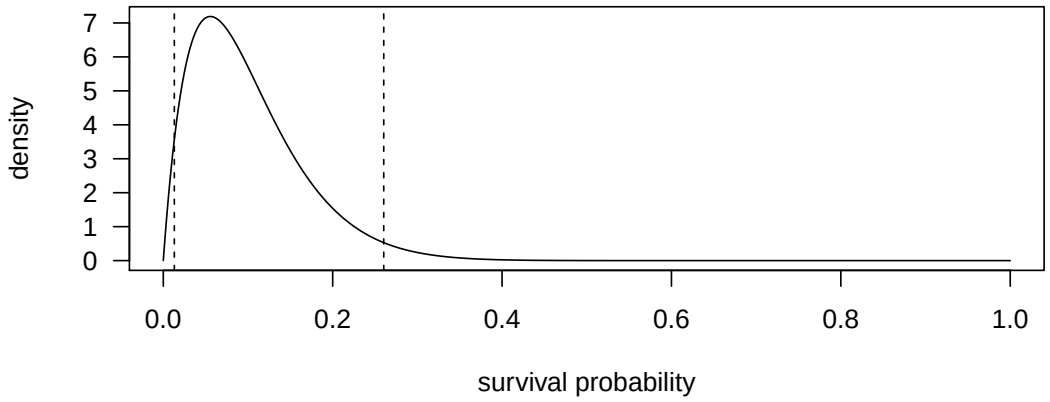

Figure 16: Our prior density on the survival probability of a mass-extinction event.

```
set.seed(12345)
tess.analysis(conifers,
              empiricalHyperPriors = FALSE,
              initialSpeciationRate = speciationPriorMu,
              speciationRatePriorMean = speciationRatePriorMean,
              speciationRatePriorStDev = speciationRatePriorStDev,
              initialExtinctionRate = extinctionPriorMu,
              extinctionRatePriorMean = extinctionRatePriorMean,
              extinctionRatePriorStDev = extinctionRatePriorStDev,
              samplingProbability = samplingFraction,
              numExpectedRateChanges = numExpectedRateChanges,
              numExpectedMassExtinctions = numExpectedMassExtinctions,
              pMassExtinctionPriorShape1 = pMassExtinctionPriorShape1,
              pMassExtinctionPriorShape2 = pMassExtinctionPriorShape2,
              MAX_ITERATIONS = 100000,
              dir = "tess_analysis")

##
## Performing CoMET analysis.
##
## Burning-in the chain ...
## 0--------25--------50--------75--------100
## ========================================
```

```
##
## Running the chain ...
## 0--------25--------50--------75--------100
## ========================================
```

Now, we process the output using the function `tess.process.output` and visualize the results using `tess.plot.output` (Figure 17).

```
output <- tess.process.output("tess_analysis",
                numExpectedRateChanges = numExpectedRateChanges,
                numExpectedMassExtinctions = numExpectedMassExtinctions)

layout.mat <- matrix(1:6,nrow=3,ncol=2,byrow=TRUE)
layout(layout.mat)
tess.plot.output(output,
                 fig.types = c("speciation rates",
                               "speciation shift times",
                               "extinction rates",
                               "extinction shift times",
                               "mass extinction Bayes factors",
                               "mass extinction times"),
                 las=2)
```

There appears to be nearly decisive support for a mass-extinction event about 27 million years ago ($2 \ln \mathrm{BF} \approx 10$), as well as strong support for a mass-extinction event about 173 million year ago ($2 \ln \mathrm{BF} \approx 6$). Additionally, there is support for an extinction-rate shift near the present; however, we caution against interpreting these rate shifts based on a single analysis and recommend assessing the sensitivity of this conclusion to different priors on the expected number of diversification-rate shifts.

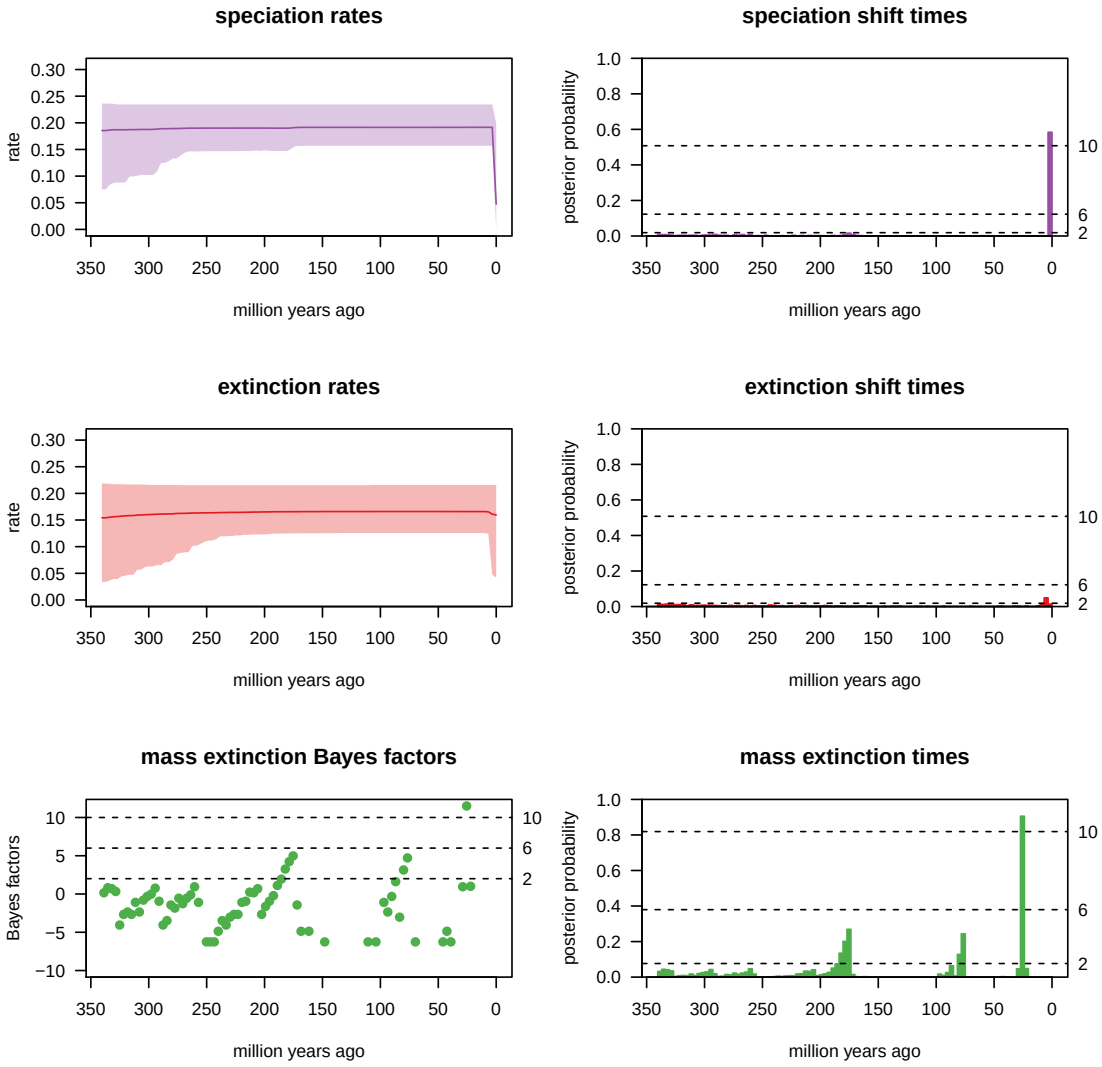

Figure 17: Visualizing the results of a `CoMET` analysis when diversification hyperpriors are specified *a priori*.

4.3.2 Empirical hyperpriors

As mentioned above, it can be difficult to specify the prior distributions for the speciation- and extinction-rate parameters. TESS implements an automatic empirical hyperprior procedure which performs an initial Bayesian MCMC analysis under a constant-rate birth-death process model to determine reasonable values for the hyperparameters of the diversification priors. To perform a hyperprior analysis, we simply have to set `empiricalHyperPriors = TRUE`. Additionally, we may omit the parameters of the lognormal distributions, since they will automatically be estimated from the data.

```
set.seed(12345)
tess.analysis(conifers,
              empiricalHyperPriors = TRUE,
              samplingProbability = samplingFraction,
              numExpectedRateChanges = numExpectedRateChanges,
              numExpectedMassExtinctions = numExpectedMassExtinctions,
              pMassExtinctionPriorShape1 = pMassExtinctionPriorShape1,
              pMassExtinctionPriorShape2 = pMassExtinctionPriorShape2,
              MAX_ITERATIONS = 100000,
              dir = "tess_analysis_empirical_hyperpriors")

## Estimating empirical hyper-parameters.
##
## Burning-in the chain ...
## 0--------25--------50--------75--------100
## ========================================
## Finished burnin period!
##
## Running the chain ...
## 0--------25--------50--------75--------100
## ========================================
## Finished MCMC!
##
## Parameter | delta | Acceptance Probability
## ========================================
## diversification | 0.192 | 0.406
## turnover | 0.007 | 0.487
##
## Performing CoMET analysis.
##
```

```
## Burning-in the chain ...
## 0--------25--------50--------75--------100
## ========================================
##
## Running the chain ...
## 0--------25--------50--------75--------100
## ========================================
```

As before, we process and visualize the output using `tess.process.output` and `tess.plot.output` (Figure 18).

```
output <- tess.process.output("tess_analysis_empirical_hyperpriors",
                numExpectedRateChanges = numExpectedRateChanges,
                numExpectedMassExtinctions = numExpectedMassExtinctions)

layout.mat <- matrix(1:6,nrow=3,ncol=2,byrow=TRUE)
layout(layout.mat)
tess.plot.output(output,
                 fig.types = c("speciation rates",
                               "speciation shift times",
                               "extinction rates",
                               "extinction shift times",
                               "mass extinction Bayes factors",
                               "mass extinction times"),
                 las=2)
```

Interestingly, these results are quite similar to those from our *a priori* analysis (Figure 17).

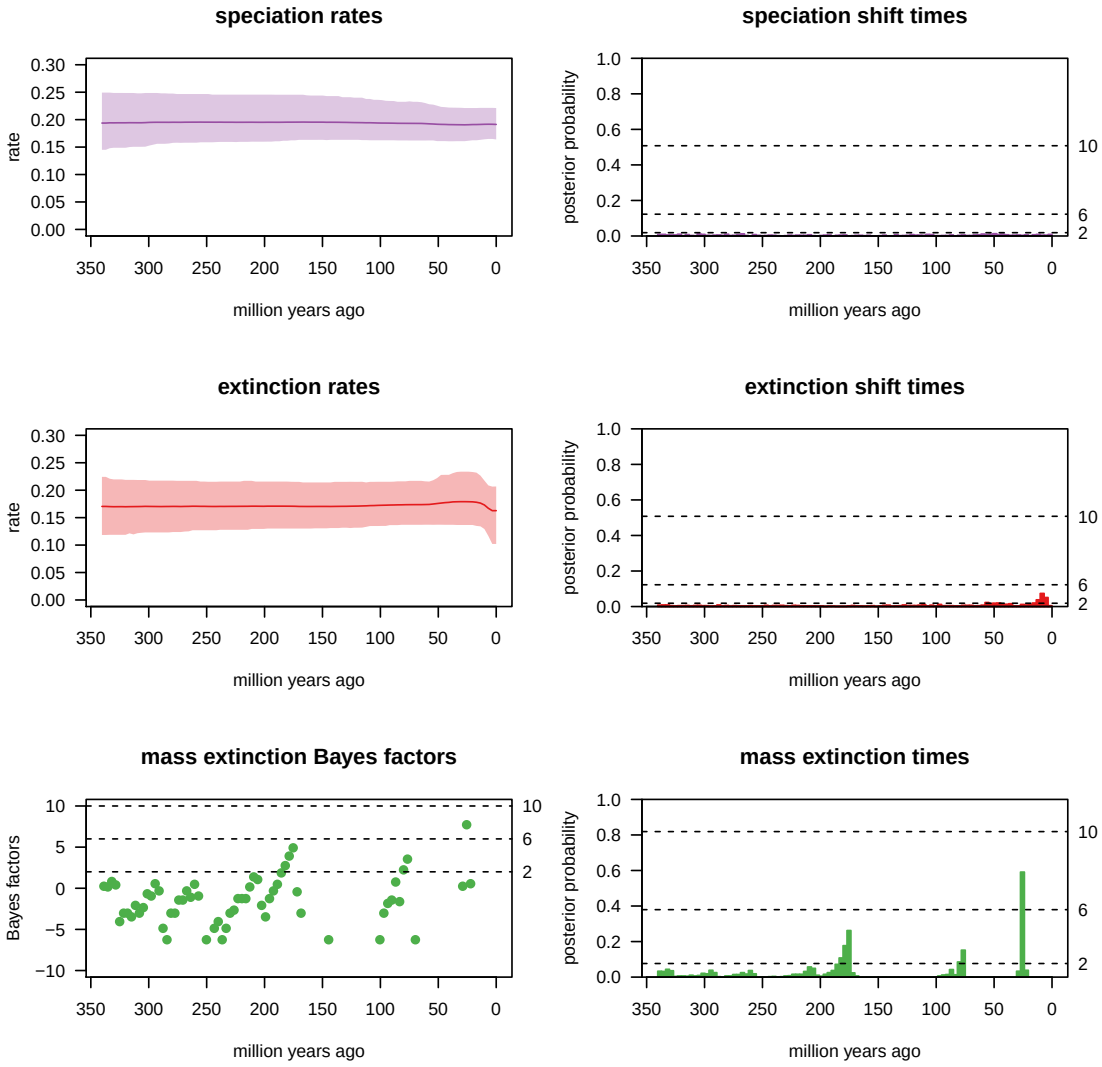

Figure 18: Visualizing the results of a `CoMET` analysis with empirically estimated diversification hyperpriors.

4.3.3 Without diversification-rate shifts

Based on the results of the Bayes factor comparisons (Section 4.1.2) and posterior-predictive tests (Section 4.2.2)—which found support for a constant-rate birth-death process—we may also be interested in performing a `CoMET` analysis without diversification-rate shifts (*i.e.,* where speciation and extinction rates are constant through time). To do so, we use the argument `estimateNumberRateChanges = FALSE`.

```
set.seed(12345)
tess.analysis(conifers,
              empiricalHyperPriors = TRUE,
              samplingProbability = samplingFraction,
              estimateNumberRateChanges = FALSE,
              numExpectedMassExtinctions = numExpectedMassExtinctions,
              pMassExtinctionPriorShape1 = pMassExtinctionPriorShape1,
              pMassExtinctionPriorShape2 = pMassExtinctionPriorShape2,
              MAX_ITERATIONS = 100000,
              dir = "tess_analysis_empirical_no_rateshifts")

## Estimating empirical hyper-parameters.
##
## Burning-in the chain ...
## 0--------25--------50--------75--------100
## ========================================
## Finished burnin period!
##
## Running the chain ...
## 0--------25--------50--------75--------100
## ========================================
## Finished MCMC!
##
## Parameter | delta | Acceptance Probability
## ========================================
## diversification | 0.192 | 0.406
## turnover | 0.007 | 0.487
##
## Performing CoMET analysis.
##
## Burning-in the chain ...
## 0--------25--------50--------75--------100
```

```
## ========================================
##
## Running the chain ...
## 0--------25--------50--------75--------100
## ========================================
```

We visualize the output as before. However, since there are no diversification-rate shifts, we only plot estimates related to mass-extinction events (Figure 19).

```
output <- tess.process.output("tess_analysis_empirical_no_rateshifts",
                numExpectedRateChanges = numExpectedRateChanges,
                numExpectedMassExtinctions = numExpectedMassExtinctions)

layout.mat <- matrix(1:2,nrow=2,ncol=1)
layout(layout.mat)
tess.plot.output(output,
                 fig.types = c("mass extinction Bayes factors",
                               "mass extinction times"),
                 las=2)
```

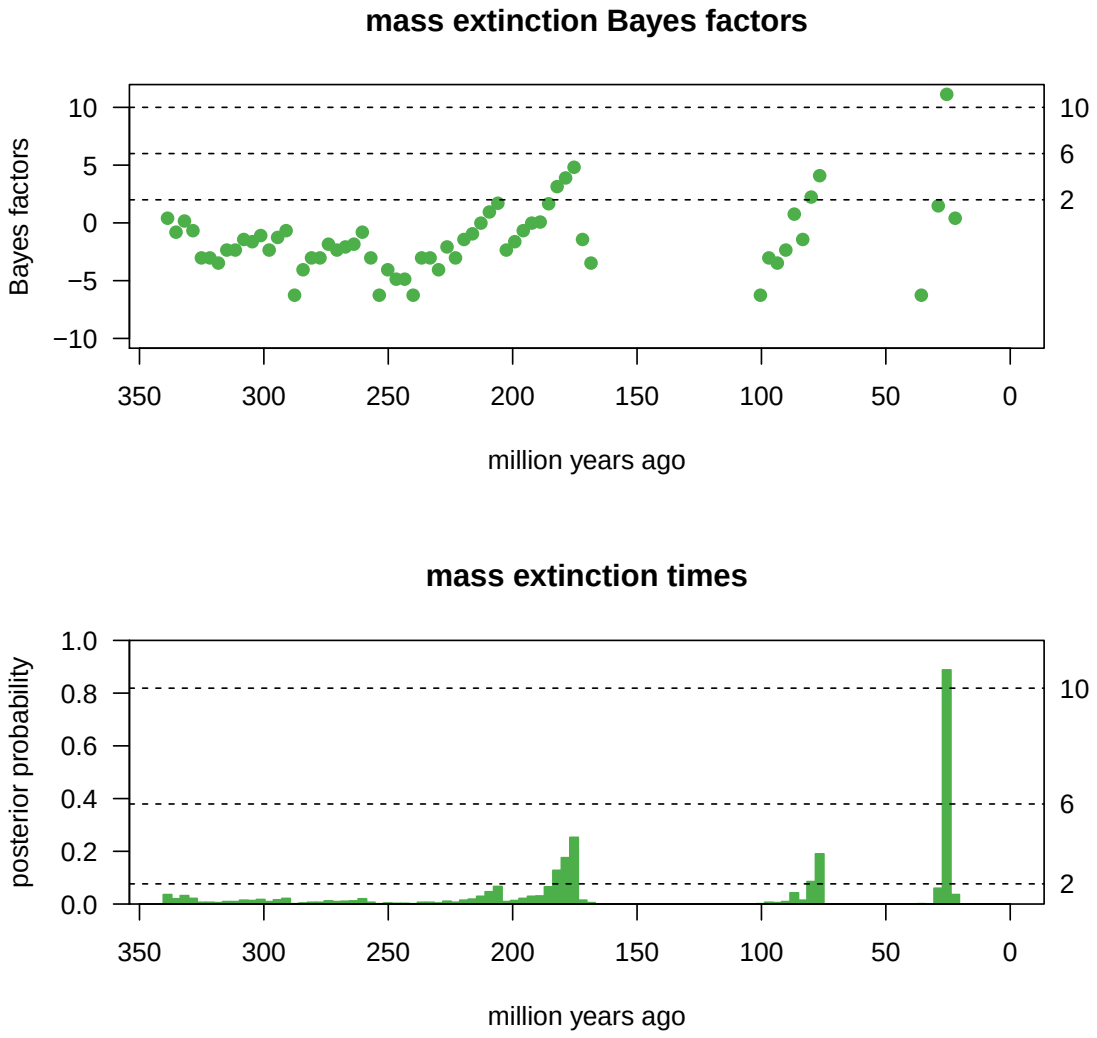

Figure 19: Visualizing the results of a `CoMET` analysis with empirically estimated diversification hyperpriors *and without* diversification rate-shifts.

4.3.4 Without mass-extinction events

Finally, we can also perform a `CoMET` analysis where mass-extinction events are disallowed. We do this with the argument `estimateNumberMassExtinctions = FALSE`. Because mass-extinction events are precluded, we can omit the corresponding parameters for survival probability, `pMassExtinctionPriorShape1` and `pMassExtinctionPriorShape2`.

```
set.seed(12345)
tess.analysis(conifers,
              empiricalHyperPriors = TRUE,
              samplingProbability = samplingFraction,
              estimateNumberMassExtinctions = FALSE,
              MAX_ITERATIONS = 100000,
              dir = "tess_analysis_empirical_no_mass_extinctions")
```

```
## Estimating empirical hyper-parameters.
##
## Burning-in the chain ...
## 0--------25--------50--------75--------100
## ========================================
## Finished burnin period!
##
## Running the chain ...
## 0--------25--------50--------75--------100
## ========================================
## Finished MCMC!
##
## Parameter | delta | Acceptance Probability
## ========================================
## diversification | 0.192 | 0.406
## turnover | 0.007 | 0.487
##
## Performing CoMET analysis.
##
## Burning-in the chain ...
## 0--------25--------50--------75--------100
## ========================================
##
## Running the chain ...
## 0--------25--------50--------75--------100
## ========================================
```

We visualize the output as before, this time omitting estimates of the number and timing of mass-extinction events (Figure 20).

```
output <- tess.process.output("tess_analysis_empirical_no_mass_extinctions",
                numExpectedRateChanges = numExpectedRateChanges,
                numExpectedMassExtinctions = numExpectedMassExtinctions)

layout.mat <- matrix(1:4,nrow=2,ncol=2,byrow=TRUE)
layout(layout.mat)
tess.plot.output(output,
                 fig.types = c("speciation rates",
                               "speciation shift times",
                               "extinction rates",
                               "extinction shift times"),
                 las=2)
```

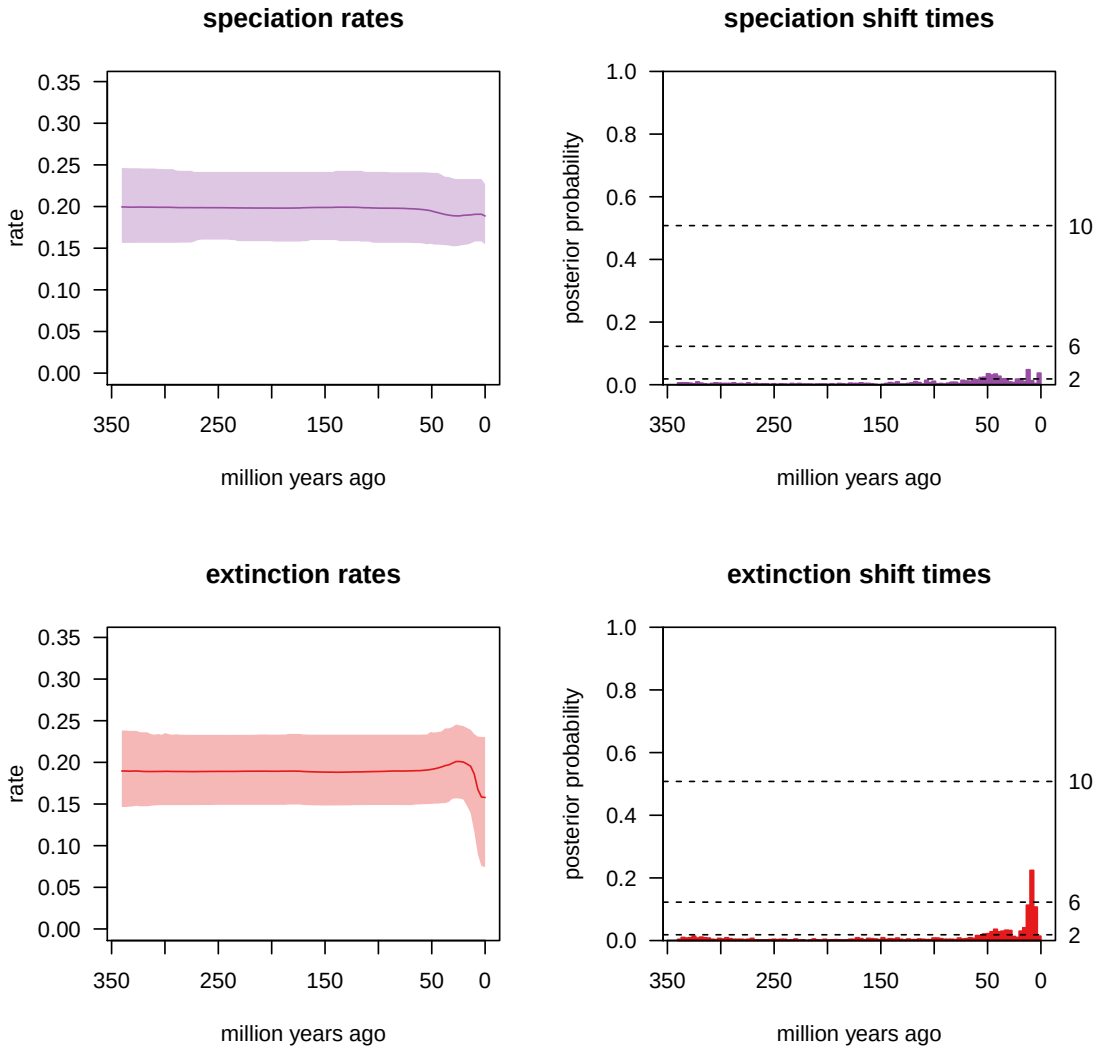

Figure 20: Visualizing the results of a `CoMET` analysis with empirically estimated diversification hyperpriors *and without* mass-extinction events.

5 MCMC Diagnosis

Model-based inference requires us to be vigilant in our choice of models, and also to rigorously assess our ability to obtain reliable estimates under the chosen model. The first issue—related to model evaluation—has been discussed in Section 4 of this guide. Here, we turn to the second issue. Bayesian inference is focussed on the joint posterior probability density of the model parameters, which must be approximated using numerical methods (MCMC simulation). It may be comforting to know that, in theory, an *appropriately constructed* and *adequately run* MCMC simulation is guaranteed to provide an arbitrarily precise description of the joint posterior probability density. In practice, however, even a given MCMC algorithm that provides reliable estimates in *most* cases will nevertheless fail in *some* cases and is not guaranteed to work for any given dataset. This raises an obvious question: "When do we know that an MCMC simulation provides reliable estimates for a given empirical analyses". The answer is simple: *Never*. Convergence can never be established with certainty, only non-convergence may be detected. We will illustrate how to assess convergence of MCMC simulations for two examples: 1) for analyses under a constant-rate birth-death process and 2) for analyses under the CoMET model.

5.1 MCMC diagnosis for a constant-rate birth-death model

In order to properly assess convergence, we need to perform at least two independent MCMC simulations. Otherwise, we run the risk in erroneously concluding that an MCMC has converged when it may have been stuck in a subregion of the joint posterior probability density.

We will use the same likelihood functions an MCMC methods described in Section 2.4. We will first run two short MCMC simulations. **For the sake of this demonstration we will use a pre-burnin value of 0, a thinning of 1, and a chain length of 200 cycles to highlight that the runs have not converged.**

```
prior_delta <- function(x) { dexp(x,rate=10.0,log=TRUE) }
prior_tau <- function(x) { dexp(x,rate=10.0,log=TRUE) }
my_priors <- c("diversification"=prior_delta,
                    "turnover"=prior_tau)
my_likelihood <- function(params) {

  speciation <- params[1] + params[2]
  extinction <- params[2]
```

```
  lnl <- tess.likelihood(times,
                         lambda = speciation,
                         mu = extinction,
                         samplingProbability = 1.0,
                         log = TRUE)

  return (lnl)

}

samples_run_1 <- tess.mcmc(likelihoodFunction = my_likelihood,
                           priors = my_priors,
                           parameters = runif(2,0,10),
                           logTransforms = c(TRUE,TRUE),
                           delta = c(1,1),
                           iterations = 200,
                           burnin = 0,
                           thinning = 1,
                           adaptive = TRUE,
                           verbose = TRUE)

## Burning-in the chain ...
## 0--------25--------50--------75--------100
##
## Finished burnin period!
##
## Running the chain ...
## 0--------25--------50--------75--------100
## ========================================
## Finished MCMC!
##
## Parameter | delta | Acceptance Probability
## ========================================
## diversification | 1.000 | 0.425
## turnover | 1.000 | 0.145

samples_run_2 <- tess.mcmc(likelihoodFunction = my_likelihood,
                           priors = my_priors,
                           parameters = runif(2,0,10),
                           logTransforms = c(TRUE,TRUE),
                           delta = c(1,1),
```

```
                              iterations = 200,
                              burnin = 0,
                              thinning = 1,
                              adaptive = TRUE,
                              verbose = TRUE)

## Burning-in the chain ...
## 0--------25--------50--------75--------100
##
## Finished burnin period!
##
## Running the chain ...
## 0--------25--------50--------75--------100
## ========================================
## Finished MCMC!
##
## Parameter | delta | Acceptance Probability
## ========================================
## diversification | 1.000 | 0.355
## turnover | 1.000 | 0.105
```

We assess convergence using three diagnostics: the effective sample size, the Geweke statistic, and the Gelman-Rubin statistic.

First, we compute the effective sample size (ESS). Samples drawn from an MCMC simulation are correlated. Accordingly, each sample is not independent, and so provides less information. We can compute the number of effectively independent samples by computing the ESS. This is important because we want to draw statistical conclusions from the samples, such as the sample mean. Higher ESS values should provide more precise inferences from the posterior sample. As a rule of thumb, the ESS should be larger than 200.

```
effectiveSize(samples_run_1)

## diversification        turnover
##        24.59086        29.87909

effectiveSize(samples_run_2)

## diversification        turnover
##        2.691079        2.553857
```

Next, we compute the Geweke diagnostic, which assesses convergence by computing the probability that the samples collected during an early window of the MCMC simulation are drawn from the same distribution as samples collected from a later window.

```
geweke.diag(samples_run_1)

##
## Fraction in 1st window = 0.1
## Fraction in 2nd window = 0.5
##
## diversification        turnover
##          0.9034          0.7872

geweke.diag(samples_run_2)

##
## Fraction in 1st window = 0.1
## Fraction in 2nd window = 0.5
##
## diversification        turnover
##         23054.2           137.7
```

Hence, we only need to test if the computed values is smaller than

```
qnorm(0.05/2)

## [1] -1.959964
```

and larger than

```
qnorm(1-0.05/2)

## [1] 1.959964
```

If this is the case, then we reject convergence. Note that we used a significance threshold of $\alpha = 0.05$, but you might use some other threshold, such as $\alpha = 0.01$.

The final MCMC diagnostic—the Gelman-Rubin test—compares samples from two independent simulations. This test compares the variance of sampled parameter values *within* each simulation to that *between* two simulations. This test effectively assesses whether we can reject the null hypothesis that samples from the two independent MCMC simulations are drawn from the same distribution.

```
gelman.diag(x=list(run1 = samples_run_1,run2 = samples_run_2),
            confidence = 0.95, transform = FALSE,
            autoburnin = FALSE, multivariate=TRUE)$mpsrf

## [1] 1.151925
```

If the two independent MCMC simulations have converged to the stationary distribution, the ratio of the within-sample variance to the between-sample variance (R) should be close to 1.0.

In general it is unlikely that those very short MCMC runs have converged. You may observe this by the very low ESS values. Depending on the starting values and your random seed, they actually may or may not have converged; and thus we can not say anything definitely. As a rule-of-thumb, you should run your MCMC simulations a bit longer, especially for more complex models than the simple constant-rate birth-death process model.

5.2 MCMC diagnosis for the CoMET model

5.2.1 Single-chain diagnostics

Proper MCMC diagnosis for a CoMET analysis requires examining the convergence and effective sample size of a large number of parameters, including the number of diversification-rate shifts, the number of mass-extinction events, and the interval-specific diversification-rate parameters.

First, we will assess the MCMC performance for the numerical parameters of the CoMET model (*i.e.,* the number of diversification-rate shifts and mass-extinction events).

```
output <- tess.process.output("tess_analysis_empirical_hyperpriors",
                numExpectedRateChanges = numExpectedRateChanges,
                numExpectedMassExtinctions = numExpectedMassExtinctions)

# Compute the effective sample size and Geweke diagnostic for
# the number of speciation-rate shifts.
effectiveSize(output$numSpeciationCategories)

##     var1
## 127.4808

geweke.diag(output$numSpeciationCategories)
```

```
##
## Fraction in 1st window = 0.1
## Fraction in 2nd window = 0.5
##
##  var1
## 2.271

# Compute the effective sample size and Geweke diagnostic for
# the number of extinction-rate shifts.
effectiveSize(output$numExtinctionCategories)

##    var1
## 17.2843

geweke.diag(output$numExtinctionCategories)

##
## Fraction in 1st window = 0.1
## Fraction in 2nd window = 0.5
##
## var1
##  7.4

# Compute the effective sample size and Geweke diagnostic for
# the number of mass-extinctionevents.
effectiveSize(output$numMassExtinctions)

##     var1
## 78.81995

geweke.diag(output$numMassExtinctions)

##
## Fraction in 1st window = 0.1
## Fraction in 2nd window = 0.5
##
##   var1
## -7.912
```

Next, we will use the function `tess.plot.singlechain.diagnostics` to examine the effective sample size and Geweke diagnostic for the interval-specific speciation- and extinction-rate parameters.

```
layout.mat <- matrix(1:6,nrow=3,ncol=2,byrow=TRUE)
layout(layout.mat)
tess.plot.singlechain.diagnostics(output,
                    parameters = c("speciation rates",
                                   "extinction rates",
                                   "mass extinction times"),
                    las=2)
```

We can see from this single run that, according to the Geweke diagnostic, the MCMC has failed to converge. Additionally, the ESS values for the interval-specific rate-parameter estimates are quite low. We therefore need to extend the length of this MCMC simulation to obtain adequate ESS values (we recommend a minimum ESS of 500) and achieve satisfactory Geweke statics.

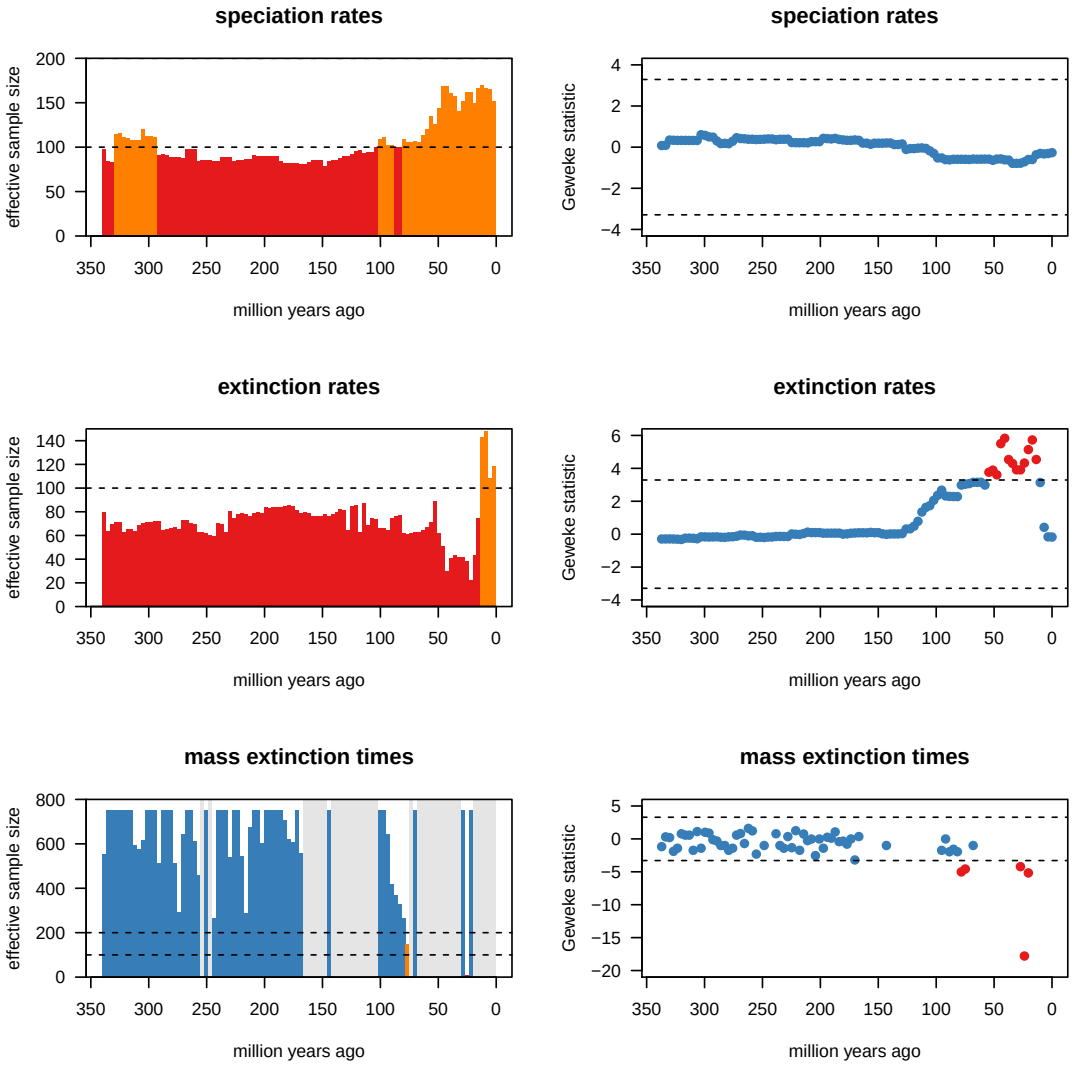

Figure 21: Visualizing the single-chain MCMC diagnostics for a `CoMET` analysis with empirically estimated diversification hyperpriors.

5.2.2 Comparing the posterior to the prior

The CoMET model is implemented in Bayesian statistical framework, where each parameter is assigned a prior probability distribution. Often, we are interested in whether there is sufficient information in the data to estimate a particular parameter. When the marginal posterior distribution of a parameter is very similar to its prior distribution, it suggests that there may not be sufficient information in the data to estimate the value of this parameter. These are referred to as 'weak parameters'. While weak parameters may not present a serious problem, we caution against making biological interpretations regarding these parameters.

We can run CoMET under the prior as follows.

```
set.seed(12345)
tess.analysis(conifers,
              empiricalHyperPriors = FALSE,
              initialSpeciationRate = speciationPriorMu,
              speciationRatePriorMean = speciationRatePriorMean,
              speciationRatePriorStDev = speciationRatePriorStDev,
              initialExtinctionRate = extinctionPriorMu,
              extinctionRatePriorMean = extinctionRatePriorMean,
              extinctionRatePriorStDev = extinctionRatePriorStDev,
              samplingProbability = samplingFraction,
              numExpectedRateChanges = numExpectedRateChanges,
              numExpectedMassExtinctions = numExpectedMassExtinctions,
              pMassExtinctionPriorShape1 = pMassExtinctionPriorShape1,
              pMassExtinctionPriorShape2 = pMassExtinctionPriorShape2,
              MAX_ITERATIONS = 100000,
              ADAPTIVE = FALSE,
              priorOnly = TRUE,
              dir = "tess_analysis_prior")

##
## Performing CoMET analysis.
##
## Burning-in the chain ...
## 0--------25--------50--------75--------100
## ========================================
##
## Running the chain ...
## 0--------25--------50--------75--------100
## ========================================
```

We now load in the data from the prior and the posterior.

```
posterior_output <- tess.process.output("tess_analysis",
                numExpectedRateChanges = numExpectedRateChanges,
                numExpectedMassExtinctions = numExpectedMassExtinctions)

prior_output <- tess.process.output("tess_analysis_prior",
                numExpectedRateChanges = numExpectedRateChanges,
                numExpectedMassExtinctions = numExpectedMassExtinctions)
```

Next, we compare the posterior distribution for the number of speciation-rate changes to the corresponding prior distribution.

```
# Plot the posterior distribution of the number of speciation
# categories (minimum 1)
plot(density(posterior_output$numSpeciationCategories,from=1),
     col = "purple",
     main = "number of speciation-rate categories",
     lwd = 2)

# Overlay the prior number of categories.
lines(density(prior_output$numSpeciationCategories,from=1),
      col = "purple",
      lwd = 2, lty = 2)
legend('topright',c("posterior","prior"),col=c("purple","purple"),lwd=c(2,2),lty=c
```

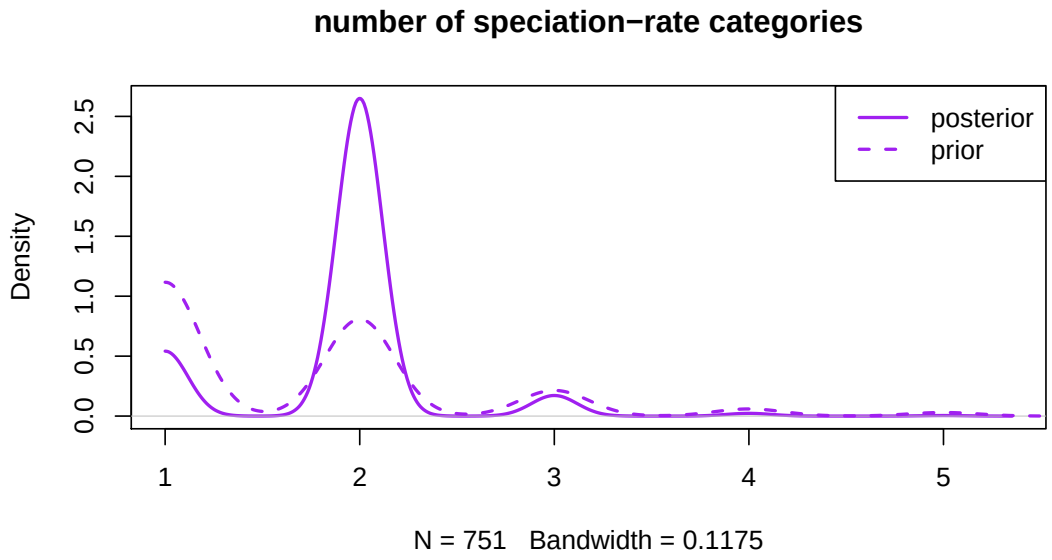

There seems to be little information in the data regarding the number of speciation-rate changes. Notice that the posterior probability of a single rate category (which corresponds to no speciation-rate changes) is slightly higher for the posterior than the prior. Now we make the same comparison for the number of mass-extinction events.

```
# Plot the posterior distribution of the number of speciation
# categories (minimum 0)
plot(density(posterior_output$numMassExtinctions,from=0),
     col = "green",
     main = "number of mass extinction events",
     lwd = 2)

# Overlay the prior number of categories.
lines(density(prior_output$numMassExtinctions,from=0),
      col = "green",
      lwd = 2, lty = 2)
legend('topright',c("posterior","prior"),col=c("green","green"),lwd=c(2,2),lty=c(1
```

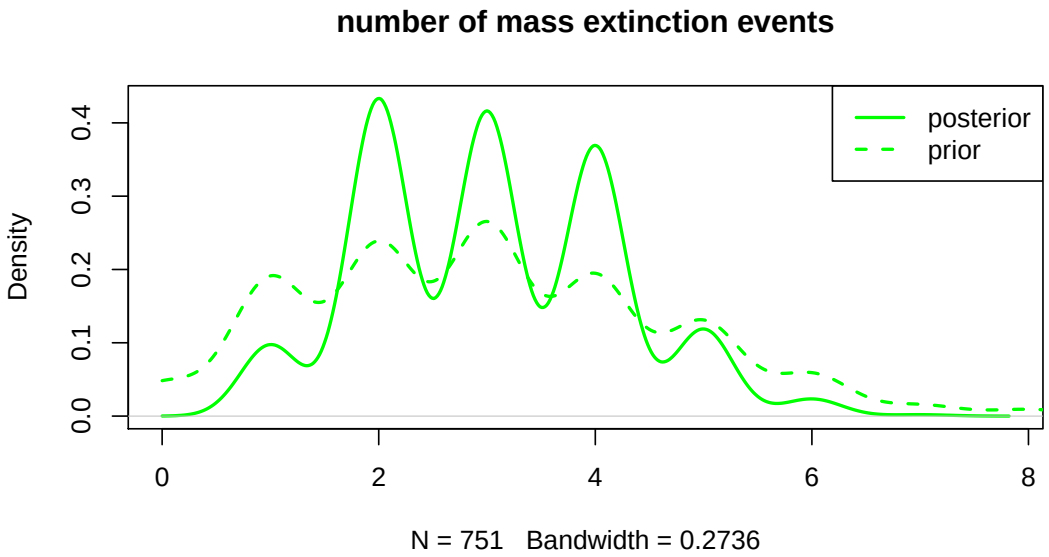

There seems to be more information in the data regarding the number of mass-extinction events. Accordingly, we should be hesitant about making conclusions about the number of speciation-rate changes, we can be more confident regariding the number of mass-extinction events (here between 2 and 4).

5.2.3 Multiple-chain diagnostics

To perform multiple-chain diagnostics, we must first obtain samples from multiple MCMC simulations. Accordingly, we will begin by repeating the above CoMET analysis with empirically estimated hyperpriors four times, each time with the same settings.

```
set.seed(12345)
posterior_directories <- paste("tess_analysis_empirical_hyperpriors_posterior_",
                               1:4,sep="")

for(dir in posterior_directories) {
  tess.analysis(conifers,
                empiricalHyperPriors = TRUE,
                samplingProbability = samplingFraction,
                numExpectedRateChanges = numExpectedRateChanges,
                numExpectedMassExtinctions = numExpectedMassExtinctions,
                pMassExtinctionPriorShape1 = pMassExtinctionPriorShape1,
                pMassExtinctionPriorShape2 = pMassExtinctionPriorShape2,
                MAX_ITERATIONS = 100000,
                dir = dir)
}
```

We use the function `tess.plot.multichain.diagnostics` to compute the Rubin-Gelman convergence diagnostic for the interval-specific parameter estimates. First, we have to process each of the CoMET outputs individually.

```
output_1 <- tess.process.output("tess_analysis_empirical_hyperpriors_posterior_1",
                numExpectedRateChanges = numExpectedRateChanges,
                numExpectedMassExtinctions = numExpectedMassExtinctions)

output_2 <- tess.process.output("tess_analysis_empirical_hyperpriors_posterior_2",
                numExpectedRateChanges = numExpectedRateChanges,
                numExpectedMassExtinctions = numExpectedMassExtinctions)

output_3 <- tess.process.output("tess_analysis_empirical_hyperpriors_posterior_3",
                numExpectedRateChanges = numExpectedRateChanges,
                numExpectedMassExtinctions = numExpectedMassExtinctions)

output_4 <- tess.process.output("tess_analysis_empirical_hyperpriors_posterior_4",
                numExpectedRateChanges = numExpectedRateChanges,
```

```
                numExpectedMassExtinctions = numExpectedMassExtinctions)
```

Next, we make a list of the MCMC outputs and use the Rubin-Gelman diagnostic on the parameters. For brevity, we analyze only a few of the parameters; you should check each parameter to make sure they have all converged to the same posterior distribution!

```
output_list <- list(output_1,output_2,output_3,output_4)

layout.mat <- matrix(1:3,nrow=3,ncol=1,byrow=TRUE)
layout(layout.mat)
tess.plot.multichain.diagnostics(output_list,
                                 parameters = c("speciation rates",
                                                "extinction rates",
                                                "mass extinction times"),
                                 las=2)
```

Generally speaking, those very short independent MCMC runs are unlikely to have converged to their posterior distributions. Depending on the starting values and your random seed, they may or may not have converged; and thus we can not say anything definitely. We recommend running the chains until the PSRF values have converged for all of the intervals.

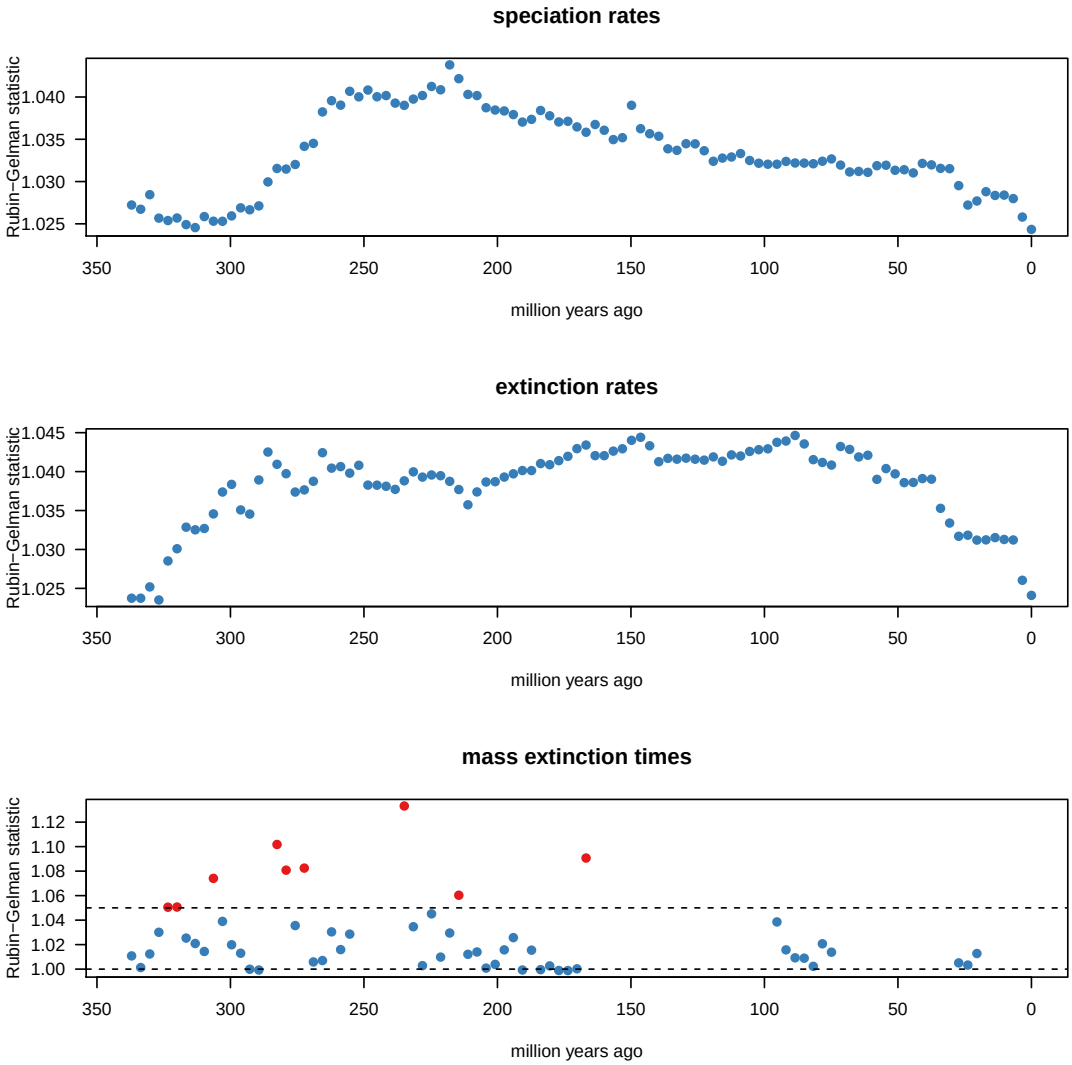

Figure 22: Visualizing the multiple-chain MCMC diagnostics for a `CoMET` analysis with empirically estimated diversification hyperpriors.

References

Baele, G., Lemey, P., Bedford, T., Rambaut, A., Suchard, M., and Alekseyenko, A. (2012). Improving the accuracy of demographic and molecular clock model comparison while accommodating phylogenetic uncertainty. *Molecular Biology and Evolution*, 29(9):2157–2167.

Baele, G., Li, W., Drummond, A., Suchard, M., and Lemey, P. (2013). Accurate Model Selection of Relaxed Molecular Clocks in Bayesian Phylogenetics. *Molecular Biology and Evolution*, 30(2):239–243.

Cusimano, N. and Renner, S. (2010). Slowdowns in diversification rates from real phylogenies may not be real. *Systematic Biology*, 59(4):458.

Fan, Y., Wu, R., Chen, M.-H., Kuo, L., and Lewis, P. O. (2011). Choosing among partition models in bayesian phylogenetics. *Molecular Biology and Evolution*, 28(1):523–532.

Heath, T. A., Huelsenbeck, J. P., and Stadler, T. (2014). The fossilized birth-death process for coherent calibration of divergence-time estimates. *Proceedings of the National Academy of Sciences*, 111(29):E2957–E2966.

Höhna, S. (2013). Fast simulation of reconstructed phylogenies under global time-dependent birth-death processes. *Bioinformatics*, 29(11):1367–1374.

Höhna, S. (2014). Likelihood Inference of Non-Constant Diversification Rates with Incomplete Taxon Sampling. *PLoS One*, 9(1):e84184.

Höhna, S. (2015). The time-dependent reconstructed evolutionary process with a key-role for mass-extinction events. *arXiv preprint arXiv:1312.2392*.

Höhna, S., Stadler, T., Ronquist, F., and Britton, T. (2011). Inferring speciation and extinction rates under different species sampling schemes. *Molecular Biology and Evolution*, 28(9):2577–2589.

Jeffreys, H. (1961). *The theory of probability*. Oxford University Press.

Kass, R. and Raftery, A. (1995). Bayes factors. *Journal of the American Statistical Association*, 90:773–795.

Kendall, D. G. (1948). On the generalized "birth-and-death" process. *The Annals of Mathematical Statistics*, 19(1):1–15.

Lambert, A. (2010). The contour of splitting trees is a lévy process. *The Annals of Probability*, 38(1):348–395.

Lartillot, N. and Philippe, H. (2006). Computing Bayes factors using thermodynamic integration. *Systematic Biology*, 55(2):195.

Leslie, A. B., Beaulieu, J. M., Rai, H. S., Crane, P. R., Donoghue, M. J., and Mathews, S. (2012). Hemisphere-scale differences in conifer evolutionary dynamics. *Proceedings of the National Academy of Sciences*, 109(40):16217–16221.

May, M. R., Höhna, S., and Moore, B. R. (2015). A Bayesian Approach for Detecting Mass-Extinction Events When Rates of Lineage Diversification Vary. *submitted*.

Nee, S., May, R. M., and Harvey, P. H. (1994). The Reconstructed Evolutionary Process. *Philosophical Transactions: Biological Sciences*, 344(1309):305–311.

Pybus, O. and Harvey, P. (2000). Testing macro–evolutionary models using incomplete molecular phylogenies. *Proceedings of the Royal Society B: Biological Sciences*, 267(1459):2267–2272.

Rabosky, D. L. and Lovette, I. (2008). Explosive evolutionary radiations: decreasing speciation or increasing extinction through time? *Evolution*, 62(8):1866–1875.

Stadler, T. (2009). On incomplete sampling under birth-death models and connections to the sampling-based coalescent. *Journal of Theoretical Biology*, 261(1):58–66.

Suchard, M. A., Weiss, R. E., and Sinsheimer, J. S. (2001). Bayesian selection of continuous-time Markov chain evolutionary models. *Molecular Biology and Evolution*, 18(6):1001–1013.

Thompson, E. (1975). *Human evolutionary trees*. Cambridge University Press Cambridge.

Xie, W., Lewis, P., Fan, Y., Kuo, L., and Chen, M. (2011). Improving marginal likelihood estimation for Bayesian phylogenetic model selection. *Systematic Biology*, 60(2):150–160.

Yule, G. (1925). A mathematical theory of evolution, based on the conclusions of dr. jc willis, frs. *Philosophical Transactions of the Royal Society of London. Series B, Containing Papers of a Biological Character*, 213:21–87.