

COMPLEX MODELS FOR GENETIC SEQUENCE DATA

NAOMI ELIZABETH HANNAFORD

Thesis submitted for the degree of
Doctor of Philosophy



*School of Mathematics, Statistics & Physics
Newcastle University
Newcastle upon Tyne
United Kingdom*

May 2021

To Mam and Dad, thank you for your love and support and for always believing in me.

Acknowledgements

Firstly, I would like to thank my supervisors, Sarah Heaps and Tom Nye, for their exceptional guidance, kindness and encouragement. I have enjoyed our time working together and could not ask for a better supervisory team. I would also like to thank Ben Allen and Tom Curtis for providing interesting data for me to work with and biological insights related to Part II of the thesis. Furthermore, I would like to thank Martin Embley and Tom Williams for their help with the biological aspects of Part I of the thesis.

I am grateful to Oonagh McGee, Andrew Turnbull and Jen Wood for providing excellent professional and personal support during my time as a PhD student. I would also like to thank Matt Forshaw, Paul Watson and Darren Wilkinson for their advice, help and kindness over the past few years. I am thankful to Michael Beaty and Stephen Dowsland for their help with computational issues. Additionally, I would like to thank my examiners Malcolm Farrow and Gavin Gibson. I am also grateful to the Engineering and Physical Sciences Research Council for the funding which made it possible to conduct this research.

During my time as a postgraduate student, I have met some wonderful people, whom I am lucky enough to call my friends. I am incredibly grateful to Antonia Kontaratou, Ashleigh McLean and Lauren Roberts for their support and for being the best company during our (possibly too) many coffee breaks. I would like to thank Jack Aiston for his mathematical help and for being such a kind and fun flatmate. I give my thanks to Tom Cooper for the many occasions that he helped fix my laptop and for tolerating me as a desk neighbour, even when I complained he was too loud. I would also like to thank Jonny Law for sharing his statistical knowledge, running advice and enthusiasm for going to the pub. I am grateful to Matt Edwards for his statistical help and for introducing me to bouldering, which really helped with PhD-related stress (and also was the biggest distraction from the PhD). With this in mind, I would like to thank Adam Cattermole and Kathryn Garside for their friendship and support both on and off the climbing wall. I would like to thank all my other friends for supporting me and I apologise for not being able to mention you all by name but there are too many great people to fit onto one page. However, whether we have battled on the badminton or squash court, climbed, frequented Wagamama or Chilli Padi, quizzed on a Friday lunchtime or chatted at the pub or lunch, I am truly thankful for your friendship.

I would like to thank Beth Cummins for her unwavering friendship and for always being there for me. Finally, I would like to thank Javier Núñez Vicandi for being the best hype man and partner anyone could ever ask for.

Abstract

In this thesis, the aim is to develop biologically motivated Bayesian models in two areas: molecular phylogenetics and time-series metagenomics. In molecular phylogenetics, the goal is generally to learn about the evolutionary history of a collection of species using molecular sequence data, for example, DNA. Evolutionary history is represented graphically using evolutionary trees, where the root of a tree represents the most recent common ancestor of all species in the tree. Substitutions in sequences are modelled through a continuous time Markov process, characterised by an instantaneous rate matrix, which standard models assume is stationary and time-reversible. These assumptions are biologically questionable and induce a likelihood function which is invariant to a tree's root position. This is detrimental to inference, since a tree's biological interpretation depends on where it is rooted. By relaxing both assumptions, we introduce two new models whose likelihoods can distinguish between rooted trees. These models are non-stationary, with step changes in the rate matrix on each branch. Each rate matrix belongs to a non-reversible family of Lie Markov models, which are closed under matrix multiplication. The two models differ in that a different non-reversible Lie Markov model is used in each. We perform our analysis in the Bayesian framework using Markov chain Monte Carlo methods. We assess the performance of our models using a simulation study, before considering an application to a *Drosophila* data set, where most models fail to identify a plausible root position.

In time-series metagenomics, counts of operational taxonomic units (OTUs), which are pragmatic proxies for microbial species, are modelled over time. We have weekly counts of different OTUs from two tanks in a wastewater treatment plant. We develop a Bayesian hierarchical vector autoregressive model to model the dynamics of the OTUs, whilst also incorporating environmental and chemical data. Clustering methods are explored to reduce the dimensionality of our data and mitigate the issue of large proportions of zero-counts in the data. We use a seasonal phase-based clustering approach and a symmetric, circulant, tri-diagonal error structure. The autoregressive coefficient matrix is assumed to be sparse, so we explore different priors that allow for sparsity by analysing simulated data sets before selecting the regularised horseshoe prior for our hierarchical model. The chemical and environmental covariates are incorporated through a time varying mean. Finally, we fit the model to the data from each tank using Hamiltonian Monte Carlo.

Contents

1	Introduction	1
1.1	Introduction and aims	1
1.2	Outline	3
2	Background	6
2.1	Introduction to phylogenetics	6
2.1.1	Statistical phylogenetics	6
2.1.2	Standard models of DNA evolution	8
2.1.3	Likelihood of a phylogenetic tree	9
2.1.4	Drawbacks of standard models	11
2.2	Introduction to metagenomics	14
2.2.1	Statistical metagenomics	14
2.2.2	Time-series metagenomics	15
2.2.3	Data overview	18
2.2.4	Motivation, aims and challenges	19
2.3	Bayesian inference	20
2.3.1	Bayes theorem	21
2.3.2	Markov chain Monte Carlo	22
2.3.3	Hamiltonian Monte Carlo	28
I	Phylogenetics	32
3	Phylogenetic models facilitating root inference	33
3.1	Non-homogeneous models	33
3.2	Non-reversible models	35
3.2.1	The RY5.6b Model	37
3.2.2	The RY8.8 Model	42
3.3	Assumption across sites	45

4	Non-homogeneous Lie Markov models	46
4.1	Non-homogeneous RY5.6b and RY8.8 models	46
4.2	Prior	47
4.3	Posterior inference via MCMC	50
4.3.1	Proposal distributions	51
4.3.2	Metropolis-within-Gibbs sampling scheme	56
5	Phylogenetics application	59
5.1	Simulation study	59
5.1.1	Different numbers of taxa and sites	60
5.1.2	Different topologies and branch lengths	65
5.2	The <i>Drosophila</i> data set	71
5.2.1	MCMC implementation	75
5.2.2	Posterior inference	76
5.2.3	Model comparison	79
5.3	Results summary	82
II	Metagenomics	84
6	Exploratory data analysis	85
6.1	Data description	85
6.1.1	Reformatting the data	85
6.1.2	Missing data	88
6.2	Exploratory data analysis	90
6.2.1	OTU tables	91
6.2.2	Taxonomy table	94
6.2.3	Analysis of the combined data	99
6.2.4	Summary	106
7	A Bayesian hierarchical model for time-series metagenomics data	108
7.1	Forming subpopulations	108
7.1.1	Clustering based on taxonomy	109
7.1.2	Time series clustering	110
7.2	Simple VAR(1) model	112
7.3	Allowing for sparsity	113
7.3.1	Spike and slab prior	114
7.3.2	Horseshoe prior	119
7.3.3	Prior for the global shrinkage parameter	122
7.3.4	Regularised horseshoe prior	137

7.3.5	Comparison of priors via simulation study	138
7.4	Error structure	148
7.5	Time varying mean	152
7.5.1	Incorporating the chemical and environmental data	152
7.5.2	Handling missing values	154
7.6	Posterior inference	154
7.6.1	Model and prior specification	154
7.6.2	Posterior inference via HMC	156
8	Metagenomics application	158
8.1	Activated sludge tank results	158
8.1.1	Time varying mean	158
8.1.2	Matrix of autoregressive coefficients	164
8.1.3	Precision matrix for errors	168
8.2	Settled sewage tank results	168
8.2.1	Time varying mean	169
8.2.2	Matrix of autoregressive coefficients	172
8.2.3	Precision matrix	173
8.3	Results summary	175
III	Conclusions	176
9	Conclusions and future work	177
9.1	Contributions	177
9.2	Conclusions	180
9.3	Future work	182
9.3.1	Phylogenetics future work	182
9.3.2	Metagenomics future work	184
A	Additional proofs and derivations	188
A.1	Constraint on trace of the RY5.6b rate matrix	188
A.1.1	Proof for boundary conditions of c	189
A.1.2	Setting the trace	190
A.2	Spike and slab prior FCDs	191
A.2.1	Derivation of the FCD of $\tilde{\beta}$	191
A.2.2	Derivation of the FCD of $\tilde{\beta}$ with $\Sigma = \tau^{-1}\mathbf{I}_K$	191
A.2.3	Derivation of the FCD of λ_{jk} with $\Sigma = \tau^{-1}\mathbf{I}_K$	194
A.2.4	Derivation of the FCD of $\tilde{\Lambda}$	195
A.2.5	Derivation of the FCD of $\tilde{\Lambda}$ with $\Sigma = \tau^{-1}\mathbf{I}_K$	197

A.2.6	The FCD of λ_{jk} (VAR(1) model)	199
A.2.7	Derivation of the FCD of λ_{jk} with $\Sigma = \tau^{-1}\mathbf{I}_K$ (VAR(1) model)	200
A.3	Horseshoe prior FCDs	201
A.3.1	Derivation of the FCD of τ_β	201
A.3.2	Derivation of the FCD of λ_{jk}	202
A.4	Prior for the global shrinkage parameter	202
A.4.1	Proof for $\bar{\beta}_j = (1 - \kappa_j)\hat{\beta}_j$	202
A.4.2	Derivation of conditional posterior for β^*	203
A.4.3	Formulating \mathcal{K}_j	204
A.4.4	Expected value for $\mathcal{K}_{j,kk}$	204
A.5	Conditional means and variances of a_{jk}	205
A.5.1	Spike and slab prior	205
A.5.2	Horseshoe prior	206
A.6	Reparameterisation of symmetric, circulant, tridiagonal precision matrix	207
B	MCMC algorithms	208
B.1	Algorithms for shrinkage parameters in VAR(1) model with horseshoe prior	208
B.1.1	Metropolis-within-Gibbs step for λ_{jk}	208
B.1.2	Metropolis-within-Gibbs step for τ_A	209
C	Code	210
C.1	Stan code	210
D	Additional tables	214
E	Additional plots	215
E.1	Exploratory plots for taxonomic ranks	215
E.2	Heatmaps of matrices used in metagenomics simulation study	222

List of Figures

2.1	Example of a phylogenetic tree with four species.	7
2.2	Examples of trace plots showing (a) a chain that has converged and (b) a chain that has not converged.	25
2.3	Diagnostic plots of two chains to assess convergence in tree space where (a) is the posterior probabilities of clades, (b) is the cumulative relative frequencies of clades and the different colours represent different clades, the solid lines represent one chain and the dashed lines represent the other chain.	29
3.1	Plots of the stationary probability π_1 of the RY5.6b model against ρ_1 for values of $\boldsymbol{\rho}$ sampled from the uniform distribution over \mathcal{S}_4 and various values of α . By symmetry, plots for the other pairs (ρ_i, π_i) , $i = 2, 3, 4$, display identical patterns (not shown).	41
4.1	Directed acyclic graph for our non-homogeneous models, where $\mathcal{Q} = \{\alpha, \varrho_1, \dots, \varrho_{B-1}\}$ under the non-homogeneous RY5.6b model and $\mathcal{Q} = \{\varrho_1, \dots, \varrho_{B-1}\}$ for the non-homogeneous RY8.8 model.	50
4.2	An illustration of the NNI move. An internal branch e is chosen uniformly at random from the set of internal branches not adjacent to the root. Either subtree T_1 or T_2 descended from the vertex v is swapped with the subtree T_R descended from v_R	52
4.3	Two possible trees resulting from the NNI move shown in Figure 4.2. In (a) the subtree T_1 is swapped with T_R . In (b) the subtree T_2 is swapped with T_R . The length of branch e^* is proposed via a log normal random walk centred on the length of e from the original tree.	53

4.4	An illustration of the SPR move. (a) During the move, the branch e_p (dashed line) and the tree T evolving from it are pruned and reattached to edge e_g . The point of attachment v_g is chosen by dividing the edge e_g using a random variable drawn from Beta(2, 2). (b) After the move is made, the vertex v_p disappears and the branches e_a and e_b are merged to form a new edge e_g^* . The grafting branch e_g is split into two new edges e_a^* and e_b^* by a new vertex v_g which is formed after reattaching the subtree T to e_g	55
5.1	Trees used to simulate alignments on (a) 6, (b) 12, (c) 24 taxa.	61
5.2	Posterior distribution over roots splits when three data sets are simulated and analysed under the non-homogeneous RY5.6b model and the number of taxa is (a) 6, (c) 12, (e) 24; and when three data sets are simulated and analysed under the non-homogeneous RY8.8 model and the number of taxa is (b) 6, (d) 12, (f) 24. The number of sites is displayed to the right of each plot. In every plot, bars are arranged in descending order of posterior probability and the correct root split is highlighted in black. In the plots for 12 and 24 taxa, bars corresponding to probabilities less than 0.01 have been removed to improve readability.	62
5.3	Posterior distribution over unrooted topologies when the data are simulated and analysed under the non-homogeneous RY5.6b model and the number of taxa is (a) 6, (c) 12, (e) 24; and when the data are simulated and analysed under the non-homogeneous RY8.8 model and the number of taxa is (b) 6, (d) 12, (f) 24. The number of sites is displayed to the right of each plot. Highlighted in black is the correct unrooted topology. In the plots for 12 and 24 taxa, bars corresponding to probabilities less than 0.01 have been removed to improve readability.	63
5.4	Marginal prior and posterior densities for the global model parameters (ϕ and α) and a random sample of branch-specific parameters in the analysis of a 24-taxon alignment simulated under the non-homogeneous RY5.6b model. Posterior densities for the branch-specific parameters are conditional on the rooted topology used to simulate the data. Indicated in the panels are the prior (—) and posterior when the number of sites is 500 (—), 1000 (—) and 2000 (—). The true values of the parameters are indicated by vertical lines.	65

5.5	Marginal prior and posterior densities for the global model parameter (ϕ) and a random sample of branch-specific parameters in the analysis of a 24-taxon alignment simulated under the non-homogeneous RY8.8 model. Posterior densities for the branch-specific parameters are conditional on the rooted topology used to simulate the data. Indicated in the panels are the prior (—) and posterior when the number of sites is 500 (—), 1000 (—) and 2000 (—). The true values of the parameters are indicated by vertical lines.	66
5.6	Number of $k : (n - k)$ root splits (on log-scale) for $n = 16$ taxa.	67
5.7	Unrooted tree used in simulation experiments to investigate the effects of different topologies and branch lengths on root inference. In the experiments, the tree is rooted at the midpoint of either branch E_1 or E_2 . The tree is depicted with branch E_1 having a “medium” length of 0.084 units. In the experiment, this is varied to 0.237 units (“long”) or 0.018 units (“short”).	68
5.8	Posterior distribution over roots splits when three data sets are simulated and analysed under the non-homogeneous RY8.8 model and the tree used for simulation is Tree (a) 1 (balanced, long root branch), (b) 2 (unbalanced, long internal branch), (c) 3 (balanced, short root branch), (d) 4 (unbalanced, short internal branch), (e) 5 (balanced, medium root branch), (f) 6 (unbalanced, medium internal branch). In every plot, bars are arranged in descending order of posterior probability and the correct root split is highlighted in black.	69
5.9	Posterior distribution over roots splits when three data sets are simulated and analysed under the non-homogeneous RY5.6b model and the tree used for simulation is Tree (a) 1 (balanced, long root branch), (b) 2 (unbalanced, long internal branch), (c) 3 (balanced, short root branch), (d) 4 (unbalanced, short internal branch), (e) 5 (balanced, medium root branch), (f) 6 (unbalanced, medium internal branch). In every plot, bars are arranged in descending order of posterior probability and the correct root split is highlighted in black.	70
5.10	Posterior distribution over unrooted topologies when three data sets are simulated and analysed under the non-homogeneous RY8.8 model and the tree used for simulation is Tree (a) 1 (balanced, long root branch), (b) 2 (unbalanced, long internal branch), (c) 3 (balanced, short root branch), (d) 4 (unbalanced, short internal branch), (e) 5 (balanced, medium root branch), (f) 6 (unbalanced, medium internal branch). In every plot, bars are arranged in descending order of posterior probability and the correct unrooted topology is highlighted in black.	71

5.11	Posterior distribution over unrooted topologies when three data sets are simulated and analysed under the non-homogeneous RY5.6b model and the tree used for simulation is Tree (a) 1 (balanced, long root branch), (b) 2 (unbalanced, long internal branch), (c) 3 (balanced, short root branch), (d) 4 (unbalanced, short internal branch), (e) 5 (balanced, medium root branch), (f) 6 (unbalanced, medium internal branch). In every plot, bars are arranged in descending order of posterior probability and the correct unrooted topology is highlighted in black.	72
5.12	Boxplots for quantiles of GC-content sampled from the prior predictive distributions of the taxa based on models (a) \mathcal{M}_4 , (b) \mathcal{M}_5 , (c) \mathcal{M}_6	75
5.13	Majority rule consensus trees under the models (a) \mathcal{M}_1 – GTR; (b) \mathcal{M}_2 – RY5.6b; (c) \mathcal{M}_3 – RY8.8; Numerical labels represent the posterior probability of the associated split (in (a)) or clade (in (b) and (c)).	77
5.14	Majority rule consensus trees under the models (a) \mathcal{M}_4 – non-homogeneous GTR; (b) \mathcal{M}_5 – non-homogeneous RY5.6b; (c) \mathcal{M}_6 – non-homogeneous RY8.8. Numerical labels represent the posterior probability of the associated clade.	78
5.15	Box-plots of the Newton and Raftery hybrid estimates of the marginal likelihood for each model \mathcal{M}_i , $i = 1, \dots, 6$, based on the output of 10 MCMC chains which were initialised at different starting points.	82
6.1	Time series plots of some OTUs (with their genus in brackets) that only appear in the OTU table of the AS and (a) mainly have zero-counts, or (b) appear at more than one or two time points. Note that the y -axis scales are different for each plot.	86
6.2	Time series plots of four OTUs (with their genus in brackets) that appear in both tanks.	87
6.3	Time series plots of total abundances for each OTU table.	91
6.4	(a) time series plots and (b) stacked bar plot for the top 12 OTUs in the AS based on median abundance	92
6.5	(a) time series plots and (b) stacked bar plot for the top 12 OTUs in the SS based on median abundance.	93
6.6	(a) time series plots and (b) stacked bar plot for the top 12 genera based on median abundance in the AS tank.	95
6.7	(a) time series plots and (b) stacked bar plot for the top 12 genera in the SS based on median abundance.	96
6.8	(a) time series plots and (b) stacked bar plot for the top 12 classes in the AS based on median abundance.	98

6.9	(a) time series plots and (b) stacked bar plot for the top 12 classes in the SS based on median abundance.	100
6.10	Heatmap of the correlations between the chemical and environmental covariates and the 12 most abundant OTUs in the AS tank.	102
6.11	Heatmap of the pairwise correlations between the chemical and environmental covariates in the AS tank.	103
6.12	Heatmaps of correlations between the chemical and environmental covariates and the 12 most abundant OTUs in the SS tank.	103
6.13	Heatmap of the pairwise correlations between the chemical and environmental covariates in the SS tank.	104
6.14	Heatmap of the correlations between the chemical and environmental covariates and the 12 most abundant genera in the AS tank.	105
6.15	Heatmaps of correlations between the chemical and environmental covariates and the 12 most abundant genera in the SS tank.	105
6.16	Heatmap of the correlations between the chemical and environmental covariates and the 12 most abundant classes in the AS tank.	106
6.17	Heatmaps of correlations between the chemical and environmental covariates and the 12 most abundant classes in the SS tank.	107
7.1	Time series plots of the scaled log counts for the 12 bins in the AS tank. . .	110
7.2	Time series plots of the scaled log counts for the 12 bins in the SS tank. . .	111
7.3	The density of the shrinkage coefficient $\kappa_j = 1/(1 + \lambda_j^2)$. The horseshoe shape of the density shows that the horseshoe prior favours values of either zero or one for κ_j , which correspond to no shrinkage or near-total shrinkage, respectively.	120
7.4	Conditional on the global shrinkage parameter τ_β and the error variance Σ , marginal densities of a diagonal and non-diagonal element of \mathcal{K}_j when Σ has compound symmetric structure and the dimension of the response vector is (a) $K = 2$, (b) $K = 8$ and (c) $K = 14$. In each case, $\sigma = 1$, $N = s_j = 1$ and $\tau_\beta = 1$. A range of values for the correlation ρ is considered.	132
7.4	Conditional on the global shrinkage parameter τ_β and error variance Σ , pairwise bivariate densities between a pair of diagonal elements of \mathcal{K}_j (Dg vs Dg); a pair of off-diagonal elements (ODg vs ODg); a diagonal and off-diagonal element within the same row (WR Dg vs ODg); and a diagonal and off-diagonal element in different rows (AR Dg vs ODg) when Σ has compound symmetric structure and the dimension of the response vector is (a) $K = 2$, (b) $K = 8$ and (c) $K = 14$. In each case, $\sigma = 1$, $N = s_j = 1$ and $\tau_\beta = 1$. A range of values for the common correlation ρ is considered.	134

7.5	Posterior means (\circ) and 95% credible intervals (—) for A_{80} , with the true values (\times), under the multivariate normal prior.	139
7.6	Posterior means (\circ) and 95% credible intervals (—) for A_{80} , with the true values (\times), under the spike and slab.	140
7.7	Posterior means (\circ) and 95% credible intervals (—) for A_{80} , with the true values (\times), under the horseshoe prior.	141
7.8	Posterior means (\circ) and 95% credible intervals (—) for A_{80} , with the true values (\times), under the regularised horseshoe prior.	141
7.9	Posterior means (\circ) and 95% credible intervals (—) for A_{50} , with the true values (\times), under the multivariate normal prior.	142
7.10	Posterior means (\circ) and 95% credible intervals (—) for A_{50} , with the true values (\times), under the spike and slab.	143
7.11	Posterior means (\circ) and 95% credible intervals (—) for A_{50} , with the true values (\times), under the horseshoe prior.	143
7.12	Posterior means (\circ) and 95% credible intervals (—) for A_{50} , with the true values (\times), under the regularised horseshoe prior.	144
7.13	Posterior means (\circ) and 95% credible intervals (—) for A_{20} , with the true values (\times), under the multivariate normal prior.	145
7.14	Posterior means (\circ) and 95% credible intervals (—) for A_{20} , with the true values (\times), under the spike and slab.	145
7.15	Posterior means (\circ) and 95% credible intervals (—) for A_{20} , with the true values (\times), under the horseshoe prior.	146
7.16	Posterior means (\circ) and 95% credible intervals (—) for A_{20} , with the true values (\times), under the regularised horseshoe prior.	147
8.1	Posterior means (\circ) and 95% credible intervals (—) for nitrate and phosphate (AS tank).	159
8.2	Posterior means (\circ) and 95% credible intervals (—) for COD, ammonia and pH (AS tank).	160
8.3	Posterior means (\circ) and 95% credible intervals (—) for β_j (AS tank).	162
8.4	Posterior means (\circ) and 95% credible intervals (—) for γ_j (AS tank).	162
8.5	Posterior means (—) and 95% credible intervals (---) for the time varying means with scaled log counts (—) for each bin (AS tank).	163
8.6	Heatmap of the posterior means of the autoregressive coefficients (AS tank).	165
8.7	Posterior means (\circ) and 95% credible intervals (—) of the autoregressive coefficients (AS tank).	166
8.8	Posterior means (\circ) and 95% credible intervals (—) of the lag- k correlations ρ_k for $k = 1, \dots, 6$ (AS tank).	168

8.9	Posterior means (\circ) and 95% credible intervals (---) for iron, fluoride and chloride (SS tank).	169
8.10	Posterior means (\circ) and 95% credible intervals (---) for silicon, phosphate, ammonia and flow (SS tank).	169
8.11	Posterior means (\circ) and 95% credible intervals (---) for β_j (SS tank).	171
8.12	Posterior means (\circ) and 95% credible intervals (---) for γ_j (SS tank).	171
8.13	Posterior means (---) and 95% credible intervals (---) for the time varying means with scaled log counts (---) for each bin (SS tank).	172
8.14	Heatmap of the posterior means of the autoregressive coefficients (SS tank).	173
8.15	Posterior means (\circ) and 95% credible intervals (---) of the autoregressive coefficients (SS tank).	174
8.16	Posterior means (\circ) and 95% credible intervals (---) of the lag- k correlations ρ_k for $k = 1, \dots, 6$ (SS tank)	174
E.1	(a) time series plots and (b) stacked bar plot for the top 12 families in the AS based on median abundance.	216
E.2	(a) time series plots and (b) stacked bar plot for the top 12 families in the SS based on median abundance.	217
E.3	(a) time series plots and (b) stacked bar plot for the top 12 orders in the AS based on median abundance.	218
E.4	(a) time series plots and (b) stacked bar plot for the top 12 orders in the SS based on median abundance.	219
E.5	(a) time series plots and (b) stacked bar plot for the top 12 phyla in the AS based on median abundance.	220
E.6	(a) time series plots and (b) stacked bar plot for the top 12 phyla in the SS based on median abundance.	221
E.7	Heatmap of A_{80}	222
E.8	Heatmap of A_{50}	223
E.9	Heatmap of A_{20}	224

List of Tables

- 5.1 Log marginal likelihoods for each model approximated using the Newton and Raftery hybrid estimator in which the prior weight in the importance density was set at $\delta = 0.05$. The models are: \mathcal{M}_1 – GTR; \mathcal{M}_2 – RY5.6b; \mathcal{M}_3 – RY8.8; \mathcal{M}_4 – non-homogeneous GTR; \mathcal{M}_5 – non-homogeneous RY5.6b; \mathcal{M}_6 – non-homogeneous RY8.8. 81
- 6.1 Chemical and environmental covariates in the AS. 88
- 6.2 Chemical and environmental covariates in the SS. 88
- 6.3 Proportions of missing data (4 d.p.) in chemical/environmental data (AS). 89
- 6.4 Proportions of missing data (4 d.p.) in chemical/environmental data (SS). . 89
- 6.5 Proportions of missing data (4 d.p.) for each taxonomic rank in the taxonomy table for each tank and overall (after OTUs with zero-counts for all time points are removed). 90
- 6.6 Number of different types for each taxonomic rank in each tank, excluding NAs. 94
- 7.1 Number of autoregressive coefficients not correctly identified in the simulation study for each matrix and prior. 148
- 8.1 Posterior means (3 d.p) of the non-zero between-bin coefficients. 165
- 8.2 Genera of the top six OTUs in each bin based on median relative abundance (AS tank). 167
- 8.3 Genera of the top six OTUs in each bin based on median relative abundance (SS tank). 170
- D.1 Chemical and environmental covariates with a measurement regarded as an outlier and removed from the data, with the exception of COD, which had three outliers. 214

Chapter 1

Introduction

1.1 Introduction and aims

The aim of this thesis is to develop biologically motivated hierarchical Bayesian models in two areas of biology: *molecular phylogenetics* and *time-series metagenomics*. In molecular phylogenetics, the general goal is to learn about the evolutionary history of a collection of individuals, for example, different species, using molecular sequence data, such as DNA or amino acids. Bifurcating trees called *evolutionary trees* (or *phylogenies*) are used to graphically represent this evolutionary history. The root of the tree represents the *most recent common ancestor* of all species in the tree and its position is fundamental to a phylogeny's biological interpretation.

Generally, substitutions in molecular sequence data are modelled using *continuous time Markov processes*, characterised by an instantaneous rate matrix. Standard substitution models assume that this rate matrix is both stationary and time-reversible, which provides mathematical convenience. However, these assumptions are biologically questionable and can be challenged empirically. Furthermore, they give rise to a likelihood function that is invariant to root position. As such, these standard models can only be used to infer unrooted trees, which depict the branching pattern of speciation events without associating direction to the branches of the tree. Models that relax at least one of the restrictive assumptions can be used to infer rooted trees, as their likelihoods depend on the position of the root.

In this thesis, we develop two phylogenetic models that are both non-stationary and non-reversible. Non-stationarity is achieved by allowing step changes in the rate matrix at each *speciation event* (internal vertex of a tree). Each rate matrix belongs to a non-reversible family of *Lie Markov models*. These models are closed under matrix multiplication, which means our models possess the conceptually appealing property that a tree and all its subtrees could have arisen from the same family of non-stationary models.

The two models are different because we use a different non-reversible Lie Markov model in each model.

We perform our analyses in the Bayesian framework using Markov chain Monte Carlo (MCMC) methods. To assess the performance of our models, we conduct an extensive simulation study. The main goal of the simulation study is to test our models' ability to identify root position. We examine the effect of different topologies and branch lengths and the effect of data dimensionality on root inference. Additionally, we investigate the identifiability of the numerical parameters in each model, as this may have an effect on the inference of the *topology* (*branching pattern*) and root position. Finally, we consider an application to a *Drosophila* data set, where the assumption of stationarity can be empirically challenged, resulting in most models failing to identify a plausible root position.

The data analysed in phylogenetics are sequences of DNA from different species. In metagenomic studies, DNA sequences are the primary data source too, but they are solely from microbial species. The sequences are then clustered based on similarity into *operational taxonomic units* (OTUs), which we consider as pragmatic proxies for species, and these OTUs are enumerated. It is these counts of OTUs that are modelled in metagenomics.

The models developed in this thesis for phylogenetic inference are theory-driven. We identify the problems and challenges associated with standard models and thus develop our models to address these issues. However, in our time-series metagenomics application, the model is data-driven and designed to solve a specific problem. We have five years of weekly counts of OTUs from two different sources in a wastewater treatment plant (WWTP): the activated sludge and the settled sewage. Further to these counts, we have taxonomic classifications for each OTU and weekly environmental and chemical measurements from each source, such as temperature, pH and ammonia concentration. Our goal is to model the *dynamics* of the OTUs (changes in their counts and interactions between them over time), whilst also incorporating the chemical and environmental data.

To achieve this goal, we develop a Bayesian hierarchical vector autoregressive (VAR) model. A VAR process models a vector of values at a current time point as a linear regression on a fixed number of its previous values. The number of previous time points used is defined by the *order* of the VAR model. We begin with a simple VAR model of order one (VAR(1)), which we adapt to our data, based on exploratory data analysis and theoretical considerations. As is commonly found in metagenomic studies, our OTU data are high-dimensional and *sparse* (meaning there is a large proportion of zeroes). We deal with both issues by first clustering the OTUs into a much smaller number of bins using a seasonal phase-based clustering method. It is then these clustered data that we model using a VAR process. A VAR(1) model is characterised by a single matrix of autoregressive coefficients and a mean vector (which can be a zero-mean vector). Given our

chosen clustering method, we allow our autoregressive matrix to be sparse through its prior specification. Via a simulation study, we explore different priors that encourage shrinkage of autoregressive coefficients towards zero before selecting a regularised horseshoe prior. The regularised horseshoe is governed by two parameters, the global shrinkage and the local shrinkage parameters. We use a hyperprior for the global shrinkage parameter that considers prior information of sparsity and error variance. Our chosen clustering method also influences our choice of error structure, which is a symmetric, circulant, tridiagonal precision matrix. The final amendment to the VAR(1) model is the inclusion of a time varying mean, which we fit with harmonic regression to capture seasonal variation in our clustered data. Furthermore, we regress the intercept term on the chemical and environmental covariates.

As in our phylogenetics work, we adopt a Bayesian approach to inference. We fit the model to the data from each source in the WWTP separately using Hamiltonian Monte Carlo (HMC) methods. More specifically, we utilise `Stan` which automatically implements HMC given a Bayesian model.

1.2 Outline

This thesis is split into three parts: one for each application area and one to summarise the conclusions of both applications. Before focussing on each area, Chapter 2 introduces terminology for both topics. The chapter provides an introduction to phylogenetics, where standard models of DNA evolution and their drawbacks are discussed. An introduction to metagenomics is also given, where existing approaches to time-series metagenomics are described. Furthermore, we show that a generalised Lotka-Volterra model, which is often used in time-series metagenomics, can be approximated with a VAR(1) model. Later, we use a VAR(1) model for our data, since linear models are easier to fit. In this chapter, we also give a brief overview of the data and the motivation for modelling such data, along with the associated aims and challenges. The chapter concludes with necessary background information regarding Bayesian inference, including MCMC and HMC.

Part I covers phylogenetics and comprises three chapters. The work for these three chapters appears in Hannaford *et al.* (2020). Chapter 3 describes phylogenetic models that facilitate root inference. We discuss existing non-homogeneous models and non-reversible models. We also introduce Lie-Markov models for DNA evolution and we describe two non-reversible Lie Markov models, the RY5.6b and RY8.8a models, including analytic forms for the stationary probabilities of each model. In addition to this, we identify a potential drawback of the homogeneous RY5.6b model, namely its additive structure, which makes it unsuitable for modelling evolutionary processes in some circumstances.

Chapter 4 concerns non-homogeneous Lie Markov models. We derive a pair of non-

homogeneous, non-stationary, non-reversible phylogenetic models by extending the two non-reversible Lie Markov models described in Chapter 3 in a parsimonious way. This is done by allowing each branch of the tree to have a different RY5.6b rate matrix (in the first model) or a different RY8.8 rate matrix (in the second model). We specify our prior for each model, which includes a prior for the branch-specific parameters that allows borrowing of strength between branches. We conclude the chapter with a description of posterior inference and give a description of the MCMC scheme for generating samples from the posterior.

In Chapter 5, we apply our non-homogeneous Lie Markov models to simulated data and biological data. The simulations assess identifiability of the root position and parameters in a number of controlled settings. We investigate the extent to which root inference depends upon the dimensions of the data. We demonstrate that the root position can be identified from the likelihood of our non-homogeneous models. Additionally, we explore the effect of different topologies and branch lengths and show that root inference for the non-homogeneous RY8.8 model remains strong, even when there is prior-data conflict (resulting from an unbalanced rooted topology). Finally, we consider an application to a real biological data set, the *Drosophila* data set, where many models fail to identify a plausible root position. Using marginal likelihood estimates, we compare the fit of our non-homogeneous models with their homogeneous counterparts, a standard (homogeneous, stationary and reversible) model and a non-homogeneous, non-stationary version of this model. We show that our non-homogeneous RY8.8 model is able to recover a biologically credible rooted tree and has the best fit to the data.

Part II covers time-series metagenomics and also comprises three chapters. In Chapter 6, we present the exploratory data analysis and a more detailed description of each data set is given. We investigate possible relationships and patterns within each data set and across the different data sets. We find evidence of seasonality in our data.

In Chapter 7, we develop our Bayesian hierarchical model for the time-series metagenomics data. In this chapter, we first discuss clustering methods to address the issue of high dimensionality in our data. Due to the evidence of seasonality that we find in Chapter 6, we use a seasonal phase-based clustering approach. We introduce a Bayesian hierarchical VAR(1) model, before adapting it to our data. The assumption that the autoregressive matrix is sparse leads to a discussion of three different shrinkage priors: the spike and slab, the horseshoe and the regularised horseshoe. We give a complete specification of our chosen prior, the regularised horseshoe, including a hyperprior for the global shrinkage parameter that allows us to incorporate prior information on the number of non-zero autoregressive coefficients. This is an extension of the work of Piironen & Vehtari (2017), who originally derived the hyperprior for multiple linear regression with a univariate response vector. To aid selection of a suitable prior, we test these priors in

a simulation study. We show that the regularised horseshoe prior is robust to varying degrees of sparsity and thus a suitable choice for our prior for the autoregressive matrix. Following this, we discuss different error structures for our model, before introducing a time varying mean, which also incorporates the chemical and environmental data. This chapter concludes with a description of posterior inference.

In Chapter 8, we apply the model to our data and present the results of our analyses. We draw conclusions from the posterior distributions of the model parameters, provide potential biological explanations for some of our findings and give valuable insights into the WWTP. Based on these findings, we remark upon the suitability of our model for these data. There is evidence to suggest that complex, non-linear relationships are present in the WWTP data, so we suggest that a generalised Lotka-Volterra model may be more suitable for this particular application.

Finally, Part III has only one chapter, Chapter 9, where we summarise the contributions and conclusions of this thesis and suggest topics for future work.

Chapter 2

Background

2.1 Introduction to phylogenetics

2.1.1 Statistical phylogenetics

Phylogenetics is the study of evolutionary history and relationships among individuals or groups of organisms, for example, different species. Evolutionary history is represented through a bifurcating tree called a *phylogenetic tree* or *phylogeny* (Yang, 2006). A tree is an acyclic connected graph. A graph with n vertices is a tree if, and only if, the graph is connected and has $n - 1$ edges. Each vertex v in a graph has a degree $\deg(v)$; it is the number of edges connected to it. The *leaves* (external vertices) of a phylogenetic tree represent *extant* species (species that currently exist) and have degree one. Trees can be rooted or unrooted. For an unrooted tree, all internal vertices have degree three. The same applies for a rooted tree, except for the root vertex, which has degree two. The root represents the *most recent common ancestor* (MRCA) of all species (or taxa) in the tree. It fixes the direction of ancestry and provides a tool for tracing the evolution of important traits along the tree. Therefore, it is fundamental to the biological interpretation of a tree.

Internal vertices represent *speciation events* or *cladogenesis*, which is the splitting of a population into two new *clades* or groups. A clade is a subset of taxa obtained by cutting a rooted tree on a branch and selecting only those leaves which are descendants of the split lineage; or, in biological terms, an ancestor and all its descendants. The corresponding concept for unrooted trees is a *split*, which is a bipartition of the taxa into two disjoint sets, induced by cutting a branch (Bryant, 2003). Phylogenetic trees are leaf labelled with species' names, for example, in Figure 2.1 there are four species, A , B , C and D . The *branches* (edges) of the tree are typically weighted and these branch lengths represent evolutionary time or duration between speciation events. Note that this is generally not proportional to clock time as the rate of evolution can vary from branch to branch. The closer two species are on a tree, the more closely related they are. In Figure 2.1, for

example, B is more closely related to D than it is to C because D and B share a more recent common ancestor than B and C . Discarding leaf labels and branch lengths from a tree gives us the *topology* (*branching pattern*).

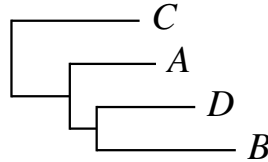


Figure 2.1: Example of a phylogenetic tree with four species.

Data usually consist of an *alignment* of DNA or amino acids from *16S/18S ribosomal RNA* (rRNA) genes. In this thesis, we focus on sequences of DNA, with its four-character alphabet $\Omega = \{A, G, C, T\}$. All living cells have *ribosomes*, which are responsible for biological protein synthesis (Lafontaine & Tollervey, 2001). Each ribosome is composed of a large subunit and a small subunit, and 16S (for prokaryotes)/18S (for eukaryotes) rRNA is a component of the small subunit. The genes coding for rRNA evolve at a slow rate (Woese & Fox, 1977). This, coupled with their universal presence, is why the 16S/18S genes are used in phylogenetics studies. An example of an alignment for n species is

Species 1	G	T	T	A	T	A	C	C	A	T	...
Species 2	A	C	A	T	C	G	C	T	A	G	...
⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	...
Species n	T	T	A	C	C	C	C	T	G	G	...

where each column is referred to as a site and we assume that each site has originated from the same nucleotide of the MRCA.

During reproduction, when an organism passes a copy of its DNA to its offspring, point mutations can occur in molecular sequences. When a point mutation becomes fixed in a population it is referred to as a *substitution*. Now consider the nucleotide (letter) at a single genomic site in the MRCA at the root of the tree. Over time, substitutions may have accumulated at that site such that the corresponding sites in the n taxa at the leaves of the tree are occupied by (possibly) different nucleotides. Equivalently, the substitutions (at site i) over time lead to (possibly) different nucleotides in column i of the alignment. The ensuing assignment of A, G, C or T to each taxon is referred to as the *DNA character* at that site. There are clearly 4^n possible DNA characters for a phylogeny on n taxa and we denote this set by Ω^n . We assume that the molecular sequences of each of the n taxa have been *aligned* such that the columns can be regarded as observations of a DNA character.

2.1.2 Standard models of DNA evolution

Denote by τ a phylogeny, with branch lengths ℓ , representing the evolutionary relationships among a collection of n taxa (species) and suppose that we have an alignment of data $\mathbf{y} = (y_{ij})$, with n rows, representing the n species at the leaves, and m columns, representing genomic sites in the MRCA. Denote by $Y(t) \in \Omega$ the nucleotide at time t at a single genomic site and consider evolution along a single branch of the phylogeny τ . Most phylogenetic models assume that substitutions can be modelled using a *continuous time Markov process* (CTMP) (Yang, 2006), characterised by an instantaneous rate matrix $\mathbf{Q} = (q_{uv})$. Over some interval of time ℓ , the transition probabilities between pairs of nucleotides are obtained by taking the matrix exponential $\mathbf{P}(\ell) = \exp(-\ell\mathbf{Q}')$, where $\mathbf{Q}' = \mathbf{Q}/(-\sum_u q_{uu}\pi_u)$ and $\boldsymbol{\pi} \in \mathcal{S}_4$, $\mathcal{S}_K = \{\mathbf{x} = (x_1, \dots, x_K) : x_i \geq 0 \forall i, \sum x_i = 1\}$, is the stationary distribution of the process, satisfying $\boldsymbol{\pi}\mathbf{Q} = \mathbf{0}^T$. This rescaling of the rate matrix allows the branch length ℓ to be interpreted as the expected number of substitutions per site. The (u, v) -element in $\mathbf{P}(\ell)$ is the probability of transitioning from nucleotide u to nucleotide v along a branch of length ℓ , $p_{uv}(\ell) = \Pr\{Y(\ell) = v | Y(0) = u\}$ for any $u, v \in \Omega$.

Standard phylogenetic models are typically based on three assumptions. The first assumption is *homogeneity* which assumes a single instantaneous rate matrix characterises the evolutionary process along every branch of the tree. The second assumption is *stationarity*, that is that the CTMP is in its stationary distribution $\boldsymbol{\pi}$ and so $\boldsymbol{\pi}$ is also the distribution at the root. It implies that the probability of each nucleotide does not change over time. The final assumption is *reversibility*, which is the assumption that the CTMP is time-reversible, that is

$$\pi_u p_{uv}(\ell) = \pi_v p_{vu}(\ell) \quad (2.1)$$

for all $u, v \in \Omega$ (Felsenstein, 1981). The instantaneous rate matrix \mathbf{Q} of a homogeneous, reversible process can be factorised as $\mathbf{Q} = \mathbf{R}\mathbf{\Pi}$ where $\mathbf{\Pi} = \text{diag}(\boldsymbol{\pi})$ is the diagonal matrix of stationary probabilities and $\mathbf{R} = (r_{ij})$ is the symmetric matrix of *exchangeability* parameters with $r_{ij} = r_{ji} \geq 0$ for $i \neq j$. We refer to a rate matrix as reversible if it permits a factorisation of this form.

In the class of reversible rate matrices, the most general is that of the general time reversible (GTR) model (Tavaré, 1986), with six distinct exchangeability parameters, whose rate matrix is given by

$$\mathbf{Q} = \begin{pmatrix} * & \kappa_1\pi_2 & \kappa_2\pi_3 & \kappa_3\pi_4 \\ \kappa_1\pi_1 & * & \kappa_4\pi_3 & \kappa_5\pi_4 \\ \kappa_2\pi_1 & \kappa_4\pi_2 & * & \kappa_6\pi_4 \\ \kappa_3\pi_1 & \kappa_5\pi_2 & \kappa_6\pi_3 & * \end{pmatrix},$$

where the values of each $*$ ensure that the rows sum to zero. Other commonly used models

are then derived as special cases. For example, ordering the nucleotides as **A**, **G**, **C**, **T**, which we will do for any rate matrix henceforth, the HKY85 model (Hasegawa *et al.*, 1985) is a special case, whose rate matrix Q is given by

$$Q = \begin{pmatrix} * & \kappa\pi_2 & \lambda\pi_3 & \lambda\pi_4 \\ \kappa\pi_1 & * & \lambda\pi_3 & \lambda\pi_4 \\ \lambda\pi_1 & \lambda\pi_2 & * & \kappa\pi_4 \\ \lambda\pi_1 & \lambda\pi_2 & \kappa\pi_3 & * \end{pmatrix}, \quad (2.2)$$

where the values of each $*$ ensure that the rows sum to zero. In this model, the reduction in the number of exchangeability parameters from six to two is biologically motivated, allowing transitions (substitutions between purines – **A** and **G** – and between pyrimidines – **C** and **T**) to occur at a different rate to transversions (substitutions between a pyrimidine and a purine). Transitions often occur at higher rates than transversions (Yang, 2006), due to the two purines having a similar chemical structure, as do the two pyrimidines (Squartini & Arndt, 2008).

To prevent arbitrary rescaling of the rate matrix Q in the transition matrix $P(\ell) = \exp(-\ell Q)$, an identifiability constraint is typically imposed. For reversible rate matrices, this often entails setting one exchangeability parameter equal to one (Zwickl & Holder, 2004) so that the others can be interpreted as relative propensities for change. For example, fixing $\lambda = 1$ in the HKY85 model in (2.2) means that κ is interpreted as the transition-transversion rate ratio. Henceforth, we drop the prime on the normalised rate matrix Q' for brevity.

2.1.3 Likelihood of a phylogenetic tree

In the next section, we discuss the drawbacks of standard phylogenetic models, which relate to the assumption that the evolutionary process is stationary and reversible. A major issue is that these assumptions render the likelihood of a tree invariant to root position. In this section, we define the likelihood of a tree and describe a useful algorithm to calculate the likelihood of a tree.

First, we consider a rooted tree τ with a single branch b of length ℓ_b . Denote by x the observed nucleotide of a single site of DNA sequence and let i be the unobserved ancestral state at the root. Assuming the process is in its stationary distribution $\boldsymbol{\pi}$, the probability that the nucleotide at the root has value i is π_i . The likelihood of the tree is given by $\pi(x|\tau) = \sum_i \pi_i p_{ix}(\ell_b)$, where $p_{ix}(\ell_b)$ is the probability of transitioning from ancestral state i to the observed state x at the leaf.

Now consider a rooted tree τ with vertices V and edges B , where the form of each edge is $b = (u, v)$ and $u, v \in V$. We denote by $X(i)$ a nucleotide at vertex i which is

only observed at the leaves and we let r be the root of the tree. If the nucleotides at the internal vertices are known, the likelihood of the tree is given by

$$p(x|X, \tau) = \pi_{X(r)} \prod_{b=(u,v)} p_{X(u),X(v)}(\ell_b),$$

where ℓ_b is the length of edge b . This is the probability associated with the root vertex multiplied by the product of the transition probabilities associated with every edge of the tree. However, in practice, the nucleotides at the internal vertices are unknown. Therefore, we average over all possible nucleotides (at the internal vertices) to obtain the likelihood. Thus, the likelihood is given by

$$\pi(x|\tau) = \sum_X \pi_{X(r)} \prod_{b=(u,v)} p_{X(u),X(v)}(\ell_b). \quad (2.3)$$

The sum is taken over all functions X from the vertices to Ω such that $X(i)$ matches data x_i at leaf vertices. This sum can be computed efficiently using a post-order traversal of the tree called Felsenstein's pruning algorithm (Felsenstein, 1973, 1981). The algorithm works by calculating the conditional probabilities at each vertex of the tree recursively from the leaves of the tree towards the root.

First, let $d_j(X_i)$ denote the character at the vertex which is the left ($j = 0$) or right ($j = 1$) immediate descendant of vertex i . Let $l(X_i) = \{l_0(X_i), l_1(X_i)\}$ denote the characters at the leaf vertices which are descendants of index i . They are partitioned according to whether they result from the left ($j = 0$) or right ($j = 1$) split at i . Then the algorithm is as follows:

1. Initialise at the leaves. For each leaf vertex i compute $\Pr(X_i = x_i | X_i = k) = \delta_{x_i, k}$, for $k \in \Omega$, where $\delta_{ab} = 1$ if $a = b$ and $\delta_{ab} = 0$ otherwise.
2. Follow the branching pattern from the leaves to the root, applying the recursion

$$\Pr(l(X_i) | X_i = k) = \Pr(l_0(X_i) | X_i = k) \Pr(l_1(X_i) | X_i = k)$$

for each internal vertex i , where

$$\Pr(l_j(X_i) | X_i = k) = \sum_m \Pr(l_j(X_i) | d_j(X_i) = m) \Pr(d_j(X_i) = m | X_i = k).$$

3. Let vertex r be the root. Compute the likelihood through

$$\Pr(l(X_r)) = \sum_k \Pr(l(X_r) | X_r = k) \Pr(X_r = k).$$

2.1.4 Drawbacks of standard models

The assumptions of stationarity and reversibility offer mathematical convenience. However, unfortunately, conditional on any particular (unrooted) tree, a mathematical consequence is that the likelihood does not depend on the root position (Felsenstein, 1981; Yang, 2006). We now consider a simple proof adapted from Felsenstein (1981) and Cherlin (2016). Suppose we have a tree with only two taxa, where we denote the leaves as l_0 and l_1 , and their characters are **A** and **C**, respectively. Furthermore, we suppose that there are two different positions for the root of this tree at two vertices U and V . The distances between U and l_0 , U and V , and V and l_1 are a , b and c , respectively. The probability of observing the tree with the root at vertex U is

$$L_1 = \sum_{U \in \Omega} \pi_U p_{UA}(a) p_{UC}(b+c)$$

and the probability of observing the tree with the root at the vertex V is

$$L_2 = \sum_{V \in \Omega} \pi_V p_{VA}(a+b) p_{VC}(c). \quad (2.4)$$

However, $p_{VA}(a+b)$ in (2.4) can be written as $\sum_{U \in \Omega} p_{VU}(b) p_{UA}(a)$. This is because the transition probability matrix P is characterised by the Chapman-Kolmogorov equations:

$$p_{ij}(t_1 + t_2) = \sum_{k \in \Omega} p_{ik}(t_1) p_{kj}(t_2).$$

This means that the probability of transitioning from state i (nucleotide) to state (nucleotide) j is a sum of probabilities of changing from state i to the intermediate state (nucleotide at the vertex) k , and then from the intermediate state k to the state j . The sum is over all possible intermediate states $k \in \Omega$. From the property of reversibility in (2.1),

$$p_{VU}(b) = \frac{1}{\pi_V} p_{UV}(b) \pi_U,$$

which means that

$$p_{VA}(a+b) = \sum_{U \in \Omega} \frac{1}{\pi_V} p_{UV}(b) \pi_U p_{UA}(a). \quad (2.5)$$

Substituting (2.5) into (2.4) gives

$$L_2 = \sum_{V \in \Omega} \pi_V \sum_{U \in \Omega} \frac{1}{\pi_V} p_{UV}(b) \pi_U p_{UA}(a) p_{VC}(c).$$

According to the Chapman-Kolmogorov equation, $\sum_{V \in \Omega} p_{UV}(b)p_{VC}(c)$ corresponds to $p_{UC}(b+c)$, so

$$L_2 = \sum_{U \in \Omega} \pi_U p_{UA}(a) p_{UC}(b+c) = L_1.$$

Therefore, we have shown that, regardless of the root position, the likelihood of the tree is the same.

Since the likelihood does not depend on the root position, traditionally, other methods have therefore been used to root evolutionary trees. The most common strategy, called *outgroup rooting*, requires inclusion of data from a set of taxa (the *outgroup*) which are known to have evolved outside the subtree of interest (the *ingroup*) (Yang, 2006). The root can then be placed on the branch connecting the outgroup to the rest of the tree. For example, Jarvis *et al.* (2014) investigated the evolutionary history of *Neoaves* (modern birds) and used two birds (an ostrich and a tinamou) that are not considered modern birds as an outgroup. Phillips *et al.* (2006) used *Ameridelphia* (marsupials that live in the Americas) as an outgroup when resolving interrelations among Australian marsupials. Unfortunately, this approach to rooting can be problematic when the outgroup is only very distantly related to the ingroup. In such cases, model assumptions become increasingly dubious and the ensuing model misspecification can result in the identification of spurious relationships with the outgroup. When a distantly related outgroup is used, there will be a long branch between the ingroup and outgroup, which can result in *long branch attraction* (Felsenstein, 1978; Tarrío *et al.*, 2000; Williams *et al.*, 2015). This is when long branches in a tree group together regardless of their true evolutionary relationships. Therefore, use of a distantly related outgroup can result in misleading inferences about the ingroup. Additionally, outgroup rooting can not be used to resolve the *universal tree of life* (the phylogenetic tree for all species on Earth) as no obvious outgroup exists (Yang, 2006; Cherlin *et al.*, 2017).

An alternative strategy for rooting, which has received relatively little attention in the phylogenetic literature, is to take a model-based approach and draw inference from a likelihood that depends on the position of the root. As the proof above shows, the likelihood's invariance to root position is a direct consequence of the assumptions of reversibility and stationarity. Therefore, relaxing at least one of these assumptions in a model results in a likelihood that does depend on the root position. Additionally, the assumptions of stationarity and reversibility typically do not hold up under the scrutiny of biological examination. For example, in a stationary process, time-reversibility implies that the direction of time is unimportant. Yet the mechanism by which substitutions in DNA occur is very complex, comprising processes of point mutation and fixation, and encompassing the effects of selection, and so on. Whilst most physical processes are time-reversible, when these sub-processes combine to produce substitutions in DNA sequences,

their complex interaction makes time-reversibility highly questionable. Correspondingly, analysis of biological data often provides evidence to rebut the reversibility assumption (Squartini & Arndt, 2008).

Similarly, there are various biological theories that are discordant with a stationary evolutionary process, for example mutational biases in DNA replication enzymes are thought to differ across the domains of life (Lind & Andersson, 2008). The assumption is also easy to challenge empirically. In any particular analysis, if the taxa had evolved according to a stationary process, one would expect the sequence composition of each taxon to be approximately the same. Yet, in analyses of deep phylogenies, the GC-content (the total proportion of G and C nucleotides) of 16S ribosomal RNA, the most widely used gene in phylogenetic analysis, varies from 45% to 74% across the diversity of sampled *Bacteria*, *Archaea* and eukaryotes (Cox *et al.*, 2008). Similarly, heterogeneity in sequence composition has also been observed in much more recent species radiations, for example, Morgan *et al.* (2013) highlighted that conflicting hypotheses regarding the root of the placental mammals can be attributed to models failing to account for variations in nucleotide composition. They summarised a number of factors that could cause these variations, such as diet, body size and disease. Phillips *et al.* (2006) showed that the third *codon* positions of the sequences were the source of compositional heterogeneity in marsupials. (A codon is a triplet of nucleotides that corresponds to a specific amino acid.) They found that by removing these codon positions or using *Y-coding* (replacing pyrimidines with a Y) or *RY-coding* (replacing pyrimidines with a Y and replacing purines with an R) reduced compositional bias in their data sets, thus allowing them to determine interrelations of three major marsupial groups. However, in some cases, removing third codon positions can be undesirable, as the first and second positions alone may contain insufficient information for resolution of short timescale relationships (Tarrío *et al.*, 2000). Heterogeneity in sequence composition is also evident in other animals, such as birds (Braun & Kimball, 2002; Phillips *et al.*, 2009) and *paraneopteran* insects (Li *et al.*, 2015). (Paraneoptera are a *superorder* of insects containing over 120000 species, and a superorder is a taxonomic category.) As a consequence, in addition to facilitating root inference, models that relax the restrictive assumptions of stationarity and/or reversibility also provide opportunities to incorporate further biological realism. Existing models that facilitate root inference are discussed in Chapter 3, before we describe our own models in Chapter 4. In Chapter 5, we test our models' ability to infer root position with simulated data and a biological data set that exhibits compositional heterogeneity.

2.2 Introduction to metagenomics

2.2.1 Statistical metagenomics

A *meta-analysis* statistically combines separate analyses and *genomics* is the wide ranging analysis of an organism's genetic material. *Metagenomics* (also called *community genomics*, *environmental genomics* and *population genomics*) can be described as the culture-independent analysis of microbial communities, which involves directly isolating genomic DNA from an environment and then cloning it in easily cultivable microorganisms (Joshi *et al.*, 2014). Gilbert *et al.* (2011) described metagenomics as directly isolating DNA from the environment and using it to characterise the taxonomy and function of the biological community in that ecosystem. In simple terms, metagenomics is the study of microbial genetic material sequenced directly from environmental samples (Jonsson *et al.*, 2016). 16S rRNA (see Section 2.1.1) sequencing, which is used in phylogenetics, where the 16S gene is targeted and read in a sample, is often used in metagenomic studies (Fisher & Mehta, 2014). However, referring to 16S rRNA sequencing as metagenomics is technically a misnomer, as the entire genomic content of a sample is not targeted (Quince *et al.*, 2017).

There are many applications of metagenomics and microorganisms that influence human life in many ways, for example, the gut *microbiome* is thought to have a role in obesity (Davis, 2016). The term microbiome refers to the community of microorganisms living within a particular environment. The microbiome of the skin prevents colonisation by pathogenic bacteria (Byrd *et al.*, 2018). However, changes in this microbiome are associated with common skin diseases, such as acne and eczema, and with chronic wounds. More pertinent to this thesis, microorganisms play a fundamental role in the treatment of wastewater (Cyzdik-Kwiatkowska & Zielińska, 2016), which is discussed in more detail in Section 2.2.4.

Joshi *et al.* (2014) summarised some important applications and outcomes from metagenomic studies. They placed applications into two categories: ecological inferences from microbial diversity estimation and biotechnology. Knowledge of how microorganisms live within a community is important. Microbial biologists seek to understand how nutrients are acquired and how energy is produced. In addition to this, they aim to learn about the community structure by understanding the physiology of the microorganisms and the roles they have within their ecosystem. Diversity analysis is a popular area of research too, with developments in metagenomics helping to overcome existing barriers to estimating microbial diversity. *Paleogenomics* is also considered an important application of metagenomics, where phylogenetic relationships are resolved between extinct and extant animals by extracting ancient DNA from the environment. Regarding biotechnology, metagenomics plays a crucial role, with the opportunity to revolutionise existing trends for

industrial production of metabolites. Already microbes have been used for the production of a multitude of products. Moreover, there is the potential to discover novel antibiotics and new antimicrobial substances, and to learn about antibiotic resistance genes in order to be able to combat them.

In metagenomic studies, DNA sequences are extracted from a microbial sample. Then, decision rules are made for grouping observed sequences into classes. Often, microbial biologists use *operational taxonomic units* (OTUs) to group similar sequences together using a clustering algorithm, where typically an OTU contains sequences that are not more than 3% different from each other (Bunge *et al.*, 2014; Xia *et al.*, 2018). It is crucial to note that OTUs do not exactly represent different species but can be thought of as pragmatic proxies for different species. Thus, the primary data structure in metagenomic studies is an OTU table which is a matrix of counts, where the rows represent different OTUs and the columns represent different samples (Xia *et al.*, 2018). Note that this is different from other fields, where rows usually represent different samples, and one can easily obtain an OTU table that follows this more “standard” convention by simply transposing it. OTU tables are high-dimensional and are typically sparse, which means they contain a lot of zeroes. In this thesis, the data are primarily counts of OTUs over time (see Section 2.2.3), so the samples are different (ordered) time points. In the following section, we discuss some of the current approaches to modelling time-series metagenomics data.

2.2.2 Time-series metagenomics

Due to recent advances in sequencing technology, there has been a rise in longitudinal studies of microbial communities from a large range of environments. Unique ecological insights into response to perturbations and community stability can be gained from such studies (Faust *et al.*, 2015). An example of a perturbation could be the introduction of an antibiotic to a gut microbiome, or a chemical spill in an aquatic environment. Typically, microbial communities tend to evolve towards a stable state, although changes in the community state can occur due to changes in the environment, such as these perturbations. In addition to this, there are complex non-linear interactions among microbes, which can result in alternative stable states. The existence of complex interactions among different microbes and between different microbes and their environment contributes significantly to microbial dynamics.

Community dynamics are often described by the generalised Lotka-Volterra (gLV) differential equations, where changes in microbial counts are modelled as a function of taxon-specific growth rates and pairwise interactions. For example, Mounier *et al.* (2008) used a generalised Lotka-Volterra model to identify interactions in species within a cheese microbial community. The Lotka-Volterra equations (Lotka, 1926; Volterra, 1926) are commonly used to describe the dynamics of biological systems with two species, one of

which is the predator, and the other is the prey. Let $y_1(t)$ be the number of prey and $y_2(t)$ be the number of some predator at time t . Under a deterministic Lotka-Volterra model, the changes in the populations of the two species over continuous time t follow the pair of equations:

$$\begin{aligned}\frac{d}{dt}y_1(t) &= \alpha y_1(t) - \beta y_1(t)y_2(t), \\ \frac{d}{dt}y_2(t) &= -\delta y_2(t) + \gamma y_1(t)y_2(t)\end{aligned}\tag{2.6}$$

where $\alpha, \beta, \gamma, \delta \in \mathbb{R}^+$ describe the growth rates and interactions of the two species.

The gLV model is an extension of (2.6) that is used to characterise the dynamics of a K -species system, where $K > 2$. Changes in population of species i are described by

$$\frac{d}{dt}y_i(t) = b_i y_i(t) + y_i(t) \sum_{j=1}^K A_{ij} y_j(t),\tag{2.7}$$

where $y_i(t)$ is the population size of species i at time t , b_i is the growth rate of species i in absence of any competition and A is a matrix of pairwise interactions.

Often, in microbiome studies, the problem of finding the growth rates and interaction terms is simplified using linearisation. Now we describe the method based on the description given in Stein *et al.* (2013), but other approaches are possible, for example, non-linear least squares (see Section 1.3.1 of Ramsay *et al.* (2007)). First we note that $\frac{d}{dt}y_i(t)$ can be written as $y'_i(t)$ and that we can divide the differential equation in (2.7) by the abundance $y_i(t)$ of OTU (species) i to give

$$\frac{y'_i(t)}{y_i(t)} = b_i + \sum_{j=1}^K A_{ij} y_j(t),$$

where K is the number of different OTUs. Note that $\frac{y'_i(t)}{y_i(t)} = \frac{d}{dt} \log(y_i(t))$ and that (2.7) can be discretised and approximated using

$$\frac{\log(y_i(t_{k+1})) - \log(y_i(t_k))}{t_{k+1} - t_k} = b_i + \sum_{j=1}^K A_{ij} y_j(t_k),\tag{2.8}$$

for successive time points $k = 1, \dots, N$. Considering all OTUs, each side of (2.8) can be concatenated into a matrix form. The problem of estimating the gLV parameters \mathbf{b} and A is then reduced to linear regression on the differenced log-transformed counts. For example, Stein *et al.* (2013) utilised ridge regression to learn about growth rates, interactions and susceptibility to external perturbations of the intestinal microbial community in mice. Additionally, Fisher & Mehta (2014) used sparse linear regression and bootstrap

aggregation to infer the gLV parameters of microorganisms present in the human gut. Furthermore, they argue against the correlation methods used in *cross-sectional studies* and show that a correlation between two species does not necessarily imply an interaction. A cross-sectional study pools samples from multiple environments (for example, when studying the gut microbiome, samples are taken from different people) and correlations between relative abundances of species are used as proxies for interactions. Bucci *et al.* (2016) took a Bayesian approach (see Section 2.3.1 for a discussion of Bayesian inference) to inference of gLV parameters, making use of a Bayesian lasso (a regularisation method) and Bayesian variable selection, which directly models presence or absence of interactions.

While most authors opt to use gLV models, Gibbons *et al.* (2017) investigated microbial dynamics in the gut by applying a vector autoregressive (VAR) model (see Section 7.2 for a general definition). They used a sparse VAR model, which benefits, in terms of quantification of uncertainty, from explicitly modelling error, unlike gLV-type models. The key difference between the two approaches is that VAR models assume *linear* dynamics, whereas gLV models assume non-linear dynamics. Moreover, under a VAR model, stationarity can be enforced. A Gaussian process $\mathbf{x} = \{x_t\}$ ($t = 1, \dots, N$) is considered stationary if its mean function μ_t does not depend on time and its *autocovariance* function $\gamma(s, t)$ depends on s and t only through the difference $|s - t|$ (Shumway & Stoffer, 2017). The autocovariance function is $\gamma(s, t) = \text{E}[(x_s - \mu_s)(x_t - \mu_t)]$ for all s and t .

A VAR model of order one (VAR(1)) can be regarded as a linear approximation to the non-linear numerical solution of a Lotka-Volterra system, which we shall now demonstrate. First consider a *stochastic* Lotka-Volterra model with two species, which is

$$\begin{aligned} \frac{d}{dt}y_1(t) &= \alpha y_1(t) - \beta y_1(t)y_2(t) + \sigma_1 dw_1(t) \\ \frac{d}{dt}y_2(t) &= -\delta y_2(t) + \gamma y_1(t)y_2(t) + \sigma_2 dw_2(t), \end{aligned} \quad (2.9)$$

where $\alpha, \beta, \gamma, \delta \geq 0$ and w_1, w_2 are independent Wiener processes (or Brownian motions). Let $\mathbf{Y}_t = (Y_{t,1}, Y_{t,2})^T$ denote an appropriate discretisation of the time series of the two species with time step Δt so that $t_N = t_0 + n\Delta t$ for some initial time t_0 , where t_N is the last time point in the time series. The Euler approximation to the solution of the stochastic differential equation driving the Lotka-Volterra model is

$$\begin{aligned} Y_{t+1,1} &= \alpha Y_{t,1} - \beta Y_{t,1}Y_{t,2} + \epsilon_{t,1} \\ Y_{t+1,2} &= -\delta Y_{t,2} + \gamma Y_{t,1}Y_{t,2} + \epsilon_{t,2}, \end{aligned} \quad (2.10)$$

where $\alpha, \beta, \gamma, \delta \geq 0$, $\boldsymbol{\epsilon}_t = (\epsilon_{t,1}, \epsilon_{t,2})^T \sim \text{N}_2(\mathbf{0}, \Sigma)$ and $\Sigma = \text{diag}(\sigma_1, \sigma_2)$. Interactions between the two population sizes $y_{t,1}$ and $y_{t,2}$ arise due to the underpinning differential equations being coupled. In the numerical solution in (2.10), the dependence of \mathbf{Y}_{t+1} on

\mathbf{Y}_t is non-linear.

Now consider a zero-mean VAR(1) model for a discrete time series with two species $\{\mathbf{Y}_t : t = 0, 1, \dots\}$. This can be expressed as

$$\begin{aligned} Y_{t+1,1} &= \alpha Y_{t,1} - \beta Y_{t,2} + \epsilon_{t,1} \\ Y_{t+1,2} &= -\delta Y_{t,2} + \gamma Y_{t,1} + \epsilon_{t,2}, \end{aligned} \tag{2.11}$$

where $\alpha, \beta, \gamma, \delta \in \mathbb{R}$ but we could impose a positivity constraint and $\epsilon_t \sim N_2(\mathbf{0}, \Sigma)$ as above for (2.10). Here the dependence of \mathbf{Y}_{t+1} on \mathbf{Y}_t is linear by construction. However, comparison of (2.11) and (2.10) reveals that the VAR(1) model has a similar structure to the Lotka-Volterra model. Thus, we can regard the VAR(1) as a linear approximation to the non-linear numerical solution of the Lotka-Volterra system. Since linear models tend to be easier to fit, we develop a VAR(1) model for our time-series metagenomics data in Chapter 7.

2.2.3 Data overview

As mentioned in Section 1.1 of Chapter 1, we take a data-driven approach to the metagenomics part of this thesis, where we have specific data sets for which we develop a model. In this section, we give a brief overview of these data sets. A description of how the data are processed is given in Chapter 6, as well as more detailed descriptions of the data and the exploratory analysis.

In this thesis, we have weekly counts of OTUs in a UK-based wastewater treatment plant (WWTP) over (roughly) five years. We have counts from two different tanks, the *activated sludge* (AS) and the *settled sewage* (SS) tanks. After wastewater enters the reactor of a WWTP, it undergoes the physical process of *primary sedimentation*. This is when large solids present in the wastewater are settled out. The wastewater that emerges from the primary sedimentation tank is the settled sewage and is the source material that is fed into an aerated tank. The content of this tank is the activated sludge and is where the wastewater is treated biologically. From the AS we have counts for 9044 different OTUs over 257 time points and from the SS we have 9678 different OTUs over 251 time points. There is a taxonomy table for the OTUs with taxonomic ranks: kingdom, phylum, class, order, family and genus. Additionally, there are weekly measurements of chemical and environmental data. These chemical and environmental covariates are often being changed, namely by changes made to the dilution rate and the solid retention time. Between the two tanks, there are 35 different metrics including temperature, pH, ammonia concentration and chemical concentrations, such as, calcium, iron and lead.

2.2.4 Motivation, aims and challenges

Considered one of the greatest engineering successes, WWTPs are fundamental components of infrastructure and are vital for sustainable urban existence. The AS plays a pivotal role in the treatment of wastewater, as it is responsible for the degradation of organic and inorganic pollutants, and comprises aerobic and anaerobic microorganisms (Shchegolkova *et al.*, 2016). Microbial communities within AS are complicated biosystems with a network of interconnected trophic links. For example, to degrade synthetic substances, such as oil products, or toxicants, many enzymes are required in a multi-stage process. Several species of microorganisms are needed for complete biodegradation. Gaining theoretical understanding of how these large biological systems work could aid in creating better biotechnological solutions more rapidly (Curtis *et al.*, 2003). The microorganisms involved in WWTPs are delicately balanced in a stable state and understanding how to maintain this stable state is paramount to sustaining and improving WWTPs' function (Cyzdik-Kwiatkowska & Zielińska, 2016). Furthermore, methods developed for WWTP microbial communities could be applied to many other microbiomes.

In the metagenomics part of this thesis, the aim is to model the OTU counts in the WWTP over time, with a particular focus on the AS. As in our phylogenetics work, we wish to use a biologically motivated stochastic model that allows for interactions between OTUs. Additionally, we wish to model the chemical and environmental effects on OTUs, as these most likely influence the microbial community. For example, Shchegolkova *et al.* (2016) found that chemical composition of wastewater played a major role in the formation of taxonomic structure of AS.

The data that we wish to model are accompanied by many statistical challenges, which will be addressed in this thesis (see Chapter 7). The first challenge is regarding the dimensions of the data, as clearly the number of different OTUs is much larger than the number of time points. Since we are interested in interactions between different OTUs, fitting joint models to the counts of all OTUs would be very difficult computationally. Furthermore, one must question whether the number of time points would be sufficiently large to detect such interactions (see Section 7.3.1). Secondly, there is the common issue of sparsity in the OTU tables, which is the presence of many zeroes. Failing to account for the excess zeroes may lead to biased parameter estimation and misleading inference (Xia *et al.*, 2018).

In standard metagenomic studies, the goal is often to determine whether OTUs are differentially abundant across samples. In these instances, the different *library sizes* need to be accounted for. The library size or sampling depth is the total number of OTUs *detected* in a sample. Library sizes can vary vastly from sample to sample, which does not reflect true biological variation but differential efficiency of the sequencing method (Weiss *et al.*, 2017). For example, in the AS tank, the smallest library size is 29298 and the

largest is more than ten times larger (692110) and in the SS tank, the library size ranges from 2564 to 542565 (more than 200 times larger). The issue of sparsity is also linked to sampling depth. OTUs may be absent from the majority of samples, but it might also be that the sampling depths were insufficient for these samples (Jonsson *et al.*, 2016).

A common approach to dealing with the issue of differing library sizes is to normalise the data by using proportions (Paulson *et al.*, 2013) or to *rarefy* the counts (McMurdie & Holmes, 2014; Weiss *et al.*, 2017). Rarefying counts works by choosing a minimum library size, say N_{min} , discarding any samples that are below this minimum and subsampling the remaining samples such that they all have size N_{min} . Typically, N_{min} is selected to be the size of the smallest library that is not considered “defective”, but deciding what constitutes a defective sample is not clear cut (McMurdie & Holmes, 2014). After normalisation, hypothesis tests are applied, such as t -tests or Wilcoxon-Mann-Whitney tests (McMurdie & Holmes, 2014; Jonsson *et al.*, 2016; Weiss *et al.*, 2017; Xia *et al.*, 2018). Alternatively, a model-based approach can be taken, for example, the raw counts are modelled by a negative binomial distribution, which explicitly accounts for differences in library sizes through a linear scaling factor (Anders & Huber, 2010; McMurdie & Holmes, 2014). Hypothesis tests can then be based on these parametric assumptions. For example, see Anders & Huber (2010).

In this thesis, we cluster OTUs into a small number of large groups (see Section 7.1) and explicitly allow the mean log absolute abundances for each group to vary over time (across samples). Therefore, the issues described in this section are less relevant in the context of this thesis.

2.3 Bayesian inference

In this thesis, we work in a Bayesian framework, so in the final section of this chapter we present an overview of Bayesian inference; see Gelman *et al.* (2014) for a more comprehensive introduction. In Bayesian inference, all unknown quantities are treated as random variables and uncertainty in inferences based on statistical analyses can be quantified using probability, where probability is interpreted as a “degree of belief” (Farrow, 2005). The first step in a Bayesian analysis is to define a full probability model, that is, a joint probability distribution must be defined for the data and the unknown quantities, which include parameters and missing observations (if applicable). Then, conditioning on the observed data, a *posterior distribution* is calculated, which is the conditional probability distribution of the unobserved quantities, given the observed data. This posterior distribution quantifies uncertainty about the parameters after observing the data and can be interpreted to make inferences. Finally, the fit of the model can be evaluated and, if necessary, the model may be altered to improve the fit.

2.3.1 Bayes theorem

Let $\boldsymbol{\theta} = (\theta_1, \dots, \theta_k)^T$ be a set of parameters, which we wish to make inferences about, and let $\mathbf{y} = (y_1, \dots, y_n)$ be the data. First, assuming that $\boldsymbol{\theta}$ is continuous, the *prior distribution* with density $\pi(\boldsymbol{\theta})$ summarises the prior belief about the parameters. The likelihood $L(\boldsymbol{\theta}|\mathbf{y}) = p(\mathbf{y}|\boldsymbol{\theta})$ contains information from the data \mathbf{y} and is the joint density (continuous \mathbf{y}) or the joint mass function (discrete \mathbf{y}) of the data given the parameters $\boldsymbol{\theta}$. Using Bayes theorem, we can calculate the posterior density, which contains the prior information, updated by the likelihood, as

$$\pi(\boldsymbol{\theta}|\mathbf{y}) = \frac{\pi(\boldsymbol{\theta})p(\mathbf{y}|\boldsymbol{\theta})}{p(\mathbf{y})},$$

where $p(\mathbf{y}) = \int \pi(\boldsymbol{\theta})p(\mathbf{y}|\boldsymbol{\theta})d\boldsymbol{\theta}$ is the *marginal likelihood*. This is a normalising constant which ensures that the posterior density integrates to one. Thus, Bayes theorem is often expressed as

$$\pi(\boldsymbol{\theta}|\mathbf{y}) \propto \pi(\boldsymbol{\theta})p(\mathbf{y}|\boldsymbol{\theta}). \quad (2.12)$$

In the case of discrete parameters $\boldsymbol{\theta}$, we speak of prior and posterior mass functions, rather than densities, and $p(\mathbf{y}) = \sum_{\boldsymbol{\theta}} p(\boldsymbol{\theta})p(\mathbf{y}|\boldsymbol{\theta})$, which is the sum over all possible values of $\boldsymbol{\theta}$.

In most cases, the marginal likelihood $p(\mathbf{y})$ is not available in closed form, and consequently neither is the posterior distribution. However, an exception to this is when the posterior distribution is *conjugate* to the prior distribution, which means that both the prior and posterior distributions have the same parametric form. To illustrate conjugacy, we consider a simple example. Suppose we have a clinical trial that aims to test the success rate of a new drug to reduce blood pressure. In the trial there are n patients and let $\theta \in (0, 1)$ be the probability that a patient's blood pressure is lower after taking the drug. If we let Y be the number of patients whose blood pressure is reduced after taking the drug (the number of successes) then we have $Y \sim \text{Bin}(n, \theta)$. We then assign a prior Beta (α, β) distribution to θ . The parameters of our prior distribution α and β are known as the *hyperparameters*. Now, let us consider the posterior density in (2.12) for θ , that is

$$\begin{aligned} p(\theta|\mathbf{y}) &\propto \theta^{\alpha-1}(1-\theta)^{\beta-1} \times \theta^y(1-\theta)^{n-y} \\ &\propto \theta^{y+\alpha-1}(1-\theta)^{n-y+\beta-1}. \end{aligned}$$

We have $\theta|\mathbf{y} \sim \text{Beta}(y+\alpha, n-y+\beta)$ and clearly the posterior distribution has the same form as the prior distribution. In this simple example, we can see how the data directly updates our beliefs about the parameter θ . Moreover, we have demonstrated the computational ease and simplicity of using a conjugate prior, since the posterior density, which includes the marginal likelihood $p(\mathbf{y})$, can be calculated analytically.

2.3.2 Markov chain Monte Carlo

Unfortunately, in practice, conjugate priors are often either not a realistic option or they are unavailable, and thus the marginal likelihood $p(\mathbf{y})$ cannot be calculated analytically. If the number of parameters is very small, it may be possible to approximate the marginal likelihood by numerical integration. However, for all the models considered in this thesis, the dimension of the parameter space is far too large for this to be computationally tractable. In phylogenetics, for example, calculating the marginal likelihood requires integration over all parameters of the substitution model and the branch lengths, and enumeration over every possible topology. Similarly, as mentioned in Section 2.2.4, we wish to model interactions between (the counts of) clusters of OTUs and interactions between (the counts of) the clusters and the many covariates that represent their environment, which naturally leads to numerous parameters in a model. Clearly for both examples, obtaining the marginal likelihood (and hence posterior distribution) is not a simple task.

A remedy to the issue of analytically intractable posterior distributions is Markov chain Monte Carlo (MCMC), which is an advanced computational technique that can be used to generate samples from complex posterior distributions with many parameters. In MCMC, a Markov chain is constructed whose stationary distribution is the posterior distribution (sometimes referred to as the *target distribution*) (Kass *et al.*, 1998). Starting from an initial point within the support of the posterior, dependent samples are drawn from this distribution by sampling realisations from the Markov chain. Since the posterior distribution is the stationary distribution, the chain must be run for long enough to obtain such samples (Farrow, 2005). In the remainder of this section, we describe two MCMC algorithms, before discussing diagnostic checks that must be made when using such methods.

Metropolis-Hastings algorithm

Our goal is to sample from the posterior distribution $\pi(\boldsymbol{\theta}|\mathbf{y})$. Suppose we have a set of parameters $\boldsymbol{\theta}$ with p components and an arbitrary transition kernel $q(\boldsymbol{\theta}^*|\boldsymbol{\theta})$, which we call the *proposal* distribution. This is the probability density (for continuous $\boldsymbol{\theta}$) or mass function (for discrete $\boldsymbol{\theta}$) of moving from state $\boldsymbol{\theta}$ to $\boldsymbol{\theta}^*$. The steps (Gelman *et al.*, 2014) of the Metropolis-Hastings algorithm (Hastings, 1970) are as follows:

1. Initialise the iteration counter to $j = 1$, and initialise the chain with $\boldsymbol{\theta}^{(0)}$ chosen from somewhere in the support of $\pi(\boldsymbol{\theta}|\mathbf{y})$.
2. Generate a proposed value $\boldsymbol{\theta}^* \sim q(\boldsymbol{\theta}^*|\boldsymbol{\theta}^{(j-1)})$.

3. Evaluate the *acceptance probability* $\alpha(\boldsymbol{\theta}^*|\boldsymbol{\theta}^{(j-1)})$ of the proposed move, which is

$$\alpha(\boldsymbol{\theta}^*|\boldsymbol{\theta}^{(j-1)}) = \min \left\{ 1, \frac{\pi(\boldsymbol{\theta}^*|\mathbf{y}) q(\boldsymbol{\theta}^{(j-1)}|\boldsymbol{\theta}^*)}{\pi(\boldsymbol{\theta}^{(j-1)}|\mathbf{y}) q(\boldsymbol{\theta}^*|\boldsymbol{\theta}^{(j-1)})} \right\}.$$

4. Set $\boldsymbol{\theta}^{(j)} = \boldsymbol{\theta}^*$ with probability $\alpha(\boldsymbol{\theta}^*|\boldsymbol{\theta}^{(j-1)})$. Otherwise set $\boldsymbol{\theta}^{(j)} = \boldsymbol{\theta}^{(j-1)}$.

5. Set $j = j + 1$ and return to step 2.

The posterior density $\pi(\boldsymbol{\theta}|\mathbf{y})$ only appears in the acceptance probability as a ratio, and so the algorithm can be used when the posterior density is only known up to proportionality, as in (2.12). Thus, the acceptance probability can be expressed as

$$\alpha(\boldsymbol{\theta}^*|\boldsymbol{\theta}^{(j-1)}) = \min \left\{ 1, \frac{\pi(\boldsymbol{\theta}^*) p(\mathbf{y}|\boldsymbol{\theta}^*) q(\boldsymbol{\theta}^{(j-1)}|\boldsymbol{\theta}^*)}{\pi(\boldsymbol{\theta}^{(j-1)}) p(\mathbf{y}|\boldsymbol{\theta}^{(j-1)}) q(\boldsymbol{\theta}^*|\boldsymbol{\theta}^{(j-1)})} \right\}.$$

Componentwise transitions

Defining a suitable proposal distribution for $\boldsymbol{\theta} = (\theta_1, \dots, \theta_p)^T$ may be difficult. Furthermore, for some problems, *full conditional* distributions (FCDs) may be available to sample from for a subset of the components of $\boldsymbol{\theta}$. Let $\pi(\theta_i|\theta_1, \dots, \theta_{i-1}, \theta_{i+1}, \dots, \theta_p, \mathbf{y})$ be the FCD for the i -th component of $\boldsymbol{\theta}$. This is the density of θ_i given all other components of $\boldsymbol{\theta}$ and the data \mathbf{y} . We can perform componentwise transitions as follows:

1. Initialise the iteration counter to $j = 1$ and initialise the chain with $\boldsymbol{\theta}^{(0)} = (\theta_1, \dots, \theta_p)^T$.
2. Obtain a new value $\boldsymbol{\theta}^{(j)}$ from $\boldsymbol{\theta}^{(j-1)}$ by successively generating values:
 - $\theta_1^{(j)} \sim \pi(\theta_1|\theta_2^{(j-1)}, \dots, \theta_p^{(j-1)}, \mathbf{y})$ using a Metropolis-Hastings step with proposal distribution $q_1(\theta_1^*|\theta_1^{(j-1)})$.
 - $\theta_2^{(j)} \sim \pi(\theta_2|\theta_1^{(j-1)}, \theta_3^{(j-1)}, \dots, \theta_p^{(j-1)}, \mathbf{y})$ using a Metropolis-Hastings step with proposal distribution $q_2(\theta_2^*|\theta_2^{(j-1)})$.
 - \vdots
 - $\theta_p^{(j)} \sim \pi(\theta_p|\theta_1^{(j-1)}, \dots, \theta_{p-1}^{(j-1)}, \mathbf{y})$ using a Metropolis-Hastings step with proposal distribution $q_p(\theta_p^*|\theta_p^{(j-1)})$.

3. Set $j = j + 1$ and return to step 2.

The Gibbs sampler

When all the FCDs are available to sample from, we can obtain the Gibbs sampler (Geman & Geman, 1984). This is a special case of the Metropolis-Hastings algorithm (Gelman *et al.*, 2014), where for each iteration j we perform p Metropolis-Hastings steps (one for each component of $\boldsymbol{\theta}$), where the acceptance probability is one.

1. Initialise the iteration counter at $j = 1$ and initialise the state of the chain to $\boldsymbol{\theta}^{(0)} = (\theta_1^{(0)}, \dots, \theta_p^{(0)})^T$.
2. Obtain a new value $\boldsymbol{\theta}^{(j)}$ from $\boldsymbol{\theta}^{(j-1)}$ by successively sampling from the conditional distributions:
 - $\theta_1^{(j)} \sim \pi(\theta_1 | \theta_2^{(j-1)}, \dots, \theta_p^{(j-1)}, \mathbf{y})$.
 - $\theta_2^{(j)} \sim \pi(\theta_2 | \theta_1^{(j-1)}, \theta_3^{(j-1)}, \dots, \theta_p^{(j-1)}, \mathbf{y})$.
 - \vdots
 - $\theta_p | \theta_p^{(j)} \sim \pi(\theta_p | \theta_1^{(j-1)}, \dots, \theta_{p-1}^{(j-1)}, \mathbf{y})$.
3. Set $j = j + 1$ and return to step 2.

It might be the case that for some components of $\boldsymbol{\theta}$ it is easy to directly simulate from the FCD, whereas for other components this may be more difficult. In such cases, we can use the *Metropolis-within-Gibbs* algorithm, where components can be updated either with a Gibbs step (updating directly from the FCD) or with a Metropolis-Hastings step.

Diagnostic checks

The simulated draws which are obtained from running such MCMC algorithms described above should (approximately) be from the posterior distribution and are used to make inferences. However, several problems can occur when using MCMC simulations, namely, issues with convergence and *mixing*. Here we discuss different diagnostic checks to assess convergence and mixing.

Often simulations can take a long time to converge to the stationary distribution and hence generate (approximate) posterior draws. Before the chain has converged, we have values that are not sampled from the posterior. These early draws are referred to as “burn-in” or “warm-up” and are discarded, that is, we do not use them to make inferences (Geyer, 2011). The number of burn-in iterations depends upon the convergence rate. Typically, convergence is assessed visually via a trace plot. Figure 2.2 shows two trace plots where convergence has and has not been achieved. If a chain has converged, the trace plot will show the sampler stabilised and moving around the mode of the distribution, as shown in Figure 2.2a. If a chain has not converged, there will be a trend over the sample space in the

trace plot (Congdon, 2019), as shown in Figure 2.2b. Unfortunately, in some instances, a chain may become stuck at a local mode without exploring the entire posterior parameter space. Therefore, it can be useful to initialise multiple chains at different starting points and run them simultaneously (Gelman *et al.*, 2014; Vehtari *et al.*, 2020). The trace plots of the output from each chain can be compared. If the trace plots for the chains fail to overlap then this suggests lack of convergence. Another method to detect multimodality in the posterior distribution is to plot the marginal posterior densities for each parameter. Multimodality in the posterior might mean that the chain(s) must be run for a long time to ensure that the parameter space has been fully explored. In this thesis, we make use of these visual checks described here to ensure that we have representative samples of our posterior distributions of interest and can make accurate inferences.

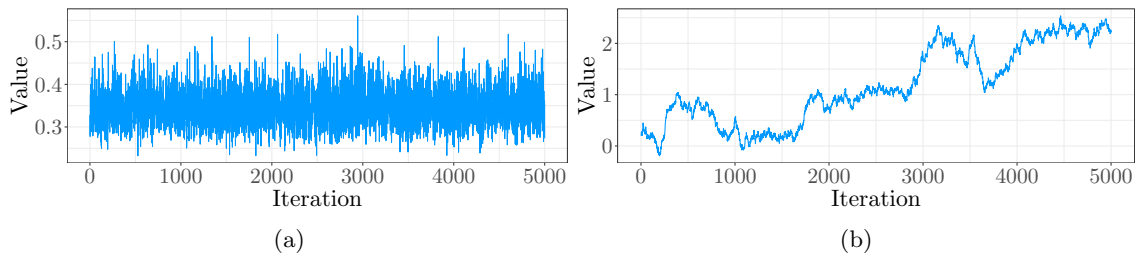


Figure 2.2: Examples of trace plots showing (a) a chain that has converged and (b) a chain that has not converged.

A sampler suffers from poor mixing when it explores the support of the posterior slowly. An indication of poor mixing is high autocorrelation over iterations, that is strong correlation between successive draws of parameter values in a chain. Exploration of the parameter space takes longer if the chain is highly correlated, so, if the sampler is not run for long enough only part of the parameter space will be explored. To assess mixing, the autocorrelation function can be computed and plotted against the lag (Geyer, 2011). The (sample) autocorrelation at lag k for a parameter θ_p is

$$\hat{\rho}_{pk} = \frac{\gamma_{pk}}{\gamma_{p0}}, \quad (2.13)$$

where γ_{p0} is the posterior variance $\text{Var}(\theta_p|\mathbf{y})$ and γ_{pk} is the autocovariance at lag k $\text{Cov}(\theta_p^{(j)}, \theta_p^{(j+k)}|\mathbf{y})$ (Congdon, 2019). The values of the autocorrelation approach zero as the lag increases if the chain is mixing well. Sometimes if an MCMC chain is mixing poorly, assessment of convergence using graphical checks described previously becomes more difficult. To address this problem, a useful tool is *thinning*, that is for some integer $k > 1$ we retain every k -th iteration. Thinning is also useful for reducing computational overheads in storing the MCMC output (Gelman *et al.*, 2014).

In the Metropolis-Hastings algorithm, the proposal distribution $q(\theta^*|\theta)$ for a parameter

θ can be defined so that it is centred on the current value of the parameter (for example, see Section 4.3.1). If the variance of the proposal distribution is large, many of the proposed values will be rejected, as they will be far from the current value. This means that the process will remain in the same state for a long time. Whereas, if the variance is small, most of the proposed values will be accepted but the sequence of accepted states will be close together. Both cases lead to poor mixing and high autocorrelation in the MCMC chain. However, the variance can be *tuned* to achieve a desirable acceptance rate ($\sim 23\%$), which optimises the efficiency of the sampler (Gelman *et al.*, 2014).

In addition to causing slow exploration of the parameter space, autocorrelation over iterations increases the uncertainty of estimates of posterior summaries, such as the mean or variance (Geyer, 2011; Stan Development Team, 2021). The *effective sample size* (ESS) (Kass *et al.*, 1998) is a measurement of how much the autocorrelation over iterations increases this uncertainty and is the equivalent number of uncorrelated observations that a chain represents. The ESS for a parameter θ_p is

$$N_{ESS} = N / \left[1 + 2 \sum_{k=0}^{\infty} \hat{\rho}_{pk} \right], \quad (2.14)$$

where $\hat{\rho}_{pk}$ is the (sample) autocorrelation at lag k as defined in (2.13) and N is the number of MCMC samples. Say we are interested in the posterior mean of a parameter which is estimated by $\hat{E}[\theta_p | \mathbf{y}] = \bar{\theta}_p = \frac{1}{N} \sum_{j=1}^N \theta_p^{(j)}$ (the sample mean of the MCMC samples). As discussed in Kass *et al.* (1998), due to positive autocorrelation over samples, an estimate of the standard error cannot simply use the sample variance $s_i^2 = \frac{1}{N} \sum_{j=1}^{N-1} (\theta_p^{(j)} - \bar{\theta}_p)^2$, as it would likely be an underestimate. Thus, a measure of uncertainty for $\bar{\theta}_p$ that takes into consideration autocorrelation over samples is a standard error based on the ESS $\sqrt{s_i^2 / N_{ESS}}$. Clearly, if there are still high autocorrelations for large k , N_{ESS} decreases, which means the standard error (and uncertainty about the estimate) increases. N_{ESS} close to N for all parameters of interest indicates efficiency in the MCMC scheme. In practice, N_{ESS} can only be estimated, as the autocorrelations $\hat{\rho}_{pk}$ are estimated from the sample (Vehtari *et al.*, 2020). Noise in the autocorrelation estimates $\hat{\rho}_{pk}$ increases as k increases, so the sum in (2.14) is calculated up to a value K . The value K is chosen based on a truncation rule (Vehtari *et al.*, 2020), for example, choosing K such that for $k > K$ the autocorrelations seem to be close to zero (Kass *et al.*, 1998; Congdon, 2019).

When running multiple chains, combined autocorrelations must be estimated to calculate the ESS. These combined autocorrelations take into consideration within-chain variance and between-chain variance (via the multi-chain variance estimate) (Gelman *et al.*, 2014; Vehtari *et al.*, 2020). Say we have \tilde{N} MCMC samples from \tilde{M} chains (after discarding the burn-in iterations) and we are interested in a parameter θ . (Note that we drop the p from θ_p hereinafter to simplify notation.) To calculate the multi-chain variance estimate,

Gelman *et al.* (2014) recommends splitting each chain in half so that there are M chains of size $N = \tilde{N}/2$ to test mixing *and* stationarity (convergence). At stationarity, both halves of a chain should be traversing the same distribution. The between-chain variance estimate is given by

$$B = \frac{N}{M-1} \sum_{m=1}^M (\bar{\theta}_{.m} - \bar{\theta}_{..})^2,$$

where $\bar{\theta}_{.m} = \frac{1}{N} \sum_{j=1}^N \theta_{jm}$, θ_{jm} is draw j of chain m , $\bar{\theta}_{..} = \frac{1}{M} \sum_{m=1}^M \bar{\theta}_{.m}$. The within-chain variance estimate is given by

$$W = \frac{1}{M} \sum_{m=1}^M s_m^2, \quad (2.15)$$

where $s_m^2 = \frac{1}{N-1} \sum_{j=1}^N (\theta_{jm} - \bar{\theta}_{.m})^2$. Thus, the multi-chain variance estimate is given by

$$\widehat{\text{var}}^+(\theta|\mathbf{y}) = \frac{W(N-1)}{N} + \frac{B}{N}. \quad (2.16)$$

According to Gelman *et al.* (2014), the combined autocorrelation at lag k is estimated by

$$\hat{\rho}_k = 1 - \frac{V_k}{2\widehat{\text{var}}^+(\theta|\mathbf{y})},$$

where $V_k = \frac{1}{M(N-k)} \sum_{m=1}^M \sum_{j=1}^N (\theta_{jm} - \theta_{j-k,m})^2$ is the *variogram* at lag k . The ESS for combined chains is

$$N_{ESS} = \frac{NM}{1 + 2 \sum_{k=1}^K \hat{\rho}_k},$$

where K is the first odd positive integer for which $\hat{\rho}_{K+1} + \hat{\rho}_{K+2}$ is negative.

Another useful convergence diagnostic to calculate, when running multiple chains, is

$$\hat{R} = \sqrt{\frac{\widehat{\text{var}}^+(\theta|\mathbf{y})}{W}},$$

where W and $\widehat{\text{var}}^+(\theta|\mathbf{y})$ were defined above in (2.15) and (2.16), respectively. This estimates the factor by which the current estimate of the marginal posterior variance of θ might be reduced if the simulations were continued in the limit $N \rightarrow \infty$. $\hat{R} \approx 1$ indicates that the chains are mixing well, where MCMC samples are deemed acceptable for some threshold, such as $\hat{R} < 1.05$ (Stan Development Team, 2020). Modified versions of N_{ESS} and \hat{R} can be obtained by using rank normalised values instead of the θ_{jm} ; see Vehtari *et al.* (2020). Note that they also define the combined autocorrelation as

$$\hat{\rho}_k = 1 - \frac{W - \frac{1}{M} \sum_{m=1}^M s_m^2 \hat{\rho}_{k,m}}{\widehat{\text{var}}^+(\theta|\mathbf{y})}.$$

Bayesian phylogenetics

Recall from Section 2.1.2 that the parameters in phylogenetic models are the tree τ with its branch lengths ℓ and the DNA substitution model parameters, which we denote by \mathcal{Q} . However, the main parameter of interest is the tree τ . Given some alignment data \mathbf{y} , the posterior probability of a tree is

$$\pi(\tau|\mathbf{y}) = \frac{\int_{\ell} \int_{\mathcal{Q}} \pi(\tau|\ell, \mathcal{Q}) \pi(\ell|\mathcal{Q}) \pi(\mathcal{Q}) p(\mathbf{y}|\tau, \mathcal{Q}, \ell) d\mathcal{Q} d\ell}{\sum_{i=1}^T \int_{\ell} \int_{\mathcal{Q}} \pi(\tau_i|\ell, \mathcal{Q}) \pi(\ell|\mathcal{Q}) \pi(\mathcal{Q}) p(\mathbf{y}|\tau_i, \mathcal{Q}, \ell) d\mathcal{Q} d\ell},$$

where T is the possible number of trees for the number of taxa in the alignment. Clearly, this is analytically intractable, as it involves integrating over all the numerical parameters for every possible topology and the number of rooted trees on n taxa is $\frac{(2n-3)!}{2^{n-2}(n-2)!}$. Thus, MCMC techniques are used to sample from the posterior distribution of the trees. In addition to checking for convergence in the numerical parameters of the model, checks for convergence in tree space must be performed too.

One such method is to compare the clade frequencies between chains initialised at different starting points. As defined in Section 2.1.1, a clade is a subset of taxa obtained by cutting a rooted tree on a branch and selecting only the leaves descending from this cut. Plots of the posterior probabilities of clades and cumulative relative frequencies in each chain (Heaps *et al.*, 2014) can also be employed to assess convergence. Figure 2.3 shows an example of where two chains have converged in tree space. In Figure 2.3a, a scatter plot of the posterior probabilities of clades for each chain suggests that the chains have converged as the points lie on a straight line. In Figure 2.3b, the cumulative clade frequencies appear to be approaching the same values, which also suggests convergence.

The main goal of phylogenetic inference is to identify the tree that best describes the evolutionary relationships amongst the species of interest. For numerical parameters, we can easily summarise the information from the MCMC samples by reporting a numerical summary, such as a posterior mean or mode. We can also quantify our uncertainty by looking at how concentrated the posterior density is or calculating posterior measures of dispersion. To summarise the posterior distribution over tree space, the most widely used summary is the majority-rule consensus tree. As a summary of a sample of trees, it includes only those splits (for unrooted trees) or clades (for rooted trees) which appear in over half of the samples (Bryant, 2003), here representing those with posterior probability greater than 0.5.

2.3.3 Hamiltonian Monte Carlo

For complicated models, the MCMC methods such as those described in Section 2.3.2 can suffer from inefficiency due to their random walk behaviour (Gelman *et al.*, 2014). This

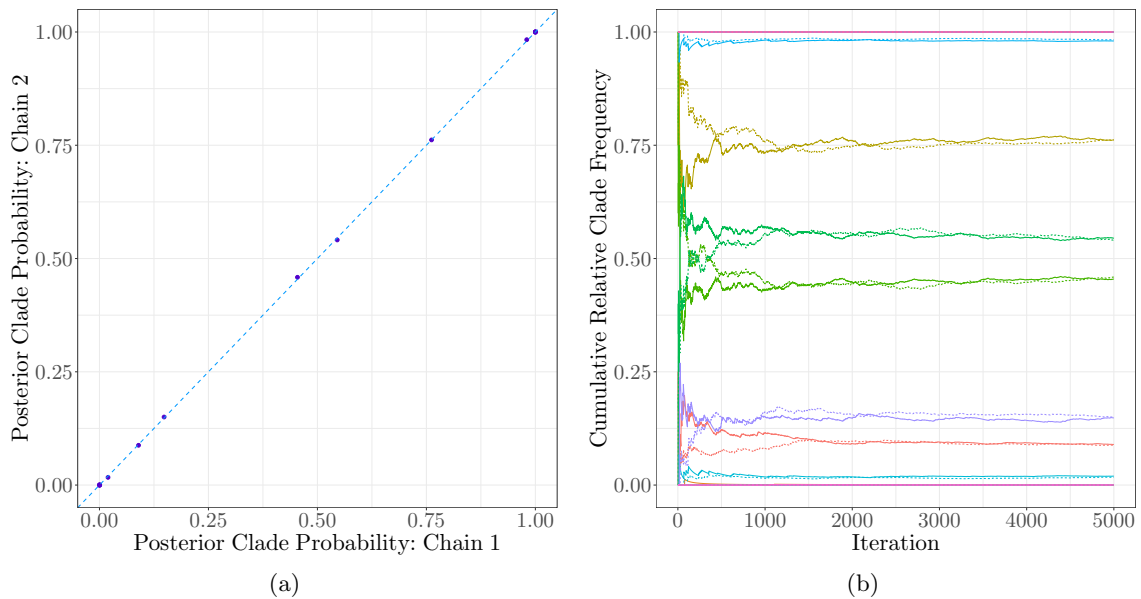


Figure 2.3: Diagnostic plots of two chains to assess convergence in tree space where (a) is the posterior probabilities of clades, (b) is the cumulative relative frequencies of clades and the different colours represent different clades, the solid lines represent one chain and the dashed lines represent the other chain.

means that it can take a long time for samplers to move around the target distribution $\pi(\boldsymbol{\theta}|\mathbf{y})$. The local random walk behaviour is particularly problematic for high-dimensional posterior distributions. Using a physical analogy of the energy of a particle moving in a frictionless space, *Hamiltonian Monte Carlo* (HMC) alleviates this issue of local random walk behaviour by generating proposals that are guided by the gradient of the posterior density. Thus, HMC algorithms can move more quickly through the posterior distribution. In this final section of Chapter 2 we briefly describe HMC, before beginning the phylogenetics part of the thesis.

Consider again a set of parameters $\boldsymbol{\theta} = (\theta_1, \dots, \theta_p)^T$ with p components in a target space and some data \mathbf{y} . HMC adds a momentum variable ϕ_j for each component θ_j . We can think of the $\boldsymbol{\theta}$ as the position of a fictitious particle in p -dimensional space and $\boldsymbol{\phi}$ as the corresponding momentum vector, where HMC algorithms simulate the trajectory of the particle. The parameters $\boldsymbol{\theta}$ and momentum variables $\boldsymbol{\phi}$ are updated together in a new Metropolis-Hastings algorithm, where proposals for $\boldsymbol{\theta}$ are largely determined by $\boldsymbol{\phi}$. The posterior density $p(\boldsymbol{\theta}|\mathbf{y})$ is augmented by the (independent) distribution $p(\boldsymbol{\theta})$ and the joint distribution $p(\boldsymbol{\theta}, \boldsymbol{\phi}|\mathbf{y})$ is simulated. However, we are only interested in $\boldsymbol{\theta}$ and $\boldsymbol{\phi}$ is an auxiliary variable that is introduced to allow the algorithm to move more quickly through the parameter space (Gelman *et al.*, 2014). The density for $\boldsymbol{\phi}$ does not depend on the parameters $\boldsymbol{\theta}$ (Stan Development Team, 2021). Generally, the momentum vector is given

a multivariate normal distribution with mean $\mathbf{0}$ and covariance matrix set to a “mass matrix” M . The term mass matrix here relates to the physical analogy of Hamiltonian dynamics, where both position and momentum evolve in continuous time.

In each iteration of an HMC algorithm, there are the following steps:

1. Update ϕ with a random draw from its posterior distribution, which is the same as its prior distribution, that is $\phi \sim N_p(\mathbf{0}, M)$.
2. Simultaneously update (θ, ϕ) using a discrete approximation to physical Hamiltonian dynamics. This involves L leapfrog steps which are scaled by a factor ϵ . One leapfrog step is as follows:

- (a) Make a “half-update” of ϕ by using the gradient of the log posterior density of θ , that is set

$$\phi = \tilde{\phi} + \frac{1}{2}\epsilon \frac{d \log p(\theta|\mathbf{y})}{d\theta},$$

where $\tilde{\phi}$ is the value of ϕ before the half-update.

- (b) Use ϕ to update θ , that is set

$$\theta = \tilde{\theta} + \epsilon M^{-1}\phi,$$

where $\tilde{\theta}$ is the value of θ before the update and ϵ is the step size and is used to tune the algorithm.

- (c) Half-update ϕ with the gradient of the log posterior density of θ , that is

$$\phi = \tilde{\phi} + \frac{1}{2}\epsilon \frac{d \log p(\theta|\mathbf{y})}{d\theta}.$$

With the exception of the first and last leapfrog step, updates (c) and (a) can be performed together.

3. Let $\theta^{(j-1)}$ and $\phi^{(j-1)}$ be the values of the parameter and momentum vectors, respectively, before the leapfrog process, and θ^* and ϕ^* be the values after the L steps. Set $\theta^{(j)} = \theta^*$ with probability

$$\frac{p(\theta^*|\mathbf{y}) p(\phi^*)}{p(\theta^{(j-1)}|\mathbf{y}) p(\phi^{(j-1)})}.$$

Otherwise, set $\theta^{(j)} = \theta^{(j-1)}$.

In addition to ϵ , HMC can be tuned via the mass matrix M , and the number of leapfrog steps L per iteration. The role of the mass matrix is to transform the target parameter space to make sampling more efficient (Betancourt, 2017; Stan Development Team, 2021).

Typically, M is chosen to be diagonal, which means that the components of ϕ are independent. It can be useful for M to roughly scale with the inverse covariance of the posterior distribution as it helps make HMC more efficient (Gelman *et al.*, 2014). If the inverse M^{-1} is a poor estimate of the posterior covariance, the step size ϵ must be kept small to maintain arithmetic precision, which means a larger number of leapfrog steps must be taken (Stan Development Team, 2021). The performance of HMC is highly sensitive to the tuning parameters ϵ and L (Hoffman & Gelman, 2014). If L is too large too much time is wasted on computation, whereas if L is too small successive samples will be close to each other, leading to undesirable random walk behaviour and slow mixing. Additionally, if ϵ is too large, the simulation will be inaccurate and yield low acceptance rates, whereas if ϵ is too small, many small steps are taken, leading to long simulation time. Tuning of these parameters requires both expertise and time (to make preliminary runs).

In the metagenomics part of this thesis, we make use of **Stan** (Stan Development Team, 2021), a computer program that automatically applies HMC, given a Bayesian model. More specifically, we use **RStan** (Stan Development Team, 2020), which is the R interface for **Stan**. **Stan** requires some data and model inputs but is responsible for computing the log posterior density and its gradients, setting the tuning parameters through a “warm-up” phase and implementing a *no-U-turn sampler* (NUTS) to move through the parameter space. In a NUTS, the number of leapfrog steps L is not fixed (nor tuned by the user) and instead is determined adaptively at each iteration. During the warm-up phase, **Stan** adaptively alters the other tuning parameters M and ϵ , where M^{-1} is set to a diagonal estimate of the covariance and ϵ is optimised to match a user defined acceptance rate target (Stan Development Team, 2021). The trajectory in each iteration continues until there is a negative value for the dot product between the momentum vector ϕ and the distance travelled from the position θ at the beginning of the iteration, that is the moment that the trajectory turns around.

In HMC, divergent transitions occur when the simulated, discretised Hamiltonian trajectory departs from the exact but analytically intractable trajectory. In practice, this is measured by the departure of the Hamiltonian value from its initial value. When the divergence is too high, the simulation becomes unreliable. If the posterior density is highly curved, the (gradient-based) simulation requires very small steps sizes to be accurate. Large step sizes (compared to the curvature) result in the simulation diverging from the true Hamiltonian. In some circumstances, a reparameterisation is needed so that the curvature of the posterior is less extreme. In other situations, it is sufficient to just increase the target acceptance rate, which will ensure the step size remains small enough to eliminate divergent transitions.

Part I

Phylogenetics

Chapter 3

Phylogenetic models facilitating root inference

The following three chapters focus on the phylogenetics part of the thesis. Additionally, the work for these three chapters appears in Hannaford *et al.* (2020). As explained in Section 2.2.2 in Chapter 2, the assumptions of stationarity and reversibility in a homogeneous model are often not justifiable from a biological perspective. Instead, they are imposed because they simplify the underpinning mathematics and computational inference. Worse still, the assumptions come at an inferential cost, giving rise to likelihood functions that are invariant to the position of the root of the tree. As such, they can only be used to infer unrooted trees, which depict the branching pattern of speciation events, without associating direction to the branches of the tree. Models which relax one or both of these assumptions can therefore offer more biological credibility, whilst also providing a likelihood function that is informative about the direction of time. In this chapter, we explore our own interpretation of some existing models of this type, before building on them to derive our own in Chapter 4.

3.1 Non-homogeneous models

Most models which allow root inference are *non-homogeneous*, which means that the process cannot be characterised by a single instantaneous rate matrix. Instead matrices from a countable set $\{Q_1, Q_2, \dots\}$ apply to different parts of the tree. In general, the Q_b all belong to the same family of rate matrices. For example, Kaehler (2017) considered a model in which the Q_b are all strand-symmetric, meaning the rate of substitution between a pair of nucleotides is the same as that between their Watson-Crick base pair complements (A pairs with T and C pairs with G). For this special class of models, a mathematical proof is provided which verifies that the root position can be identified from the likelihood.

However, inferential methodology to fit the model to data has not yet been developed. More often in the literature, the Q_b all belong to a family of *reversible* rate matrices, such as HKY85 or GTR (see Section 2.1.2), so that $Q_b = R_b \Pi_b$ for $b = 1, 2, \dots$. We refer to such a non-homogeneous process as *locally* reversible. The resulting models are generally *non-stationary*, with $\Pi_b \neq \Pi_{b'}$ for $b \neq b'$, and hence allow step changes in the theoretical stationary distribution, sometimes termed the *composition vector*, across the tree.

Some authors also allow variation in the exchangeability parameters (Dutheil & Bousau, 2008), although these are often constant, with $R_b = R$ for all b . For example, Yang & Roberts (1995) investigated two models based on the HKY85 model, N1 and N2, in which the exchangeability parameters are constant over the tree, but the composition vectors in the Π_b vary from branch to branch. The N2 model allows for a different composition vector on every branch of the tree. For a tree with n taxa, the N1 model assumes $n + 2$ composition vectors: one per branch leading to each terminal node (taxon), one for the internal branches and one for the root. In their application, their models were able to identify a widely accepted root for taxa with different base frequencies, although it must be noted that the model was applied to an alignment of only four taxa. Heaps *et al.* (2014) investigated a similar model in a Bayesian framework, with a prior that assumes positive correlation among the set of composition vectors, thereby allowing information to be shared between branches. More details on their prior and its relevance to our work are given in Section 4.2 in the next chapter. In both of their applications, where there was substantial evidence of compositional heterogeneity in the data, their model provided credible topological inferences. Compositional heterogeneity means that the nucleotide compositions (the proportions of **A,G,C** and **T**) change across the tree. Due to their assumption of the same rate matrix (and thus stationary distribution) across the tree, standard homogeneous models cannot account for this.

Other approaches intended to reduce the variance of parameter estimates in complex models of this form have been largely based on the idea of dimension reduction. For instance, Foster (2004) considered a mixture model in which the B branches of the tree are allocated to one of $K \ll B$ mixture components, with branches in the same component sharing a composition vector. In an application to a problematic data set of five bacterial 16S genes (Embley *et al.*, 1993), a model with $K = 2$ mixture components was used and was able to infer the most biologically plausible tree. Similarly, Blanquart & Lartillot (2006) introduced a model in which the step changes in the theoretical stationary distribution occur according to a Poisson process, independently of the speciation events which determine the tree's branching structure. Their model allows for "break points" to be added to or removed from the topology. At each break point, "compositional shift" events occur, where a new composition vector π' is proposed independently of the composition vector before the break point π , which means that the successive compositional shifts

essentially follow a 0-th order Markov process. Their model performed well when applied to several data sets. However, one criticism highlighted by the authors was that the use of the 0-th order Markov process was not realistic but mathematically convenient. They suggested that using a first order Markov process would be an improvement but would increase computational time.

The main difficulty with these mixture-type models is that the dimension of the problem – determined by the number of mixture components in the former case and the number of break points in the latter – are not known *a priori*, which substantially complicates computational inference. When dimension-changing moves are included in MCMC algorithms, issues with convergence and mixing can commonly occur. Foster (2004) acknowledged that the adding or removing of parameters could be incorporated into the MCMC, but they instead fitted models with different values of K mixture components and compared the fit of the models to decide the optimum number of mixture components to use. The work of Heaps *et al.* (2014) does not suffer from the aforementioned MCMC-related problems, as the dimension of their model is fixed. Foster (2004) and Blanquart & Lartillot (2006) do not advocate using “branchwise” models and argue that overparameterisation can lead to poor fitting models or that models with many more parameters do not substantially improve model fit. However, these issues of overparameterisation are avoided in the model of Heaps *et al.* (2014) because the prior is structured to allow information to be shared between branches (Williams *et al.*, 2015).

3.2 Non-reversible models

Motivated by the rooting problem, Huelsenbeck *et al.* (2002) investigated stationary but non-reversible substitution models in a Bayesian framework. Their model is based on an instantaneous rate matrix which is structurally unconstrained, representing the so-called *general Markov model* of DNA evolution (Barry & Hartigan, 1987). Since a single rate matrix is applied to the whole tree, this model is more parsimonious than the models discussed in the previous section. They assessed the model using a non-informative uniform prior for the off-diagonal elements of the rate matrix, with several simulated data sets of eight taxa, with varying degrees of non-reversibility. Their model was able to identify the correct root position with a fixed unrooted topology in the data sets when the substitution process used to simulate the data was highly non-reversible.

Cherlin *et al.* (2017) built on the work of Huelsenbeck *et al.* (2002). Firstly, instead of fixing the unrooted topology, Cherlin *et al.* (2017) allowed the rooted tree to be unknown, assigned it a prior and fitted the model using an MCMC algorithm that included topological moves. They also considered larger data sets, including a data set with 36 taxa. Finally, instead of using a uniform prior for the off-diagonal elements q_{ij} of the rate matrix,

they developed two hierarchical priors that are centred on a standard reversible matrix but allow non-reversible perturbations of the elements. Each prior treats each q_{ij} as a log-normal perturbation of the corresponding element of the unknown rate matrix of an HKY85 model (shown in (2.2)). Their two models, the NR and NR2 models, differ in their number of perturbation components, with the NR model having one perturbation component and the NR2 model utilising two components. Letting q_{ij}^H denote an off-diagonal element of an HKY85 rate matrix, the NR model for $i \neq j$ is

$$\log q_{ij} = \log q_{ij}^H + \epsilon_{ij},$$

where the ϵ_{ij} are independent $N(0, \sigma^2)$ quantities. The extent to which Q departs from a HKY85 structure is represented by the perturbation standard deviation σ , with a larger value implying a greater degree of departure. Under the NR model, for any given data set, if σ is found to be large *a posteriori*, it is not necessarily evidence of non-reversibility. This is because departures from an HKY85 model could lead to a non-reversible model or a GTR model and these two deviations are confounded. Thus, the NR2 model was developed, which uses a two-stage process to perturb the underlying HKY85 rate matrix. The first perturbation is within the space of GTR matrices and leads to a reversible rate matrix Q^R , which is then perturbed within the space of general rate matrices, leading to a general non-reversible rate matrix. Note that the first perturbation results in a GTR matrix with probability zero, since the parameters of the perturbed matrix are modelled as continuous quantities. Its prior simply has more mass around GTR structures. With this model, the two deviations from an HKY85 model can be decoupled, with the extent of each deviation represented by the variance parameters of the perturbations. Similar to Huelsenbeck *et al.* (2002), their models were able to provide sensible root inferences on data simulated with a large degree of non-reversibility. Analysis of empirical data also recovered roots with wide biological consensus.

Simulation experiments and application to biological data sets have suggested that the stationary, but non-reversible models described in this section can produce sensible root inferences when the model assumptions are clearly supported. However, root inference was found to be very sensitive to model misspecification, especially violation of the assumption of stationarity (Williams *et al.*, 2015). This limits the utility of such models in application to data sets of biological interest, where it is common to see variation in sequence composition across taxa due to lineage-specific compositional change.

Lie Markov models for DNA evolution have the property of closure under matrix multiplication (Sumner *et al.*, 2012a). Let P_1 and P_2 be transition matrices obtained by taking the matrix exponential of two rate matrices from a family of Markov models \mathcal{M} . \mathcal{M} is multiplicatively closed if and only if for all such P_1, P_2 , the product $P_1 P_2$ is obtainable as the matrix exponential of another rate matrix from the family \mathcal{M} . Woodhams *et al.* (2015)

defined a hierarchy of Lie Markov models capable of distinguishing pairs of nucleotides, such as purines and pyrimidines. The general Markov model, with 12 degrees of freedom, represents the most complex family of Lie Markov models. However, all others can be represented by 10 degrees of freedom or fewer. Some are non-reversible and, like families of reversible rate matrices, have a biologically interpretable structure. It is therefore possible to combine the ideas from this section and Section 3.1 and build parsimonious models that are non-stationary, non-homogeneous and (locally) non-reversible by using an appropriate set of rate matrices from a family of Lie Markov models. This is the main focus of Chapter 4.

For now, we discuss two families of non-reversible Lie Markov models which, in the terminology of Woodhams *et al.* (2015), are referred to as the RY5.6b and RY8.8 Lie Markov families. In each case, we derive a new parameterisation of the underpinning rate matrix Q , and the relationship between the new parameters and the theoretical stationary probabilities $\boldsymbol{\pi}$. These two non-reversible models are the starting points for our own models that are non-stationary, non-homogeneous and (locally) non-reversible, which we will discuss in Chapter 4.

3.2.1 The RY5.6b Model

Motivated by its simplicity and similarity to the widely used HKY85 model defined in (2.2) in Section 2.1.2, the first Lie Markov model we consider is the (non-reversible) RY5.6b model. Following the formulation presented in Woodhams *et al.* (2015), its rate matrix Q can be represented as

$$Q = \begin{pmatrix} * & \alpha + \rho_2 & \beta + \rho_3 & \beta + \rho_4 \\ \alpha + \rho_1 & * & \beta + \rho_3 & \beta + \rho_4 \\ \beta + \rho_1 & \beta + \rho_2 & * & \alpha + \rho_4 \\ \beta + \rho_1 & \beta + \rho_2 & \alpha + \rho_3 & * \end{pmatrix}$$

where $\alpha, \beta, \rho_1, \rho_2, \rho_3, \rho_4 \geq 0$ and the *s ensure that the rows sum to zero. As indicated by the prefix of its name, the model has the symmetry condition of purine-pyrimidine (RY) pairing, with rates of change for transversions sharing a parameter and rates of change for transitions sharing a different parameter. However, the six parameters are plainly not identifiable since we can replace α and β with $\alpha + \delta$ and $\beta + \delta$, and ρ_i with $\rho_i - \delta$ for $i = 1, 2, 3, 4$, and obtain exactly the same rate matrix. In the reversible HKY85 case, the off-diagonal elements in each column of the rate matrix share a stationary probability π_i with $\boldsymbol{\pi} \in \mathcal{S}_4$. In the RY5.6b model, they each share a parameter ρ_i . By choosing the analogous constraint, $\boldsymbol{\rho} = (\rho_1, \rho_2, \rho_3, \rho_4)^T \in \mathcal{S}_4$, we can eliminate the parameter redundancy. We note that the 5 and 6 in the name of the RY5.6b model arise from it

being a five-dimensional model whose rate matrices form a polyhedral cone with six rays (Fernández-Sánchez *et al.*, 2015). These allow it to be expressed through six non-negative parameters.

Although the simplex constraint removes the additive identifiability issue, the overall scale of the rate matrix is still arbitrary since it appears only in its normalised form, $Q' = Q/(-\sum_u q_{uu}\pi_u)$, in the transition matrix. To resolve this problem, it is convenient to fix the scale of the rate matrix by constraining its trace to be equal to -7 as this limits the support of the remaining parameters so that $(\alpha, 2\beta) \in \mathcal{S}_2$ or, equivalently, $\alpha \in [0, 1]$ and then $\beta = (1 - \alpha)/2$. This, in turn, simplifies the process of specifying a prior. See Appendix A for the proof that this additional constraint on the trace of Q is required.

The stationary probabilities π associated with this rate matrix are given by

$$\begin{aligned}\pi_1 &= \frac{-\alpha^2 + (5 - \alpha)\rho_1 + (3\alpha - 1)\rho_2 - \alpha + 2}{2(3 - 2\alpha)(\alpha + 2)}, \\ \pi_2 &= \frac{-\alpha^2 + (5 - \alpha)\rho_2 + (3\alpha - 1)\rho_1 - \alpha + 2}{2(3 - 2\alpha)(\alpha + 2)}, \\ \pi_3 &= \frac{-\alpha^2 + (5 - \alpha)\rho_3 + (3\alpha - 1)\rho_4 - \alpha + 2}{2(3 - 2\alpha)(\alpha + 2)}, \\ \pi_4 &= \frac{-\alpha^2 + (5 - \alpha)\rho_4 + (3\alpha - 1)\rho_3 - \alpha + 2}{2(3 - 2\alpha)(\alpha + 2)}.\end{aligned}\tag{3.1}$$

This is found using formulae provided in the supplementary material of Woodhams *et al.* (2015) for the stationary probabilities in terms of their parameterisation. First, we note that all rate matrices for the Lie Markov models can be expressed using a linear combination (or subset) of 12 basis matrices, which are listed in Table 1 of Woodhams *et al.* (2015). Thus, the RY5.6b rate matrix is also written as

$$Q = \begin{pmatrix} * & a + a_1 + d - e_1 & a - d + e_2 & a - d - e_2 \\ a + a_1 + d + e_1 & * & a - d + e_2 & a - d - e_2 \\ a + d + e_1 & a + d - e_1 & * & a + a_1 - d - e_2 \\ a + d + e_1 & a + d - e_1 & a + a_1 - d + e_2 & * \end{pmatrix}\tag{3.2}$$

under their canonical parameterisation which expresses Q as a linear combination of basis matrices A, A_1, D, E_1, E_2 . They work with the columns summing to zero rather than the rows, so for the RY5.6b model we have $Q^T = aA + a_1A_1 + dD + e_1E_1 + e_2E_2$, which we transpose to get Q in (3.2). Then we let

$$\begin{aligned}\rho_1 &= d + e_1 + z, & \rho_3 &= -d + e_2 + z, \\ \rho_2 &= d - e_1 + z, & \rho_4 &= -d - e_2 + z, \\ \alpha &= a + a_1 - z, & \beta &= a - z.\end{aligned}$$

Here $\sum_{i=1}^4 \rho_i = 4z$, but we chose the constraint that $\sum_{i=1}^4 \rho_i = 1$ so we have

$$4z = 1 \Leftrightarrow z = \frac{1}{4}.$$

Now $\alpha + 2\beta = a + a_1 - z + 2a - 2z = 3a + a_1 - 3z$ and we chose the constraint $\alpha + 2\beta = 1$, so we have

$$\begin{aligned} 3a + a_1 - 3z &= 1 \\ 3a + a_1 - \frac{3}{4} &= 1 \\ 3a + a_1 &= \frac{7}{4} \\ \Leftrightarrow a_1 &= \frac{7}{4} - 3a. \end{aligned}$$

So we have

$$\rho_1 = d + e_1 + \frac{1}{4}, \quad \rho_2 = d - e_1 + \frac{1}{4}, \quad \rho_3 = -d + e_2 + \frac{1}{4}$$

and

$$\begin{aligned} \alpha &= a + \left(a_1 - \frac{1}{4}\right) \\ &= a + \frac{7}{4} - 3a - \frac{1}{4} \\ &= -2a + \frac{3}{2}. \end{aligned}$$

Solving for d, e_1, e_2 and a gives

$$\begin{aligned} a &= -\frac{\alpha}{2} + \frac{3}{4}, \\ d &= \frac{1}{2}(\rho_1 + \rho_2) - \frac{1}{4}, \\ e_1 &= \frac{1}{2}(\rho_1 - \rho_2), \\ e_2 &= \frac{1}{2}(2\rho_3 + \rho_1 + \rho_2 - 1). \end{aligned}$$

These values need to be substituted into

$$\begin{aligned} p &= 4a^3 + 4a^2a_1 + aa_1^2, & r &= 4a^2e_1 + 2aa_1e_1, \\ q &= 4a^2d + 4aa_1d + a_1^2d, & s &= 4a^2e_2 + 2aa_1e_2 \end{aligned}$$

to find

$$\begin{aligned} \pi_1 &= \frac{1}{4} + \frac{q+r}{4p}, & \pi_3 &= \frac{1}{4} + \frac{-q+s}{4p}, \\ \pi_2 &= \frac{1}{4} + \frac{q-r}{4p}, & \pi_4 &= \frac{1}{4} + \frac{-q-s}{4p}. \end{aligned} \tag{3.3}$$

Substituting our values into the above formulae gives us the stationary probabilities given in (3.1). For example, the first stationary probability is

$$\begin{aligned}
 \pi_1 &= \frac{1}{4} + \frac{4a^2d + 4aa_1d + a_1^2d + 4a^2e_1 + 2aa_1e_1}{16a^3 + 16a^2a_1 + 4aa_1^2} \\
 &= \frac{1}{4} + \frac{(2a + a_1)^2 d + (2a + a_1) 2ae_1}{4a(2a + a_1)^2} \\
 &= \frac{1}{4} + \frac{1}{4(2a + a_1)} \left(2d + \frac{da_1}{a} + 2e_1 \right) \\
 &= \frac{1}{4} + \frac{1}{7 - 4a} \left(-d + \frac{7d}{4a} + 2e_1 \right) \quad (\text{substituting in } a_1) \\
 &= \frac{1}{4} + \frac{1}{(4a - 7)4a} (-7d + 4ad - 8ae_1) \\
 &= \frac{(4a - 7)a + (-7d + 4ad - 8ae_1)}{4a(4a - 7)4a} \\
 &= \frac{4a^2 - 7a - 7d + 4ad - 8ae_1}{(4a - 7)4a},
 \end{aligned}$$

which we can substitute the values of a, d and e_1 into to obtain

$$\begin{aligned}
 \pi_1 &= \frac{4\left(\frac{3}{4} - \frac{\alpha}{2}\right)^2 + \frac{7\alpha}{2} - \frac{21}{4} - 7\frac{(\rho_1 + \rho_2)}{2} + \frac{7}{4} + 4\left(\frac{3}{4} - \frac{\alpha}{2}\right)\left(\frac{(\rho_1 + \rho_2)}{2} - \frac{1}{4}\right) - 8\left(\frac{3}{4} - \frac{\alpha}{2}\right)\frac{(\rho_1 - \rho_2)}{2}}{4\left(\frac{3}{4} - \frac{\alpha}{2}\right)(-2\alpha - 4)} \\
 &= \frac{\left(\frac{3}{2} - \alpha\right)^2 + \frac{7(\alpha - \rho_1 - \rho_2)}{2} - \frac{14}{4} + (3 - 2\alpha)\left(\frac{(\rho_1 + \rho_2)}{2} - \frac{1}{4}\right) - (6 - 4\alpha)\frac{(\rho_1 - \rho_2)}{2}}{-2(3 - 2\alpha)(\alpha + 2)} \\
 &= \frac{\alpha^2 - 3\alpha + \frac{7(\alpha - \rho_1 - \rho_2)}{2} - \frac{5}{4} - \alpha(\rho_1 + \rho_2) + \frac{\alpha}{2} + \frac{3(\rho_1 + \rho_2)}{2} - \frac{3}{4} + (2\alpha - 3)(\rho_1 - \rho_2)}{-2(3 - 2\alpha)(\alpha + 2)} \\
 &= \frac{-\alpha^2 - \alpha + 2 + 5\rho_1 - \rho_2 - \alpha\rho_1 + 3\alpha\rho_2}{2(3 - 2\alpha)(\alpha + 2)},
 \end{aligned}$$

which is equivalent to π_1 in (3.1). The remaining stationary probabilities π_2, π_3 and π_4 are found in the same way.

For ease of interpretation, it might seem more natural to reparameterise the model directly in terms of $\alpha \in [0, 1]$ and the stationary distribution $\boldsymbol{\pi} \in \mathcal{S}_4$. However, given a fixed value for α , the mapping $\pi_\alpha : \mathcal{S}_4 \rightarrow \mathcal{S}_4$, where $\pi_\alpha(\boldsymbol{\rho}) = \boldsymbol{\pi}$, is not surjective. For example, if $\alpha = 0.5$ and $\pi_1 = 1$, then we require

$$8.75 = 4.5\rho_1 + 0.5\rho_2,$$

which is not possible since $\boldsymbol{\rho} \in \mathcal{S}_4$. A value of $\pi_1 = 1$ is not possible for any fixed value of α . Therefore, for any fixed value of α , not all points on the simplex can be represented as the stationary distribution under the RY5.6b model. Parameterisation using α and $\boldsymbol{\pi}$

would substantially complicate inference and so we retain the original parameterisation, in terms of α and ρ .

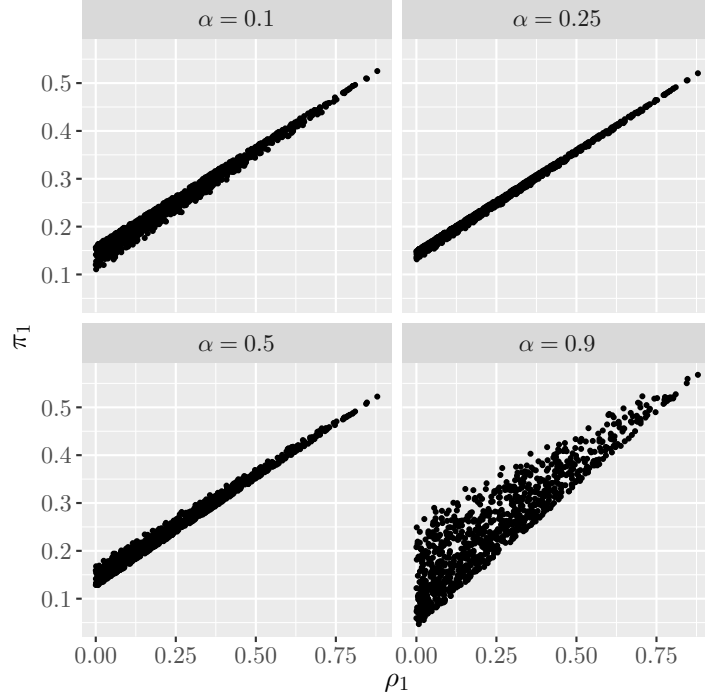


Figure 3.1: Plots of the stationary probability π_1 of the RY5.6b model against ρ_1 for values of ρ sampled from the uniform distribution over \mathcal{S}_4 and various values of α . By symmetry, plots for the other pairs (ρ_i, π_i) , $i = 2, 3, 4$, display identical patterns (not shown).

Seeking an interpretation of ρ , the relationship between each π_i and the corresponding ρ_i is complicated by the simplex constraints, which preclude isolation of the effect of a change in ρ_i on π_i , whilst all the other ρ_j remain fixed. However, from (3.1), because $5 - \alpha > 3\alpha - 1$, it is clear that for any fixed $\alpha \in [0, 1]$, there is a positive linear relationship between, say, ρ_1 and π_1 . The slope and intercept depend on how a simplex-preserving decrease in $\rho_2 + \rho_3 + \rho_4$ is shared between ρ_2 and $\rho_3 + \rho_4$ when ρ_1 is increased. To illustrate the relationships numerically, we simulate a sample of ρ vectors from a uniform distribution over \mathcal{S}_4 and then compute the corresponding stationary distribution π for various values of α . Plots of π_i against ρ_i are displayed in Figure 3.1, where they show a strong positive relationship. We therefore interpret the parameter vector ρ as playing a role similar to the stationary distribution π in the HKY85 model. The parameter α then allows for differences between the rates of transition and transversion.

3.2.2 The RY8.8 Model

The structure of the RY5.6b model is biologically appealing because of its simplicity and parallels with the widely used HKY85 model. However, it suffers a number of drawbacks. First, the model only has five degrees of freedom, which makes it inflexible compared with more complex Lie Markov models. Second, the additive structure of the instantaneous rates of change in the RY5.6b, as well as various other Lie Markov models, can cause problems in the analysis of biological data. In many alignments, the empirical proportions of A, G, C and T are all reasonably close to 0.25 (Bohlin *et al.*, 2017). If we imagine that the data arose from a stationary CTMP, this would demand $\pi_i \simeq 1/4$ for all $i = 1, \dots, 4$. With reference to the RY5.6b model, with stationary distribution (3.1), arguments of symmetry imply that for any $\alpha \in [0, 1]$, we can only achieve $\pi_i = 1/4$ for all i if $\rho_i = 1/4$ for all i . In this case, the ratio of the rates of change for transitions and transversions is given by

$$(\alpha + 1/4)/\{(1 - \alpha)/2 + 1/4\} = (4\alpha + 1)/(3 - 2\alpha) \leq 5.$$

However, experience suggests that for some alignments, we would expect a value much larger than this (Rosenberg *et al.*, 2003). In mammalian genomes, for example, this can occur due to *5-methylcytosine deamination* to thymine at some sites, causing high rates of C to T point mutation (Hodgkinson & Eyre-Walker, 2011). (5-methylcytosine is a methylated form of cytosine which has several biological roles including regulation of gene transcription. Deamination is the removal of an amino acid from a compound or amino acid group and when this happens to 5-methylcytosine it becomes thymine.) This provides a possible explanation for the conclusions drawn in Woodhams *et al.* (2015), based on analyses of seven data sets, that the fit of the RY5.6b model is notably worse than that of the structurally similar HKY85 model. We therefore investigate a second (non-reversible) Lie Markov model, the RY8.8 model, which is more highly parameterised than RY5.6b and free from its additive structure. It has also been found to fit well in analyses of biological data (Woodhams *et al.*, 2015).

As its RY-prefix suggests, the RY8.8 rate matrix is based on the symmetry condition of purine-pyrimidine pairing and has eight degrees of freedom which can be represented by eight non-negative parameters. A representation of its rate matrix is given by

$$Q = \begin{pmatrix} * & \rho_2 & \tilde{\rho}_7 & \tilde{\rho}_8 \\ \rho_1 & * & \tilde{\rho}_7 & \tilde{\rho}_8 \\ \tilde{\rho}_5 & \tilde{\rho}_6 & * & \rho_4 \\ \tilde{\rho}_5 & \tilde{\rho}_6 & \rho_3 & * \end{pmatrix} \quad (3.4)$$

where $\rho_1, \rho_2, \rho_3, \rho_4, \tilde{\rho}_5, \tilde{\rho}_6, \tilde{\rho}_7, \tilde{\rho}_8 \geq 0$ and the values of each * ensure that the rows sum to zero and $\rho_1, \rho_2, \rho_3, \rho_4, \tilde{\rho}_5, \tilde{\rho}_6, \tilde{\rho}_7, \tilde{\rho}_8 \geq 0$. In order to fix the scale of the rate matrix, it is

convenient to fix the trace as -1, then we can take $\rho_i = 2\tilde{\rho}_i$ for $i = 5, \dots, 8$, and restrict $\boldsymbol{\rho} \in \mathcal{S}_8$.

The analytic forms for the stationary probabilities $\boldsymbol{\pi}$ are given by

$$\begin{aligned} \pi_1 &= \frac{(2\rho_1 + \rho_7 + \rho_8)\rho_5 + 2\rho_1\rho_6}{k_1k_2}, & \pi_2 &= \frac{(2\rho_2 + \rho_7 + \rho_8)\rho_6 + 2\rho_2\rho_5}{k_1k_2}, \\ \pi_3 &= \frac{(2\rho_3 + \rho_5 + \rho_6)\rho_7 + 2\rho_3\rho_8}{k_2k_3}, & \pi_4 &= \frac{(2\rho_4 + \rho_5 + \rho_6)\rho_8 + 2\rho_4\rho_7}{k_2k_3}, \end{aligned} \quad (3.5)$$

where

$$k_1 = \rho_8 + 2\rho_2 + \rho_7 + 2\rho_1, \quad k_2 = \rho_5 + \rho_6 + \rho_7 + \rho_8, \quad k_3 = 2\rho_3 + 2\rho_4 + \rho_5 + \rho_6.$$

As with the stationary probabilities for the RY5.6b model, these are found using the formulae in the supplementary material of Woodhams *et al.* (2015). To use these formulae we must write the original parameters in terms of our parameters. First, we note that the original rate matrix given by the linear combination of basis matrices A, A₁, D, D₁, E₁, E₂, F₁ and F₂ (see Table 1 in Woodhams *et al.* (2015)) is

$$Q = \begin{pmatrix} * & \tilde{a} + d + d_1 - e_1 - f_1 & a - d + e_2 + f_2 & a - d - e_2 - f_2 \\ \tilde{a} + d + d_1 + e_1 + f_1 & * & a - d + e_2 + f_2 & a - d - e_2 - f_2 \\ a + d + e_1 - f_1 & a + d - e_1 + f_1 & * & \tilde{a} - d - d_1 - e_2 + f_2 \\ a + d + e_1 - f_1 & a + d - e_1 + f_1 & \tilde{a} - d - d_1 + e_2 - f_2 & * \end{pmatrix}, \quad (3.6)$$

where $\tilde{a} = a + a_1$ and the *s ensure that the rows sum to zero. Setting (3.4) equal to (3.6) and using the fact that $\rho_i = 2\tilde{\rho}_i$ for $i = 5, \dots, 8$, we get

$$\begin{aligned} \rho_1 &= a + a_1 + d + d_1 + e_1 + f_1, & \rho_5 &= 2(a + d + e_1 - f_1), \\ \rho_2 &= a + a_1 + d + d_1 - e_1 - f_1, & \rho_6 &= 2(a + d - e_1 + f_1), \\ \rho_3 &= a + a_1 - d - d_1 + e_2 - f_2, & \rho_7 &= 2(a - d + e_2 + f_2), \\ \rho_4 &= a + a_1 - d - d_1 - e_2 + f_2, & \rho_8 &= 2(a - d - e_2 - f_2). \end{aligned}$$

Due to the simplex constraint on $\boldsymbol{\rho}$, we get

$$a_1 = \frac{1}{4} - 3a \quad (3.7)$$

and calculating $\sum_{i=5}^8 \rho_i$ gives us

$$\begin{aligned} 8a &= \rho_5 + \rho_6 + \rho_7 + \rho_8 \\ \Rightarrow a &= \frac{1}{8}(\rho_5 + \rho_6 + \rho_7 + \rho_8). \end{aligned} \quad (3.8)$$

It can be shown that

$$e_1 = \frac{1}{8} (2\rho_1 - 2\rho_2 + \rho_5 - \rho_6), \quad (3.9)$$

$$e_2 = \frac{1}{8} (2\rho_3 - 2\rho_4 + \rho_7 - \rho_8). \quad (3.10)$$

Substituting (3.9) into $\rho_1 - \rho_2$ and (3.10) into $\rho_4 - \rho_3$ gives us

$$f_1 = \frac{1}{8} (2\rho_1 - 2\rho_2 - \rho_5 + \rho_6) \quad (3.11)$$

$$f_2 = \frac{1}{8} (2\rho_4 - 2\rho_3 + \rho_7 - \rho_8).$$

To find d we substitute (3.8), (3.9) and (3.11) into ρ_5 to get

$$d = \frac{1}{8} (\rho_5 + \rho_6 - \rho_7 - \rho_8).$$

Finally, to find d_1 , we substitute our values for a, a_1, d, e_1 and f_1 into ρ_1 to get

$$d_1 = \frac{1}{8} (2(\rho_1 + \rho_2 - \rho_3 - \rho_4) - \rho_5 - \rho_6 + \rho_7 + \rho_8).$$

For the RY8.8 model we use

$$\begin{aligned} p &= 4a^3 + 4a^2a_1 + aa_1^2 - ad_1^2 \\ q &= 4a^2d + 4aa_1d + a_1^2d - dd_1^2 \\ r &= 4a^2e_1 + 2aa_1e_1 - 2ad_1e_1 + 4adf_1 + 2a_1df_1 - 2dd_1f_1 \\ s &= 4a^2e_2 + 2aa_1e_2 + 2ad_1e_2 + 4adf_2 + 2a_1df_2 + 2dd_1f_2 \end{aligned}$$

in (3.3). Using the `simplify()` function in the software `Maple` (Maplesoft, 2020) and a few algebraic simplifications we obtain the stationary probabilities given in (3.5). For example, the formula for the first stationary probability yields

$$\begin{aligned} \pi_1 &= \frac{-\rho_5^2 + (3\rho_1 - \rho_2 - \rho_4 - 2\rho_6 + \rho_7 + \rho_8 + 1)\rho_5}{2(\rho_1 + \rho_2 - \rho_3 - \rho_4 - \rho_5 - \rho_6 + 1)(\rho_5 + \rho_6 + \rho_7 + \rho_8)} \\ &\quad + \frac{(3\rho_1 - (\rho_2 + \rho_3 + \rho_4 + \rho_6 + \rho_7 + \rho_8 + 1))\rho_6}{2(\rho_1 + \rho_2 - \rho_3 - \rho_4 - \rho_5 - \rho_6 + 1)(\rho_5 + \rho_6 + \rho_7 + \rho_8)}, \end{aligned}$$

which can be simplified to the form shown in (3.5) by exploiting the constraint $\boldsymbol{\rho} \in \mathcal{S}_8$.

Similar to the RY5.6b model, direct parameterisation in terms of $\boldsymbol{\pi}$ and, say, the rates of transition $(\rho_1, \rho_2, \rho_3, \rho_4)$, or the (scaled) rates of transversion $(\rho_5, \rho_6, \rho_7, \rho_8)$, would complicate inference because for fixed $(\rho_{i+1}, \rho_{i+2}, \rho_{i+3}, \rho_{i+4})$ where $i = 0$ or $i = 4$, we cannot invert the mapping from the remaining elements in $\boldsymbol{\rho} \in \mathcal{S}_8$ to $\boldsymbol{\pi} \in \mathcal{S}_4$. However,

the parameters in the RY8.8 model have clear interpretations as instantaneous rates of change between different pairs of nucleotides and so we parameterise the model in terms of the single interpretable stochastic vector ρ .

3.3 Assumption across sites

The previous sections have described substitution models for evolution at a single genomic site. In order to extend this to a joint model for the whole alignment, sites are generally assumed to evolve independently, but with their own rates γ_i , $i = 1, \dots, m$, which scale the normalised rate matrix Q' linearly. Biologically, this reflects the idea that rates of evolution vary according to functional or structural pressures acting at a site: important sites are subject to higher selective constraints and hence evolve more slowly (Yang, 1996). These site-specific parameters γ_i are modelled as multiplicative random effects, where γ_i is often given a gamma distribution (Yang, 1993). This is $\gamma_i|\phi \sim \text{Ga}(\phi, \phi)$ for $i = 1, \dots, m$, where the common shape and rate ϕ give the distribution a unit mean. The value of ϕ controls the manner and extent to which evolutionary rates vary across sites. For example, $\phi > 1$ results in a bell-shaped distribution, which means that most sites have intermediate rates, while few sites have very high or low rates. Meanwhile, $\phi \leq 1$ gives an L-shaped distribution, implying that most sites have very low substitution rates (almost “invariable”) and a few sites have very high rates (Yang, 1996).

During model-fitting, discretising the continuous gamma distribution allows intermediate likelihood calculations to be cached, which substantially speeds up computation. Therefore, in keeping with standard practice in the phylogenetic literature, we adopt a discrete approximation to the gamma distribution with four rate classes (Yang, 1994). Under this distribution, the rate γ_i at site i is equal to $r_k(\phi)$, $k = 1, 2, 3, 4$, with probability $p_k = 1/4$ and $r_k(\phi)$ taken as the $(k - 0.5)/4$ quantile in the $\text{Ga}(\phi, \phi)$ distribution.

Chapter 4

Non-homogeneous Lie Markov models

In this chapter, we derive a pair of non-homogeneous, non-stationary, non-reversible phylogenetic models. These models extend the RY5.6b and RY8.8 models described in the previous chapter by incorporating compositional heterogeneity. We discuss the prior specification for each model, before finally describing posterior inference.

4.1 Non-homogeneous RY5.6b and RY8.8 models

As explained in Section 2.1.1, there are often both theoretical and empirical arguments for building non-stationarity into models for substitutions in molecular sequences. We therefore propose non-homogeneous and non-stationary extensions of the (non-reversible) RY5.6b and RY8.8 models outlined in the previous chapter. A bifurcating rooted tree on n taxa has $B = 2n - 2$ branches and its underpinning unrooted topology has one fewer. We construct a non-homogeneous RY5.6b model by allowing evolution along every branch b of the associated unrooted topology to be controlled by its own rate matrix Q_b which belongs to the RY5.6b family. Rooting this tree on branch r of the unrooted topology divides the branch into two. The rate matrix Q_r is associated with the two new branches on either side of the root, whilst its stationary distribution is used as the distribution at the root. We define our non-homogeneous RY8.8 model in an analogous fashion. It is worth mentioning that an alternative, though less parsimonious, way to formulate the models would be to allow the branches on either side of the root to have their own rate matrix with a simplex-valued parameter describing the distribution at the root of the tree.

Computational inference is greatly simplified if the number of parameters which vary from branch to branch is kept small. For example, when Heaps *et al.* (2014) and Williams *et al.* (2015) developed non-homogeneous, non-stationary extensions of the (reversible)

HKY85 and GTR models, this was achieved by keeping the exchangeability parameters fixed across the tree, so that only the stationary probabilities varied. For the RY5.6b model, we take a similar approach, allowing only the parameter $\boldsymbol{\rho} \in \mathcal{S}_4$, which serves as a proxy for the stationary distribution, to vary across branches. The parameter α , controlling the differences between the rates of transition and transversion, is held constant. For the RY8.8 model, there is no corresponding partition of the parameters, and so we allow all parameters in $\boldsymbol{\rho} \in \mathcal{S}_8$ to vary from branch to branch.

This yields non-homogeneous RY5.6b and RY8.8 models in which a set of branch-specific simplex-valued parameters $\{\boldsymbol{\rho}_1, \dots, \boldsymbol{\rho}_{B-1}\}$ induce corresponding heterogeneity in the theoretical stationary distribution across branches. The models are therefore non-stationary, with step-changes in the stationary distribution at each speciation event. This enables us to capture changes in sequence composition over evolutionary time.

Our non-homogeneous, non-stationary, locally non-reversible models offer two main advantages over locally reversible counterparts, such as the non-homogeneous HKY85 model. First, as we investigate further in Section 5.2 of Chapter 5, the property of non-reversibility can provide an additional source of likelihood information about the direction of time, and hence the position of the root. Second, if we prune n_0 taxa from a tree on n -species, the non-homogeneous Lie Markov model on DNA characters in Ω^n induces a distribution on the reduced DNA characters in Ω^{n-n_0} . Due to Lie Markov models being closed under matrix multiplication, this distribution could, in *most* cases, have been constructed directly from a non-homogeneous Lie Markov model over the $(n - n_0)$ -taxa subtree. (We note that this cannot be guaranteed in *all* cases because it is theoretically possible for the product of two Lie Markov rate matrices $Q_1, Q_2 \in \mathcal{M}$, to yield a rate matrix $Q = \log\{\exp(Q_1)\exp(Q_2)\} \in \mathcal{M}$ which is not stochastic; see Woodhams *et al.* (2015) for an empirical investigation.) Non-homogeneous, non-stationary but locally reversible models generally lack this property of mathematical consistency, for example, the GTR model does not possess this property (Sumner *et al.*, 2012a,b).

4.2 Prior

In a homogeneous model, the instantaneous rate matrix which characterises the evolutionary process is the same on all branches of the phylogeny. In our non-homogeneous models, it can change from branch to branch. Letting $K = 4$ and $K = 8$ for the RY5.6b and RY8.8 models, respectively, we adopt a prior in which the branch-specific parameter vectors, $\boldsymbol{\rho}_1, \dots, \boldsymbol{\rho}_{B-1} \in \mathcal{S}_K$, are positively correlated. This provides flexibility, whilst retaining some of the benefits of the homogeneous model, by allowing information to be shared between branches. As we move from one branch to its descendants, we do not anticipate a substantial change in the evolutionary process. We therefore build explicit

dependence on recent ancestors into our joint prior through the assignment of a stationary, first order autoregression over a reparameterised set of vectors $\boldsymbol{\rho}_b \in \mathbb{R}^{K-1}$, $b = 1, \dots, B-1$, each of which is related to the corresponding $\boldsymbol{\rho}_b \in \mathcal{S}_K$ through a linear mapping, followed by multinomial logit transformation. This prior was developed by Heaps *et al.* (2014), where it was used for the composition vectors $\boldsymbol{\pi}_b$ on each branch of a tree.

Following their formulation but with $\boldsymbol{\rho}_b$ rather than $\boldsymbol{\pi}_b$, we have

$$p(\boldsymbol{\rho}_r, \dots, \boldsymbol{\rho}_B | \tau) = p(\boldsymbol{\rho}_r | \tau) \prod_{b=1}^B p(\boldsymbol{\rho}_b | \boldsymbol{\rho}_{a(b)}, \tau),$$

where r is the index of the rooting branch and $a(b)$ is the index of the branch (or root) which is ancestral to branch b and τ is the topology. To construct a prior distribution with this structure and which is exchangeable over the components of $\boldsymbol{\rho}$, a multinomial logit reparameterisation is used in which, for branch j

$$\rho_{bk} = \frac{e^{\alpha_{bk}}}{\sum_{m=1}^K e^{\alpha_{bm}}}, \quad k = 1, \dots, K,$$

where $\sum_{k=1}^K \alpha_{bk} = 0$ and $\alpha_{bk} \in \mathbb{R}$ for $k = 1, \dots, K$. An exchangeable prior for $\boldsymbol{\rho}$ implies that the elements of $\boldsymbol{\rho}$ have common mean, common variance and common covariance. This is a desirable property because it means there is nothing in our prior beliefs to distinguish between **A**, **G**, **C** and **T** in the non-homogeneous RY5.6b model and rates of change in the non-homogeneous RY8.8 model. To construct an exchangeable prior for the elements of $\boldsymbol{\rho}_b = (\rho_{b1}, \dots, \rho_{bK})^T$ an exchangeable prior must be used for the elements of $\boldsymbol{\alpha}_b = (\alpha_{b1}, \dots, \alpha_{bK})^T$. It is difficult to construct an exchangeable prior for $\boldsymbol{\alpha}_b$ because of the constrained nature of its space. Therefore, we use new parameters $\boldsymbol{\rho}_b = (\rho_{b1}, \dots, \rho_{b,K-1})^T \in \mathbb{R}^{K-1}$ through the linear mapping $\boldsymbol{\alpha}_b = \mathbf{H}\boldsymbol{\rho}_b$, where \mathbf{H} is a $K \times (K-1)$ matrix with (j, k) -th entry

$$h_{jk} = \begin{cases} 0 & \text{if } j < k, \\ d_k & \text{if } j = k \\ -d_k/(K-k) & \text{if } j > k, \end{cases} \quad (4.1)$$

for $j = 1, \dots, K$, $k = 1, \dots, K-1$. Here $d_1 = 1$ and $d_k = d_{k-1} \sqrt{1 - 1/(K-k+1)^2}$ for $k = 2, \dots, K-1$. For example, when $K = 4$, the corresponding matrix would be

$$\mathbf{H} = \begin{pmatrix} 1 & 0 & 0 \\ -1/3 & 2\sqrt{2}/3 & 0 \\ -1/3 & -\sqrt{2}/3 & \sqrt{2/3} \\ -1/3 & -\sqrt{2}/3 & -\sqrt{2/3} \end{pmatrix}.$$

Noting that $\text{AR}(p)$ denotes an autoregressive process of order p , we take independent stationary $\text{AR}(1)$ processes for each of the collections $(\varrho_{rk}, \dots, \varrho_{Bk})$, $k = 1, \dots, K - 1$, so that

$$p(\boldsymbol{\varrho}_r, \dots, \boldsymbol{\varrho}_B | \tau) = \prod_{k=1}^{K-1} \left[p(\varrho_{rk} | \tau) \prod_{b=1}^B p(\varrho_{bk} | \varrho_{a(b),k}, \tau) \right], \quad (4.2)$$

where

$$\varrho_{rk} | \tau \sim \text{N}(0, v_\varrho / (1 - p_\varrho^2)) \quad \text{and} \quad \varrho_{bk} | \varrho_{a(b),k}, \tau \sim \text{N}(p_\varrho \varrho_{a(b),k}, v_\varrho)$$

in which $p_\varrho \in [0, 1]$ and $v_\varrho \in \mathbb{R}^+$ are fixed hyperparameters that control the marginal variances and covariances of the $\boldsymbol{\varrho}$ and hence $\boldsymbol{\rho}$. This is a prior for $\boldsymbol{\varrho}_b$ that is exchangeable over its elements and given the topology τ , $\varrho_{b1}, \dots, \varrho_{b,K-1}$ have zero prior mean and are uncorrelated with variance $v_\varrho / (1 - p_\varrho^2)$. This prior with the choice of H defined above in (4.1) is an exchangeable prior on the elements of $\boldsymbol{\alpha}_b$ and hence on those of $\boldsymbol{\rho}_b$.

In the RY5.6b model, the instantaneous rate matrix \mathbf{Q}_b on branch b depends on the parameter $\alpha \in [0, 1]$ in addition to the stochastic vector $\boldsymbol{\rho}_b \in \mathcal{S}_4$. Conditional on τ , we factorise the joint prior of α and the $\boldsymbol{\rho}_b$ as $\pi(\alpha, \boldsymbol{\rho}_1, \dots, \boldsymbol{\rho}_{B-1} | \tau) = \pi(\alpha) \pi(\boldsymbol{\rho}_1, \dots, \boldsymbol{\rho}_{B-1} | \tau)$ and assign a flat distribution to α , that is, $\alpha \sim \text{Beta}(1, 1)$. Additionally, both models also have the branch lengths $\boldsymbol{\ell} = (\ell_1, \dots, \ell_B)^T \in \mathbb{R}_+^B$, the shape parameter $\phi \in \mathbb{R}_+$ in the discretised gamma distribution for the rate variation across sites (see Section 3.3) and the rooted tree topology τ . We assign a prior $\ell_b \sim \text{Exp}(\lambda)$ to the branch lengths and choose $\lambda = 10$, which reflects a strong prior belief that the expected number of substitutions per site is 0.1 (and thus the branch lengths are short). The joint prior for the branch lengths is $\pi(\boldsymbol{\ell}) = \prod_{i=1}^B \pi(\ell_i)$. We assign a distribution $\phi \sim \text{Ga}(g_\phi, h_\phi)$, with $g_\phi = h_\phi = 10$, for the shape parameter in the discretised gamma distribution. This gives a marginal mean of 1 for the site-specific substitution rates. Conditionally on ϕ , the square of their coefficient of variation is $\text{IG}(10, 10)$. This is roughly centred at 1 with most (97%) of its mass below 2. The rooted topology is given a prior according to the Yule model of speciation, which assumes that at any given time each of the species is equally likely to undergo a speciation event. Defining a root type of size $j : (n - j)$, $j \in \{1, \dots, \lfloor n/2 \rfloor\}$, as the set of all rooted trees with j taxa on one side of the root and $n - j$ on the other, the Yule model generates a distribution in which near equal probability is assigned to root types of all sizes (Cherlin *et al.*, 2017). Under the Yule model, the probability of generating an n -taxa tree T is

$$\frac{2^{n-1}}{n!} \left(\prod_{v \in V_0} \lambda_v \right)^{-1}, \quad (4.3)$$

where λ_v is the number of internal vertices that are descendants of v (including v), which is one less than the number of leaves descending from v , and V_0 is the set of internal vertices (Steel & McKenzie, 2001). This probability is the number of associated labelled histories

$(n-1)!/\prod_{v \in V_0} \lambda_v$ for a rooted binary tree with n taxa divided by the total number of labelled histories for n taxa $n!(n-1)!/2^{n-1}$.

4.3 Posterior inference via MCMC

The unknowns in the model comprise the rooted tree topology τ , the branch lengths $\boldsymbol{\ell} = (\ell_1, \dots, \ell_B)^T \in \mathbb{R}_+^B$ and the shape parameter $\phi \in \mathbb{R}_+$ in the discretised gamma distribution for rate variation across sites. Additionally, we have the set of substitution model parameters, denoted by \mathcal{Q} , where $\mathcal{Q} = \{\alpha, \boldsymbol{\varrho}_1, \dots, \boldsymbol{\varrho}_{B-1}\}$ for the non-homogeneous RY5.6b model and $\mathcal{Q} = \{\boldsymbol{\varrho}_1, \dots, \boldsymbol{\varrho}_{B-1}\}$ for the non-homogeneous RY8.8 model. These parameters govern the distribution at the root of the tree, say $\boldsymbol{\pi}_0$, and the instantaneous rate matrices, Q_1, \dots, Q_B , on each branch. A directed acyclic graph for our hierarchical models is shown in Figure 4.1. As this graph demonstrates, the rooted topology τ , the branch lengths $\boldsymbol{\ell}$ and the gamma shape parameter ϕ are treated as independent in the prior distribution. The individual branch lengths ℓ_1, \dots, ℓ_B are independent and identically distributed. Note that α is conditionally independent of $\boldsymbol{\varrho}_1, \dots, \boldsymbol{\varrho}_{B-1}$ given τ in the prior distribution of the RY5.6b model.

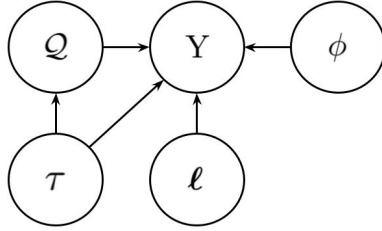


Figure 4.1: Directed acyclic graph for our non-homogeneous models, where $\mathcal{Q} = \{\alpha, \boldsymbol{\varrho}_1, \dots, \boldsymbol{\varrho}_{B-1}\}$ under the non-homogeneous RY5.6b model and $\mathcal{Q} = \{\boldsymbol{\varrho}_1, \dots, \boldsymbol{\varrho}_{B-1}\}$ for the non-homogeneous RY8.8 model.

The posterior distribution for the unknowns can be expressed as

$$\pi(\tau, \boldsymbol{\ell}, \phi, \mathcal{Q} | \mathbf{y}) \propto p(\mathbf{y} | \tau, \boldsymbol{\ell}, \phi, \mathcal{Q}) \pi(\tau, \boldsymbol{\ell}, \phi, \mathcal{Q}),$$

in which $p(\mathbf{y} | \tau, \boldsymbol{\ell}, \phi, \mathcal{Q})$ is the likelihood of the alignment \mathbf{y} and $\pi(\tau, \boldsymbol{\ell}, \phi, \mathcal{Q})$ is the prior density, which can be written as $\pi(\mathcal{Q} | \tau) \pi(\tau) \pi(\boldsymbol{\ell}) \pi(\phi)$. The likelihood is calculated as $p(\mathbf{y} | \tau, \boldsymbol{\ell}, \phi, \mathcal{Q}) = \prod_{i=1}^m p(\mathbf{Y}_i = \mathbf{y}_i | \tau, \boldsymbol{\ell}, \phi, \mathcal{Q})$ in which $\mathbf{Y}_i \in \Omega^n$ is the DNA character at site i and m is the total number of sites. The probability of the observed character \mathbf{y}_i at site i is given by

$$\Pr(\mathbf{Y}_i = \mathbf{y}_i | \tau, \boldsymbol{\ell}, \phi, \mathcal{Q}) = \frac{1}{4} \sum_{k=1}^4 \sum_X \pi_{0,X(0)} \prod_{\text{edges } b=(v,w)} p_{b,X(v),X(w)} \{r_k(\phi) \ell_b\}.$$

Here v and w are the vertices (nodes) at the two ends of edge b with length ℓ_b , $X(u)$ is the character at vertex u , $u = 0$ denotes the root vertex and $P_b\{r_k(\phi)\ell_b\} = [p_{bhi}\{r_k(\phi)\ell_b\}] = \exp\{r_k(\phi)\ell_b Q'_b\}$ is the transition matrix associated with edge b for discretised site rate category k . The outer sum is over the four rate categories of the discretised gamma distribution for rate variation across sites. The inner sum is the likelihood of observing the characters at the leaves at site i , which was given in (2.3) (without the consideration of different rate categories) in Section 2.1.3 and can be calculated using Felsenstein's pruning algorithm (Felsenstein, 1973, 1981) (also given in Section 2.1.3).

The posterior density $\pi(\tau, \ell, \phi, \mathcal{Q}|y)$ is not available analytically. Therefore, we build up a numerical approximation by generating samples from the posterior using a Metropolis-within-Gibbs sampling scheme, which iterates through a series of updates for each unknown. A general description of such a scheme was given in Section 2.3.2 and we describe the specific scheme for our models in Section 4.3.2. Before doing so, we describe the various proposal distributions used in the scheme.

4.3.1 Proposal distributions

Parameters which lie in \mathbb{R} or \mathbb{R}_+ are updated using standard proposal distributions. Gaussian random walks are used for the reparameterised branch-specific parameters $\boldsymbol{\varrho}_b$. We have $\boldsymbol{\varrho}_b^*|\boldsymbol{\varrho}_b \sim N_{K-1}(\boldsymbol{\varrho}_b, s^2 I_{K-1})$, where I_{K-1} is a $(K-1) \times (K-1)$ identity matrix, $\boldsymbol{\varrho}_b$ is the current value and s^2 is the innovation standard deviation, which is a tuning parameter. For the parameter $\alpha \in [0, 1]$ in the non-homogeneous RY5.6b model we generate proposals α^* from a Beta distribution which is roughly centred at the current value α , namely $\alpha^*|\alpha \sim \text{Beta}(s_1\alpha + s_2, s_1(1 - \alpha) + s_2)$. Here $s_1 \in \mathbb{R}_+$ and $s_2 \in \mathbb{R}_+$ are tuning parameters. The first affects the variance of the proposal and should be tuned to adjust the acceptance rate. The second helps to prevent the sampler from sticking at the boundaries of the unit interval and should be set close to zero; for example, $s_2 = 0.005$ (Germain, 2010). Both the branch lengths ℓ and the shape parameter ϕ in the discretised gamma distribution for the rate variation across sites are updated using log normal random walk proposals. That is, a new branch length is proposed with a $\ell_b^*|\ell_b \sim \text{LN}(\log(\ell_b), s_\ell^2)$ distribution, where ℓ_b is the current branch length's value and s_ℓ is the tuning parameter. Similarly, a new shape parameter value is proposed via a $\phi^*|\phi \sim \text{LN}(\log(\phi), s_\phi^2)$ distribution, where ϕ is the current value and s_ϕ is the tuning parameter. Both s_ℓ and s_ϕ are used to alter acceptance rates in their respective proposals.

Finally, the rooted topology τ is updated using standard proposals for topological moves, the *nearest neighbour interchange* (NNI), *subtree prune and regraft* (SPR) and proposals to alter the root position. We base our descriptions of these proposals on the descriptions given in Heaps *et al.* (2014) and Cherlin (2016). The NNI algorithm works by swapping two subtrees on the two sides of a branch. First, an internal branch e is

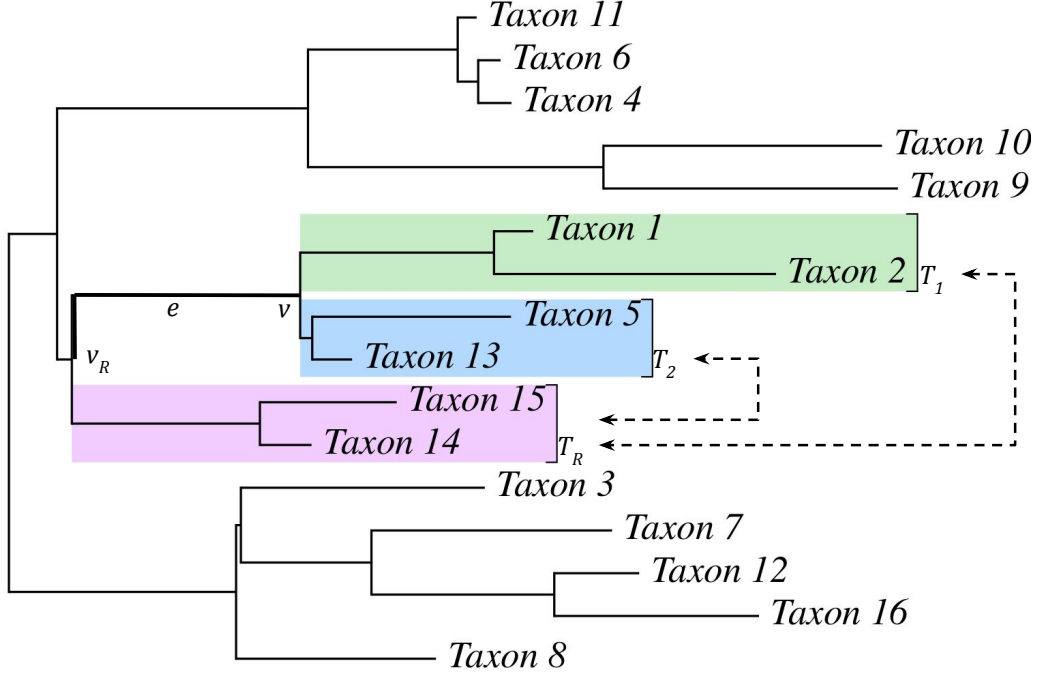
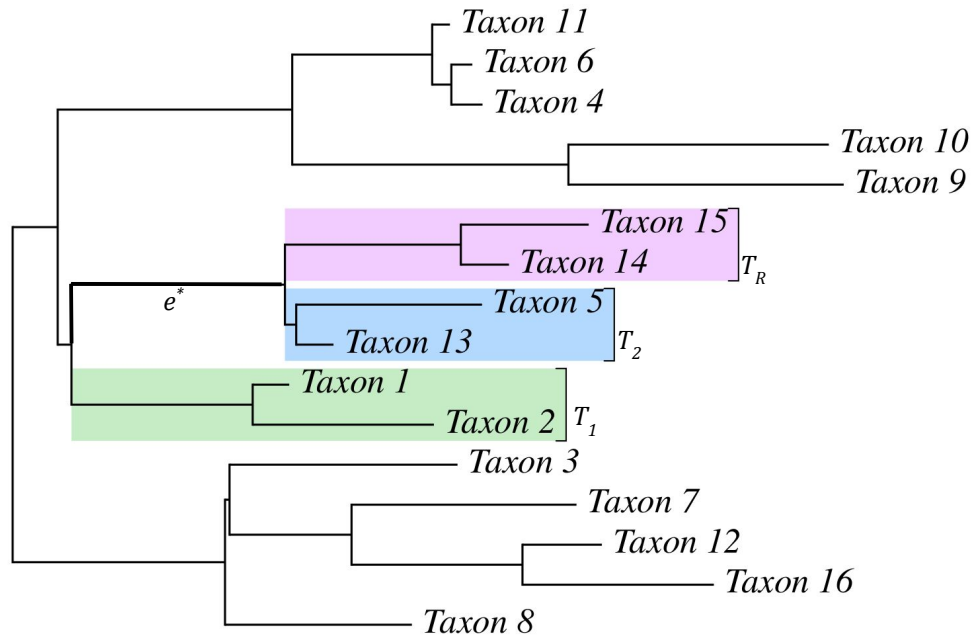


Figure 4.2: An illustration of the NNI move. An internal branch e is chosen uniformly at random from the set of internal branches not adjacent to the root. Either subtree T_1 or T_2 descended from the vertex v is swapped with the subtree T_R descended from v_R .

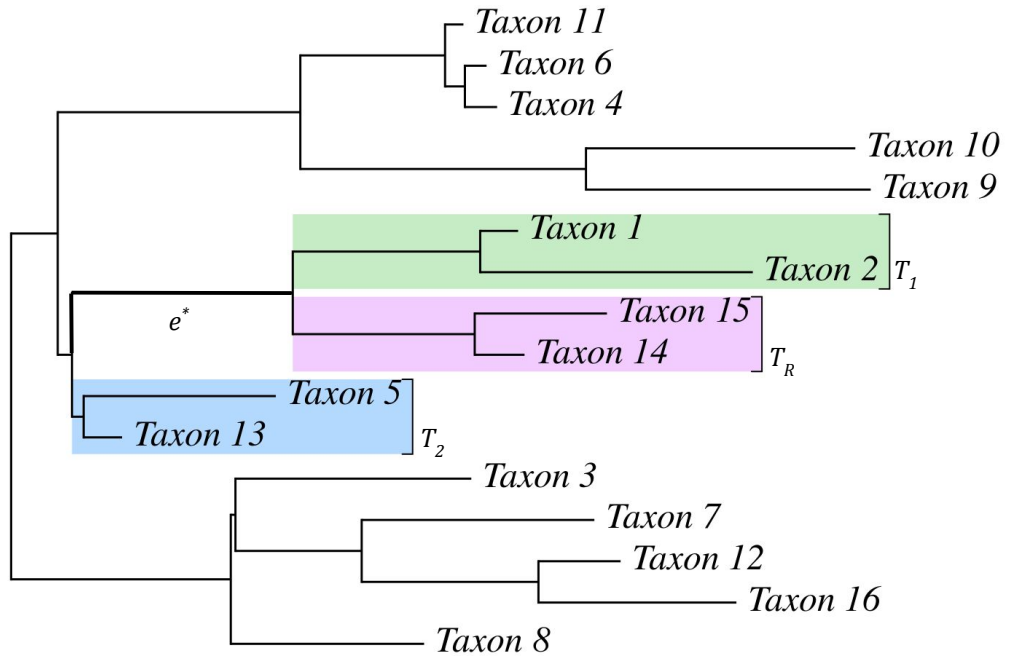
selected uniformly at random, excluding the two branches adjacent to the root of the tree. We denote the vertex on e which is closest to the root as v_R and the vertex closest to the leaves as v . There are two subtrees descending from v , which we denote as T_1 and T_2 , and a subtree T_R descending from v_R . In NNI, there are two possible moves, each occurring with probability of 0.5. Either subtree T_2 is swapped with T_R , or subtree T_1 is swapped with T_R resulting in a new rooted tree topology τ^* . This is illustrated in Figure 4.2 and the two possible resulting trees are shown in Figure 4.3. The branch e is replaced with a new branch e^* , where its length ℓ_{e^*} is proposed via a log normal random walk. We also propose a new value for the branch-specific parameter $\boldsymbol{\varrho}_{e^*}$ using a Gaussian random walk. All other branch lengths and branch-specific parameters $\boldsymbol{\varrho}_b$ ($b \neq e$) remain unchanged. For the NNI proposal, the acceptance probability is the product of the prior ratio, the observed data likelihood ratio and the proposal ratio. Thus, the probability of accepting an NNI move is $\min\{1, A\}$, where

$$\begin{aligned}
 A &= \frac{\pi(\tau^*)}{\pi(\tau)} \times \frac{\pi(\boldsymbol{\varrho}_{e^*}|\tau^*)}{\pi(\boldsymbol{\varrho}_e|\tau)} \times \frac{\pi(\boldsymbol{\ell}^*)}{\pi(\boldsymbol{\ell})} \times \frac{p(y|\tau, \boldsymbol{\ell}^*, \phi, \mathcal{Q})}{p(y|\tau, \boldsymbol{\ell}, \phi, \mathcal{Q})} \times \frac{q(\ell_e|\ell_{e^*})}{q(\ell_{e^*}|\ell_e)} \times \frac{q(\boldsymbol{\varrho}_e|\boldsymbol{\varrho}_{e^*})}{q(\boldsymbol{\varrho}_{e^*}|\boldsymbol{\varrho}_e)} \\
 &= \frac{\pi(\tau^*)}{\pi(\tau)} \times \frac{\pi(\boldsymbol{\varrho}_{e^*}|\tau^*)}{\pi(\boldsymbol{\varrho}_e|\tau)} \times \exp\{\lambda(\ell_e - \ell_{e^*})\} \times \frac{p(y|\tau, \boldsymbol{\ell}^*, \phi, \mathcal{Q})}{p(y|\tau, \boldsymbol{\ell}, \phi, \mathcal{Q})} \times \frac{\ell_e^*}{\ell_e}. \quad (4.4)
 \end{aligned}$$

Note that q denotes a proposal density and that we have calculated A above using the



(a)



(b)

Figure 4.3: Two possible trees resulting from the NNI move shown in Figure 4.2. In (a) the subtree T_1 is swapped with T_R . In (b) the subtree T_2 is swapped with T_R . The length of branch e^* is proposed via a log normal random walk centred on the length of e from the original tree.

general rate parameter λ for the prior exponential distribution for the branch lengths for completeness, rather than substituting in our chosen value of $\lambda = 10$. The prior ratio for the topology τ is calculated using (4.3) in Section 4.2. Additionally, the proposal ratio given above in (4.4) does not include the topology because every tree topology has the same number of neighbouring topologies obtained by a single NNI operation (Heaps *et al.*, 2014; Allen & Steel, 2001), that is $q(\tau^*|\tau) = q(\tau|\tau^*)$, so their ratio is 1.

The SPR topological operation involves pruning off a subtree and grafting it to a different branch of the tree. An illustration of the SPR move is shown in Figure 4.4. Firstly, two internal branches which are not adjacent to the root nor each other are chosen uniformly at random. We denote by e_p the branch which is farther away from the root and let e_g be the branch closer to the root. Additionally, we let v_p be the vertex closest to the root on the branch e_p and let e_a and e_b be the branches containing the vertex v_p . We denote by T the tree evolving from the vertex v_p , where T includes the branch e_p . The subtree T is pruned from the tree and reattached to a point v_g on branch e_g . The reattachment of T divides the branch e_g into two new branches: e_a^* and e_b^* and thus v_g becomes a new vertex which is shared by these two branches. The branches e_a and e_b merge to form a new edge e_g^* and vertex v_p disappears. New branch lengths are proposed for e_a^* and e_b^* subject to the constraint $l_{e_a^*} + l_{e_b^*} = l_{e_g}$. First, a random variable is sampled $u \sim \text{Beta}(2, 2)$ and we set $l_{e_a^*} = u \times l_{e_g}$. The length of branch e_b^* is then set so that the overall branch length (l_{e_g}) is preserved, that is $l_{e_b^*} = (1-u) \times l_{e_g}$. Our choice of parameters for the beta distribution means that $E[u] = 0.5$, which means the expected regrafting point is the middle of branch e_g . New values are proposed for ϱ_{e_x} for $x \in \{a, b, g, p\}$ as described above for the NNI proposal. The acceptance probability of the SPR move is $\min\{1, A\}$, where

$$\begin{aligned}
 A &= \frac{\pi(\tau^*)}{\pi(\tau)} \times \frac{\pi(\varrho_{e_a^*}, \varrho_{e_b^*}, \varrho_{e_g^*}, \varrho_{e_p} | \tau^*)}{\pi(\varrho_{e_a}, \varrho_{e_b}, \varrho_{e_g}, \varrho_{e_p} | \tau)} \times \frac{\pi(\ell^*)}{\pi(\ell)} \times \frac{p(y|\tau, \ell^*, \phi, \mathcal{Q})}{p(y|\tau, \ell, \phi, \mathcal{Q})} \times \frac{q(u^*)}{q(u)} \\
 &\quad \times \left| \frac{\partial(l_{e_a^*}, l_{e_b^*}, l_{e_g^*}, u^*)}{\partial(l_{e_a}, l_{e_b}, l_{e_g}, u)} \right| \times \frac{q(\varrho_{e_a}, \varrho_{e_b}, \varrho_{e_g}, \varrho_{e_p} | \varrho_{e_a^*}^*, \varrho_{e_b^*}^*, \varrho_{e_g^*}^*, \varrho_{e_p}^*, \tau^*)}{q(\varrho_{e_a}^*, \varrho_{e_b}^*, \varrho_{e_g}^*, \varrho_{e_p}^* | \varrho_{e_a}, \varrho_{e_b}, \varrho_{e_g}, \varrho_{e_p}, \tau)} \\
 &= \frac{\pi(\tau^*)}{\pi(\tau)} \times \frac{\pi(\varrho_{e_a^*}, \varrho_{e_b^*}, \varrho_{e_g^*}, \varrho_{e_p} | \tau^*)}{\pi(\varrho_{e_a}, \varrho_{e_b}, \varrho_{e_g}, \varrho_{e_p} | \tau)} \times \frac{\pi(l_{e_a^*}, l_{e_b^*}, l_{e_g^*})}{\pi(l_{e_a}, l_{e_b}, l_{e_g})} \times \frac{p(y|\tau, \ell^*, \phi, \mathcal{Q})}{p(y|\tau, \ell, \phi, \mathcal{Q})} \times \frac{u^*(1-u^*)}{u(1-u)} \\
 &\quad \times \frac{l_{e_g}}{l_{e_a} + l_{e_b}} \times \frac{q(\varrho_{e_a}, \varrho_{e_b}, \varrho_{e_g}, \varrho_{e_p} | \varrho_{e_a}^*, \varrho_{e_b}^*, \varrho_{e_g}^*, \varrho_{e_p}^*, \tau^*)}{q(\varrho_{e_a}^*, \varrho_{e_b}^*, \varrho_{e_g}^*, \varrho_{e_p}^* | \varrho_{e_a}, \varrho_{e_b}, \varrho_{e_g}, \varrho_{e_p}, \tau)}, \tag{4.5}
 \end{aligned}$$

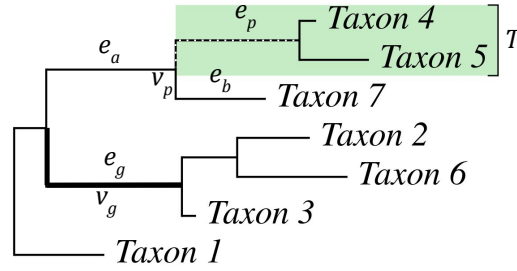
$u^* = l_{e_a}/(l_{e_a} + l_{e_b})$ is the auxiliary variable for the reverse move and

$$\left| \frac{\partial (l_{e_a^*}, l_{e_b^*}, l_{e_g^*}, u^*)}{\partial (l_{e_a}, l_{e_b}, l_{e_g}, u)} \right| = \frac{l_{e_g}}{l_{e_a} + l_{e_b}}$$

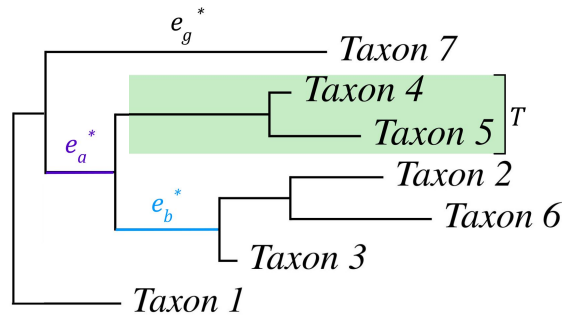
is the Jacobian. We also note that

$$\begin{aligned} \frac{\pi(l_{e_a^*}, l_{e_b^*}, l_{e_g^*})}{\pi(l_{e_a}, l_{e_b}, l_{e_g})} &= \exp\{\lambda(l_{e_a} - l_{e_a^*})\} \times \exp\{\lambda(l_{e_b} - l_{e_b^*})\} \times \exp\{\lambda(l_{e_g} - l_{e_g^*})\} \\ &= \exp\left\{\lambda(l_{e_a} + l_{e_b} + l_{e_g} - l_{e_a^*} - l_{e_b^*} - l_{e_g^*})\right\}. \end{aligned}$$

Again, the proposal ratio does not include the topology in (4.5) above because every tree topology has the same number of neighbouring topologies obtained by a single SPR operation (Heaps *et al.*, 2014; Allen & Steel, 2001). Again, the prior ratio for the topology τ is calculated using (4.3) in Section 4.2.



(a)



(b)

Figure 4.4: An illustration of the SPR move. (a) During the move, the branch e_p (dashed line) and the tree T evolving from it are pruned and reattached to edge e_g . The point of attachment v_g is chosen by dividing the edge e_g using a random variable drawn from Beta(2, 2). (b) After the move is made, the vertex v_p disappears and the branches e_a and e_b are merged to form a new edge e_g^* . The grafting branch e_g is split into two new edges e_a^* and e_b^* by a new vertex v_g which is formed after reattaching the subtree T to e_g .

The root proposal is very similar to the SPR proposal. A branch e_g is selected uniformly at random, with the exception of the two branches adjacent to the root, which we denote by e_a and e_b . We form a new root position by insertion of a new degree two vertex somewhere on e_g and replacing e_g with two new branches e_a^* and e_b^* . Then the branches e_a and e_b are merged to form a single branch e_g^* . New branch lengths are proposed for e_a^* and e_b^* in the identical way as in the SPR proposal. Recall from Section 4.1 that the branches on each side of the root have the same rate matrix Q_r . Therefore, $\boldsymbol{\varrho}_{e_a^*} = \boldsymbol{\varrho}_{e_b^*} (= \boldsymbol{\varrho}_r)$ and we propose new values for $\boldsymbol{\varrho}_{e_a^*}$ and $\boldsymbol{\varrho}_{e_b^*}$ in the same way as in the SPR proposal. The acceptance probability is also calculated in a similar way as in the SPR proposal in (4.5).

In practice, the acceptance rates for all three moves (NNI, SPR and root) are very low and the proposals cannot be tuned to achieve the desirable acceptance rate (see Section 2.3.2). As a result of this, the posterior samples have high autocorrelations, leading to poor mixing of the chain(s). To allow the sampler to move more per overall iteration, we perform multiple NNI, SPR and root moves per MCMC iteration. Acceptance rates are often about an order of magnitude lower for topological moves than the desired acceptance rate for random walk type moves (see Section 2.3.2). Thus, if we perform ten of each topological move per fixed sweep of our algorithm and make the (implausible) assumption that the outcomes of proposals are independent of each other, we can approximately match the expected number of topological moves and random walk type moves.

4.3.2 Metropolis-within-Gibbs sampling scheme

Before applying our models to simulated data and empirical data in Chapter 5, we conclude this chapter by describing our Metropolis-within-Gibbs sampling scheme, which we implement in Java. At each iteration of our MCMC algorithm, we update the substitution model parameters \mathcal{Q} and ϕ , the topology τ and the branch lengths $\boldsymbol{\ell}$. Recall that under the non-homogenous RY5.6b model $\mathcal{Q} = \{\alpha, \boldsymbol{\varrho}_1, \dots, \boldsymbol{\varrho}_{B-1}\}$ and under the non-homogeneous RY8.8 model $\mathcal{Q} = \{\boldsymbol{\varrho}_1, \dots, \boldsymbol{\varrho}_{B-1}\}$.

Our MCMC scheme is a series of Metropolis-Hastings steps (one for each parameter), where one iteration of the entire scheme is as follows:

1. Update $\boldsymbol{\varrho}_b$ one at a time, so for $b = 1, \dots, B - 1$:
 - (a) Propose a new value $\boldsymbol{\varrho}_b^* | \boldsymbol{\varrho}_b \sim N(\boldsymbol{\varrho}_b, s^2 I_{K-1})$.
 - (b) Accept $\boldsymbol{\varrho}_b^*$ with probability $\min\{1, A\}$, where

$$\begin{aligned} A &= \frac{\pi(\boldsymbol{\varrho}_b^* | \tau)}{\pi(\boldsymbol{\varrho}_b | \tau)} \times \frac{p(y | \tau, \boldsymbol{\ell}, \phi, \mathcal{Q}^*)}{p(y | \tau, \boldsymbol{\ell}, \phi, \mathcal{Q})} \times \frac{q(\boldsymbol{\varrho}_b | \boldsymbol{\varrho}_b^*)}{q(\boldsymbol{\varrho}_b^* | \boldsymbol{\varrho}_b)} \\ &= \frac{\pi(\boldsymbol{\varrho}_b^* | \tau)}{\pi(\boldsymbol{\varrho}_b | \tau)} \times \frac{p(y | \tau, \boldsymbol{\ell}, \phi, \mathcal{Q}^*)}{p(y | \tau, \boldsymbol{\ell}, \phi, \mathcal{Q})}, \end{aligned}$$

\mathcal{Q} contains the current value ϱ_b , \mathcal{Q}^* contains the proposed value ϱ_b^* and the prior density $\pi(\varrho_b|\tau)$ can be found using (4.2). Note that the proposal ratio is 1 as we have a symmetric proposal, that is $q(\varrho_b|\varrho_b^*) = q(\varrho_b^*|\varrho_b)$.

2. If the non-homogeneous RY5.6b model is used, update α , otherwise move to step 3.

- (a) Propose a new value $\alpha^*|\alpha \sim \text{Beta}(s_1\alpha + s_2, s_1(1 - \alpha) + s_2)$.
- (b) Accept α^* with probability $\min\{1, A\}$, where

$$\begin{aligned} A &= \frac{\pi(\alpha^*)}{\pi(\alpha)} \times \frac{p(y|\tau, \ell, \phi, \mathcal{Q}^*)}{p(y|\tau, \ell, \phi, \mathcal{Q})} \times \frac{q(\alpha|\alpha^*)}{q(\alpha^*|\alpha)} \\ &= \frac{p(y|\tau, \ell, \phi, \mathcal{Q}^*)}{p(y|\tau, \ell, \phi, \mathcal{Q})} \\ &\quad \times \frac{\alpha^{s_1\alpha^* + s_2 - 1} (1 - \alpha)^{s_1(1 - \alpha^*) + s_2} B(s_1\alpha + s_2, s_1(1 - \alpha) + s_2)}{\alpha^{s_1\alpha + s_2 - 1} (1 - \alpha)^{s_1(1 - \alpha) + s_2} B(s_1\alpha^* + s_2, s_1(1 - \alpha^*) + s_2)}, \end{aligned}$$

\mathcal{Q} contains the current value α , \mathcal{Q}^* contains the proposed value α^* and $B(a, b) = \frac{\Gamma(a)\Gamma(b)}{\Gamma(a+b)}$. Here Γ denotes the gamma function, where $\Gamma(n) = (n - 1)!$. Since the prior distribution is $\alpha \sim \text{Beta}(1, 1)$, both prior densities in the prior ratio are 1 and hence the prior ratio is 1.

3. Update ϕ .

- (a) Propose a new value $\phi^*|\phi \sim \text{LN}(\log(\phi), s_\phi^2)$.
- (b) Accept ϕ^* with probability $\min\{1, A\}$, where

$$\begin{aligned} A &= \frac{\pi(\phi^*)}{\pi(\phi)} \times \frac{p(y|\tau, \ell, \phi^*, \mathcal{Q})}{p(y|\tau, \ell, \phi, \mathcal{Q})} \times \frac{q(\phi|\phi^*)}{q(\phi^*|\phi)} \\ &= \left(\frac{\phi^*}{\phi}\right)^{g_\phi - 1} \exp\{h_\phi(\phi - \phi^*)\} \times \frac{p(y|\tau, \ell, \phi^*, \mathcal{Q})}{p(y|\tau, \ell, \phi, \mathcal{Q})} \times \frac{\phi^*}{\phi} \end{aligned}$$

and recalling that $g_\phi = h_\phi = 10$.

4. Update ℓ one branch at a time. For $b = 1, \dots, B$:

- (a) Propose a new value $\ell_b^*|\ell_b \sim \text{LN}(\log(\ell_b), s_\ell^2)$.
- (b) Accept ℓ_b^* with probability $\min\{1, A\}$, where

$$\begin{aligned} A &= \frac{\pi(\ell_b^*)}{\pi(\ell_b)} \times \frac{p(y|\tau, \ell^*, \phi, \mathcal{Q})}{p(y|\tau, \ell, \phi, \mathcal{Q})} \times \frac{q(\ell_b|\ell_b^*)}{q(\ell_b^*|\ell_b)} \\ &= \exp\{\lambda(\ell_b - \ell_b^*)\} \times \frac{p(y|\tau, \ell, \phi^*, \mathcal{Q})}{p(y|\tau, \ell, \phi, \mathcal{Q})} \times \frac{\ell_b^*}{\ell_b} \end{aligned}$$

ℓ^* is the vector of branch lengths containing ℓ^* and recalling that $\lambda = 10$.

5. Update τ .

- (a) Propose multiple, say ten, NNI moves as described in Section 4.3.1, where each move is accepted with probability $\min\{1, A\}$ and A was given in (4.4).
- (b) Propose multiple, say ten, SPR moves as described in Section 4.3.1, where each move is accepted with probability $\min\{1, A\}$ and A was given in (4.5).
- (c) Propose multiple, say ten, root moves as described in Section 4.3.1, where each move is accepted with probability $\min\{1, A\}$ and A is calculated similarly to (4.5).

Chapter 5

Phylogenetics application

Now that our non-homogeneous, non-stationary, non-reversible phylogenetic models have been fully specified, along with a detailed description of posterior inference, we can apply the models to some data. To show that our models are able to identify root position, we first test our models with simulated data and present the results in this chapter. Finally, we conclude this part of the thesis with a discussion of the results for a biological data set that exhibits compositional heterogeneity.

5.1 Simulation study

For the results of model-based inference on the root position to offer biological insight, the position of the root has to be identifiable under the likelihood. Proving that this is the case for models that are non-stationary, non-reversible, or both, is extremely challenging, except in very special cases (Kaehler, 2017). On the other hand, carefully designed simulation experiments can readily be used to provide empirical evidence of identifiability, and to investigate the conditions under which inference more closely reflects the data-generating mechanism. Therefore, we adopt a simulation-based approach to investigate the identifiability of the root position and underlying topology in our non-homogeneous RY5.6b and RY8.8 models. Specifically, we consider the effects of: (i) different numbers of taxa and sites; (ii) different topologies and branch lengths.

The prior for both models was described in Section 4.2 of Chapter 4. For each analysis, we use the MCMC algorithm described in Section 4.3.2. Two chains are initialised at different starting points and run for 1M iterations. The first 500K are discarded as burn-in and to reduce computational overheads the remaining output is thinned to obtain every 100-th iteration. The standard graphical and numerical diagnostics used in phylogenetic inference (Lartillot *et al.*, 2004) (discussed in Section 2.3.2) are used to assess convergence and mixing.

5.1.1 Different numbers of taxa and sites

To assess the extent to which root inference depends upon the dimensions of the data being analysed, we simulate alignments under the non-homogeneous RY5.6b and RY8.8 models, varying the number of taxa (6, 12, 24) and the number of sites (500, 1000, 2000). First an unrooted tree on 24 taxa is simulated by random resolution of a *star tree*. (A star tree for n taxa has a single vertex with n branches.) This is then rooted to form a balanced tree, that is, a tree with an equal number of taxa on either side of the root, and then branch lengths are sampled from a $\text{Ga}(2, 20)$ distribution. The branch-specific parameter vectors ρ_b are simulated from Dirichlet $\mathcal{D}_4(10, 10, 10, 10)$ and Dirichlet $\mathcal{D}_8(3, 3, 3, 3, 3, 3, 3, 3)$ distributions for the RY5.6b and RY8.8 models, respectively. This gave a degree of heterogeneity which is consistent with what we have inferred in analyses of biological data. For the non-homogeneous RY5.6b model, α is set as 0.5, the mean of its symmetric prior. Likewise, the shape parameter ϕ in the discretised gamma distribution for rate variation across sites is set to the mode of its prior, 0.9. Using this rooted tree and these parameter values, three alignments of 2000 sites were simulated under each model. Taxa, sites, or taxa and sites are then removed from the alignments to give three data sets for each combination of sites and taxa specified above. The taxa that are pruned are chosen uniformly at random, but constrained so that the corresponding tree for each resulting alignment is balanced. This is to avoid any potential confounding with the effect of balance in the rooted topology, which we examine separately in Section 5.1.2. The rooted topologies on 6, 12 and 24 taxa are displayed in Figure 5.1.

For the trees on 6, 12 and 24 taxa, Figures 5.2a, 5.2c and 5.2e, respectively, display the posterior distribution over root splits for the alignments simulated and analysed under the non-homogeneous RY5.6b model. Recall from Section 4.2 that a root type of size $j : (n - j)$, $j \in \{1, \dots, \lfloor n/2 \rfloor\}$ is the set of all rooted trees with j taxa on one side of the root and $n - j$ on the other. For example, a root type of size $1 : (n - 1)$ represents a root split on a pendant branch (a branch connected to a leaf), whereas a root type of size $3 : (n - 3)$ represents a root split between three taxa and the remaining taxa. A root split can inform us of *which* taxa lie on each side of the root, for example, the correct root split for the six-taxa alignment is $(2, 16, 8) : (10, 22, 15)$. Figures 5.3a, 5.3c and 5.3e show the analogous plots for the posterior distribution over unrooted topologies. The black bars highlight the (true) root split or unrooted topology from the tree used to simulate the data. The corresponding plots for the non-homogeneous RY8.8 model are shown in Figures 5.2b, 5.2d and 5.2f and Figures 5.3b, 5.3d and 5.3f respectively. As expected, irrespective of the model or number of taxa, the posterior support for the correct root split tends to increase as the number of sites increases, and the correct root split is more frequently identified as the posterior mode. The same is true for unrooted topologies. Indeed, when there are 2000 sites in the alignment and the non-homogeneous

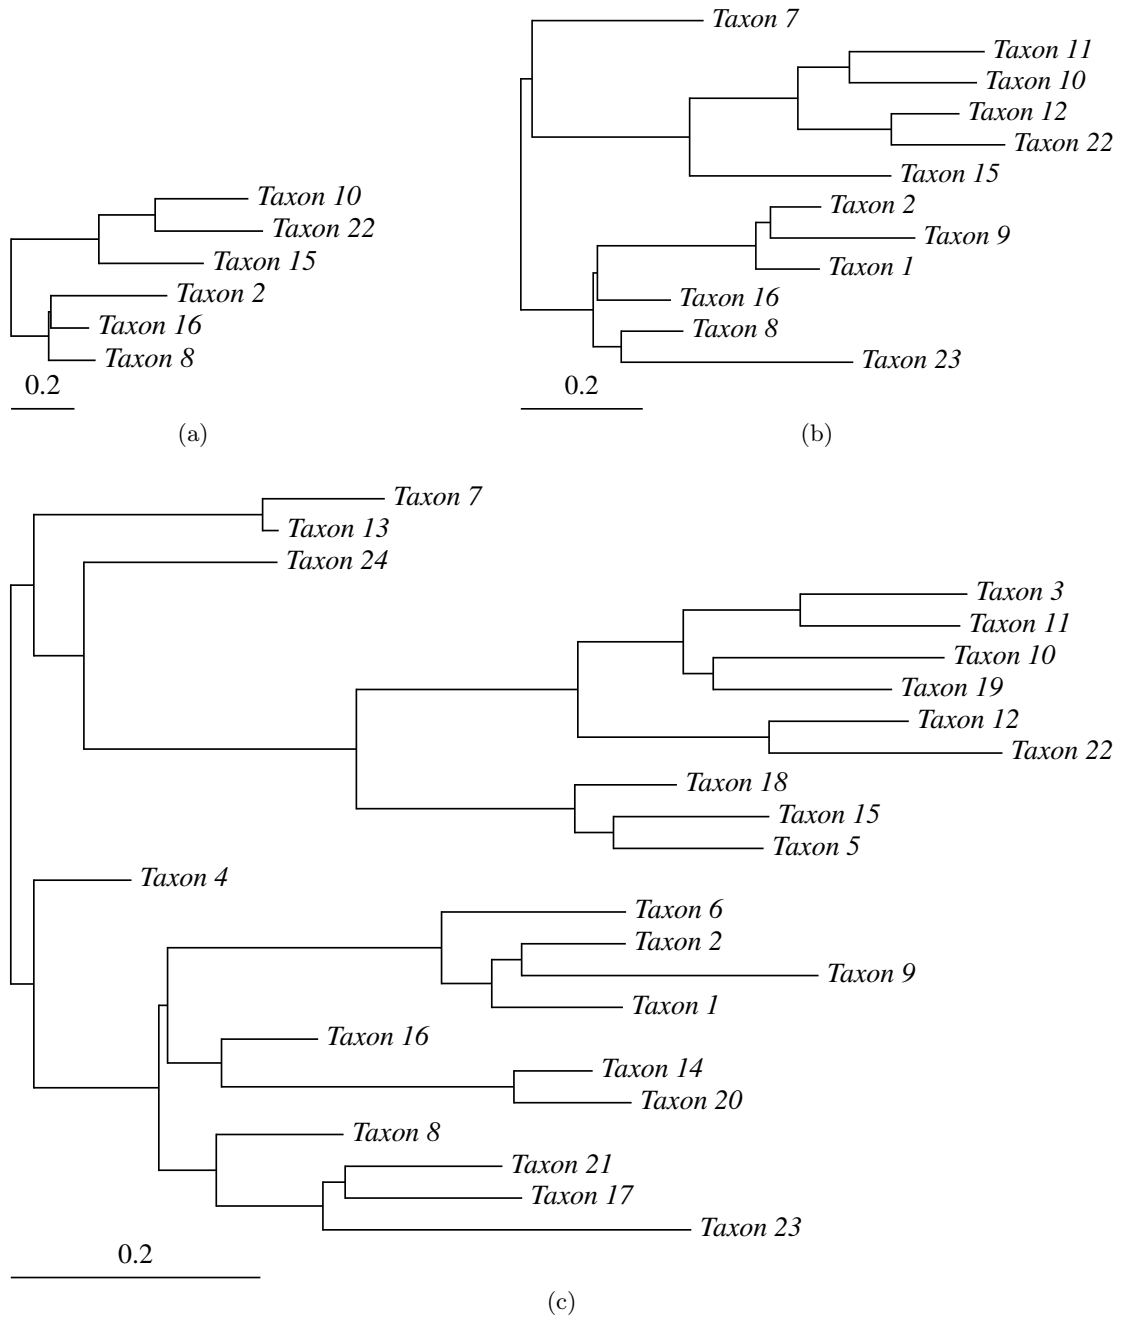


Figure 5.1: Trees used to simulate alignments on (a) 6, (b) 12, (c) 24 taxa.

RY8.8 model is used, the posterior probability for the correct root split is close to one for all tree sizes and the posterior probability of the correct unrooted topology is 0.758 on average across the three 24-taxa alignments. Although increasing the number of taxa leads to quadratic growth in the number of possible root splits and super-exponential growth in the number of possible unrooted topologies, it does not seem to have a detrimental

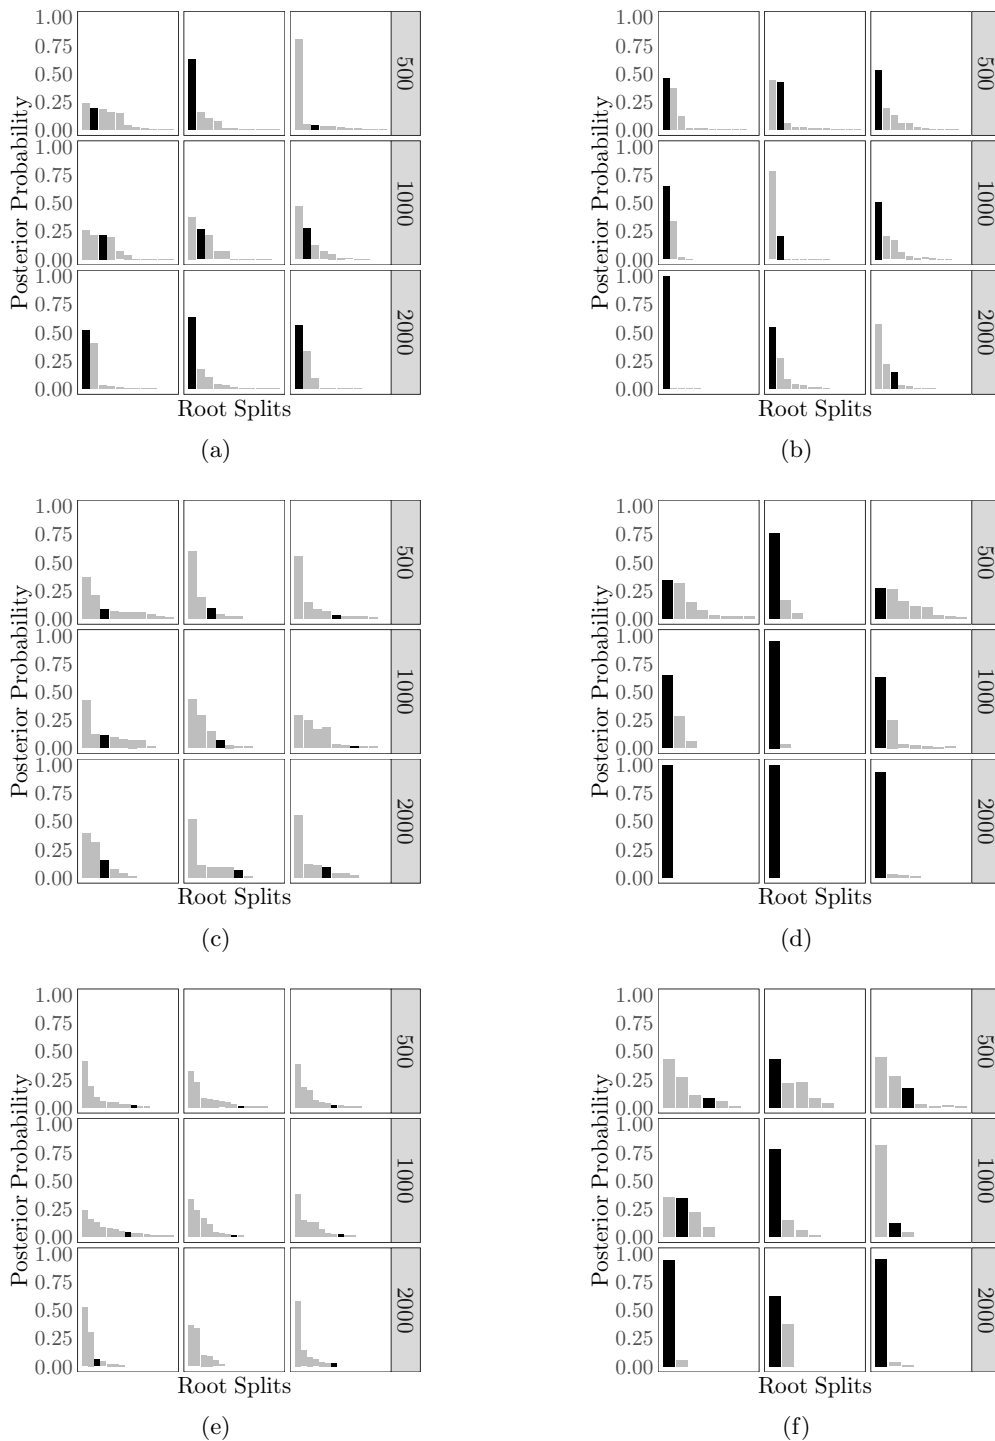


Figure 5.2: Posterior distribution over roots splits when three data sets are simulated and analysed under the non-homogeneous RY5.6b model and the number of taxa is (a) 6, (c) 12, (e) 24; and when three data sets are simulated and analysed under the non-homogeneous RY8.8 model and the number of taxa is (b) 6, (d) 12, (f) 24. The number of sites is displayed to the right of each plot. In every plot, bars are arranged in descending order of posterior probability and the correct root split is highlighted in black. In the plots for 12 and 24 taxa, bars corresponding to probabilities less than 0.01 have been removed to improve readability.

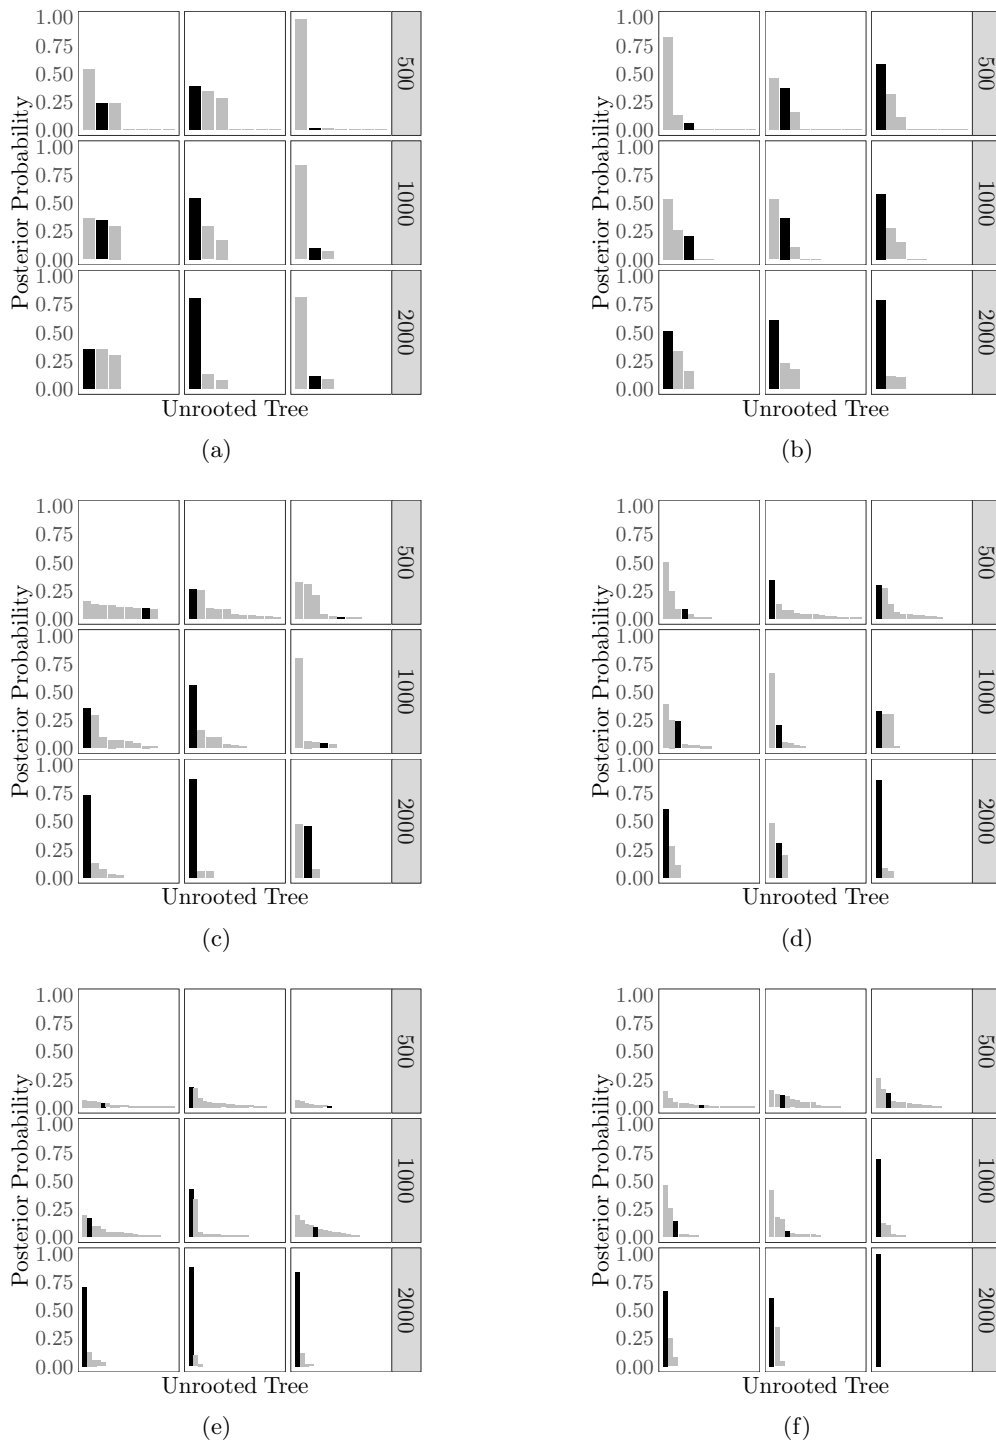


Figure 5.3: Posterior distribution over unrooted topologies when the data are simulated and analysed under the non-homogeneous RY5.6b model and the number of taxa is (a) 6, (c) 12, (e) 24; and when the data are simulated and analysed under the non-homogeneous RY8.8 model and the number of taxa is (b) 6, (d) 12, (f) 24. The number of sites is displayed to the right of each plot. Highlighted in black is the correct unrooted topology. In the plots for 12 and 24 taxa, bars corresponding to probabilities less than 0.01 have been removed to improve readability.

effect on inferential performance for the non-homogeneous RY8.8 model. Unfortunately, the same is not true for the RY5.6b model, under which inference of the root position is generally worse, particularly for larger trees. For example, when the number of taxa is 12 or 24, the correct root split is not recovered as the posterior mode in any simulations. The better rooting performance of the RY8.8 model is likely explained by two factors. First, the model has more parameters that can vary across the tree and induce non-stationary behaviour. Second, identifiability of branch-specific parameters in each model also plays a role, which we explore in the following section.

Parameter identifiability

By design, the non-homogeneous Lie Markov models are highly parameterised. Although of secondary interest in its own right, identifiability of the quantitative model parameters in the posterior is likely to impact on inference of the root position and topology. In the most highly parameterised case, where data are simulated under a 24-taxon tree, we therefore investigate the extent to which the true values of the model parameters can be identified in the posterior. For one of the three alignments, results are summarised in Figures 5.4 and 5.5 for the non-homogeneous RY5.6b and RY8.8 models, respectively, and show the effect of varying the number of sites from 500 through 1000 to 2000. Results for the other two alignments show the same patterns and are not shown. The posterior densities for the global parameters ϕ , in the discretised gamma distribution for rate variation across sites, and α , from the RY5.6b rate matrix, are based on draws from the joint posterior of all unknowns. The branch lengths ℓ_1, \dots, ℓ_B and branch-specific parameter vectors $\boldsymbol{\rho}_1, \dots, \boldsymbol{\rho}_{B-1}$ are only meaningfully labelled on the tree in Figure 5.1c used to simulate the data, say τ_{true} . The densities for the branch-specific parameters are therefore based on draws from the conditional posterior of the model parameters given the topology $\tau = \tau_{\text{true}}$. Posterior densities are visualised for a representative, random selection of branch-specific parameters.

For both models, the true values of the global parameters and branch lengths are identified with high posterior support, even with only 500 sites in the alignment. In nearly all cases, the true values of the branch-specific ρ_{bk} lie within the main body of the posterior density, and the posterior becomes more concentrated around the true value as the number of sites in the alignment increases. However, it is clear that the posteriors for the ρ_{bk} are much more concentrated under the non-homogeneous RY8.8 model than the RY5.6b model. This may be because the additive structure of the RY5.6b rate matrix makes the likelihood less sensitive to changes in the $\boldsymbol{\rho}_b$ and is likely to be partly responsible for the poorer rooting performance of the non-homogeneous RY5.6b model.

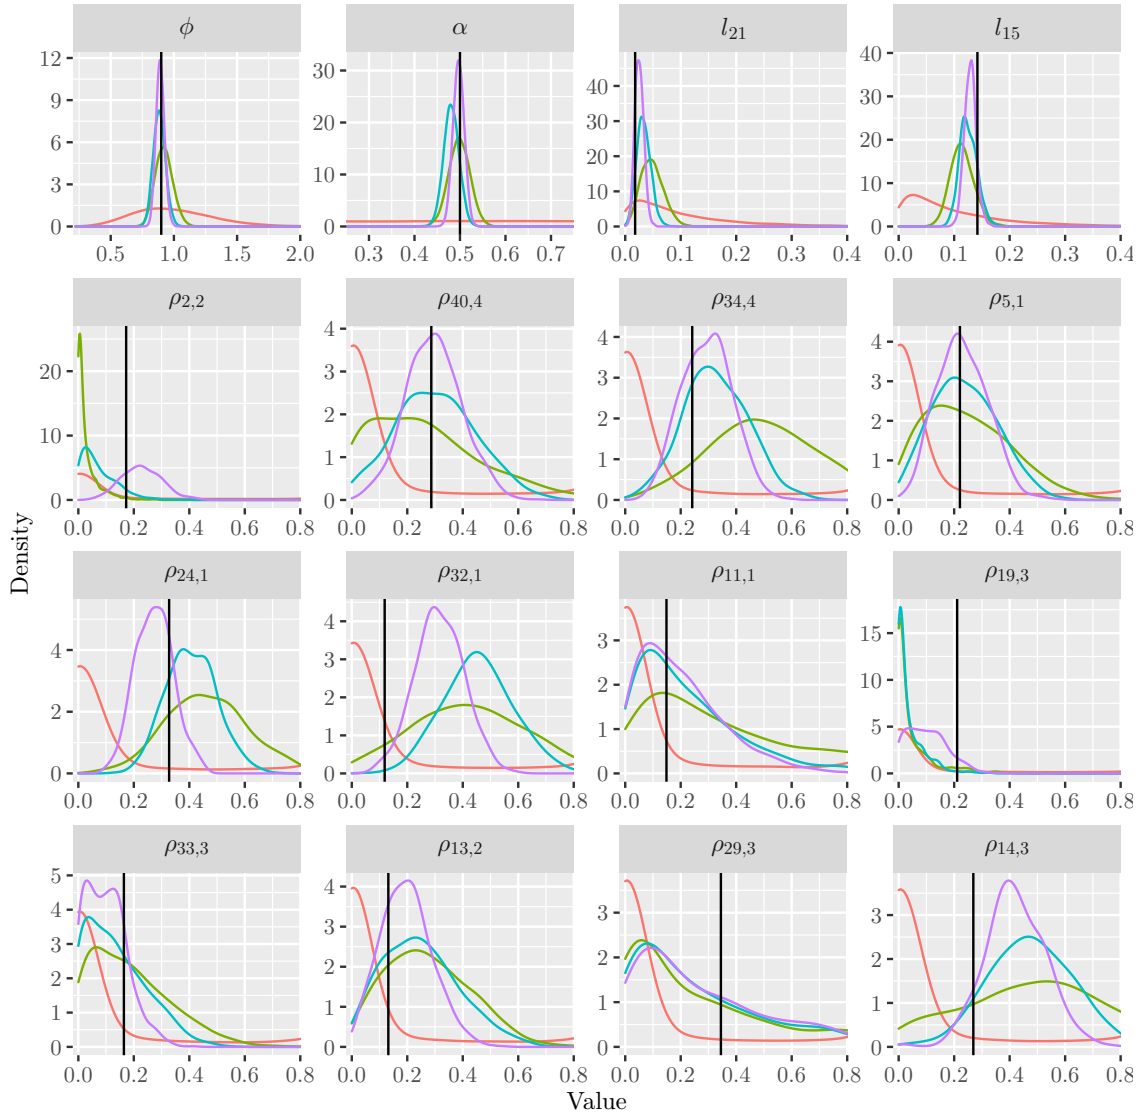


Figure 5.4: Marginal prior and posterior densities for the global model parameters (ϕ and α) and a random sample of branch-specific parameters in the analysis of a 24-taxon alignment simulated under the non-homogeneous RY5.6b model. Posterior densities for the branch-specific parameters are conditional on the rooted topology used to simulate the data. Indicated in the panels are the prior (—) and posterior when the number of sites is 500 (—), 1000 (—) and 2000 (—). The true values of the parameters are indicated by vertical lines.

5.1.2 Different topologies and branch lengths

When investigating their homogeneous, stationary, non-reversible model, Cherlin *et al.* (2017) found root inference to be sensitive to some of the prior-data conflicts that occur commonly in the analysis of biological data. Typically these arise due to incongruent prior and likelihood information about branch lengths and the rooted topology. In our analyses we adopt the near ubiquitous prior for the set of branch lengths, which structures beliefs

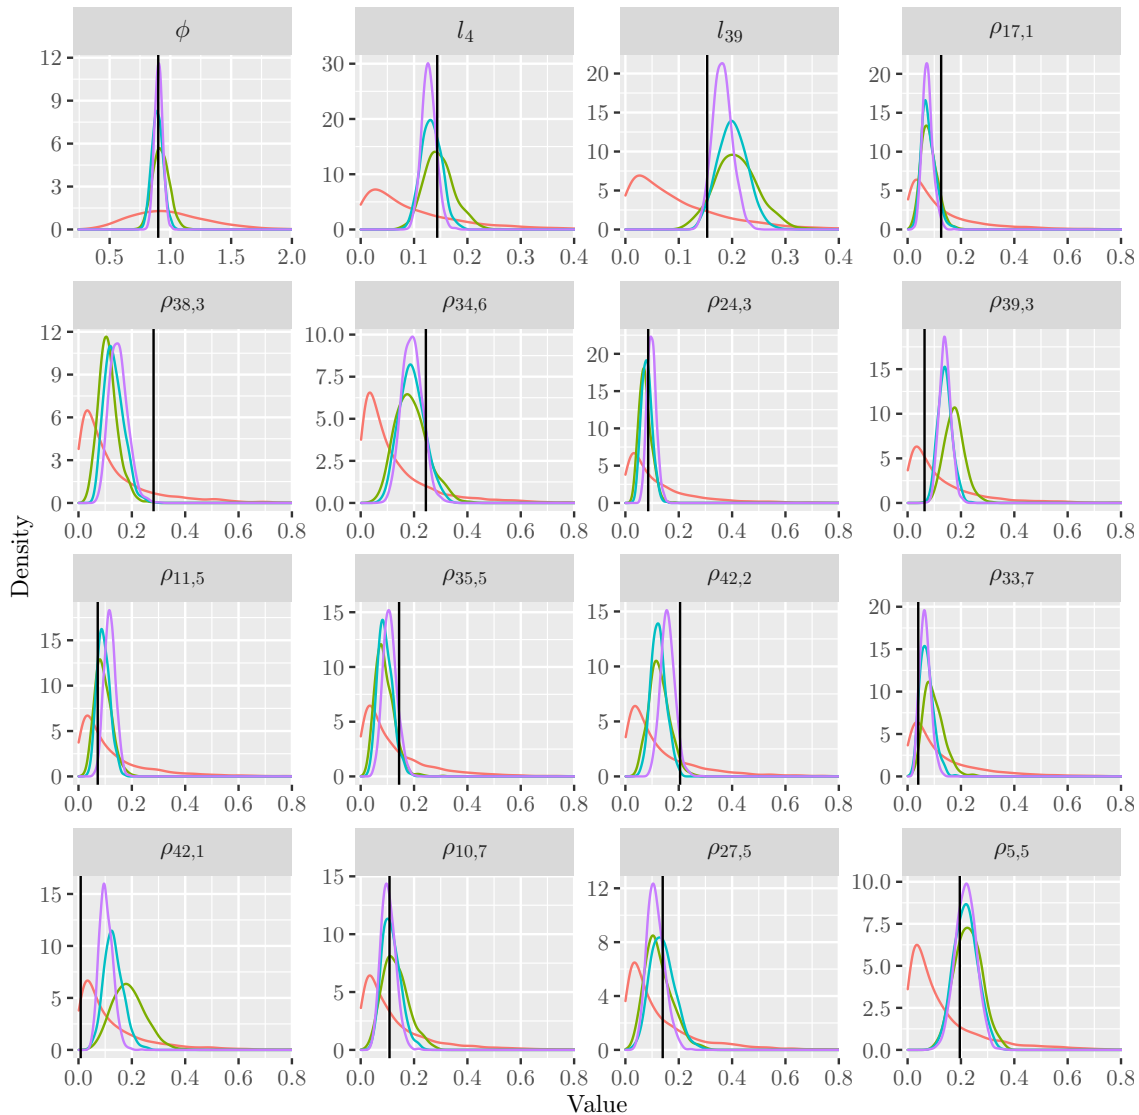


Figure 5.5: Marginal prior and posterior densities for the global model parameter (ϕ) and a random sample of branch-specific parameters in the analysis of a 24-taxon alignment simulated under the non-homogeneous RY8.8 model. Posterior densities for the branch-specific parameters are conditional on the rooted topology used to simulate the data. Indicated in the panels are the prior (—) and posterior when the number of sites is 500 (—), 1000 (—) and 2000 (—). The true values of the parameters are indicated by vertical lines.

as independent $\text{Exp}(10)$ distributions. This prior places 99.9% of its mass below 0.691 and so asserts a strong belief that branch lengths are reasonably short. As a consequence, given an unrooted topology that contains a long branch, the prior supports rooting on this branch in order to split it into two shorter edges. In our analyses we adopt a Yule prior over rooted topologies. As discussed previously in Section 4.2, one of the compelling properties of the Yule distribution is that it assigns near equal probability to root types

of all sizes. However, a combinatorial consequence of this property is that more support is assigned to balanced than unbalanced trees. For n taxa, if n is odd, all root types have the same prior probability, and if n is even, all root types receive the same prior probability except a root type of size $n/2 : n/2$, which has half the prior probability of all other root types. For $k \neq n/2$, the number of trees with a $k : (n - k)$ root split is $\binom{n}{k} \times \frac{(2k-3)!}{2^{(k-2)}(k-2)!} \times \frac{(2(n-k)-3)!}{2^{(n-k-2)}(n-k-2)!}$, that is, the number of ways of choosing k taxa from n multiplied by the number of possible trees for k taxa and $n - k$ taxa. Figure 5.6 shows the number of trees for each root split with $n = 16$ taxa. Clearly, as $k \rightarrow n/2$, the number of trees decreases, and there are many more (unbalanced) trees with a root split on a pendant edge. However, since the prior probability of any particular root type is the same, more prior mass is assigned to any particular balanced tree. In this section, we analyse simulated data to explore posterior sensitivity to prior-data conflicts that arise because of long branches in the underlying unrooted tree or an unbalanced rooted topology.

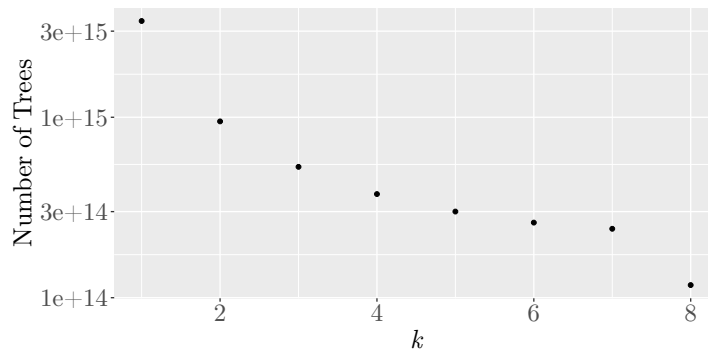


Figure 5.6: Number of $k : (n - k)$ root splits (on log-scale) for $n = 16$ taxa.

We set up this simulation experiment in the same way as Cherlin *et al.* (2017). Simulations are based on the unrooted tree on 16 taxa depicted in Figure 5.7 whose topology was simulated through random resolution of a star tree. Its branch lengths were simulated from a $\text{Ga}(2, 20)$ distribution. Based on this unrooted tree, we construct six different rooted trees by varying the root position, which is placed at the midpoint either of branch E_1 or branch E_2 , and the length of the branch E_1 , which can be the 95%, 50% or 5% quantile of the $\text{Ga}(2, 20)$ distribution:

Tree 1: balanced (rooted on E_1), long root branch (length 0.237);

Tree 2: unbalanced (rooted on E_2), long internal branch (length 0.237);

Tree 3: balanced (rooted on E_1), short root branch (length 0.018);

Tree 4: unbalanced (rooted on E_2), short internal branch (length 0.018);

Tree 5: balanced (rooted on E_1), medium root branch (length 0.084);

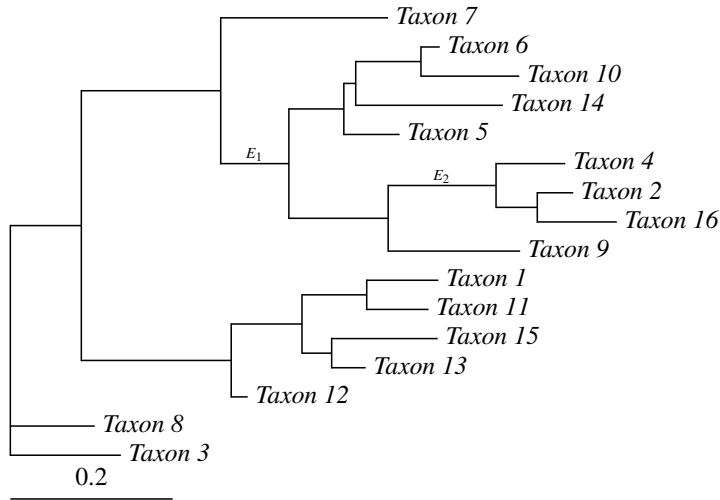


Figure 5.7: Unrooted tree used in simulation experiments to investigate the effects of different topologies and branch lengths on root inference. In the experiments, the tree is rooted at the midpoint of either branch E_1 or E_2 . The tree is depicted with branch E_1 having a “medium” length of 0.084 units. In the experiment, this is varied to 0.237 units (“long”) or 0.018 units (“short”).

Tree 6: unbalanced (rooted on E_2), medium internal branch (length 0.084).

As indicated above, Trees 1, 3 and 5 have a balanced rooted topology, with root type 8 : 8, whilst Trees 2, 4 and 6 are unbalanced, with root type 3 : 13. By substituting $n = 16$ into (4.3) with $k = 8$ and $k = 3$, we can calculate the prior probabilities of the balanced and unbalanced topologies, respectively. In doing so, we find that the Yule prior offers more than six times more mass to the balanced tree and hence the prior and likelihood are likely to be in conflict when the tree is unbalanced. In the unrooted tree associated with Trees 1 and 2, branch E_1 is the longest, whilst for Trees 3 and 4 it is among the shortest. Given the unrooted topology depicted in Figure 5.7, the prior support for a root on edge E_1 increases as the branch becomes longer, and hence will increasingly conflict with the likelihood if E_1 is not the root edge.

For each of the six trees, three 2000-site alignments are simulated and analysed under both non-homogeneous Lie Markov models. The posterior distributions over root splits for the RY8.8 model are shown in Figure 5.8. In general, root inference is good, with the true root recovered as the posterior mode in most cases. This suggests the posterior is reasonably robust to prior-data conflict concerning the rooted topology and branch lengths. Moreover, for Trees 3 – 6, whose unrooted trees do not contain any very long edges, the absence of a marked difference between the results for balanced Trees 3 and 5 and unbalanced Trees 4 and 6 suggests that the prior over rooted topologies imparts little influence over the posterior. It is interesting to note that this was not the case in the work by Cherlin *et al.* (2017) based on their homogeneous, stationary and non-reversible

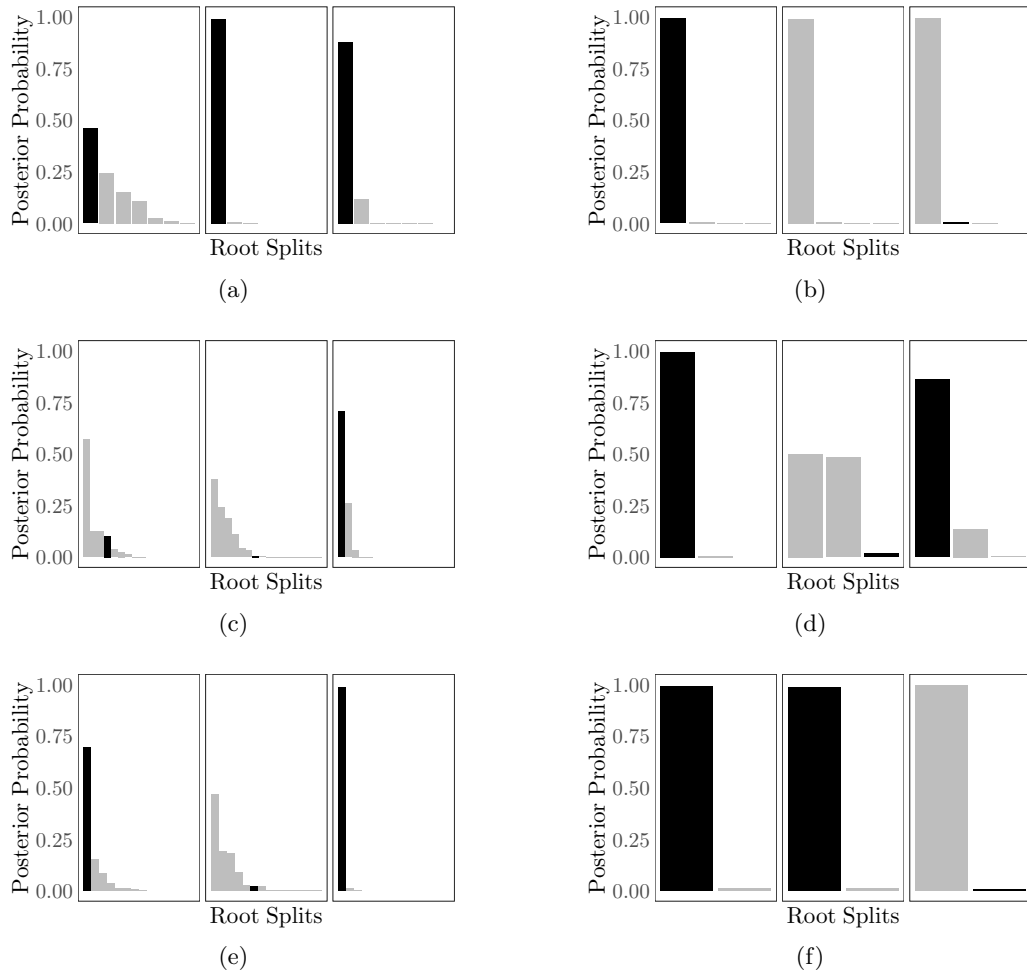


Figure 5.8: Posterior distribution over roots splits when three data sets are simulated and analysed under the non-homogeneous RY8.8 model and the tree used for simulation is Tree (a) 1 (balanced, long root branch), (b) 2 (unbalanced, long internal branch), (c) 3 (balanced, short root branch), (d) 4 (unbalanced, short internal branch), (e) 5 (balanced, medium root branch), (f) 6 (unbalanced, medium internal branch). In every plot, bars are arranged in descending order of posterior probability and the correct root split is highlighted in black.

model. However, comparisons between the results for these four trees and Trees 1 and 2, which do contain a very long edge, suggest that long branches in the unrooted tree can influence posterior inference of the root position. When the long edge is the root edge (Tree 1), the posterior is concentrated around the true root position in the analyses of all three alignments; see Figure 5.8a. However, when the long edge is not the root edge (Tree 2), prior-data conflict arises and the true root only has appreciable posterior support in the analysis of one of the three alignments; see Figure 5.8b.

The corresponding plots for the non-homogeneous RY5.6b model are shown in Figure 5.9. As we found for the results in Section 5.1.1 for the larger trees on 12 or 24 taxa,

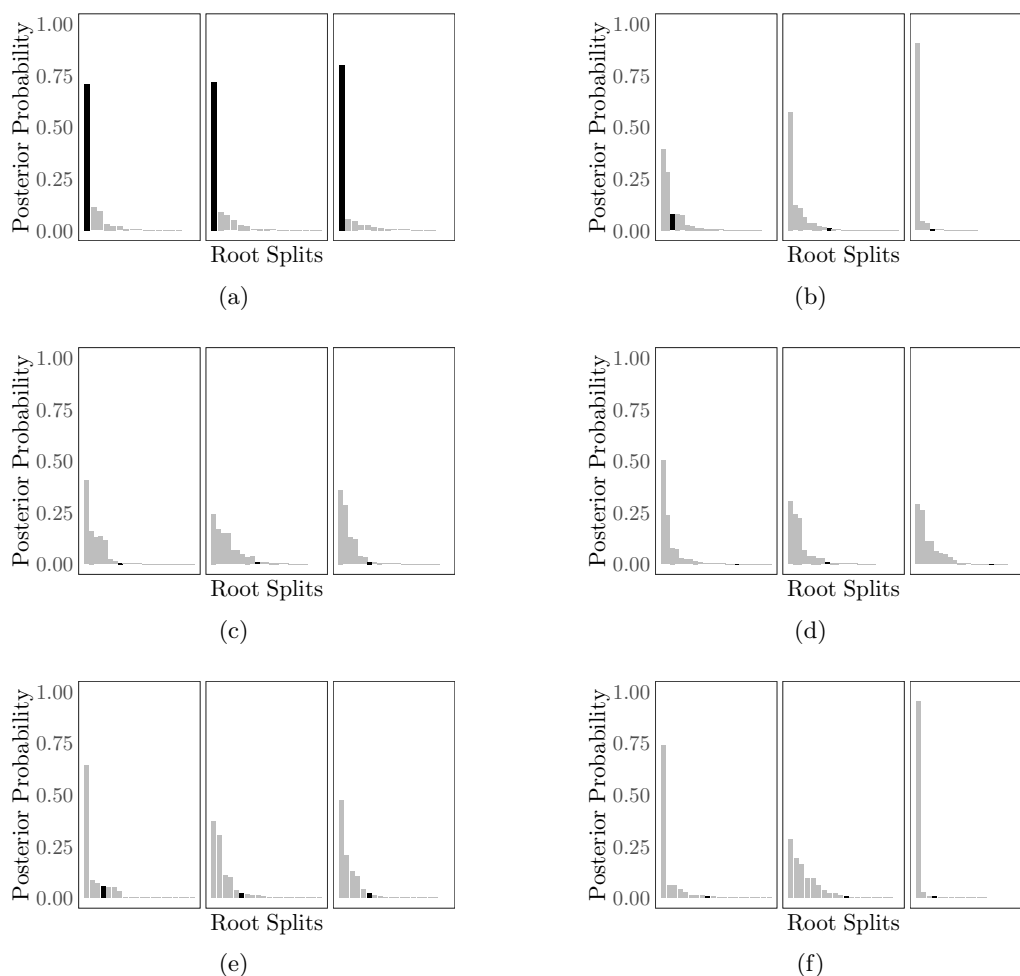


Figure 5.9: Posterior distribution over roots splits when three data sets are simulated and analysed under the non-homogeneous RY5.6b model and the tree used for simulation is Tree (a) 1 (balanced, long root branch), (b) 2 (unbalanced, long internal branch), (c) 3 (balanced, short root branch), (d) 4 (unbalanced, short internal branch), (e) 5 (balanced, medium root branch), (f) 6 (unbalanced, medium internal branch). In every plot, bars are arranged in descending order of posterior probability and the correct root split is highlighted in black.

the true root rarely receives particularly appreciable posterior support. In fact, the only cases where the true root was recovered as the posterior mode were the analyses of the three alignments simulated under Tree 1, where the root edge is a long branch. This further suggests that for alignments of around 2000 sites on a modest number of taxa, the likelihood of a non-homogeneous RY5.6b model does not itself clearly identify the position of the root. However, it is worth noting that for both models and all trees, inference of the unrooted topology was excellent, with the true unrooted topology identified as the posterior mode, with high support, in all cases (see Figures 5.10 and 5.11).

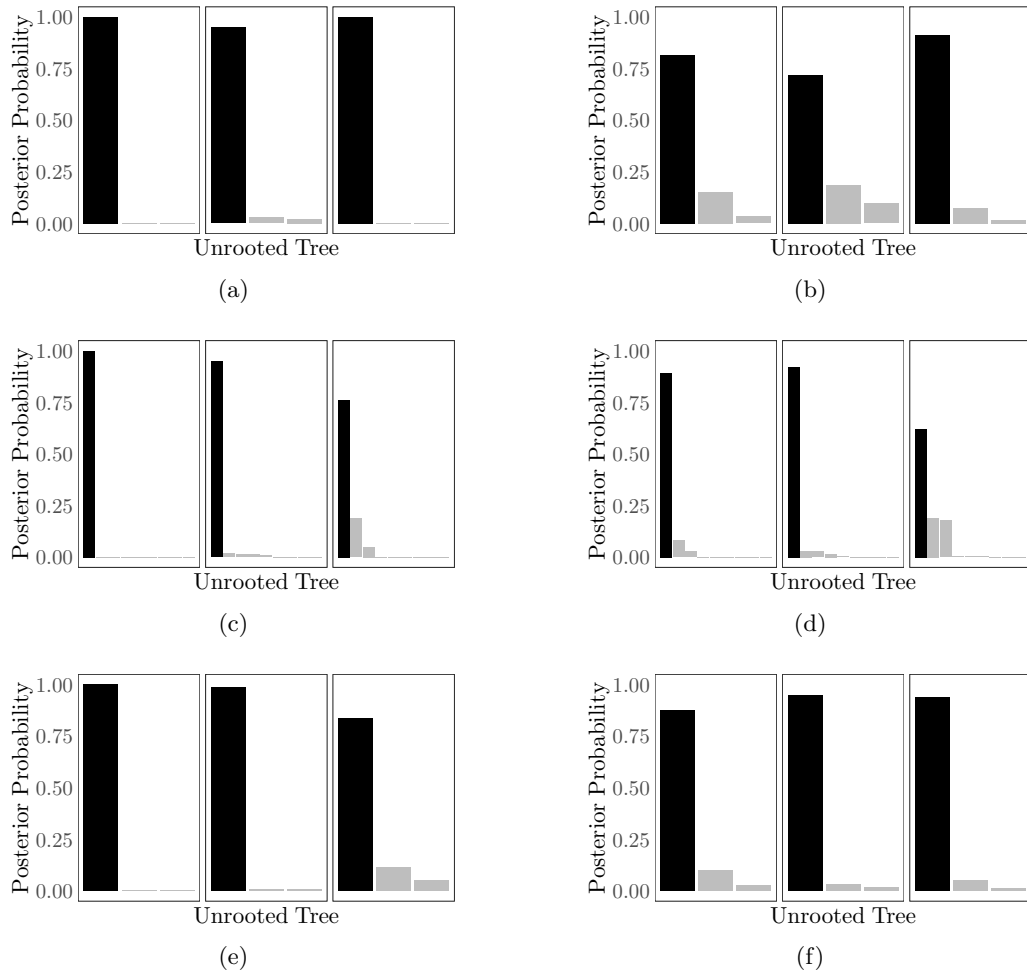


Figure 5.10: Posterior distribution over unrooted topologies when three data sets are simulated and analysed under the non-homogeneous RY8.8 model and the tree used for simulation is Tree (a) 1 (balanced, long root branch), (b) 2 (unbalanced, long internal branch), (c) 3 (balanced, short root branch), (d) 4 (unbalanced, short internal branch), (e) 5 (balanced, medium root branch), (f) 6 (unbalanced, medium internal branch). In every plot, bars are arranged in descending order of posterior probability and the correct unrooted topology is highlighted in black.

5.2 The *Drosophila* data set

To illustrate the benefit of our non-homogeneous model in its facility to bring two sources of information to bear on the rooting problem, we consider an application to a *Drosophila* data set, taken from Tarrío *et al.* (2000). Most models fail to identify a plausible root position. One exception is the model developed by Foster (2004), discussed in Section 3.1. Recall that their model allows the branches of the tree to be allocated to one of $K \ll B$ mixture components, where each component has a different composition vector. They fitted their model with $K = 2, 3, 4, 5$ and found that when more than two mixture components were used the most biologically plausible root was identified.

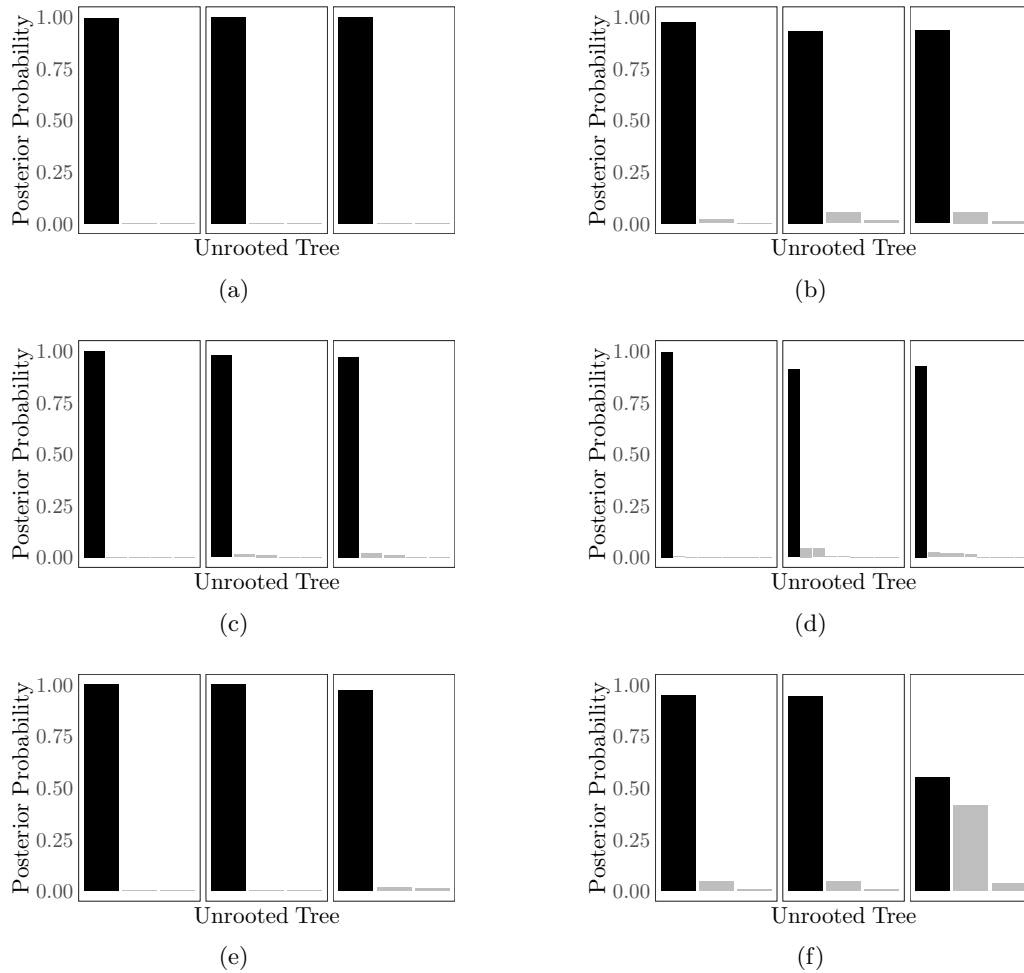


Figure 5.11: Posterior distribution over unrooted topologies when three data sets are simulated and analysed under the non-homogeneous RY5.6b model and the tree used for simulation is Tree (a) 1 (balanced, long root branch), (b) 2 (unbalanced, long internal branch), (c) 3 (balanced, short root branch), (d) 4 (unbalanced, short internal branch), (e) 5 (balanced, medium root branch), (f) 6 (unbalanced, medium internal branch). In every plot, bars are arranged in descending order of posterior probability and the correct unrooted topology is highlighted in black.

The alignment contains 2085 nucleotides (sites) from the xanthine dehydrogenase (Xdh) gene of 17 different species of *Drosophila*. *D. saltans*, *D. prosaltans*, *D. neocordata*, *D. emarginata*, *D. sturtevantii* and *D. subsaltans* form a clade of *saltans*. Three of the species form an outgroup: *D. melanogaster*, *D. virilis* and *D. pseudoobscura*. The remaining species form a clade of *willistoni*. Recall that the terms clade and outgroup were defined in Section 2.1.1. Due to differences between nucleotide compositions in the Xdh gene of the outgroup and ingroup, many models fail to find a biologically plausible root (Tarrío *et al.*, 2000). We compare inferences obtained under six different models:

\mathcal{M}_1 : homogeneous, stationary, reversible GTR model;

\mathcal{M}_2 : homogeneous, stationary, non-reversible RY5.6b model;

\mathcal{M}_3 : homogeneous, stationary, non-reversible RY8.8 model;

\mathcal{M}_4 : non-homogeneous, non-stationary, locally reversible GTR model;

\mathcal{M}_5 : non-homogeneous, non-stationary, locally non-reversible RY5.6b model;

\mathcal{M}_6 : non-homogeneous, non-stationary, locally non-reversible RY8.8 model.

We note that the likelihood for our baseline model \mathcal{M}_1 is invariant to the position of the root and so can only distinguish between unrooted trees. However, as the data set contains an outgroup, we can apply the standard approach of outgroup rooting to polarise the relationships on the unrooted trees with the highest posterior support.

For the non-homogeneous Lie Markov models, \mathcal{M}_5 and \mathcal{M}_6 , we use the priors described in Section 4.2. For the remaining models, \mathcal{M}_i , $i = 1, \dots, 4$, we give the prior specification here. For models \mathcal{M}_i , $i = 1, \dots, 4$, the unknowns comprise the tree topology τ , branch lengths $\ell = (\ell_1, \dots, \ell_B)^T \in \mathbb{R}_+^B$, substitution model parameters \mathcal{Q}_i for model i and the shape parameter $\phi \in \mathbb{R}_+$ in the discretised gamma distribution for rate variation across sites. For model \mathcal{M}_1 , τ belongs to the set \mathcal{U}_n of unrooted tree topologies on n species and there are $B = 2n - 3$ branches. For models $\mathcal{M}_2 - \mathcal{M}_4$, τ belongs to the set of rooted tree topologies on n species and there are $B = 2n - 2$ branches.

Our prior takes the form $\pi(\tau, \ell, \phi, \mathcal{Q}_i) = \pi(\tau)\pi(\phi)\pi(\mathcal{Q}_i|\tau) \prod_{b=1}^B \pi(\ell_b)$ for model \mathcal{M}_i , where the dependence of \mathcal{Q}_i on τ is dropped for the homogeneous models $\mathcal{M}_1 - \mathcal{M}_3$. As we did for \mathcal{M}_5 and \mathcal{M}_6 in Section 4.2, we assign priors $\ell_b \sim \text{Exp}(10)$ to the branch lengths, expressing the belief that there will be $E(\ell_b) = 0.1$ substitutions per site on average. Likewise, for the shape parameter we assign $\phi \sim \text{Ga}(10, 10)$ to give a distribution which is modestly concentrated about $E(\phi) = 1$, conveying the belief that the multiplicative random effects across sites will equal 1 on average. For models $\mathcal{M}_2 - \mathcal{M}_4$ we assign the rooted topology a prior according to the Yule model of speciation. For the GTR model \mathcal{M}_1 (see Section 2.1.2) we assign the unrooted topology τ a prior which is uniform over \mathcal{U}_n , expressing prior indifference with respect to the topology.

The substitution model parameters \mathcal{Q}_1 in \mathcal{M}_1 comprise the theoretical stationary distribution $\boldsymbol{\pi} \in \mathcal{S}_4$ and the exchangeability parameters $\boldsymbol{\kappa} = (\kappa_{12}, \kappa_{13}, \kappa_{14}, \kappa_{23}, \kappa_{24})^T \in \mathbb{R}_+^5$ to which we assign the prior

$$\pi(\mathcal{Q}_1) = \pi(\boldsymbol{\pi}) \prod_{i=1}^2 \prod_{j=i+1}^4 \pi(\kappa_{ij}), \quad \text{where } \boldsymbol{\pi} \sim \mathcal{D}_4(1, 1, 1, 1), \quad \kappa_{ij} \sim \text{Ga}(1, 1).$$

In the homogeneous RY5.6b model \mathcal{M}_2 , the substitution model parameters consist of the stochastic vector $\boldsymbol{\rho} \in \mathcal{S}_4$, which controls the theoretical stationary distribution, and the

parameter $\alpha \in \mathbb{R}_+$. We assign the prior

$$\pi(\mathcal{Q}_2) = \pi(\boldsymbol{\rho})\pi(\alpha), \quad \text{where } \boldsymbol{\rho} \sim \mathcal{D}_4(1, 1, 1, 1), \quad \alpha \sim \text{Beta}(1, 1).$$

For the homogeneous RY8.8 model \mathcal{M}_3 , the substitution model parameters simply comprise the stochastic vector $\boldsymbol{\rho} \in \mathcal{S}_8$ and we assign the prior $\boldsymbol{\rho} \sim \mathcal{D}_8(\mathbf{1}_8)$, where $\mathbf{1}_n$ denotes an n -vector of 1s.

For the non-homogeneous GTR model \mathcal{M}_4 , the substitution model parameters \mathcal{Q}_4 comprise the branch-specific composition vectors $\boldsymbol{\pi}_b \in \mathcal{S}_4$ for $b = 1, \dots, B-1$ and the shared set of GTR exchangeability parameters $\boldsymbol{\kappa} = (\kappa_{12}, \kappa_{13}, \kappa_{14}, \kappa_{23}, \kappa_{24})^T \in \mathbb{R}_+^5$. We assign the prior

$$\pi(\mathcal{Q}_4|\tau) = \pi(\boldsymbol{\pi}_1, \dots, \boldsymbol{\pi}_{B-1}|\tau) \prod_{i=1}^2 \prod_{j=i+1}^4 \pi(\kappa_{ij})$$

in which $\kappa_{ij} \sim \text{Ga}(1, 1)$ and the $\boldsymbol{\pi}_b$ are assigned the same joint, conditional distribution as the $\boldsymbol{\rho}_b$ in \mathcal{M}_5 and \mathcal{M}_6 , described in Section 4.2.

The choices of the parameters p_ϱ and v_ϱ in our priors for the branch-specific simplex-valued parameters in models $\mathcal{M}_4 - \mathcal{M}_6$ are:

$$\begin{aligned} \mathcal{M}_4 : \quad & p_\varrho = 0.94, \quad v_\varrho = 0.31^2; \\ \mathcal{M}_5 : \quad & p_\varrho = 0.95, \quad v_\varrho = 2.00^2; \\ \mathcal{M}_6 : \quad & p_\varrho = 0.94, \quad v_\varrho = 0.38^2. \end{aligned}$$

For the non-homogeneous GTR model \mathcal{M}_4 , we use the method described in Heaps *et al.* (2014) to choose hyperparameters. This involves a process of trial-and-improvement, varying p_ϱ and v_ϱ until samples from the ensuing prior predictive distribution of the empirical sequence composition matches quantiles informed by expert biological judgement. We consider samples of the minimum, lower quartile, median, upper quartile and maximum of the GC-content of the *Drosophila* data set. For the non-homogeneous Lie Markov models, \mathcal{M}_5 and \mathcal{M}_6 , we adopt the same iterative procedure. Figure 5.12 shows boxplots for the samples for each model. A biologically credible range for mean GC-content is about 13%–75% (Romiguier & Roux, 2017) and the prior predictive distributions of each model with our chosen hyperparameters seem to be consistent with this range. The prior predictive distribution of \mathcal{M}_5 appears to be less diffuse than the prior predictive distributions of \mathcal{M}_4 and \mathcal{M}_6 . Closer inspection, via a density plot (not shown), reveals some evidence of bimodality. A possible reason for this may be that in \mathcal{M}_5 there is only one parameter, α , determining the degree of imbalance in the rate matrix, whereas the other models have several such parameters. Testing with various values for the hyperparameters during this trial-and-improvement procedure reveals that this feature of the prior predictive distribution remains unchanged.

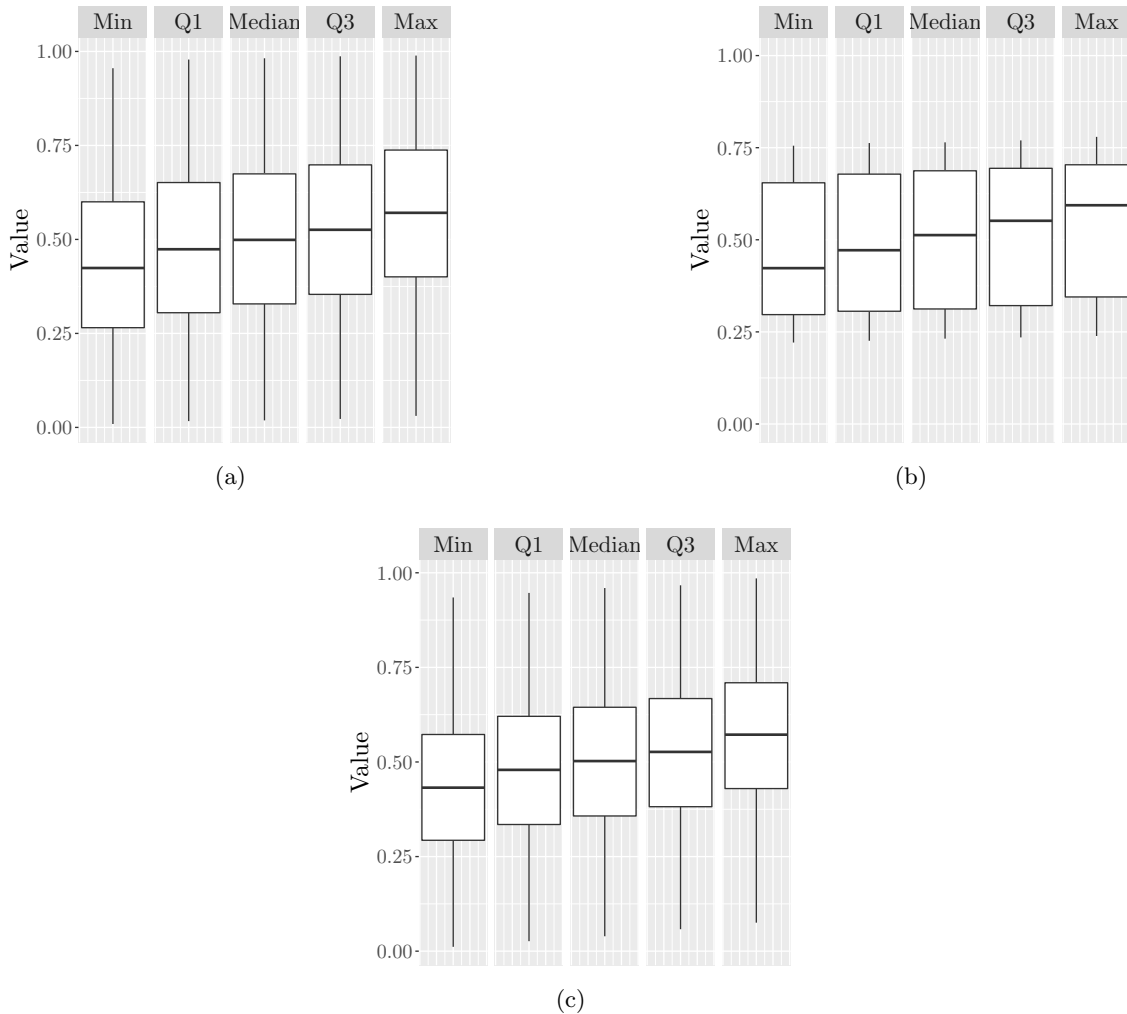


Figure 5.12: Boxplots for quantiles of GC-content sampled from the prior predictive distributions of the taxa based on models (a) \mathcal{M}_4 , (b) \mathcal{M}_5 , (c) \mathcal{M}_6 .

5.2.1 MCMC implementation

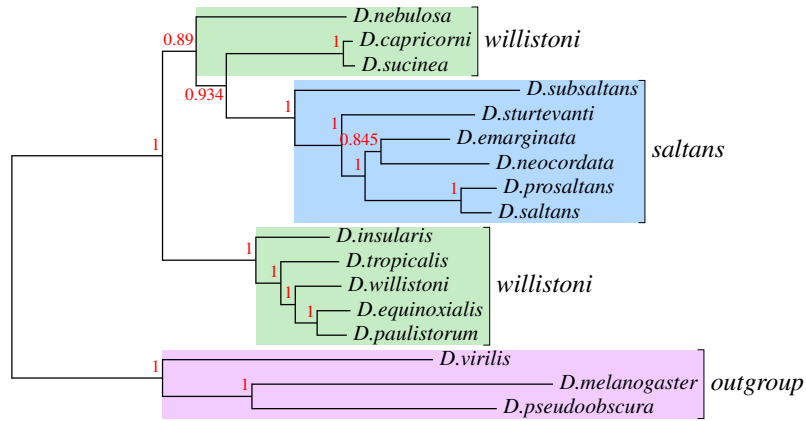
For each model, we run the MCMC algorithm for at least 300K iterations, omitting all but the last 100K as burn-in and thinning the remaining output to retain every 100-th iteration so as to reduce computational overheads. To rigorously assess convergence and mixing we follow the methods utilised by Heaps *et al.* (2014). For each analysis, we run two chains initialised at different starting states. We then consider standard graphical diagnostics, such as trace and density plots, for the quantitative parameters and assess mixing and convergence in tree space using plots of the cumulative relative frequencies of sampled splits (for model \mathcal{M}_1) or clades (for models $\mathcal{M}_2 - \mathcal{M}_6$) over the course of the MCMC run. Examples of these diagnostic plots were given in Section 2.3.2. These checks give no evidence of any lack of convergence and thinning to every 100-th iterate seems to

produce near-uncorrelated posterior samples.

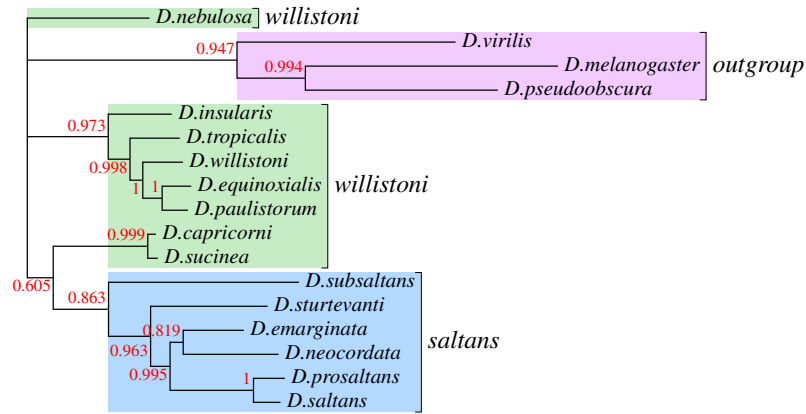
5.2.2 Posterior inference

The majority-rule consensus trees (defined in Section 2.3.2) for the homogeneous and stationary models $\mathcal{M}_1 - \mathcal{M}_3$ are shown in Figures 5.13a – 5.13c, in which numerical labels represent the posterior probability of the associated split (\mathcal{M}_1) or clade ($\mathcal{M}_2, \mathcal{M}_3$). The majority-rule consensus tree obtained using the GTR model \mathcal{M}_1 is unrooted but has been visualised with the root at the midpoint of the branch leading to the outgroup, in accordance with the method of outgroup rooting. We see immediately that the branching structure of the underlying unrooted topology prevents identification of a monophyly (clade) of *willistoni*. Apart from some lack of resolution in the RY5.6b consensus tree, those for the two homogeneous and stationary Lie Markov models, \mathcal{M}_2 and \mathcal{M}_3 , represent the same unrooted topology as the majority-rule consensus tree for the GTR model \mathcal{M}_1 . As such, they also fail to isolate the *willistoni* as a clade. We note, however, that whilst the RY8.8 consensus tree has a biologically implausible branching structure, its root position, within the outgroup, represents a credible evolutionary hypothesis. Indeed under the RY8.8 model \mathcal{M}_3 , the posterior probability for a root position within the outgroup or on its parent branch is 0.999, compared to only 0.343 for the simpler RY5.6b model \mathcal{M}_2 .

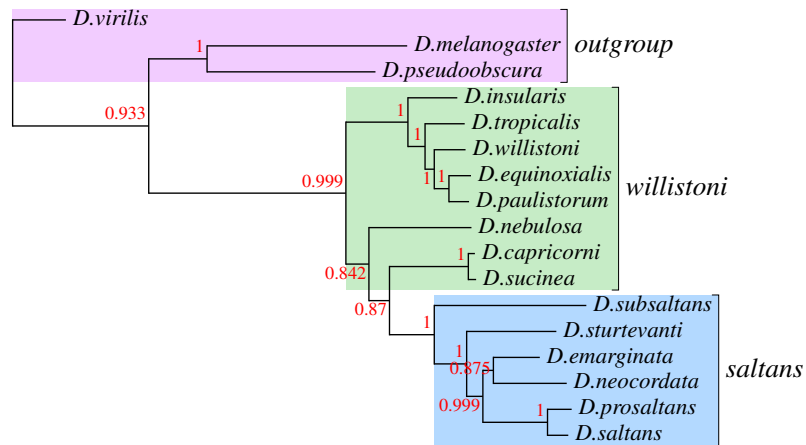
The majority-rule consensus trees for the non-homogeneous and non-stationary models $\mathcal{M}_4 - \mathcal{M}_6$ are shown in Figures 5.14a – 5.14c. All three trees depict the same underlying unrooted topology. This differs from that obtained under the three homogeneous models, and is now biologically plausible, with the *willistoni*, *saltans* and *outgroup* species forming a tripartition, induced by cutting two edges. However, only the non-homogeneous RY8.8 model \mathcal{M}_6 identifies a credible root position, with the root on the consensus tree appearing inside the outgroup, and the marginal posterior probability for a root within the outgroup, or on its parent branch, equal to 1.000. In contrast, the roots on the consensus trees for the non-homogeneous GTR and RY5.6b models, \mathcal{M}_4 and \mathcal{M}_5 , split the *willistoni*, whilst the marginal posterior probability for a root within the outgroup, or on its parent branch, is equal to 0.000 in each case. It is interesting to note that under model \mathcal{M}_6 , the posterior is not only centred on a plausible tree, it is also concentrated in its vicinity, with the posterior for rooted trees notably less diffuse than the distribution obtained under other models. For instance, the rooted topology of the consensus tree, depicted in Figure 5.14c, has posterior probability equal to 0.9235, compared to posterior probabilities of at most 0.6870 for the modes in other cases. This greater concentration of the posterior for the unknowns in \mathcal{M}_6 compared with those in \mathcal{M}_5 is consistent with the results from the analyses of simulated data in Section 5.1.



(a)

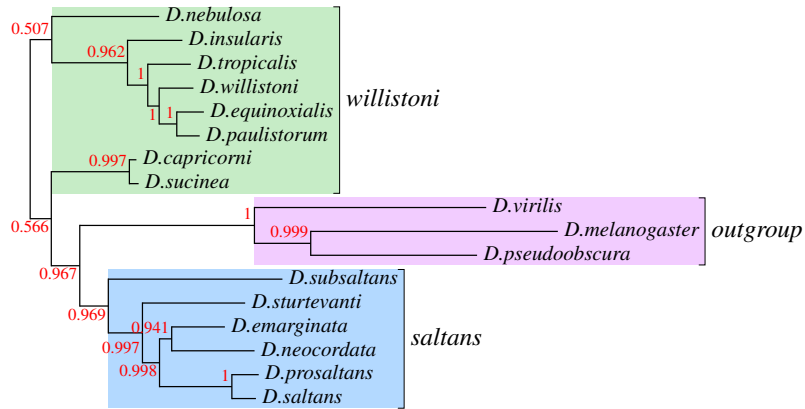


(b)

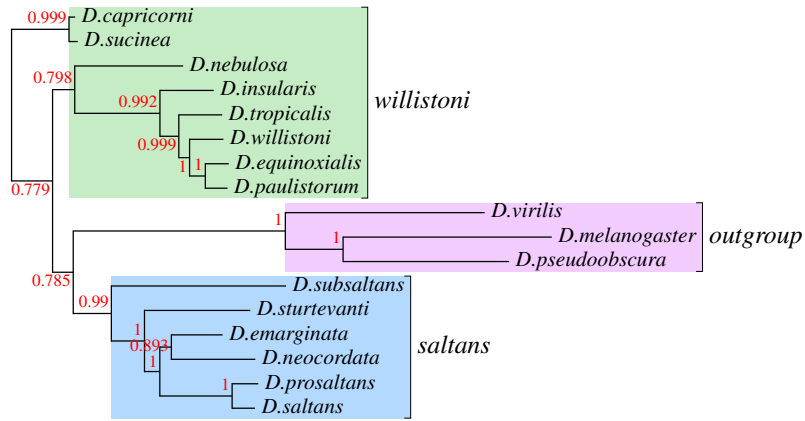


(c)

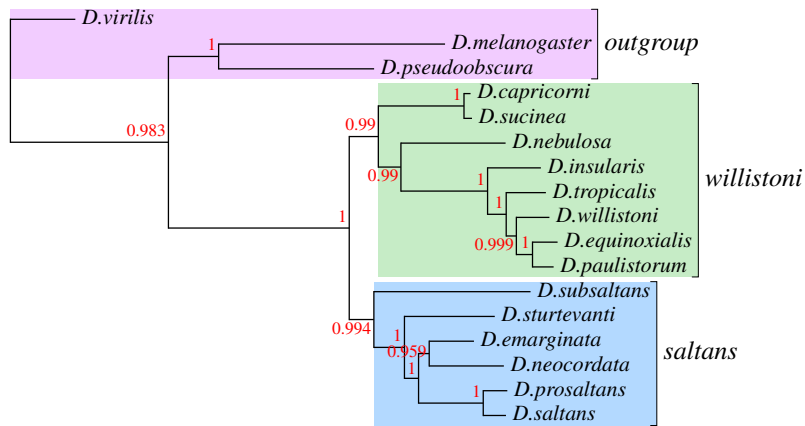
Figure 5.13: Majority rule consensus trees under the models (a) \mathcal{M}_1 – GTR; (b) \mathcal{M}_2 – RY5.6b; (c) \mathcal{M}_3 – RY8.8; Numerical labels represent the posterior probability of the associated split (in (a)) or clade (in (b) and (c)).



(a)



(b)



(c)

Figure 5.14: Majority rule consensus trees under the models (a) \mathcal{M}_4 – non-homogeneous GTR; (b) \mathcal{M}_5 – non-homogeneous RY5.6b; (c) \mathcal{M}_6 – non-homogeneous RY8.8. Numerical labels represent the posterior probability of the associated clade.

5.2.3 Model comparison

As is commonly observed in statistical phylogenetics (Foster, 2004; Heaps *et al.*, 2014; Williams *et al.*, 2015; Cherlin *et al.*, 2017), our phylogenetic inferences are sensitive to the choice of substitution model. One way to arbitrate this inconsistency is through comparisons of the fits of different models; notionally, we have less reason to refute the conclusions of a model which shows a better fit to the data. There are several ways to assess the fit of a model. For example, there are methods based on predictive distributions, like cross-validation and posterior predictive checks (Gelman *et al.*, 2014), where at each iteration of the MCMC algorithm the data are simulated using the accepted parameter values at that particular iteration and compared to the known real data. Another approach is to use information criteria, such as the Bayesian information criterion (BIC) (Schwarz, 1978) or deviance information criterion (DIC) (Spiegelhalter *et al.*, 2002). The BIC and DIC are based on the likelihood, whose maximum increases (or at least remains the same) with every parameter that is added into a model (Oaks *et al.*, 2019). Thus, a penalty is incurred when a parameter is added. The definition of the penalty is where the two information criteria differ.

A more natural measure of model uncertainty in the Bayesian framework is the posterior mass function over models, in this case $\Pr(\mathcal{M}_i|y) \propto p(y|\mathcal{M}_i)\Pr(\mathcal{M}_i)$ for $i = 1, \dots, 6$, which reduces to $\Pr(\mathcal{M}_i|y) \propto p(y|\mathcal{M}_i)$ in the case of equal prior probabilities, $\Pr(\mathcal{M}_i) = 1/6$. For each model \mathcal{M}_i , the crucial component is therefore the marginal, or integrated, likelihood $p(y|\mathcal{M}_i)$, given by

$$p(y|\mathcal{M}_i) = \sum_{\tau} \int_{\Theta_i} p(y|\tau, \Theta_i, \mathcal{M}_i) \pi(\tau, \Theta_i|\mathcal{M}_i) d\Theta_i.$$

Here Θ_i denotes the collection of continuous-valued model parameters, ℓ , α and \mathcal{Q}_i , for model \mathcal{M}_i .

Numerical calculation of the marginal likelihood is a notoriously difficult computational challenge. This is particularly true in phylogenetics due to the discrete nature of tree space; see Oaks *et al.* (2019) for a recent review. Many techniques for approximating the marginal likelihood are based on importance sampling or reciprocal importance sampling. For the *Drosophila* application, these methods would use the identities

$$p(y|\mathcal{M}_i) = \mathbb{E}_q \left\{ \frac{p(y|\tau, \Theta_i, \mathcal{M}_i) \pi(\tau, \Theta_i|\mathcal{M}_i)}{q(\tau, \Theta_i|\mathcal{M}_i)} \right\}, \quad (5.1)$$

or

$$p(y|\mathcal{M}_i) = \left[\mathbb{E}_{\pi} \left\{ \frac{q(\tau, \Theta_i|\mathcal{M}_i)}{p(y|\tau, \Theta_i, \mathcal{M}_i) \pi(\tau, \Theta_i|\mathcal{M}_i)} \right\} \right]^{-1}, \quad (5.2)$$

where \mathbb{E}_q and \mathbb{E}_{π} denote expectation with respect to the importance density $q(\tau, \Theta_i|\mathcal{M}_i)$

and the posterior $\pi(\tau, \Theta_i|y, \mathcal{M}_i)$, respectively. Taking $q(\tau, \Theta_i|\mathcal{M}_i)$ as the prior $\pi(\tau, \Theta_i|\mathcal{M}_i)$ in (5.1) and (5.2) leads to the widely used Monte Carlo and harmonic mean estimators, respectively. The main advantages of these techniques are their computational simplicity, requiring only a sample from the prior in the former case, or posterior in the latter case. However, both are prone to large Monte Carlo error, essentially because the prior and posterior are generally very different, with the former being substantially more diffuse. Motivated by this observation, a number of techniques have been developed that build up the approximation using sequences of intermediate densities that form a bridge between the prior and posterior. Methods that have received considerable attention in the phylogenetic literature are thermodynamic integration (Lartillot & Philippe, 2006), also called the power posterior approach (Friel & Pettitt, 2008), and the stepping-stone method (Xie *et al.*, 2011). The problem with these methods is that they require samples from each of the intermediate densities and so implementation of the algorithms is highly computationally expensive. Indeed, we have found them to be practically infeasible for the more complex non-homogeneous models introduced here.

Newton & Raftery (1994) discussed the use of sampling importance resampling in conjunction with their weighted likelihood bootstrap to obtain simulation consistent samples from a posterior distribution. Further to this, they discussed methods that use such posterior samples to estimate the marginal likelihood and identified the aforementioned issues with the Monte Carlo and harmonic mean estimators. In an effort to stabilise the variance of the harmonic mean estimator, they suggested a hybrid estimator, based on combined samples from the prior and posterior. In practice, it has been found to perform better than the Monte Carlo or harmonic mean estimators (Green, 2003). Starting with the simulation consistent marginal likelihood estimator

$$\frac{\sum_{m=1}^M \pi(\tau^{[m]}, \Theta_i^{[m]}|\mathcal{M}_i)/q(\tau^{[m]}, \Theta_i^{[m]}|\mathcal{M}_i) \times p(y|\tau^{[m]}, \Theta_i^{[m]}, \mathcal{M}_i)}{\sum_{m=1}^M \pi(\tau^{[m]}, \Theta_i^{[m]}|\mathcal{M}_i)/q(\tau^{[m]}, \Theta_i^{[m]}|\mathcal{M}_i)}$$

and taking the importance density to be

$$q(\tau, \Theta_i|\mathcal{M}_i) = \delta\pi(\tau, \Theta_i|\mathcal{M}_i) + (1 - \delta)\pi(\tau, \Theta_i|y, \mathcal{M}_i),$$

\mathcal{M}_1	\mathcal{M}_2	\mathcal{M}_3	\mathcal{M}_4	\mathcal{M}_5	\mathcal{M}_6
-14719.55	-14750.98	-14719.95	-14531.11	-14625.97	-14508.22

Table 5.1: Log marginal likelihoods for each model approximated using the Newton and Raftery hybrid estimator in which the prior weight in the importance density was set at $\delta = 0.05$. The models are: \mathcal{M}_1 – GTR; \mathcal{M}_2 – RY5.6b; \mathcal{M}_3 – RY8.8; \mathcal{M}_4 – non-homogeneous GTR; \mathcal{M}_5 – non-homogeneous RY5.6b; \mathcal{M}_6 – non-homogeneous RY8.8.

with $0 < \delta < 1$ and δ small, leads to the approximation

$$\hat{p}(y|\mathcal{M}_i) = \frac{\sum_{m=1}^M \frac{p(y|\tau^{[m]}, \Theta_i^{[m]}, \mathcal{M}_i)}{\delta \hat{p}(y|\mathcal{M}_i) + (1 - \delta)p(y|\tau^{[m]}, \Theta_i^{[m]}, \mathcal{M}_i)}}{\sum_{m=1}^M \{\delta \hat{p}(y|\mathcal{M}_i) + (1 - \delta)p(y|\tau^{[m]}, \Theta_i^{[m]}, \mathcal{M}_i)\}^{-1}},$$

which can be computed by a standard iterative scheme using draws $(\tau^{[m]}, \Theta_i^{[m]})$, $m = 1, \dots, M$, from $q(\tau, \Theta_i|\mathcal{M}_i)$. The scheme requires samples from both the prior and the posterior. For each model, we already have 1000 samples from the posterior generated from our MCMC analyses and can augment them with $\delta \times 1000/(1 - \delta)$ samples from the prior. We initialise the scheme with $\hat{p}(y|\mathcal{M}_i)^{(0)}$ estimated using the harmonic mean estimator given in (5.2) and run the scheme for 1000 iterations to ensure convergence.

For models $\mathcal{M}_1 - \mathcal{M}_6$ in the *Drosophila* application, the log marginal likelihoods obtained using the Newton and Raftery hybrid estimator (with $\delta = 0.05$) are displayed in Table 5.1. The superior model fit afforded by the three non-homogeneous and non-stationary models is immediately apparent. For instance, if we perform pairwise comparisons between each homogeneous model and its non-homogeneous counterpart, then the log Bayes factor ranges from 125.01 to 211.73 in favour of the non-homogeneous model. Of the three non-homogeneous models, the non-stationary RY8.8 model seems to give the best fit to the data. This is also the only model whose posterior supports a biologically credible rooted tree. Reasons for its superiority over the simpler RY5.6b variant were discussed from a theoretical and practical perspective in Sections 3.2.2 and 5.1.1, respectively, whilst the improvement over the non-homogeneous GTR model may be attributable to the additional source of root information gained through the non-reversible structure of the RY8.8 rate matrix.

Although the Newton and Raftery hybrid estimator tends to be more stable than the harmonic mean estimator (Newton & Raftery, 1994), its variance can still be large. For each of the models $\mathcal{M}_1 - \mathcal{M}_6$, we therefore run ten MCMC chains, initialised at different starting points, and repeat the calculation of the marginal likelihood. The results are summarised in Figure 5.15, where the box-plots for the homogeneous and non-homogeneous

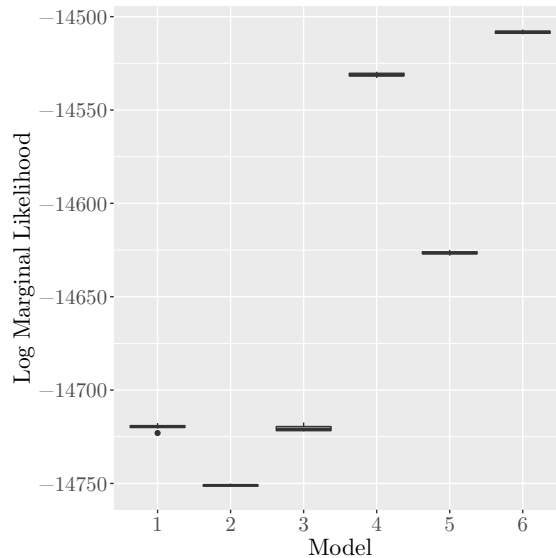


Figure 5.15: Box-plots of the Newton and Raftery hybrid estimates of the marginal likelihood for each model \mathcal{M}_i , $i = 1, \dots, 6$, based on the output of 10 MCMC chains which were initialised at different starting points.

models are well separated and similarly for the three non-homogeneous models $\mathcal{M}_4 - \mathcal{M}_6$. Even after allowing for Monte Carlo error, therefore, these results do not give any cause to question the ranking of models, and the subsequent conclusions drawn here.

In addition to model fit, we compare the computational times of the six models. Taking as a baseline, the computational time to run a fixed number of MCMC iterations for model \mathcal{M}_1 , the computational times for \mathcal{M}_2 and \mathcal{M}_3 are broadly consistent with this time across a range of alignment sizes. Furthermore, the three non-homogeneous models, \mathcal{M}_4 , \mathcal{M}_5 and \mathcal{M}_6 , have computational times that are broadly consistent with each other and are around double the times of their homogeneous counterparts. This shows evidence to suggest that the additional parameters introduced with the non-homogeneous models do not have a detrimental effect on computational time, even in the case of \mathcal{M}_6 which has the most parameters. The doubling of computational time for \mathcal{M}_6 (compared to the homogeneous, stationary and reversible model \mathcal{M}_1) seems like a reasonable sacrifice to make in order to achieve a better model fit and credible inferences.

5.3 Results summary

In this chapter we have presented the results of analyses of simulated data and biological data (the *Drosophila* application). From the results of the simulation study, we have shown that the non-homogeneous RY8.8 model is able to identify the root position from the likelihood and increasing the number of sites in the alignments tends to lead to more

accurate and precise inference of all unknowns. This was true for the non-homogeneous RY5.6b model when the number of taxa was small. However, for larger trees, root inference was generally poor. Root inference under the non-homogeneous RY8.8 model remained strong even with the presence of prior-data conflict arising from an unbalanced rooted topology. However, inference can be sensitive to the presence of long branches in the unrooted topology. Interesting future work could investigate how to incorporate more biological insight into the prior for the branch lengths ℓ (see Section 9.3.1 in Chapter 9). Inference of the unrooted topology for both models was strong, with the posterior mode being the true unrooted topology in every analysis of our simulation study.

After applying our models to the *Drosophila* data set, where there is evidence of compositional heterogeneity, we found that the non-homogeneous RY8.8 model was able to identify a plausible rooted tree. This was not the case for the non-homogeneous RY5.6b model, most likely due to it having fewer parameters than the non-homogeneous RY8.8 model and its rate matrix having an additive structure. This property may explain why we found (branch-specific) parameter identifiability to be weaker for the non-homogeneous RY5.6b model. Finally, we compared our models with their homogeneous counterparts, the GTR model and the non-homogeneous GTR model. We found evidence to suggest, via marginal likelihoods, that our non-homogeneous RY8.8 model has a better fit to the *Drosophila* data set for a relatively low computational cost.

Part II

Metagenomics

Chapter 6

Exploratory data analysis

Now we focus our attention on the time series metagenomics part of the thesis. As discussed in Section 2.2.4 of Chapter 2, our goal is to model five years of weekly counts of operational taxonomic units (OTUs) from a wastewater treatment plant (WWTP) and their interactions with each other, whilst also incorporating chemical and environmental data. Before a model can be developed, we need to have an in-depth understanding of the data, which is the purpose of this chapter. In Section 2.2.3 of Chapter 2, we gave a brief overview of the metagenomics data, but here we look at the data in more detail. We discuss data reformatting and we present our exploratory analysis of the data sets. The findings of the exploratory data analysis will be used to inform some of our model choices in the next chapter.

6.1 Data description

6.1.1 Reformatting the data

The raw data consist of three tables: an OTU table, a chemical and environmental data set and a taxonomy table, detailed in the sections which follow.

OTU tables

The OTU table comprises weekly counts of OTUs from 1st June 2011 to 1st June 2015 in each of two tanks, the activated sludge (AS) tank and settled sewage (SS) tank of a UK-based WWTP (see Section 2.2.3). We split this table into two OTU tables, one for each tank. There are 9044 different OTUs sampled from the AS and 9678 different OTUs sampled from the SS. Note that some OTUs are present in both tanks, whereas some only appear in one of the tanks. Figure 6.1 shows time series plots of some OTUs only present in the AS. We show examples of OTUs that have mainly zero-counts (Figure 6.1a), which are most likely OTUs that enter the system randomly and then die out quickly. However,

as discussed in Section 2.2.4, the zeroes could be attributed to insufficient sampling depths (see also Section 7.1). Additionally, we show examples of OTUs that seem to be in the system for more than one or two time points (Figure 6.1b). For the OTUs that only appear in the OTU table of the SS, time series profiles like those mentioned above can also be seen but are not shown.

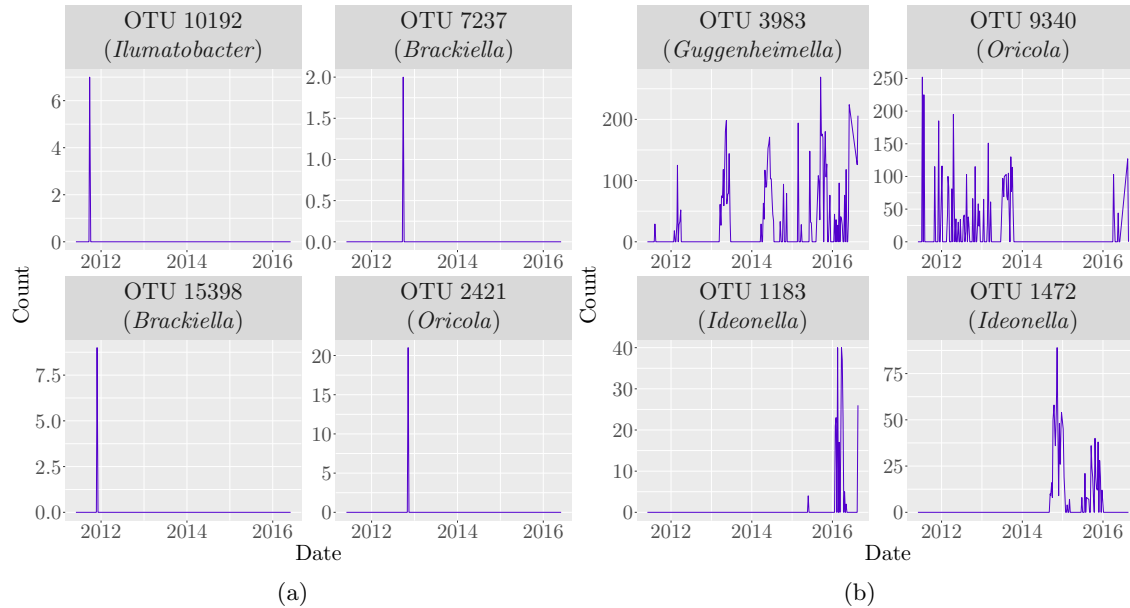


Figure 6.1: Time series plots of some OTUs (with their genus in brackets) that only appear in the OTU table of the AS and (a) mainly have zero-counts, or (b) appear at more than one or two time points. Note that the y -axis scales are different for each plot.

There are 3323 OTUs present in both tanks. Figure 6.2 shows time series plots for the counts of four different OTUs in each tank. OTU 10 from the genus *Romboutsia* is an example of an OTU that is present in both tanks throughout, although it seems to be slightly more abundant in the SS. OTU 103 from the genus *Hirschia* seems to be present most of the time in the AS but mostly absent from the SS, apart from appearances every now and then, with a particularly high count in April 2016. OTU 174 from an unknown genus is mostly absent from the AS but is present most of the time in the SS. Finally, OTU 7001 from the genus *Leadbetterella* is an example of an OTU that is mostly absent from both tanks.

Chemical and environmental data

As we did with the OTU table, we split the chemical and environmental data set into two tables, one for each tank. For the AS tank, we have 14 chemical and environmental covariates measured at 257 time points and for the SS tank we have 29 covariates measured

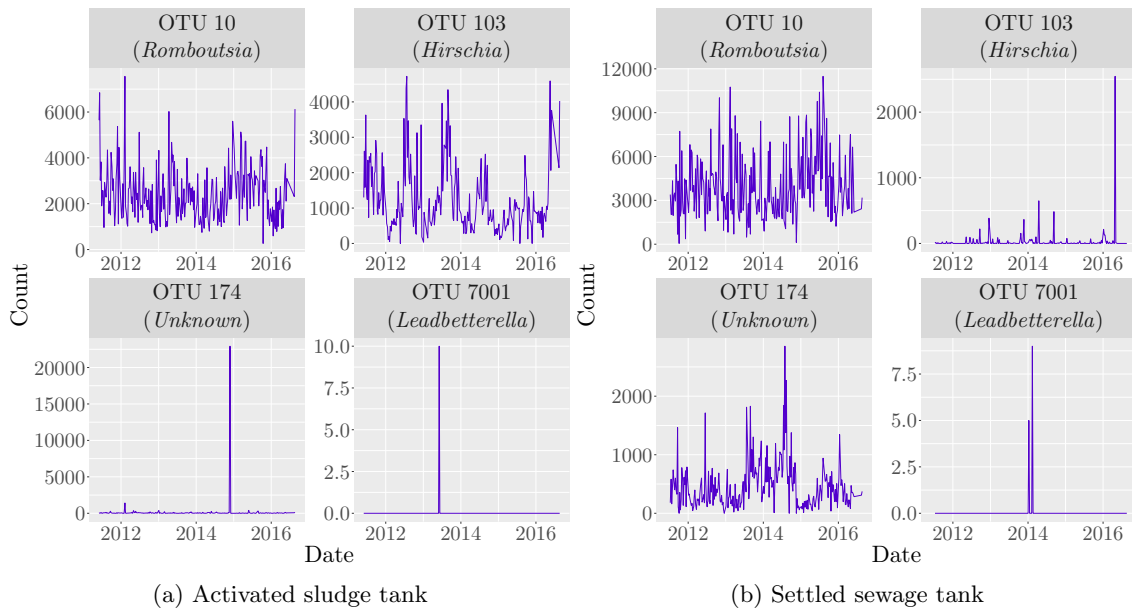


Figure 6.2: Time series plots of four OTUs (with their genus in brackets) that appear in both tanks.

at 257 time points. The covariates for each tank are listed in Tables 6.1 and 6.2. Chemical oxygen demand (COD) is an indicative measure of the amount of oxygen that can be consumed by reactions in the tank. It is used to quantify the oxidisable pollutants in water. The mixed liquor suspended solids (MLSS) represent the concentration of suspended solids in the tank. These suspended solids are mainly microorganisms and non-biodegradable suspended matter. The mixed liquor volatile suspended solids (MLVSS) make up the proportion of the MLSS that consume the food sources in the tank. These will be mostly microorganisms and organic matter, implying that the MLVSS are approximately equal to the amount of microorganisms in the water and can be used to indicate if there are enough microorganisms present for water purification. Dissolved oxygen (DO) is a measure of dissolved oxygen in the wastewater available to sustain life, i.e., the microorganisms, in the tank. Total nitrogen is the sum of all sources of nitrogen in the tank, so the sum of the nitrate-nitrogen ($\text{NO}_3\text{-N}$), nitrite-nitrogen ($\text{NO}_2\text{-N}$), ammonia-nitrogen ($\text{NH}_3\text{-N}$) and organically bonded nitrogen. The flow is the rate at which water enters the SS tank.

Taxonomy table

The taxonomy table contains the kingdom, phylum, class, order, family and genus for every OTU present in the two tanks and does not require reformatting.

Covariate	Unit
Ammonia, Chloride, COD, DO, Fluoride, MLSS, MLVSS, Nitrate, Nitrite, Phosphate, Sulphate	mg/L
pH	-
Sludge Age	days
Temperature	Celsius

Table 6.1: Chemical and environmental covariates in the AS.

Covariate	Unit
Aluminium, Ammonia, Arsenic, Cadmium, Calcium, Chloride, Chromium, COD, Copper, Fluoride, Iron, Lead, Magnesium, Manganese, MLSS, MLVSS, Nickel, Nitrate, Nitrite, Phosphorus, Phosphate, Potassium, Silicon, Sulphate, Sulphur, Total nitrogen, Zinc	mg/L
pH	-
Flow	L/s

Table 6.2: Chemical and environmental covariates in the SS.

6.1.2 Missing data

Now we discuss the missing data in each data set.

OTU tables

In the OTU tables, there are 257 counts per OTU in the AS tank and 251 counts per OTU in the SS tank. The first six weeks in the SS tank are omitted and so could be regarded as missing. However, we do not fit a joint model to the data for the two tanks and instead simply acknowledge that we have a shorter time series for the SS tank. Each year, measurements for one week are missing during the Christmas period. We treat the data from these weeks as missing at random, that is, the distribution of the missing data mechanism does not depend on the missing values (Rubin, 1976; Gelman *et al.*, 2014; Little & Rubin, 2019). In other words, the probability of a measurement being missing does not depend upon its value. These missing data are accounted for in our analysis; see Section 7.5.2.

Chemical and environmental data

Tables 6.3 and 6.4 show the proportions of missing data for each covariate in each tank. We discard any covariates where more than 50% of the data are missing, as allowing for uncertainty in their values would substantially complicate computational inference during model fitting whilst potentially yielding only very imprecise inferences about their influence on OTU abundance. This means that we disregard the sludge age for the AS tank, whereas

for the SS tank, we disregard pH, phosphorous and total nitrogen. We account for the small amount of missing data in the remaining variables later in Section 7.5.2, where we describe how the missing chemical and environmental data are handled. Note that, as with the OTU data, there also is a week missing from each year over the Christmas period for the chemical and environmental data and this is also taken into consideration.

Covariate	Proportion	Covariate	Proportion
COD	0.0661	Ammonia	0.0661
MLSS	0.0000	MLVSS	0.0000
Nitrate	0.0039	Nitrite	0.0039
Sulphate	0.0039	Phosphate	0.0039
Fluoride	0.0039	Chloride	0.0039
pH	0.0039	Temperature	0.0039
DO	0.0039	Sludge Age	0.5058

Table 6.3: Proportions of missing data (4 d.p.) in chemical/environmental data (AS).

Covariate	Proportion	Covariate	Proportion
COD	0.0661	Ammonia	0.0661
MLSS	0.0000	MLVSS	0.0000
Nitrate	0.0039	Nitrite	0.0039
Sulphate	0.0039	Phosphate	0.0039
Fluoride	0.0039	Chloride	0.0039
pH	0.8482	Flow	0.0039
Phosphorus	0.7977	Total Nitrogen	0.7977
Calcium	0.0506	Magnesium	0.0506
Potassium	0.0506	Iron	0.0545
Manganese	0.0506	Aluminium	0.0506
Zinc	0.0506	Lead	0.0506
Copper	0.0506	Arsenic	0.0506
Cadmium	0.0506	Sulphur	0.0506
Silicon	0.0506	Nickel	0.0506
Chromium	0.0506		

Table 6.4: Proportions of missing data (4 d.p.) in chemical/environmental data (SS).

There were a few measurements which were orders of magnitude larger than the rest. After discussion with a biological expert, these were deemed biologically implausible, removed and treated as missing. A table detailing the affected covariates is shown in Table D.1 in Appendix D.

Taxonomy table

Table 6.5 shows the proportions of missing data for each taxonomic rank in each tank and for both tanks considered together. We can see that there is a larger proportion of missing data in the AS tank than in the SS tank. The finer the taxonomic rank is, the more missing data there are. Missing values here could be due to errors, but it is more likely due to the known sequence database used to classify the OTUs being incomplete. The taxonomy assignment used to give the taxonomy table presented to us is based on a set of known sequences from 16S genes. The Ribosomal Database Project (RDP) classifier (Wang *et al.*, 2007) was used, which is a naïve Bayesian classifier for taxonomic assignments from domain to genus. The most recent release of the RDP classifier was in September 2016, three months after the last measurement was recorded at the sewage treatment plant. This release has 3356809 16S rRNA gene sequences and 125525 Fungal 28S rRNA gene sequences (Ribosomal Database Project, 2016). An OTU in our data set was classified as NA if its sequence did not match any of these (roughly) 3.5 million sequences. This problem of missing taxonomic information becomes more pertinent when we discuss clustering methods later in Section 7.1.1.

Taxonomic Rank	Proportion		
	AS	SS	Both
Kingdom	0.0108	0.0033	0.0072
Phylum	0.0532	0.0282	0.0440
Class	0.1202	0.0741	0.1042
Order	0.2021	0.1342	0.1800
Family	0.3014	0.2364	0.2856
Genus	0.5409	0.4656	0.5257

Table 6.5: Proportions of missing data (4 d.p.) for each taxonomic rank in the taxonomy table for each tank and overall (after OTUs with zero-counts for all time points are removed).

6.2 Exploratory data analysis

In this section we discuss our exploratory data analysis of the OTU tables and the taxonomy table. We then look at all the data together, with a particular focus on finding possible relationships between some of the chemical and environmental covariates and relative abundances. As mentioned at the start of this chapter, our findings here help to guide our decision making when developing our model in Chapter 7.

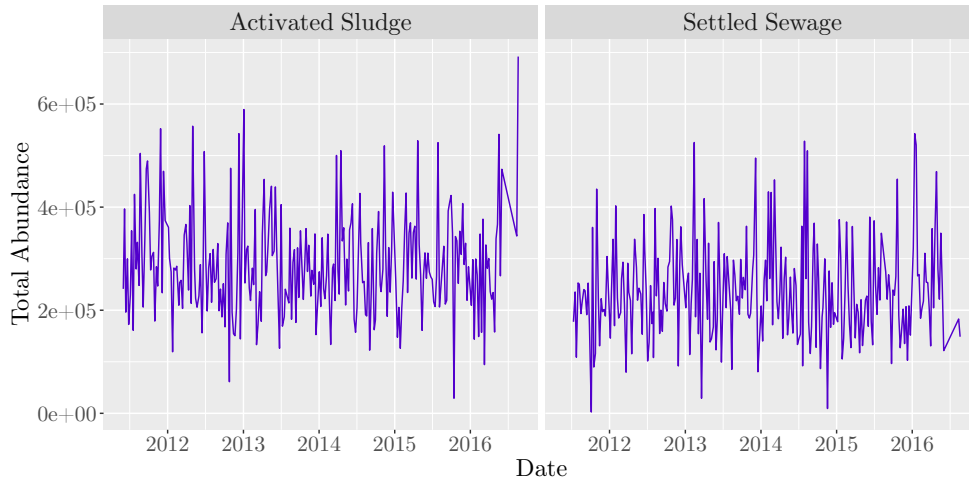


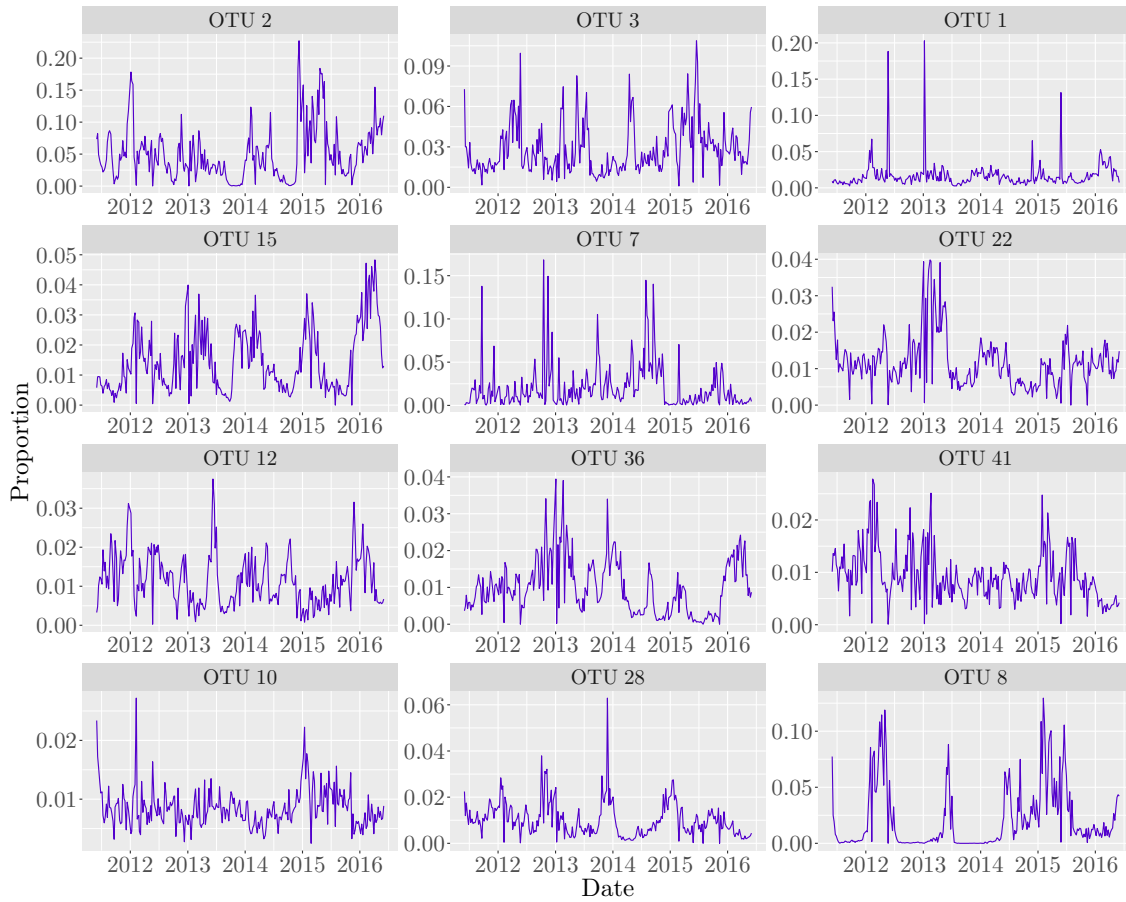
Figure 6.3: Time series plots of total abundances for each OTU table.

6.2.1 OTU tables

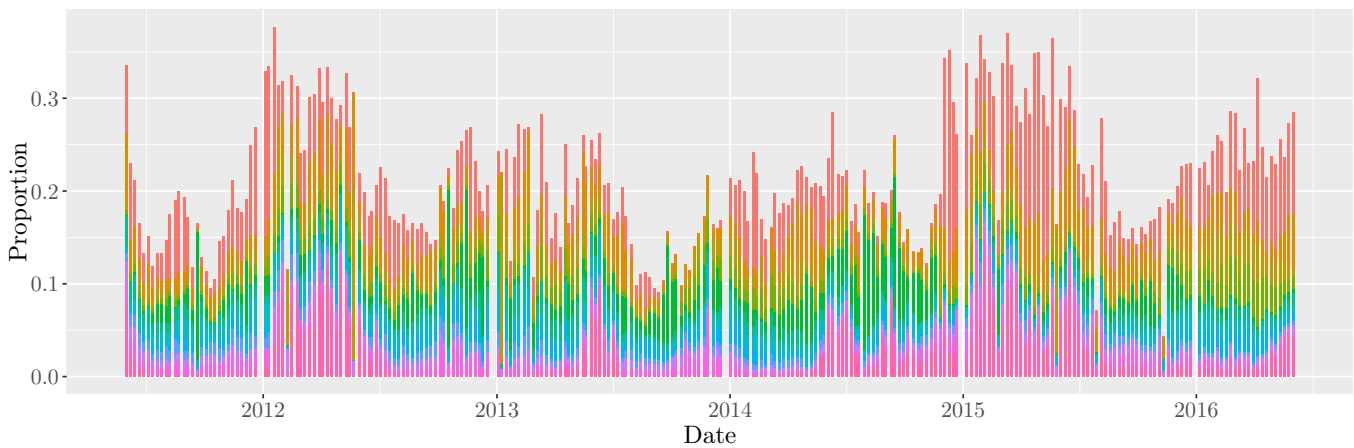
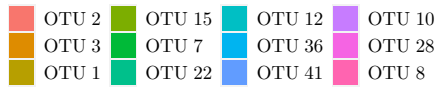
Figure 6.3 shows the time series plots of the total abundances for each OTU table. There is an absence of trend in both plots, except for a possible upwards trend in the last few time points of the AS tank. The total abundance in the AS tank seems to be higher than the SS tank on average and there is more variability in the AS tank too. However, as discussed in Section 2.2.4 of Chapter 2, we note that total abundance at a particular time point here is not the *actual* total abundance of OTUs. It is the total number of OTUs (sequences) detected in a sample at a particular time point.

Now we look at the 12 most abundant OTUs in each tank based on median abundance. Figure 6.4a shows time series plots of the relative abundances for the 12 most abundant OTUs in the AS tank. OTU 15, OTU 28 and OTU 8 demonstrate clear seasonality with peaks appearing roughly once a year. There does not appear to be a trend over time for the top 12 OTUs, with the *possible* exception of OTU 15, whose relative abundance may be increasing over time. OTU 1 has a proportion of around 0.02 most of the time but has the occasional spike. For example, its relative abundance is about 0.19 in May 2012 and over 0.2 at the start of 2013. This OTU also seems to display seasonality, which is clearer on the log-scale (not shown). Figure 6.4b shows a stacked bar plot for the top 12 OTUs in the AS tank over time. When we aggregate the 12 most abundant OTUs we can clearly see seasonality with annual peaks. We also observe that these top 12 OTUs represent about 20% of the relative abundance on average. Additionally, this plot can help to identify when different OTUs tend to dominate. For example, at the time points when OTU 1 spikes, the other OTUs tend to have smaller proportions.

Figure 6.5a and Figure 6.5b respectively show the time series plots and the stacked

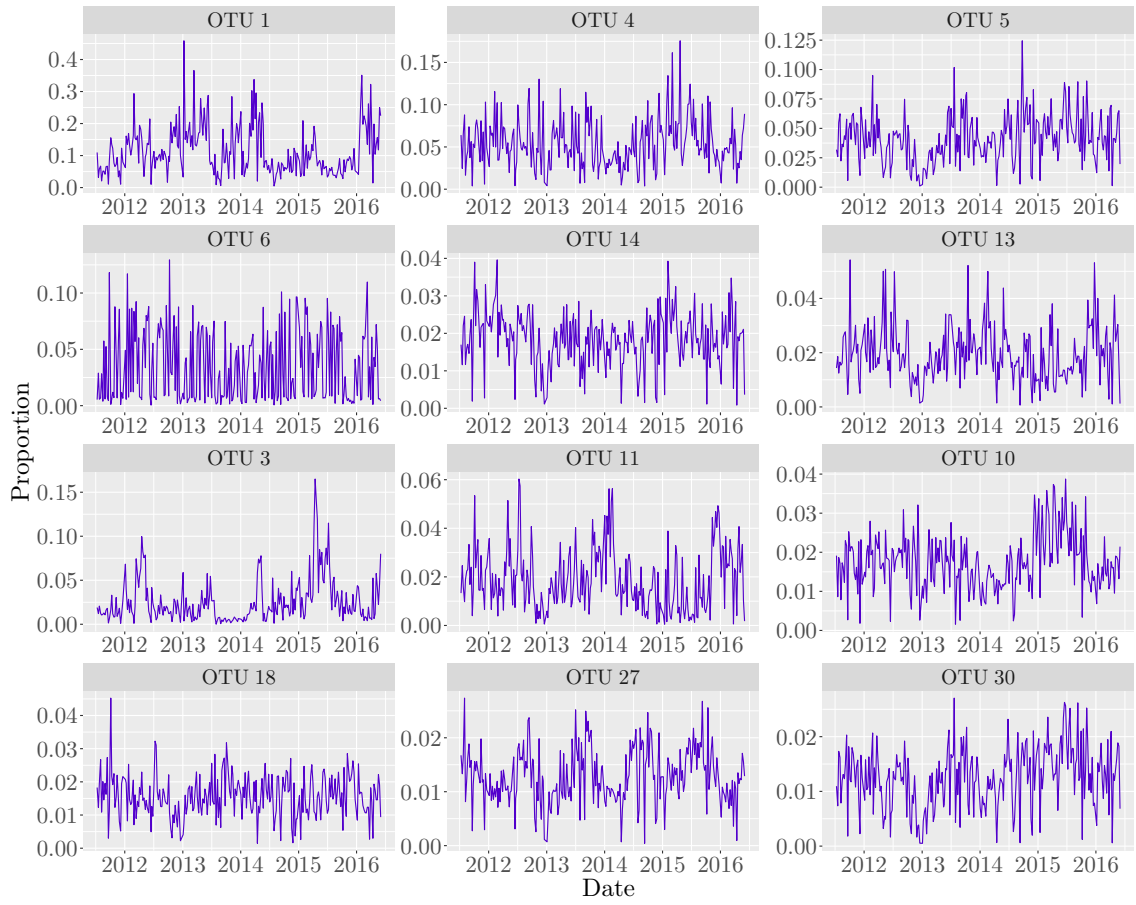


(a)

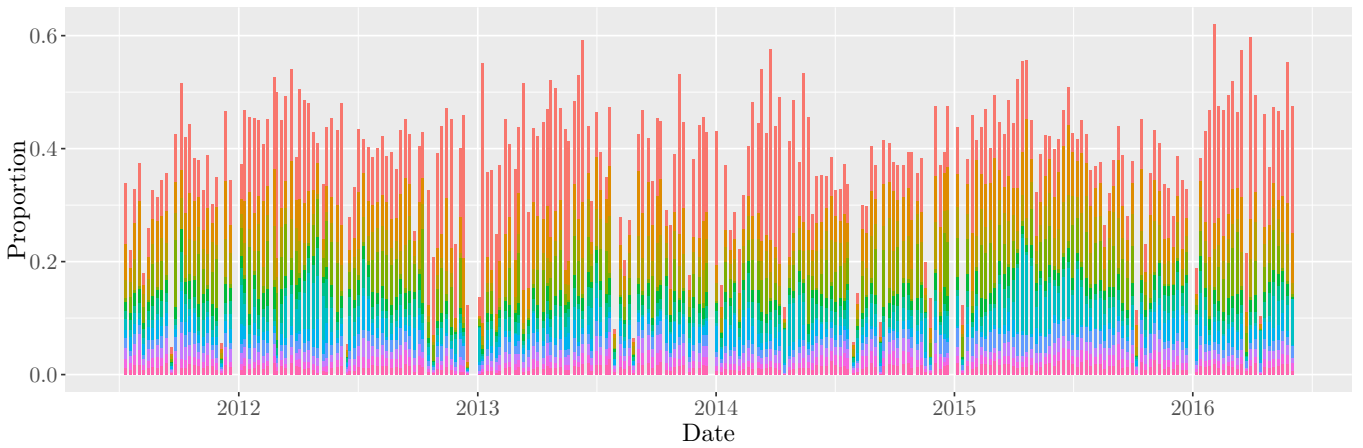


(b)

Figure 6.4: (a) time series plots and (b) stacked bar plot for the top 12 OTUs in the AS based on median abundance



(a)



(b)

Figure 6.5: (a) time series plots and (b) stacked bar plot for the top 12 OTUs in the SS based on median abundance.

bar plot for the 12 most abundant OTUs in the SS tank. Here, seasonality is not as clear as we saw in the AS tank. The time series plots do not show any evidence of trends and do not have any obvious annual peaks, with the exception of OTU 27 and possibly OTU 1, which do seem to peak roughly once per year. There is more variability in the time series plots here compared to in the AS, suggesting that perhaps the OTU population sizes are changing too quickly in the SS to be captured by weekly measurements. OTU 1, OTU 3 and OTU 10 are the only OTUs to be in the 12 most abundant OTUs for both tanks. Since the two tanks perform different functions in the WWTP, it is likely that the conditions in each tank are different, which could suggest that OTUs that thrive in one tank will not thrive in the other. In fact, only 18 of the top 100 OTUs for each tank appear in both tanks. The stacked bar plot for the SS tank shows that the OTUs' proportions stay roughly the same over time without any OTU obviously dominating at any particular time. The top 12 OTUs in the SS tank have a larger relative abundance (~ 0.3) on average than the top 12 OTUs in the AS tank.

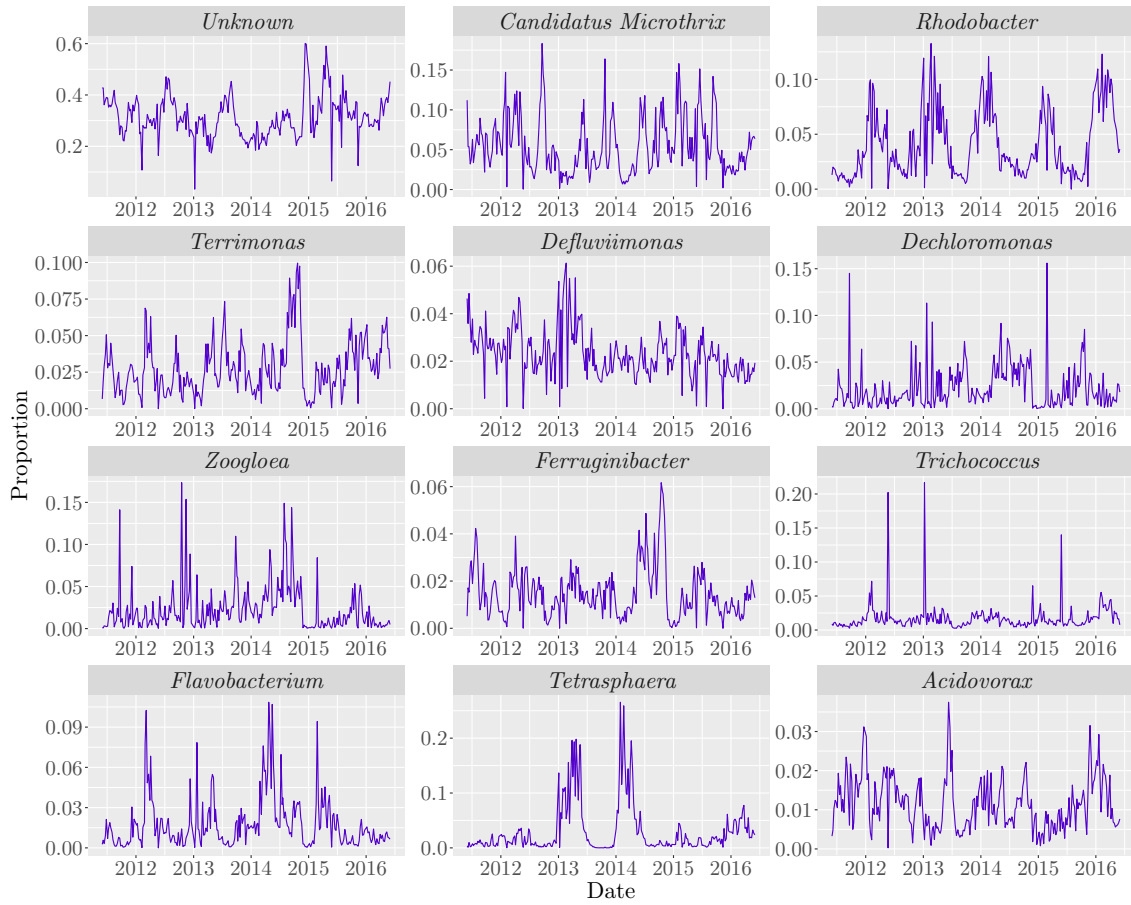
6.2.2 Taxonomy table

Taxonomic Rank	AS	SS
Kingdom	3	3
Phylum	42	43
Class	84	89
Order	132	139
Family	272	276
Genus	713	831

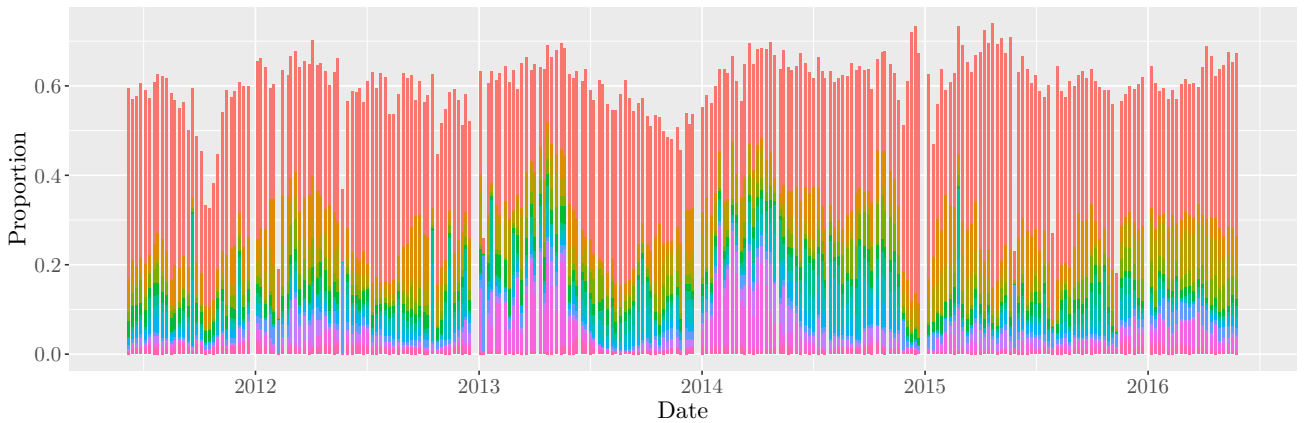
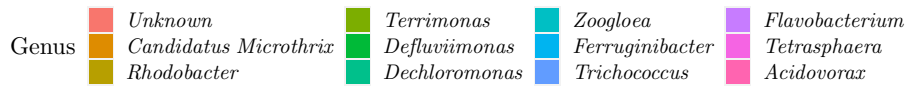
Table 6.6: Number of different types for each taxonomic rank in each tank, excluding NAs.

Table 6.6 shows the unique number of kingdoms, phyla, classes, orders, families and genera in each tank. The OTUs in the SS tank appear to be more diverse, although it is worth remembering that there are more OTUs in this tank. With the exception of kingdom, there are too many in each taxonomic rank to look at in great detail, so instead we just look at the 12 most abundant genera and 12 most abundant classes based on median abundance. For completeness, plots for the 12 most abundant families, orders and phyla are shown in Figures E.1a to E.6b in Appendix E but are not discussed.

Figure 6.6a shows the time series plots and Figure 6.6b shows the stacked bar plot for the top 12 genera in the AS based on median abundance. The corresponding plots for the top 12 genera in the SS are shown in Figure 6.7a and Figure 6.7b respectively. Note that in both tanks, missing genera, which are grouped together in a single “Unknown” group per tank, represent a large proportion of the total abundance. McIlroy *et al.* (2015) surveyed 20 full-scale sewage treatment plants in Denmark and found that the top five genera in

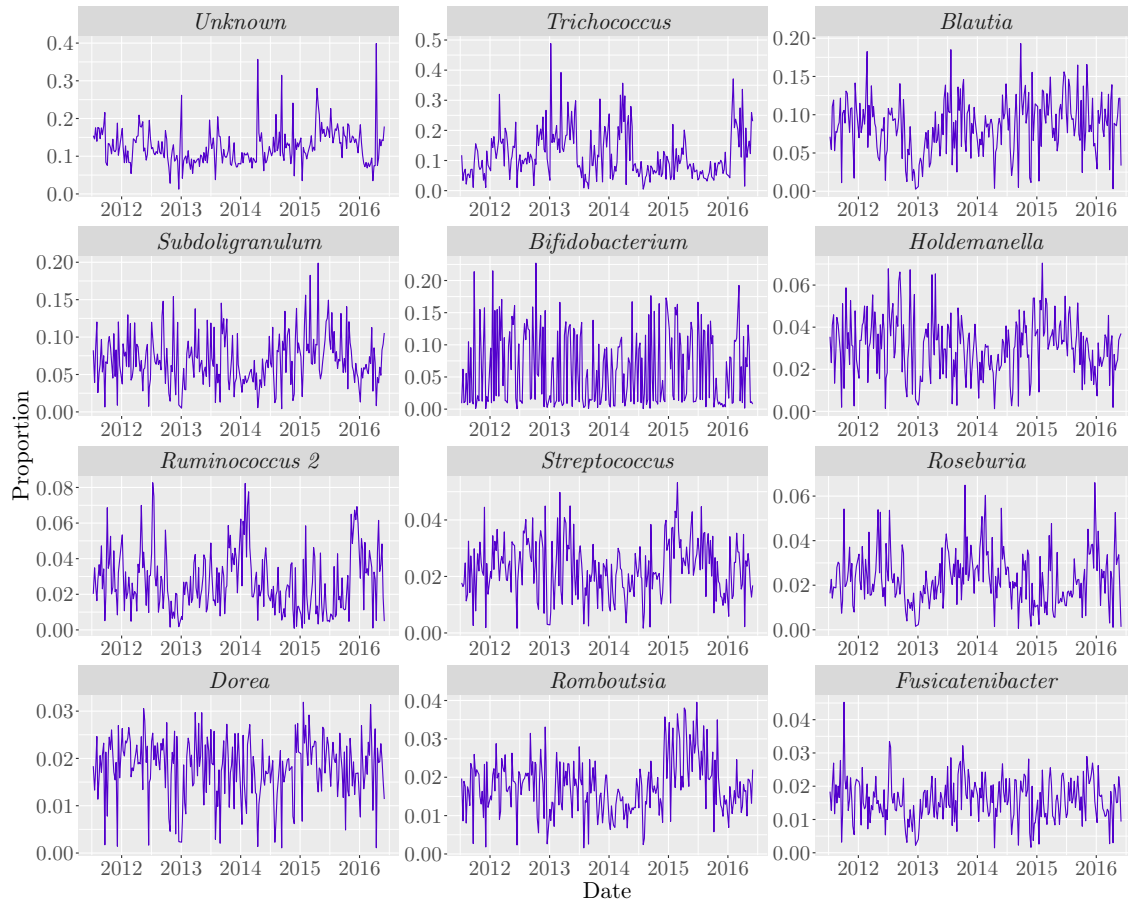


(a)

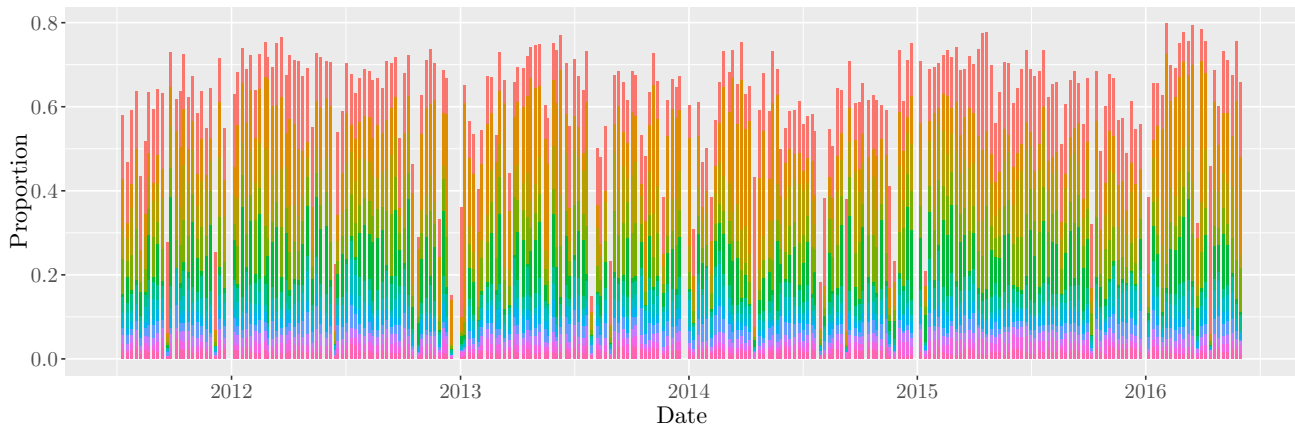


(b)

Figure 6.6: (a) time series plots and (b) stacked bar plot for the top 12 genera based on median abundance in the AS tank.



(a)



(b)

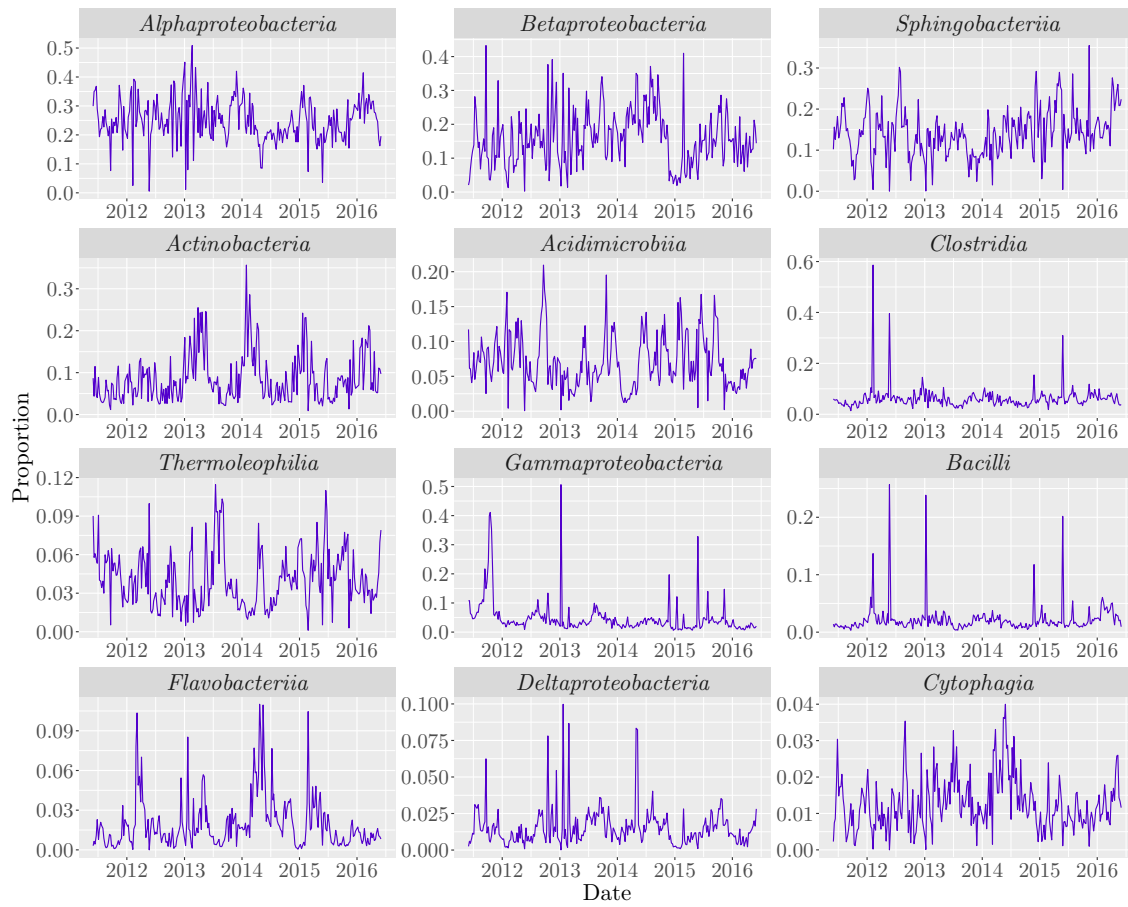
Figure 6.7: (a) time series plots and (b) stacked bar plot for the top 12 genera in the SS based on median abundance.

the AS, based on median abundance, were *Tetrasphaera*, *Trichococcus*, *Ca. Microthrix*, *Rhodofera*x and *Rhodobacter*. All of these, except for *Rhodofera*x, appear in our top 12 genera. In addition to this, *Dechloromonas*, *Flavobacterium* and *Ferruginibacter* appear in the top 50 genera found by McIlroy *et al.* (2015). There are many possibilities as to why some of our top genera do not appear, for example, the difference in location, as our sewage treatment plant is located in the UK. Regarding *Rhodofera*x, this genus is present in the AS tank but is not one of the 12 most abundant. Additionally, the genus could have a larger abundance as it could also be represented by OTUs with an unknown genus. Furthermore, our data are from just one plant with weekly measurements, whereas their data consisted of 574 samples from several sewage treatment plants with recordings taken up to four times a year for eight years.

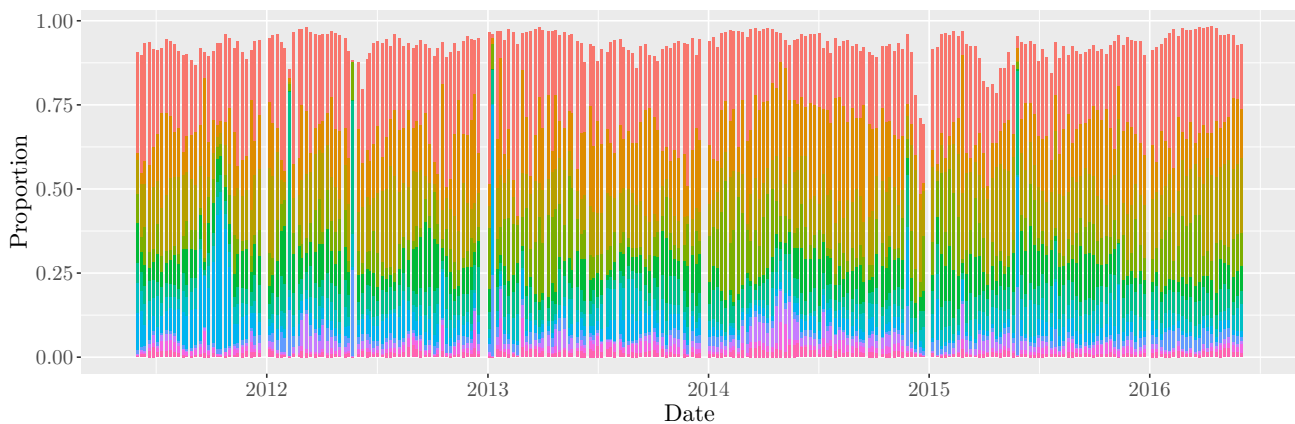
In the AS, *Rhodobacter* clearly shows seasonality with annual peaks in late February/early March. *Flavobacterium*, *Ferruginibacter* and *Trichococcus* also seem to display seasonality, although for the latter the seasonality is clearer on the log-scale (not shown). Turning our attention to the stacked bar plot for the AS tank in Figure 6.6b, we can see that the 12 genera account for about 60% of the total abundance on average at each time point. However, there are times when their relative abundance is particularly small, for example, in early 2012, when it is under 0.2. The stacked bar plot also hints at the presence of seasonality, although it is less clear, as the genera do not all peak in the same weeks. This plot also suggests that when *Trichococcus* has a particularly large relative abundance, for example, at the start of 2013, where it accounts for more than 20% of the total abundance, the remaining genera have low proportions. This suggests that *Trichococcus* (and some of the remaining 701 genera) might thrive where the other top 11 genera do not. Comparing Figure 6.4a and Figure 6.6a we can deduce that the abundance of OTU 1 represents most of the overall abundance for the genus *Trichococcus* in the AS tank.

Trichococcus is the only genus that appears in the top 12 genera for both tanks. From Figure 6.7a we see that its relative abundance in the SS tank is a lot higher than in the AS tank, with it being the most abundant genus in the SS tank (with the exception of unknown genera). As we saw with the AS tank, the abundance of OTU 1 represents most of the overall abundance for *Trichococcus*, thus explaining why it also displays hints of seasonality. Otherwise, seasonality in the SS is less evident than in the AS. There seem to be more fluctuations in the time series plots of the top 12 genera in the SS, suggesting that the population sizes are changing at a quicker rate than in the AS. Based on Figure 6.7b, it seems that the top 12 genera in the SS roughly represent $\sim 60\%$ of the total abundance.

Figure 6.8a shows the time series plots for the 12 most abundant classes in the AS based on median abundance. There are hints of seasonality in some of the classes, even at this fairly coarse taxonomic rank, for example, *Flavobacteriia*, *Actinobacteria* and *Deltapro-*



(a)



(b)

Figure 6.8: (a) time series plots and (b) stacked bar plot for the top 12 classes in the AS based on median abundance.

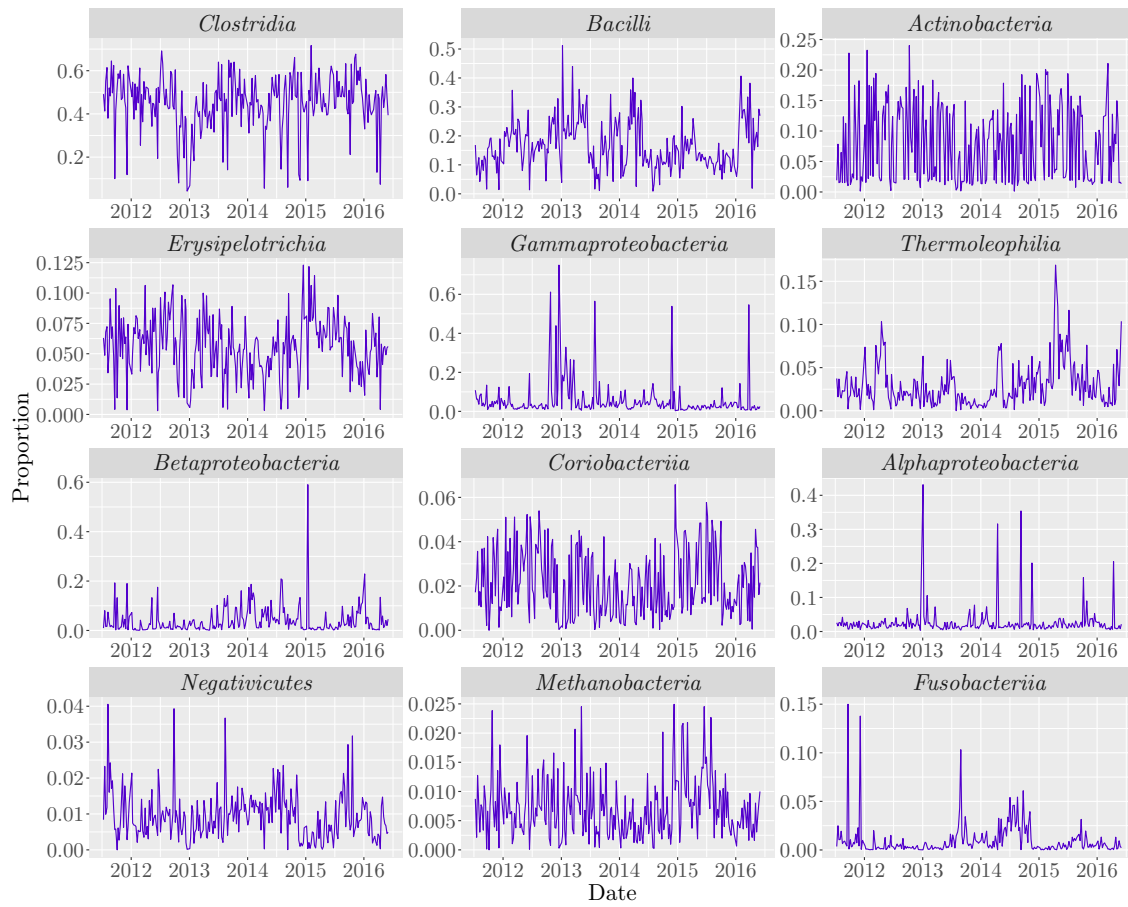
teobacteria show rough annual peaks. *Clostridia* and *Bacilli* also show seasonal behaviour, which is clearer on the log-scale (not shown). The time series plot for *Bacilli* is very similar to the plot for OTU 1 in Figure 6.4a. After further inspection, we find that *Bacilli* is represented by 223 OTUs in the AS tank, one of which is OTU 1 from the genus *Trichococcus*. On average this OTU represents about 77% of the abundance of *Bacilli*, thus explaining why their profiles are so similar and why we see evidence of seasonality in the class. Plotting *Gammaproteobacteria* on the log-scale (not shown) also reveals seasonal behaviour. The most abundant class is *Alphaproteobacteria* and its time series profile is very noisy without any obvious annual peaks. *Alphaproteobacteria* form one of the most abundant groups of bacteria on the planet and are extremely diverse (Williams *et al.*, 2007), so, perhaps it is unsurprising that this class is the most abundant in the AS tank. The diversity of *Alphaproteobacteria* is also reflected here, as there are 1238 different OTUs from the class present in the AS. It seems reasonable to assume that a class as diverse as this would have species, or OTUs in our case, that prefer different conditions and hence have peaks in population size at different times of the year.

Figure 6.8b shows the stacked bar plot for the top 12 classes in the AS tank. We can see that these classes represent most of the abundance in the tank and thus seasonality is not easy to detect. The dominance of *Alphaproteobacteria* and *Betaproteobacteria* is clear in this plot, although spikes in *Gammaproteobacteria*, *Bacilli* and *Clostridia* are also visible.

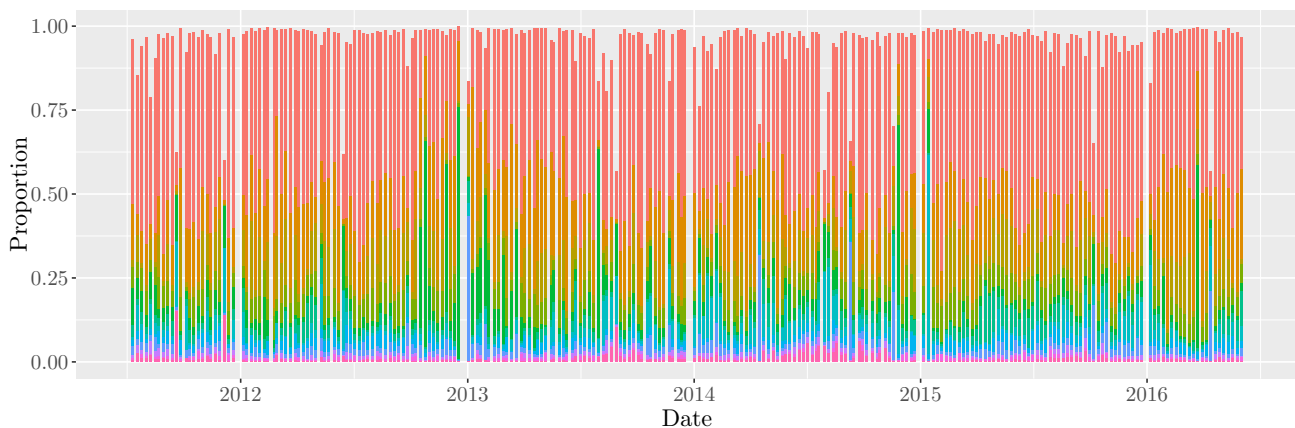
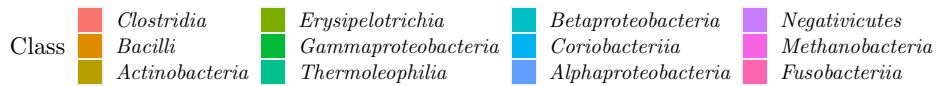
Figures 6.9a and 6.9b show the time series plots and stacked bar plot for the top 12 classes in the SS, respectively. As we saw with the OTUs and genera, there do not seem to be obvious signs of seasonality and there seems to be more variability in the SS, again suggesting that weekly intervals are possibly not small enough to capture the changes in population sizes here. However, *Bacilli* shows vague signs of seasonality, which can be explained by the fact that on average, of the 323 OTUs in the class, OTU 1 represents 64% of the class' total abundance. Nevertheless, clearly the microbial populations in the two tanks differ greatly, with different classes dominating. For example, only seven of the 12 top classes are in both tanks. Furthermore, *Clostridia*, *Bacilli* and *Actinobacteria* are the three most abundant classes in the SS, whereas in the AS we have *Alphaproteobacteria*, *Betaproteobacteria* and *Sphingobacteriia* as the top three classes.

6.2.3 Analysis of the combined data

In this section we further our understanding of the data by combining the chemical and environmental data with the other data sets. First, we identify potential relationships between the 12 most abundant OTUs with the chemical and environmental covariates in each tank. We conclude this section, and this chapter, by drawing together all three sources of data and considering the chemical and environmental data with the relative



(a)



(b)

Figure 6.9: (a) time series plots and (b) stacked bar plot for the top 12 classes in the SS based on median abundance.

abundances of the top 12 genera and the top 12 classes.

Relationships with 12 most abundant OTUs

Figure 6.10 shows a heatmap of the correlations between the covariates and the relative abundances of the top 12 OTUs in the AS. OTU 15 shows a strong negative correlation with temperature and its time series plots in Figure 6.4a showed seasonal behaviour. OTU 8 and OTU 28 also displayed seasonality, although their correlations with temperature are very weak. Meanwhile, a lot of other OTUs seem to have stronger correlations with temperature, for example, OTU 7 and OTU 36. Reviewing their time series plots again, it would perhaps be reasonable to say that these OTUs show seasonality too.

There appear to be some correlations between some of the OTUs and chloride, nitrite, COD, DO, phosphate, MLSS and MLVSS. Figure 6.11 shows a heatmap of the pairwise correlations between the chemical and environmental covariates in the AS tank. From this heatmap it would seem as if these covariates are potentially correlated with temperature. With the exception of temperature and chloride, these correlations are fairly weak though so it would be naïve to attribute all of the possible relationships that we see here to an indirect relationship with temperature. It is important to remember that the microorganisms in our tanks are very diverse and interact with, feed on and utilise chemical compounds in many different ways. However, it does seem that in general for these top 12 OTUs in the AS, if the correlation with temperature is weaker, then the correlations with other covariates tend to be weaker too.

Figure 6.12 shows the heatmaps of the correlations between the top 12 OTUs and the chemical and environmental covariates in the SS. Note that we have split the covariates into two groups, putting the first 13 covariates that appear in the data set in one heatmap and the remaining 13 covariates into another. There are quite a lot of fairly strong correlations present in the SS. For example, OTU 5 and OTU 30 seem to be correlated with ammonia, phosphate, potassium, calcium, flow, magnesium and nitrate. Fluoride also seems to be weakly correlated with half of the top 12 OTUs.

Again, the correlations amongst the chemical covariates need to be considered here. A heatmap of these correlations is shown in Figure 6.13. Nitrate has a strong positive correlation with calcium, flow and magnesium. With their negative correlation with nitrate, we see that phosphate and ammonia have a negative correlation with flow, calcium and magnesium. Phosphate and potassium are positively correlated. Potassium also has a negative correlation with flow. Negative correlations between chemical covariates and flow may simply demonstrate dilution, since flow is the rate at which water enters the tank. Some of the covariates do not appear to be correlated with any of the top 12 OTUs, for example, COD, aluminium, cadmium, copper, sulphate, arsenic, nickel and manganese, although it is worth noting that some of these covariates have a large number of values as

zero. For example, cadmium has a value of zero 90.66% of the time.

We do not have temperature recordings in the SS tank. However, it would be reasonable to assume that chloride could be used as a rough proxy for temperature, given that temperature and chloride are negatively correlated in the AS tank. From the heatmap, it seems as if chloride does not have any strong correlations with the relative abundances of the top 12 OTUs in the SS. We already saw that the time series plots for these OTUs did not show any obvious signs of seasonality, other than OTU 27 and OTU 1. So, perhaps this was to be expected.

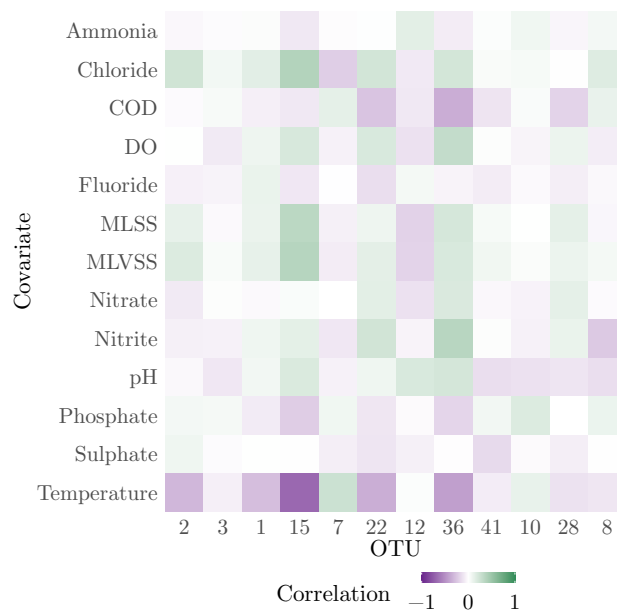


Figure 6.10: Heatmap of the correlations between the chemical and environmental covariates and the 12 most abundant OTUs in the AS tank.

Relationships with 12 most abundant genera

As we did with the 12 top OTUs, we now look at heatmaps of correlations between the covariates and the 12 most abundant genera in each tank. Figure 6.14 shows the heatmap for the AS tank. Most of the genera have at least a weak correlation with temperature, with the exception of *Acidovorax*, *Ca. Microthrix* and *Deftuviimonas*. Three of the genera appear to have a strong correlation with temperature, *Rhodobacter*, *Ferruginibacter* and *Tetrasphaera*, which we already identified as showing seasonal behaviour in Section 6.2.2.

Here we see that pH has a weak correlation with *Acidovorax*, *Ferruginibacter* and *Flavobacterium*. There are weak correlations present between: nitrite and *Tetrasphaera*; nitrite and *Rhodobacter*; and nitrate and *Tetrasphaera*. COD has weak correlations with a few of the genera, as does chloride, with a stronger positive correlation with *Rhodobacter*.

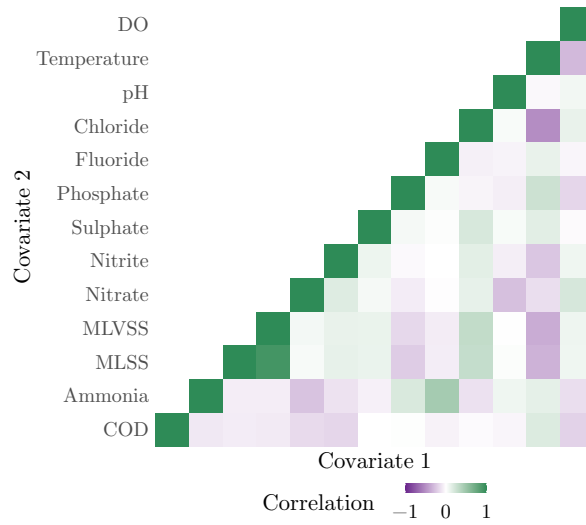


Figure 6.11: Heatmap of the pairwise correlations between the chemical and environmental covariates in the AS tank.

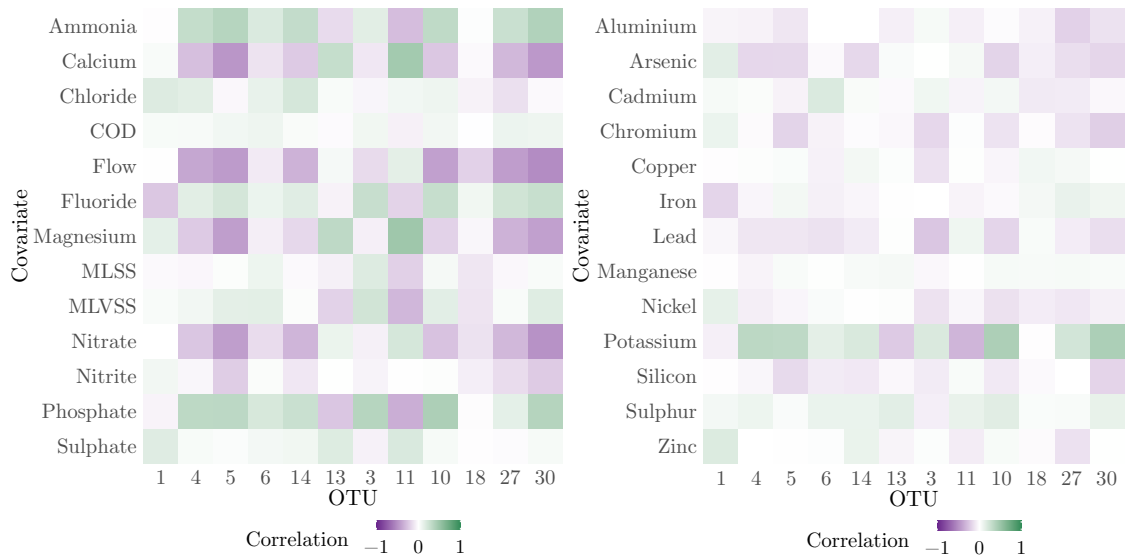


Figure 6.12: Heatmaps of correlations between the chemical and environmental covariates and the 12 most abundant OTUs in the SS tank.

This can possibly be attributed to the strong negative correlation this genus has with temperature and the negative correlation that we have already observed between temperature and chloride.

Heatmaps for the correlations between the 12 most abundant genera and the chemical

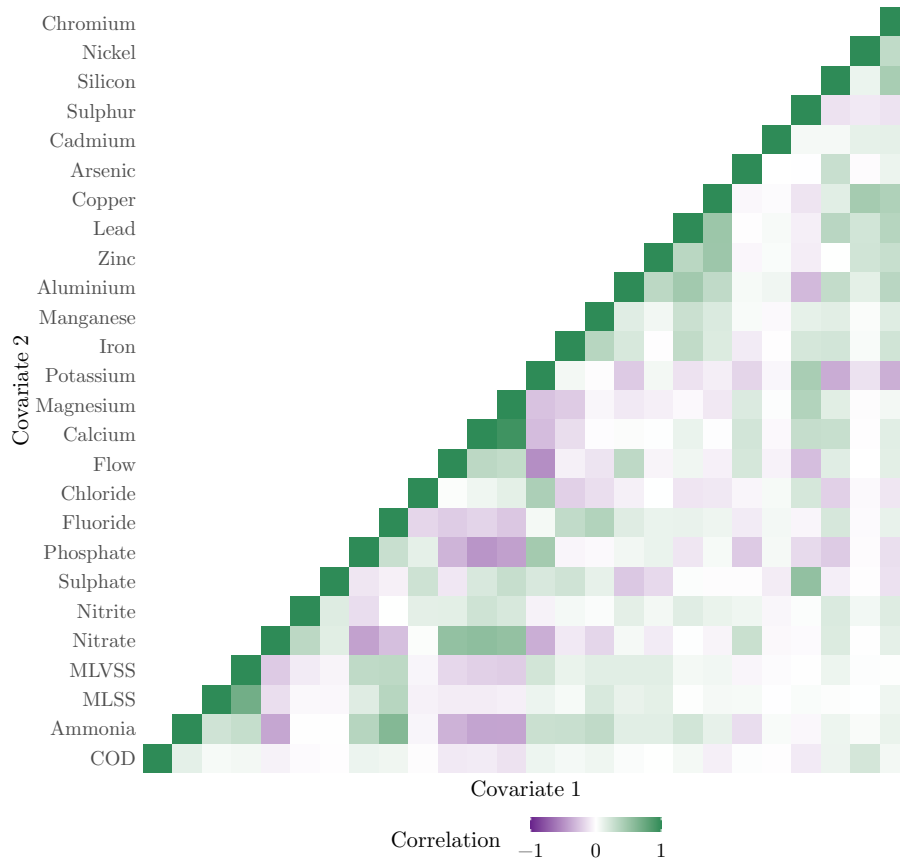


Figure 6.13: Heatmap of the pairwise correlations between the chemical and environmental covariates in the SS tank.

and environmental covariates in the SS tank are shown in Figure 6.15. The unknown genera have a positive correlation with fluoride and a weak negative correlation with chloride. *Trichococcus* has a negative correlation with fluoride but does not seem to be correlated with any of the other covariates. Several of the genera are correlated with ammonia, nitrate, calcium, flow, potassium, magnesium and phosphate. Additionally, *Streptococcus* has a positive correlation with sulphur and *Blautia* has a weak negative correlation with nitrite. Finally, MLVSS is negatively correlated with *Roseburia* and *Ruminococcus 2*.

Relationships with 12 most abundant classes

Now we look at the correlations between the top 12 classes in the AS tank and the chemical and environmental covariates shown in a heatmap in Figure 6.16. Here the seasonality we saw in Section 6.2.2 is confirmed with most classes having at least a weak correlation with temperature. *Actinobacteria* has a fairly strong negative correlation with temperature suggesting that the bacteria of this class present in the AS tank peak in

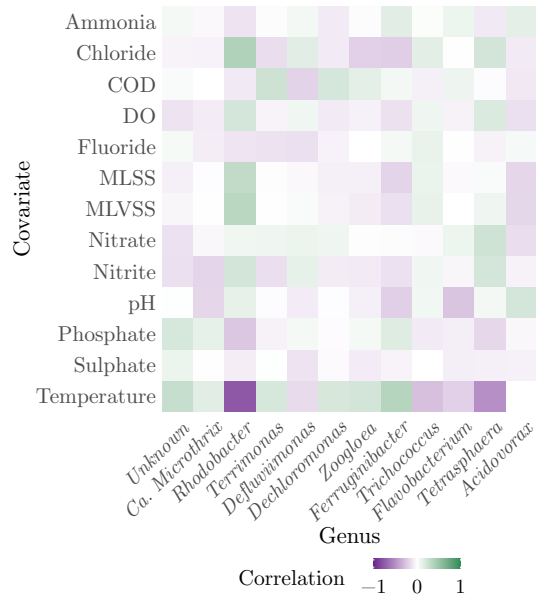


Figure 6.14: Heatmap of the correlations between the chemical and environmental covariates and the 12 most abundant genera in the AS tank.

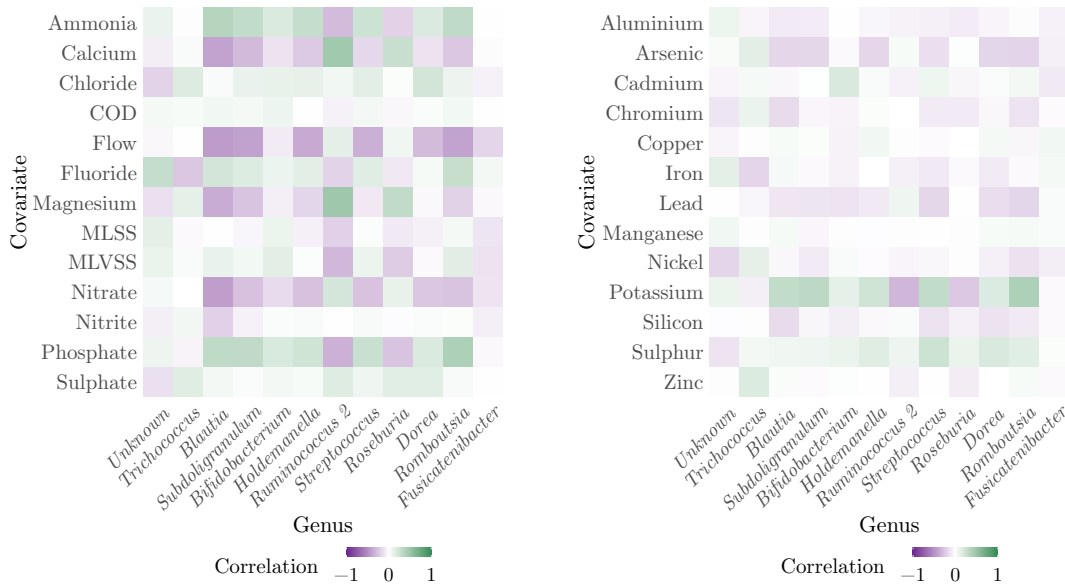


Figure 6.15: Heatmaps of correlations between the chemical and environmental covariates and the 12 most abundant genera in the SS tank.

winter. *Flavobacteriia* appears to have a negative correlation with pH and *Acidimicrobiia* appears to have a weak negative correlation with pH too. Sulphate also seems to have a weak positive correlation with *Thermoleophilia*. The remaining covariates do not seem to

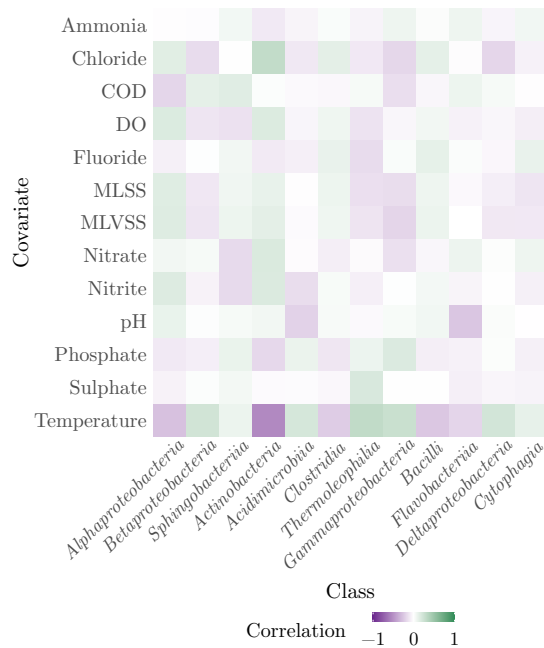


Figure 6.16: Heatmap of the correlations between the chemical and environmental covariates and the 12 most abundant classes in the AS tank.

be correlated with the classes.

Figure 6.17 shows the heatmaps of the correlations in the SS tank between the 12 most abundant classes and the covariates. As we saw with the top 12 OTUs and top 12 genera, there seem to be stronger correlations in general. Most of the classes seem to be correlated with flow, calcium, fluoride, phosphate, ammonia and nitrate and potassium. With the exception of *Negativicutes*, there does not seem to be any correlations between the top 12 classes and chloride. This might suggest a tenuous positive relationship between temperature and *Negativicutes*, although the time series plot for this class did not show clear evidence of seasonality. Iron has weak correlations with *Bacilli* and *Fusobacteriia*. Magnesium is weakly correlated with *Gammaproteobacteria*, *Methanobacteria* and *Fusobacteriia*. MLVSS has a weak positive correlation with *Coriobacteriia* and *Thermoleophilia*. Both silicon and chromium have positive correlations with *Gammaproteobacteria*. For the remaining covariates, there does not seem to be any obvious relationships with any of the classes.

6.2.4 Summary

From this exploratory analysis, we have identified that there are relationships between some of the chemical and environmental covariates and the relative abundances of some of the top 12 OTUs in the AS and SS. Relationships can also be seen at the coarser



Figure 6.17: Heatmaps of correlations between the chemical and environmental covariates and the 12 most abundant classes in the SS tank.

taxonomic ranks of genera and classes. We have also seen that some of the chemical and environmental covariates are correlated with each other. Finally, we have observed signs of seasonality and absence of time trend in the relative abundances at both fine and coarse taxonomic ranks, although seasonality is less obvious in the SS than in the AS. All of these observations will help to inform decisions when developing our model in the following chapter and they will also aid interpretation of our results in Chapter 8.

Chapter 7

A Bayesian hierarchical model for time-series metagenomics data

In this chapter, due to the dimensions of our data and the presence of sparsity, we first discuss clustering methods, before selecting a phase-based approach to clustering. We then introduce a Bayesian hierarchical vector autoregressive model, which is the starting point for modelling our clustered data. In the sections that follow this, we discuss how we extend this basic model, based on ideas from the current literature and on the findings of the exploratory data analysis in the previous chapter. Allowing our matrix of autoregressive coefficients to be sparse, leads to an exploration into and testing of different shrinkage priors for this matrix-valued parameter. The starting model assumes a diagonal error covariance matrix with common variance. However, we discuss more complicated error structures before selecting a symmetric, circulant, tri-diagonal *precision* matrix for the errors. Finally, we allow for a time varying mean which incorporates the chemical and environmental data via regression, before discussing posterior inference for our model.

7.1 Forming subpopulations

For each tank, the data are counts of OTUs, where some OTUs have counts in the thousands and others have (mostly) counts of zero throughout time (for example, see Figure 6.2 in Section 6.1.1). Kaul *et al.* (2017) argue that in microbiome data, zeroes can arise for structural reasons (“hard zeroes”) or due to lack of sampling depth (“soft zeroes”), as discussed in Section 2.2.4. As such, we would expect an excess of zeroes over Poisson variation. Indeed, in our data, 91.515% and 93.716% of the counts are zero in the AS tank and SS tank, respectively. Therefore, the most natural approach would be to use a zero-inflated model. See, for example, Lee *et al.* (2018) and Xia *et al.* (2018). However, there are over 9000 OTUs per tank, indicating that our model would have to allow for

over 81 million pairwise interactions. To make model-fitting more manageable we instead choose to cluster the data, following analyses by other authors (see Eiler *et al.* (2012); Stein *et al.* (2013); David *et al.* (2014); Dam *et al.* (2016)). Choosing a small enough number of clusters removes the problem of zeroes and allows us to make the simplifying assumption that our data can be modelled as continuous.

7.1.1 Clustering based on taxonomy

An approach that reduces the dimensionality of the data but also retains all of the OTUs is clustering based on taxonomy. This involves taking the n most abundant taxa at each time point that represent a high percentage, say 90%, of the total abundance and grouping the remaining taxa into an “other” category. The term taxa here could refer to any taxonomic rank. For example, Stein *et al.* (2013) grouped OTUs into the top ten genera and an “other” category in their work to infer gut microbiota ecology in mice and predict the temporal dynamics under time-dependent external perturbations (introduction of an antibiotic).

For our data, if we let $g_t(k)$ be the relative abundance for the k -th most abundant genus at time t , then we require n such that $\sum_{k=1}^n g_t(k) \approx 90\%$ for each t . However, there are two issues with this approach. Firstly, to capture 90% of abundance, $n = 30$, resulting in 187 genera for the AS tank and 173 genera in the SS tank, in the union over all time points. Secondly, unknown genera are among the most abundant, more specifically 53 of the 187 genera are unknown in the AS tank and 29 of the 173 genera are unknown in the SS tank. This is unsurprising, given that in our exploratory analysis we saw that 54.09% of the OTUs in the AS tank and 46.56% in the SS tank are missing their genus in the taxonomy table (see Table 6.5 in Section 6.1.2).

A solution to the first issue is to let $g(k)$ be the relative abundance for the k -th most abundant taxon over *all* time points and then find the smallest possible n such that $\sum_{k=1}^n g(k) \approx 90\%$. Another option would be to let $g(k)$ be the median relative abundance of the k -th most abundant taxon over all time points. For the problem of missing genera, we could instead consider coarser taxonomic ranks. The finest taxonomic rank that we can use without having an unknown as a group is class, which is possibly too coarse. This is because, as we noted in Section 6.2.2, when we looked at the top 12 classes in each tank, *Alphaproteobacteria* was the most abundant class in the AS tank (with 1238 different OTUs) but this class is known to be extremely diverse in general (Williams *et al.*, 2007). Modelling the change in its abundance over time and its interactions with other classes and the environment seems counter-intuitive, given that the different OTUs within the class may prefer different conditions. With no reconciliation for these problems, we choose not to cluster using taxonomic ranks.

7.1.2 Time series clustering

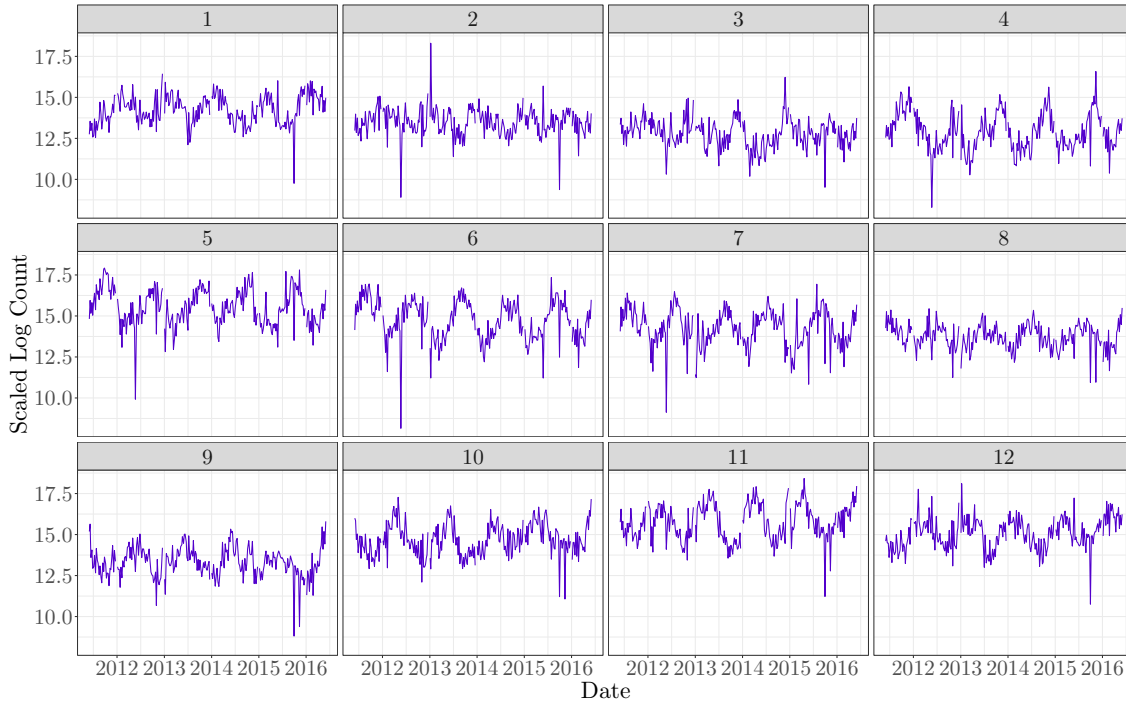


Figure 7.1: Time series plots of the scaled log counts for the 12 bins in the AS tank.

Dam *et al.* (2016) researched dynamic models of the complex microbial metapopulation of Lake Mendota in Wisconsin, USA and suggested that, for characterising interaction dynamics, clustering by taxonomy is not an effective strategy. They looked at the top seven phyla, orders, classes and genera in their data and found that the abundance profiles of OTUs in each ‘cluster’ varied widely, which is what we have found with our data (for example, the *Alphaproteobacteria* in the AS tank). They proposed an alternative method of clustering OTUs, where they define peak profiles, which involves identifying positions in time where each OTU has its largest abundance(s). They superimposed their eleven years of data which resulted in a single ‘collective year period’ and smoothed the data by computing the mean value of each 30 day interval. They then identified the position of the maximum value of each OTU to define the peak profiles. The OTUs were then clustered into “subcommunities” based on these profiles with remaining OTUs placed in an additional group. The rationale is that the subcommunities represent OTUs with similar dynamics perhaps because of symbiotic relationships or shared dependence on the environment. Since we have clear evidence of seasonality in our data we adopt a similar approach here. Our method is as follows:

1. Calculate the scaled weekly means of each OTU.

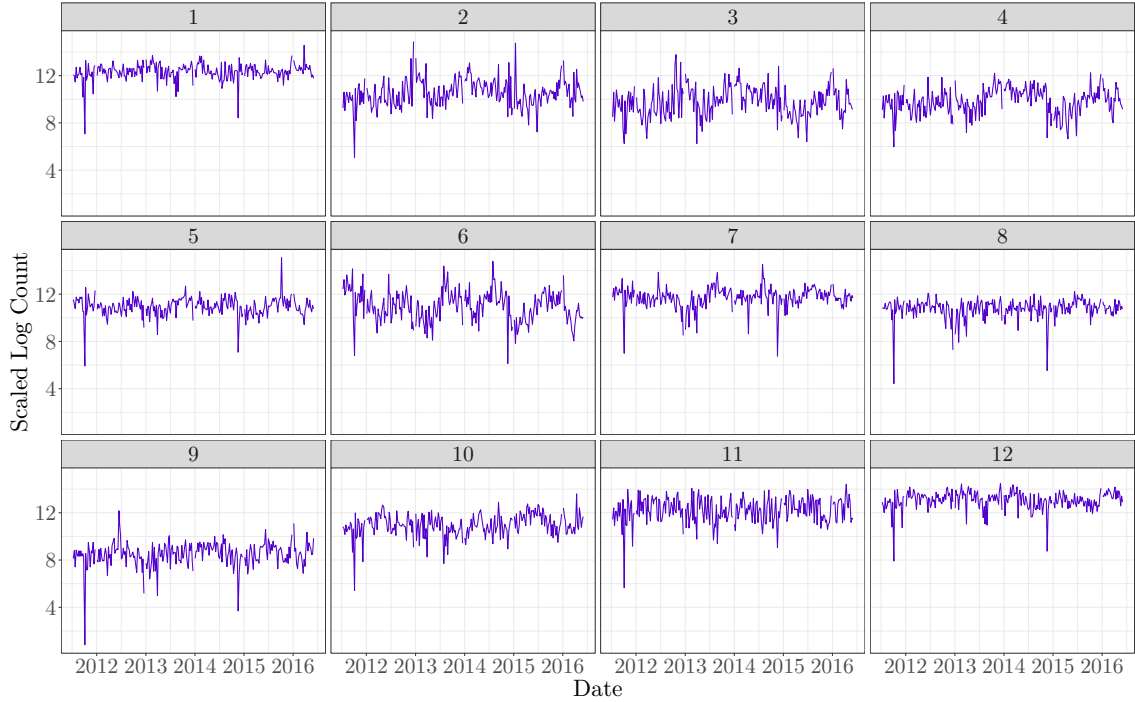


Figure 7.2: Time series plots of the scaled log counts for the 12 bins in the SS tank.

2. Represent the annual series for each OTU using a Fourier basis and calculate the phase and amplitude of each harmonic, where the frequency is $2\pi kt/51$ for the k -th harmonic and time t . Denote the first phase of OTU i by ϕ_i .
3. Divide the interval $[-\pi, \pi]$ into 12 equally sized intervals and assign each OTU i to the interval in which ϕ_i lies for all i . This gives 12 clusters, which we call “bins”. Let $\tilde{w}_{t,i}$ be the count of OTU i at time point t . The set of OTUs in bin j is S_j and $w_{t,j} = \sum_{i \in S_j} \tilde{w}_{t,i}$ is the count for bin j at time t . The counts of each bin peak once per year, with different bins peaking in different months.

We plot a histogram (not shown) of $w_{t,j}$ over all t , for each bin, and find that the counts of the bins are positively skewed. Therefore, we log-transform the counts of the bins and set $\tilde{y}_{t,j} = \log(w_{t,j})$. We then scale the log-counts of each bin so that their variance is roughly one (see Section 7.3.3 and the end of Section 7.3.5). This is done by scaling the $\tilde{y}_{t,j}$, for each j , by dividing by the average (time series) standard deviation. Let s_j be the standard deviation of $\tilde{y}_{1:N,j}$ and $\bar{s} = (\sum_{j=1}^{12} s_j)/12$. We have $y_{t,j} = \tilde{y}_{t,j}/\bar{s}$ for all t and $j = 1, \dots, 12$, where $y_{t,j}$ denotes the scaled log-counts. Figures 7.1 and 7.2 show the time series plots of the scaled log counts for the 12 bins in the AS tank and SS tank, respectively. In the AS, each bin clearly shows seasonal behaviour with a peak every year, with the exception of bins 2 and 3, where the peaks are not as obvious. We can see that for each bin the annual peaks are different, for example, bin 6 seems to peak around September

each year, whereas bin 12 seems to peak around March each year. Seasonality in the bins of the SS is not so obvious, as we observed with the 12 most abundant OTUs, genera and classes in our exploratory data analysis in Chapter 6. With the exception of bin 6, annual peaks are hard to detect in the bins of the SS tank. As we suggested previously with the top 12 OTUs, genera and classes, this is likely due to weekly intervals not being small enough to capture the rapid changes in the microbial populations.

7.2 Simple VAR(1) model

As this chapter’s preamble states, the foundation of our model is a vector autoregressive model, and so we give its general definition. Let \mathbf{y}_t be a multivariate, K -dimensional process. We say that \mathbf{y}_t follows a zero-mean vector autoregressive model of order p , typically denoted as VAR(p), if we can write \mathbf{y}_t as a linear regression on its last p values. This is

$$\mathbf{y}_t = A_1\mathbf{y}_{t-1} + A_2\mathbf{y}_{t-2} + \cdots + A_p\mathbf{y}_{t-p} + \boldsymbol{\epsilon}_t, \quad t = 1, \dots, N$$

where the A_j s are $K \times K$ matrices of autoregressive coefficients for $j = 1, \dots, p$, and the error terms are $\boldsymbol{\epsilon}_t \sim N_K(\mathbf{0}, \Sigma)$. The error terms are assumed to be time independent, that is, for any times s and t , where $s \neq t$, $\boldsymbol{\epsilon}_s \perp\!\!\!\perp \boldsymbol{\epsilon}_t$ (where $\perp\!\!\!\perp$ means “is independent of”). As discussed in Section 2.2.2 of Chapter 2, a VAR(1) model can be regarded as a linear approximation to a Lotka-Volterra model, which is commonly used in time-series metagenomic studies to model non-linear microbial dynamics. This, coupled with its simplicity, makes a VAR(1) model a sensible starting point for our model. A stationary VAR(1) model with mean $\boldsymbol{\mu}$ has the form

$$\mathbf{y}_t = \boldsymbol{\mu} + A(\mathbf{y}_{t-1} - \boldsymbol{\mu}) + \boldsymbol{\epsilon}_t, \quad t = 1, \dots, N \quad (7.1)$$

where A is the single matrix of autoregressive coefficients. Technically, to enforce stationarity, we would need to impose the constraint that the spectral radius of A is less than one, that is $\max\{|\eta_1|, \dots, |\eta_K|\} < 1$, where η_i is the i -th eigenvalue of A . In our work we leave A unconstrained but we observe in our later analyses in Chapter 8 that all posterior mass ultimately lies within the stationary region.

In the completely general case, the error variance matrix Σ simply needs to lie in the space of symmetric, positive definite matrices. However, very often a parsimonious parametric form is adopted, for example, $\Sigma = \tau^{-1}I_K$. Under this model specification, a

typical (and semi-conjugate) prior would be

$$\begin{aligned}\boldsymbol{\mu} &\sim N_K(\mathbf{m}_\mu, \mathbf{V}_\mu), \\ \tau &\sim \text{Ga}(g, h), \\ a_{jk} &\stackrel{iid}{\sim} N(c, d^2).\end{aligned}\tag{7.2}$$

We note that in our case, each element a_{jk} of \mathbf{A} can be thought of as determining the influence of the abundance of bin k at the previous time point $y_{t-1,k}$ on the abundance of bin j at the current time point y_{tj} .

7.3 Allowing for sparsity

In our clustered data we have 257 observations for 12 bins in the AS tank and 251 observations for 12 bins in the SS tank. Given the circular time-ordering of the bins, it is unlikely that the previous abundances for all bins influence the current abundance of any particular bin. For example, a bin that peaks in May is probably more influenced by bins that peak in neighbouring months April and June than bins that peak five to seven months later (and have low abundances in May). A sparse autoregressive matrix would contain a lot of zeroes and reflects this idea that not all bins influence each other. A zero-coefficient represents no influence, given autoregression on $y_{t-1,k}$, of the count of bin k at time $t - 1$ on the count of bin j at current time t , which is likely to aid in understanding the multivariate dynamics implied by our results in Chapter 8. The prior for \mathbf{A} described in (7.2) is perhaps the most simple but it does not encourage shrinkage of autoregressive coefficients towards zero. In this section, we discuss priors that allow for sparsity in the autoregressive matrix \mathbf{A} . Each prior is introduced in the context of simple multiple linear regression, before being extended to the autoregressive matrix for a VAR(1) process. This extension is straightforward since a VAR(1) model can be written as a multiple linear regression with a multivariate response vector.

First we consider priors for $\boldsymbol{\beta}$ in

$$y_i = \boldsymbol{\beta}^T \mathbf{x}_i + \epsilon_i, \quad \epsilon_i \sim N(0, \sigma^2), \quad i = 1, \dots, N,\tag{7.3}$$

where y_i is a univariate response, \mathbf{x}_i is a p -variate vector of explanatory variables for experimental unit i , $\boldsymbol{\beta}$ is the vector of corresponding regression coefficients and σ^2 is the error variance. The problem of identifying regression coefficients as zero or non-zero is essentially a variable selection problem for the covariates in the model. In the case of our VAR(1) model, it is selecting if the count (at the previous time point) of a particular bin is used to model the count of another bin or not. The first prior that we discuss is the spike

and slab prior, a firm favourite for Bayesian variable selection. The second prior is the horseshoe prior and the third is an extension of the horseshoe, the regularised horseshoe prior. Finally, we conclude this section on sparsity, by comparing our priors and testing our inferential procedures via a simulation study.

7.3.1 Spike and slab prior

The spike and slab (Mitchell & Beauchamp, 1988) is a popular shrinkage prior; its name comes from its density with a spike at zero and a flat slab elsewhere. It is often written as a two-component mixture of normal distributions (Piironen & Vehtari, 2017). In our simple multiple linear regression example, we have

$$\begin{aligned}\beta_j | \lambda_j, d, \epsilon &\sim \lambda_j N(0, d^2) + (1 - \lambda_j) N(0, \epsilon^2), \\ \lambda_j &\sim \text{Bern}(\psi_j), \quad j = 1, \dots, p,\end{aligned}\tag{7.4}$$

where the indicator variable $\lambda_j \in \{0, 1\}$ denotes presence ($\lambda_j = 1$) or absence ($\lambda_j = 0$) of coefficient β_j , ψ_j is the prior inclusion probability and $\epsilon \ll d$. Often ϵ is chosen to be zero, resulting in a delta spike at the origin, which is the approach we adopt for our work. The conditional density of β_j , given λ_j , is

$$\lambda_j f_N(\beta_j; 0, d^2) + (1 - \lambda_j) \delta_{\beta_j},$$

where $f_N(x; 0, d^2) = d^{-1} \phi(x/d)$ is the density of a normal $N(0, d^2)$ distribution and δ_x is the Dirac delta function, which takes the value 0 everywhere except at $x = 0$ but which integrates to 1. We denote this as

$$\begin{aligned}\beta_j | \lambda_j, d &\sim \lambda_j N(0, d^2) + (1 - \lambda_j) \delta_0, \\ \lambda_j &\sim \text{Bern}(\psi_j), \quad j = 1, \dots, p.\end{aligned}\tag{7.5}$$

O'Hara & Sillanpää (2009) give a comprehensive review of some Bayesian variable selection methods in regression with spike and slab priors and suggest that, for such methods, it is safe to only consider problems where there are at most 10 to 15 times more candidates than observations. That is, effects can still be detected by models, even when the number of candidates is 10 to 15 times larger than the number of observations. However, this upper limit's existence is problem specific and also depends on the level of correlation between candidates.

When using certain methods to implement the spike and slab, O'Hara & Sillanpää (2009) highlight the need for a second auxiliary variable, the effect size, which we will call $\tilde{\beta}_j$. Now we define our regression coefficient as $\beta_j = \lambda_j \tilde{\beta}_j$. The indicators and effects are assumed to be independent *a priori*, which is the approach of Kuo & Mallick (1998),

that is $P(\tilde{\beta}_j, \lambda_j) = P(\tilde{\beta}_j)P(\lambda_j)$. Additionally, $\lambda_j \perp\!\!\!\perp \lambda_k$ for $j \neq k$, $\tilde{\beta}_j \perp\!\!\!\perp \tilde{\beta}_k$ for $j \neq k$ and $\lambda_j \perp\!\!\!\perp \tilde{\beta}_k \forall j, k$. In the simple multiple linear regression case the prior distributions are $\tilde{\beta}_j \sim N(0, d^2)$ and $\lambda_j \sim \text{Bern}(\psi_j)$, where ψ_j is the chosen prior inclusion probability.

Now we consider the spike and slab prior for multiple linear regression with a *multivariate* response vector. Let $\mathbf{y}_i = (y_{i1}, \dots, y_{iK})^T$ be a K -variate response vector and \mathbf{x}_i be a p -variate vector of explanatory variables for experimental unit i . We have a multiple linear regression model

$$\mathbf{y}_i = \mathbf{B}^T \mathbf{x}_i + \boldsymbol{\epsilon}_i, \quad \boldsymbol{\epsilon}_i \sim N_K(\mathbf{0}, \Sigma), \quad i = 1, \dots, N, \quad (7.6)$$

where $\boldsymbol{\epsilon}_i = (\epsilon_{i1}, \dots, \epsilon_{iK})^T$ is the vector of errors and $\mathbf{B} = (\beta_{jk})$ is a $p \times K$ matrix of regression coefficients. The elements (y_{ij}) of \mathbf{y}_i can be modelled as a linear combination of the elements in \mathbf{x}_i , subject to normal error ϵ_{ij} . The errors are conditionally independent across units i , but may be correlated across variables j . The model in (7.6) can be written in the matrix form

$$\mathbf{Y} = \mathbf{X}\mathbf{B} + \mathbf{E}, \quad (7.7)$$

where \mathbf{Y} is a $N \times K$ response matrix with i -th row \mathbf{y}_i^T , \mathbf{X} is a $N \times p$ covariate matrix with i -th row \mathbf{x}_i^T and \mathbf{E} is a $N \times K$ matrix of errors with i -th row $\boldsymbol{\epsilon}_i^T$. A spike and slab prior extending the formulation in (7.4) is

$$\begin{aligned} \beta_{jk} | \lambda_{jk}, d^2 &\sim \lambda_{jk} N(0, d^2) + (1 - \lambda_{jk}) \delta_0 \\ \lambda_{jk} &\sim \text{Bern}(\psi_{jk}), \end{aligned} \quad (7.8)$$

where the random variables λ_{jk} are independent of each other.

Alternatively, using the approach of Kuo & Mallick (1998), we have

$$\begin{aligned} \beta_{jk} &= \lambda_{jk} \tilde{\beta}_{jk}, \\ \tilde{\beta}_{jk} | d^2 &\sim N(0, d^2), \\ \lambda_{jk} &\sim \text{Bern}(\psi_{jk}), \end{aligned} \quad (7.9)$$

for $j = 1, \dots, p$, $k = 1, \dots, K$ and λ_{jk} are independent of each other.

Lei *et al.* (2011) considered the problem of learning the parameters and structure of a VAR(1) model with relatively short, high-dimensional time series. For simulated data and a time course microarray data set, their methods performed reasonably well. Their autoregressive matrix \mathbf{A} was assumed to be sparse, for which they implemented a spike and slab prior using reversible jump MCMC. They used the form of prior shown in (7.8) for a_{jk} , the elements of \mathbf{A} . In their simulation study, they assessed their algorithm's ability to detect and infer the size of non-zero elements of an autoregressive matrix with both

a small number of variables ($K = 5$) and a fairly large number of variables ($K = 100$). In data-poor scenarios, when the numbers of observations were small, for example, when $K = 5$ and $N = 10$, or $K = 100$ and $N = 30$, their algorithm did not perform well. However, when the number of observations was increased performance improved. For example, with $K = 100$ and $N = 150$ they found that 86% of the posterior mean values for non-zero elements were within 0.1 of their true values. The “positive predictive values” (number of true non-zero elements divided by number of non-zero elements detected) were 99% and 100% when using a probability threshold of 0.5 and 0.7 respectively, and even when a low threshold of 0.1 was used, the positive predictive value was 78%. Here the term threshold relates to the posterior inclusion probabilities $\Psi|Y$. If the posterior inclusion probability $\psi_{jk}|Y$ of a particular element a_{jk} is greater than the threshold, then $a_{jk}|Y$ is considered to be a non-zero autoregressive coefficient. It must be noted that, based on their results, it would seem that their spike and slab prior made it difficult to identify small non-zero elements. For example, in the $K = 5$ case, three of the elements were given the value 0.1, where the error variance was 0.25, and even in the most data-rich scenario with 1000 observations, one of the three was not detected unless the probability threshold was reduced to 0.35. Additionally, in the $K = 100$ case a fairly “strong” value of 0.4 was given to all of the non-zero elements. It is possible that smaller values for the non-zero elements may not have yielded such good results. Since we have $K \ll N$ for our clustered data, and in light of these findings, it seems we are in a “data-rich” scenario and should be able to identify a sparse structure if there is evidence of it in the data.

Full conditional distributions

To assess this prior later in our simulation study in Section 7.3.5, the full conditional distributions (FCDs) are required, as we will use a Gibbs sampler to fit the model with a spike and slab prior. We derive the FCDs for our multiple linear regression (with multivariate response vector) case before stating the prior and FCD for the VAR(1) model.

The rows of Y and B^T can be stacked into NK - and Np - vectors, respectively, to give us

$$\begin{aligned} \mathbf{y}^* &= \text{vec}(Y^T) = (y_{11}, \dots, y_{1K}, y_{21}, \dots, y_{NK})^T \text{ and} \\ \boldsymbol{\beta}^* &= \text{vec}(B) = (\beta_{11}, \dots, \beta_{p1}, \beta_{12}, \dots, \beta_{pK})^T. \end{aligned}$$

Now (7.6) can be written as

$$\mathbf{y}_i = W_i \boldsymbol{\beta}^* + \boldsymbol{\epsilon}_i, \quad (7.10)$$

where $W_i = I_K \otimes \mathbf{x}_i^T$ and \otimes is the Kronecker product. Then we can rewrite (7.7) as

$$\mathbf{y}^* = W \boldsymbol{\beta}^* + \mathbf{e}, \quad (7.11)$$

where

$$\mathbf{W} = \begin{pmatrix} \mathbf{W}_1 \\ \vdots \\ \mathbf{W}_N \end{pmatrix}, \quad \mathbf{e} = (\boldsymbol{\epsilon}_1^T, \dots, \boldsymbol{\epsilon}_N^T)^T.$$

First, we consider $\tilde{\mathbf{B}}$, or in its vectorised form $\tilde{\boldsymbol{\beta}}$, and note that

$$\boldsymbol{\beta}^* = \left(\tilde{\beta}_{11}\lambda_{11}, \dots, \tilde{\beta}_{p1}\lambda_{p1}, \tilde{\beta}_{12}\lambda_{12}, \dots, \tilde{\beta}_{pK}\lambda_{pK} \right)^T.$$

If we let $\mathbf{W}^* = \mathbf{W}\boldsymbol{\Lambda}^*$, where $\boldsymbol{\Lambda}^* = \text{diag}(\lambda_{11}, \dots, \lambda_{p1}, \lambda_{12}, \dots, \lambda_{pK})$ and note that $(\mathbf{I}_N \otimes \boldsymbol{\Sigma})^{-1} = (\mathbf{I}_N \otimes \boldsymbol{\Sigma}^{-1})$ then it can be shown that

$$\tilde{\boldsymbol{\beta}}|\mathbf{y}^*, \boldsymbol{\Lambda}, \boldsymbol{\Sigma} \sim N_{pK}(\mathbf{M}_{\tilde{\mathbf{B}}}, \mathbf{V}_{\tilde{\mathbf{B}}}), \quad (7.12)$$

where

$$\begin{aligned} \mathbf{V}_{\tilde{\mathbf{B}}} &= \left(d^{-2}\mathbf{I}_{pK} + \mathbf{W}^{*T} (\mathbf{I}_N \otimes \boldsymbol{\Sigma}^{-1}\mathbf{W}^*) \right)^{-1} \\ \mathbf{M}_{\tilde{\mathbf{B}}} &= \mathbf{V}_{\tilde{\mathbf{B}}}\mathbf{b}_{\tilde{\mathbf{B}}}, \text{ where} \\ \mathbf{b}_{\tilde{\mathbf{B}}} &= \mathbf{W}^{*T} (\mathbf{I}_N \otimes \boldsymbol{\Sigma}^{-1}) \mathbf{y}^*. \end{aligned}$$

The derivation for this is shown in Appendix A.

In the special case, where $\boldsymbol{\Sigma} = \tau^{-1}\mathbf{I}_K$, it can be shown that the columns of $\tilde{\mathbf{B}}$ are independent in the posterior with the j -th column having mean

$$\left(d^{-2}\mathbf{I}_p + \tau\mathbf{X}_j^{*T}\mathbf{X}_j^* \right)^{-1} \tau\mathbf{X}_j^{*T} \mathbf{y}_{1:N,j} \quad (7.13)$$

and variance

$$\left(d^{-2}\mathbf{I}_p + \tau\mathbf{X}_j^{*T}\mathbf{X}_j^* \right)^{-1}, \quad (7.14)$$

where $\mathbf{y}_{1:N,j}$ is column j of \mathbf{Y} , $\mathbf{X}_j^* = (\lambda_{1j}\mathbf{X}_{1:N,1}, \dots, \lambda_{pj}\mathbf{X}_{1:N,p}) = \mathbf{X}\boldsymbol{\Lambda}_j^*$, where $\boldsymbol{\Lambda}_j^* = \text{diag}(\lambda_{1j}, \dots, \lambda_{pj})$ and $\mathbf{X}_{1:N,j} = (\mathbf{X}_{1j}, \dots, \mathbf{X}_{Nj})^T$. This is the multivariate analogue of the univariate case with the assumption of independent errors with a common variance derived by Kuo & Mallick (1998). A complete derivation of this result is given in Appendix A.

The FCD for the indicator parameters is

$$\begin{aligned}\lambda_{jk}|\mathbf{y}^*, \tilde{\mathbf{B}}, \Sigma &\sim \text{Bern}\left(\tilde{\psi}_{jk}\right), \\ \tilde{\psi}_{jk} &= c_{jk}/(c_{jk} + d_{jk}), \\ c_{jk} &= \psi_{jk} \exp\left\{-\frac{1}{2}(\mathbf{y}^* - \mathbf{W}\boldsymbol{\beta}^{**})^T (\mathbf{I}_N \otimes \Sigma^{-1}) (\mathbf{y}^* - \mathbf{W}\boldsymbol{\beta}^{**})\right\}, \\ d_{jk} &= (1 - \psi_{jk}) \exp\left\{-\frac{1}{2}(\mathbf{y}^* - \mathbf{W}\boldsymbol{\beta}^{***})^T (\mathbf{I}_N \otimes \Sigma^{-1}) (\mathbf{y}^* - \mathbf{W}\boldsymbol{\beta}^{***})\right\},\end{aligned}\tag{7.15}$$

where $\boldsymbol{\beta}^{**}$ is $\boldsymbol{\beta}^*$ with its $\{(k-1)p + j\}$ -th element set to $\tilde{\beta}_{jk}$ and $\boldsymbol{\beta}^{***}$ is $\boldsymbol{\beta}^*$ with its $\{(k-1)p + j\}$ -th element set to 0. This FCD is easily derived by noting that

$$\begin{aligned}\psi_{jk} &= \Pr(\lambda_{jk} = 1 | \mathbf{Y}, \cdot) \\ &= \frac{\pi(\mathbf{Y} | \lambda_{jk} = 1, \cdot) \Pr(\lambda_{ij} = 1)}{\pi(\mathbf{Y} | \lambda_{jk} = 1, \cdot) \Pr(\lambda_{jk} = 1) + \pi(\mathbf{Y} | \lambda_{jk} = 0, \cdot) \Pr(\lambda_{jk} = 0)}\end{aligned}$$

and using the formulation of the model in (7.11).

In the special case, when $\Sigma = \tau^{-1}\mathbf{I}_K$, it can be shown (see Appendix A) that

$$c_{jk} \propto \psi_{jk} \exp\left\{-\frac{\tau}{2}(\mathbf{Y}_{1:N,k} - \mathbf{X}\boldsymbol{\beta}_k^{**})^T (\mathbf{Y}_{1:N,k} - \mathbf{X}\boldsymbol{\beta}_k^{**})\right\}\tag{7.16}$$

and

$$d_{jk} \propto (1 - \psi_{jk}) \exp\left\{-\frac{\tau}{2}(\mathbf{Y}_{1:N,k} - \mathbf{X}\boldsymbol{\beta}_k^{***})^T (\mathbf{Y}_{1:N,k} - \mathbf{X}\boldsymbol{\beta}_k^{***})\right\},\tag{7.17}$$

where $\boldsymbol{\beta}_k^{***} = (\beta_{1k}^*, \dots, \beta_{pk}^*)^T$ with its j -th element set to $\tilde{\beta}_{jk}$ and $\boldsymbol{\beta}_k^{**} = (\beta_{1k}^*, \dots, \beta_{pk}^*)^T$ with its j -th element set to 0. Now, we have a version of the Kuo & Mallick (1998) derivation for a multivariate response vector when we assume independent errors with common variance.

Extension to VAR(1) model

Extending this prior to the autoregressive matrix \mathbf{A} of the VAR(1) model in (7.1) with K bins, we have

$$\begin{aligned}a_{jk} &= \lambda_{jk}\tilde{a}_{jk}, \\ \tilde{a}_{jk}|d^2 &\sim \text{N}(0, d^2), \\ \lambda_{jk} &\sim \text{Bern}(\psi_{jk}),\end{aligned}$$

for $j, k = 1, \dots, K$. The derivation of the FCDs for a VAR(1) model is very similar to the derivation for a multiple linear regression with multivariate response vector, given that a

VAR(1) model can be written as a multiple linear regression model. The FCDs are shown in Appendix A, along with their derivations. We find in the special case, when $\Sigma = \tau^{-1}\mathbf{I}_K$, we obtain a result for $\tilde{\mathbf{A}}$ similar to the result found for $\tilde{\mathbf{B}}$. However, for $\tilde{\mathbf{A}}$ it is the *rows* that exhibit independence, rather than the columns. Likewise, we find a similar result for the posterior of Λ too (see Appendix A for the derivations of the FCDs in the special case).

7.3.2 Horseshoe prior

A horseshoe prior (Carvalho *et al.*, 2009, 2010) for the regression coefficients β in (7.3) is as follows:

$$\begin{aligned}\beta_j|\lambda_j, \tau_\beta &\sim \text{N}(0, \tau_\beta^2 \lambda_j^2), \\ \lambda_j &\sim \text{C}^+(0, 1), \quad j = 1, \dots, p, \\ \tau_\beta &\sim \text{C}^+(0, \tau_{\beta_0}),\end{aligned}\tag{7.18}$$

where τ_β is the *global* shrinkage parameter, the λ_j s are the *local* shrinkage parameters, $\text{C}^+(0, c)$ is a half-Cauchy prior with scale parameter c and typically $\tau_{\beta_0} = 1$. The role of τ_β is to estimate the overall sparsity level, whereas the role of the λ_j s is to identify non-zero coefficients. The heavy tails of $\pi(\lambda_j)$ facilitate this. The horseshoe prior's name comes from the density of its shrinkage coefficient $\kappa_j = 1/(1 + \lambda_j^2)$ in the model $\mathbf{y}|\beta \sim \text{N}(\beta, \sigma^2\mathbf{I}_p)$, where \mathbf{y} is a p -dimensional vector and β_j has the prior specified in (7.18). The posterior mean of β_j when $\tau_{\beta_0}^2 = \sigma^2 = 1$ is $\text{E}[\beta_j|\mathbf{y}] = \int_0^1 (1 - \kappa_j) y_j \text{P}(\kappa_j|\mathbf{y}) d\kappa_j = (1 - \text{E}[\kappa_j|\mathbf{y}]) \mathbf{y}$, where $\text{E}[\kappa_j|\mathbf{y}]$ can be interpreted as the amount of shrinkage towards zero, *a posteriori*. Due to the half-Cauchy prior on λ_j , the shrinkage coefficient κ_j has a horseshoe shaped Beta $(\frac{1}{2}, \frac{1}{2})$ prior, hence the name of the prior. Figure 7.3 shows the density of κ_j . When $\kappa_j \approx 0$, there is virtually no shrinkage and it describes signals (β_j is non-zero), whereas when $\kappa_j \approx 1$, there is near-total shrinkage and it describes noise (β_j is close to zero). For our multivariate linear regression model with multivariate response vector in (7.6), a horseshoe prior for the regression coefficients is given by

$$\begin{aligned}\beta_{jk}|\lambda_{jk}, \tau_\beta &\sim \text{N}(0, \tau_\beta^2 \lambda_{jk}^2), \\ \lambda_{jk} &\sim \text{C}^+(0, 1), \\ \tau_\beta &\sim \text{C}^+(0, \tau_{\beta_0}),\end{aligned}\tag{7.19}$$

for $j = 1, \dots, p$, $k = 1, \dots, K$, where the λ_{jk} s are independent from each other.

The horseshoe prior is particularly well suited to sparse problems because its flat, Cauchy-like tails allow strong signals to remain large (unshrunk) *a posteriori* but its infinitely tall spike at the origin provides severe shrinkage for the zero elements of β .

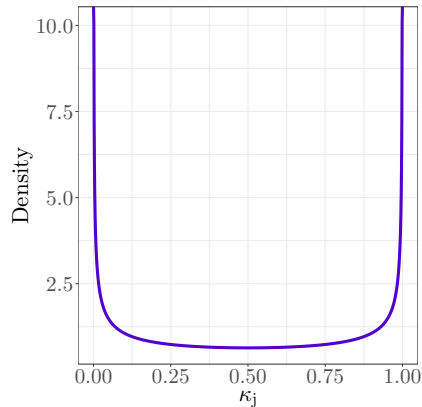


Figure 7.3: The density of the shrinkage coefficient $\kappa_j = 1/(1 + \lambda_j^2)$. The horseshoe shape of the density shows that the horseshoe prior favours values of either zero or one for κ_j , which correspond to no shrinkage or near-total shrinkage, respectively.

Carvalho *et al.* (2009) showed how other commonly used shrinkage priors, such as the Laplacian prior and Student- t prior, do not have these properties. Additionally, they demonstrated that in a variety of situations (experiments) the horseshoe performed similarly to Bayesian model averaging across discrete mixture models, which is considered the “gold standard” for sparse estimation and prediction.

However, Piironen & Vehtari (2017) argue that there has been no agreement on how to perform inference for the global shrinkage parameter τ_β . As its name suggests, this parameter determines the overall shrinkage to zero for the parameter of interest, which can be thought of as determining the overall sparsity in the vector regression coefficients β (in linear regression) or the matrix of coefficients B in multivariate linear regression. They dispute use of the commonly employed $C^+(0, 1)$ hyperprior for the global shrinkage parameter τ_β , as it places too much mass on large values of τ_β , which means it favours solutions with most of the coefficients unshrunk. Thus, when we assume only a small number of variables are relevant, as is often the case, τ_β needs to be strongly identified by the data for sensible inference to occur. Therefore, Piironen & Vehtari (2017) define a systematic way of defining a prior for the global shrinkage parameter based on prior information of sparsity. We discuss this in more detail in Section 7.3.3 and derive a similar result for the multiple linear regression case with a multivariate response vector.

Additionally, Piironen & Vehtari (2017) suggest that it would be useful to be able to control the amount of shrinkage in the largest coefficients. Parameters that the likelihood suggests could be far away from zero typically have $\kappa_j \approx 0$ and so are not regularised. This can be considered disadvantageous, especially when parameters are only weakly identified by the data. The example they give is in logistic regression with separable data, where the likelihood is flat. A remedy to this issue is given by their regularised horseshoe prior, which we discuss in Section 7.3.4.

Full conditional distributions

It can be shown (see Appendix A) that the full conditional density for τ_β is

$$\pi(\tau_\beta | \mathbf{B}, \cdot) = \frac{2}{\pi \tau_{\beta_0} \left(1 + \frac{\tau_\beta^2}{\tau_{\beta_0}}\right)} |2\pi \tau_\beta^2 \Lambda^*|^{-\frac{1}{2}} \exp \left\{ -\frac{1}{2} \left(\boldsymbol{\beta}^{*T} \frac{1}{\tau_\beta^2} \Lambda^{*-1} \boldsymbol{\beta}^* \right) \right\}, \quad (7.20)$$

where $\Lambda^* = \text{diag} \left(\lambda_{11}^2, \lambda_{21}^2, \dots, \lambda_{p1}^2, \lambda_{12}^2, \dots, \lambda_{pK}^2 \right)$ and here

$$\begin{aligned} \boldsymbol{\beta}^* &= \text{vec}(\mathbf{B}) \\ &= (\beta_{11}, \beta_{21}, \dots, \beta_{p1}, \beta_{12}, \dots, \beta_{pK})^T. \end{aligned}$$

This is a non-standard density, which means that a Metropolis-Hastings step needs to be incorporated into the Gibbs sampler (Metropolis-within-Gibbs) for this model.

Similarly, a Metropolis-within-Gibbs step is required for the local shrinkage parameter λ_{jk} , as it has a non-standard conditional density given by

$$\pi(\lambda_{jk} | \mathbf{B}, \cdot) \propto \frac{2}{\pi (1 + \lambda_{jk}^2)} (\lambda_{jk}^2)^{-\frac{1}{2}} \exp \left\{ -\frac{1}{2} \left(\frac{\beta_{jk}^2}{\tau_\beta^2 \lambda_{jk}^2} \right) \right\}. \quad (7.21)$$

The derivation of this density is given in Appendix A.

Extension to VAR(1) model

A horseshoe prior on the autoregressive matrix \mathbf{A} in the VAR(1) model in (7.1) with K bins is given by

$$\begin{aligned} a_{jk} | \lambda_{jk}, \tau_A &\sim \mathcal{N}(0, \tau_A^2 \lambda_{jk}^2) \\ \lambda_{jk} &\sim \mathcal{C}^+(0, 1), \\ \tau_A &\sim \mathcal{C}^+(0, \tau_0), \end{aligned} \quad (7.22)$$

for $j, k = 1, \dots, K$. As we found for the multivariate linear regression model, the global shrinkage parameter τ_A and local shrinkage parameters λ_{jk} have non-standard full conditional densities. Therefore, we must include Metropolis-within-Gibbs steps in our sampler to make posterior draws of these shrinkage parameters. For τ_A , we have

$$\pi(\tau_A | \mathbf{A}, \cdot) = \frac{2}{\pi \tau_0 \left(1 + \frac{\tau_A^2}{\tau_0}\right)} |2\pi \tau_A^2 \Lambda^*|^{-\frac{1}{2}} \exp \left\{ -\frac{1}{2} \left(\mathbf{a}^T \frac{1}{\tau_A^2} \Lambda^{*-1} \mathbf{a} \right) \right\}, \quad (7.23)$$

where $\Lambda^* = \text{diag}(\lambda_{11}, \lambda_{12}, \dots, \lambda_{1K}, \lambda_{21}, \dots, \lambda_{KK})$

$$\begin{aligned} \mathbf{a} &= \text{vec}(\mathbf{A}^T) \\ &= (a_{11}, a_{12}, \dots, a_{1K}, a_{21}, \dots, a_{KK})^T. \end{aligned}$$

For λ_{jk} we have

$$\pi(\lambda_{jk} | \mathbf{A}, \cdot) \propto \frac{2}{\pi(1 + \lambda_{jk}^2)} (\lambda_{jk}^2)^{-\frac{1}{2}} \exp \left\{ -\frac{1}{2} \left(\frac{a_{jk}^2}{\tau_A^2 \lambda_{jk}^2} \right) \right\}. \quad (7.24)$$

The derivation of these densities is very similar to the derivations shown in Appendix A for multivariate linear regression. We use **Stan** to fit the model with a horseshoe prior on \mathbf{A} and so the FCDs and Metropolis algorithms are not needed. However, for completeness we have stated the FCDs here and the algorithms are given in Appendix B.

7.3.3 Prior for the global shrinkage parameter

In this section, we define a prior for the global shrinkage parameter in the horseshoe prior for β_{jk} for our multivariate linear regression model in (7.6), based on prior information of sparsity. To define our hyperprior, we first derive a general shrinkage factor matrix, which indicates how much a set of regression coefficients are shrunk. This is found by considering the conditional posterior distribution for the regression coefficients in (7.6). Since the shrinkage factor matrix depends on the error covariance matrix Σ , we then formulate the shrinkage factor matrix for some different cases of Σ . From the shrinkage factor matrix, we define the effective number of non-zero coefficients m_{eff} , which is based on the assumption that the shrinkage factor matrices tend to be diagonal with their diagonal elements typically taking values of zero or one. We then consider the prior expectation of the effective number of non-zero coefficients, which depends on the global shrinkage parameter τ_β . By setting this prior expectation to some prior guess of the number of non-zero coefficients e_0 and setting $\tau_\beta = \tau_{\beta_0}$, we can solve for τ_{β_0} . This can be used as the scale parameter for the global shrinkage parameter and thus prior information on the number of non-zero coefficients can be incorporated into the prior.

The work in this section is an extension of the work of Piironen & Vehtari (2017). Their work is primarily in the context of simple multiple linear regression with a *univariate* response variable, defined in (7.3). For β_j under the standard horseshoe prior in (7.18) they define a *shrinkage factor* κ_j for coefficient β_j and we include their formulation here for completeness. The conditional posterior distribution for the coefficients β given the

hyperparameters, τ_β and Λ , and the data \mathbf{y} , is

$$\boldsymbol{\beta} | \Lambda, \tau_\beta, \sigma^2, \mathbf{y} \sim \mathcal{N}(\bar{\boldsymbol{\beta}}, \Sigma), \quad (7.25)$$

where

$$\begin{aligned} \bar{\boldsymbol{\beta}} &= \tau_\beta^2 \Lambda \left(\tau_\beta^2 \Lambda + \sigma^2 (\mathbf{X}^T \mathbf{X})^{-1} \right)^{-1} \hat{\boldsymbol{\beta}}, \\ \Sigma &= \left(\tau_\beta^{-2} \Lambda^{-1} + \frac{1}{\sigma^2} \mathbf{X}^T \mathbf{X} \right)^{-1}, \end{aligned}$$

with $\Lambda = \text{diag}(\lambda_1^2, \dots, \lambda_p^2)$, and $\hat{\boldsymbol{\beta}} = (\mathbf{X}^T \mathbf{X})^{-1} \mathbf{X}^T \mathbf{y}$ is the familiar maximum likelihood estimate. If we assume that the covariates are uncorrelated and have zero means and sample variances s_1^2, \dots, s_p^2 , then $\mathbf{X}^T \mathbf{X} \approx N \text{diag}(s_1^2, \dots, s_p^2)$, and we can approximate

$$\bar{\beta}_j = (1 - \kappa_j) \hat{\beta}_j, \quad (7.26)$$

where

$$\kappa_j = \frac{1}{1 + N \sigma^{-2} \tau_\beta^2 s_j^2 \lambda_j^2} \quad (7.27)$$

which is the shrinkage factor for β_j . The shrinkage factor quantifies how much a coefficient β_j is shrunk from its maximum likelihood estimate towards zero, with no shrinkage represented by a value of zero and complete shrinkage represented by a value of one.

The interpretation of κ_j , as defined in (7.27), as a shrinkage factor for β_j relies on the assumption that $\mathbf{X}^T \mathbf{X} \approx N \text{diag}(s_1^2, \dots, s_p^2)$. If $\mathbf{X}^T \mathbf{X}$ is not diagonal, the conditional posterior mean $\bar{\beta}_j$ no longer depends only on the least squares estimate $\hat{\beta}_j$. Thus, a closed form, interpretable expression relating $\bar{\beta}_j$ to $\hat{\beta}_j$ is difficult to obtain, even under the assumption of a very simple non-diagonal structure for $\mathbf{X}^T \mathbf{X}$. However, we note that it is always possible to make the covariates uncorrelated by replacing them with the set of p principal components.

Shrinkage factor matrix

Now we extend the shrinkage factor to a shrinkage factor matrix by considering the horse-shoe prior in (7.19) for multiple linear regression with a multivariate response vector, which is defined in (7.6) and in its matrix form in (7.7). The rows of \mathbf{Y} and \mathbf{B} can be stacked into NK - and Np - vectors, respectively, to give us

$$\begin{aligned} \mathbf{y}^* &= \text{vec}(\mathbf{Y}^T) = (y_{11}, \dots, y_{1K}, y_{21}, \dots, y_{NK})^T \text{ and} \\ \boldsymbol{\beta}^* &= \text{vec}(\mathbf{B}^T) = (\beta_{11}, \dots, \beta_{1K}, \beta_{21}, \dots, \beta_{pK})^T. \end{aligned}$$

It can be shown (see Appendix A) that the conditional posterior for the regression coefficients $\boldsymbol{\beta}^*$ given the shrinkage parameters, $\Lambda^* = \text{diag}(\lambda_{11}, \dots, \lambda_{1K}, \lambda_{21}, \dots, \lambda_{pK})$ and τ_β , and the error variance Σ , is given by

$$\boldsymbol{\beta}^* | \Lambda^*, \tau_\beta, \Sigma, \mathbf{y}^* \sim N_{pK}(\mathbf{m}^*, V^*) \quad (7.28)$$

where

$$\begin{aligned} \mathbf{m}^* &= \tau_\beta^2 \Lambda^* \left[\tau_\beta^2 \Lambda^* + \left\{ (\mathbf{X}^T \mathbf{X})^{-1} \otimes \Sigma \right\} \right]^{-1} \hat{\boldsymbol{\beta}}^* \\ V^* &= \left\{ \tau_\beta^{-2} \Lambda^{*-1} + (\mathbf{X}^T \mathbf{X} \otimes \Sigma^{-1}) \right\}^{-1} \end{aligned}$$

in which

$$\hat{\boldsymbol{\beta}}^* = \text{vec}(\hat{\mathbf{B}}^T) = \left\{ (\mathbf{X}^T \mathbf{X})^{-1} \mathbf{X}^T \otimes \mathbf{I}_K \right\} \mathbf{y}^*$$

and $\hat{\mathbf{B}}$ is the least squares estimator of \mathbf{B} .

Under the assumptions that the explanatory variables are uncorrelated with zero mean and sample variances s_1^2, \dots, s_p^2 , $\mathbf{X}^T \mathbf{X} \simeq N \text{diag}(s_1^2, \dots, s_p^2)$ and so

$$\begin{aligned} & \tau_\beta^2 \Lambda^* (\tau_\beta^2 \Lambda^* + \{ (\mathbf{X}^T \mathbf{X})^{-1} \otimes \Sigma \})^{-1} \\ &= \text{blockdiag} \left\{ \tau_\beta^2 \Lambda_1 \left(\tau_\beta^2 \Lambda_1 + \frac{1}{N s_1^2} \Sigma \right)^{-1}, \dots, \tau_\beta^2 \Lambda_p \left(\tau_\beta^2 \Lambda_p + \frac{1}{N s_p^2} \Sigma \right)^{-1} \right\} \\ &= \text{blockdiag} (\mathbf{I}_K - \mathcal{K}_1, \dots, \mathbf{I}_K - \mathcal{K}_p), \end{aligned}$$

where $\Lambda_j = \text{diag}(\lambda_{j1}^2, \dots, \lambda_{jK}^2)$ is the j -th diagonal block of Λ^* and

$$\mathcal{K}_j = (\mathbf{I}_K + N s_j^2 \tau_\beta^2 \Lambda_j \Sigma^{-1})^{-1}, \quad j = 1, \dots, p. \quad (7.29)$$

It then follows that

$$\mathbf{m}^* = \text{blockdiag} (\mathbf{I}_K - \mathcal{K}_1, \dots, \mathbf{I}_K - \mathcal{K}_p) \hat{\boldsymbol{\beta}}^* = \text{vec}(\mathbf{M}^T),$$

where $\mathbf{M} = (m_{jk}) = \text{E}(\mathbf{B} | \Lambda^*, \tau_\beta, \Sigma, \mathbf{Y})$.

If we define $\boldsymbol{\beta}_j = (\beta_{j1}, \dots, \beta_{jK})^T$, $\hat{\boldsymbol{\beta}}_j = (\hat{\beta}_{j1}, \dots, \hat{\beta}_{jK})^T$ and $\mathbf{m}_j = (m_{j1}, \dots, m_{jK})^T$ for $j = 1, \dots, p$ as the (transposed) rows of \mathbf{B} , $\hat{\mathbf{B}}$ and \mathbf{M} , respectively, then it can be seen that

$$\mathbf{m}_j = (\mathbf{I}_K - \mathcal{K}_j) \hat{\boldsymbol{\beta}}_j, \quad j = 1, \dots, p$$

and so the posterior mean \mathbf{M}^T of \mathbf{B}^T can be constructed column-wise. The j^{th} column, corresponding to the coefficients of covariate j in the linear predictors of Y_1 through Y_K , is a linear transformation of column j of the (transposed) least squares estimator $\hat{\mathbf{B}}^T$.

Since Λ_j and Σ are real and positive definite, the eigenvalues, η_1, \dots, η_K , of $ns_j^2\tau_\beta^2\Lambda_j\Sigma^{-1}$ must be real and positive. The eigenvalues of $\mathcal{K}_j^{-1} = \mathbf{I}_K + Ns_j^2\tau_\beta^2\Lambda_j\Sigma^{-1}$ are therefore $1 + \eta_j > 1$ for $j = 1, \dots, K$, and hence the eigenvalues of \mathcal{K}_j , $1/(1 + \eta_j)$, must lie between 0 and 1, making it a convergent matrix. Therefore, we can regard \mathcal{K}_j as the *shrinkage factor matrix* for coefficients β_j of covariate j . The extent to which the coefficients β are shrunk towards zero is determined by the size of the eigenvalues of \mathcal{K}_j . Since the η_j are directly proportional to τ_β^2 , as $\tau_\beta \rightarrow 0$, all eigenvalues of \mathcal{K}_j approach 1 and we have $\mathcal{K}_j \rightarrow \mathbf{I}_K$ and hence complete shrinkage. When $\tau_\beta \rightarrow \infty$ all eigenvalues of \mathcal{K}_j approach 0 and we have $\mathcal{K}_j \rightarrow 0_K$ and hence no shrinkage. Hence, we have a result for multiple linear regression with a multivariate response vector analogous to (7.27).

Special cases of Σ

Now we formulate the shrinkage factor matrix for special cases of Σ .

Diagonal

When $\Sigma = \text{diag}(\sigma_1^2, \dots, \sigma_K^2)$, the shrinkage factor matrix in (7.29) reduces to

$$\mathcal{K}_j = \text{diag} \left(\frac{1}{1 + Ns_j^2\tau_\beta^2\lambda_{j1}^2\sigma_1^{-2}}, \dots, \frac{1}{1 + Ns_j^2\tau_\beta^2\lambda_{jK}^2\sigma_K^{-2}} \right), \quad j = 1, \dots, p \quad (7.30)$$

so that each component of the posterior mean can be expressed as a product of a single shrinkage factor and the corresponding element of the least squares estimator

$$m_{jk} = \left(1 - \frac{1}{1 + Ns_j^2\tau_\beta^2\lambda_{jk}^2\sigma_k^{-2}} \right) \hat{\beta}_{jk}, \quad j = 1, \dots, p, \quad k = 1, \dots, K.$$

Note that the elements in \mathcal{K}_j have essentially the same form as κ_j in equation (7.27). Since \mathcal{K}_j is a diagonal matrix, it follows that its eigenvalues are simply its diagonal entries.

Compound symmetric

The compound symmetric variance matrix can be written as $\Sigma = \rho\sigma^2(\mathbf{1}_K\mathbf{1}_K^T + (1-\rho)/\rho\mathbf{I}_K)$, where $\sigma^2 > 0$ and $\mathbf{1}_K$ is a K -vector of 1s. The eigenvalues of $\mathbf{1}_K\mathbf{1}_K^T$ are K and 0 and so the eigenvalues of Σ are $\rho\sigma^2(K + (1-\rho)/\rho)$ and $\sigma^2(1-\rho)$. Therefore, Σ is positive definite if and only if

$$\rho \left(K + \frac{1-\rho}{\rho} \right) > 0 \iff \rho > -\frac{1}{K-1}$$

and

$$1 - \rho > 0 \iff \rho < 1,$$

i.e. $-1/(K-1) < \rho < 1$.

To derive the shrinkage factor matrix in the case of a compound symmetric variance matrix, we need to first calculate the precision matrix Σ^{-1} . This is found by application of the Sherman-Morrison formula, that is

$$(\mathbf{A} + \mathbf{u}\mathbf{v}^T)^{-1} = \mathbf{A}^{-1} - \frac{\mathbf{A}^{-1}\mathbf{u}\mathbf{v}^T\mathbf{A}^{-1}}{1 + \mathbf{v}^T\mathbf{A}^{-1}\mathbf{u}}, \quad (7.31)$$

where $\mathbf{A} \in \mathbb{R}^{K \times K}$ is an invertible square matrix and $\mathbf{u}, \mathbf{v} \in \mathbb{R}^K$ are vectors. Here $\mathbf{u}\mathbf{v}^T$ is the outer product of \mathbf{u} and \mathbf{v} and $(1 + \mathbf{v}^T\mathbf{A}^{-1}\mathbf{u}) \neq 0$. We note that the variance matrix can be written as $\Sigma = \rho\sigma^2\mathbf{1}_K\mathbf{1}_K^T + \sigma^2(1-\rho)\mathbf{I}_K$. Substituting $\mathbf{u} = \rho\sigma^2\mathbf{1}_K$, $\mathbf{v} = \mathbf{1}_K$ and $\mathbf{A} = \sigma^2(1-\rho)\mathbf{I}_K$ into equation (7.31) above, we can show that the precision matrix is

$$\begin{aligned} \Sigma^{-1} &= \frac{1}{\sigma^2(1-\rho)}\mathbf{I}_K - \frac{\left(\frac{1}{\sigma^2(1-\rho)}\right)\left(\frac{\rho\sigma^2}{\sigma^2(1-\rho)}\right)\mathbf{1}_K\mathbf{1}_K^T}{1 + \mathbf{1}_K^T\frac{1}{\sigma^2(1-\rho)}\mathbf{I}_K\rho\sigma^2\mathbf{1}_K} \\ &= \frac{1}{\sigma^2(1-\rho)}\mathbf{I}_K - \frac{\frac{\rho}{\sigma^2(1-\rho)^2}\mathbf{1}_K\mathbf{1}_K^T}{1 + \frac{K\rho}{1-\rho}} \\ &= \frac{1}{\sigma^2(1-\rho)}\mathbf{I}_K - \left(\frac{\rho}{\sigma^2(1-\rho)^2}\right)\left(\frac{1-\rho}{1-\rho+K\rho}\right)\mathbf{1}_K\mathbf{1}_K^T \\ &= \frac{1}{\sigma^2(1-\rho)}\left\{\mathbf{I}_K - \frac{\rho}{1+\rho(K-1)}\mathbf{1}_K\mathbf{1}_K^T\right\}. \end{aligned}$$

Now we can write the shrinkage factor matrix as

$$\begin{aligned} \mathcal{K}_j &= (\mathbf{I}_K + Ns_j^2\tau_\beta^2\Lambda_j\Sigma^{-1})^{-1} \\ &= \left[\mathbf{I}_K + \frac{Ns_j^2\tau_\beta^2}{\sigma^2(1-\rho)}\Lambda_j\left(\mathbf{I}_K - \frac{\rho}{1+\rho(K-1)}\mathbf{1}_K\mathbf{1}_K^T\right)\right]^{-1} \\ &= \left[\mathbf{I}_K + \frac{Ns_j^2\tau_\beta^2}{\sigma^2(1-\rho)}\Lambda_j - \frac{Ns_j^2\tau_\beta^2\rho}{\sigma^2(1-\rho)(1+\rho(K-1))}\Lambda_j\mathbf{1}_K\mathbf{1}_K^T\right]^{-1} \\ &= \left[\left(\mathbf{I}_K + \frac{Ns_j^2\tau_\beta^2}{\sigma^2(1-\rho)}\Lambda_j\right) - \frac{Ns_j^2\tau_\beta^2\rho}{\sigma^2(1-\rho)(1+\rho(K-1))}\boldsymbol{\lambda}_j\mathbf{1}_K^T\right]^{-1}, \end{aligned}$$

which is in the form of the Sherman-Morrison formula in (7.31), where $\boldsymbol{\lambda}_j = (\lambda_{j1}^2, \dots, \lambda_{jK}^2)^T$.

Here we have

$$\begin{aligned} \mathbf{A} &= \mathbf{I}_K + \frac{Ns_j^2\tau_\beta^2}{\sigma^2(1-\rho)}\Lambda_j \\ &= \text{diag}\left(1 + \frac{Ns_j^2\tau_\beta^2}{\sigma^2(1-\rho)}\lambda_{j1}^2, \dots, 1 + \frac{Ns_j^2\tau_\beta^2}{\sigma^2(1-\rho)}\lambda_{jK}^2\right) \end{aligned}$$

$$\begin{aligned}
 &= \text{diag} \left(\frac{\sigma^2(1-\rho) + N s_j^2 \tau_\beta^2 \lambda_{j1}^2}{\sigma^2(1-\rho)}, \dots, \frac{\sigma^2(1-\rho) + N s_j^2 \tau_\beta^2 \lambda_{jK}^2}{\sigma^2(1-\rho)} \right), \\
 \mathbf{u} &= -\frac{N s_j^2 \tau_\beta^2 \rho}{\sigma^2(1-\rho)(1+\rho(K-1))} \boldsymbol{\lambda}_j \text{ and} \\
 \mathbf{v} &= \mathbf{1}_K.
 \end{aligned}$$

Thus, we have

$$\mathbf{A}^{-1} = \sigma^2(1-\rho) \text{diag} \left\{ \frac{1}{\sigma^2(1-\rho) + N s_j^2 \tau_\beta^2 \lambda_{j1}^2}, \dots, \frac{1}{\sigma^2(1-\rho) + N s_j^2 \tau_\beta^2 \lambda_{jK}^2} \right\}, \quad (7.32)$$

$$\begin{aligned}
 \mathbf{u} \mathbf{v}^T \mathbf{A}^{-1} &= -\frac{N s_j^2 \tau_\beta^2 \rho}{(1+\rho(K-1))} \boldsymbol{\lambda}_j \mathbf{1}_K^T \\
 &\quad \times \text{diag} \left\{ \frac{1}{\sigma^2(1-\rho) + N s_j^2 \tau_\beta^2 \lambda_{j1}^2}, \dots, \frac{1}{\sigma^2(1-\rho) + N s_j^2 \tau_\beta^2 \lambda_{jK}^2} \right\} \quad (7.33)
 \end{aligned}$$

and

$$\begin{aligned}
 &1 + \mathbf{v}^T \mathbf{A}^{-1} \mathbf{u} \\
 &= 1 - \frac{N s_j^2 \tau_\beta^2 \rho}{(1+\rho(K-1))} \mathbf{1}_K^T \text{diag} \left\{ \frac{1}{\sigma^2(1-\rho) + N s_j^2 \tau_\beta^2 \lambda_{j1}^2}, \dots, \frac{1}{\sigma^2(1-\rho) + N s_j^2 \tau_\beta^2 \lambda_{jK}^2} \right\} \boldsymbol{\lambda}_j \\
 &= 1 - \frac{N s_j^2 \tau_\beta^2 \rho}{(1+\rho(K-1))} \mathbf{1}_K^T \left(\frac{\lambda_{j1}^2}{\sigma^2(1-\rho) + N s_j^2 \tau_\beta^2 \lambda_{j1}^2}, \dots, \frac{\lambda_{jK}^2}{\sigma^2(1-\rho) + N s_j^2 \tau_\beta^2 \lambda_{jK}^2} \right)^T \\
 &= 1 - \frac{N s_j^2 \tau_\beta^2 \rho}{(1+\rho(K-1))} \sum_{k=1}^K \frac{\lambda_{jk}^2}{\sigma^2(1-\rho) + N s_j^2 \tau_\beta^2 \lambda_{jk}^2} \\
 &= \frac{(1+\rho(K-1)) - N s_j^2 \tau_\beta^2 \rho \sum_{k=1}^K \frac{\lambda_{jk}^2}{\sigma^2(1-\rho) + N s_j^2 \tau_\beta^2 \lambda_{jk}^2}}{(1+\rho(K-1))} \quad (7.34)
 \end{aligned}$$

Substituting in the values found in (7.32) to (7.34) into the Sherman-Morrison formula we have

$$\begin{aligned}
 \mathcal{K}_j &= \sigma^2(1-\rho) \text{diag} \left\{ \frac{1}{\sigma^2(1-\rho) + N s_j^2 \tau_\beta^2 \lambda_{j1}^2}, \dots, \frac{1}{\sigma^2(1-\rho) + N s_j^2 \tau_\beta^2 \lambda_{jK}^2} \right\} \\
 &\quad \times \left[\mathbf{I}_K + \frac{N s_j^2 \tau_\beta^2 \rho}{(1+\rho(K-1)) - N s_j^2 \tau_\beta^2 \rho \sum_{k=1}^K \frac{\lambda_{jk}^2}{\sigma^2(1-\rho) + N s_j^2 \tau_\beta^2 \lambda_{jk}^2}} \right] \\
 &\quad \times \text{diag} \left\{ \frac{1}{\sigma^2(1-\rho) + N s_j^2 \tau_\beta^2 \lambda_{j1}^2}, \dots, \frac{1}{\sigma^2(1-\rho) + N s_j^2 \tau_\beta^2 \lambda_{jK}^2} \right\}
 \end{aligned}$$

$$\begin{aligned}
 &= \text{diag} \left\{ \frac{1}{1 + N s_j^2 \tau_\beta^2 \lambda_{j1}^2 (\sigma^2(1 - \rho))^{-1}}, \dots, \frac{1}{1 + N s_j^2 \tau_\beta^2 \lambda_{jK}^2 (\sigma^2(1 - \rho))^{-1}} \right\} \\
 &\quad \times \left[\mathbf{I}_K + \frac{N s_j^2 \tau_\beta^2 \rho}{1 + \rho(K - 1) - N s_j^2 \tau_\beta^2 \rho \sum_{k=1}^K \frac{\lambda_{jk}^2}{\sigma^2(1 - \rho) + N s_j^2 \tau_\beta^2 \lambda_{jk}^2}} \boldsymbol{\lambda}_j \mathbf{1}_K^T (\sigma^2(1 - \rho))^{-1} \right. \\
 &\quad \left. \times \text{diag} \left\{ \frac{1}{1 + N s_j^2 \tau_\beta^2 \lambda_{j1}^2 (\sigma^2(1 - \rho))^{-1}}, \dots, \frac{1}{1 + N s_j^2 \tau_\beta^2 \lambda_{jK}^2 (\sigma^2(1 - \rho))^{-1}} \right\} \right] \\
 &= \text{diag} \left(\frac{1}{1 + d_j^2 \lambda_{j1}^2}, \dots, \frac{1}{1 + d_j^2 \lambda_{jK}^2} \right) \\
 &\quad \times \left\{ \mathbf{I}_K + \frac{\rho d_j^2}{1 - \rho + \rho \sum_{k=1}^K \frac{1}{1 + d_j^2 \lambda_{jk}^2}} \boldsymbol{\lambda}_j \mathbf{1}_K^T \text{diag} \left(\frac{1}{1 + d_j^2 \lambda_{j1}^2}, \dots, \frac{1}{1 + d_j^2 \lambda_{jK}^2} \right) \right\},
 \end{aligned}$$

where $d_j^2 = N s_j^2 \tau_\beta^2 (\sigma^2(1 - \rho))^{-1}$ and $\boldsymbol{\lambda}_j = (\lambda_{j1}^2, \dots, \lambda_{jK}^2)$. Therefore, the (k, ℓ) -th element of \mathcal{K}_j is

$$\mathcal{K}_{j,k\ell} = \begin{cases} \frac{1}{1 + d_j^2 \lambda_{jk}^2} \left\{ 1 + \frac{\rho d_j^2 \lambda_{jk}^2}{\left(1 - \rho + \rho \sum_{m=1}^K \frac{1}{1 + d_j^2 \lambda_{jm}^2} \right) (1 + d_j^2 \lambda_{jk}^2)} \right\}, & \text{if } k = \ell, \\ \frac{\rho d_j^2 \lambda_{jk}^2}{\left(1 - \rho + \rho \sum_{m=1}^K \frac{1}{1 + d_j^2 \lambda_{jm}^2} \right) (1 + d_j^2 \lambda_{jk}^2) (1 + d_j^2 \lambda_{j\ell}^2)}, & \text{otherwise.} \end{cases}$$

We see that \mathcal{K}_j is a rank-one update of a diagonal matrix. Thus, there is no closed form solution for the eigenvalues. As $\rho \rightarrow 0$, we obtain \mathcal{K}_j when the error variance matrix is diagonal, shown in (7.30), and as $\rho \rightarrow 1$ we get

$$\mathcal{K}_j = \frac{1}{N s_j^2 \tau_\beta^2 \sigma^{-2} + \sum_{k=1}^K \frac{1}{\lambda_{jk}^2}} \begin{pmatrix} 1/\lambda_{j1}^2 & \dots & 1/\lambda_{jK}^2 \\ 1/\lambda_{j1}^2 & \dots & 1/\lambda_{jK}^2 \\ \vdots & \ddots & \vdots \\ 1/\lambda_{j1}^2 & \dots & 1/\lambda_{jK}^2 \end{pmatrix}$$

which is a rank-one matrix. Its eigenvalues are

$$0 \text{ and } \sum_{k=1}^K (1/\lambda_{jk}^2) / \left\{ N s_j^2 \tau_\beta^2 \sigma^{-2} + \sum_{k=1}^K (1/\lambda_{jk}^2) \right\}.$$

First order autoregressive

Consider the variance matrix of a first order autoregressive time series of length K

$$\Sigma = \sigma^2 \begin{pmatrix} 1 & \rho & \rho^2 & \cdots & \rho^{K-1} \\ \rho & 1 & \rho & \cdots & \rho^{K-2} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ \rho^{K-1} & \rho^{K-2} & \rho^{K-3} & \cdots & 1 \end{pmatrix}$$

for which the stationarity condition is $|\rho| < 1$. The inverse is given by the symmetric, tridiagonal matrix

$$\Sigma^{-1} = \frac{1}{\sigma^2(1-\rho^2)} \begin{pmatrix} 1 & -\rho & 0 & \cdots & 0 & 0 & 0 \\ -\rho & 1+\rho^2 & -\rho & \cdots & 0 & 0 & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots & \vdots \\ 0 & 0 & 0 & \cdots & -\rho & 1+\rho^2 & -\rho \\ 0 & 0 & 0 & \cdots & 0 & -\rho & 1 \end{pmatrix}.$$

Therefore, letting $d_j^2 = N s_j^2 \tau_\beta^2 \{\sigma^2(1-\rho^2)\}^{-1}$, we have

$$\mathcal{K}_j^{-1} = \begin{pmatrix} 1 + d_j^2 \lambda_{j1}^2 & -d_j^2 \lambda_{j1}^2 \rho & 0 & \cdots & 0 & 0 & 0 \\ -d_j^2 \lambda_{j2}^2 \rho & 1 + d_j^2 \lambda_{j2}^2 (1 + \rho^2) & -d_j^2 \lambda_{j2}^2 \rho & \cdots & 0 & 0 & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots & \vdots \\ 0 & 0 & 0 & \cdots & -d_j^2 \lambda_{j,K-1}^2 \rho & 1 + d_j^2 \lambda_{j,K-1}^2 (1 + \rho^2) & -d_j^2 \lambda_{j,K-1}^2 \rho \\ 0 & 0 & 0 & \cdots & 0 & -d_j^2 \lambda_{jK}^2 \rho & 1 + d_j^2 \lambda_{jK}^2 \end{pmatrix}$$

which is tridiagonal. As it is not Toeplitz, there is no closed-form solution for its eigenvalues. A closed-form for the inverse \mathcal{K}_j is available but its hugely complicated structure prohibits further analytic manipulation.

Circular first order autoregressive

This special case is defined in terms of the precision matrix Σ^{-1} which is a symmetric, circulant, tridiagonal matrix taking the form

$$\Sigma^{-1} = \begin{pmatrix} \omega_0 & \omega_1 & 0 & 0 & \cdots & 0 & 0 & 0 & \omega_1 \\ \omega_1 & \omega_0 & \omega_1 & 0 & \cdots & 0 & 0 & 0 & 0 \\ 0 & \omega_1 & \omega_0 & \omega_1 & \cdots & 0 & 0 & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots & \vdots & \vdots \\ 0 & 0 & 0 & 0 & \cdots & 0 & \omega_1 & \omega_0 & \omega_1 \\ \omega_1 & 0 & 0 & 0 & \cdots & 0 & 0 & \omega_1 & \omega_0 \end{pmatrix}. \quad (7.35)$$

This special case is of particular interest, as we look at adopting such a precision matrix for our errors in Section 7.4. In this case, letting $d_j^2 = Ns_j^2\tau_\beta^2\omega_0\omega_1$ and $\tilde{\lambda}_{jk} = d_j\lambda_{jk}$ for $k = 1, \dots, K$ we have

$$\mathcal{K}_j^{-1} = \begin{pmatrix} 1 + \omega_1^{-1}\tilde{\lambda}_{j1}^2 & \omega_0^{-1}\tilde{\lambda}_{j1}^2 & 0 & 0 & \cdots & 0 & 0 & \omega_0^{-1}\tilde{\lambda}_{j1}^2 \\ \omega_0^{-1}\tilde{\lambda}_{j2}^2 & 1 + \omega_1^{-1}\tilde{\lambda}_{j2}^2 & \omega_0^{-1}\tilde{\lambda}_{j2}^2 & 0 & \cdots & 0 & 0 & 0 \\ 0 & \omega_0^{-1}\tilde{\lambda}_{j3}^2 & 1 + \omega_1^{-1}\tilde{\lambda}_{j3}^2 & \omega_0^{-1}\tilde{\lambda}_{j3}^2 & \cdots & 0 & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots & \vdots \\ 0 & 0 & 0 & 0 & \cdots & 0 & \omega_0^{-1}\tilde{\lambda}_{j,K-1}^2 & 1 + \omega_1^{-1}\tilde{\lambda}_{j,K-1}^2 & \omega_0^{-1}\tilde{\lambda}_{j,K-1}^2 \\ \omega_0^{-1}\tilde{\lambda}_{jK}^2 & 0 & 0 & 0 & \cdots & 0 & 0 & \omega_0^{-1}\tilde{\lambda}_{jK}^2 & 1 + \omega_1^{-1}\tilde{\lambda}_{jK}^2 \end{pmatrix}$$

which is a tridiagonal matrix with corners. Again, as it is not Toeplitz, there is not a closed-form solution for its eigenvalues. Moreover, there is no closed-form solution for the inverse, \mathcal{K}_j .

Conditional prior

When the error variance matrix $\Sigma = \text{diag}(\sigma_1^2, \dots, \sigma_K^2)$ is diagonal, we saw in the previous section that the shrinkage factor matrix \mathcal{K}_j is diagonal, with the k -th diagonal element given by

$$\mathcal{K}_{j,kk} = \frac{1}{1 + N\sigma_k^{-2}\tau_\beta^2 s_j^2 \lambda_{jk}^2}, \quad k = 1, \dots, K$$

which depends only on the local shrinkage parameter λ_{jk} . Conditional on Σ and the global shrinkage parameter τ_β , it therefore follows from the results in the univariate case (Piironen & Vehtari, 2017) that the diagonal elements of \mathcal{K}_j are independent *a priori* with horseshoe-shaped densities given by

$$\frac{1}{\pi} \frac{a_{jk}}{(a_{jk}^2 - 1)\mathcal{K}_{j,kk} + 1} \frac{1}{\sqrt{\mathcal{K}_{j,kk}} \sqrt{1 - \mathcal{K}_{j,kk}}},$$

where $a_{jk} = \tau_\beta \sigma_k^{-1} \sqrt{N} s_j$. Now we consider the non-trivial case when Σ is not diagonal. Conditional on Σ and τ_β , the mapping from K -dimensional $\boldsymbol{\lambda}_j$ to K^2 -dimensional \mathcal{K}_j is dimension-increasing, and so \mathcal{K}_j must lie on a K -dimensional manifold of \mathbb{R}^{K^2} . As such, the independent, unit-median half Cauchy distributions for the elements of $\boldsymbol{\lambda}_j$ induce a joint distribution for \mathcal{K}_j for which a density function does not exist. However, for some parametric forms of Σ , it may, in principle, be possible to derive analytic expressions for the marginal densities for the $\mathcal{K}_{j,kl}$ along with pairwise joint densities. Unfortunately, for the special cases mentioned above, we cannot find a closed-form solution. Similarly, although the mapping from K -dimensional $\boldsymbol{\lambda}_j$ to the K eigenvalues of \mathcal{K}_j is dimension-preserving, we are unable to derive analytic expressions for the joint density of the eigenvalues for any of the previously considered special cases, due to closed-form solutions for the eigenvalues

being unavailable. Instead, we therefore explore the marginal and pairwise joint densities for the elements of the $\mathcal{K}_{j,kl}$ by simulation.

For the simplest non-diagonal special case of compound symmetry, the marginal distributions for the diagonal elements of \mathcal{K}_j will be identical and likewise for its off-diagonal elements. Figure 7.4 shows kernel density estimates of the marginal densities for a diagonal and off-diagonal element when $K \in \{2, 8, 14\}$, $\sigma = 1$ and the correlation parameter ρ takes a range of values:

$$\begin{aligned} K = 2 : & \quad - 0.50 \ 0.00 \ 0.50 \ 0.95 \\ K = 8 : & \quad - 0.12 \ 0.00 \ 0.50 \ 0.95 \\ K = 14 : & \quad - 0.07 \ 0.00 \ 0.50 \ 0.95. \end{aligned}$$

The negative correlation is close to its smallest permissible value for each K . In all cases, we take $N = s_j = 1$ and $\tau_\beta = 1$.

There are four distinct joint densities, between: a pair of diagonal elements; a pair of off-diagonal elements; a diagonal and off-diagonal element within the same row; and a diagonal and off-diagonal element in different rows. For each value of K and ρ , scatter plots of draws from the joint density are shown in Figure 7.4. For every combination of dimension K and correlation ρ , it is evident from the marginal and bivariate plots, that the diagonal elements of \mathcal{K}_j retain the horseshoe shape from the univariate ($K = 1$) case, with most of the density for the off-diagonal elements lying near the mode at 0. The horseshoe shape is evident in the marginal density plots of the diagonal elements. In the bivariate density plots of two diagonal elements, there are fewer points for values not near 0 or 1. Moreover, it seems that the off-diagonal elements are more likely to be close to zero, if the diagonal element in any particular row is close to one. For example, we can see in Figure 7.4 for all combinations of K and ρ that when the diagonal element is close to 1, there is no density at values of the off-diagonal element greater than 0. When the diagonal element is close to 0, the off-diagonal element can take a value anywhere between 0 and ρ . This suggests that there is a high prior probability for values of \mathcal{K}_j lying in the vicinity of diagonal matrices with ones and zeros on the diagonal.

Effective number of non-zero coefficients

In the case of univariate linear regression, we have a single shrinkage factor κ_j for each regression coefficient β_j . The prior for the shrinkage factors can be derived in closed-form and is bimodal, with most of the density concentrated around its two modes at zero and one. Therefore *a priori*, the shrinkage factors are either close to zero or close to one, and

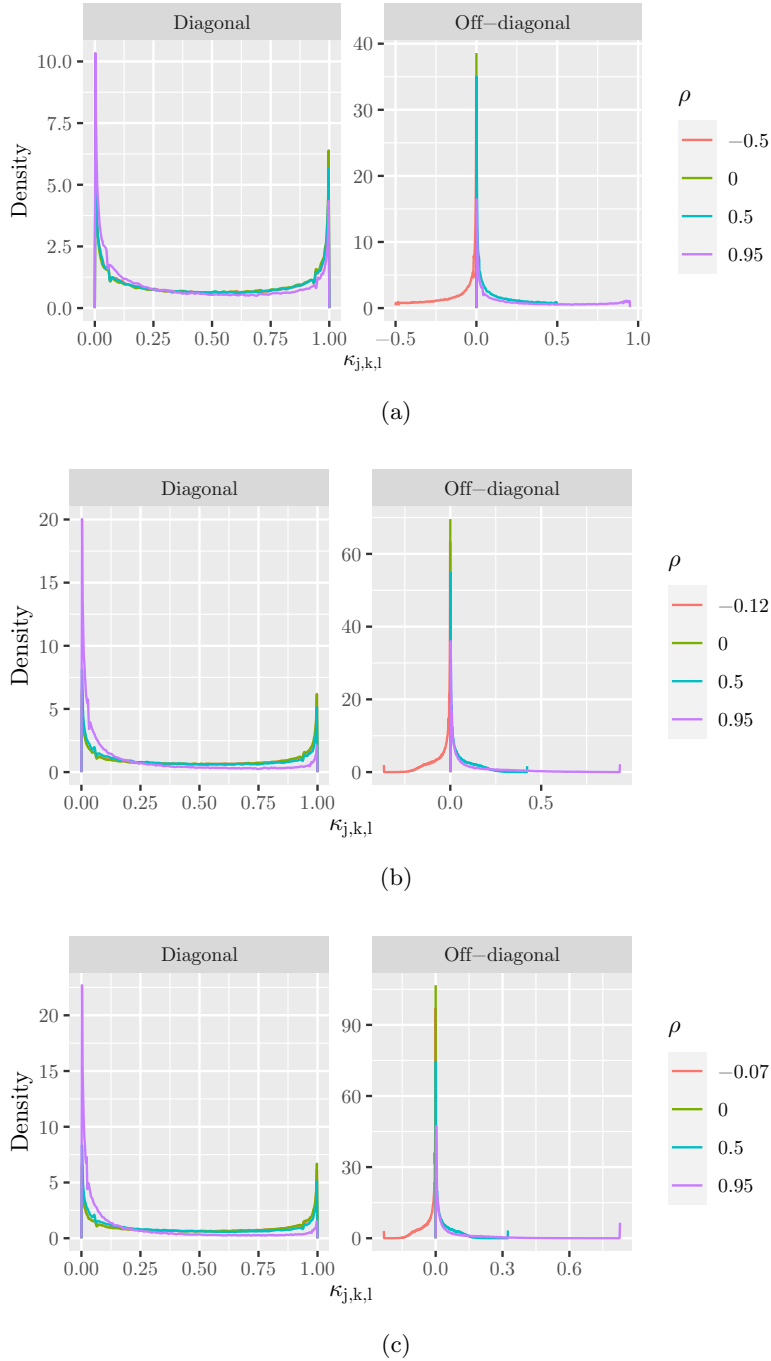
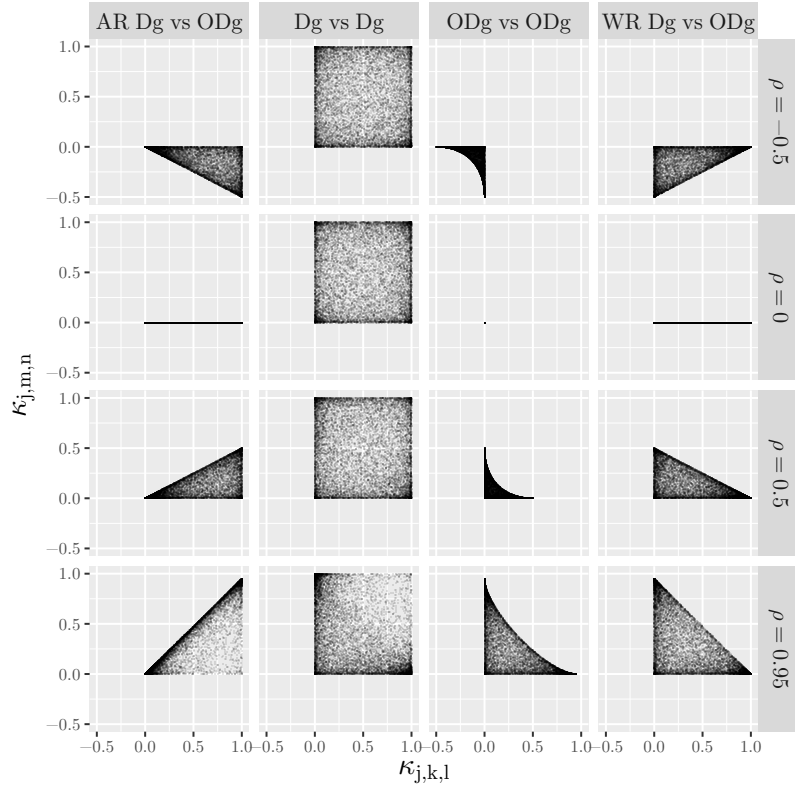
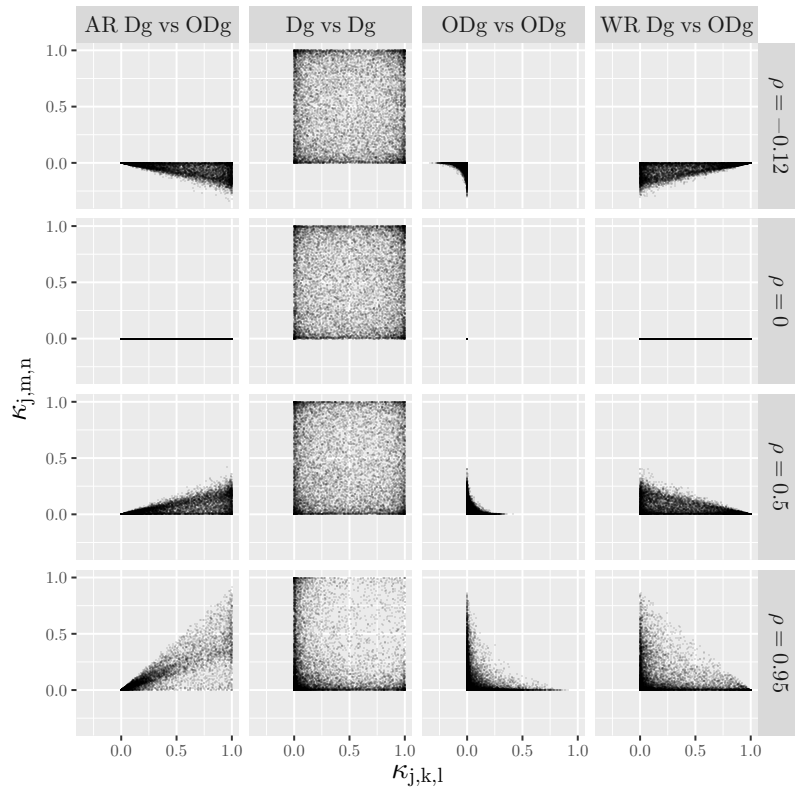


Figure 7.4: Conditional on the global shrinkage parameter τ_β and the error variance Σ , marginal densities of a diagonal and non-diagonal element of \mathcal{K}_j when Σ has compound symmetric structure and the dimension of the response vector is (a) $K = 2$, (b) $K = 8$ and (c) $K = 14$. In each case, $\sigma = 1$, $N = s_j = 1$ and $\tau_\beta = 1$. A range of values for the correlation ρ is considered.



(a)



(b)

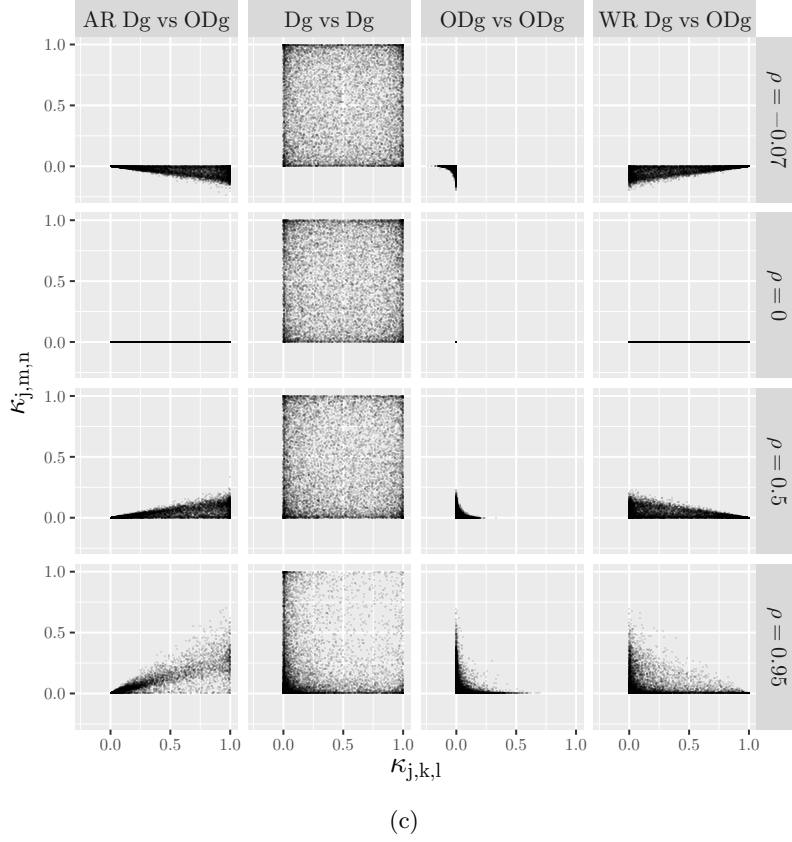


Figure 7.4: Conditional on the global shrinkage parameter τ_β and error variance Σ , pairwise bivariate densities between a pair of diagonal elements of \mathcal{K}_j (Dg vs Dg); a pair of off-diagonal elements (ODg vs ODg); a diagonal and off-diagonal element within the same row (WR Dg vs ODg); and a diagonal and off-diagonal element in different rows (AR Dg vs ODg) when Σ has compound symmetric structure and the dimension of the response vector is (a) $K = 2$, (b) $K = 8$ and (c) $K = 14$. In each case, $\sigma = 1$, $N = s_j = 1$ and $\tau_\beta = 1$. A range of values for the common correlation ρ is considered.

so we can loosely interpret

$$m_{\text{eff}} = \sum_{j=1}^p (1 - \kappa_j)$$

as the effective number of non-zero coefficients. As discussed earlier in this section, in the multivariate case with diagonal Σ , the shrinkage factor matrices \mathcal{K}_j are also diagonal, with bimodal priors for the diagonal elements, concentrated around zero or one. Thus, the same logic applies and we can interpret

$$m_{\text{eff}} = \sum_{j=1}^p \text{tr}(\mathbf{I}_K - \mathcal{K}_j) \quad (7.36)$$

as the effective number of non-zero coefficients.

In the case of the general error variance Σ , the (unattainable) joint distribution for \mathcal{K}_j was discussed earlier. In the special case of compound symmetry, we saw that the prior distribution assigned high probability to diagonal binary matrices. Therefore, it seems reasonable, to continue to interpret m_{eff} as defined in equation (7.36), as the effective number of non-zero coefficients.

Expected value

Conditional on τ_β and Σ , the prior expectation of the effective number of non-zero coefficients m_{eff} is given by

$$\begin{aligned} \mathbb{E}_{\Lambda|\tau_\beta, \Sigma}(m_{\text{eff}}) &= \sum_{j=1}^p \mathbb{E}_{\Lambda|\tau_\beta, \Sigma} \{ \text{tr}(\mathbf{I}_K - \mathcal{K}_j) \} \\ &= pK - \sum_{j=1}^p \sum_{k=1}^K \mathbb{E}_{\Lambda|\tau_\beta, \Sigma}(\mathcal{K}_{j,kk}) \end{aligned}$$

where

$$\begin{aligned} \mathbb{E}_{\Lambda|\tau_\beta, \Sigma}(\mathcal{K}_{j,kk}) &= \int \cdots \int \mathcal{K}_{j,kk} \prod_{\ell=1}^K \pi(\lambda_{j\ell} | \tau_\beta, \Sigma) d\lambda_{j1} \cdots, d\lambda_{jK} \\ &= \int \cdots \int \mathcal{K}_{j,kk} \prod_{\ell=1}^K \frac{2}{\pi(1 + \lambda_{j\ell}^2)} d\lambda_{j1} \cdots, d\lambda_{jK} \end{aligned}$$

for $j = 1, \dots, p$. Suppose we have a guess for the effective number of non-zero coefficients, e_0 . If we can solve $e_0 = \mathbb{E}_{\Lambda|\tau_\beta, \Sigma}(m_{\text{eff}})$ for τ_β , this gives us a candidate value, say τ_{β_0} , to take as some central value in the prior for τ_β .

In the simple case where $\Sigma = \text{diag}(\sigma_1^2, \dots, \sigma_K^2)$, the shrinkage factor matrix is given by (7.30) and so

$$\begin{aligned} \mathbb{E}_{\Lambda|\tau_\beta, \Sigma}(\mathcal{K}_{j,kk}) &= \int \cdots \int \frac{1}{1 + N s_j^2 \tau_\beta^2 \lambda_{jk}^2 \sigma_k^{-2}} \prod_{\ell=1}^K \frac{2}{\pi(1 + \lambda_{j\ell}^2)} d\lambda_{j1} \cdots, d\lambda_{jK} \\ &= \int \frac{1}{1 + N s_j^2 \tau_\beta^2 \lambda_{jk}^2 \sigma_k^{-2}} \frac{2}{\pi(1 + \lambda_{jk}^2)} d\lambda_{jk} \\ &= \int \frac{1}{1 + a_{jk}^2 \lambda_{jk}^2} \frac{2}{\pi(1 + \lambda_{jk}^2)} d\lambda_{jk}, \end{aligned}$$

where $a_{jk} = \sqrt{N} s_j \tau_\beta \sigma_k^{-1}$. By expressing the integrand as the sum of its partial fractions,

it can then be shown (see Appendix A) that

$$\mathbb{E}_{\Lambda|\tau_{\beta},\Sigma}(\mathcal{K}_{j,kk}) = \frac{1}{1 + a_{jk}}. \quad (7.37)$$

Therefore,

$$\begin{aligned} \mathbb{E}_{\Lambda|\tau_{\beta},\Sigma}(m_{\text{eff}}) &= pK - \sum_{j=1}^p \sum_{k=1}^K \frac{1}{1 + a_{jk}} \\ &= \sum_{j=1}^p \sum_{k=1}^K \frac{a_{jk}}{1 + a_{jk}} \\ &= \sum_{j=1}^p \sum_{k=1}^K \frac{\sqrt{N}s_j\tau_{\beta}\sigma_k^{-1}}{1 + \sqrt{N}s_j\tau_{\beta}\sigma_k^{-1}} \\ &= p \sum_{k=1}^K \frac{\sqrt{N}\tau_{\beta}\sigma_k^{-1}}{1 + \sqrt{N}\tau_{\beta}\sigma_k^{-1}} \end{aligned} \quad (7.38)$$

where the last line holds if we assume that all of the explanatory variables have been standardised to have variance equal to one, that is, $s_j^2 = 1$ for $j = 1, \dots, p$. A closed-form expression for τ_{β} which solves

$$e_0 = p \sum_{k=1}^K \frac{\sqrt{N}\tau_{\beta}\sigma_k^{-1}}{1 + \sqrt{N}\tau_{\beta}\sigma_k^{-1}}$$

cannot be found. However, if we make the further assumption that the diagonal elements of Σ are all the same, say $\sigma_k^2 = \sigma^2$ for $k = 1, \dots, K$, then we can simplify the expression in (7.38) to obtain

$$\mathbb{E}_{\Lambda|\tau_{\beta},\Sigma}(m_{\text{eff}}) = \frac{\sqrt{N}\tau_{\beta}\sigma^{-1}}{1 + \sqrt{N}\tau_{\beta}\sigma^{-1}}pK.$$

Setting $\mathbb{E}_{\Lambda|\tau_{\beta},\Sigma}(m_{\text{eff}}) = e_0$, $\tau_{\beta} = \tau_{\beta_0}$ and solving for τ_{β_0} gives

$$\tau_{\beta_0} = \frac{e_0}{pK - e_0} \frac{\sigma}{\sqrt{N}} \quad (7.39)$$

as we would expect from the univariate case. We can then choose as our prior for τ_{β}

$$\tau_{\beta}|\sigma \sim \text{C}^+(0, \tau_{\beta_0}^2) \quad (7.40)$$

which has median equal to τ_{β_0} .

When $\Sigma \neq \sigma^2\mathbf{I}_K$, it is not generally possible to compute $\mathbb{E}_{\Lambda|\tau_{\beta},\Sigma}(\mathcal{K}_{j,kk})$ in closed-form. Therefore, we introduce a hyperparameter σ such that, say, $\mathbb{E}(\Sigma|\sigma) = \sigma^2\mathbf{I}_K$ or

$E(\Sigma^{-1}|\sigma) = \sigma^{-2}\mathbf{I}_K$ and construct our prior for $(\tau_\beta, \Sigma, \sigma)$ or $(\tau_\beta, \Sigma^{-1}, \sigma)$ hierarchically so that

$$\pi(\tau_\beta, \Sigma, \sigma) = \pi(\tau_\beta|\sigma)\pi(\Sigma|\sigma)\pi(\sigma) \quad \text{or} \quad \pi(\tau_\beta, \Sigma^{-1}, \sigma) = \pi(\tau_\beta|\sigma)\pi(\Sigma^{-1}|\sigma)\pi(\sigma)$$

with the conditional distribution for $\tau_\beta|\sigma$ specified in (7.40). The distribution that we give $\Sigma|\sigma$ (or $\Sigma^{-1}|\sigma$) will be considered when we specify our error structure in Section 7.4.

7.3.4 Regularised horseshoe prior

As discussed in Section 7.3.2, a major drawback of the “standard” horseshoe prior is that the parameters far away from zero are not regularised, which is particularly problematic when parameters are only weakly identified by the data. Furthermore, there can also be sampling issues with the standard horseshoe prior because of the posterior possessing an extreme funnel shape, which is problematic for MCMC methods (Piironen & Vehtari, 2017). These sampling issues manifest as divergent transitions when using the NUTS algorithm (see Section 2.3.3). To tackle these problems with the horseshoe prior, Piironen & Vehtari (2017) proposed the regularised horseshoe prior, which for the regression coefficients β in (7.3) is

$$\begin{aligned} \beta_j|\lambda_j, \tau_\beta, c &\sim N\left(0, \tau_\beta^2 \tilde{\lambda}_j^2\right), \\ \tilde{\lambda}_j^2 &= \frac{c^2 \lambda_j^2}{c^2 + \tau_\beta^2 \lambda_j^2}, \\ \lambda_j &\sim C^+(0, 1), \quad j = 1, \dots, p, \\ \tau_\beta &\sim C^+(0, \tau_{\beta_0}^2), \end{aligned} \tag{7.41}$$

where $c > 0$ can be fixed or given a prior, for example,

$$c^2 \sim \text{IG}(\alpha, \beta), \quad \alpha = \nu/2, \quad \beta = \nu s^2/2. \tag{7.42}$$

When $\tau_\beta^2 \lambda_j^2 \ll c^2$, which implies that the coefficient β_j is near to zero, then $\tilde{\lambda}_j^2 \rightarrow \lambda_j^2$ and (7.41) approaches the standard horseshoe. However, when $\tau_\beta^2 \lambda_j^2 \gg c^2$, which means the coefficient is far from zero, then $\tilde{\lambda}_j^2 \rightarrow c^2/\tau_\beta^2$ and (7.41) approaches $N(0, c^2)$. Therefore, the prior can shrink irrelevant coefficients as the horseshoe does, whilst also being able to regularise even the largest coefficients as a normal slab with variance c^2 (see Piironen & Vehtari (2017) for further explanation). The inverse-gamma prior for c^2 is a reasonable choice because it possesses both a heavy right tail and a light left tail, which prevents a lot of mass accumulating near zero. For the coefficients far from zero, it translates to a Student- $t_\nu(0, s^2)$.

In the case of multiple linear regression with a multivariate response vector, the regularised horseshoe is given by

$$\begin{aligned}\beta_{jk}|\lambda_j, \tau_\beta, c &\sim N\left(0, \tau_\beta^2 \tilde{\lambda}_{jk}^2\right), \\ \tilde{\lambda}_{jk}^2 &= \frac{c^2 \lambda_{jk}^2}{c^2 + \tau_\beta^2 \lambda_{jk}^2}, \\ \lambda_{jk} &\sim C^+(0, 1), \quad j = 1, \dots, p, \quad k = 1, \dots, K, \\ \tau_\beta &\sim C^+(0, \tau_{\beta_0}^2)\end{aligned}\tag{7.43}$$

where c^2 has the prior defined in (7.42) and τ_{β_0} is defined in (7.39). Extending this to a VAR(1) model gives us

$$\begin{aligned}a_{jk}|\lambda_{jk}, \tau_A, c &\sim N\left(0, \tau_A^2 \tilde{\lambda}_{jk}^2\right), \\ \tilde{\lambda}_{jk}^2 &= \frac{c^2 \lambda_{jk}^2}{c^2 + \tau_A^2 \lambda_{jk}^2}, \\ \lambda_{jk} &\sim C^+(0, 1), \quad j = 1, \dots, K, \quad k = 1, \dots, K, \\ \tau_A &\sim C^+(0, \tau_0^2),\end{aligned}\tag{7.44}$$

$$\tag{7.45}$$

where c^2 has the prior defined in (7.42),

$$\tau_0 = \frac{e_0}{K^2 - e_0} \frac{\sigma}{\sqrt{N}}$$

and e_0 is the expected number of non-zero autoregressive coefficients *a priori*. We fit the model with a regularised horseshoe prior on A in Stan and so we do not need to derive the FCDs.

7.3.5 Comparison of priors via simulation study

In the following section, we simulate some data to test and compare the performance of the shrinkage priors (discussed above) and corresponding inference methods. Additionally, we perform inference with the multivariate normal prior specified in (7.2) to see the effect of choosing a prior that does not encourage shrinkage. We simulate three autoregressive coefficient matrices constrained to the stationary region, with 12 groups ($K = 12$), and varying degrees of sparsity, A_{80} , A_{50} and A_{20} , where the subscripts indicate the rough percentage of zero-coefficients in the matrix. Heatmaps of A_{80} , A_{50} and A_{20} are shown in Figures E.7 to E.9 in Appendix E. Our simulated data have 250 time points so that the dimensions roughly match the dimensions of our clustered WWTP data.

We assume a time invariant mean and adopt the simple diagonal error structure with

common variance $\tau^{-1} = 1$. Hyperparameters for the priors on A are chosen so that the conditional expectations and variances are the same (or similar). All of our priors have a conditional expectation of zero. It can be shown that the conditional variance under the spike and slab prior is $\text{Var}(a_{jk}) = d^2\psi_{jk}$ and under both the standard horseshoe and regularised horseshoe prior is undefined (see Appendix A). The variance under the multivariate normal prior is simply d^2 . Therefore, we choose $\psi_{jk} = 0.5$ for all j, k and $d^2 = 200$ in the spike and slab prior and $d^2 = 100$ in the multivariate normal prior, resulting in reasonably large (conditional) variances of 100. **Stan** is used to fit the majority of models, with the exception of the spike and slab prior, as **Stan** is only suitable for fitting models with continuous valued parameters. Instead, a Gibbs sampler is written in **R** to fit the VAR(1) model with a spike and slab prior for A . The results presented are based on 4000 near-uncorrelated samples from four chains (1000 per chain) initialised at different starting points.

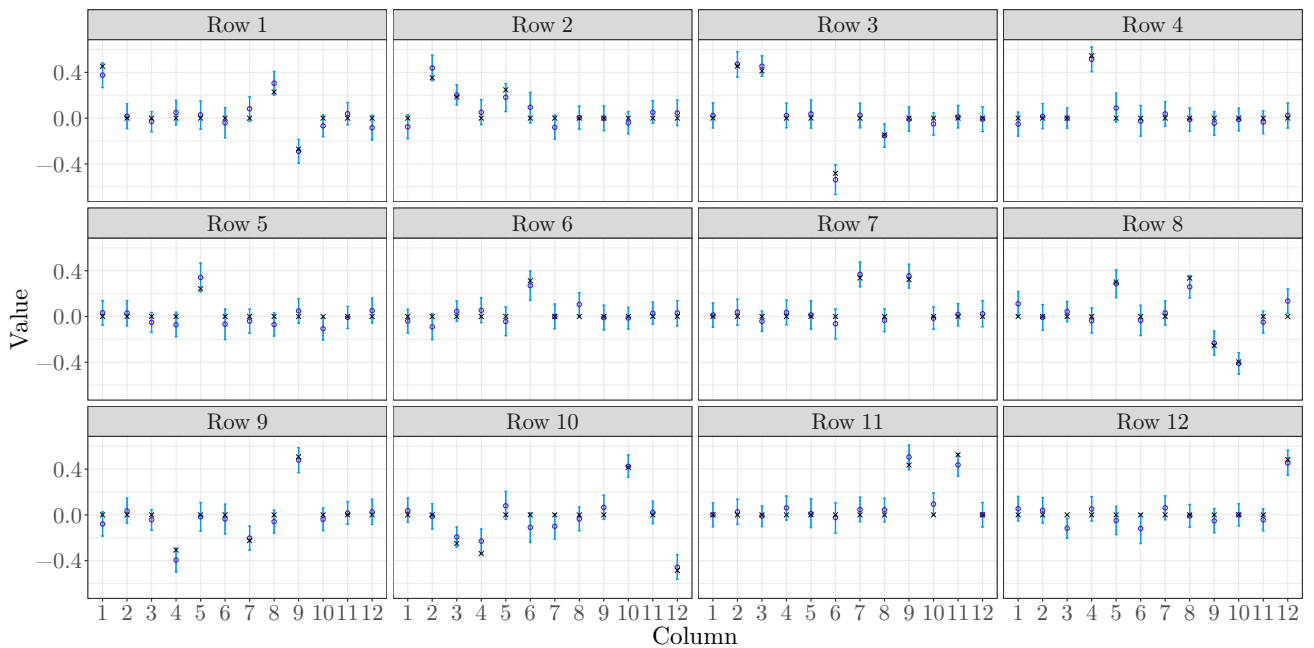


Figure 7.5: Posterior means (\circ) and 95% credible intervals (—) for A_{80} , with the true values (\times), under the multivariate normal prior.

First we look at the results for the data simulated using the most sparse matrix A_{80} . We make comparisons with plots of the posterior means and 95% credible intervals (CIs) for the elements of A_{80} , in the case of each prior. With the multivariate normal prior, we can see in Figure 7.5 that most of the true values for the elements of A_{80} are captured in the 95% CIs. However, for autoregressive coefficients $a_{8,12}$ and $a_{12,3}$, their true values of zero are not within the 95% CI. In addition to these, several other coefficients with true values of zero have CIs that only just include zero, for example, see $a_{10,7}$, $a_{11,10}$ and $a_{12,6}$.

Given that there are a total of 144 parameters to infer, perhaps it is not surprising that the true values of some coefficients are not captured in the CIs.

Figure 7.6 shows the plot for the spike and slab prior for the 80% sparse matrix. Here most of the elements appear to be captured within their CIs, although $a_{2,5}$ does not have its true values contained within its CI. There is a suggestion of poor mixing in the non-zero elements, as $\lambda_{jk} = 1$ for 100% of the posterior samples. This is a suspiciously high percentage that may simply be the result of the sampler getting stuck at $\lambda_{jk} = 1$. The CIs are also extremely narrow for most of the zero elements. For the majority of these elements, λ_{jk} is typically zero in $\sim 99\%$ of the posterior samples, which also indicates poor mixing.

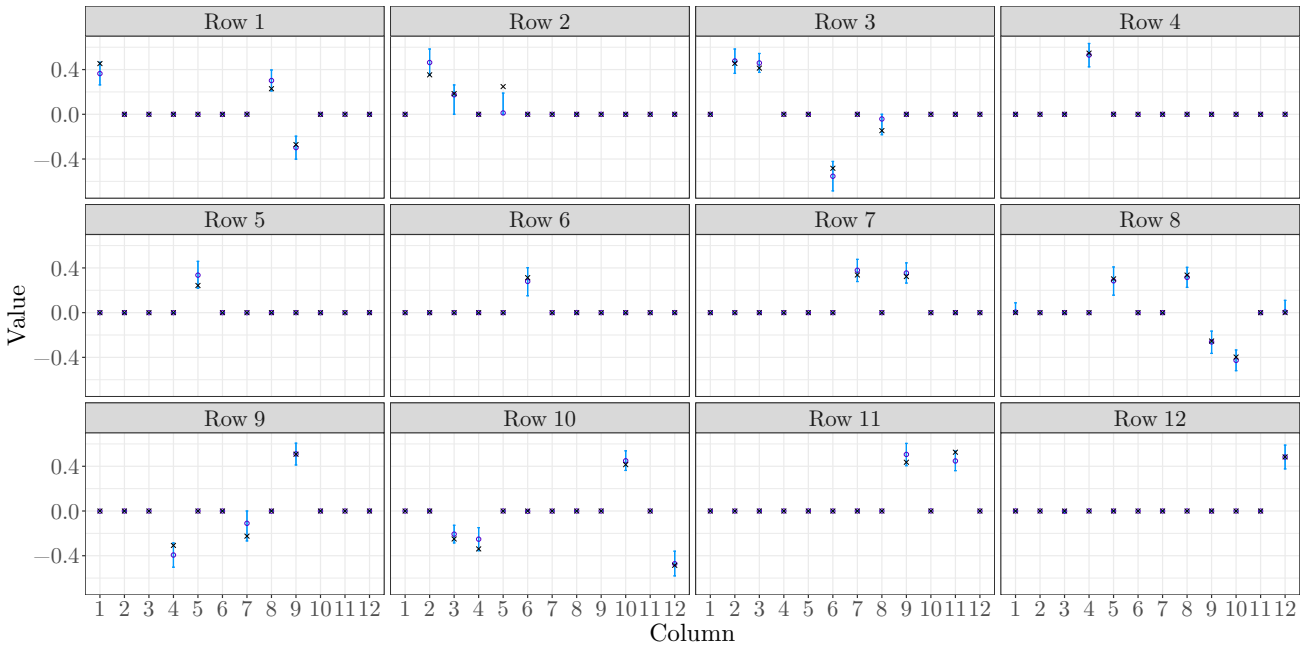


Figure 7.6: Posterior means (\circ) and 95% credible intervals (—) for A_{80} , with the true values (\times), under the spike and slab.

The results for the horseshoe prior are shown in Figure 7.7. All the true values for the autoregressive coefficients in A_{80} have been captured in the CIs, except for one element, $a_{2,5}$, where the true value is just outside of the CI and zero is included. For many of the elements, the posterior means are very close to the true values, indicating a good performance by the horseshoe prior. However, even with the acceptance rate as 0.999 (which significantly reduces the step size; see Section 2.3.3), there were seven divergent transitions. Such a low number of divergent transitions is not overly concerning, but the small step size causes an increase in computational time. This was not an issue for the regularised horseshoe prior, as larger step sizes could be used without yielding any divergent transitions. Figure 7.8 shows the results for the regularised horseshoe, which

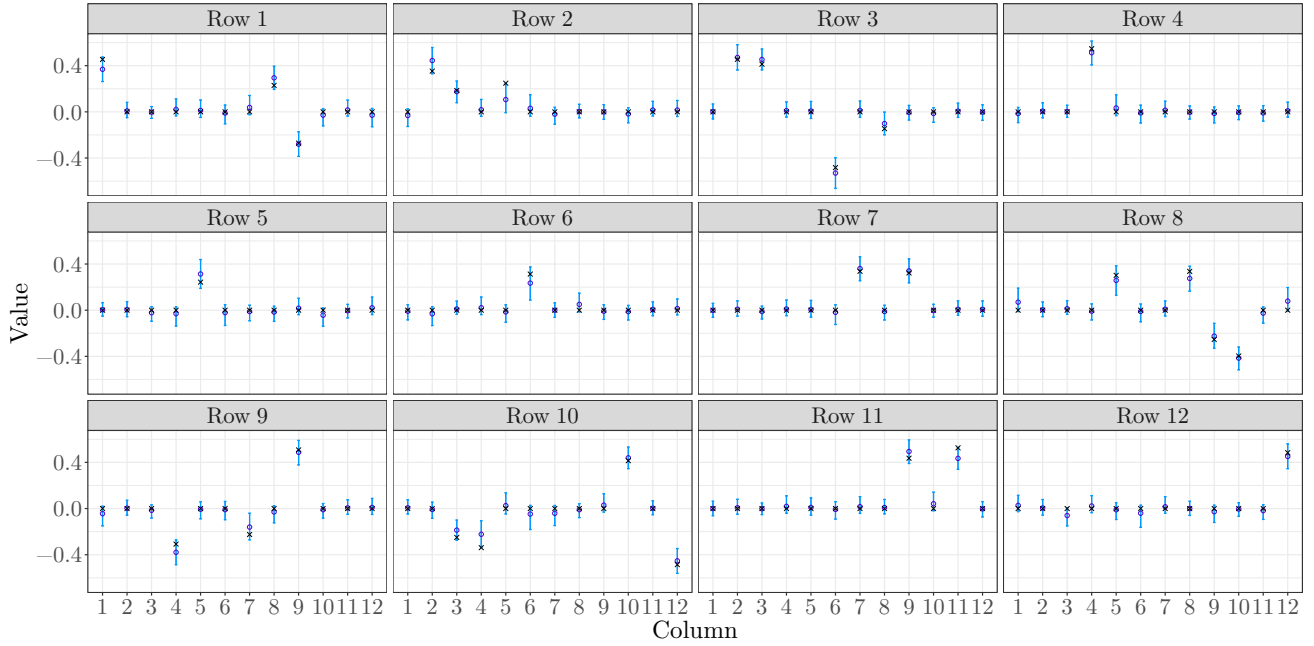


Figure 7.7: Posterior means (\circ) and 95% credible intervals (—) for A_{80} , with the true values (\times), under the horseshoe prior.

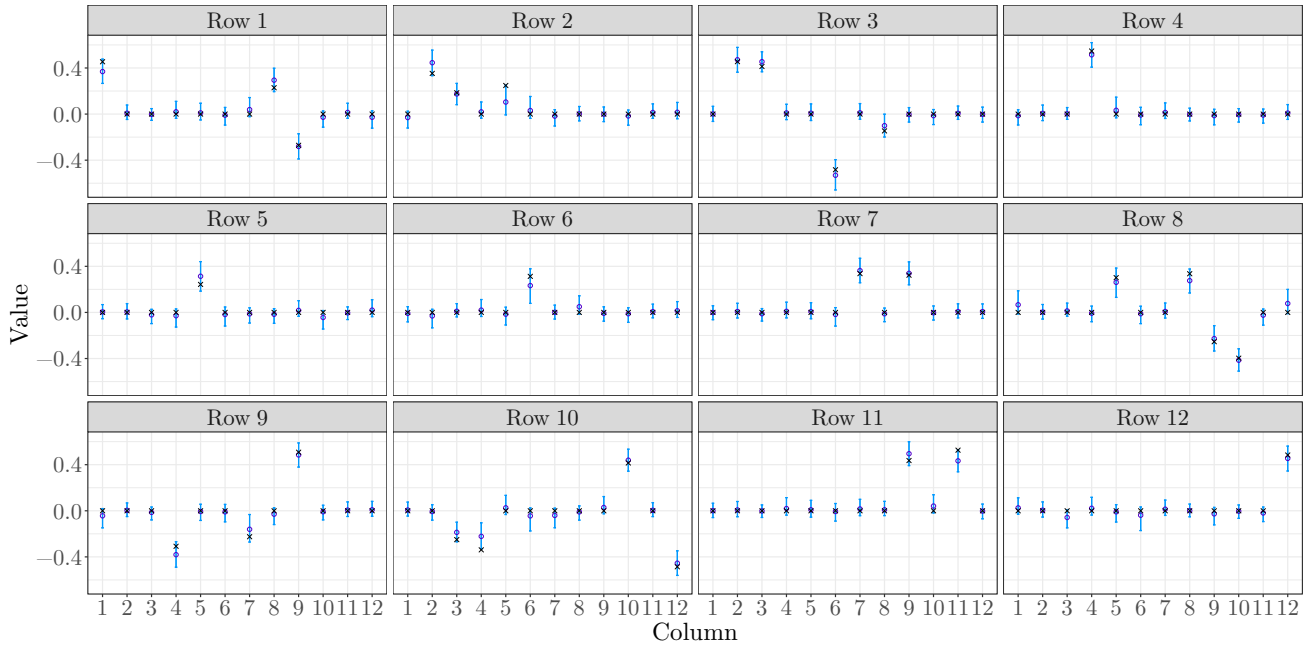


Figure 7.8: Posterior means (\circ) and 95% credible intervals (—) for A_{80} , with the true values (\times), under the regularised horseshoe prior.

are very similar to the results obtained for the standard horseshoe prior. Although not obvious from the plot, the CI for $a_{2,5}$ contains the true value of the element, but zero is

included in the CI too. The obvious advantage here for using the regularised horseshoe instead of the standard horseshoe is computational speed. For example, running the model with the standard horseshoe with a smaller step size takes 26.5 minutes to perform 2000 iterations, whereas the regularised horseshoe with its larger step size (that results in zero divergent transitions) takes 10.5 minutes to perform the same number of iterations, which is more than twice as fast.

Next we compare the results of each prior for A_{50} and begin again with the results of the multivariate normal prior, which are shown in Figure 7.9. Most elements have their true values captured in their posterior 95% CIs but we also note that $a_{1,12}, a_{2,11}$ and $a_{6,2}$ have their true values of zero outside of their CIs. This is most likely an effect of the prior not encouraging shrinkage. Although not obvious from the plot, the true value for $a_{9,2}$ also lies *just* outside of its CI.

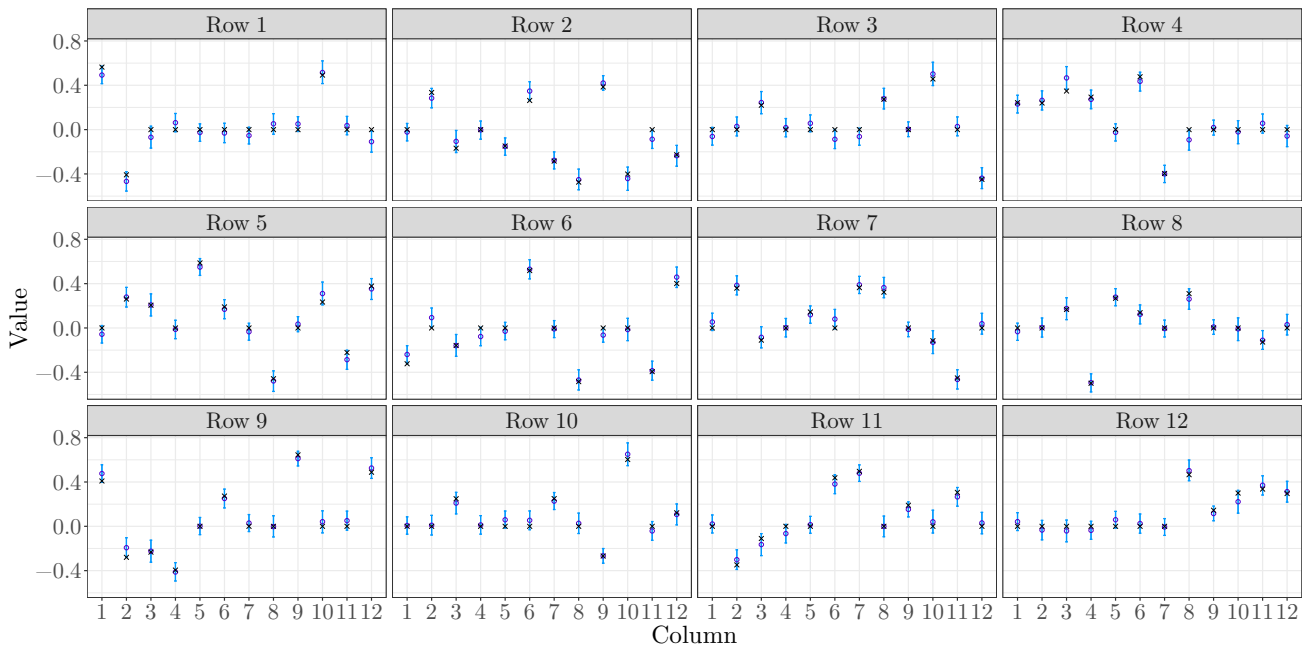


Figure 7.9: Posterior means (\circ) and 95% credible intervals (—) for A_{50} , with the true values (\times), under the multivariate normal prior.

The result for A_{50} under the spike and slab prior are shown in Figure 7.10. The non-zero autoregressive coefficient $a_{11,3}$ contains zero in its CI. Additionally, the true values of $a_{2,3}$ (-0.168), $a_{8,11}$ (-0.127) and $a_{10,12}$ (0.122) are not captured in the posterior CIs, with their posterior means as zero. This could be due to their true values representing weak signals, which seemed to be an issue for Lei *et al.* (2011), as mentioned in Section 7.3.1. Almost 100% of the posterior samples for these elements are zero, for example, $\lambda_{10,12} = 0$ in 100%, 100%, 100% and 98.8% of the posterior samples of each chain, respectively. Again, this may indicate “stickiness” of the sampler and hence poor mixing, as we saw with A_{80} .

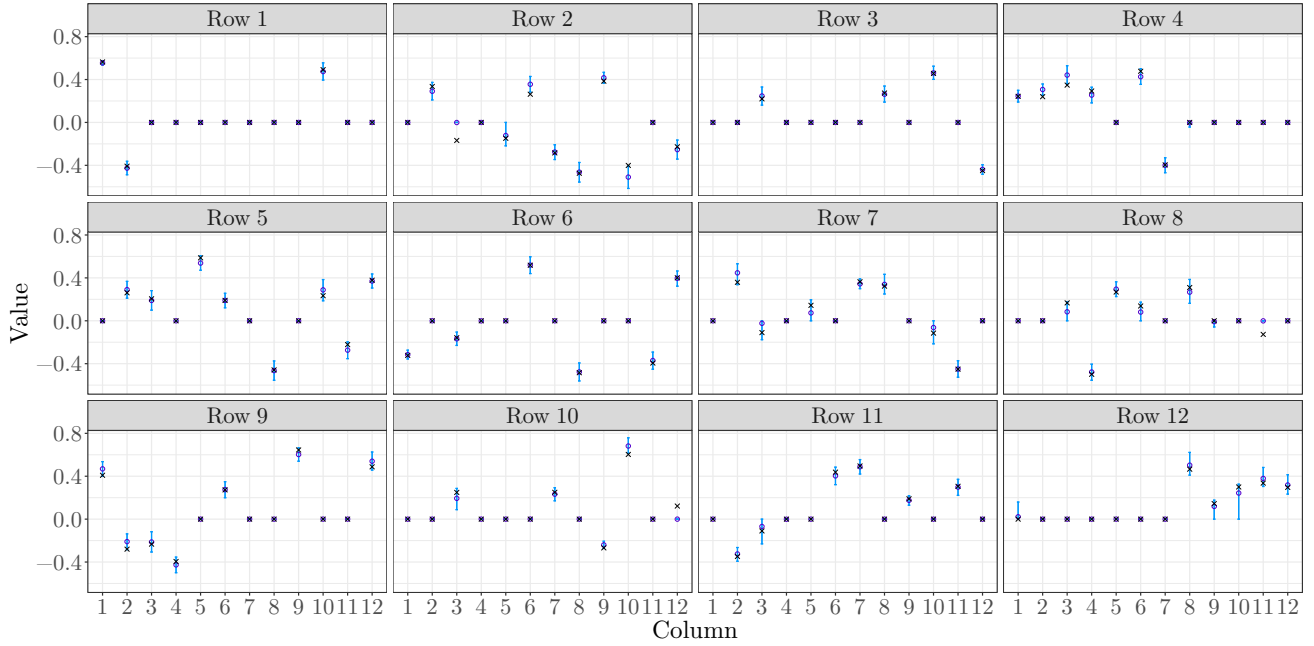


Figure 7.10: Posterior means (\circ) and 95% credible intervals (—) for A_{50} , with the true values (\times), under the spike and slab.

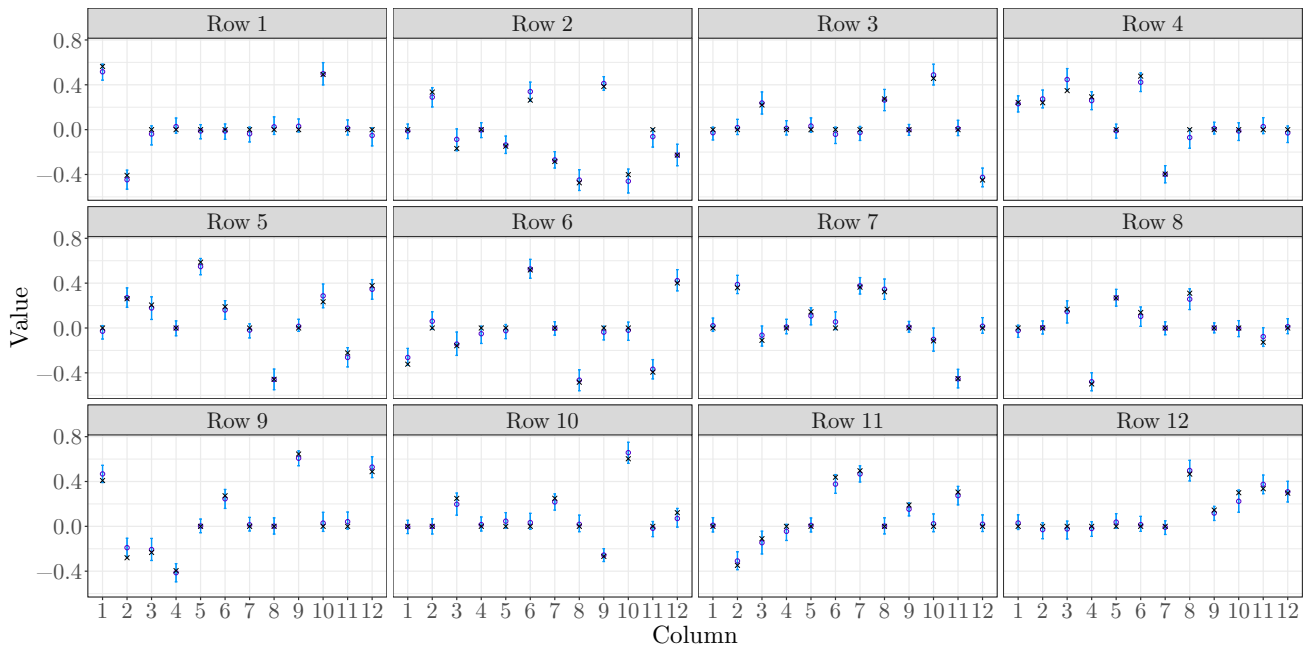


Figure 7.11: Posterior means (\circ) and 95% credible intervals (—) for A_{50} , with the true values (\times), under the horseshoe prior.

Furthermore, many of the other elements of A_{50} suffer from this problem, for example, the majority of the remaining non-zero autoregressive coefficients have their corresponding

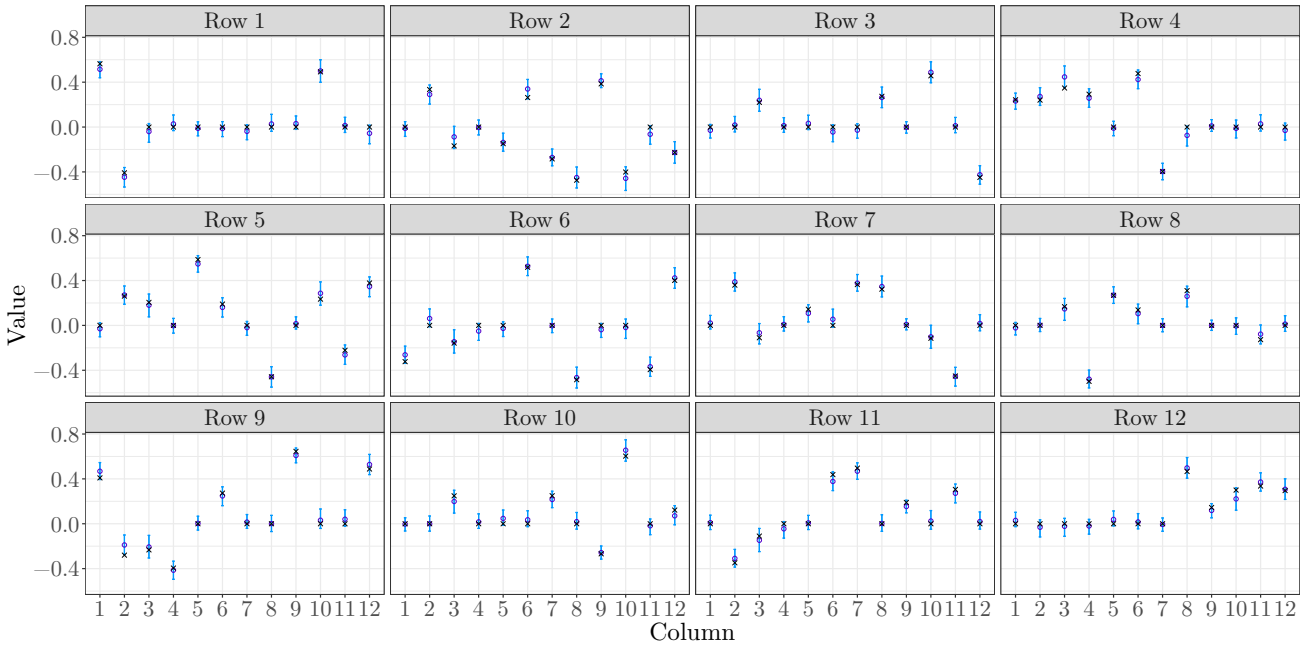


Figure 7.12: Posterior means (\circ) and 95% credible intervals (—) for A_{50} , with the true values (\times), under the regularised horseshoe prior.

$\lambda_{jk} = 1$ for 100% of their posterior samples.

Figure 7.11 shows the results for A_{50} under the standard horseshoe prior. Again, the horseshoe seems to perform reasonably well and captures the true values of the autoregressive coefficients within the CIs, with the exception of $a_{9,2}$ and $a_{4,3}$. The true values of these elements lie just outside of the CIs. Comparing the results for A_{50} under the regularised horseshoe (Figure 7.12) to the standard horseshoe, we can see that it performs similarly to the standard horseshoe. However, the CI for $a_{9,2}$ does contain the true value. Additionally, we do not have any divergent transitions in our posterior samples when using the regularised horseshoe, whereas for the standard horseshoe, there were 41 divergent transitions.

Finally, we discuss the results for the simulation study with the least sparse matrix A_{20} . Figure 7.13 shows the results under the multivariate normal prior. Seven of the autoregressive coefficients ($a_{2,11}$, $a_{3,1}$, $a_{3,3}$, $a_{6,2}$, $a_{8,12}$, $a_{10,4}$ and $a_{10,6}$) do not have their true values in their CIs, although the true values lie just outside of the intervals. The multivariate normal prior only fails to “find” one of the zero-coefficients ($a_{10,4}$) but overall performs fairly well, as we might expect for this particular matrix.

Figure 7.14 shows the results for A_{20} under the spike and slab prior. Clearly from the plots, the true values of many autoregressive coefficients have not been found, for example, $a_{4,6}$, $a_{8,9}$, $a_{8,12}$, $a_{11,4}$ and $a_{11,12}$ are all non-zero coefficients but zero is included in their CIs, meanwhile their true values are not included. Closer inspection reveals that 16 different

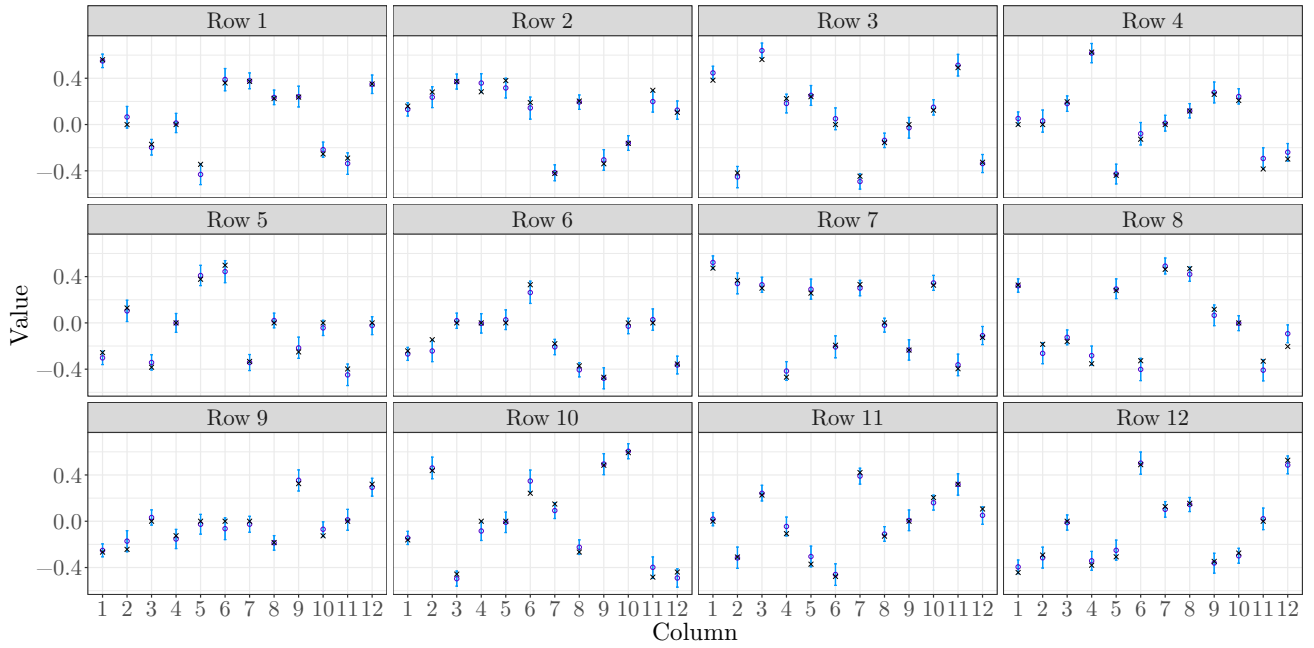


Figure 7.13: Posterior means (\circ) and 95% credible intervals (—) for A_{20} , with the true values (\times), under the multivariate normal prior.

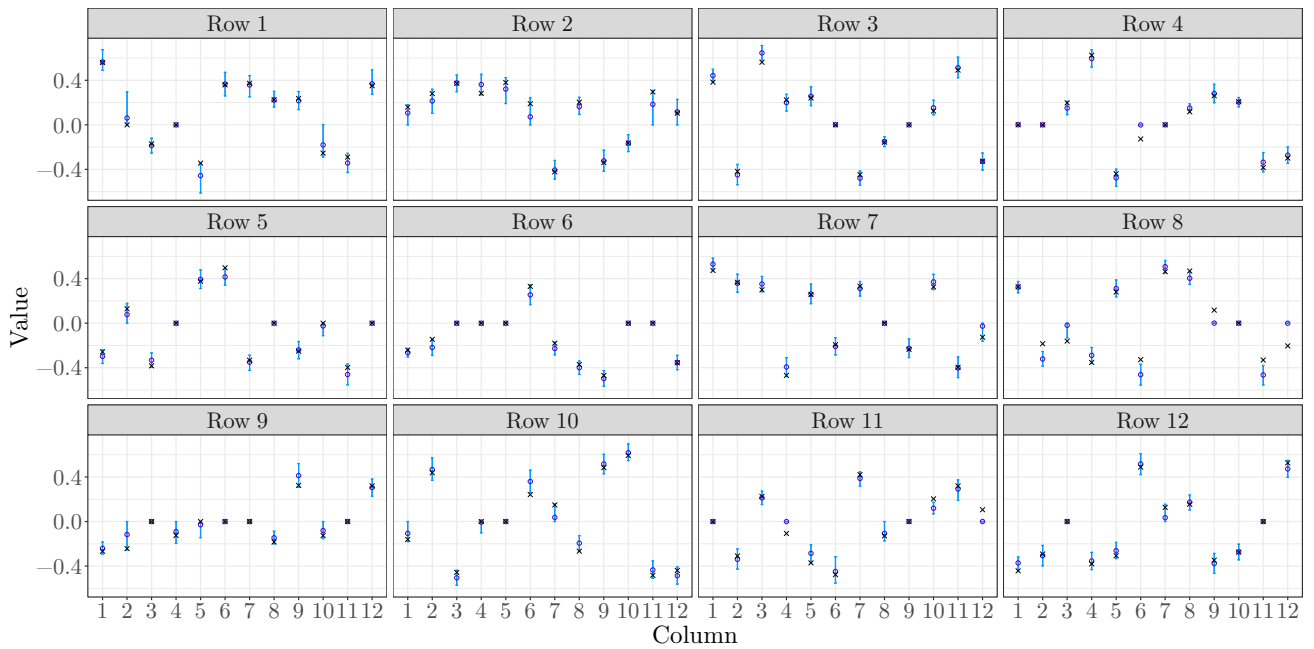


Figure 7.14: Posterior means (\circ) and 95% credible intervals (—) for A_{20} , with the true values (\times), under the spike and slab.

autoregressive coefficients do not contain their true values in their CIs and 11 non-zero autoregressive coefficients include zero in their CIs. This is evidence to suggest that this

prior is not well suited to data that break the assumption of sparsity. Furthermore, mixing issues are also likely to be a problem here, with many of the λ_{jk} taking the same value for 100% of posterior samples in all four chains.

The standard horseshoe seems to perform better than the spike and slab prior when the data are simulated using A_{20} . However, inspecting the posterior means and CIs in Figure 7.15 clearly shows that $a_{8,12}$'s true value of -0.204 is not within the 95% CI. Further investigation reveals that seven other autoregressive coefficients do not have their true values contained in their corresponding CIs and $a_{11,12}$'s CI contains both the true value (0.106) and zero. As we found with the more sparse matrices, there were some divergent transitions (seven) present when using the standard horseshoe prior for A_{20} .

Finally, we discuss the results for A_{20} under the regularised horseshoe prior. As we found with the more sparse autoregressive matrices, performance was similar to the standard horseshoe, although only six autoregressive coefficients did not have the correct value contained in their CIs. An almost identical result was also found for $a_{11,12}$ as the result found under the standard horseshoe, both the true non-zero value and zero were in the CI. For both this prior and the standard horseshoe, where there was evidence to suggest that the true values had not been identified by the model, most true values were very close to the edges of their CIs, which can be seen in Figure 7.16.

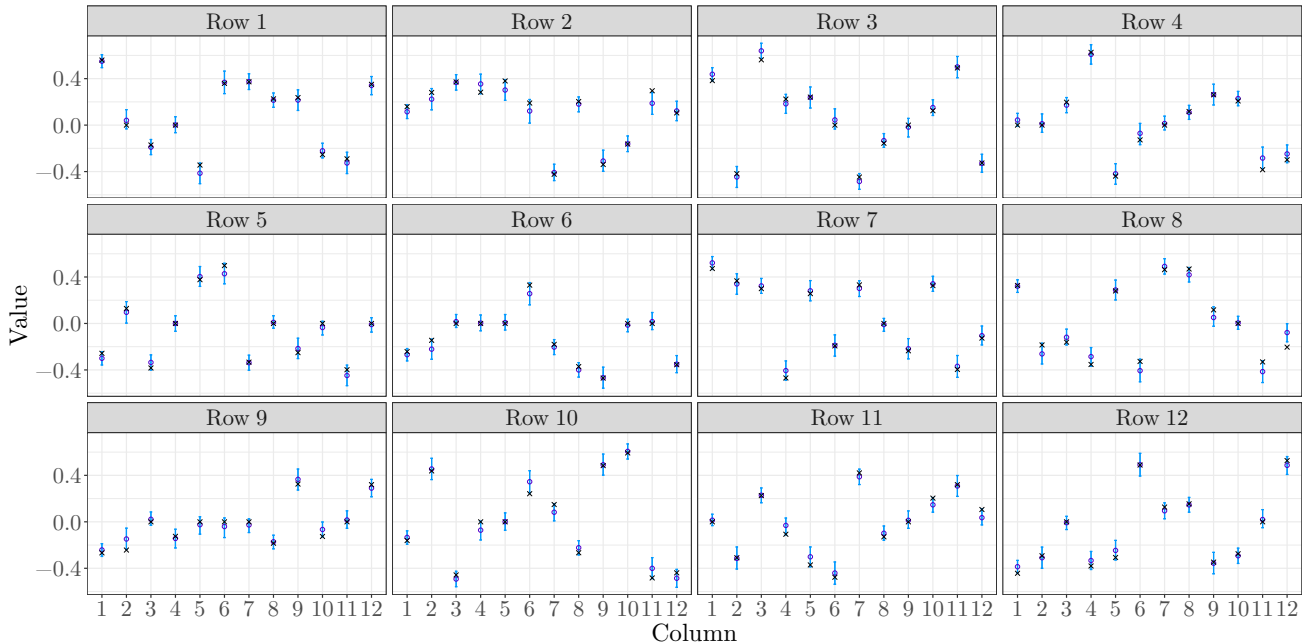


Figure 7.15: Posterior means (\circ) and 95% credible intervals (—) for A_{20} , with the true values (\times), under the horseshoe prior.

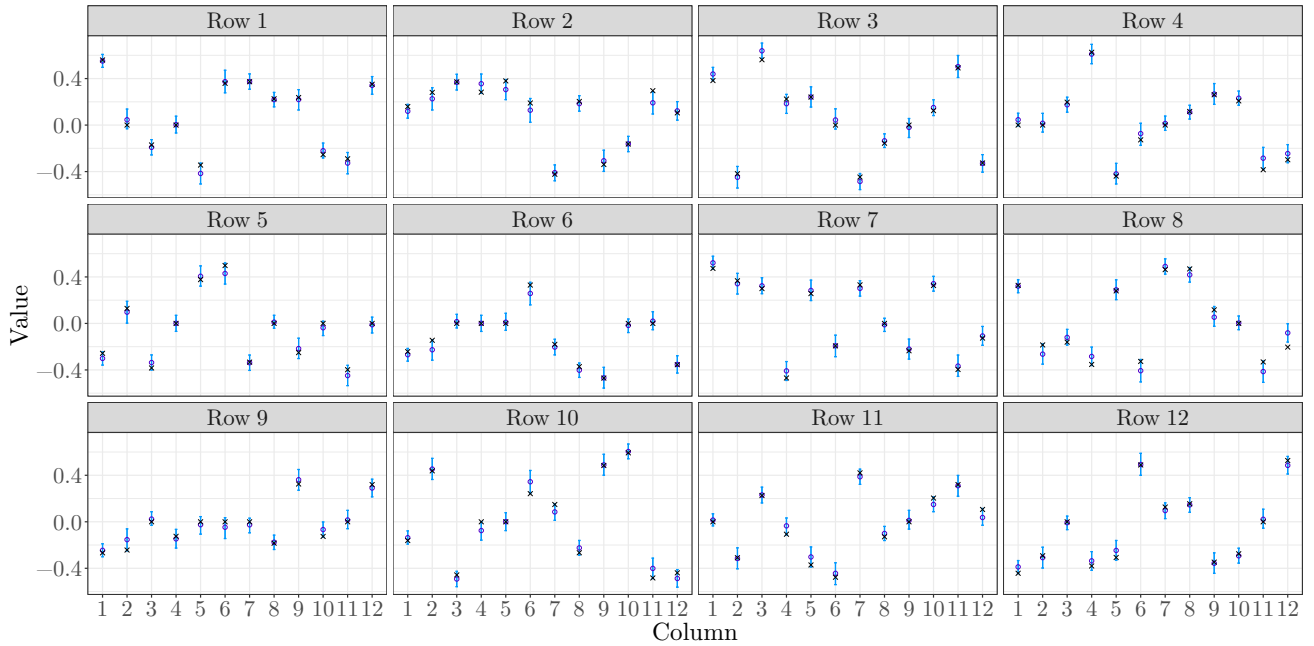


Figure 7.16: Posterior means (\circ) and 95% credible intervals (—) for A_{20} , with the true values (\times), under the regularised horseshoe prior.

Summary

Table 7.1 shows the number of incorrectly identified autoregressive coefficients for each matrix and prior in the simulation study. We consider a coefficient incorrectly identified if its CI does not contain its true value or for a non-zero coefficient, if zero is contained in the CI. Solely based on these results, if we have a matrix of autoregressive coefficients A in a VAR(1) model, where the level of sparsity is unknown (as we have with our WWTP data) then it would seem sensible to choose a regularised horseshoe prior for A . In the case where the data are simulated using A_{50} , the regularised horseshoe outperforms the other priors. For the data simulated with the least sparse matrix A_{20} , the regularised horseshoe performs as well as the multivariate normal. For the most sparse matrix, the spike and slab, standard horseshoe and regularised horseshoe only fail to correctly identify one coefficient $a_{2,5}$. However, the regularised horseshoe does not have the sampling issues that the spike and slab prior has, nor does it suffer from the problem of divergent transitions that the standard horseshoe has. Furthermore, the regularised horseshoe performs computationally faster than the standard horseshoe, because the step size in the NUTS algorithm can be larger (see Section 2.3.3). The obvious choice for the prior for A in our model is the regularised horseshoe because of these benefits and also due to its flexibility with regards to varying degrees of sparseness in the autoregressive matrix.

Therefore, for A we use the prior specified in (7.44) with the hyperprior for the global

	Prior			
	Normal	Spike & Slab	Horseshoe	Regularised Horseshoe
A ₈₀	2	1	1	1
A ₅₀	4	4	2	1
A ₂₀	7	27	8	7

Table 7.1: Number of autoregressive coefficients not correctly identified in the simulation study for each matrix and prior.

shrinkage τ_A specified in (7.45). We set the scale of τ_A to

$$\tau_0 = \frac{e_0}{K^2 - e_0} \frac{\sigma}{\sqrt{N}}, \quad (7.46)$$

where N is the number of the time points, K is the number of bins and e_0 is the estimated number of non-zero autoregressive coefficients. In Section 7.3.3, the definition of the effective number of non-zero coefficients relies on the assumption that the covariates are standardised to have a variance of one. This assumption was made by Piironen & Vehtari (2017) in the case of multiple linear regression with a univariate response vector. Thus, for τ_0 in (7.46) to be a sensible value, we must scale the data to have a variance of one (see Section 7.1.2). We note, however, that the justification for this value also relies on the elements of \mathbf{y}_{t-1} being uncorrelated. Although this will not be satisfied here, specifying τ_0 according to (7.46) is still a principled, pragmatic choice.

7.4 Error structure

Now that we have chosen a suitable prior for the matrix of autoregressive coefficients \mathbf{A} , we can make another alteration to the simple VAR(1) model specified in (7.1). Although convenient, the diagonal error structure with common variance for Σ is perhaps too simple. However, a prior for Σ must also consider our prior for \mathbf{A} , as the scale of the global shrinkage parameter in (7.46) depends on σ . We adopt the same approach for our VAR(1) model as discussed at the end of Section 7.3.3 for the multivariate linear regression case. When $\Sigma \neq \sigma^2 \mathbf{I}_K$, we let σ be a hyperparameter such that, $\text{E}(\Sigma|\sigma) = \sigma^2 \mathbf{I}_K$ or $\text{E}(\Sigma^{-1}|\sigma) = 1/\sigma^2 \mathbf{I}_K$ and construct a prior for (τ_A, Σ, σ) or $(\tau_A, \Sigma^{-1}, \sigma)$ so that

$$\pi(\tau_A, \Sigma, \sigma) = \pi(\tau_A|\sigma) \pi(\Sigma|\sigma) \pi(\sigma) \quad \text{or} \quad \pi(\tau_A, \Sigma^{-1}, \sigma) = \pi(\tau_A|\sigma) \pi(\Sigma^{-1}|\sigma) \pi(\sigma).$$

If Σ was unstructured, there would be many options for the conditional distribution for $\Sigma|\sigma$. One simple possibility is the inverse Wishart distribution, that is

$$\Sigma|\sigma \sim \text{IW}_K \{ \sigma^2(\nu - K - 1)\mathbf{I}_K, \nu \},$$

where $\nu > K + 3$ to ensure that the distribution has the required mean and finite variance. As it is a conjugate prior, the inverse Wishart is computationally convenient. However, it is very restrictive and inflexible (Daniels & Pourahmadi, 2002). The main issue is that a single hyperparameter ν is used to control the distributional properties of all elements in Σ .

A better option for an unconstrained Σ would be to decompose $\Sigma = \text{DRD}$ where $D = \text{diag}(\sigma_1, \dots, \sigma_K)$ is a diagonal matrix of standard deviations $\sigma_i = \sqrt{\Sigma_{ii}}$ and $R = (R_{ij})$ is a matrix of correlations $R_{ij} = \Sigma_{ij} / (\sigma_i \sigma_j)$. Then we could assign the prior

$$\pi(\Sigma|\sigma) = \pi(R) \prod_{i=1}^K \pi(\sigma_i^2|\sigma),$$

giving the σ_i some distribution on \mathbb{R}^+ with $E(\sigma_i^2|\sigma) = \sigma^2$ and R some distribution over the space of $K \times K$ correlation matrices such that $E(R_{ij}) = 0$, for all $i \neq j$. For instance, we could have

$$\sigma_i \sim \text{LN} \{ \log(a), \log(\sigma^2/a) \},$$

where a is the median of σ_i and assign to R the (proper) jointly uniform distribution, $\pi(R) \propto 1$. In other words, we could assign equal prior density to all possible $K \times K$ correlation matrices. Barnard *et al.* (2000) found that the drawback of using this prior is that the marginal priors for the corresponding individual correlations, r_{ij} are not uniform. They found that as K increases, the marginal priors for individual correlations tend to be concentrated around zero.

Instead of specifying a prior for the covariance matrix, we could instead specify a prior for the precision matrix Σ^{-1} . Similarly, when Σ^{-1} is unstructured, there is a variety of options for the conditional distribution for $\Sigma^{-1}|\sigma$. A reasonable approach would be to reparameterise Σ^{-1} in terms of its square-root free Cholesky decomposition (Daniels & Pourahmadi, 2002) $\Sigma^{-1} = \text{T}^T \text{D}^{-1} \text{T}$, where $\text{D} = \text{diag}(\sigma_1, \dots, \sigma_K)$ and T is a lower triangular matrix with 1s on the diagonal and its (i, j) -th entry is $-\phi_{i,j}$ for $i > j$. This gives a new set of real valued parameters that are unconstrained, unlike the original covariance matrix, which must be symmetric and positive definite. However, for the parameters to be interpretable, the data must have some order.

Recall that our clustered data have a circular time-ordering, with the OTUs in neighbouring bins (or bins 12 and 1), typically peaking in abundance in neighbouring months. Therefore, the square-root free Cholesky decomposition might be sensible for our data. However, we choose to make our model more parsimonious by assuming a parametric form for Σ^{-1} , motivated by the structure of the data. Specifically, we take Σ^{-1} to be a symmetric, circulant, tridiagonal matrix like that described in (7.35) in Section 7.3.3. Under this structure, we want $E(\omega_0|\sigma) = \frac{1}{\sigma^2}$ and $E(\omega_1|\sigma) = 0$. To ensure that Σ^{-1} is

positive definite, it is convenient to reparameterise in terms of $\omega_0 = (\varpi_0 + \varpi_1) / \sqrt{2}$ and $\omega_1 = (\varpi_0 - \varpi_1) / 2\sqrt{2}$. The precision matrix is then positive definite if and only if $\varpi_i > 0$ for $i = 0, 1$. The full derivation of this is given in Appendix A. Therefore, using these new parameters, we want

$$\mathbb{E}(\omega_0|\sigma) = \frac{\mathbb{E}(\varpi_0|\sigma) + \mathbb{E}(\varpi_1|\sigma)}{\sqrt{2}} = \frac{1}{\sigma^2} \quad \text{and} \quad \mathbb{E}(\omega_1|\sigma) = \frac{\mathbb{E}(\varpi_0|\sigma) - \mathbb{E}(\varpi_1|\sigma)}{2\sqrt{2}} = 0$$

and hence

$$\mathbb{E}(\varpi_0|\sigma) = \mathbb{E}(\varpi_1|\sigma) = \frac{\sqrt{2}}{2\sigma^2}.$$

For example, for $i = 0, 1$, independently, we could take

$$\varpi_i|\sigma \sim \text{Ga}\left(\frac{1}{c_\omega^2}, \frac{\sqrt{2}\sigma^2}{c_\omega^2}\right).$$

Or, perhaps, we could set the median, rather than mean equal to $\sqrt{2}/(2\sigma^2)$ and take

$$\varpi_i|\sigma \sim \text{C}^+\left(0, \frac{1}{2\sigma^4}\right).$$

This particular precision structure appears in Gaussian conditional autoregression, which can be used to model spatial dependence in data. We now discuss this model (see Chapter 13 of Gelfand *et al.* (2010) for a more detailed discussion), as it aids interpretation of our errors. Suppose we have a random vector $\mathbf{X} = (X_1, \dots, X_n)^T$, where each component X_i is located at a fixed “site” $i \in 1, \dots, n$. These sites can represent a particular point in two-dimensional (or higher) space, or a particular geographical location, for example. If we wish to specify a joint distribution $p(\mathbf{x})$ for \mathbf{X} , in the temporal context, under a Markov assumption, a first-order autoregressive model $p(\mathbf{x}) = p(x_1)p(x_2|x_1) \dots p(x_n|x_{n-1})$ can be used, or can be generalised for higher dimensions. However, if we are handling spatial data, such a model is not useful, since the indices are arbitrary and could easily be permuted. Instead the full conditional distribution $p(x_i|\mathbf{x}_{-i})$ can be specified, that is the conditional distribution of X_i at a particular site, given the values at all other sites. The Gaussian conditional autoregression is a well studied model, which assumes that $p(\mathbf{x})$ is multivariate normal and $p(x_i|\mathbf{x}_{-i})$ is univariate normal. With the inclusion of a Markov property, they are also known as Gaussian Markov random fields.

Suppose that, $X_i|\mathbf{x}_{-i}$ for $i = 1, \dots, n$ is normal with conditional mean and variance

$$\mathbb{E}[X_i|\mathbf{x}_{-i}] = \mu_i + \sum_{j \neq i} \beta_{ij}(x_j - \mu_j), \tag{7.47}$$

$$\text{Var}(X_i|\mathbf{x}_{-i}) = \gamma_i^{-1}. \tag{7.48}$$

Typically, μ_i is modelled by regression but here we assume that $\mu_i = 0$ for all $i \in \{1, \dots, n\}$. Under the additional assumption of symmetry, that is $\gamma_i \beta_{ij} = \gamma_j \beta_{ji}$ for all $i \neq j$, the above conditional distributions then correspond to a multivariate joint Gaussian distribution with mean $\mathbf{0}$ and precision matrix Σ^{-1} , with (i, i) -th element γ_i and (i, j) -th element $-\gamma_i \beta_{ij}$ ($i \neq j$), assuming that Σ^{-1} is symmetric and positive definite.

For many applications, the coefficients β_{ij} will be non-zero for only a few “neighbours” of X_i . With the zero mean assumption, we can write (7.47) as

$$\mathbb{E}[X_i | \mathbf{x}_{-i}] = \sum_{j \in d_i} \beta_{ij} x_j,$$

which illustrates that the conditional mean of X_i depends only on its neighbours d_i . Now we take the example of a circular first-order autoregressive model from Gelfand *et al.* (2010), that allows the X_i s to follow a zero-mean Gaussian conditional autoregression with

$$\mathbb{E}[X_i | \mathbf{x}_i] = \phi \begin{cases} 0.5(x_2 + x_n) & \text{for } i = 1 \\ 0.5(x_{i-1} + x_{i+1}) & \text{for } 1 < i < n, \\ 0.5(x_1 + x_{n-1}) & \text{for } i = n \end{cases}$$

where $\phi \in [0, 1)$ and $\text{Var}(X_i | \mathbf{x}_{-i}) = \gamma^{-1}$. The precision matrix of \mathbf{X} is

$$\Sigma^{-1} = \frac{\gamma}{2} \begin{pmatrix} 2 & -\phi & 0 & 0 & \cdots & 0 & 0 & 0 & -\phi \\ -\phi & 2 & -\phi & 0 & \cdots & 0 & 0 & 0 & 0 \\ 0 & -\phi & 2 & -\phi & \cdots & 0 & 0 & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots & \vdots & \vdots \\ 0 & 0 & 0 & 0 & \cdots & 0 & -\phi & 2 & -\phi \\ -\phi & 0 & 0 & 0 & \cdots & 0 & 0 & -\phi & 2 \end{pmatrix},$$

which has the same structure as our precision matrix in (7.35). The first column (or row) of Σ^{-1} is the base of Σ^{-1} , which we will denote as \mathbf{e}_1 . A base of a circular matrix is a single vector that you can create the matrix from. The first column is set to the base then each subsequent column is formed by performing the forward shift operator on the current column. A forward shift moves every element of a vector one space down (or to the right, for a row vector), for example, a forward shift of a vector $\mathbf{x} = (x_1, x_2, \dots, x_{N-1}, x_N)^T$ is $(x_N, x_1, x_2, \dots, x_{N-2}, x_{N-1})^T$. The covariance matrix Σ of \mathbf{x} is also circular, where its base \mathbf{e}_2 can be calculated via a transformation involving the discrete Fourier transform of \mathbf{e}_1 . \mathbf{e}_2 is the autocovariance function of \mathbf{X} and from this the autocorrelations can be easily found. Naturally, the autocorrelation function must be symmetric, for example, the correlation between x_1 and x_2 must be equal to the correlation between x_1 and x_n . Recalling that y_{tk} is the count of bin k at time t and taking the correlation matrix that

corresponds to our symmetric, circular, tridiagonal precision matrix of the error terms in (7.35), at any time $t \in \{1, \dots, N\}$ we can interpret the following: (i) for $j \neq k$, the errors of y_{tk} have stronger correlations with the errors of y_{tj} when j is closer to k ; (ii) the correlation between the errors of y_{t1} and y_{t2} is the same as the correlation between the errors of y_{t1} and y_{tK} .

We select the symmetric, circulant, tridiagonal precision matrix for our errors defined in (7.35) because it complements our chosen clustering method well (see Section 7.1.2). Such a precision matrix requires that an assumption of circular data holds and our clustering method that incorporates seasonality gives the resulting subpopulations an order that satisfies this assumption.

7.5 Time varying mean

From our exploratory data analysis and the plots of our bins in Figures 7.1 and 7.2 in Section 7.1.2, a time invariant mean does not seem realistic, so we adopt a time varying mean. Now our model is

$$\mathbf{y}_t = \boldsymbol{\mu}_t + \mathbf{A} (\mathbf{y}_{t-1} - \boldsymbol{\mu}_{t-1}) + \boldsymbol{\epsilon}_t, \quad (7.49)$$

where $\boldsymbol{\epsilon}_t$ is normally distributed with zero-mean and has the precision matrix described in (7.35) in Section 7.4 and \mathbf{A} is unstructured, but assumed sparse and given a regularised horseshoe prior described in (7.44). To capture the seasonal variation for each bin, we use a harmonic regression to fit a time varying mean, that is

$$\boldsymbol{\mu}_t = \boldsymbol{\alpha} + \sum_{j=1}^J \boldsymbol{\beta}_j \sin\left(\frac{2\pi t j}{52}\right) + \boldsymbol{\gamma}_j \cos\left(\frac{2\pi t j}{52}\right), \quad (7.50)$$

where J is the number of harmonics and $\boldsymbol{\beta}_j \sim N_K(\mathbf{0}, \mathbf{V}_\beta)$ and $\boldsymbol{\gamma}_j \sim N_K(\mathbf{0}, \mathbf{V}_\gamma)$. After fitting the model with $J = 1, \dots, 4$ harmonics, we select $J = 2$ for our final model because there was little evidence of $\boldsymbol{\beta}_j$ and $\boldsymbol{\gamma}_j$ being non-zero for $j = 3, 4$.

7.5.1 Incorporating the chemical and environmental data

Now we wish to incorporate the chemical and environmental data into our model. A sensible approach is to regress the intercept term of the mean $\boldsymbol{\mu}_t$ on the chemical and environmental covariates. Under the guidance of a biological expert, we regress the intercept term at time t on the measurements of the covariates at the previous time point $t - 1$, as the effect of any environmental conditions is unlikely to be instantaneous. We let \mathbf{X} be the $N \times L$ matrix of the covariates, where N is the number of time points and L is the number of covariates. We find that the covariate data are skewed, so we apply a

square-root transformation. In addition to this, we standardise each column (covariate) and denote the resulting matrix as \tilde{X} . Now the intercept term α in equation (7.50) above becomes α_t , where

$$\alpha_{tk} = B_{0k} + B_{1k}\tilde{x}_{t-1,1} + \cdots + B_{Lk}\tilde{x}_{t-1,L}, \quad (7.51)$$

for $t = 2, \dots, N$ and $k = 1, 2, \dots, K$ where $\tilde{x}_{t\ell}$ is the measurement of covariate ℓ at time t and $B_{\ell k}$ is the regression coefficient for bin k and covariate ℓ , noting that B_{0k} is the intercept term for bin k .

For each $\ell = 0, 1, \dots, L$, we adopt a hierarchical prior for $B_{\ell k}$, such that

$$\begin{aligned} B_{\ell k} | B_{\ell}, \sigma_{B_{\ell}}^2 &\sim N(B_{\ell}, \sigma_{B_{\ell}}^2), \\ B_{\ell} &\sim N(a_{\alpha}, b_{\alpha}^2), \\ \sigma_{B_{\ell}}^2 &\sim \text{IG}(c_{\alpha}, d_{\alpha}) \end{aligned}$$

and each $B_{\ell k}$ is conditionally *iid*. Using the law of total expectation, law of total variance and law of total covariance, it can be shown that, marginally,

$$\begin{aligned} \text{E}[B_{\ell k}] &= a_{\alpha}, & \text{Var}(B_{\ell k}) &= \frac{d_{\alpha}}{c_{\alpha} - 1} + b_{\alpha}^2, \\ \text{Cov}(B_{\ell j}, B_{\ell k}) &= b_{\alpha}^2, & \text{Corr}(B_{\ell j}, B_{\ell k}) &= \frac{b_{\alpha}^2}{d_{\alpha}(c_{\alpha} - 1)^{-1} + b_{\alpha}^2}, \quad j, k = 1, \dots, K, \quad j \neq k, \end{aligned}$$

for each $\ell = 0, 1, \dots, L$. We select values so that the expectation of $B_{\ell k}$ is zero and $\text{Corr}(B_{\ell j}, B_{\ell k}) = 0.95$, so we have $a_{\alpha} = 0$, $b_{\alpha} = \sqrt{95}$, $c_{\alpha} = 2.25$ and $d_{\alpha} = 6.25$.

To select which covariates to include, we fit the model without any chemical and environmental covariates in (7.49) with μ_t in (7.50). Then we check to see which covariates have a lag-one correlation with the (mean of the) model residuals at each time point. These are calculated by simulating from the posterior predictive distribution for each time point, at each iteration, and subtracting the true values. This is a convenient yet simple way to select our covariates. A more rigorous approach might be to use a variable selection method instead. For example, we could make use of the regularised horseshoe prior again (see Section 9.3.2). In the AS tank, five covariates seem to be correlated with the residuals and are selected: nitrate, chemical oxygen demand (COD), ammonia, pH and phosphate. This selection of covariates is supported by our exploratory data analysis. In Section 6.2.3 of Chapter 6, we found that several of the top 12 OTUs in the AS appeared to have a (contemporary) correlation with COD, ammonia and phosphate. We also found that one of the top 12 genera seemed to be correlated with nitrate and two of the top 12 classes had weak negative (contemporary) correlations with pH. In the SS tank, eight covariates are selected: iron, chloride, fluoride, silicon, phosphate, ammonia, flow and MLVSS. As we found with the AS tank, this selection of covariates for the SS tank also agrees with

the findings of the exploratory data analysis in Section 6.2.3.

7.5.2 Handling missing values

As stated in Section 2.2.3 there are missing data for some of the environmental and chemical covariates. Like the missing values of Y_t , we therefore treat these missing values as unknowns and average over our uncertainty in their values, which requires specification of a model for the (transformed) covariates \tilde{X} . We define $\tilde{\mathbf{x}}_t = (\tilde{x}_{t1}, \dots, \tilde{x}_{tL})^T$ and have

$$\tilde{\mathbf{x}}_t = \Phi \tilde{\mathbf{x}}_{t-1} + \tilde{\boldsymbol{\epsilon}}_t \quad (7.52)$$

for $t = 2, \dots, N$, where

$$\begin{aligned} \Phi &= \text{diag}(\phi_1, \dots, \phi_L), \\ \phi_\ell &\sim \text{Beta}(\alpha_\phi, \beta_\phi), \text{ for } \ell = 1, \dots, L, \\ \tilde{\boldsymbol{\epsilon}}_t &\sim \text{N}(\mathbf{0}, \Sigma_{\tilde{X}}), \\ \Sigma_{\tilde{X}} &\sim \mathcal{W}^{-1}(\Psi_{\tilde{X}}, \nu_{\tilde{X}}). \end{aligned}$$

Essentially, we are using a first-order autoregression to model the missing values.

7.6 Posterior inference

7.6.1 Model and prior specification

In this chapter, we have developed our Bayesian hierarchical VAR(1) model for modelling the counts of the OTUs in our WWTP. Before we describe posterior inference for our model, we summarise the key elements of our model and prior. The OTUs are clustered using a seasonal phase-based approach to form 12 bins, with OTUs in different bins peaking at different months and the data are logged to make an assumption of normality more plausible. After a comprehensive simulation study, we select a regularised horseshoe prior for the matrix of autoregressive coefficients A which allows this matrix to be sparse. As a consequence of the scale chosen for our half-Cauchy hyperprior for the global shrinkage parameter of A , we scale our data to have a variance of around one. We adopt a symmetric, circulant, tridiagonal precision matrix for the errors, as such a matrix assumes circular data, which holds for our clustered data. Finally, we use harmonic regression to fit a time varying mean $\boldsymbol{\mu}_t$ and incorporate the chemical and environmental covariates in an additive fashion.

For either the AS tank or the SS tank, let \mathbf{y}_t be the scaled log counts of our $K = 12$ bins at time t and let \tilde{X} be our $(N \times L)$ matrix of standardised, square rooted chosen

chemical and environmental covariates. Thus, our final model specification is as follows:

$$\mathbf{y}_t = \boldsymbol{\mu}_t + \mathbf{A}(\mathbf{y}_{t-1} - \boldsymbol{\mu}_{t-1}) + \boldsymbol{\epsilon}_t, \quad \text{for } t = 2, \dots, N,$$

where

$$\boldsymbol{\mu}_t = \boldsymbol{\alpha}_t + \sum_{j=1}^2 \boldsymbol{\beta}_j \sin\left(\frac{2\pi t j}{52}\right) + \boldsymbol{\gamma}_j \cos\left(\frac{2\pi t j}{52}\right) \text{ and}$$

$$\alpha_{tk} = B_{0k} + B_{1k}\tilde{x}_{t1} + \dots + B_{Lk}\tilde{x}_{tL}$$

with $\boldsymbol{\beta}_j \sim \mathbf{N}_K(\mathbf{0}, 100\mathbf{I}_K)$, $\boldsymbol{\gamma}_j \sim \mathbf{N}_K(\mathbf{0}, 100\mathbf{I}_K)$ and for $\ell = 0, 1, \dots, L$, $B_{\ell k} | B_\ell, \sigma_{B_\ell}^2 \sim \mathbf{N}(B_\ell, \sigma_{B_\ell}^2)$, $B_\ell \sim \mathbf{N}(0, 95)$ and $\sigma_{B_\ell}^2 \sim \text{IG}(2.25, 6.25)$. For the matrix of autoregressive coefficients, for $j = 1, \dots, K$ and $k = 1, \dots, K$, we have

$$a_{jk} | \lambda_{jk}, \tau_A, c \sim \mathbf{N}\left(0, \tau_A^2 \tilde{\lambda}_{jk}^2\right),$$

$$\tilde{\lambda}_{jk}^2 = \frac{c^2 \lambda_{jk}^2}{c^2 + \tau_A^2 \lambda_{jk}^2}$$

$$\lambda_{jk} \sim \text{C}^+(0, 1),$$

$$\tau_A | \sigma \sim \text{C}^+(0, \tau_0^2),$$

where $c^2 \sim \text{IG}(2, 8)$ and $\tau_0 = \frac{e_0}{K^2 - e_0} \frac{\sigma}{\sqrt{N}}$, where e_0 is our prior guess for the number of non-zero autoregressive coefficients. We select $e_0 = 12$ but several other plausible values of e_0 were tested with little effect on the results (omitted). The errors $\boldsymbol{\epsilon}_t$ are normally distributed with zero-mean and precision matrix

$$\Sigma^{-1} = \begin{pmatrix} \omega_0 & \omega_1 & 0 & 0 & \cdots & 0 & 0 & 0 & \omega_1 \\ \omega_1 & \omega_0 & \omega_1 & 0 & \cdots & 0 & 0 & 0 & 0 \\ 0 & \omega_1 & \omega_0 & \omega_1 & \cdots & 0 & 0 & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots & \vdots & \vdots \\ 0 & 0 & 0 & 0 & \cdots & 0 & \omega_1 & \omega_0 & \omega_1 \\ \omega_1 & 0 & 0 & 0 & \cdots & 0 & 0 & \omega_1 & \omega_0 \end{pmatrix},$$

with a reparameterisation where $\omega_0 = (\varpi_0 + \varpi_1) / \sqrt{2}$ and $\omega_1 = (\varpi_0 - \varpi_1) / 2\sqrt{2}$ and for $i = 1, 2$, $\varpi_i | \sigma \sim \text{Ga}(1, \sqrt{2}\sigma^2)$, where $\sigma \sim \text{LN}(0, 10)$.

7.6.2 Posterior inference via HMC

There are several parameters in the model. First, there is the matrix of autoregressive coefficients A with its global shrinkage parameter $\tau_A \in \mathbb{R}_+$, the matrix of local shrinkage parameters Λ and the regularising parameter c^2 . Additionally, there are the parameters for the precision matrix Σ^{-1} of the errors, denoted by $\boldsymbol{\omega} = (\omega_0, \omega_1) \in \mathbb{R}_+^2$ and the parameter σ . There are the harmonic regression coefficients for each $\boldsymbol{\mu}_t$, denoted by $\boldsymbol{\theta} = \{\boldsymbol{\beta}_1, \boldsymbol{\beta}_2, \gamma_1, \gamma_2\}$. We also have the coefficients of the chemical and environmental covariates \tilde{X} and the intercepts B , with mean B_ℓ and variance $\sigma_{B_\ell}^2$ for each B_{lk} in B . Finally, we have the parameters Φ and $\Sigma_{\tilde{X}}$ in the missing data model.

The prior density is given by

$$\begin{aligned} \pi(A, \tau_A, \Lambda, c^2, \boldsymbol{\omega}, \sigma, \boldsymbol{\theta}, B, B_\ell, \sigma_{B_\ell}^2) &= \pi(A|\Lambda, \tau_A, c^2) \pi(\tau_A|\sigma) \pi(\sigma) \pi(\Lambda) \pi(c^2) \pi(\boldsymbol{\omega}|\sigma) \\ &\quad \times \pi(\boldsymbol{\theta}) \pi(B|B_\ell, \sigma_{B_\ell}^2) \pi(B_\ell) \pi(\sigma_{B_\ell}^2) \\ &\quad \times \pi(\Phi) \pi(\Sigma_{\tilde{X}}). \end{aligned}$$

The posterior of the unknowns is therefore given by

$$\begin{aligned} \pi(A, \tau_A, \Lambda, c^2, \boldsymbol{\omega}, \sigma, \boldsymbol{\theta}, B, B_\ell, \sigma_{B_\ell}^2, \Phi, \Sigma_{\tilde{X}}|Y) &\propto \pi(A, \tau_A, \Lambda, c^2, \boldsymbol{\omega}, \sigma, \boldsymbol{\theta}, B, B_\ell, \sigma_{B_\ell}^2, \Phi, \Sigma_{\tilde{X}}) \\ &\quad \times \pi(Y|A, \boldsymbol{\theta}, B, \boldsymbol{\omega}), \end{aligned}$$

where $\pi(Y|A, \boldsymbol{\theta}, B, \boldsymbol{\omega})$ is the data likelihood. Since our model has a Markovian structure, the likelihood can be written as

$$\pi(Y|A, \boldsymbol{\theta}, B, \boldsymbol{\omega}) = \pi(\mathbf{y}_1|A, \boldsymbol{\theta}, B, \boldsymbol{\omega}) \prod_{t=2}^N \pi(\mathbf{y}_t|\mathbf{y}_{t-1}, A, \boldsymbol{\theta}, B, \boldsymbol{\omega}).$$

Following the same arguments presented by Lei *et al.* (2011), we ignore the contribution of the marginal model $\pi(\mathbf{y}_1|A, \boldsymbol{\theta}, B, \boldsymbol{\omega})$, as we have large enough sample sizes for each tank that little information will be lost in doing so. Our likelihood is conditioned on the observed value of \mathbf{y}_1 and is

$$\pi(Y|A, \boldsymbol{\theta}, B, \boldsymbol{\omega}) = \prod_{t=2}^N \pi(\mathbf{y}_t|\mathbf{y}_{t-1}, A, \boldsymbol{\theta}, B, \boldsymbol{\omega}),$$

where the conditional distributions $\mathbf{y}_t|\mathbf{y}_{t-1}, A, \boldsymbol{\theta}, B, \boldsymbol{\omega} \sim N_K(\boldsymbol{\mu}_t + A(\mathbf{y}_{t-1} - \boldsymbol{\mu}_{t-1}), \Sigma)$ and $K = 12$.

This distribution is analytically intractable, so we make use of MCMC methods to generate samples from the posterior. More specifically, we fit the model in **Stan** (see Appendix C for the **Stan** model) which uses the NUTS algorithm (see Section 2.3.3). The

model is run for each data set for 10K iterations with a warm-up period of 3000 iterations. In the interests of saving memory, the output is thinned to leave us with 1000 samples from the posterior. The usual graphical and numerical diagnostic checks (see Section 2.3.2) gave no evidence of any lack of convergence and mixing was good. We present the results in the next chapter based on these posterior samples.

Chapter 8

Metagenomics application

Now that we have presented the details of our model, we present the results of our analyses. Recall that for the AS tank and the SS tank we have two different (transformed) data sets of scaled log counts for 12 subpopulations of OTUs, or 12 bins, as we described in Section 7.1.2, where bin 1 peaks in February, bin 2 peaks in January, bin 3 peaks in December and so on. For each tank we have a different set of standardised, square rooted chemical and environmental covariates in our model. The selection process for these covariates was described and the rationale for transforming the covariate data was explained in Section 7.5.1. In this chapter, we look at the posterior distributions for the parameters in our model and provide biological explanations for some of our findings.

8.1 Activated sludge tank results

First we look at the results for the AS tank. Regarded as the tank of more biological interest (see Section 2.2.4), we discuss this tank in more detail than the SS tank.

8.1.1 Time varying mean

We begin with the time varying mean $\boldsymbol{\mu}_t$ of our model, which is modelled with harmonic regression with two harmonics. Recall that we have

$$\boldsymbol{\mu}_t = \boldsymbol{\alpha}_t + \sum_{j=1}^2 \boldsymbol{\beta}_j \sin\left(\frac{2\pi jt}{52}\right) + \sum_{j=1}^2 \boldsymbol{\gamma}_j \cos\left(\frac{2\pi jt}{52}\right),$$

where

$$\begin{aligned}\boldsymbol{\alpha}_t &= \mathbf{B}^T \mathbf{X}_t \\ &= \mathbf{B}_0 + \mathbf{B}_1 X_{1t} + \dots + \mathbf{B}_L X_{Lt}\end{aligned}$$

and \mathbf{X}_t are the transformed measurements of our chosen chemical and environmental covariates at time t .

Chemical and environmental covariates

In the AS tank, we have five covariates in our model: nitrate, chemical oxygen demand (COD), ammonia, pH and phosphate. Figure 8.1 and Figure 8.2 show the posterior means and the 95% credible intervals (CIs) for the regression coefficients for each bin in the model. When looking at the results of our analyses, we take zero lying in the CI of a particular parameter as a quick discriminator to suggest that the parameter’s value may be (close to) zero. However, this does not necessarily mean that there is not considerable support for a positive or negative coefficient (for example, see phosphate below). For all 12 bins, the CIs for nitrate and phosphate shown in Figure 8.1 all include zero, suggesting that neither of these covariates has a linear relationship with the time varying mean of any of our bins, although for phosphate we note that the CIs for bins 2 and 3 only just overlap zero. Despite all of the CIs overlapping zero, there is a clear pattern for phosphate, with “winter blooming” bins showing a positive relationship and “summer blooming” bins showing a negative relationship, though this might just be an artefact of how the data were binned.

From Figure 8.2, bins 4 to 10 all have positive regression coefficients with COD, with bin 7 having the largest regression coefficient with a posterior mean of 0.2136 (4 d.p.). Ammonia has a positive regression coefficient with four bins (4 to 7) and, as we saw with COD, bin 7 has the largest regression coefficient with a posterior mean of 0.1412 (4 d.p.). Finally, bin 12 seems to have a positive relationship with pH, with its regression coefficient having a posterior mean of 0.1167 (4 d.p.).

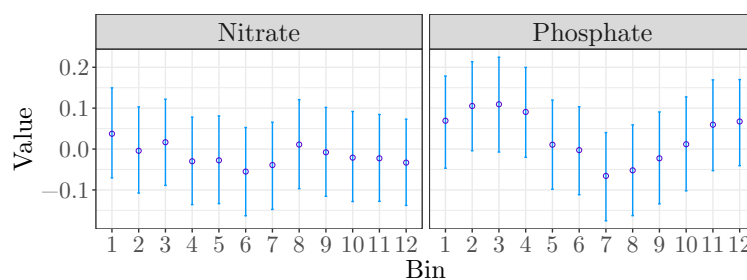


Figure 8.1: Posterior means (○) and 95% credible intervals (—) for nitrate and phosphate (AS tank).

As mentioned in Section 6.1.1, COD is an approximation of the amount of oxygen that can be consumed by reactions in water. In other words, it indicates the mass of dissolved oxygen which is required by microorganisms for degradation of organic and some inorganic compounds. COD is used to quantify oxidisable pollutants in wastewater. Ammonia is

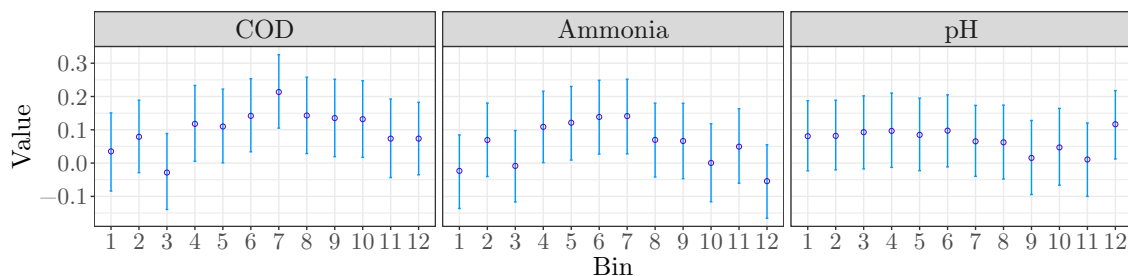


Figure 8.2: Posterior means (\circ) and 95% credible intervals (—) for COD, ammonia and pH (AS tank).

an (inorganic) pollutant that is oxidisable, and so, it creates oxygen demand. Therefore, it is perhaps unsurprising that some of the bins have non-zero regression coefficients with both COD and ammonia. Removal of ammonia and other pollutants is essential in the treatment of wastewater and ammonia is removed through nitrification by bacteria. It may be possible, that some covariates such as COD and ammonia could be driven by the bin counts (OTU abundances). Thus, these potential relationships between the bins and COD and ammonia identified here could suggest the presence of certain metabolic reactions in the AS.

In the nitrogen cycle, nitrification is a two-step process of ammonia oxidation then nitrite oxidation. Bacteria from the genus *Nitrosomonas* can oxidise ammonia to nitrite (Wetzel, 2001), although there are other ammonia oxidising microorganisms (AOM) too. (There are ammonia oxidising bacteria and *Archaea*, hence the use of the word ‘microorganisms’ in the term rather than bacteria.) *Nitrobacter* bacteria from the same phylum as *Nitrosomonas* oxidise nitrite to nitrate but are difficult to detect *in-situ*. Wagner *et al.* (1996) suggested that this could be because they have a minor role in WWTPs and although Alawi *et al.* (2009) agreed that their role is small, they also noted that lack of detection does not necessarily mean lack of presence. In the AS tank, no *Nitrobacter* counts are recorded. This could suggest that in our WWTP other nitrite oxidising bacteria (NOB), for example, *Nitrospira*, are responsible for nitrite removal or that the *Nitrobacter* bacteria simply have not been detected, as seen in the literature.

Until recently, *Nitrospira* were considered solely NOB (Mehrani *et al.*, 2020). Daims *et al.* (2015) and van Kessel *et al.* (2015) independently discovered a single microorganism from the genus *Nitrospira* that can carry out complete nitrification through the *comammox* (complete oxidation of ammonia to nitrate) process. Additionally, it has been found that there is a reciprocal feeding interaction between nitrifiers. Some species of *Nitrospira* are able to convert urea to ammonia and carbon dioxide, which means they can supply AOM with ammonia and in return receive nitrite produced by ammonia oxidation (Koch *et al.*, 2015). We do not know exactly which microorganisms our OTUs represent but an OTU

from the genus *Nitrospira* is one of the most abundant OTUs within bin 4, so this could explain why we see the positive coefficient for COD and bin 4. Additionally, this OTU could be capable of comammox which could provide a reasonable explanation as to why there is a positive coefficient for ammonia and bin 4 too. In addition to this, an OTU from the genus *Terrimonas* is the most abundant in bin 4, with a median within-bin relative abundance of around 22.78%. Bacteria from this genus are strictly aerobic and assimilate organic compounds such as sugars and proteins (McIlroy *et al.*, 2015). This provides another possible explanation as to why bin 4 has a positive relationship with COD.

To understand why some of the other bins may have a relationship with ammonia and COD, we look at the top 12 OTUs within some of the bins. Of the 1274 OTUs in bin 5, the most abundant OTU based on median within-bin relative abundance ($\sim 9.1\%$) is from the genus *Zoogloea*. Bacteria from this genus are highly active oxidisers of organic compounds (Dugan, 1981). Recalling that our covariates are incorporated into the model via lag-one regression, the transformed COD measurement from the previous time point is used to model the intercept of the time varying mean at the current time point. If COD is high then this would suggest that there is a larger amount of organic compounds available for the *Zoogloea* bacteria to oxidise for energy and grow, thus explaining the positive coefficient between the bin containing *Zoogloea* and COD. However, this could result in the amount of organic compounds (and COD) decreasing which in turn could eventually slow the growth rate of the *Zoogloea* bacteria. More organic compounds can migrate into the system as more wastewater enters the WWTP which could then cause the COD to rise again. This describes a predator-prey-like dynamic and demonstrates that the relationships between the covariates and bins (of OTUs) are unlikely to be simple.

An OTU from the genus *Leptothrix* is the most abundant OTU in bin 7 based on median within-bin relative abundance ($\sim 10.4\%$). Species from this genus typically oxidise iron and manganese (McIlroy *et al.*, 2015). The second most abundant OTU is from the genus *Dechloromonas* with a median within-bin relative abundance of around 10%. As part of our exploratory analysis, we saw that some of the top genera were correlated with COD (Figure 6.14, Section 6.2.3), where *Dechloromonas* had a fairly weak positive correlation and *Leptothrix* did not appear in the top 12 genera in the AS tank. Some species of *Dechloromonas* are *polyphosphate-accumulating organisms* (PAOs) and some species have a role in *denitrification* (McIlroy *et al.*, 2015). PAOs are bacteria that aid the removal of organic compounds containing phosphorus from wastewater, under certain conditions, in a process called *enhanced biological phosphorus removal*. Denitrification is the reduction of nitrate to the eventual product of nitrogen gas, following a series of intermediate gaseous nitrogen oxide products. Ammonia, nitrate and phosphorus all contribute to the COD of wastewater. Applying logic similar to that discussed for the *Zoogloea* bacteria in bin 5, a positive and likely non-linear relationship between COD and bin 7 seems sensible.

Finally, we focus on bin 12, which is the only bin that has a non-zero (positive) coefficient with pH. All from the genus *Rhodobacter*, OTUs 15, 33 and 65 are in the top 12 OTUs of bin 12 and together they represent about 23.90% of bin 12 on average. Most strains of *Rhodobacter* grow at an optimal pH range of 6.5 - 7.5 (Imhoff, 2015). In the AS tank, the pH range is from 5.02 to 7.5 with a median of 6.53, thus providing a possible explanation as to why bin 12 has a positive relationship with pH. Looking back at the heatmaps in Section 6.2.3, we see that *Rhodobacter* does not seem to have a correlation with pH (Figure 6.14). However, we see that OTU 15 possibly has a weak positive correlation with pH (Figure 6.10). It is also important to remember that the heatmaps show correlations not lag-one correlations. Calculating both the correlation (0.1155) and lag-one correlation (0.1679) between pH and *Rhodobacter*, we see that the lag-one correlation is stronger, thus corroborating our results. Furthermore, this relationship remains after allowing for other things, which highlights the benefit of the model, as otherwise this relationship may go unnoticed.

Harmonic regression coefficients

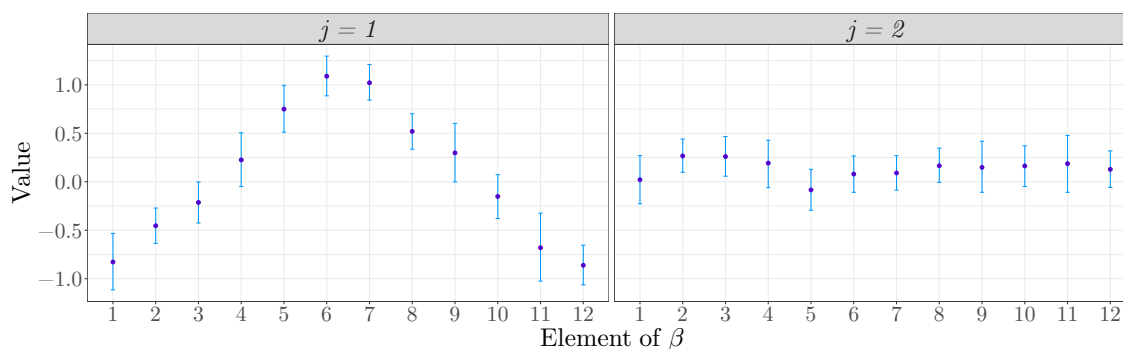


Figure 8.3: Posterior means (\circ) and 95% credible intervals (—) for β_j (AS tank).

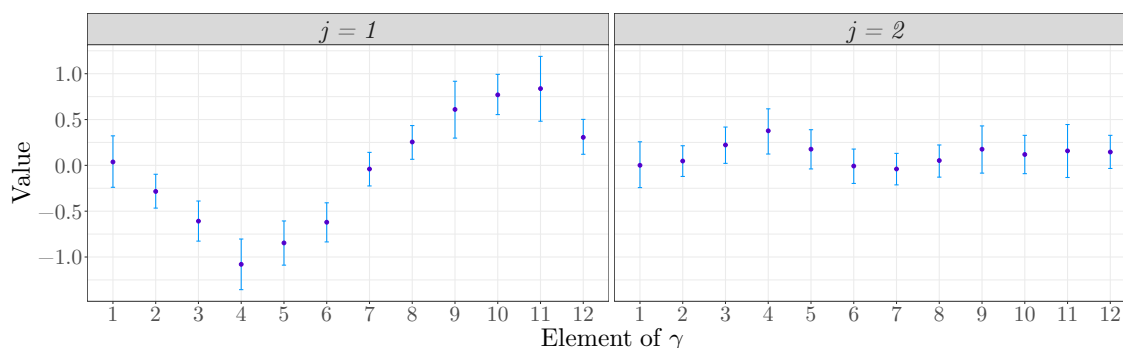


Figure 8.4: Posterior means (\circ) and 95% credible intervals (—) for γ_j (AS tank).

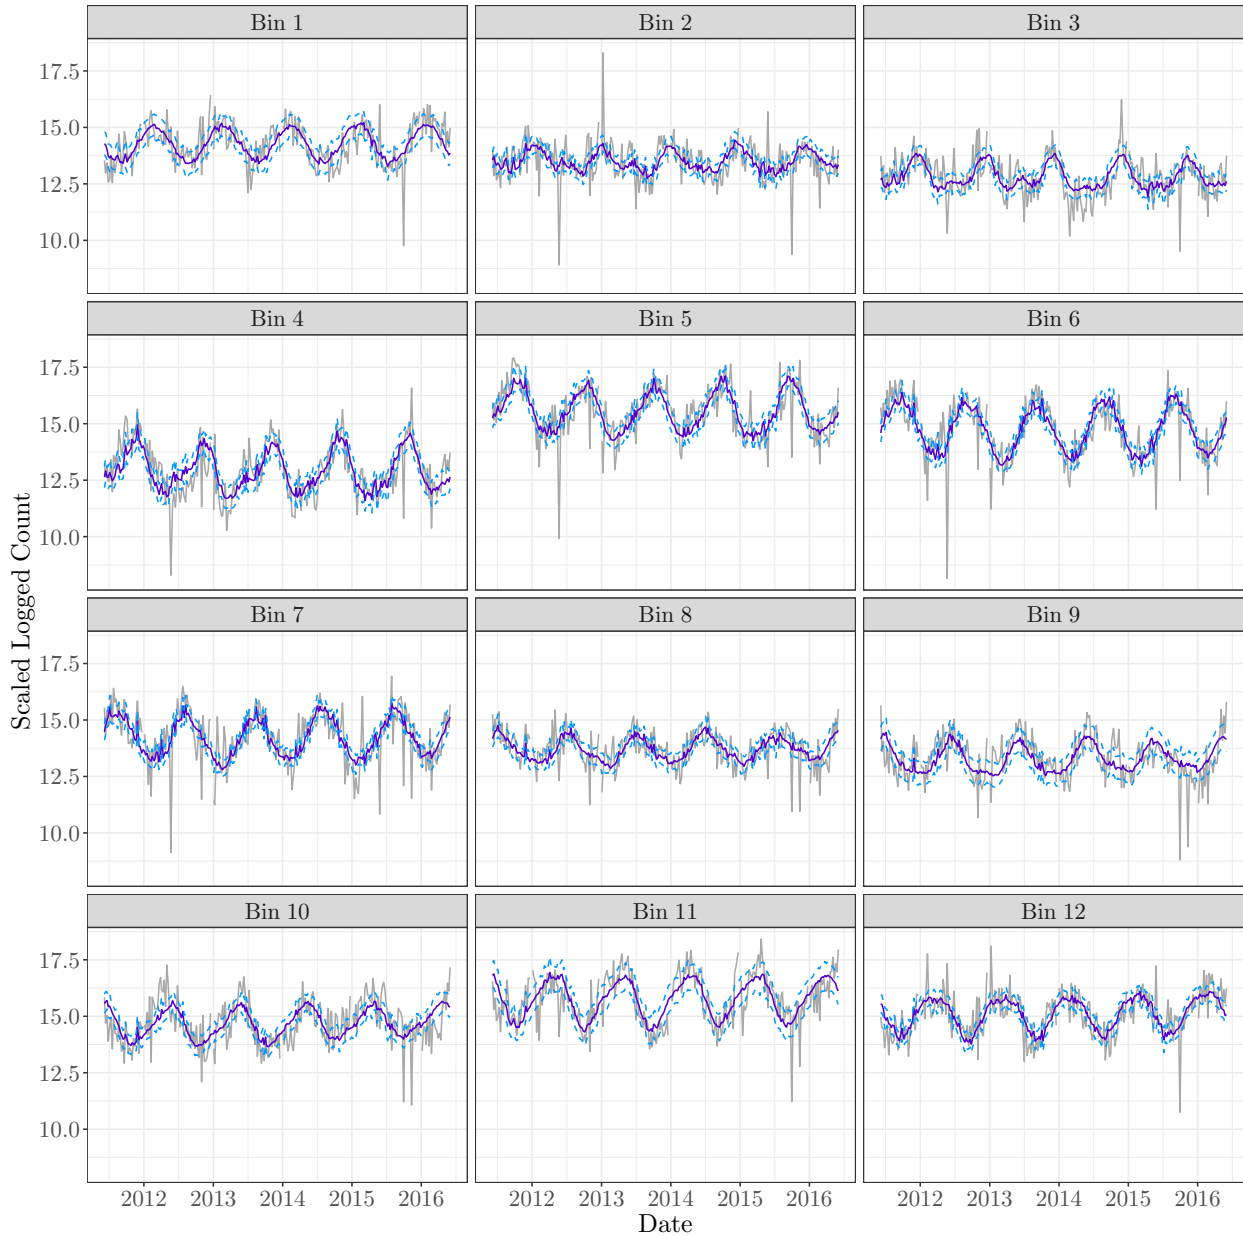


Figure 8.5: Posterior means (—) and 95% credible intervals (---) for the time varying means with scaled log counts (—) for each bin (AS tank).

Now we look at the harmonic regression coefficients of our model. Figures 8.3 and 8.4 show the posterior means and 95% CIs for our harmonic regression coefficients β_j and γ_j for $j = 1, 2$. The change in the values of the β_{jk} and γ_{jk} across bins, $k = 1, \dots, 12$, for the first harmonic ($j = 1$) can be explained by our chosen clustering method, which is based on the idea that the OTUs display seasonal variation and peak in different months. Recalling from Section 7.1.2, OTUs in bin 1 peak in February, OTUs in bin 2 peak in

January, OTUs in bin 3 peak in December and so on. Based on the CIs, it would seem that only bins 2, 3 and 4 seem to have non-zero coefficients for the second harmonic, suggesting that their scaled log counts do not follow a pattern as simple as a sinusoid. Recall that in the time series plots of the bins in Figure 7.1 (Section 7.1.2), we saw that the annual peaks were not as obvious in bins 2 and 3, suggesting a sinusoid may not be such a good descriptor. Perhaps, this is why we have non-zero coefficients for the second harmonics for these two bins. Reviewing the time series plots again, we can also see that in bin 4, there seem to be two peaks within 2013, with a smaller peak in the middle of the year and a larger peak around October. This might explain why we have a non-zero coefficient for the second harmonic in bin 4. Figure 8.5 shows posterior means of the time varying means μ_t and the 95% CIs plotted over the scaled log counts for each bin. The seasonal patterns of each bin seem to have been captured fairly well.

8.1.2 Matrix of autoregressive coefficients

The matrix of autoregressive coefficients can inform us of relationships between our bins. Recall that a_{jk} tells us how the scaled log count of bin k at the previous time point $y_{t-1,k}$ contributes to the scaled log count of bin j at the current time point $y_{t,j}$. We note that for both the AS tank and the SS tank, we find that the posterior masses of A lie within the stationary region. This is done by checking that all the eigenvalues of A for every posterior sample have a size smaller than one. The posterior means of the autoregressive coefficients are shown in a heatmap in Figure 8.6 and they are also shown in Figure 8.7 with their corresponding 95% CIs. From the heatmap, we can see that the matrix of autoregressive coefficients based on posterior means is fairly sparse. With the exception of bin 12, all the bins have a positive autoregressive coefficient with themselves. In other words, the scaled log count of the previous time point seems to have a positive relationship with the scaled log count at the current time point, which seems sensible. Bins 1, 4, 5, 9 and 11 have particularly large “within-bin” autoregressive coefficients with posterior means larger than 0.5. It is surprising that the $a_{12,12}$ is a near-zero coefficient, with a posterior mean of 0.019 (3 d.p.). It could be that $y_{t,12}$ is better explained by $y_{t-1,1}$ than $y_{t-1,12}$. Bin 1 peaks in February and bin 12 peaks in March and the posterior mean of $a_{12,1}$ is positive (0.276), so this does not seem unreasonable.

In addition to the within-bin autoregressive coefficients a_{kk} , we see from the CIs in Figure 8.7 that there is evidence for a few non-zero “between-bin” posterior autoregressive coefficients $a_{jk}, j \neq k$. The posterior means for these coefficients whose CIs do not overlap zero are listed in Table 8.1. Apart from $a_{12,1}$, the non-zero between-bin coefficients are smaller than all the non-zero within-bin coefficients.

Now we try to understand why we have non-zero coefficients between particular bins by looking at the most abundant OTUs in each bin again. As mentioned above, the most

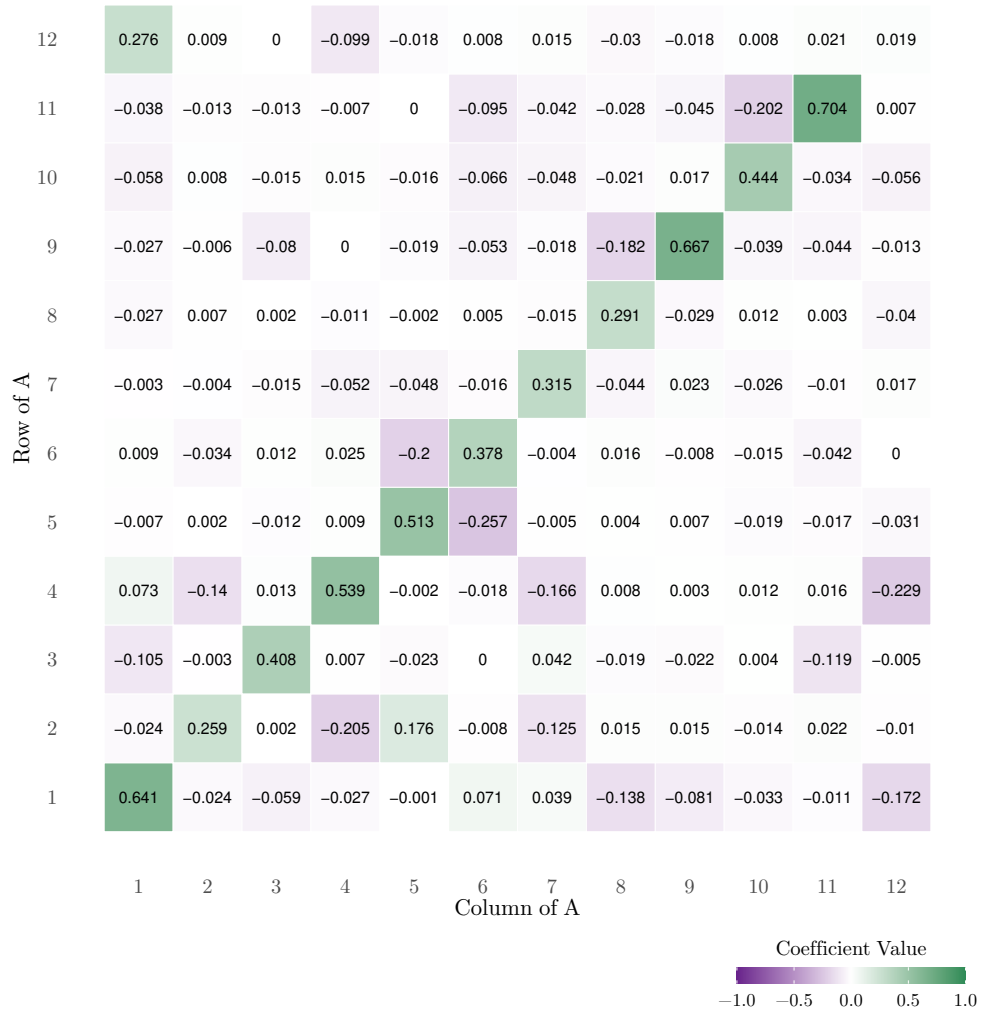


Figure 8.6: Heatmap of the posterior means of the autoregressive coefficients (AS tank).

Coefficient	Posterior Mean
$a_{2,4}$	-0.204
$a_{2,5}$	0.176
$a_{3,11}$	-0.119
$a_{4,2}$	-0.14
$a_{4,7}$	-0.166
$a_{4,12}$	-0.229
$a_{5,6}$	-0.257
$a_{6,5}$	-0.200
$a_{11,10}$	-0.202
$a_{12,1}$	0.276

Table 8.1: Posterior means (3 d.p) of the non-zero between-bin coefficients.

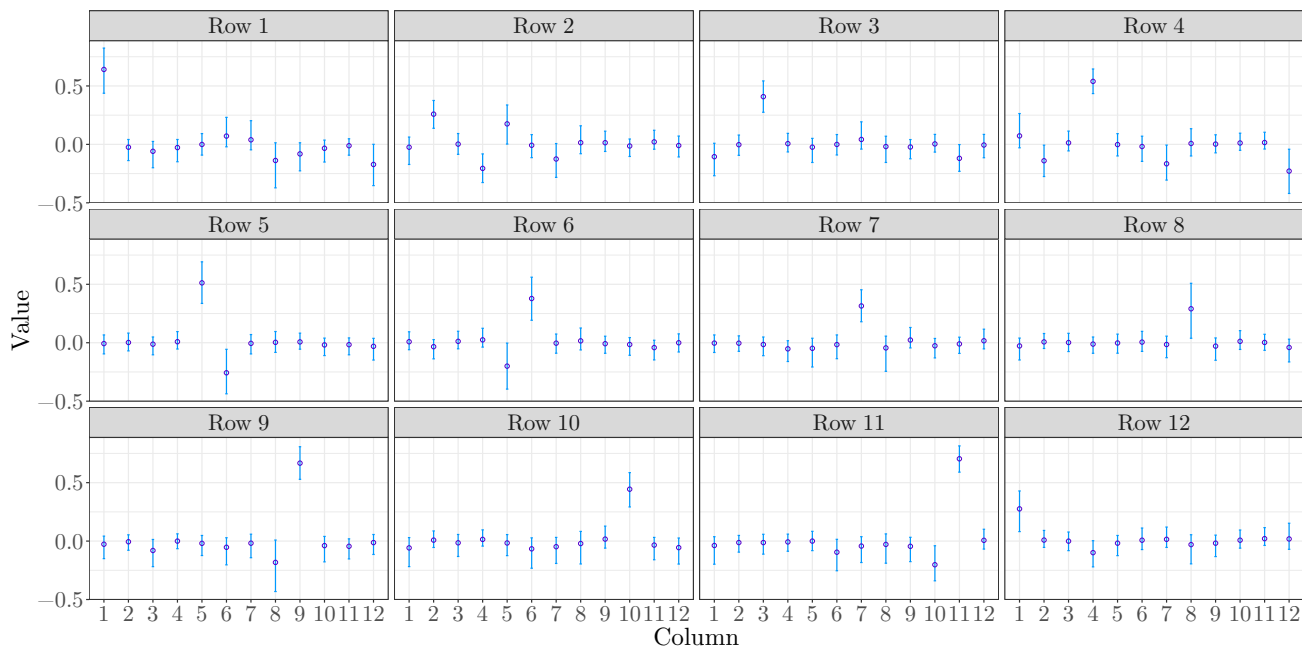


Figure 8.7: Posterior means (\circ) and 95% credible intervals (—) of the autoregressive coefficients (AS tank).

abundant OTU in bin 5 is from the genus *Zoogloea*. The second most abundant OTU is from the genus *Acidovorax*. In bin 6, the second most abundant OTU is from the genus *Dechloromonas*, which as mentioned above is capable of nitrite reduction, as well as sulphate reduction. This is also true for *Zoogloea* and *Acidovorax* bacteria (McIlroy *et al.*, 2015). Perhaps these bacteria amongst others that are not in the most abundant OTUs are competing for resources such as nitrite and sulphate, hence why we see the negative autoregressive coefficients.

As we already stated above, an OTU from the strictly aerobic genus *Terrimonas* is the most abundant in bin 4. The second most abundant OTU is from the genus *Ca. Microthrix*, which is also described as aerobic in McIlroy *et al.* (2015). The top two OTUs in bin 2 are from the family *Rhodobacteraceae* with unknown genera. According to Pujalte *et al.* (2014), at the time of their publication, there were 288 known species from 99 genera in the family *Rhodobacteraceae*, any of which the top two OTUs could be from. However, the third most abundant OTU, representing on average 14.50% of bin 2, is from the genus *Haematobacter* from the same family, which are aerobic bacteria. As we cannot determine the genera of the top two OTUs in bin 2, it is impossible to determine whether they are aerobic or not but they *could* be, especially as most *Rhodobacteraceae* are aerobic (Pujalte *et al.*, 2014). Perhaps there are negative autoregressive coefficients between bins 2 and 4 because these aerobic OTUs in both bins are competing for oxygen.

Exploring these other bins and their top 12 OTUs, it is hard to identify “obvious”

explanations as to why some of the autoregressive coefficients are non-zero. However, we can see that the scaled log counts of bins 7 and 12 at the previous time point seem to have a negative effect on the scaled log count of bin 4 at the current time points. Additionally, bin 5 seems to have a positive effect on bin 2 at the next time point, as does bin 1 with bin 12 at the next time point. Finally, bin 10 seems to have a negative effect on bin 11 at the next time point. The genera of the top six OTUs in each bin are listed below in Table 8.2. Perhaps there are some relationships between these genera corresponding to the bins with non-zero autoregressive coefficients. However, it could be that other OTUs or collections of OTUs that are not as abundant within the bins are responsible for the non-zero autoregressive coefficients. Although it would be interesting to delve deeper into the biological interpretation of our results, it is beyond the scope of this thesis.

Bin	Genera
1	<i>Tabrizicola</i> , <i>Rhodobacter</i> , <i>Methylorosula</i> , <i>Ruminococcus</i> 2, <i>Fusicatenibacter</i> , <i>Albidiferax</i>
2	Unknown (Family: <i>Rhodobacteraceae</i> , 2), <i>Haematobacter</i> , <i>Polymorphobacter</i> , <i>Thermomonas</i> , <i>Novosphingobium</i>
3	<i>Amaricoccus</i> , Unknown (Family: <i>Microbacteriaceae</i>), <i>Paracoccus</i> (2), CL500-29 marine group, <i>Ca. Microthrix</i>
4	<i>Terrimonas</i> , <i>Ca. Microthrix</i> , <i>Zymomonas</i> , <i>Acinetobacter</i> , <i>Phenyllobacterium</i> , <i>Nitrospira</i>
5	<i>Zoogloea</i> , <i>Acidovorax</i> , <i>Ca. Microthrix</i> , <i>Afipia</i> , <i>Simplicispira</i> , <i>Iamia</i>
6	Unknown (Family: <i>Alcaligenaceae</i>), Unknown (Family: MN67), <i>Dokdonella</i> , <i>Dechloromonas</i> , <i>Piscinibacter</i> , <i>Rhizobium</i>
7	<i>Leptothrix</i> , <i>Dechloromonas</i> , Unknown (Family: <i>Cytophagaceae</i> , <i>Hirschia</i> , Unknown (Family: <i>Saprospiraceae</i>), <i>Iamia</i>
8	<i>Romboutsia</i> , Unknown (Family: <i>Comamonadaceae</i>), <i>Devosia</i> , <i>Lautropia</i> , Unknown (Family: Unknown), <i>Lacibacter</i>
9	<i>Defluviimonas</i> , <i>Intestinibacter</i> , <i>Nitrosomonas</i> , <i>Hyphomicrobium</i> , <i>Ferruginibacter</i> (2)
10	Unknown (Family: Gsoil-1167), <i>Ca. Microthrix</i> , <i>Rhizobacter</i> , <i>Defluviimonas</i> , Unknown (Family: <i>Rhodobacteraceae</i>), <i>Ca. Nitrotoga</i>
11	Unknown (Family: <i>Saprospiraceae</i> , 2), <i>Defluviimonas</i> , <i>Terrimonas</i> (2), <i>Rhodofera</i>
12	<i>Trichococcus</i> , <i>Rhodobacter</i> (3), <i>Ornithinibacter</i> , <i>Subdoligranulum</i>

Table 8.2: Genera of the top six OTUs in each bin based on median relative abundance (AS tank).

8.1.3 Precision matrix for errors

Recall that the errors in our model ϵ_t follow a $N(0, \Sigma)$ distribution and we have a symmetric, tridiagonal, circulant precision matrix for the errors, that is

$$\Sigma^{-1} = \begin{pmatrix} \omega_0 & \omega_1 & 0 & 0 & \cdots & 0 & 0 & 0 & \omega_1 \\ \omega_1 & \omega_0 & \omega_1 & 0 & \cdots & 0 & 0 & 0 & 0 \\ 0 & \omega_1 & \omega_0 & \omega_1 & \cdots & 0 & 0 & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots & \vdots & \vdots \\ 0 & 0 & 0 & 0 & \cdots & 0 & \omega_1 & \omega_0 & \omega_1 \\ \omega_1 & 0 & 0 & 0 & \cdots & 0 & 0 & \omega_1 & \omega_0 \end{pmatrix}.$$

The posterior means for ω_0 and ω_1 are 6.7354 and -3.2183 (to 4 d.p.) respectively, with standard deviations 0.1926 and 0.0987 (to 4 d.p.). If it is invertible, it can easily be shown that the inverse of a symmetric, tridiagonal, circulant matrix is a symmetric, circulant matrix. Thus the covariance matrix for the errors Σ is a symmetric, circulant matrix. The correlation matrix associated with Σ is therefore defined by the lag- k correlations ρ_k for $k = 1, \dots, 6$. Figure 8.8 shows the posterior means and 95% CIs for ρ_1, \dots, ρ_6 . All of the CIs lie above zero which provides evidence of between-bin correlation in the errors.

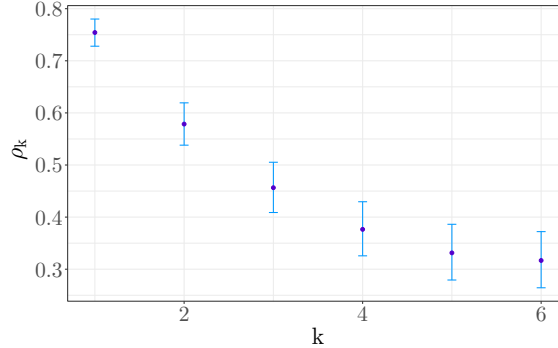


Figure 8.8: Posterior means (\circ) and 95% credible intervals (—) of the lag- k correlations ρ_k for $k = 1, \dots, 6$ (AS tank).

8.2 Settled sewage tank results

Now we discuss the results for the SS tank and make some comparisons with our findings for the AS tank. The same model that was used for the AS tank is used for the SS tank, with the exception of the chemical and environmental covariates selected to model the intercept term of the time varying mean.

8.2.1 Time varying mean

Chemical and environmental covariates

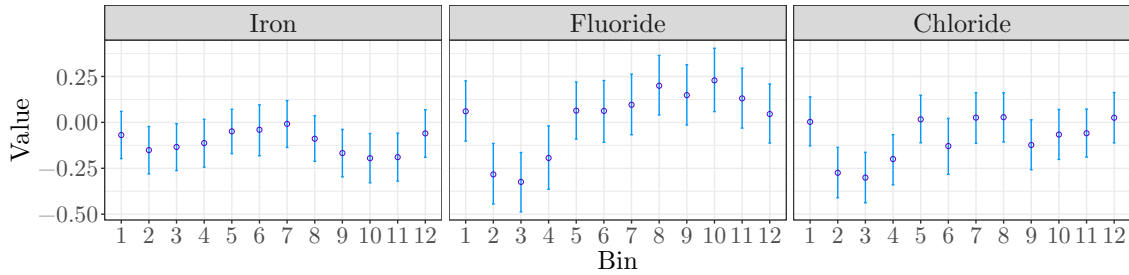


Figure 8.9: Posterior means (\circ) and 95% credible intervals (—) for iron, fluoride and chloride (SS tank).

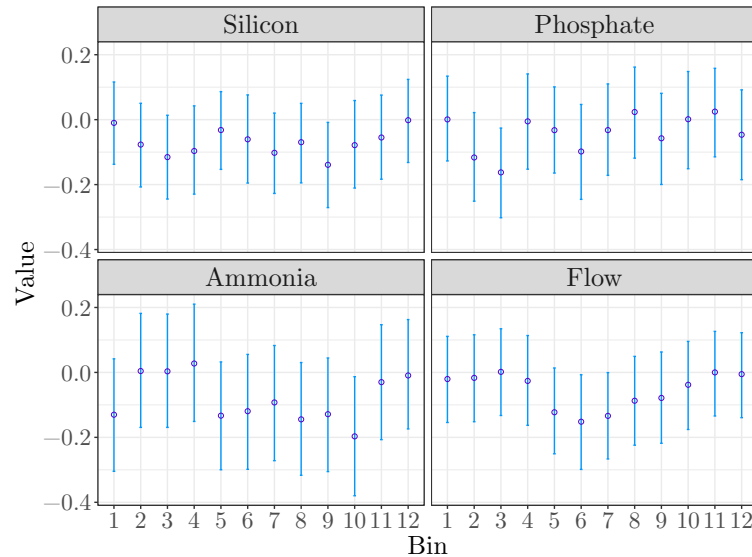


Figure 8.10: Posterior means (\circ) and 95% credible intervals (—) for silicon, phosphate, ammonia and flow (SS tank).

In the model for the SS tank, we have eight covariates: iron, chloride, fluoride, silicon, phosphate, ammonia, flow and MLVSS. The posterior means and 95% CIs for iron, chloride and fluoride are shown in Figure 8.9. These covariates have non-zero coefficients with multiple bins. Iron has a negative coefficient with bins 2, 3, 9, 10 and 11. Fluoride has a negative coefficient with bins 2 to 4 and a positive coefficients with bins 8 and 10. Fluoride is well known for its use in combatting oral bacteria and is toxic to many different bacteria, but not all bacteria (Ochoa-Herrera *et al.*, 2009), so it seems plausible that the coefficient for fluoride varies from bin to bin. For example, perhaps a bin with a negative coefficient contains more bacteria that are more susceptible to fluoride toxicity. Chloride

has a negative coefficient with bins 2 to 4. Figure 8.10 shows the posterior means and 95% CIs for covariates that have a non-zero coefficient with just one of the 12 bins. Silicon has a negative coefficient with bin 9, phosphate has a negative coefficient with bin 3, ammonia has a negative coefficient with bin 10 and flow has a negative coefficient with bin 6. The plot of the posterior means and CIs for MLVSS is omitted, as the CIs for all the bins overlap zero.

In our exploratory analysis (Chapter 6), we identified fluoride, chloride, flow, ammonia and phosphate, as being correlated with some of the 12 most abundant OTUs, genera and classes. We also found that iron had weak negative correlations with some of the top classes. Silicon was not identified as a correlated covariate but MLVSS was. These slight discrepancies between the findings of our exploratory analysis and results after fitting our model are perfectly reasonable, given that the data are grouped differently, contemporaneous correlations rather than lag-one correlations were used for exploratory analysis and our model consists of many components.

Bin	Genera
1	<i>Subdoligranulum</i> , <i>Rumiococcus</i> 2 (2), <i>Anaerostipes</i> , <i>Coprococcus</i> 3, <i>Ruminiclostridium</i> 5
2	<i>Proteocatella</i> , <i>Faecalibacterium</i> , Unknown (Family: <i>Ruminococcaceae</i>), <i>Lachnospira</i> , <i>Simplicispira</i> , <i>Psychrobacter</i>
3	<i>Acidovorax</i> , <i>Faecalibacterium</i> , <i>Pseudorhodofera</i> , <i>Streptococcus</i> , <i>Ferruginibacter</i> , <i>Acinetobacter</i>
4	<i>Faecalibacterium</i> , <i>Christensenellaceae</i> R-7 group, Unknown (Family: Unknown), <i>Lachnospira</i> , <i>Coprococcus</i> 2, <i>Paludibacter</i>
5	<i>Fusicatenibacter</i> , <i>Escherichia/Shigella</i> , <i>Paracoccus</i> , <i>Lactobacillus</i> , Unknown (Family: TM146), <i>Ruminococcaceae</i> UCG-002
6	<i>Leptotrichia</i> , <i>Acinetobacter</i> , <i>Arcobacter</i> , Unknown (Family: <i>Enterobacteriaceae</i>), <i>Shuttleworthia</i> , <i>Rickettsiella</i>
7	<i>Blautia</i> , <i>Subdoligranulum</i> (2), Unknown (Family: <i>Lachnospiraceae</i>), <i>Zoogloea</i> , <i>Acetobacterium</i>
8	<i>Blautia</i> (2), <i>Methanobrevibacter</i> , Unknown (Family: <i>Coriobacteriaceae</i>), <i>Ruminococcaceae</i> UCG-013, <i>Lactococcus</i>
9	<i>Mogibacterium</i> , Unknown (Family: <i>Rhodocyclaceae</i>), <i>Ruminococcaceae</i> NK4A214 group, Unknown (Family: <i>Gemmatimonadaceae</i>), <i>Sphingomonas</i> , <i>Sellimonas</i>
10	Unknown (Family: Gsoil-1167), <i>Romboutsia</i> , <i>Lactobacillus</i> , Unknown (Family: <i>Lachnospiraceae</i>), <i>Synergistes</i> , Unknown (Family: <i>Planctomycetaceae</i>)
11	<i>Bifidobacterium</i> (2), <i>Erysipelotrichaceae</i> UCG-003, <i>Holdemanella</i> (2), <i>Streptococcus</i>
12	<i>Trichococcus</i> , <i>Blautia</i> , <i>Roseburia</i> , <i>Holdemanella</i> , <i>Dorea</i> , Unknown (Family: <i>Lachnospiraceae</i>)

Table 8.3: Genera of the top six OTUs in each bin based on median relative abundance (SS tank).

Table 8.3 shows the genera of the six most abundant OTUs in each bin. Using this table and biological expertise or research, one can possibly infer and explain some of the links between the different covariates and some of the genera. However, we do not investigate any further here, as it is beyond the scope of this thesis and the SS is of less interest than the AS.

Harmonic regression coefficients

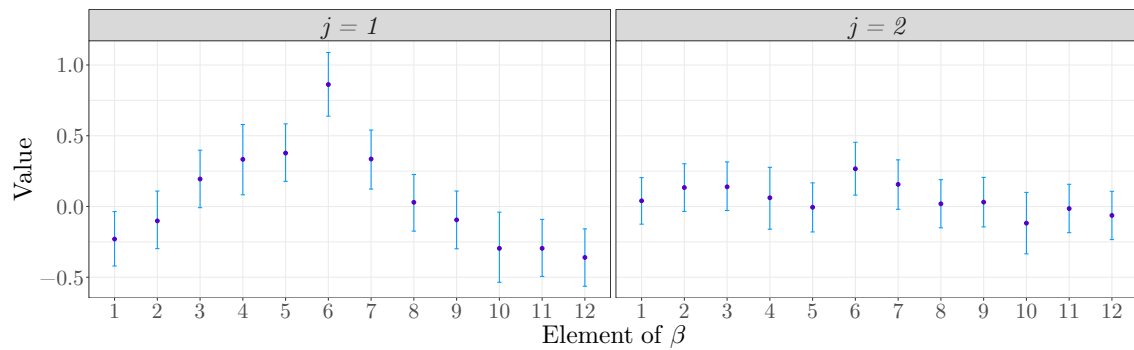


Figure 8.11: Posterior means (\circ) and 95% credible intervals (—) for β_j (SS tank).

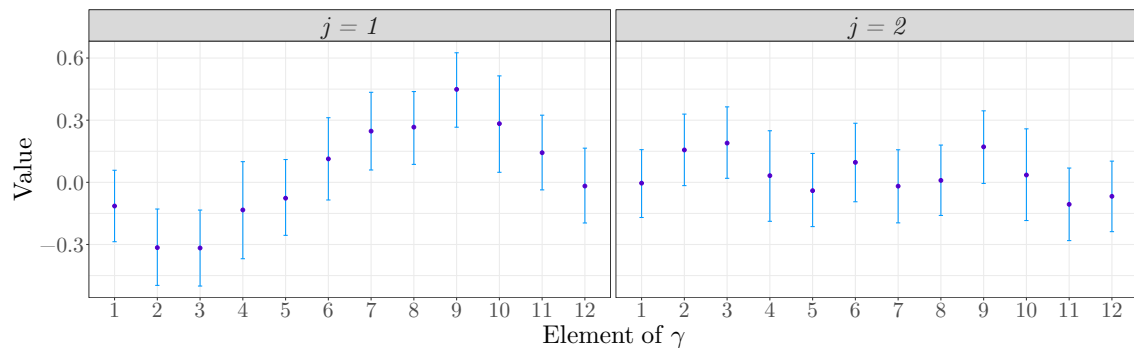


Figure 8.12: Posterior means (\circ) and 95% credible intervals (—) for γ_j (SS tank).

Figures 8.11 and 8.12 show the posterior means and CIs for the harmonic regression coefficients for the SS tank. For the first harmonics, we see a sinusoidal pattern again due to our clustering method. There are very small non-zero coefficients for the second harmonics in bins 3 and 6. Figure 8.13 shows the posterior means and 95% CIs for the time varying means μ_t plotted over the scaled log counts of each bin. The general pattern in each bin seems to be captured fairly within the CIs. For most bins, the harmonic regression coefficients (γ_j in particular) are smaller than in the AS tank, which is unsurprising since we did not see strong evidence for seasonality in general in the SS tank during our exploratory analysis.

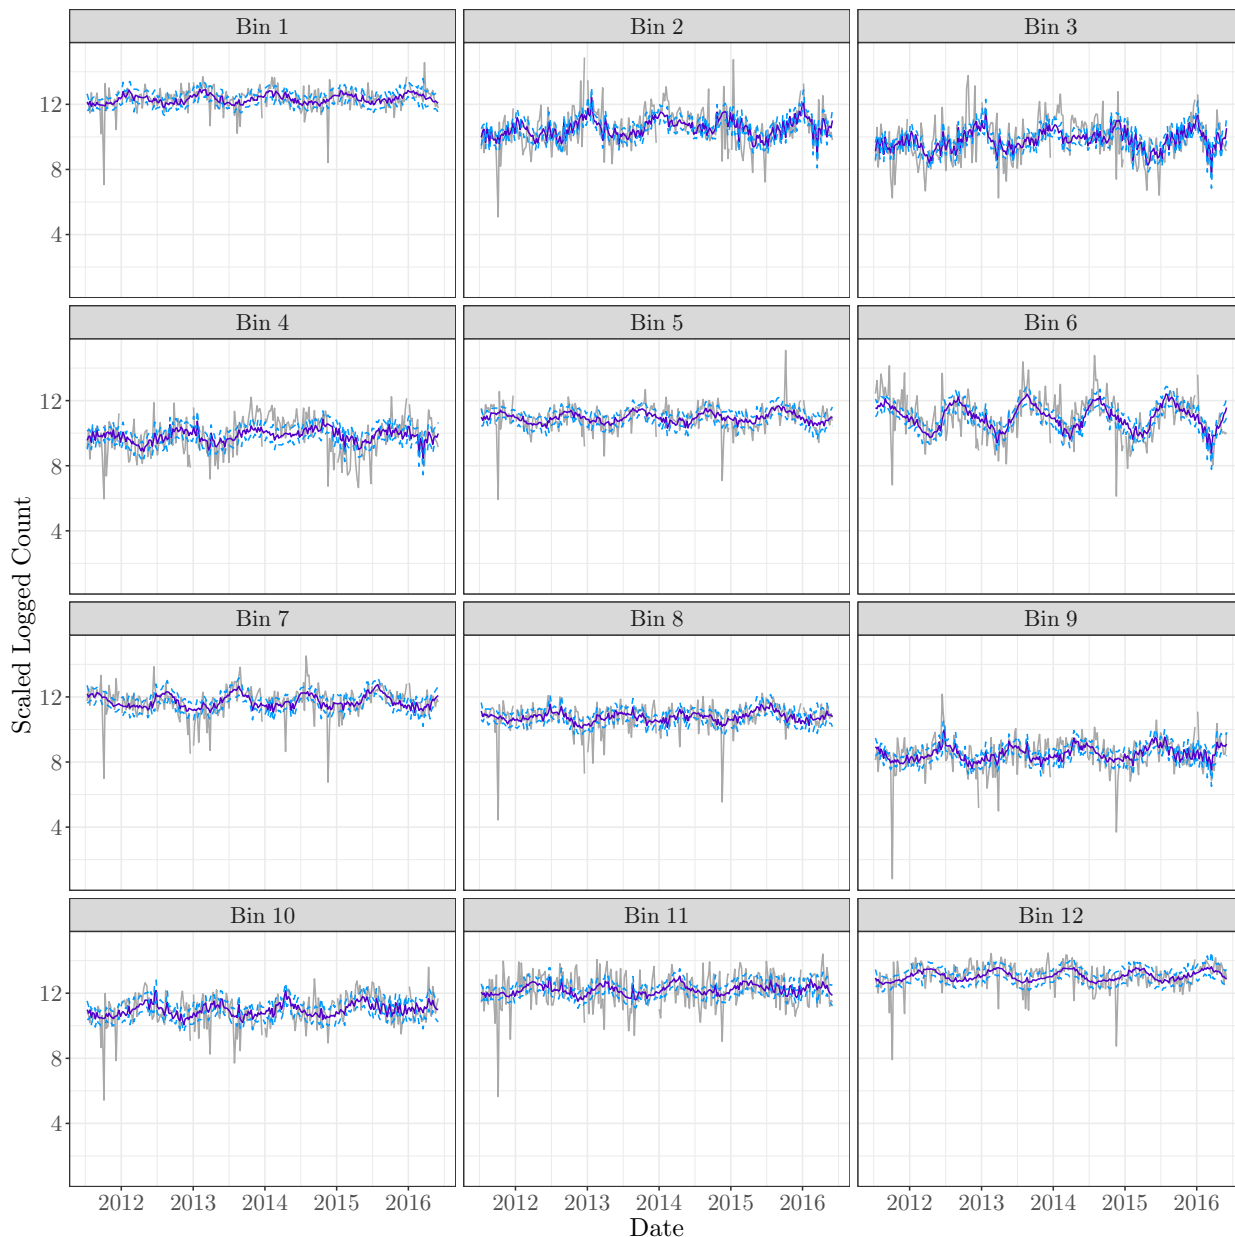


Figure 8.13: Posterior means (—) and 95% credible intervals (---) for the time varying means with scaled log counts (—) for each bin (SS tank).

8.2.2 Matrix of autoregressive coefficients

Figure 8.14 shows a heatmap of the posterior means of the autoregressive coefficients a_{jk} for the SS tank and they are also shown with their 95% CIs in Figure 8.15. Here we have a very sparse matrix with most coefficients being zero. Closer inspection of Figure 8.15 reveals that only two autoregressive coefficients have CIs that do not overlap zero, $a_{4,4}$ and $a_{10,10}$ with respective posterior means of 0.311 and 0.303. It seems counter-intuitive that only

two of the within-bin coefficients are non-zero. Perhaps this and the overall sparsity of the matrix that we see is due to the populations changing too quickly for the autocorrelations to be detectable from weekly measurements, as we suggested in the exploratory analysis in Chapter 6. We also have more chemical and environmental covariates in our model with more non-zero regression coefficients than in the AS tank. These covariates possibly explain more about the scaled log counts, although as we saw in Section 8.2.1 some of the bins seem to have no relationship with any of the covariates.

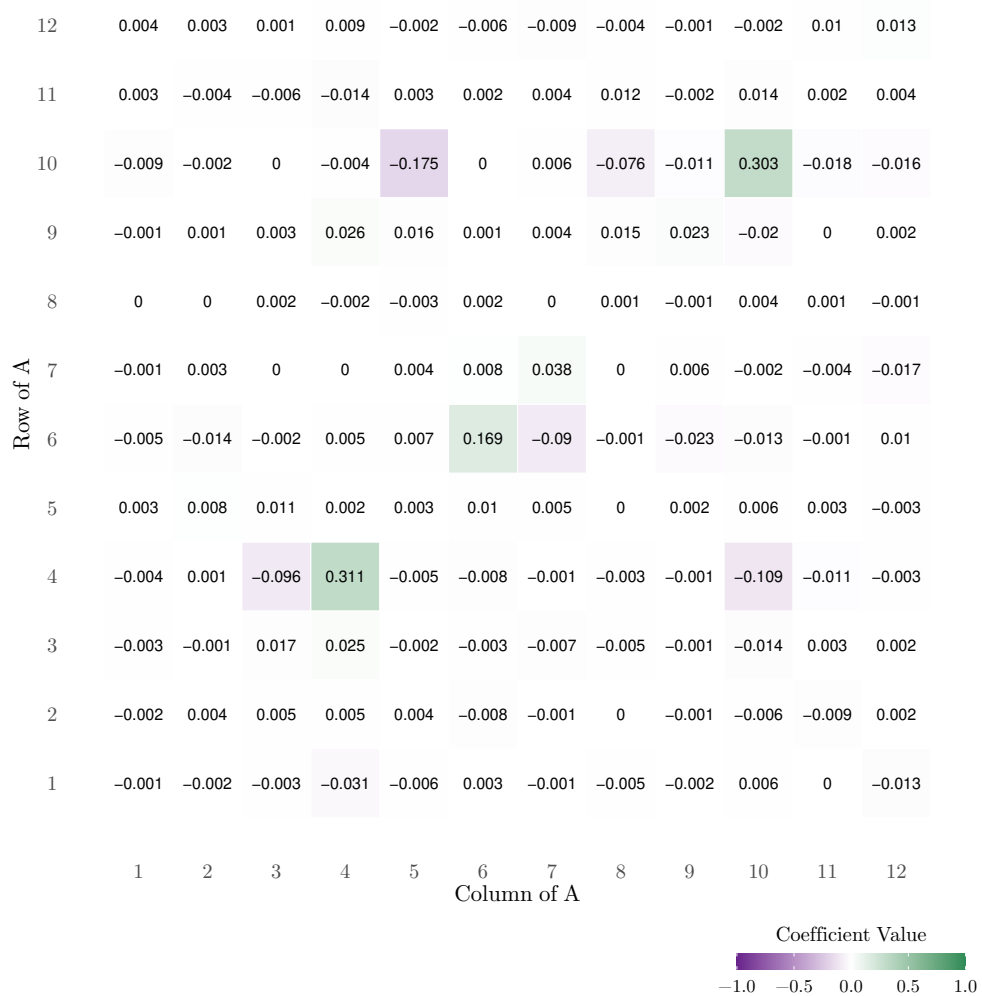


Figure 8.14: Heatmap of the posterior means of the autoregressive coefficients (SS tank).

8.2.3 Precision matrix

The posterior means for ω_0 and ω_1 in the precision matrix of the errors for the SS tank are 3.6906 and -1.7422 (to 4 d.p.) respectively with standard deviations 0.1081 and 0.0556

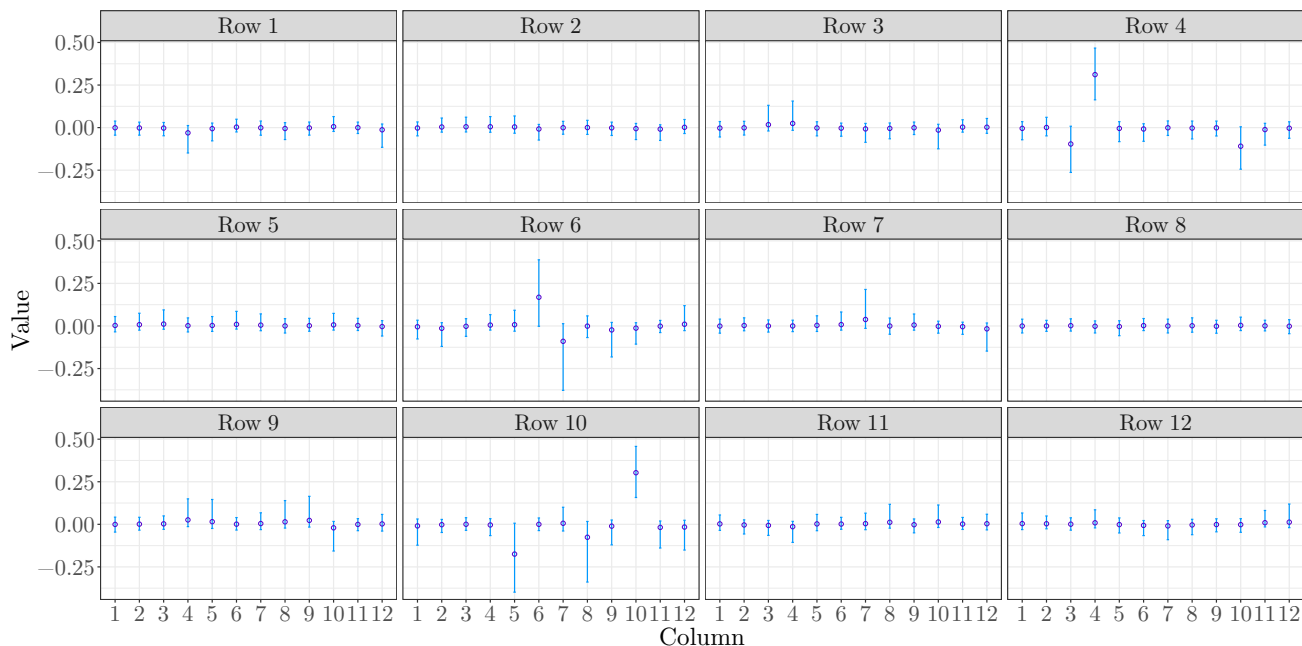


Figure 8.15: Posterior means (\circ) and 95% credible intervals (—) of the autoregressive coefficients (SS tank).

(to 4 d.p.). As we did with the AS tank, we plot the posterior means and 95% CIs for the lag- k correlations ρ_k , for $k = 1, \dots, 6$, of the errors (Figure 8.16). The CIs are all above zero which is evidence to suggest that there is between-bin correlation in the errors. The posterior means are slightly smaller than those observed for the AS. For the SS, the posterior mean for the within-bin error variance σ_{kk} is 0.8528 (to 4 d.p) with a standard deviation of 0.0413 (to 4 d.p.). This is larger than the posterior mean of the within-bin error variance for the AS tank (0.5337 to 4 d.p.). Perhaps this is because regression on the previous week's counts explains lots of variation in the AS, but not much in the SS.

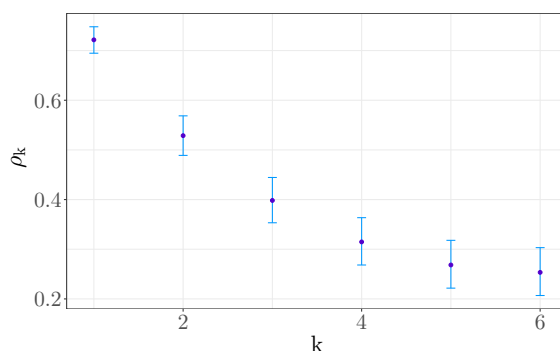


Figure 8.16: Posterior means (\circ) and 95% credible intervals (—) of the lag- k correlations ρ_k for $k = 1, \dots, 6$ (SS tank)

8.3 Results summary

In this chapter we have presented the results of the analysis on each tank and we now summarise the key findings in this section. In the AS tank, we found evidence to indicate that COD, ammonia and pH have positive relationships with some of the bins. After looking at the most abundant genera in each bin and using findings from research in biology, it would be reasonable to assume that these relationships are likely to be complex and non-linear. For example, we suggested that a potential predator-prey-like dynamic may exist between COD and some of the bins. In the SS tank, there was evidence to suggest that iron, fluoride and chloride have relationships with several bins and silicon, phosphate, ammonia and flow have a relationship with one bin. Overall, these findings were supported by our exploratory analysis.

Generally, in the AS tank the seasonal variation observed for each bin was captured fairly well by the time varying mean of our model. However, based on the posterior distributions for the second harmonics, we found that for some bins a simple sinusoidal shape may not be enough to describe their (transformed) counts. In the SS tank, the general pattern for each bin was captured fairly well by the time varying mean. We observed smaller harmonic regression coefficients too, which seems sensible given that there seemed to be less evidence of seasonality in our exploratory analysis.

As we might expect, within-bin positive autoregressive coefficients were found for all bins (except for bin 12) in the AS tank and we also identified several between-bin relationships. These relationships can potentially be explained by OTUs in each bin competing with each other for resources, suggesting that relationships between bins may also be complex. However, in the SS tank, we found the autoregressive matrix based on posterior means to be extremely sparse. We suspect that this is a result of the microbial populations changing too rapidly for autocorrelations to be detected from weekly counts.

Additionally, we saw that there was evidence for between-bin correlations in the errors. The within-bin error variance for the SS tank was found to be higher than for the AS tank. This is probably because regression on the previous week's counts explained a lot of variation in the AS but not in the SS, given how sparse the posterior autoregressive matrix was for the SS. Based on our results, it would seem that our model is not as useful for modelling the microbial populations in the SS as in the AS.

We opted to use a VAR(1) model instead of a Lotka-Volterra model because of the simplicity it offers. However, our results have shown us that both non-linear and predator-prey-like interactions may exist within the data, suggesting that a gLV model (see Section 2.2.2) might be more suitable for our data. This is discussed in more detail in the future work section (Section 9.3.2) of the following and final chapter.

Part III

Conclusions

Chapter 9

Conclusions and future work

The purpose of this final chapter is to highlight this thesis' contributions to the literature, summarise the overall conclusions and discuss potential future work.

9.1 Contributions

The main objective of this thesis was to develop biologically motivated stochastic models in two different applications: molecular phylogenetics and time-series metagenomics. We begin with a discussion of the contributions made by the phylogenetics part of the thesis. The work for this part of the thesis can also be found in Hannaford *et al.* (2020).

The CTMP that defines standard substitution models of DNA evolution in phylogenetics is typically assumed to be reversible and in its stationary distribution. These assumptions are made primarily for mathematical convenience, despite being refutable (by experimental evidence and biological theory) and restrictive in generating a root-invariant likelihood (Section 2.1.4). Both issues can be addressed by relaxing one or both simplifying assumptions. Among models in the literature which facilitate root inference, the most biologically credible are those which allow variation in sequence composition over time.

Heaps *et al.* (2014) introduced a class of non-homogeneous and non-stationary models with a locally reversible structure. Conditional on a given tree, each branch of the underpinning unrooted topology was associated with its own matrix from a class of reversible rate matrices. The distribution at the root of the tree was taken as the stationary distribution on the rooting branch. In this thesis, we advanced this idea so that each rate matrix comes from a class of non-reversible Lie Markov models; either RY5.6b or RY8.8 (Section 4.1). For both models, we provided a new parameterisation and gave an interpretation of the parameters (Sections 3.2.1 and 3.2.2). For the homogeneous RY5.6b model, we showed that the additive structure of the rate matrix makes it ill-suited to modelling evolutionary processes where the long-run proportions of each nucleotide are similar and

the transition-transversion rate ratio is high (Section 3.2.2). This provides an explanation for the poor fit that is reported for the RY5.6b model (Woodhams *et al.*, 2015). To our knowledge, this has hitherto gone unnoticed in the literature.

Our non-homogeneous Lie Markov models have a number of strengths. With fewer parameters than an analogous non-homogeneous general Markov model, they provide a parsimonious way of introducing local non-reversible structure into a non-stationary model. This yields an extra source of information about the root position, whilst retaining computational tractability in model-fitting. Moreover, because Lie Markov models are closed under matrix multiplication, our non-homogeneous extensions are mathematically consistent, meaning the distributions over DNA characters induced by a tree, and all its subtrees, could have arisen from the same family of non-homogeneous Lie Markov models.

Taking a Bayesian approach to inference, we describe a prior for the branch-specific parameters that encourages borrowing of strength between edges (Section 4.2). This has a regularising effect on the posterior distribution. We additionally describe an MCMC scheme for generating samples from the posterior (Section 4.3.2). Through extensive simulation experiments, we demonstrated, empirically, that the root position can be identified from the likelihood of our non-homogeneous models, and that increasing the number of sites in the alignment tends to lead to more accurate and precise inferences of all unknowns (Section 5.1.1). Whilst root inference for the non-homogeneous RY5.6b model was generally poor for larger trees, we showed that root inference under the non-homogeneous RY8.8 model remains strong, even in the face of prior-data conflict arising from an unbalanced rooted topology, though inference can be sensitive to the presence of long branches in the unrooted topology (Section 5.1.2).

We utilised our model and inferential procedures in a biological application concerning a challenging data set of *Drosophila*, in which simpler models typically fail to identify a plausible root position. In this analysis, our non-homogeneous RY8.8 model identified a rooted tree that was biologically credible (Section 5.2.2). We showed that this model had the highest marginal likelihood, indicating better fit to the data, while only doubling the computational time, compared to the homogeneous, stationary, reversible GTR model (Section 5.2.3).

In the metagenomics part of the thesis, we were given an OTU table for the AS and SS of a WWTP. The main goal was to model the counts of the OTUs and the interactions between them over time, whilst allowing for chemical and environmental effects, with particular interest in the AS tank. Microorganisms in the AS are responsible for biologically treating wastewater and are delicately balanced in a stable state. Understanding how this stable state is maintained is important to ensure a WWTP can continue functioning or, better still, be improved (Cyzdik-Kwiatkowska & Zielińska, 2016). As is commonly found in metagenomic studies, our data suffer from high-dimensionality and sparsity. Owing

to the evidence of seasonality in our data (Section 6.2) and the unsuitability of clustering by taxonomy (Section 7.1.1), we chose a seasonal phase-based clustering approach (Section 7.1.2) to address both issues.

Often, in time-series metagenomics, gLV differential equations are used to model non-linear dynamics of the microbial communities of interest. However, we chose a more parsimonious option and developed a Bayesian hierarchical VAR(1) model for our clustered data (bins), which is a simple first-order approximation to a gLV model (Section 2.2.2). The counts of the clustered data were positively skewed so we chose to log-transform the data, since a VAR(1) model assumes normality. The autoregressive matrix was assumed to be sparse, due to the circular time-ordering of the bins. Three priors that induce sparsity were considered: the spike and slab, the horseshoe and the regularised horseshoe. Piironen & Vehtari (2017) derived a systematic way of defining a prior that considers prior sparsity information for the global shrinkage parameter in the horseshoe (and the regularised horseshoe) for simple multiple linear regression. We extended their work by deriving a similar result in the case of a multivariate response vector (Section 7.3.3). Through simulations, we demonstrated that the regularised horseshoe was robust to varying degrees of sparsity and a sensible choice for our prior for the autoregressive coefficient matrix (Section 7.3.5). This hyperprior was chosen for the global shrinkage parameter, since a VAR(1) model can be expressed as a multiple linear regression with a multivariate response vector. Using this global shrinkage hyperprior necessitated scaling the log counts to have a variance roughly equal to one. We gave the errors of our model a symmetric, circulant, tri-diagonal precision matrix to complement our chosen clustering method (Section 7.4). To capture the seasonal variation in each bin, we used a harmonic regression to fit a time varying mean (Section 7.5). The chemical and environmental data were incorporated through the time varying mean.

We fitted the model to the data from each tank. We identified possible relationships amongst bins (Sections 8.1.2 and 8.2.2) and between bins and chemical and environmental covariates (Sections 8.1.1 and 8.2.1) by inspecting the posterior distributions obtained for the parameters in the model. After looking at the most abundant genera in each bin and findings from biological studies, we found evidence to suggest that these relationships are likely complex and non-linear. Our model was found to be more useful for modelling the microbial communities in the AS than in the SS. Our analyses also provided evidence to suggest that a gLV model may be better suited to modelling the data, which is discussed in Section 9.3.2. Nevertheless, our model was still able to provide an interesting insight into the dynamics of the microbial communities present in the AS and SS of the WWTP, which was the overarching aim of this part of the thesis.

9.2 Conclusions

In Chapter 2, we discussed standard phylogenetic models and their drawbacks (Section 2.1.4), namely their inability to infer the root position of a tree, and common approaches to modelling metagenomics data. We showed VAR(1) models to be a suitable choice for modelling our WWTP data, since they can be regarded as a linear approximation to gLV models (Section 2.2.2), which are commonly used in time-series metagenomic studies. An introduction to Bayesian inference was given, including MCMC methods, Bayesian phylogenetics (Section 2.3.2) and a brief description of HMC and the computer program *Stan* (Section 2.3.3).

Chapter 3 contained a discussion of existing models that facilitate root inference with a focus on two non-reversible Lie Markov models, the RY5.6b and RY8.8 models. These two models were extended in Chapter 4 by allowing each branch to have a different rate matrix, resulting in two non-homogeneous models. After applying our models to simulated data in Chapter 5, we showed that both models were able to infer root position, for a small number of taxa (six) but when we increased the number of taxa (to 12 and 24), inference under the non-homogeneous RY5.6b model was poor (Section 5.1.1). Root inference under the non-homogeneous RY8.8 model remained strong, even with the presence of prior-data conflict (Section 5.1.2). However, we found that inference can be sensitive to the presence of long branches in the underlying unrooted topology, suggesting investigation into branch length priors could be an interesting avenue of research (see Section 9.3.1). We also found that the posterior densities for the non-homogeneous RY8.8 model were more concentrated than the non-homogeneous RY5.6b model (Section 5.1.1) and we postulated that the additive structure of the RY5.6b rate matrix may be partially responsible for this. Finally, both models were shown to be useful for inferring unrooted trees as the posterior mode was the true unrooted topology in every analysis of the simulation study.

When applying our models to the challenging *Drosophila* data set (Section 5.2), we found the posterior densities for the non-homogeneous RY5.6b model to be more diffuse than for the non-homogeneous RY8.8 model, as we found in the simulation studies. After comparing models with marginal likelihood estimates (Section 5.2.3), we found that our non-homogeneous RY8.8 model had the best fit to the *Drosophila* data set. The non-homogeneous GTR model had the second highest marginal likelihood, followed by the non-homogeneous RY5.6b model with the third highest marginal likelihood, showing that non-homogeneous models have a better fit to this data set. Our results show that accounting for compositional bias in data by using non-homogeneous models can provide credible inferences and a reasonable fit to the data, as has been previously indicated in the literature (Yang & Roberts, 1995; Foster, 2004; Heaps *et al.*, 2014; Williams *et al.*, 2015). However, the non-homogeneous RY8.8 model was the only model to infer a biologically plausible

rooted tree, with a posterior probability of 0.9235. The non-homogeneous RY5.6b model is inflexible because it only has five degrees of freedom and a rate matrix with an additive structure (Section 3.2.2), which may explain why it was unable to infer a plausible tree. Meanwhile, the non-homogeneous GTR model has reversible rate matrices on each branch. This suggests that perhaps the non-reversible structure of the RY8.8 matrix was able to provide an additional source of root information. Thus, we have demonstrated that relaxing both assumptions of reversibility and stationarity can enable us to find biologically credible rooted trees for data that exhibits compositional heterogeneity.

In Chapter 6, we performed a detailed exploratory analysis of the WWTP data, where we identified possible relationships between chemical and environmental covariates and the relative abundances of the most abundant OTUs, genera and classes in the AS and SS. Temperature seemed to be correlated with many of the relative abundances of the AS and there was evidence of seasonality in the OTU data (Section 6.2.3). This was also observed in time-series plots of the relative abundances (Sections 6.2.1 and 6.2.2). However, seasonality was not as obvious in the SS data. The OTUs in the AS and the SS showed a time-varying mean but properties that seemed otherwise independent of time, suggesting it would be reasonable to fit a VAR model to the detrended data.

Chapter 7 commenced with a discussion of clustering methods and we concluded that a seasonal phase-based clustering approach was most suitable for our data (Section 7.1.2). Clustering based on taxonomy was deemed inappropriate (Section 7.1.1) because, after accounting for missing taxonomic information, the finest taxonomic rank we could cluster with was class. This is a coarse taxonomic rank, meaning each class can encompass many OTUs. The example we gave was the *Alphaproteobacteria*, which was one of the most abundant classes present in the AS. This class is known to be extremely diverse (Williams *et al.*, 2007), which suggests that the OTUs clustered into this class (and OTUs clustered by taxonomy, generally (Dam *et al.*, 2016)) will have very different abundance profiles.

We then gave a detailed discussion of the development of our Bayesian hierarchical VAR(1) model, where each advancement from the simple VAR(1) model (Section 7.2) was justified. In our simulation study (Section 7.3.5), which investigated the performance of different shrinkage priors for the matrix of autoregressive coefficients, we showed that the regularised horseshoe was robust to varying degrees of sparsity and, thus, a suitable prior for our model. As discussed in Section 9.1, our chosen clustering method influenced our choice of error precision matrix and model for the time varying mean.

After fitting the model to the clustered AS and SS data in Chapter 8 we inspected the posterior distributions of the parameters. We identified evidence to suggest that some of the bins in the AS have positive relationships with COD and ammonia (Section 8.1.1). We speculated on why these relationships may occur by identifying the most abundant OTUs in each bin and using findings from research on their corresponding genera in biology. Our

findings also suggested that predator-prey-like dynamics may occur between some of the bins and these two covariates and that the relationships may be complex and non-linear in nature. There was also evidence to suggest that a positive relationship between pH and bin 12 may exist. In the SS, we found evidence for several chemical and environmental covariates having a relationship with some of the bins (Section 8.2.1). We did not delve too deeply into why these relationships may exist because, as explained in Section 2.2.4, the microbial community in AS is of more interest.

We also found evidence indicating the existence of relationships amongst bins. As we did for the relationships between bins and chemical and environmental covariates, we looked at the most abundant OTUs in each bin and considered relevant research in the biology literature to provide potential explanations for these relationships (Sections 8.1.2). It is possible that negative autoregressive coefficients between bins can be attributed to the most abundant OTUs competing for resources. For example, evidence of a negative relationship between bin 5 and 6 of the AS was found. We suggested that this may be due to both bins containing (a high proportion of) OTUs from genera that reduce nitrite and sulphate, which they may be competing for. Not all relationships could be explained in this way but knowledge of their (potential) existence still provides valuable insight into the ecology of AS. *A posteriori* the matrix of autoregressive coefficients for the SS was found to be very sparse (Section 8.2.1), which we suspect could be due to the microbial populations changing too quickly to be detected by sampling every week.

Overall we found our model to be unsuitable for modelling the SS data but it provided insights into the microbial community dynamics of the AS. We found evidence of non-linear relationships in our data, which ultimately suggests that a stochastic Lotka-Volterra model (see Section 9.3.2) may be more appropriate, where the findings from our VAR(1) model could be used to simplify model structure.

9.3 Future work

In the final section of this thesis we discuss areas of further work that may be worth exploring. We first discuss possible future work for the phylogenetics part of the thesis before concluding with potential areas where our time-series metagenomics work can be extended.

9.3.1 Phylogenetics future work

There are a few possible ways the work from the phylogenetics part of the thesis could be extended. Firstly, in our non-homogeneous RY5.6b model, we allowed the ρ_b to vary from branch to branch while treating α as a global parameter and fixing it across the tree (Section 4.1). This model could be adapted by also allowing α to vary from branch to branch

so that under the non-homogeneous RY5.6b model we have $\mathcal{Q} = \{\boldsymbol{\alpha}, \boldsymbol{q}_1, \dots, \boldsymbol{q}_{B-1}\}$, where $\boldsymbol{\alpha} = (\alpha_1, \dots, \alpha_B)^T$. We could then adopt a prior for $\boldsymbol{\alpha}$ with an autoregressive structure to facilitate more borrowing of information between neighbouring branches. We suspect that this adjustment would not result in a better performing model than the non-homogeneous RY8.8 model, as our simulation study in Chapter 5 revealed that branch-specific parameter identifiability was weaker for the non-homogeneous RY5.6b model (Section 5.1.1), which we attributed to the additive structure of its rate matrix. Nevertheless, it would still be an interesting direction to explore.

The simulation study also revealed that root inference under the non-homogeneous RY8.8 model can be sensitive to the presence of long branches in the unrooted topology (Section 5.1.2). Standard phylogenetic priors for branch lengths assign little prior mass to long branches. For example, we gave our branch lengths an $\text{Exp}(10)$ prior, which means that the expected branch lengths are short (0.1) and 99.9% of its mass is below 0.691. Consequently, such priors favour rooting on long branches as it results in splitting a long branch into two shorter branches. Therefore, it would be worth investigating a joint prior for the branch lengths $\boldsymbol{\ell}$ and rooted topology τ that incorporates biological insight, for example, a prior that considers total distance from the root to the leaves, or a prior that places atoms of mass on biologically plausible trees, with a distribution that decays smoothly with increasing distance from these trees. Nye (2020) explored such a distribution by constructing a normal-like distribution in tree space. Currently, evaluation of the mass function is computationally challenging which would make its use as a prior difficult without further work.

In this thesis, we used rate matrices from two non-reversible Lie Markov families. However, there are 99 Lie Markov models in total (Woodhams *et al.*, 2015), many of which are also non-reversible. Further work could entail fitting non-homogeneous versions of other models that Woodhams *et al.* (2015) found to perform well. They found the most successful models tended to be parameter-rich, with at least eight parameters, although only seven data sets were used to test the models, so this does not necessarily mean that other models with fewer parameters should be immediately discounted. One of the better performing models was the Lie Markov model with the most parameters (12), the general Markov model, which is not the most parsimonious model. Nevertheless, it would be worthwhile to develop a model that has a general Markov rate matrix on each branch of the tree and compare it with our non-homogeneous models in applications to real and simulated data.

Finally, a relatively easy extension of our work would be to use the other *variants* of the RY5.6b and RY8.8 matrices. These two matrices are *RY-variants*, which means that they possess the symmetry condition of purine-pyrimidine pairing and their rows and columns are labelled in the order A, G, C, T. There are also *WS-variants* and *MK-*

variants, where the former orders the bases as A, T, G, C and the latter orders them as A, C, G, T. WS (Weak/Strong) models distinguish the Watson-Crick pairs AT and GC. According to Watson-Crick base-pairing adenine pairs with thymine with two hydrogen bonds and guanine pairs to cytosine with three hydrogen bonds, which means that the pair GC has a stronger bond than the pair AT. MK (aMino,Keto) models distinguish AC and GT (Woodhams *et al.*, 2015). The terms amino and keto refer to forms that DNA bases can take. It would be straightforward to fit MK- and WS-variants of our models, as it would simply require reordering of the rows and columns in the rate matrices.

9.3.2 Metagenomics future work

There are two main ways that the model developed in the metagenomics part of the thesis could be improved. Firstly, we could adapt our prior for the matrix of autoregressive coefficients to include more levels of shrinkage. Here we consider two possible approaches to developing the model in this way. In (7.29) in Section 7.3.3 we defined a shrinkage factor matrix \mathcal{K}_j for row j of the matrix of regression coefficients B in multiple linear regression with a multivariate response vector. Since a VAR(1) model can be written as a multiple linear regression, our findings from Section 7.3.3 hold for our VAR(1) model. Thus, \mathcal{K}_j can be associated to the j -th row of the matrix of autoregressive coefficients A and the extent to which the coefficients are shrunk is determined by the eigenvalues of \mathcal{K}_j . These eigenvalues are directly proportional to the global shrinkage parameter τ_A . As τ_A approaches zero, all eigenvalues of \mathcal{K}_j approach one and $\mathcal{K}_j \rightarrow \mathbf{I}_K$, which is complete shrinkage. Conversely, as $\tau_A \rightarrow \infty$, all eigenvalues of \mathcal{K}_j approach zero and $\mathcal{K}_j \rightarrow \mathbf{0}_K$, which represents no shrinkage. It might be beneficial to explore a model that allows each row of A to have its own global shrinkage parameter. This would mean that if one bin's (scaled log) counts did not seem to be affected by the other bins' (scaled log) counts from the previous time point, it would not affect the overall shrinkage of the autoregressive coefficients of all bins whereas, with our current model, the value of τ_A could be influenced by such a bin.

Griffin & Brown (2017) investigated hierarchical priors for regression models, where relationships between predictor covariates can be assumed and regression coefficients can be arranged in levels. They argued that regression coefficients at higher levels add complexity to a model and should be more aggressively shrunk towards zero. Furthermore, they should depend on the importance of regression coefficients at lower levels. The priors they developed allowed regression coefficients at one level to depend on the subset of the effect sizes at lower levels. One example they looked at was linear regression with interaction terms with *strong* or *weak heredity*. Strong heredity means that an interaction term is only included if both corresponding main effects are included, whereas weak heredity assumes that an interaction term is included if at least one main effect is included. They

described the main effects as the first level of regression coefficients and the interactions as the second level. Let y_i be a response, which is observed with covariates X_{i1}, \dots, X_{ip} . A linear model with interaction terms can be written as

$$y_i = \alpha + \sum_{j=1}^p X_{ij}\beta_j + \sum_{j=1}^p \sum_{k=1}^{j-1} X_{ij}X_{ik}\gamma_{jk} + \epsilon_i, \quad \text{for } i = 1, \dots, n,$$

where $\epsilon_i \sim N(0, \sigma^2)$. In the case of strong heredity, they use the prior

$$\beta_j \sim N(0, a_1\tau_j) \text{ and } \gamma_{jk} \sim N(0, a_2\lambda_{jk}\tau_j\tau_k), \quad (9.1)$$

where a_1 and a_2 are scale parameters. The prior variance of γ_{jk} is small if at least one of τ_j , τ_k and λ_{jk} is small, which means that the interaction term γ_{jk} will tend to be small if either λ_{jk} is small or at least one of β_j and β_k are small. In weak heredity, $a_2\lambda_{jk}\tau_j\tau_k$ in (9.1) is replaced with $a_2\lambda_{jk}\frac{1}{2}(\tau_j + \tau_k)$, which means that the interaction terms are small if λ_{jk} is small or both τ_j and τ_k are small.

We could introduce more levels of shrinkage following this approach, for example, we could let the diagonal elements of A be the first level of our autoregressive coefficients and non-diagonal elements be the second level. Then assuming strong heredity, we would have

$$a_{jj} \sim N(0, a_1\tau_j) \text{ and } a_{jk} \sim N(0, a_2\lambda_{jk}\tau_j\tau_k) \text{ for } j \neq k.$$

This would mean that the influence of the count of bin k on the count of bin j is small if one of λ_{jk} , τ_j or τ_k is small. In other words, a_{jk} is far from zero only if λ_{jk} , a_{jj} and a_{kk} are all far from zero. Under weak heredity, we would have

$$a_{jj} \sim N(0, a_1\tau_j) \text{ and } a_{jk} \sim N\left(0, a_2\lambda_{jk}\frac{1}{2}(\tau_j + \tau_k)\right) \text{ for } j \neq k.$$

This would mean that the influence of the count of bin k on the count of bin j is small if λ_{jk} is small or both τ_j and τ_k are small. The autoregressive coefficient between bins j and k is far from zero if either bin j or bin k has an autoregressive coefficient with itself that is far from zero (and λ_{jk} is far from zero). Different hyperpriors for τ_j , λ_{jk} , a_1 and a_2 could be explored and then tested in a simulation study similar to the study carried out in Section 7.3.5. Adopting a prior with this hierarchical structure would add more flexibility into our model and could lead to more adaptive shrinkage of autoregressive coefficients. This would become even more important if we chose to have a larger number of smaller-sized bins.

A potential area for improvement in our metagenomics work was the method we used to select chemical and environmental covariates to include in our model. As discussed in

Section 7.5.1, we selected covariates based on lag-one correlations between covariates and model residuals after fitting our model without any chemical environmental covariates to each data set. Instead, we could fit a model for each data set (AS and SS) that includes all covariates and use a shrinkage prior, such as the regularised horseshoe, to aid in identifying which covariates might be involved in the microbial community dynamics of the WWTP. We avoided doing this in this thesis to reduce the computational requirements of fitting the model.

Our findings in Chapter 8 indicated that complex, non-linear dynamics are present in the microbial communities of the AS and SS. This would suggest that a VAR(1) model, which assumes linear dynamics and is a linear approximation to a gLV model, could be improved upon. Moreover, our model did not prove to be useful in modelling the microbial populations in the SS. Therefore, instead of adapting our VAR(1) model in the ways described earlier in this section, a worthy direction of exploration would be to fit a stochastic gLV model (defined in Section 2.2.2), which assumes non-linear dynamics. A stochastic model, rather than a deterministic model, seems more appropriate for modelling the microbial community dynamics in the WWTP. This is because we do not have a complete model for the biological system and there are many other factors which can cause biological systems to behave unpredictably (Wilkinson, 2011). For example, there may be other environmental factors that have not been measured or microbes that have not been detected or sampled that could influence the system.

In Section 2.2.2 we gave the stochastic gLV model for a two species system in (2.9) and we gave the deterministic gLV model for K species in (2.7). Now we give a stochastic gLV model for our K bins of OTUs that would form the starting point for this area of further work. As well as the form given in (2.7), a deterministic gLV model can be written in vector notation as

$$\frac{d}{dt}\mathbf{y}(t) = \text{diag}(y_1(t), \dots, y_K(t)) [\mathbf{b} + \mathbf{A}\mathbf{y}(t)], \quad (9.2)$$

where $\mathbf{y}(t) = (y_1(t), \dots, y_K(t))^T$ with $y_i(t)$ being the (scaled log) count of bin i , \mathbf{b} is the vector of growth rates and \mathbf{A} is the matrix of pairwise interactions between bins. A stochastic gLV model for K bins of OTUs is given by

$$d\mathbf{y}(t) = \text{diag}(y_1(t), \dots, y_K(t)) [f(\mathbf{y}(t)) dt + g(\mathbf{y}(t)) d\mathbf{B}(t)], \quad (9.3)$$

where $\mathbf{B}(t) = (B_1(t), \dots, B_M(t))^T$ is an M -dimensional Brownian motion and $f : \mathbb{R}^K \rightarrow \mathbb{R}^K$ and $g : \mathbb{R}^K \rightarrow \mathbb{R}^{K \times M}$ are continuous functions (see Section 1 of Yin *et al.* (2009)). By stochastically perturbing \mathbf{b} by $\mathbf{b} \rightarrow \mathbf{B} + \Delta \dot{\mathbf{B}}(t)$ and $a_{ij} \rightarrow a_{ij} + \sigma_{ij} \dot{B}_j(t)$, where $\Delta = (\delta_{ij})$ and $\Sigma = (\sigma_{ij})$ are two $K \times K$ matrices and $\dot{\mathbf{B}}(t)$ is K -dimensional white noise, (9.2) becomes a stochastic Lotka-Volterra system with special linear functions f and g , as

determined by (9.3).

The model in (9.3) does not take into consideration environmental factors, such as our chemical and environmental covariates. Stein *et al.* (2013) introduced perturbation terms into their deterministic gLV model for intestinal microbes in mice and Dam *et al.* (2016) introduced interaction terms between environmental conditions and OTUs sampled from a lake. Our chemical and environmental covariates could be added into the model in a similar fashion. In a deterministic setting this would give us the following model:

$$\frac{d}{dt}y_i(t) = y_i(t) \left[b_i + \sum_{j=1} A_{ij}y_j(t) + \sum_{\ell=1}^L \psi_{i\ell}x_\ell(t) \right],$$

where $x_\ell(t)$ is the measurement of covariate ℓ at time t and Ψ is a $K \times L$ matrix of interaction terms between the covariates and bins. We can yield a stochastic version of this model in a number of ways. For example, we could, again, stochastically perturb \mathbf{b} and \mathbf{A} in a similar way as described above for the model without environmental factors. Inference would be more complex for this model than our VAR(1) model because we would need to introduce and tune approximations in the solution to the stochastic differential equations.

Appendix A

Additional proofs and derivations

A.1 Constraint on trace of the RY5.6b rate matrix

Recall that our instantaneous rate matrix for the RY5.6b model is given by

$$Q = \begin{pmatrix} * & \alpha + \rho_2 & \beta + \rho_3 & \beta + \rho_4 \\ \alpha + \rho_1 & * & \beta + \rho_3 & \beta + \rho_4 \\ \beta + \rho_1 & \beta + \rho_2 & * & \alpha + \rho_4 \\ \beta + \rho_1 & \beta + \rho_2 & \alpha + \rho_3 & * \end{pmatrix}. \quad (\text{A.1})$$

with the constraint $\boldsymbol{\rho} = (\rho_1, \rho_2, \rho_3, \rho_4)^T \in \mathcal{S}_4$. Suppose we set the parameters in (A.1) as $\alpha = b_1$, $\beta = b_2$ and $\boldsymbol{\rho} = (r_1, r_2, r_3, r_4)^T$, where we have imposed the constraint $\boldsymbol{\rho} \in \mathcal{S}_4$.

Suppose, without loss of generality (wlog), that the branch length is $t = 1$. Taking $\alpha = a_1$, $\beta = a_2$, $\boldsymbol{\rho} = (p_1, p_2, p_3, p_4)^T$ and branch length $t = \frac{1}{c}$ would then result in the same (valid) branch-length-scaled rate matrix (i.e. tQ) if we set

$$\begin{aligned} a_1 &= cb_1 + \frac{c}{4} - \frac{1}{4} \\ a_2 &= cb_2 + \frac{c}{4} - \frac{1}{4} \\ p_i &= \frac{1}{4} - \frac{c}{4} + cr_i, \quad i = 1, \dots, 4, \end{aligned} \quad (\text{A.2})$$

where $c \in ([4 \min\{b_1, b_2\} + 1]^{-1}, [1 - 4 \min\{r_1, r_2, r_3, r_4\}]^{-1})$. Thus, an additional constraint is needed.

A.1.1 Proof for boundary conditions of c

Upper bound

Suppose, wlog, $r_1 = \min\{r_1, r_2, r_3, r_4\}$ then (A.2) becomes

$$p_1 = \frac{1}{4} - \frac{c}{4} + cr_1.$$

By definition $p_1 \in [0, 1]$, so we have $\frac{1}{4} - \frac{c}{4} + cr_1 \in [0, 1]$.

We require

$$c < (1 - 4 \min\{r_1, r_2, r_3, r_4\})^{-1} = (1 - 4r_1)^{-1}. \quad (\text{A.3})$$

$$\begin{aligned} p_1 \geq 0 &\implies \frac{1}{4} - \frac{c}{4} + cr_1 \geq 0 \\ &\implies \frac{1}{4} \geq \frac{c}{4} - cr_1 \\ &\implies 1 \geq c - 4cr_1 \\ &\implies 1 \geq c(1 - 4r_1) \\ &\implies (1 - 4r_1)^{-1} \geq c, \text{ as required.} \end{aligned}$$

Lower bound (proof by contradiction)

Suppose $c \leq (4 \min\{b_1, b_2\} + 1)^{-1}$ and suppose that, wlog, $b_1 = \min\{b_1, b_2\}$. In other words,

$$c \leq (4b_1 + 1)^{-1} \implies c^{-1} \geq 4b_1 + 1. \quad (\text{A.4})$$

$a_1 = cb_1 + \frac{c}{4} - \frac{1}{4}$ and by definition $a_1 \geq 0$.

$$\begin{aligned} \text{So } cb_1 + \frac{c}{4} - \frac{1}{4} &\geq 0 \\ \implies cb_1 + \frac{c}{4} &\geq \frac{1}{4} \\ \implies 4cb_1 + c &\geq 1 \\ \implies c(4b_1 + 1) &\geq 1 \\ \implies c &\geq (4b_1 + 1)^{-1}. \quad (b_1 \geq 0 \text{ so } 4b_1 + 1 \neq 0) \end{aligned}$$

This is a contradiction to (A.4), hence $c > (4 \min\{b_1, b_2\} + 1)^{-1}$.

Understanding the range

We have $\alpha = b_1, \beta = b_2, \boldsymbol{\rho} = (r_1, r_2, r_3, r_4)^T \in \mathcal{S}_4$ and $b_1, b_2, \geq 0$. In (A.3), $\min\{r_1, r_2, r_3, r_4\} \in [0, 0.25]$ and suppose, wlog, $r_1 = \min\{r_1, r_2, r_3, r_4\}$. For the case where $r_1 = 0$, we have $(1 - 4r_1)^{-1} = (1 - 0)^{-1} = 1$. For the case where $r_1 = 0.25$, we have $(1 - 4r_1)^{-1} = (1 - 4 \times 0.25)^{-1} \rightarrow \infty$. So the range for the upper bound is $[1, \infty)$. Now take $(4 \min\{b_1, b_2\} + 1)^{-1}$, where $b_1, b_2 \in [0, \infty)$ and suppose, wlog, $b_1 = \min\{b_1, b_2\}$. We then have $(4b_1 + 1)^{-1}$. For the case where $b_1 = 0$, we have $(4b_1 + 1)^{-1} = (0 + 1)^{-1} = 1$. For the case where $b_1 \rightarrow \infty$, we have $(4b_1 + 1)^{-1} \rightarrow 0$. The range for the lower bound is $(0, 1]$. Therefore we have

$$\begin{aligned} (4 \min\{b_1, b_2\} + 1)^{-1} &< (1 - 4 \min\{r_1, r_2, r_3, r_4\})^{-1} \\ \implies (4 \min\{b_1, b_2\} + 1)^{-1} &< c < (1 - 4 \min\{r_1, r_2, r_3, r_4\})^{-1}. \end{aligned}$$

A.1.2 Setting the trace

We apply a constraint on the trace to fix the scale of \mathbf{Q} . This prevents compensatory rescaling of the branch lengths and overall scale of \mathbf{Q} . Under our constraint that $\boldsymbol{\rho} \in \mathcal{S}_4$, we have,

$$\text{tr}(\mathbf{Q}) = -(4\alpha + 8\beta + 3). \tag{A.5}$$

Since $\alpha, \beta \geq 0$, $\text{tr}(\mathbf{Q}) \leq 3$, we impose the constraint that the $\text{tr}(\mathbf{Q}) = -7$ so that

$$4\alpha + 8\beta = 4. \tag{A.6}$$

Now we can replace β in (A.1) with $\frac{1 - \alpha}{2}$. To summarise, we have $\boldsymbol{\rho}$ lying on the four dimensional simplex and $\alpha + 2\beta$ lying on the two dimensional simplex.

A.2 Spike and slab prior FCDs

A.2.1 Derivation of the FCD of $\tilde{\beta}$

The prior density for $\tilde{\beta}$ is

$$\pi(\tilde{\beta}) = (2\pi d^2)^{-\frac{pK}{2}} \exp\left\{-\frac{1}{2}\tilde{\beta}^T (d^{-2}\mathbf{I}_{pK}) \tilde{\beta}\right\}.$$

The likelihood is given by

$$\pi(\mathbf{y}^*|\tilde{\beta}, \Lambda, \Sigma) = (2\pi)^{-\frac{N}{2}} |\mathbf{I}_N \otimes \Sigma|^{-\frac{N}{2}} \exp\left\{-\frac{1}{2}(\mathbf{y}^* - \mathbf{W}^*\tilde{\beta})^T (\mathbf{I}_N \otimes \Sigma^{-1}) (\mathbf{y}^* - \mathbf{W}^*\tilde{\beta})\right\}.$$

Thus, the full conditional density is given by

$$\begin{aligned} \pi(\tilde{\beta}|\mathbf{y}^*, \Lambda, \Sigma) &\propto \pi(\tilde{\beta}) \pi(\mathbf{y}^*|\tilde{\beta}, \Lambda, \Sigma) \\ &\propto \exp\left\{-\frac{1}{2}\tilde{\beta}^T (d^{-2}\mathbf{I}_{pK}) \tilde{\beta}\right\} \\ &\quad \times \exp\left\{-\frac{1}{2}(\mathbf{y}^* - \mathbf{W}^*\tilde{\beta})^T (\mathbf{I}_N \otimes \Sigma^{-1}) (\mathbf{y}^* - \mathbf{W}^*\tilde{\beta})\right\} \\ &\propto \exp\left\{-\frac{1}{2}\tilde{\beta}^T \left(d^{-2}\mathbf{I}_{pK} + \mathbf{W}^{*T} (\mathbf{I}_N \otimes \Sigma^{-1}) \mathbf{W}^*\right) \tilde{\beta} - 2\tilde{\beta}^T \mathbf{W}^{*T} (\mathbf{I}_N \otimes \Sigma^{-1}) \mathbf{y}^*\right\}. \end{aligned}$$

This is (proportional to) the full conditional density corresponding to the FCD given in (7.12) in Section 7.3.1.

A.2.2 Derivation of the FCD of $\tilde{\beta}$ with $\Sigma = \tau^{-1}\mathbf{I}_K$

In the special case when $\Sigma = \tau^{-1}\mathbf{I}_q$, $\mathbf{I}_n \otimes \Sigma^{-1} = \tau\mathbf{I}_{NK}$. Then

$$\begin{aligned} \mathbf{V}_{\tilde{\beta}} &= \left(d^{-2}\mathbf{I}_{pK} + \tau\mathbf{W}^{*T}\mathbf{W}^*\right)^{-1} \\ &= \left(d^{-2}\mathbf{I}_{pK} + \tau\Lambda^{*T}\mathbf{W}^{*T}\mathbf{W}^*\Lambda^*\right)^{-1}. \end{aligned}$$

Now,

$$\begin{aligned} \mathbf{W}^T\mathbf{W} &= (\mathbf{W}_1^T, \dots, \mathbf{W}_N^T) \begin{pmatrix} \mathbf{W}_1 \\ \vdots \\ \mathbf{W}_N \end{pmatrix} \\ &= \sum_{i=1}^N \mathbf{W}_i^T \mathbf{W}_i \end{aligned}$$

$$\begin{aligned}
&= \sum_{i=1}^N (\mathbf{I}_K \otimes \mathbf{x}_i^T)^T (\mathbf{I}_K \otimes \mathbf{x}_i^T) \\
&= \sum_{i=1}^N (\mathbf{I}_K \otimes \mathbf{x}_i) (\mathbf{I}_K \otimes \mathbf{x}_i^T) \\
&= \sum_{i=1}^N \mathbf{I}_K \otimes \mathbf{x}_i \mathbf{x}_i^T \\
&= \mathbf{I}_K \otimes \sum_{i=1}^N \mathbf{x}_i \mathbf{x}_i^T \quad (\text{using } (\mathbf{A} \otimes \mathbf{B})(\mathbf{C} \otimes \mathbf{D}) = \mathbf{AC} \otimes \mathbf{BD}) \\
&= \mathbf{I}_K \otimes \mathbf{X}^T \mathbf{X} \\
&= \text{blockdiag} \left(\underbrace{\mathbf{X}^T \mathbf{X}, \dots, \mathbf{X}^T \mathbf{X}}_{K \text{ times}} \right).
\end{aligned}$$

Let $\Lambda_j^* = \text{diag}(\lambda_{1j}, \dots, \lambda_{pj})$, so that $\Lambda^* = \text{blockdiag}(\Lambda_1^*, \dots, \Lambda_K^*)$. Then

$$\begin{aligned}
\Lambda^{*T} \mathbf{W}^{*T} \mathbf{W}^* \Lambda^* &= \text{blockdiag}(\Lambda_1^*, \dots, \Lambda_K^*) \text{blockdiag}(\mathbf{X}^T \mathbf{X}, \dots, \mathbf{X}^T \mathbf{X}) \text{blockdiag}(\Lambda_1^*, \dots, \Lambda_K^*) \\
&= \text{blockdiag}(\Lambda_1^* \mathbf{X}^T \mathbf{X} \Lambda_1^*, \dots, \Lambda_K^* \mathbf{X}^T \mathbf{X} \Lambda_K^*).
\end{aligned}$$

Therefore

$$\begin{aligned}
V_{\bar{B}} &= \text{blockdiag} \left\{ (d^{-2} \mathbf{I}_p + \tau \Lambda_1^* \mathbf{X}^T \mathbf{X} \Lambda_1^*)^{-1}, \dots, (d^{-2} \mathbf{I}_p + \tau \Lambda_K^* \mathbf{X}^T \mathbf{X} \Lambda_K^*)^{-1} \right\} \\
&= \text{blockdiag} \left\{ (d^{-2} \mathbf{I}_p + \tau \mathbf{X}_1^{*T} \mathbf{X}_1^*)^{-1}, \dots, (d^{-2} \mathbf{I}_p + \tau \mathbf{X}_p^{*T} \mathbf{X}_p^*)^{-1} \right\},
\end{aligned}$$

where

$$\begin{aligned}
\mathbf{X}_j^* &= (\lambda_{1j} \mathbf{X}_{1:N,1}, \dots, \lambda_{pj} \mathbf{X}_{1:N,p}) \\
&= \mathbf{X} \Lambda_j^*,
\end{aligned}$$

and $\mathbf{X}_{1:N,j} = (\mathbf{X}_{1j}, \dots, \mathbf{X}_{Nj})^T$.

Similarly

$$\begin{aligned}
\mathbf{b}_{\bar{B}} &= \tau \mathbf{W}^{*T} \mathbf{y}^* \\
&= \tau \Lambda^{*T} \mathbf{W}^T \mathbf{y}^*,
\end{aligned}$$

where

$$\begin{aligned}
 \mathbf{W}^T \mathbf{y}^* &= (\mathbf{W}_1^T \dots, \mathbf{W}_N^T) \begin{pmatrix} \mathbf{y}_1 \\ \vdots \\ \mathbf{y}_N \end{pmatrix} \\
 &= \sum_{i=1}^N \mathbf{W}_i^T \mathbf{y}_i \\
 &= \sum_{i=1}^N (\mathbf{I}_K \otimes \mathbf{x}_i^T)^T \mathbf{y}_i \\
 &= \sum_{i=1}^N (\mathbf{I}_K \otimes \mathbf{x}_i) (\mathbf{y}_i \times 1) \\
 &= \sum_{i=1}^N (\mathbf{y}_i \otimes \mathbf{x}_i) \quad (\text{using } (\mathbf{A} \otimes \mathbf{B})(\mathbf{C} \otimes \mathbf{D}) = (\mathbf{AC} \otimes \mathbf{BD})) \\
 &= \begin{bmatrix} \sum_{i=1}^N y_{i1} \mathbf{x}_i \\ \vdots \\ \sum_{i=1}^N y_{iK} \mathbf{x}_i \end{bmatrix} \\
 &= \begin{bmatrix} \mathbf{X}^T \mathbf{y}_{1:N,1} \\ \vdots \\ \mathbf{X}^T \mathbf{y}_{1:N,K} \end{bmatrix},
 \end{aligned}$$

where $\mathbf{y}_{1:N,j}$ is column j of \mathbf{Y} , and so

$$\begin{aligned}
 \mathbf{b}_{\tilde{B}} &= \tau \begin{bmatrix} \Lambda_1^* \mathbf{X}^T \mathbf{y}_{1:N,1} \\ \vdots \\ \Lambda_K^* \mathbf{X}^T \mathbf{y}_{1:N,K} \end{bmatrix} \\
 &= \tau \begin{bmatrix} \mathbf{X}_1^{*T} \mathbf{y}_{1:N,1} \\ \vdots \\ \mathbf{X}_K^{*T} \mathbf{y}_{1:N,K} \end{bmatrix}.
 \end{aligned}$$

Hence

$$\mathbf{M}_{\tilde{B}} = \begin{bmatrix} \left(d^{-2} \mathbf{I}_p \tau \mathbf{X}_1^{*T} \mathbf{X}_1^* \right)^{-1} \tau \mathbf{X}_1^{*T} \mathbf{y}_{1:N,1} \\ \vdots \\ \left(d^{-2} \mathbf{I}_p \tau \mathbf{X}_K^{*T} \mathbf{X}_K^* \right)^{-1} \tau \mathbf{X}_K^{*T} \mathbf{y}_{1:N,K} \end{bmatrix}.$$

Recall that $\boldsymbol{\beta}^* = \left(\tilde{\beta}_{11}\lambda_{11}, \dots, \tilde{\beta}_{p1}\lambda_{p1}, \tilde{\beta}_{12}\lambda_{12}, \dots, \tilde{\beta}_{pK}\lambda_{pK} \right)^T$, so

$$\begin{aligned} \tilde{\boldsymbol{\beta}} &= \left(\tilde{\beta}_{11}, \dots, \tilde{\beta}_{p1}, \tilde{\beta}_{12}, \dots, \tilde{\beta}_{pK} \right)^T \\ &= \text{vec} \left(\tilde{\mathbf{B}} \right), \end{aligned}$$

thus we have shown that the columns of $\tilde{\boldsymbol{\beta}}$ are independent with mean given in (7.13) and variance given in (7.14) shown in Section 7.3.1.

A.2.3 Derivation of the FCD of λ_{jk} with $\Sigma = \tau^{-1}\mathbf{I}_K$

In the special case, when $\Sigma = \tau^{-1}\mathbf{I}_K$, we can write $(\mathbf{y}^* - \mathbf{W}\boldsymbol{\beta}^*)^T (\mathbf{I}_N \otimes \Sigma^{-1}) (\mathbf{y}^* - \mathbf{W}\boldsymbol{\beta}^*)$ as $\tau (\mathbf{y}^* - \mathbf{W}\boldsymbol{\beta}^*)^T (\mathbf{y}^* - \mathbf{W}\boldsymbol{\beta}^*)$, where

$$\begin{aligned} \mathbf{y}^* - \mathbf{W}\boldsymbol{\beta}^* &= \begin{pmatrix} \mathbf{y}_1 \\ \vdots \\ \mathbf{y}_N \end{pmatrix} - \begin{pmatrix} \mathbf{W}_1\boldsymbol{\beta}^* \\ \vdots \\ \mathbf{W}_N\boldsymbol{\beta}^* \end{pmatrix} \\ &= \begin{pmatrix} \mathbf{y}_1 \\ \vdots \\ \mathbf{y}_N \end{pmatrix} - \begin{bmatrix} \mathbf{x}_1^T\boldsymbol{\beta}_1^* \\ \vdots \\ \mathbf{x}_1^T\boldsymbol{\beta}_K^* \\ \vdots \\ \mathbf{x}_N^T\boldsymbol{\beta}_1^* \\ \vdots \\ \mathbf{x}_N^T\boldsymbol{\beta}_K^* \end{bmatrix} \\ &= \begin{bmatrix} y_{11} - \mathbf{x}_1^T\boldsymbol{\beta}_1^* \\ \vdots \\ y_{1K} - \mathbf{x}_1^T\boldsymbol{\beta}_K^* \\ y_{21} - \mathbf{x}_2^T\boldsymbol{\beta}_1^* \\ \vdots \\ y_{2K} - \mathbf{x}_2^T\boldsymbol{\beta}_K^* \\ \vdots \\ y_{NK} - \mathbf{x}_N^T\boldsymbol{\beta}_K^* \end{bmatrix}, \end{aligned}$$

in which

$$\boldsymbol{\beta}^* = \left(\boldsymbol{\beta}_1^{*T}, \dots, \boldsymbol{\beta}_K^{*T} \right)^T \text{ and } \boldsymbol{\beta}_k^* = (\beta_{1k}, \beta_{2k}, \dots, \beta_{pk})^T.$$

Now $\tau (\mathbf{y}^* - \mathbf{W}\boldsymbol{\beta}^*)^T (\mathbf{y}^* - \mathbf{W}\boldsymbol{\beta}^*)$ can be written as $\tau \sum_{i=1}^N \sum_{\ell=1} (y_{i\ell} - \mathbf{x}_i^T \boldsymbol{\beta}_\ell^*)^2$. Returning to the FCD of λ_{jk} , we have

$$\begin{aligned} c_{jk} &= \psi_{jk} \exp \left\{ -\frac{\tau}{2} (\mathbf{y}^* - \mathbf{W}\boldsymbol{\beta}^{**})^T (\mathbf{y}^* - \mathbf{W}\boldsymbol{\beta}^{**}) \right\} \\ &= \psi_{jk} \exp \left\{ \sum_{i=1}^N (y_{ik} - \mathbf{x}_i^T \boldsymbol{\beta}_k^{**})^2 \right\} \times \exp \left\{ -\frac{\tau}{2} \sum_{i=1}^N \sum_{\ell \neq k} (y_{i\ell} - \mathbf{x}_i^T \boldsymbol{\beta}_\ell^{**})^2 \right\}, \end{aligned}$$

where $\boldsymbol{\beta}^{**}$ is $\boldsymbol{\beta}^*$ with $\tilde{\beta}_{jk}$ as the $\{(k-1)p + j\}$ -th term and $\boldsymbol{\beta}_k^{**}$ is $\boldsymbol{\beta}_k^*$ with its j -th component as $\tilde{\beta}_{jk}$. Similarly, we have

$$\begin{aligned} d_{jk} &= (1 - \psi_{jk}) \exp \left\{ -\frac{\tau}{2} (\mathbf{y}^* - \mathbf{W}\boldsymbol{\beta}^{***})^T (\mathbf{y}^* - \mathbf{W}\boldsymbol{\beta}^{***}) \right\} \\ &= (1 - \psi_{jk}) \exp \left\{ \sum_{i=1}^N (y_{ik} - \mathbf{x}_i^T \boldsymbol{\beta}_k^{***})^2 \right\} \\ &\quad \times \exp \left\{ -\frac{\tau}{2} \sum_{i=1}^N \sum_{\ell \neq k} (y_{i\ell} - \mathbf{x}_i^T \boldsymbol{\beta}_\ell^{***})^2 \right\}, \end{aligned}$$

where $\boldsymbol{\beta}^{***}$ is $\boldsymbol{\beta}^*$ with 0 as the $\{(k-1)p + j\}$ -th term and $\boldsymbol{\beta}_k^{***}$ is $\boldsymbol{\beta}_k^*$ with its j -th component set to 0. Thus,

$$c_{jk} \propto \psi_{jk} \exp \left\{ \frac{\tau}{2} \sum_{i=1}^N (y_{ik} - \mathbf{x}_i^T \boldsymbol{\beta}_k^{**})^2 \right\},$$

which is equivalent to (7.16) shown in Section 7.3.1, and

$$d_{jk} \propto (1 - \psi_{jk}) \exp \left\{ \frac{\tau}{2} \sum_{i=1}^N (y_{ik} - \mathbf{x}_i^T \boldsymbol{\beta}_k^{***})^2 \right\},$$

which is equivalent to (7.17) shown in Section 7.3.1.

A.2.4 Derivation of the FCD of $\tilde{\mathbf{A}}$

To calculate the FCDs for $\tilde{\mathbf{A}}$ and λ_{jk} , it is convenient to write our model in (7.1) in a different way. First, we let $\mathbf{z}_t = \mathbf{y}_t - \boldsymbol{\mu}$, so that our model can be written, for $t = 2, \dots, N$, as

$$\mathbf{z}_t = \mathbf{A}\mathbf{z}_{t-1} + \boldsymbol{\epsilon}_t, \quad \boldsymbol{\epsilon}_t \sim \mathbf{N}_K(\mathbf{0}, \boldsymbol{\Sigma}) \quad (\text{A.7})$$

which can be written in matrix form as

$$\mathbf{Z}_t = \mathbf{Z}_{t-1}\mathbf{A}^T + \mathbf{E}, \quad (\text{A.8})$$

where

$$\mathbf{Z}_t = \begin{pmatrix} \mathbf{z}_2^T \\ \vdots \\ \mathbf{z}_N^T \end{pmatrix}, \quad \mathbf{Z}_{t-1} = \begin{pmatrix} \mathbf{z}_1^T \\ \vdots \\ \mathbf{z}_{N-1}^T \end{pmatrix}, \quad \mathbf{E} = \begin{pmatrix} \boldsymbol{\epsilon}_2^T \\ \vdots \\ \boldsymbol{\epsilon}_N^T \end{pmatrix}.$$

Note that (A.7) can be written as

$$\mathbf{z}_t = \mathbf{W}_t \mathbf{a} + \boldsymbol{\epsilon}_t, \quad (\text{A.9})$$

where

$$\mathbf{a} = \text{vec}(\mathbf{A}^T) = (a_{11}, a_{12}, \dots, a_{1K}, a_{21}, \dots, a_{KK}) \text{ and} \\ \mathbf{W}_t = \mathbf{I}_K \otimes \mathbf{z}_{t-1}^T.$$

We can rewrite (A.8) as

$$\mathbf{Z} = \mathbf{W} \mathbf{a} + \mathbf{e}, \quad (\text{A.10})$$

where

$$\mathbf{Z} = (\mathbf{z}_2^T, \dots, \mathbf{z}_N^T)^T, \quad \mathbf{W} = \begin{pmatrix} \mathbf{W}_2 \\ \vdots \\ \mathbf{W}_N \end{pmatrix}, \quad \mathbf{e} = (\boldsymbol{\epsilon}_2^T, \dots, \boldsymbol{\epsilon}_N^T)^T.$$

Firstly, note that $\mathbf{a} = (\tilde{a}_{11}\lambda_{11}, \dots, \tilde{a}_{1K}\lambda_{1K}, \tilde{a}_{21}\lambda_{21}, \dots, \tilde{a}_{KK}\lambda_{KK})^T$. Letting $\mathbf{W}^* = \mathbf{W}\boldsymbol{\Lambda}^*$, where $\boldsymbol{\Lambda}^* = \text{diag}(\lambda_{11}, \dots, \lambda_{1K}, \dots, \lambda_{21}, \dots, \lambda_{KK})$, and using the fact that $(\mathbf{I}_{N-1} \otimes \boldsymbol{\Sigma})^{-1} = (\mathbf{I}_{N-1} \otimes \boldsymbol{\Sigma}^{-1})$, the likelihood can be written as

$$\pi(\mathbf{Z} | \tilde{\mathbf{a}}, \boldsymbol{\lambda}, \boldsymbol{\Sigma}) = (2\pi)^{-\frac{(N-1)K}{2}} |\mathbf{I}_{N-1} \otimes \boldsymbol{\Sigma}|^{-\frac{(N-1)K}{2}} \\ \times \exp \left\{ -\frac{1}{2} (\mathbf{Z} - \mathbf{W}^* \tilde{\mathbf{a}})^T (\mathbf{I}_{N-1} \otimes \boldsymbol{\Sigma}^{-1}) (\mathbf{Z} - \mathbf{W}^* \tilde{\mathbf{a}}) \right\}. \quad (\text{A.11})$$

The prior density for $\tilde{\mathbf{a}}$ is

$$(2\pi d^2)^{-\frac{K^2}{2}} \exp \left\{ -\frac{1}{2} \tilde{\mathbf{a}}^T (d^{-2} \mathbf{I}_{K^2}) \tilde{\mathbf{a}} \right\}. \quad (\text{A.12})$$

The FCD for $\tilde{\mathbf{a}}$ is given by

$$\tilde{\mathbf{a}} | \mathbf{Z}, \boldsymbol{\Lambda}, \boldsymbol{\Sigma} \sim \mathbf{N}_{K^2}(\mathbf{M}_{\tilde{\mathbf{a}}}, \mathbf{V}_{\tilde{\mathbf{a}}}), \text{ where} \quad (\text{A.13}) \\ \mathbf{V}_{\tilde{\mathbf{a}}} = \left(d^{-2} \mathbf{I}_{K^2} + \mathbf{W}^{*T} (\mathbf{I}_{T-1} \otimes \boldsymbol{\Sigma}^{-1}) \mathbf{W}^* \right)^{-1}, \\ \mathbf{M}_{\tilde{\mathbf{a}}} = \mathbf{V}_{\tilde{\mathbf{a}}} \mathbf{b}_{\tilde{\mathbf{a}}}, \\ \mathbf{b}_{\tilde{\mathbf{a}}} = \mathbf{W}^{*T} (\mathbf{I}_{N-1} \otimes \boldsymbol{\Sigma}^{-1}) \mathbf{Z}.$$

Proof

The full conditional density for $\tilde{\mathbf{a}}$ is given by

$$\begin{aligned}
 \pi(\tilde{\mathbf{a}}|\mathbf{Z}, \cdot) &\propto \pi(\tilde{\mathbf{a}})\pi(\mathbf{Z}|\tilde{\mathbf{a}}, \cdot) \\
 &\propto \exp\left\{-\frac{1}{2}\tilde{\mathbf{a}}^T(d^{-2}I_{K^2})\tilde{\mathbf{a}}\right\} \exp\left\{-\frac{1}{2}(\mathbf{Z} - \mathbf{W}^*\tilde{\mathbf{a}})^T(I_{N-1} \otimes \Sigma^{-1})(\mathbf{Z} - \mathbf{W}^*\tilde{\mathbf{a}})\right\} \\
 &\propto \exp\left\{-\frac{1}{2}(\tilde{\mathbf{a}}^T(d^{-2}I_{K^2})\tilde{\mathbf{a}})\right\} \\
 &\quad \times \exp\left\{-\frac{1}{2}\left(\tilde{\mathbf{a}}^T\mathbf{W}^{*T}(I_{N-1} \otimes \Sigma^{-1})\mathbf{W}^*\tilde{\mathbf{a}} - 2\tilde{\mathbf{a}}^T\mathbf{W}^{*T}(I_{N-1} \otimes \Sigma^{-1})\mathbf{Z}\right)\right\} \\
 &= \exp\left\{-\frac{1}{2}\left[\tilde{\mathbf{a}}^T(d^{-2}I_{K^2} + \mathbf{W}^{*T}(I_{N-1} \otimes \Sigma^{-1})\mathbf{W}^*)\tilde{\mathbf{a}} - 2\tilde{\mathbf{a}}^T\mathbf{W}^{*T}(I_{N-1} \otimes \Sigma^{-1})\mathbf{Z}\right]\right\}.
 \end{aligned}$$

From the full conditional density of $\tilde{\mathbf{a}}$ we can see that we have the multivariate normal FCD described above in (A.13).

A.2.5 Derivation of the FCD of $\tilde{\Lambda}$ with $\Sigma = \tau^{-1}I_K$

In the special case when $\Sigma = \tau^{-1}I_K$, we have

$$I_{N-1} \otimes \Sigma^{-1} = \tau I_{K(N-1)}$$

and so

$$\begin{aligned}
 V_{\tilde{\mathbf{a}}} &= (d^{-2}I_{K^2} + \tau\mathbf{W}^{*T}\mathbf{W}^*)^{-1} \\
 &= (d^{-2}I_{K^2} + \tau\Lambda^{*T}\mathbf{W}^T\mathbf{W}\Lambda^*)^{-1}.
 \end{aligned}$$

Now,

$$\begin{aligned}
 \mathbf{W}^T\mathbf{W} &= (\mathbf{W}_2^T, \dots, \mathbf{W}_N^T) \begin{pmatrix} \mathbf{W}_2 \\ \vdots \\ \mathbf{W}_N \end{pmatrix} \\
 &= \sum_{t=2}^N \mathbf{W}_t^T\mathbf{W}_t \\
 &= \sum_{t=2}^N (\mathbf{I}_K \otimes \mathbf{z}_{t-1}^T)^T (\mathbf{I}_K \otimes \mathbf{z}_{t-1}) \\
 &= \sum_{t=2}^N (\mathbf{I}_K \otimes \mathbf{z}_{t-1}) (\mathbf{I}_K \otimes \mathbf{z}_{t-1}^T)
 \end{aligned}$$

$$\begin{aligned}
 &= \sum_{t=2}^N \mathbf{I}_K \otimes \mathbf{z}_{t-1} \mathbf{z}_{t-1}^T \\
 &= \mathbf{I}_K \otimes \sum_{t=2}^N \mathbf{z}_{t-1} \mathbf{z}_{t-1}^T \\
 &= \mathbf{I}_K \otimes \mathbf{Z}_{t-1}^T \mathbf{Z}_{t-1} \\
 &= \text{blockdiag} \left(\underbrace{\mathbf{Z}_{t-1}^T \mathbf{Z}_{t-1}, \dots, \mathbf{Z}_{t-1}^T \mathbf{Z}_{t-1}}_{K \text{ times}} \right).
 \end{aligned}$$

Let $\Lambda_j^* = \text{diag}(\lambda_{j1}, \dots, \lambda_{jK})$ so that $\Lambda^* = \text{blockdiag}(\Lambda_1^*, \dots, \Lambda_K^*)$. Then

$$\begin{aligned}
 \Lambda^{*T} \mathbf{W}^T \mathbf{W} \Lambda^* &= \text{blockdiag}(\Lambda_1^*, \dots, \Lambda_K^*) \text{blockdiag}(\mathbf{Z}_{t-1}^T \mathbf{Z}_{t-1}, \dots, \mathbf{Z}_{t-1}^T \mathbf{Z}_{t-1}) \\
 &\quad \times \text{blockdiag}(\Lambda_1^*, \dots, \Lambda_K^*) \\
 &= \text{blockdiag}(\Lambda_1^* \mathbf{Z}_{t-1}^T \mathbf{Z}_{t-1} \Lambda_1^*, \dots, \Lambda_K^* \mathbf{Z}_{t-1}^T \mathbf{Z}_{t-1} \Lambda_K^*).
 \end{aligned}$$

Therefore

$$\begin{aligned}
 \mathbf{V}_{\bar{a}} &= \text{blockdiag} \left\{ \left(d^{-2} \mathbf{I}_K + \tau \Lambda_1^* \mathbf{Z}_{t-1}^T \mathbf{Z}_{t-1} \Lambda_1^* \right)^{-1}, \dots, \left(d^{-2} \mathbf{I}_K + \tau \Lambda_K^* \mathbf{Z}_{t-1}^T \mathbf{Z}_{t-1} \Lambda_K^* \right)^{-1} \right\} \\
 &= \text{blockdiag} \left\{ \left(d^{-2} \mathbf{I}_K + \tau \mathbf{Z}_{t-1}^{1*T} \mathbf{Z}_{t-1}^{1*} \right)^{-1}, \dots, \left(d^{-2} \mathbf{I}_K + \tau \mathbf{Z}_{t-1}^{K*T} \mathbf{Z}_{t-1}^{K*} \right)^{-1} \right\},
 \end{aligned}$$

where

$$\begin{aligned}
 \mathbf{Z}_{t-1}^{j*} &= (\lambda_{j1} \mathbf{Z}_{1:N-1,1}, \dots, \lambda_{jK} \mathbf{Z}_{1:N-1,K}) \\
 &= \mathbf{Z}_{t-1}^* \Lambda_j^*.
 \end{aligned}$$

Note that here $\mathbf{Z}_{1:N-1,k} = (\mathbf{Z}_{1k}, \dots, \mathbf{Z}_{N-1,k})^T$.

Similarly

$$\begin{aligned}
 \mathbf{b}_{\bar{a}} &= \tau \mathbf{W}^{*T} \mathbf{Z} \\
 &= \tau \Lambda^{*T} \mathbf{W}^T \mathbf{Z},
 \end{aligned}$$

where

$$\begin{aligned}
 \mathbf{W}^T \mathbf{Z} &= (\mathbf{W}_2^T, \dots, \mathbf{W}_N^T) \begin{pmatrix} \mathbf{z}_2 \\ \vdots \\ \mathbf{z}_N \end{pmatrix} \\
 &= \sum_{t=2}^N \mathbf{W}_t^T \mathbf{z}_t
 \end{aligned}$$

$$\begin{aligned}
 &= \sum_{t=2}^N (\mathbf{I}_K \otimes \mathbf{z}_{t-1}^T)^T \mathbf{z}_t \\
 &= \sum_{t=2}^N (\mathbf{I}_K \otimes \mathbf{z}_{t-1}) \mathbf{z}_t \\
 &= \sum_{t=2}^N (\mathbf{I}_K \otimes \mathbf{z}_{t-1}) (\mathbf{z}_t \otimes \mathbf{1}) \\
 &= \sum_{t=2}^N (\mathbf{z}_t \otimes \mathbf{z}_{t-1}) \\
 &= \begin{bmatrix} \sum_{t=2}^N z_{t1} z_{t-1} \\ \vdots \\ \sum_{t=2}^N z_{tK} z_{t-1} \end{bmatrix} = \begin{bmatrix} \mathbf{Z}_{t-1}^T \mathbf{z}_{2:N,1} \\ \vdots \\ \mathbf{Z}_{t-1}^T \mathbf{z}_{2:N,K} \end{bmatrix}
 \end{aligned}$$

and so

$$\mathbf{b}_{\tilde{\mathbf{a}}} = \tau \begin{bmatrix} \Lambda_1^* \mathbf{Z}_{t-1}^T \mathbf{z}_{2:N,1} \\ \vdots \\ \Lambda_K^* \mathbf{Z}_{t-1}^T \mathbf{z}_{2:N,K} \end{bmatrix} = \tau \begin{bmatrix} \mathbf{Z}_{t-1}^{1*T} \mathbf{z}_{2:N,1} \\ \vdots \\ \mathbf{Z}_{t-1}^{K*T} \mathbf{z}_{2:N,K} \end{bmatrix}.$$

Hence

$$\mathbf{M}_{\tilde{\mathbf{a}}} = \begin{bmatrix} \left(d^{-2} \mathbf{I}_K + \tau \mathbf{Z}_{t-1}^{1*T} \mathbf{Z}_{t-1}^{1*} \right)^{-1} \tau \mathbf{Z}_{t-1}^{1*T} \mathbf{z}_{2:N,1} \\ \vdots \\ \left(d^{-2} \mathbf{I}_K + \tau \mathbf{Z}_{t-1}^{K*T} \mathbf{Z}_{t-1}^{K*} \right)^{-1} \tau \mathbf{Z}_{t-1}^{K*T} \mathbf{z}_{2:N,K} \end{bmatrix},$$

and so the rows of $\tilde{\mathbf{A}}$ are independent in the posterior distribution with the j -th row having mean $\left(d^{-2} \mathbf{I}_K + \tau \mathbf{Z}_{t-1}^{j*T} \mathbf{Z}_{t-1}^{j*} \right)^{-1} \tau \mathbf{Z}_{t-1}^{j*T} \mathbf{z}_{2:N,j}$ and variance $\left(d^{-2} \mathbf{I}_K + \tau \mathbf{Z}_{t-1}^{j*T} \mathbf{Z}_{t-1}^{j*} \right)^{-1}$. Hence we have a result for the VAR(1) model that is similar to what we found for the multivariate linear regression, except here, instead of the columns, it is the *rows* of $\tilde{\mathbf{A}}$ that exhibit independence.

A.2.6 The FCD of λ_{jk} (VAR(1) model)

The FCD for the (j, k) -th indicator parameter is given by

$$\lambda_{jk} | \mathbf{Z} \sim \text{Bern}(\tilde{p}_{jk}), \text{ with} \tag{A.14}$$

$$\tilde{p}_{jk} = c_{jk} / (c_{jk} + d_{jk}), \text{ where}$$

$$c_{jk} = p_{jk} \exp \left\{ -\frac{1}{2} (\mathbf{Z} - \mathbf{W} \mathbf{a}^*)^T (\mathbf{I}_{N-1} \otimes \Sigma^{-1}) (\mathbf{Z} - \mathbf{W} \mathbf{a}^*) \right\}$$

$$d_{jk} = (1 - p_{jk}) \exp \left\{ -\frac{1}{2} (\mathbf{Z} - \mathbf{W} \mathbf{a}^{**})^T (\mathbf{I}_{N-1} \otimes \Sigma^{-1}) (\mathbf{Z} - \mathbf{W} \mathbf{a}^{**}) \right\},$$

where \mathbf{a}^* is \mathbf{a} with its $\{(j-1)K+j\}$ -th element set to \tilde{a}_{jk} and \mathbf{a}^{**} is \mathbf{a} with its $\{(j-1)K+k\}$ -th element set to 0.

A.2.7 Derivation of the FCD of λ_{jk} with $\Sigma = \tau^{-1}\mathbf{I}_K$ (VAR(1) model)

Similarly, in the special case, when $\Sigma = \tau^{-1}\mathbf{I}_K$, we can write

$$(\mathbf{Z} - \mathbf{W}\mathbf{a})^T (\mathbf{I}_{N-1} \otimes \Sigma^{-1}) (\mathbf{Z} - \mathbf{W}\mathbf{a}) = \tau (\mathbf{Z} - \mathbf{W}\mathbf{a})^T (\mathbf{Z} - \mathbf{W}\mathbf{a}),$$

where

$$\begin{aligned} \mathbf{Z} - \mathbf{W}\mathbf{a} &= \begin{pmatrix} \mathbf{z}_2 \\ \vdots \\ \mathbf{z}_N \end{pmatrix} - \begin{pmatrix} \mathbf{W}_2\mathbf{a} \\ \vdots \\ \mathbf{W}_N\mathbf{a} \end{pmatrix} \\ &= \begin{pmatrix} \mathbf{z}_2 \\ \vdots \\ \mathbf{z}_N \end{pmatrix} - \begin{bmatrix} \mathbf{z}_1^T \mathbf{a}_1 \\ \vdots \\ \mathbf{z}_1^T \mathbf{a}_K \\ \vdots \\ \mathbf{z}_{N-1}^T \mathbf{a}_1 \\ \vdots \\ \mathbf{z}_{N-1}^T \mathbf{a}_K \end{bmatrix} \\ &= \begin{bmatrix} Z_{21} - \mathbf{z}_1^T \mathbf{a}_1 \\ \vdots \\ Z_{2K} - \mathbf{z}_1^T \mathbf{a}_K \\ Z_{31} - \mathbf{z}_1^T \mathbf{a}_1 \\ \vdots \\ Z_{NK} - \mathbf{z}_{N-1}^T \mathbf{a}_K \end{bmatrix}, \end{aligned}$$

in which $\mathbf{a} = (\mathbf{a}_1^T, \dots, \mathbf{a}_K^T)$ and $\mathbf{a}_j = (a_{j1}, a_{j2}, \dots, a_{jk})^T$. And so

$$\tau (\mathbf{Z} - \mathbf{W}\mathbf{a})^T (\mathbf{Z} - \mathbf{W}\mathbf{a}) = \tau \sum_{t=2}^N \sum_{\ell=1}^K (Z_{t\ell} - \mathbf{z}_{t-1}^T \mathbf{a}_\ell)^2.$$

Suppose we are interested in the FCD for λ_{jk} . Then we have

$$\begin{aligned} c_{jk} &= p_{jk} \exp \left\{ -\frac{\tau}{2} (\mathbf{Z} - \mathbf{W}\mathbf{a}^*)^T (\mathbf{Z} - \mathbf{W}\mathbf{a}^*) \right\} \\ &= p_{jk} \exp \left\{ -\frac{\tau}{2} \sum_{t=2}^N (Z_{tj} - \mathbf{z}_{t-1}^T \mathbf{a}_j^*)^2 \right\} \exp \left\{ -\frac{\tau}{2} \sum_{t=2}^N \sum_{\ell \neq j} (Z_{t\ell} - \mathbf{z}_{t-1}^T \mathbf{a}_\ell)^2 \right\}, \end{aligned}$$

where \mathbf{a}^* is \mathbf{a} with \tilde{a}_{jk} in its $\{(j-1)K+k\}$ -th term. Note that this only differs from \mathbf{a} in the k -th component of \mathbf{a}_j .

We also have

$$\begin{aligned} d_{jk} &= (1 - p_{jk}) \exp \left\{ -\frac{\tau}{2} (\mathbf{Z} - \mathbf{W}\mathbf{a}^{**})^T (\mathbf{Z} - \mathbf{W}\mathbf{a}^{**}) \right\} \\ &= (1 - p_{jk}) \exp \left\{ -\frac{\tau}{2} \sum_{t=2}^N (Z_{tj} - \mathbf{z}_{t-1}^T \mathbf{a}_j^{**})^2 \right\} \exp \left\{ -\frac{\tau}{2} \sum_{t=2}^N \sum_{\ell \neq j} (Z_{t\ell} - \mathbf{z}_{t-1}^T \mathbf{a}_\ell)^2 \right\}, \end{aligned}$$

where \mathbf{a}^{**} is \mathbf{a} with 0 in its $\{(j-1)K+k\}$ -th term. Note that this only differs from \mathbf{a} in the k -th component of \mathbf{a}_j .

Thus

$$\begin{aligned} c_{jk} &\propto p_{jk} \exp \left\{ -\frac{\tau}{2} \sum_{t=2}^N (Z_{tj} - \mathbf{z}_{t-1}^T \mathbf{a}_j^*)^2 \right\} \\ &= p_{jk} \exp \left\{ -\frac{\tau}{2} (\mathbf{Z}_{2:N,j} - \mathbf{Z}_{t-1} \mathbf{a}_j^*)^T (\mathbf{Z}_{2:N,j} - \mathbf{Z}_{t-1} \mathbf{a}_j^*) \right\} \end{aligned}$$

and

$$d_{jk} \propto (1 - p_{jk}) \exp \left\{ -\frac{\tau}{2} (\mathbf{Z}_{2:N,j} - \mathbf{Z}_{t-1} \mathbf{a}_j^{**})^T (\mathbf{Z}_{2:N,j} - \mathbf{Z}_{t-1} \mathbf{a}_j^{**}) \right\}.$$

This is analogous to the result found for multivariate linear regression.

A.3 Horseshoe prior FCDs

A.3.1 Derivation of the FCD of τ_β

To calculate the full conditional density of τ_β , we treat the density $\pi(\mathbf{B}|\tau_\beta, \cdot)$ as the likelihood and let $\boldsymbol{\beta}^* = \text{vec}(\mathbf{B})$. Then we have

$$\pi(\boldsymbol{\beta}^*|\tau_\beta, \cdot) = |2\pi\tau_\beta^2\Lambda^*|^{-\frac{1}{2}} \exp \left\{ -\frac{1}{2} \left(\boldsymbol{\beta}^{*T} \frac{1}{\tau_\beta^2} \Lambda^{*-1} \boldsymbol{\beta}^* \right) \right\}, \quad (\text{A.15})$$

where $\Lambda^* = \text{diag}(\lambda_{11}^2, \lambda_{21}^2, \dots, \lambda_{p1}^2, \lambda_{12}^2, \dots, \lambda_{pK}^2)$. The prior density for τ_β is

$$\pi(\tau_\beta) = \frac{2}{\pi\tau_{\beta_0} \left(1 + \frac{\tau_\beta^2}{\tau_{\beta_0}} \right)}. \quad (\text{A.16})$$

Substituting (A.15) and (A.16) into $\pi(\tau_\beta|\mathbf{B}, \cdot) \propto \pi(\tau_\beta) \pi(\boldsymbol{\beta}^*|\tau_\beta, \cdot)$ gives the density in (7.20) in Section 7.3.2. In the case of our VAR(1) model in (7.1) with a horseshoe prior for the matrix of autoregressive coefficients, the derivation for the FCD of τ_A shown in (7.23) is very similar to the derivation given here.

A.3.2 Derivation of the FCD of λ_{jk}

As we did for the global shrinkage parameter τ_β , we treat the density $\pi(\mathbf{B}|\lambda_{jk}, \cdot)$ as the likelihood. However, instead of using the vectorised form of \mathbf{B} , we can think of each element β_{jk} as being independent univariate normal quantities. Thus, we have

$$\pi(\mathbf{B}|\cdot) = \prod_{j=1}^p \prod_{k=1}^K (2\pi\tau_\beta^2\lambda_{jk}^2)^{-\frac{1}{2}} \exp\left\{-\frac{1}{2} \frac{\beta_{jk}^2}{\tau_\beta^2\lambda_{jk}^2}\right\}.$$

The prior for density for λ_{jk} is given by

$$\pi(\lambda_{jk}) = \frac{2}{\pi(1 + \lambda_{jk}^2)}.$$

The full conditional density for λ_{jk} , for $j = 1, \dots, p$, and $k = 1, \dots, K$, is

$$\begin{aligned} \pi(\lambda_{jk}|\mathbf{B}, \cdot) &\propto \pi(\lambda_{jk}) \pi(\mathbf{B}|\cdot) \\ &= \frac{2}{\pi(1 + \lambda_{jk}^2)} \prod_{a=1}^p \prod_{b=1}^K (2\pi\tau_\beta^2\lambda_{ab}^2)^{-\frac{1}{2}} \exp\left\{-\frac{1}{2} \frac{\beta_{ab}^2}{\tau_\beta^2\lambda_{ab}^2}\right\}, \end{aligned}$$

which is proportional to the density shown in (7.21) in Section 7.3.2.

In the case of our VAR(1) model in (7.1) with a horseshoe prior for the matrix of autoregressive coefficients, the derivation for the FCD of λ_{jk} shown in (7.24) is very similar to the derivation given here.

A.4 Prior for the global shrinkage parameter

A.4.1 Proof for $\bar{\beta}_j = (1 - \kappa_j)\hat{\beta}_j$

Here we show the proof for (7.26) in Section 7.3.3.

$$\begin{aligned} \bar{\beta}_j &= \tau_\beta^2\lambda_j^2 \left(\tau_\beta^2\lambda_j^2 + \sigma^2N^{-1}s_j^{-2}\right)^{-1} \hat{\beta}_j \\ &= \tau_\beta^2\lambda_j^2 \left(\frac{1}{\tau_\beta^2\lambda_j^2 + \sigma^2N^{-1}s_j^{-2}}\right) \hat{\beta}_j \\ &= \left(1 - \frac{\sigma^2N^{-1}s_j^{-2}}{\tau_\beta^2\lambda_j^2 + \sigma^2N^{-1}s_j^{-2}}\right) \hat{\beta}_j \\ &= \left(1 - \frac{1}{\tau_\beta^2\lambda_j^2\sigma^{-2}Ns_j^2 + 1}\right) \hat{\beta}_j \\ &= (1 - \kappa_j)\hat{\beta}_j. \end{aligned}$$

A.4.2 Derivation of conditional posterior for β^*

The likelihood for Y is given by

$$\begin{aligned}\pi(Y|B, \Sigma) &= \pi(E|\Sigma) \\ &= (2\pi)^{-pK/2} |\Sigma|^{-N/2} \exp \left\{ -\frac{1}{2} \text{tr} \left[\Sigma^{-1} (Y - XB)^T (Y - XB) \right] \right\} \\ &= (2\pi)^{-pK/2} |\Sigma|^{-N/2} \\ &\quad \times \exp \left\{ -\frac{1}{2} \text{tr} \left[\Sigma^{-1} (Y - X\hat{B})^T (Y - X\hat{B}) + \Sigma^{-1} (B - \hat{B})^T X^T X (B - \hat{B}) \right] \right\},\end{aligned}$$

using the fact that $\hat{B} = (X^T X)^{-1} X^T Y$. The prior density for B is

$$\begin{aligned}\pi(B) &= \pi(\beta^*) \\ &= (2\pi)^{-pK/2} |\tau_\beta^2 \Lambda^*|^{-\frac{1}{2}} \exp \left\{ -\frac{1}{2} \beta^{*T} (\tau_\beta^2 \Lambda^*)^{-1} \beta^* \right\}.\end{aligned}$$

The posterior density for B is proportional to $\pi(B)\pi(Y|B, \Sigma)$. Note that

$$\pi(Y|B, \Sigma) \propto (2\pi)^{-pK/2} |\Sigma|^{-N/2} \times \exp \left\{ -\frac{1}{2} \text{tr} \left[\Sigma^{-1} (B - \hat{B})^T X^T X (B - \hat{B}) \right] \right\}.$$

We wish to find $\pi(\beta^*|Y, \Sigma, \Lambda^*, \tau_\beta) \propto \pi(\beta^*)\pi(Y|\beta^*, \Sigma, \Lambda^*, \tau_\beta)$. Up to a constant of proportionality the likelihood $\pi(Y|\beta^*, \Sigma, \Lambda^*, \tau_\beta)$ can be written as

$$\exp \left\{ -\frac{1}{2} (\beta^* - \hat{\beta}^*)^T \left((X^T X)^{-1} \otimes \Sigma \right)^{-1} (\beta^* - \hat{\beta}^*) \right\}.$$

Thus we have

$$\begin{aligned}\pi(Y|\beta^*, \Sigma, \Lambda^*, \tau_\beta) \\ \propto \exp \left\{ -\frac{1}{2} \left(\beta^{*T} \left((X^T X)^{-1} \otimes \Sigma \right)^{-1} \beta^* \right) - 2\beta^{*T} \left((X^T X)^{-1} \otimes \Sigma \right)^{-1} \hat{\beta}^* \right\}.\end{aligned}$$

The posterior density for β^* is given by

$$\begin{aligned}
 \pi(\beta^* | Y, \Sigma, \Lambda^*, \tau_\beta) & \propto \exp \left\{ -\frac{1}{2} \beta^{*T} (\tau_\beta^2 \Lambda^*)^{-1} \beta^* \right\} \\
 & \quad \times \exp \left\{ -\frac{1}{2} \left(\beta^{*T} \left((X^T X)^{-1} \otimes \Sigma \right)^{-1} \beta^* \right) - 2 \beta^{*T} \left((X^T X)^{-1} \otimes \Sigma \right)^{-1} \hat{\beta}^* \right\}. \\
 & = \exp \left\{ -\frac{1}{2} \left[\beta^{*T} \left((\tau_\beta^2 \Lambda^*)^{-1} + \left((X^T X)^{-1} \otimes \Sigma \right)^{-1} \right) \beta^* \right. \right. \\
 & \quad \left. \left. - 2 \beta^{*T} \left((X^T X)^{-1} \otimes \Sigma \right)^{-1} \hat{\beta}^* \right] \right\}.
 \end{aligned}$$

This is the density corresponding to the posterior distribution described in (7.28) in Section 7.3.3.

A.4.3 Formulating \mathcal{K}_j

$$\begin{aligned}
 \tau_\beta^2 \Lambda_j (\tau_\beta^2 \Lambda_j + \frac{1}{N s_j^2} \Sigma)^{-1} & = \tau_\beta^2 \Lambda_j \Sigma^{-1} \left(\tau_\beta^2 \Lambda_j \Sigma^{-1} + \frac{1}{N s_j^2} \mathbf{I}_K \right)^{-1} \\
 & = N s_j^2 \tau_\beta^2 \Lambda_j \Sigma^{-1} \left(\tau_\beta^2 \Lambda_j \Sigma^{-1} N s_j^2 + \mathbf{I}_K \right)^{-1} \\
 & = \mathbf{I}_K - \left(\mathbf{I}_K + N s_j^2 \tau_\beta^2 \Lambda_j \Sigma^{-1} \right)^{-1}
 \end{aligned}$$

A.4.4 Expected value for $\mathcal{K}_{j,kk}$

We wish to find the prior expected value for the diagonal elements $\mathcal{K}_{j,kk}$ conditional on τ_β and Σ . Recall that we have

$$\mathbb{E}_{\Lambda | \tau_\beta, \Sigma} (\mathcal{K}_{j,kk}) = \int \frac{1}{(1 + a_{jk}^2 \lambda_{jk}^2)} \frac{2}{\pi (1 + \lambda_{jk}^2)} d\lambda_{jk},$$

where $a_{jk} = \sqrt{N} s_j \tau_\beta \sigma_k^{-1}$. This can be found by writing the integrand as a sum of its partial fractions. First note that we can write the integrand as $\frac{1}{(1 + a_{jk}^2 \lambda_{jk}^2)} \frac{2/\pi}{(1 + \lambda_{jk}^2)}$. We need to find $P(\lambda_{jk})$ and $Q(\lambda_{jk})$ such that

$$\frac{P(\lambda_{jk})}{(1 + a_{jk}^2 \lambda_{jk}^2)} + \frac{Q(\lambda_{jk})}{(1 + \lambda_{jk}^2)} = \frac{1}{(1 + a_{jk}^2 \lambda_{jk}^2)} \frac{2/\pi}{(1 + \lambda_{jk}^2)}.$$

It then follows that

$$P(\lambda_{jk}) + Q(\lambda_{jk}) + [P(\lambda_{jk}) + a_{jk}^2 Q(\lambda_{jk})] \lambda_{jk}^2 = \frac{2}{\pi},$$

which gives us

$$P(\lambda_{jk}) + Q(\lambda_{jk}) = \frac{2}{\pi} \text{ and} \tag{A.17}$$

$$P(\lambda_{jk}) + a_{jk}^2 Q(\lambda_{jk}) = 0. \tag{A.18}$$

Subtracting (A.18) from (A.17) we get

$$(1 - a_{jk}^2) Q(\lambda_{jk}) = \frac{2}{\pi} \implies Q(\lambda_{jk}) = \frac{2}{\pi(1 - a_{jk}^2)}.$$

From this, we get

$$\begin{aligned} P(\lambda_{jk}) &= \frac{2}{\pi} \left(1 - \frac{1}{1 - a_{jk}^2} \right) \\ &= \frac{2}{\pi} \left(\frac{-a_{jk}^2}{1 - a_{jk}^2} \right). \end{aligned}$$

Now we can write our integrand as a partial fraction and solve the integral to find our prior expected value for $\mathcal{K}_{j,kk}$ as follows:

$$\begin{aligned} \mathbb{E}_{\Lambda|\tau_\beta, \Sigma}(\mathcal{K}_{j,kk}) &= \frac{2}{\pi(1 - a_{jk}^2)} \left\{ \int_0^\infty \frac{1}{1 + \lambda_{jk}^2} d\lambda_{jk} - \int_0^\infty \frac{a_{jk}^2}{1 - a_{jk}^2 \lambda_{jk}^2} d\lambda_{jk} \right\} \\ &= \frac{2}{\pi(1 - a_{jk}^2)} \left\{ \left[\arctan(\lambda_{jk}) \right]_0^\infty - \left[a_{jk} \arctan(a_{jk} \lambda_{jk}) \right]_0^\infty \right\} \\ &= \frac{2}{\pi(1 - a_{jk}^2)} \left(\frac{\pi}{2}(1 + a_{jk}) \right) \\ &= \frac{1}{1 + a_{jk}}, \end{aligned}$$

which is result (7.37) stated in Section 7.3.3.

A.5 Conditional means and variances of a_{jk}

A.5.1 Spike and slab prior

If we have $X|Y \sim F(Y)$, then the conditional expectation and variance of X given Y are

$$\mathbb{E}[X] = \mathbb{E}_Y\{\mathbb{E}_X[X|Y]\} \text{ and} \tag{A.19}$$

$$\text{Var}(X) = \mathbb{E}_Y\{\text{Var}_X(X|Y)\} + \text{Var}_Y\{\mathbb{E}_Y[X|Y]\}. \tag{A.20}$$

Applying (A.19) to a_{jk} under a spike and slab prior we get the following conditional expectation for a_{jk} :

$$\begin{aligned} \mathbb{E}[a_{jk}] &= \mathbb{E}_{\lambda_{jk}} [\mathbb{E}_{a_{jk}}(a_{jk}|\lambda_{jk})] \\ &= \mathbb{E}(a_{jk}|\lambda_{jk} = 0) \Pr(\lambda_{jk} = 0) + \mathbb{E}(a_{jk}|\lambda_{jk} = 1) \Pr(\lambda_{jk} = 1) \\ &= 0 \end{aligned}$$

The conditional expectation of a_{jk} is zero under the spike and slab prior. Using (A.20) we get the following conditional variance for a_{jk} :

$$\begin{aligned} \text{Var}(a_{jk}) &= \mathbb{E}_{\lambda_{jk}} \{\text{Var}_{a_{jk}}(a_{jk}|\lambda_{jk})\} + \text{Var}_{\lambda_{jk}} \{\mathbb{E}_{\lambda_{jk}}[a_{jk}|\lambda_{jk}]\} \\ &= \text{Var}(a_{jk}|\lambda_{jk} = 0) \Pr(\lambda_{jk} = 0) + \text{Var}(a_{jk}|\lambda_{jk} = 1) \Pr(\lambda_{jk} = 1) + \text{Var}(0) \\ &= 0 + d^2 p + 0 \\ &= d^2 p \end{aligned}$$

A.5.2 Horseshoe prior

The conditional expectation for A_{ij} under the horseshoe prior is

$$\begin{aligned} \mathbb{E}[a_{jk}] &= \mathbb{E}_V [E_{a_{jk}}(a_{jk}|V)], \text{ where } V = (\lambda_{jk}\tau_A) \\ &= \mathbb{E}_V [0] \\ &= 0. \end{aligned}$$

The conditional variance for a_{jk} under the horseshoe prior is:

$$\begin{aligned} \text{Var}(a_{jk}) &= \mathbb{E}_V \{\text{Var}_{a_{jk}}(a_{jk}|V)\} + \text{Var}_V \{\mathbb{E}_V[a_{jk}|V]\} \\ &= \mathbb{E}_V [V^2] + \text{Var}_V(0) \\ &= \mathbb{E}[(\lambda_{jk}\tau_A)^2], \end{aligned}$$

which is undefined. The conditional variance is undefined because the expectation and variance of a Cauchy distribution is undefined. The same is also true for the regularised horseshoe due to its Cauchy hyperpriors for τ_A and λ_{jk} .

A.6 Reparameterisation of symmetric, circulant, tridiagonal precision matrix

Recall that a $(K \times K)$ symmetric, circulant, tridiagonal precision matrix has the following form:

$$\Sigma^{-1} = \begin{pmatrix} \omega_0 & \omega_1 & 0 & 0 & \cdots & 0 & 0 & 0 & \omega_1 \\ \omega_1 & \omega_0 & \omega_1 & 0 & \cdots & 0 & 0 & 0 & 0 \\ 0 & \omega_1 & \omega_0 & \omega_1 & \cdots & 0 & 0 & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots & \vdots & \vdots \\ 0 & 0 & 0 & 0 & \cdots & 0 & \omega_1 & \omega_0 & \omega_1 \\ \omega_1 & 0 & 0 & 0 & \cdots & 0 & 0 & \omega_1 & \omega_0 \end{pmatrix}.$$

Its eigenvalues are $\eta_j = \omega_0 + 2\omega_1 \cos(2\pi j/K)$ for $j = 0, 1, \dots, K-1$. It is positive definite if

$$\eta_j > 0 \Leftrightarrow \omega_0 > -2\omega_1 \cos(2\pi j/K) \quad \text{for all } j.$$

Suppose that $\omega_1 < 0$, then we require

$$-\omega_0/(2\omega_1) > \cos(2\pi j/K).$$

This is true if $-\omega_0/(2\omega_1) > 1$, which can only occur when $\omega_0 > 0$. Now suppose $\omega_1 > 0$. Then we require

$$-\omega_0/(2\omega_1) < \cos(2\pi j/K).$$

This is true if $-\omega_0/(2\omega_1) < -1$, i.e. $\omega_0/(2\omega_1) > 1$, which can only occur when $\omega_0 > 0$. Altogether, this means we require $\omega_0/(2|\omega_1|) > 1$, i.e. $\omega_0 > 2|\omega_1|$, whilst $\omega_1 \in \mathbb{R}$. Parameterisations over fixed regions are generally nicer, so we parameterise in terms of $\bar{\omega}_0$ and $\bar{\omega}_1$, where

$$\begin{aligned} \omega_0 &= (\varpi_0 + \varpi_1) / \sqrt{2} \\ \omega_1 &= (\varpi_0 - \varpi_1) / 2\sqrt{2} \\ \varpi_0 &= \sqrt{2}(\omega_0 + 2\omega_1) / 2 \\ \varpi_1 &= \sqrt{2}(\omega_0 - 2\omega_1) / 2, \end{aligned}$$

where $-\infty < \omega_1 < \infty, \omega_0 > 2|\omega_1|$, which is equivalent to $\varpi_0 > 0, \varpi_1 > 0$.

Appendix B

MCMC algorithms

B.1 Algorithms for shrinkage parameters in VAR(1) model with horseshoe prior

B.1.1 Metropolis-within-Gibbs step for λ_{jk}

1. Initialise chain with $\lambda_{jk}^{(0)}$. Set counter $i = 1$.
2. Generate a proposed value $\lambda_{jk}^* \sim \text{LN}\left(\log(\lambda_{jk}^{(i-1)}), \sigma_\lambda^2\right)$.
3. Evaluate acceptance probability

$$\alpha\left(\lambda_{jk}^* | \lambda_{jk}^{(i-1)}\right) = \min \left\{ 1, \log \left(\frac{\pi(\lambda_{jk}^* | \mathbf{A}) q(\lambda_{jk}^{(i-1)} | \lambda_{jk}^*)}{\pi(\lambda_{jk}^{(i-1)} | \mathbf{A}) q(\lambda_{jk}^* | \lambda_{jk}^{(i-1)})} \right) \right\}.$$

Dropping the $(i-1)$ notation, we have $\alpha\left(\lambda_{jk}^* | \lambda_{jk}^{(i-1)}\right) = \min\{1, A\}$, where

$$\begin{aligned} A &= \log \left(\frac{\frac{2}{\pi(1+\lambda_{jk}^{*2})} (\lambda_{jk}^*)^{-1} \exp \left\{ -\frac{1}{2} \left(\frac{a_{jk}}{\tau_A \lambda_{jk}^*} \right)^2 \right\} \frac{1}{\lambda_{jk} \sigma_\lambda \sqrt{2\pi}} \exp \left\{ \frac{-(\log(\lambda_{jk}) - \log(\lambda_{jk}^*))^2}{2\sigma_\lambda^2} \right\}}{\frac{2}{\pi(1+\lambda_{jk}^2)} (\lambda_{jk})^{-1} \exp \left\{ -\frac{1}{2} \left(\frac{a_{jk}}{\tau_A \lambda_{jk}} \right)^2 \right\} \frac{1}{\lambda_{jk}^* \sigma_\lambda \sqrt{2\pi}} \exp \left\{ \frac{-(\log(\lambda_{jk}^*) - \log(\lambda_{jk}))^2}{2\sigma_\lambda^2} \right\}} \right) \\ &= -\log(1 + \lambda_{jk}^{*2}) + \log(1 + \lambda_{jk}^2) + \frac{1}{2} \left[\left(\frac{a_{jk}}{\tau_A \lambda_{jk}} \right)^2 - \left(\frac{a_{jk}}{\tau_A \lambda_{jk}^*} \right)^2 \right] \end{aligned}$$

4. Set $\lambda_{jk}^{(j)} = \lambda_{jk}^*$ with probability $\alpha\left(\lambda_{jk}^* | \lambda_{jk}^{(i-1)}\right)$, otherwise set $\lambda_{jk}^{(j)} = \lambda_{jk}^{(i-1)}$.
5. Set $i = i + 1$. Go to step 2.

B.1.2 Metropolis-within-Gibbs step for τ_A

1. Initialise chain with $\tau_A^{(0)}$. Set counter $i = 1$.
2. Generate a proposed value $\tau_A^* \sim \text{LN}(\log(\tau_A^{(i-1)}), \sigma_{\tau_A}^2)$.
3. Evaluate the acceptance probability

$$\alpha(\tau_A^* | \tau_A^{(i-1)}) = \min \left\{ 1, \log \left(\frac{\pi(\tau_A^* | A) q(\tau_A | \tau_A^*)}{\pi(\tau_A | A) q(\tau_A^* | \tau_A)} \right) \right\}$$

Dropping the $(i-1)$ notation, we have $\alpha(\tau_A^* | \tau_A^{(i-1)}) = \min \{1, A\}$

$$\begin{aligned} A &= \log \left(\frac{2}{\pi \tau_0 (1 + \tau_A^{*2} / \tau_0)} \right) - \log \left(\frac{2}{\pi \tau_0 (1 + \tau_A^2 / \tau_0)} \right) - 0.5 \log \left(\prod_{k=1}^{K^2} \Lambda_{kk}^* \right) \\ &\quad - \frac{K^2}{2} \log(2\pi \tau_A^{*2}) + 0.5 \log \left(\prod_{k=1}^{K^2} \Lambda_{kk}^* \right) - 0.5 \left(\frac{1}{\tau_A^{*2}} \mathbf{a}^T \Lambda^{*-1} \mathbf{a} \right) \\ &\quad + 0.5 \left(\frac{1}{\tau_A^2} \mathbf{a}^T \Lambda^{*-1} \mathbf{a} \right) + \log \left(\frac{1}{\tau_A \sigma_{\tau_A} \sqrt{2\pi}} \right) - \log \left(\frac{1}{\tau_A \sigma_{\tau_A^*} \sqrt{2\pi}} \right) \\ &= -\log \left(1 + \frac{\tau_A^{*2}}{\tau_0} \right) + \log \left(1 + \frac{\tau_A^2}{\tau_0} \right) + \frac{K^2}{2} \left(\log(\tau_A^2) - \log(\tau_A^{*2}) \right) \\ &\quad - \log(\tau_A) + \log(\tau_A^*) + 0.5 \left(\frac{1}{\tau_A^2} - \frac{1}{\tau_A^{*2}} \right) \left(\mathbf{a}^T \Lambda^{*-1} \mathbf{a} \right) \Big\} \\ &= \log \left(1 + \frac{\tau_A^2}{\tau_0} \right) - \log \left(1 + \frac{\tau_A^{*2}}{\tau_0} \right) + (K^2 - 1) \left(\log(\tau_A) - \log(\tau_A^*) \right) \\ &\quad + 0.5 \left(\frac{1}{\tau_A^2} - \frac{1}{\tau_A^{*2}} \right) \left(\mathbf{a}^T \Lambda^{*-1} \mathbf{a} \right). \end{aligned}$$

4. Set $\tau_A^{(i)} = \tau_A^*$ with probability $\alpha(\tau_A^* | \tau_A^{(i-1)})$, otherwise set $\tau_A^{(i)} = \tau_A^{(i-1)}$.
5. Set $i = i + 1$.

Appendix C

Code

C.1 Stan code

Listing C.1: Code for the VAR(1) model used for the AS tank. The code for the SS tank is very similar.

```
functions { // No functions
}

data {
  int<lower = 1> K; // No. of bins / columns in data
  int<lower = 1> N; // No. of time points / rows in data
  int<lower = 1> N_miss; // No. time points missing from data
  matrix[N-N_miss,K] y_obs; // Observed data
  int<lower = 1, upper = N> which_y_miss[N_miss]; // Indices of missing time points
  int<lower = 1, upper = N> which_y_obs[N-N_miss]; // Indices of observed time points
  row_vector[N] intercept_X; // Intercept term for design matrix
  int<lower=1> L; // No. of chemical/environmental
  // covariates + 1 (for intercept)
  int<lower=0> N_Nitrate_obs; // Number of observed Nitrate values
  int<lower=0> N_COD_obs; // Number of observed COD values
  int<lower=0> N_Ammonia_obs; // Number of observed Ammonia values
  int<lower=0> N_pH_obs; // Number of observed pH values
  int<lower=0> N_Phosphate_obs; // Number of observed phosphate values

  row_vector[N_Nitrate_obs] Nitrate_obs; // Observed values for Nitrate
  row_vector[N_COD_obs] COD_obs; // Observed values for COD
  row_vector[N_Ammonia_obs] Ammonia_obs; // Observed values for Ammonia
  row_vector[N_pH_obs] pH_obs; // Observed values for pH
  row_vector[N_Phosphate_obs] Phosphate_obs; // Observed values for phosphate

  // Indices of missing values for each covariate
  int <lower = 1, upper = N> which_Nitrate_miss[N-N_Nitrate_obs];
  int <lower = 1, upper = N> which_COD_miss[N-N_COD_obs];
  int <lower = 1, upper = N> which_Ammonia_miss[N-N_Ammonia_obs];
  int <lower = 1, upper = N> which_pH_miss[N-N_pH_obs];
  int <lower = 1, upper = N> which_Phosphate_miss[N-N_Phosphate_obs];
  // Indices of observed values for Phosphate
  int <lower = 1, upper = N> which_Nitrate_obs[N_Nitrate_obs];
  int <lower = 1, upper = N> which_COD_obs[N_COD_obs];
  int <lower = 1, upper = N> which_Ammonia_obs[N_Ammonia_obs];
  int <lower = 1, upper = N> which_pH_obs[N_pH_obs];
  int <lower = 1, upper = N> which_Phosphate_obs[N_Phosphate_obs];

  int<lower=1> Nh; // Number of harmonics for mu
  real<lower = 0> global_scale; // tau_0 (scale for global shrinkage)
  real<lower=0> slab_scale; // Slab scale for the regularised horseshoe
}
```

```

real<lower=0> slab_df;          // Slab DOF for the regularised horseshoe
real m_beta; // Prior mean for beta
real<lower=0> v_beta; // Prior variance for beta
real m_gamma; // Prior mean for gamma
real<lower=0> v_gamma // Prior variance for gamma
real<lower=1> nu_local; // Degrees of freedom for half-t prior for tau.A
real<lower=1> nu_global; // Degrees of freedom for half-t prior for lambda
// nu_global = nu_local = 1 gives half-Cauchy.

real<lower=0> s_sigma; // Prior sd for lognormal for sigma
real<lower=0> c_error; // Prior coefficient of variation for precision parameters.
real a_B; // Prior mean for means_B
real<lower=0> b_B; // Prior sd for means_B
real<lower=2> c_B; // Prior shape for sigma-sq-B
real<lower=0> d_B; // Prior scale for sigma-sq-B
real<lower=0> alpha_phi; // Shape 1 for beta prior for phi
real<lower=0> beta_phi; // Shape 2 for beta prior for phi
real<lower = (L-2)> nu; // Prior degrees of freedom for Sigma_X
cov_matrix[(L-1)] Sigma_Sigma_X; // Prior scale matrix for Sigma_X
}

transformed data {
// sin. and cos. values for mu
matrix[Nh,N] sin_mat;
matrix[Nh,N] cos_mat;
real pi_const = (2*pi())/52;
for(j in 1:Nh) {
for(t in 1:N) {
sin_mat[j,t] = pi_const*j*t;
cos_mat[j,t] = sin_mat[j,t];
}
}
sin_mat = sin(sin_mat);
cos_mat = cos(cos_mat);
}

parameters {
matrix[K,K] z;
matrix[Nh,K] beta; // Beta params (sin. fourier coefficient) for mu
matrix[Nh,K] gamma; // Gamma params (cos. fourier coefficient) for mu
matrix[N_miss, K] y_miss; // Missing y values
real<lower=0> d0;
real<lower=0> d1;
real<lower=0> caux;
real<lower=0> aux1_global;
real<lower=0> aux2_global;
matrix<lower=0>[K,K] aux1_local;
matrix<lower=0>[K,K] aux2_local;
real<lower=0> sigma;
vector[L] means_B; // Prior means for alpha
real<lower=0> sigma_sq_B[L]; // Prior variance parameters for alphas
matrix[L,K] B_tilde;
row_vector[N-N.Nitrate_obs] Nitrate_miss;
row_vector[N-N.COD_obs] COD_miss;
row_vector[N-N.Ammonia_obs] Ammonia_miss;
row_vector[N-N.pH_obs] pH_miss;
row_vector[N-N.Phosphate_obs] Phosphate_miss;
vector<lower=0, upper=1>[L-1] phi;
cov_matrix[L-1] Sigma_X;
}

transformed parameters {
real<lower=0> tau_A; // Global shrinkage parameter
matrix<lower=0>[K,K] lambda; // Local shrinkage parameter
matrix<lower=0>[K,K] lambda_tilde; // 'Truncated' local shrinkage parameter
real<lower=0> c;
matrix[K,K] A; // Autoregressive coefficient matrix
real<lower=0> c0;
real c1;
matrix[K,K] prec_mat; // Precision matrix of errors
matrix[N,K] alpha; // Alpha (intercept for mu)
matrix[N,K] mu; // Time varying mean
matrix[L,K] B; // Regression coefficients for alpha
matrix[L,N] X; // Design matrix (chem. and env. covariates)
matrix[N,K] y; // Data constructed from y_obs and y_miss
}

```

```

y[which_y_obs,] = y_obs;
y[which_y_miss,] = y_miss;

X[2,which_Nitrate_miss] = Nitrate_miss;
X[2,which_Nitrate_obs] = Nitrate_obs;
X[3,which_COD_miss] = COD_miss;
X[3,which_COD_obs] = COD_obs;
X[4,which_Ammonia_miss] = Ammonia_miss;
X[4,which_Ammonia_obs] = Ammonia_obs;
X[5,which_pH_miss] = pH_miss;
X[5,which_pH_obs] = pH_obs;
X[6,which_Phosphate_miss] = Phosphate_miss;
X[6,which_Phosphate_obs] = Phosphate_obs;
X[1,] = intercept_X;

// Error precision matrix
c0 = (d0+d1)/sqrt(2); // c0 main diagonal of precision
c1 = (d0-d1)/(2*sqrt(2)); // c1 super-, sub-diagonals of precision.
// top right, bottom left corners

prec_mat = rep.matrix(0.0, K,K);
prec_mat[1,K] = c1;
prec_mat[K,1] = c1;
prec_mat[K,K] = c0;
for(j in 1:(K-1)) {
  prec_mat[j,j] = c0;
  prec_mat[j,(j+1)] = c1;
  prec_mat[(j+1),j] = c1;
}

// mu
for(l in 1:L) {
  B[1,] = B_tilde[1,] * sigma_sq_B[1];
}
for(i in 1:N) {
  alpha[i,] = (B'*X[,i])';
  mu[i,] = alpha[i,];
  for(j in 1:Nh) {
    mu[i,] += (beta[j,]*sin_mat[j,i] + gamma[j,]*cos_mat[j,i]);
  }
}

// A
c = slab_scale * sqrt(caux);
tau_A = aux1_global * sqrt(aux2_global) * global_scale * sigma;
for(i in 1:K) {
  lambda[i,] = aux1_local[i,] .* sqrt(aux2_local[i,]);
  lambda_tilde[i,] = sqrt(c^2*square(lambda[i,]) ./ (c^2 +tau_A^2*square(lambda[i,])));
  A[i,] = z[i,] .* lambda_tilde[i,]*tau_A;
}
}

model {
  real c_inv = 1/c_error^2;
  vector[L-1] mean_X;

  // Prior for error precision matrix (reparameterisation)
  sigma ~ lognormal(0, s_sigma);
  d0 ~ gamma(c_inv, sqrt(2)*sigma^2*c_inv);
  d1 ~ gamma(c_inv, sqrt(2)*sigma^2*c_inv);

  // Chem. / env. covariates
  phi ~ beta(alpha_phi, beta_phi);
  Sigma_X ~ inv_wishart(nu, Sigma_Sigma_X);
  for(i in 2:N) {
    mean_X = phi .* X[2:L,i-1];
    X[2:L,i] ~ multi_normal(mean_X, Sigma_X);
  }

  // Prior for B (alpha)
  means_B ~ normal(a_B, b_B);

```

```

sigma_sq_B ~ inv_gamma(c.B, d.B);
for(l in 1:L) {
  B_tilde[l,] ~ normal(means.B[l],1);
}

// Prior for beta and gamma
for(j in 1:Nh) {
  beta[j,] ~ normal(m_beta, sqrt(v_beta));
  gamma[j,] ~ normal(m_gamma, sqrt(v_gamma));
}

// Half t-priors for lambdas and tau, and inverse-gamma for c^2
// Half-Cauchy if nu_local = nu_global = 1
for(i in 1:K) {
  z[i,] ~ normal(0,1);
  aux1_local[i,] ~ normal(0,1);
  aux2_local[i,] ~ inv_gamma(0.5*nu_local, 0.5*nu_local);
}
aux1_global ~ normal(0,1);
aux2_global ~ inv_gamma(0.5*nu_global, 0.5*nu_global);

caux ~ inv_gamma(0.5*slab_df, 0.5*slab_df);

// Mu and Likelihood
for(i in 2:N) {
  y[i,] ~ multi_normal_prec(mu[i,]'+A*(y[(i-1),]'-mu[i-1,]'), prec_mat);
}
}

```

Appendix D

Additional tables

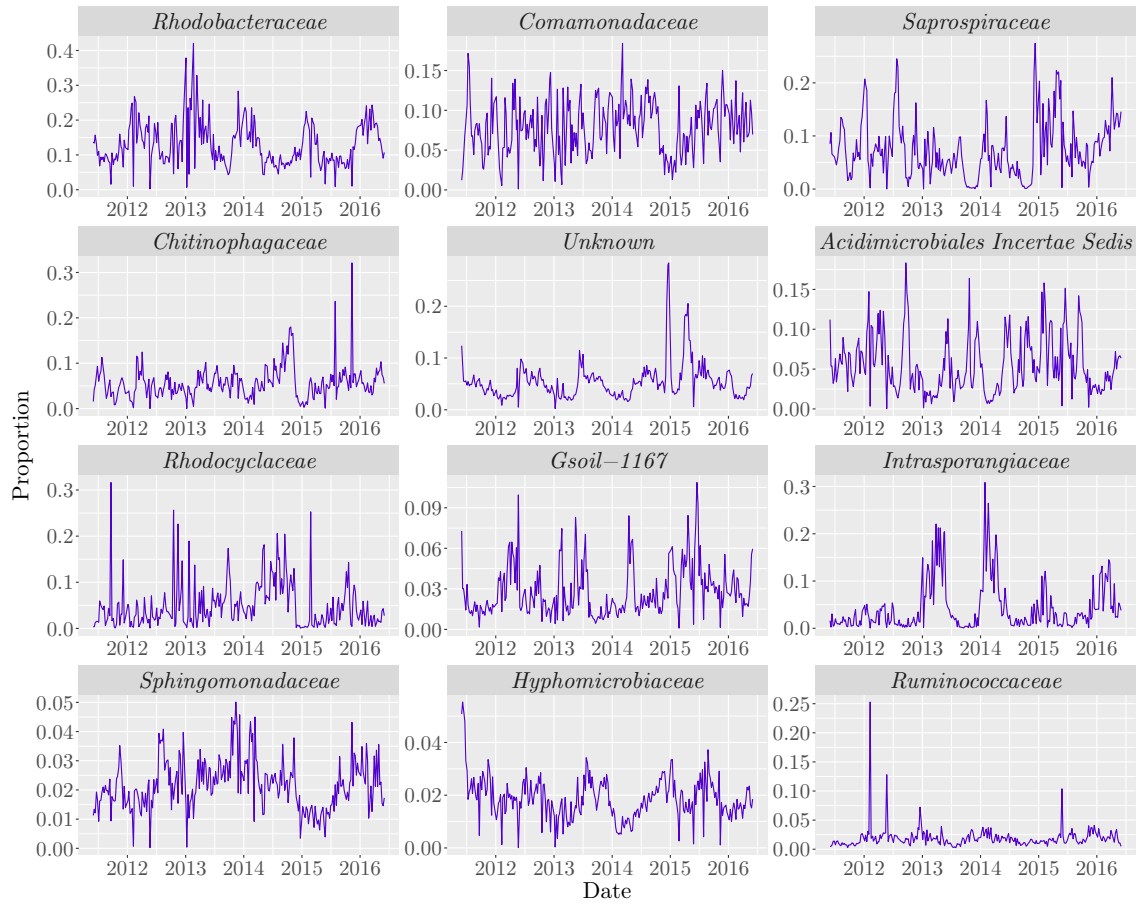
Covariate	Tank
COD	AS
Fluoride	AS
Phosphate	AS
Chloride	AS
Cadmium	SS
Lead	SS
Manganese	SS

Table D.1: Chemical and environmental covariates with a measurement regarded as an outlier and removed from the data, with the exception of COD, which had three outliers.

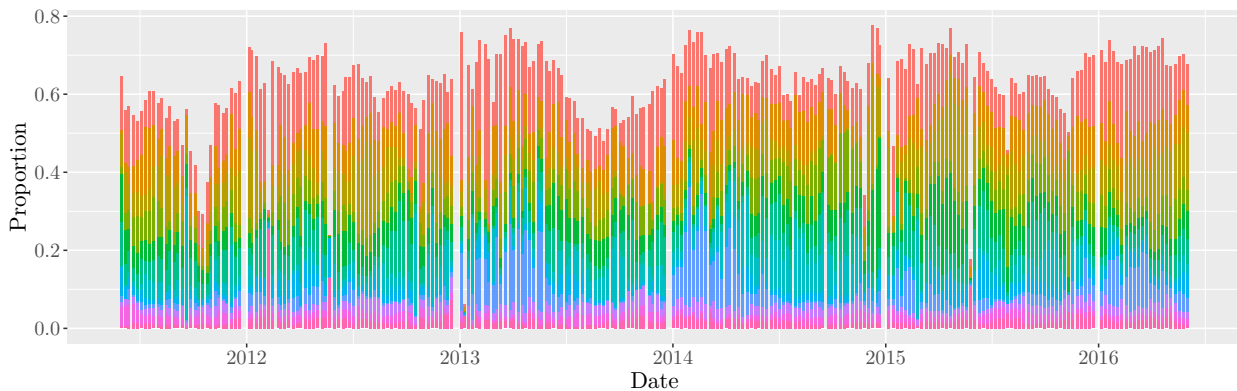
Appendix E

Additional plots

E.1 Exploratory plots for taxonomic ranks

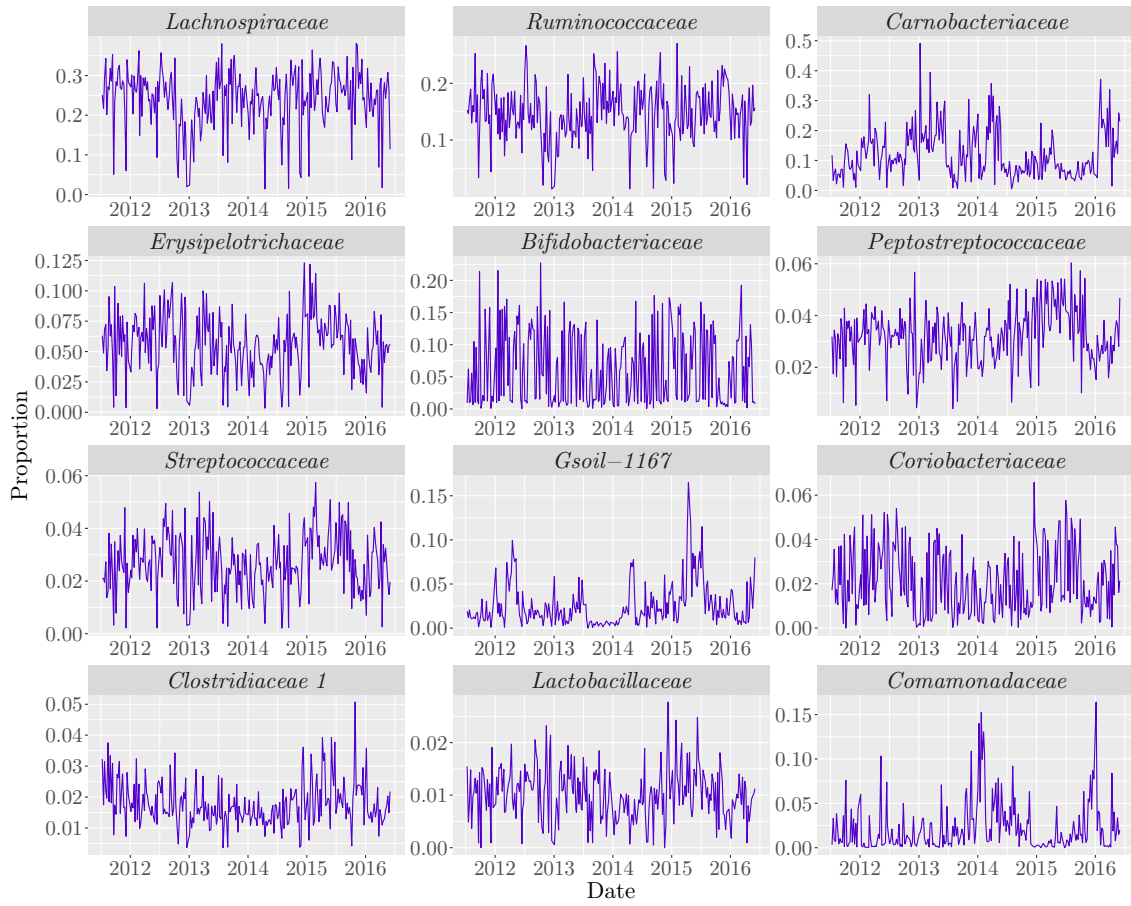


(a)

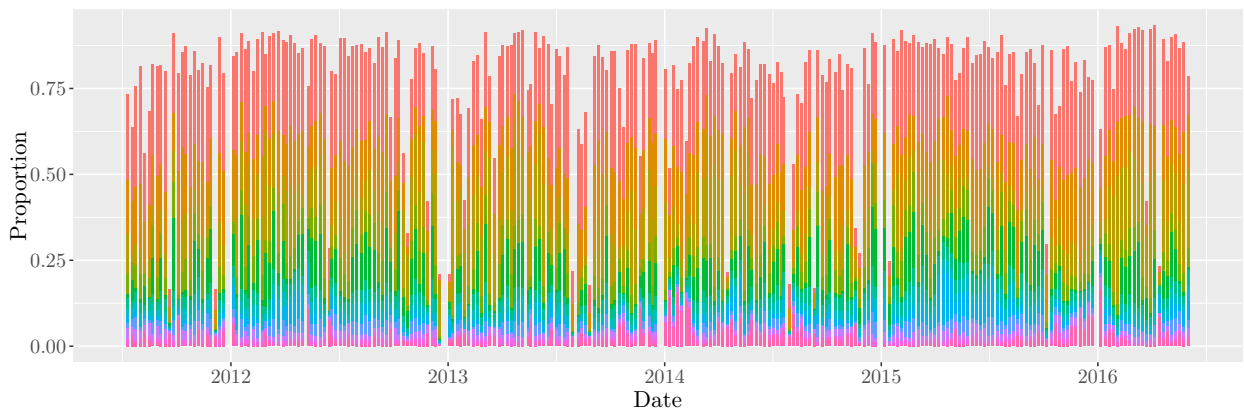
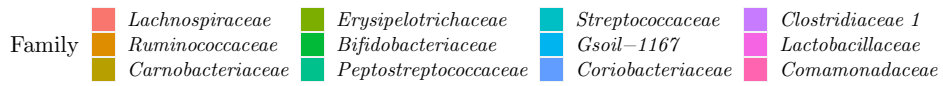


(b)

Figure E.1: (a) time series plots and (b) stacked bar plot for the top 12 families in the AS based on median abundance.

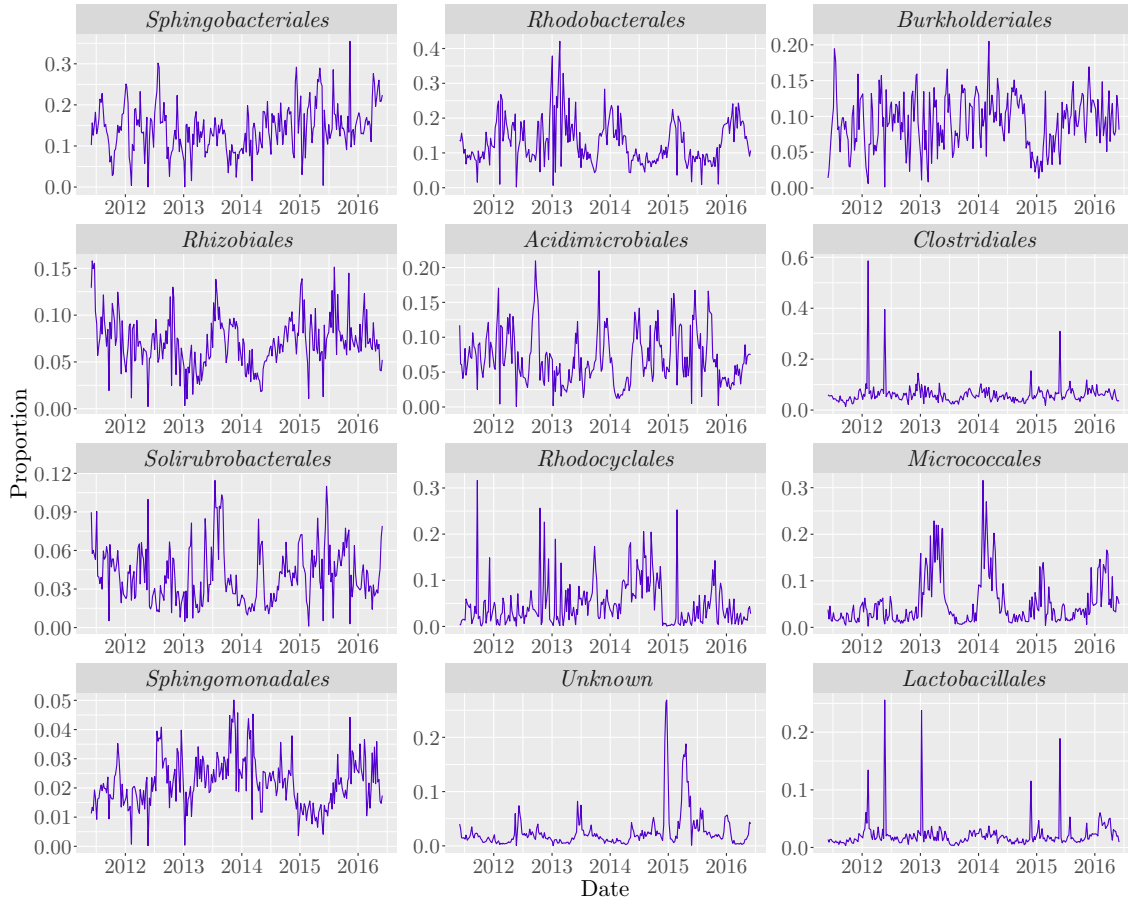


(a)

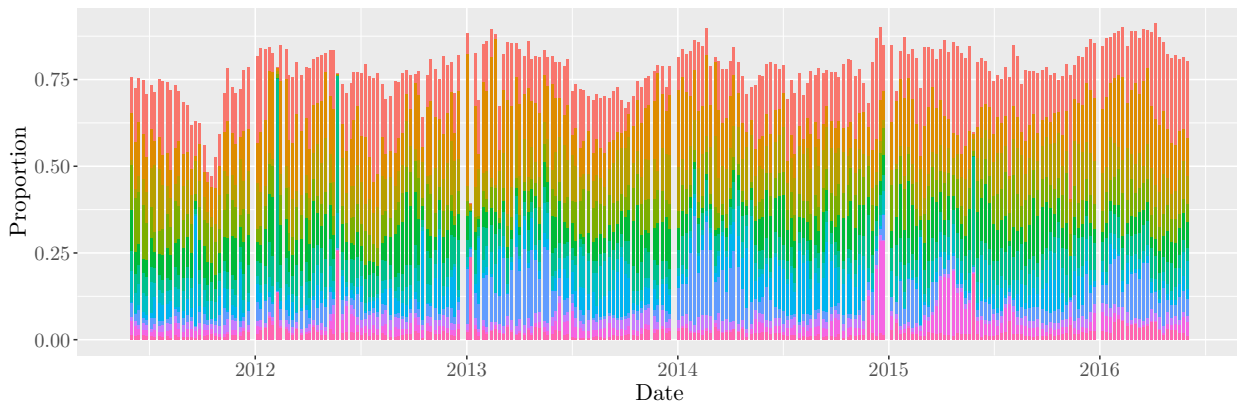


(b)

Figure E.2: (a) time series plots and (b) stacked bar plot for the top 12 families in the SS based on median abundance.

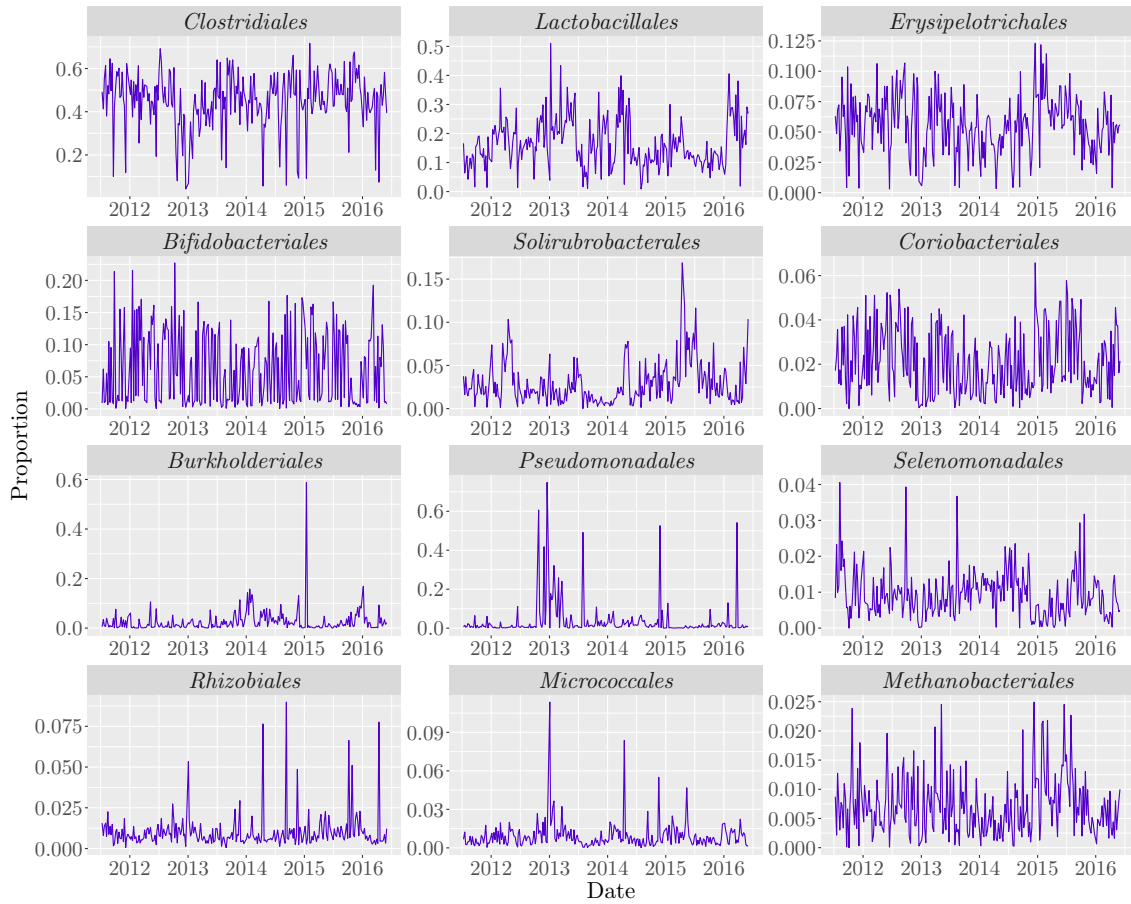


(a)

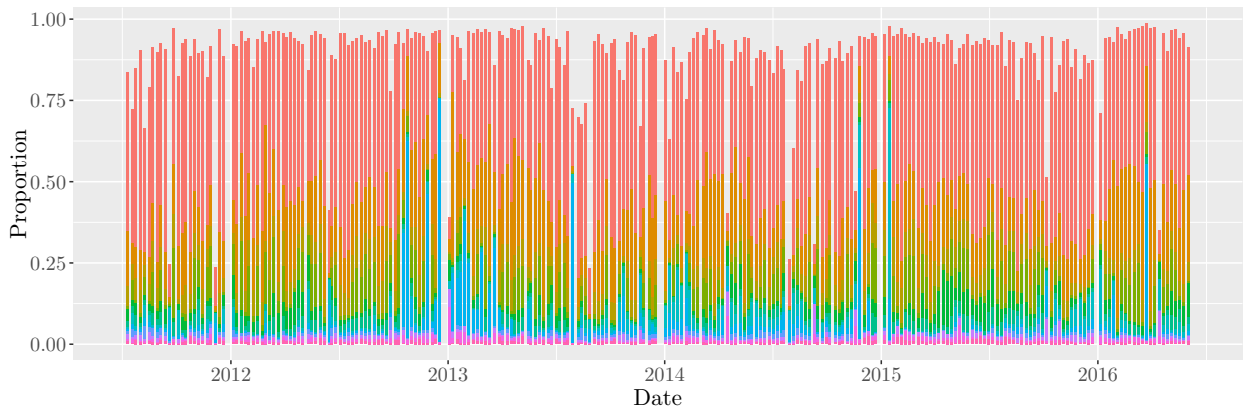


(b)

Figure E.3: (a) time series plots and (b) stacked bar plot for the top 12 orders in the AS based on median abundance.

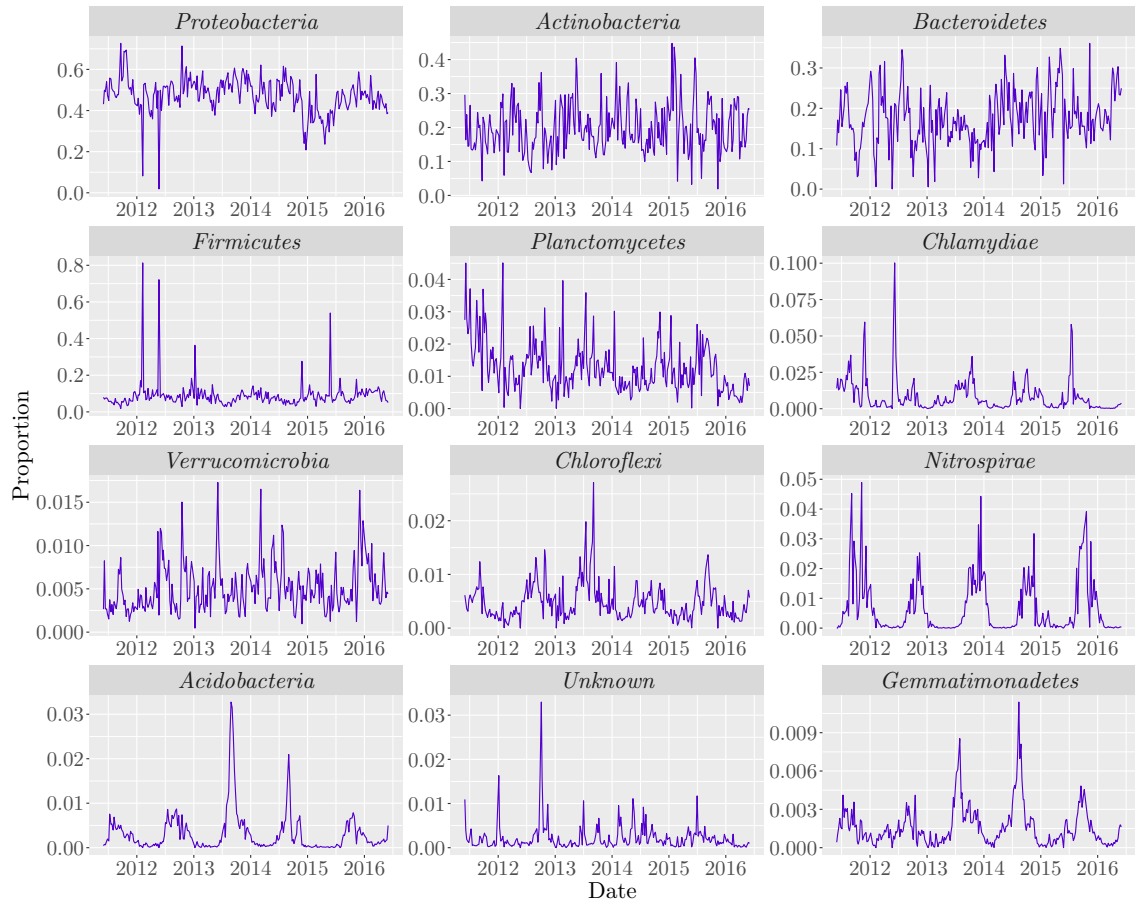


(a)

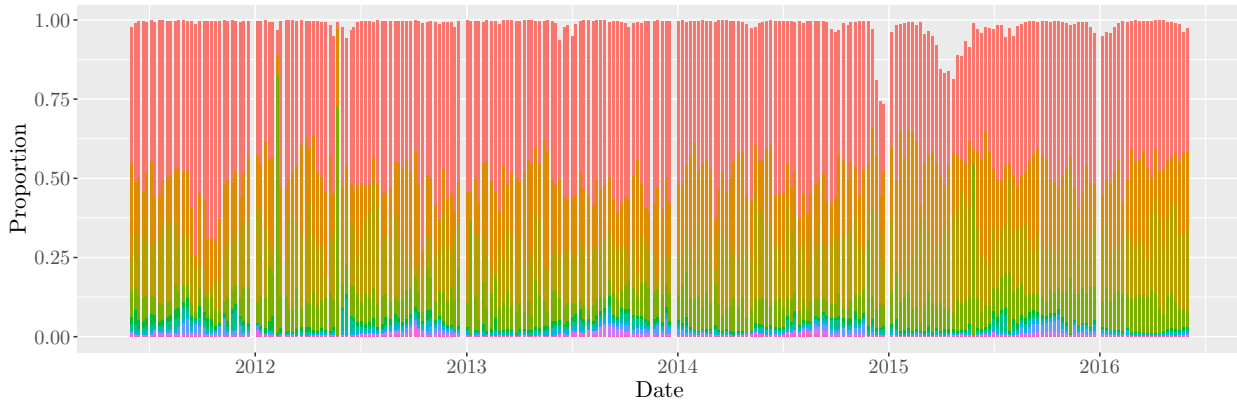


(b)

Figure E.4: (a) time series plots and (b) stacked bar plot for the top 12 orders in the SS based on median abundance.

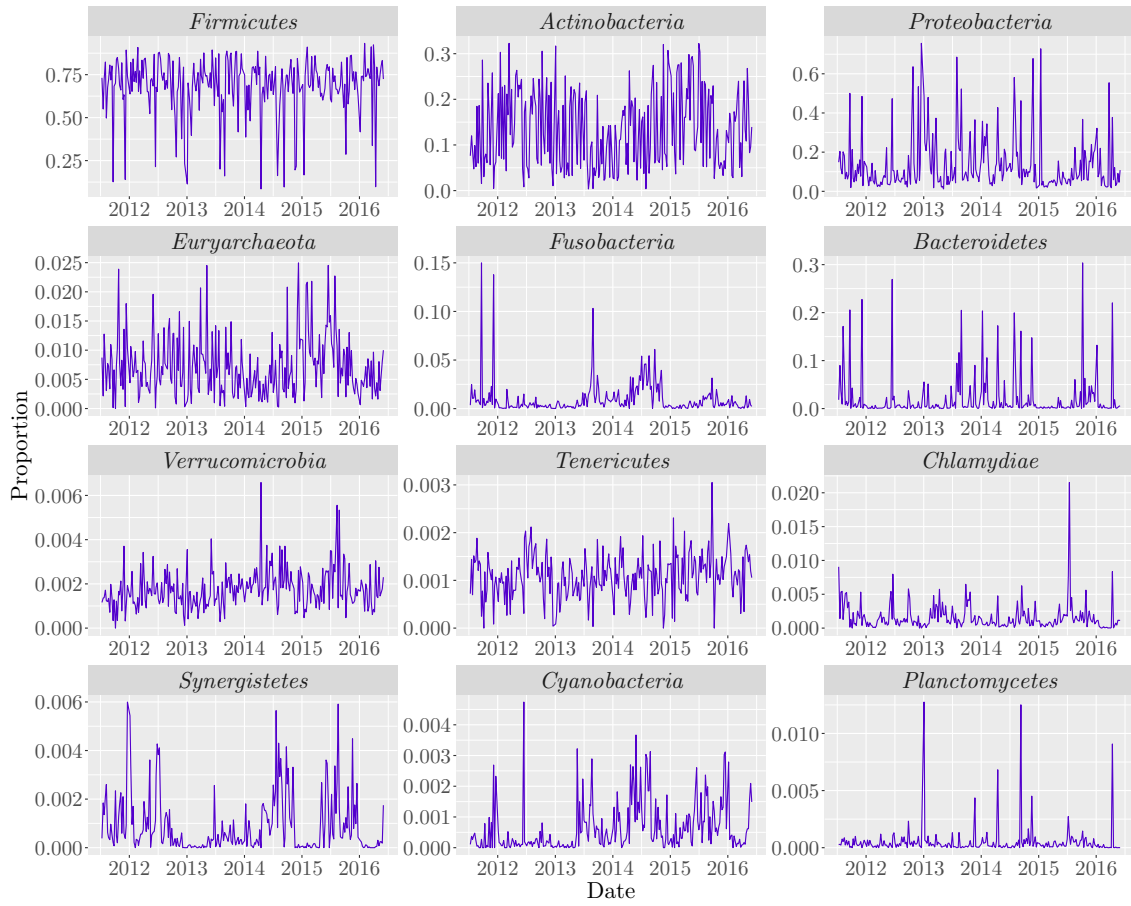


(a)

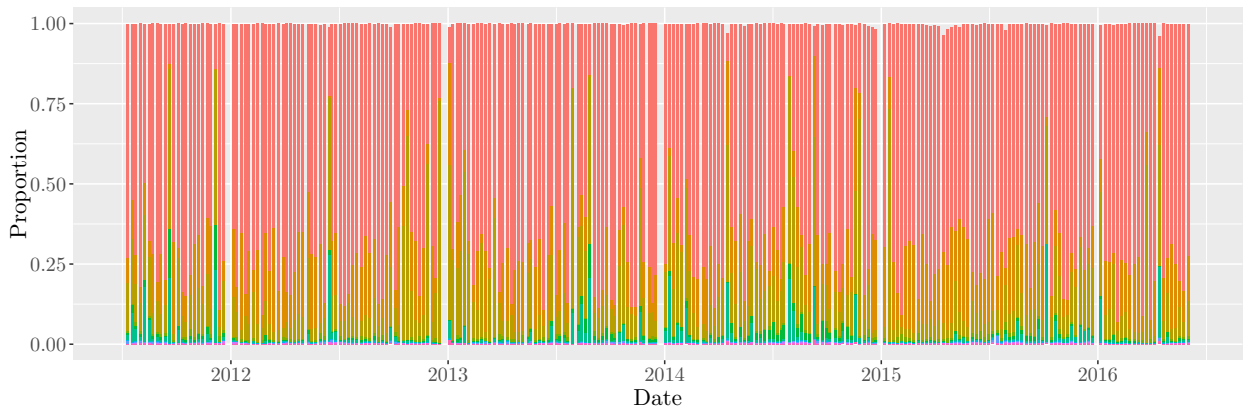


(b)

Figure E.5: (a) time series plots and (b) stacked bar plot for the top 12 phyla in the AS based on median abundance.



(a)



(b)

Figure E.6: (a) time series plots and (b) stacked bar plot for the top 12 phyla in the SS based on median abundance.

E.2 Heatmaps of matrices used in metagenomics simulation study

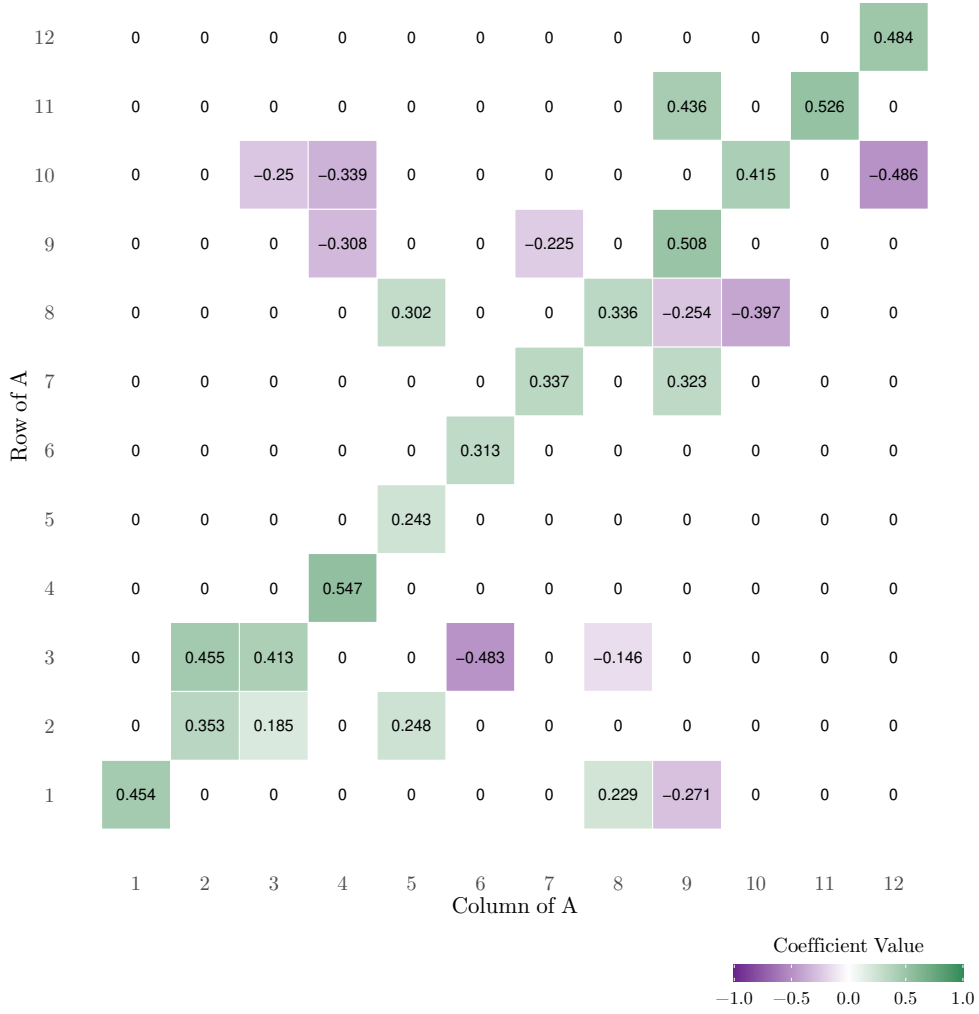


Figure E.7: Heatmap of A_{80} .

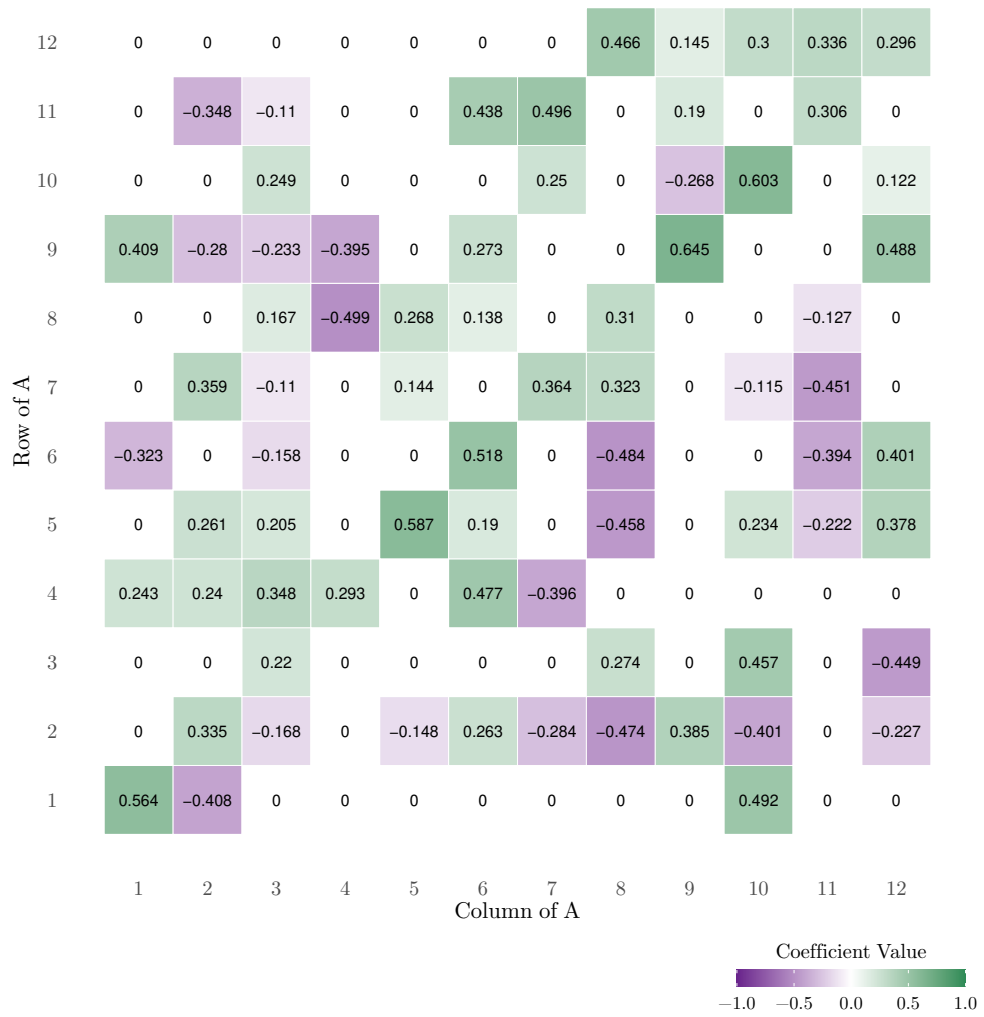


Figure E.8: Heatmap of A_{50} .

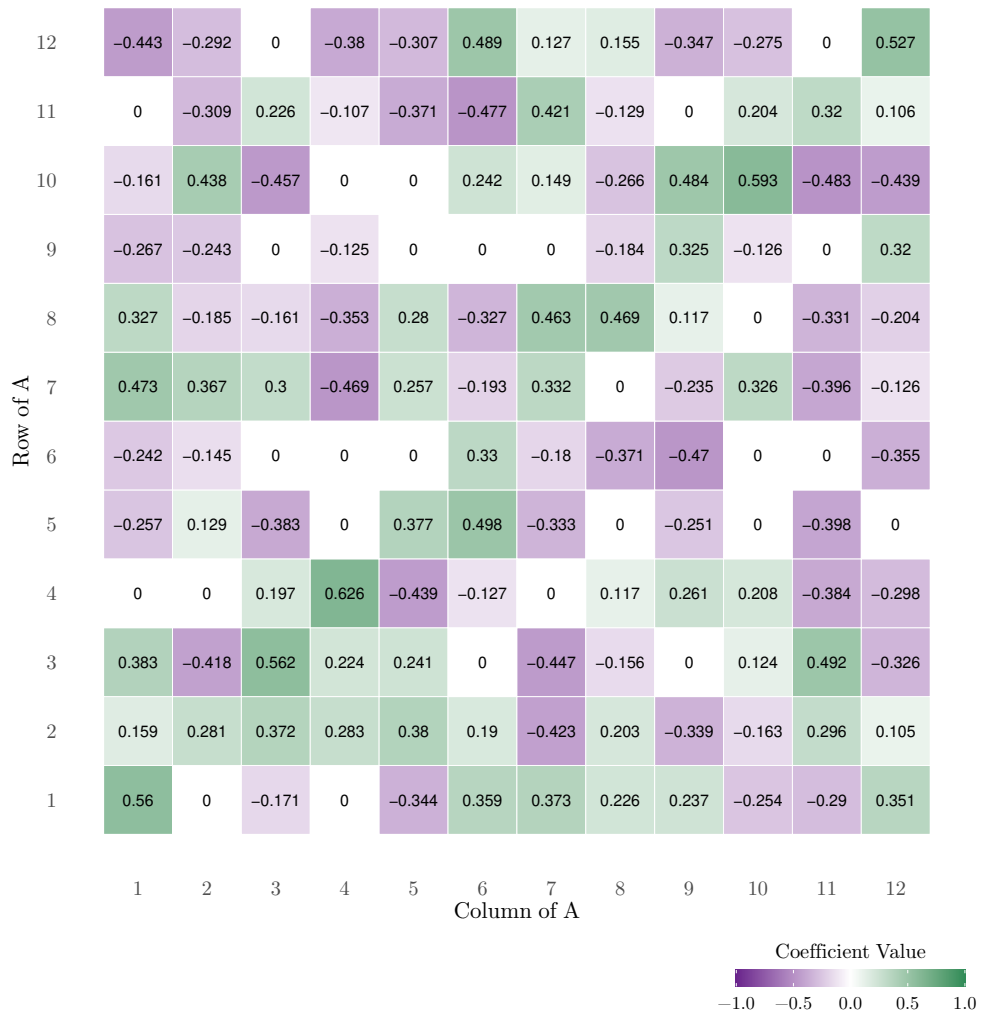


Figure E.9: Heatmap of A_{20} .

Bibliography

- ALAWI, M., OFF, S., KAYA, M. & SPIECK, E. 2009 Temperature influences the population structure of nitrite-oxidizing bacteria in activated sludge. *Environmental Microbiology Reports* **1** (3), 184–190.
- ALLEN, B. L. & STEEL, M. 2001 Subtree transfer operations and their induced metrics on evolutionary trees. *Annals of Combinatorics* **5**, 1–15.
- ANDERS, S. & HUBER, W. 2010 Differential expression analysis for sequence count data. *Genome Biology* **11**, 106.
- BARNARD, J., MCCULLOCH, R. & MENG, X.-L. 2000 Modelling covariance matrices in terms of standard deviations and correlations, with application to shrinkage. *Statistica Sinica* **10**, 1281–1311.
- BARRY, D. & HARTIGAN, J. A. 1987 Statistical analysis of hominoid molecular evolution. *Statistical Science* **2**, 191–210.
- BETANCOURT, M. 2017 A Conceptual Introduction to Hamiltonian Monte Carlo. *arXiv e-prints* p. arXiv:1701.02434.
- BLANQUART, S. & LARTILLOT, N. 2006 A Bayesian compound stochastic process for modeling nonstationary and nonhomogeneous sequence evolution. *Molecular Biology and Evolution* **23** (11), 2058–2071.
- BOHLIN, J., ELDHOLM, V., PETERSSON, J. H. O., BRYNILDSRUD, O. & SNIPEN, L. 2017 The nucleotide composition of microbial genomes indicates differential patterns of selection on core and accessory genomes. *BMC Genomics* **18** (1), 151.
- BRAUN, E. L. & KIMBALL, R. T. 2002 Examining basal avian divergences with mitochondrial sequences: Model complexity, taxon sampling, and sequence length. *Systematic Biology* **51** (4), 614–625.
- BRYANT, D. 2003 A classification of consensus methods for phylogenies. In *Bioconsensus, DIMACS Series* (ed. M. Janowitz, F. J. Lapointe, F. R. McMorris, B. Mirkin & F. S. Roberts), pp. 163–184. Providence, Rhode Island: American Mathematical Society.

- BUCCI, V., TZEN, B., LI, N., SIMMONS, M., TANOUE, T., BOGART, E., ... & GERBER, G. 2016 MDSINE: Microbial dynamical systems inference engine for microbiome time-series analyses. *Genome Biology* **17** (1), 121.
- BUNGE, J., WILLIS, A. & WALSH, F. 2014 Estimating the number of species in microbial diversity studies. *Annual Review of Statistics and Its Application* **1** (1), 427–445.
- BYRD, A. L., BELKAID, Y. & SEGRE, J. A. 2018 The human skin microbiome. *Nature Reviews Microbiology* **16**, 143–155.
- CARVALHO, C. M., POLSON, N. G. & SCOTT, J. G. 2009 Handling sparsity via the horseshoe. *Journal of Machine Learning Research* **5**, 73–80.
- CARVALHO, C. M., POLSON, N. G. & SCOTT, J. G. 2010 The horseshoe estimator for sparse signals. *Biometrika* **97** (2), 465–480.
- CHERLIN, S. 2016 Rooting major cellular radiations using statistical phylogenetics. PhD thesis, Newcastle University.
- CHERLIN, S., HEAPS, S., NYE, T., BOYS, R., WILLIAMS, T. & EMBLEY, T. 2017 The effect of non-reversibility on inferring rooted phylogenies. *Molecular Biology and Evolution* **35** (4), 984–1002.
- CONGDON, P. D. 2019 *Bayesian hierarchical models: with applications using R*, 2nd edn. New York, USA: Taylor & Francis.
- COX, C. J., FOSTER, P. G., HIRT, R. P., HARRIS, S. & EMBLEY, T. M. 2008 The archaeobacterial origin of eukaryotes. *Proceedings of the National Academy of Sciences* **105** (51), 20356–20361.
- CURTIS, T. P., HEAD, I. M. & GRAHAM, D. W. 2003 Peer reviewed: Theoretical ecology for engineering biology. *Environ. Sci. Technol.* **37**, 64–70.
- CYDZIK-KWIATKOWSKA, A. & ZIELIŃSKA, M. 2016 Bacterial communities in full-scale wastewater treatment systems. *World Journal of Microbiology and Biotechnology* **32**, 66.
- DAIMS, H., LEBEDEVA, E. V., PJEVAC, P., HAN, P., HERBOLD, C., ALBERTSEN, M., ... & WAGNER, M. 2015 Complete nitrification by *Nitrospira* bacteria. *Nature* **528**, 504–509.
- DAM, P., FONSECA, L. L., KONSTANTINIDIS, K. T. & VOIT, E. O. 2016 Dynamic models of the complex microbial metapopulation of Lake Mendota. *Npj Systems Biology And Applications* **2**, 16007.

- DANIELS, M. J. & POURAHMADI, M. 2002 Bayesian analysis of covariance matrices and dynamic models for longitudinal data. *Biometrika* **89** (3), 553–566.
- DAVID, L. A., MATERNA, A. C., FRIEDMAN, J., CAMPOS-BAPTISTA, M. I., BLACKBURN, M. C., PERROTTA, A., ERDMAN, S. E. & ALM, E. J. 2014 Host lifestyle affects human microbiota on daily timescales. *Genome Biology* **15**, R89.
- DAVIS, C. D. 2016 The gut microbiome and its role in obesity. *Nutrition Today* **51**, 167–174.
- DUGAN, P. R. 1981 The genus Zoogloea. In *The Prokaryotes: A Handbook on Habitats, Isolation, and Identification of Bacteria* (ed. M. P. Starr, H. Stolp, H. G. Trüper, A. Balows & H. G. Schlegel), pp. 764–770. Berlin, Heidelberg: Springer Berlin Heidelberg.
- DUTHEIL, J. & BOUSSAU, B. 2008 Non-homogeneous models of sequence evolution in the Bio++ suite of libraries and programs. *BMC Evolutionary Biology* **28**, 255.
- EILER, A., HEINRICH, F. & BERTILSSON, S. 2012 Coherent dynamics and association networks among lake bacterioplankton taxa. *The ISME journal* **6**, 330–342.
- EMBLEY, T. M., THOMAS, R. H. & WILLIAMS, R. A. D. 1993 Reduced thermophilic bias in the 16S rDNA sequence from *Thermus ruber* provides further support for a relationship between *Thermus* and *Deinococcus*. *Systematic and Applied Microbiology* **16** (1), 25–29.
- FARROW, M. 2005 A short introduction to Bayesian analysis. *Tech. Rep.*. School of Mathematics, Statistic & Physics, Newcastle University, Newcastle upon Tyne, UK.
- FAUST, K., LAHTI, L., GONZE, D., DE VOS, W. M. & RAES, J. 2015 Metagenomics meets time series analysis: unraveling microbial community dynamics. *Current Opinion in Microbiology* **25**, 56–66.
- FELSENSTEIN, J. 1973 Maximum likelihood and minimum-steps methods for estimating evolutionary trees from data on discrete characters. *Systematic Zoology* **22**, 240–249.
- FELSENSTEIN, J. 1978 Cases in which parsimony or compatibility methods will be positively misleading. *Systematic Zoology* **27** (4), 401–410.
- FELSENSTEIN, J. 1981 Evolutionary trees from DNA sequences: A maximum likelihood approach. *Journal of Molecular Evolution* **17**, 368–376.
- FERNÁNDEZ-SÁNCHEZ, J., SUMNER, J. G., JARVIS, P. D. & WOODHAMS, M. D. 2015 Lie Markov models with purine/pyrimidine symmetry. *Journal of Mathematical Biology* **70** (4), 855–891.

- FISHER, C. K. & MEHTA, P. 2014 Identifying keystone species in the human gut microbiome from metagenomic timeseries using sparse linear regression. *PLOS ONE* **9** (7), 1–10.
- FOSTER, P. G. 2004 Modeling compositional heterogeneity. *Systematic Biology* **53** (3), 485.
- FRIEL, N. & PETTITT, A. N. 2008 Marginal likelihood estimation via power posteriors. *Journal of the Royal Statistical Society: Series B* **70**, 589–607.
- GELFAND, A. E., FUENTES, M., GUTTORP, P. & DIGGLE, P. 2010 *Handbook of Spatial Statistics*. Taylor & Francis.
- GELMAN, A., CARLIN, J. B., STERN, H. S., DUNSON, D. B., VEHTARI, A. & RUBIN, D. B. 2014 *Bayesian Data Analysis*, 3rd edn. Taylor & Francis.
- GEMAN, S. & GEMAN, D. 1984 Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **PAMI-6** (6), 721–741.
- GERMAIN, S. E. 2010 Bayesian spatiotemporal modelling of rainfall through nonhomogeneous hidden markov models. PhD thesis, Newcastle University.
- GEYER, C. J. 2011 Introduction to Markov chain Monte Carlo. In *Handbook of Markov Chain Monte Carlo*, pp. 3–48. Chapman & Hall London.
- GIBBONS, S. M., KEARNEY, S. M., SMILLIE, C. S. & ALM, E. J. 2017 Two dynamic regimes in the human gut microbiome. *PLOS Computational Biology* **13** (2), 1–20.
- GILBERT, J. A., O’DOR, R., KING, N. & VOGEL, T. M. 2011 The importance of metagenomic surveys to microbial ecology: or why Darwin would have been a metagenomic scientist. *Microbial Informatics and Experimentation* **1** (1), 5.
- GREEN, P. J. 2003 Trans-dimensional Markov chain Monte Carlo. In *Highly Structured Stochastic Systems* (ed. P. J. Green, N. L. Hjort & S. Richardson), *Oxford Statist. Sci. Ser.*, vol. 27, pp. 179–206. Oxford: Oxford Univ. Press.
- GRIFFIN, J. & BROWN, P. 2017 Hierarchical shrinkage priors for regression models. *Bayesian Analysis* **12** (1), 135–159.
- HANNAFORD, N. E., HEAPS, S. E., NYE, T. M. W., WILLIAMS, T. A. & EMBLEY, T. M. 2020 Incorporating compositional heterogeneity into Lie Markov models for phylogenetic inference. *Ann. Appl. Stat.* **14** (4), 1964–1983.

- HASEGAWA, M., KISHINO, H. & YANO, T. 1985 Dating of the human-ape splitting by a molecular clock of mitochondrial DNA. *Journal of Molecular Evolution* **22** (2), 160–174.
- HASTINGS, W. K. 1970 Monte Carlo sampling methods using Markov chains and their applications. *Biometrika* **57** (1), 97–109.
- HEAPS, S. E., NYE, T. M. W., BOYS, R. J., WILLIAMS, T. A. & EMBLEY, T. M. 2014 Bayesian modelling of compositional heterogeneity in molecular phylogenetics. *Statistical Applications in Genetics and Molecular Biology* **1**, 1–21.
- HODGKINSON, A. & EYRE-WALKER, A. 2011 Variation in the mutation rate across mammalian genomes. *Nature Reviews Genetics* **12** (3), 756–766.
- HOFFMAN, M. D. & GELMAN, A. 2014 The No-U-Turn Sampler: Adaptively setting path lengths in Hamiltonian Monte Carlo. *Journal of Machine Learning Research* **15**, 1593–1623.
- HUELSENBECK, J. P., BOLLBACK, J. P. & LEVINE, A. M. 2002 Inferring the root of a phylogenetic tree. *Systematic Biology* **51** (1), 32–43.
- IMHOFF, J. F. 2015 Rhodobaca. In *Bergey's Manual of Systematics of Archaea and Bacteria*, pp. 1–4. American Cancer Society.
- JARVIS, E. D., MIRARAB, S., ABERER, A. J., LI, B., HOUDE, P. & LI, CAI ... ZHANG, G. 2014 Whole-genome analyses resolve early branches in the tree of life of modern birds. *Science* **346** (6215), 1320–1331.
- JONSSON, V., ÖSTERLUND, T., NERMAN, O. & KRISTIANSOON, E. 2016 Statistical evaluation of methods for identification of differentially abundant genes in comparative metagenomics. *BMC Genomics* **17**, 78.
- JOSHI, G. K., JUGRAN, J. & BHATT, J. P. 2014 Metagenomics: The exploration of unculturable microbial world. In *Advances in Biotechnology*, pp. 105–115. New Delhi: Springer India.
- KAEHLER, B. D. 2017 Full reconstruction of non-stationary strand-symmetric models on rooted phylogenies. *Journal of Theoretical Biology* **420**, 144–151.
- KASS, R. E., CARLIN, B. P., GELMAN, A. & NEAL, R. M. 1998 Markov chain Monte Carlo in practice: A roundtable discussion. *The American Statistician* **52** (2), 93–100.
- KAUL, A., MANDAL, S. & DAVIDOV, O. 2017 Analysis of microbiome data in the presence of excess zeros. *Frontiers in Microbiology* **8**, 2114.

- VAN KESSEL, M. A. H. J., SPETH, D. R., ALBERTSEN, M., NIELSEN, P. H., OP DEN CAMP, H. J. M., KARTAL, B., JETTEN, M. S. M. & LÜCKER, S. 2015 Complete nitrification by a single microorganism. *Nature* **528**, 555–559.
- KOCH, H., LÜCKER, S., ALBERTSEN, M., KITZINGER, K., HERBOLD, C., SPIECK, E., ... & DAIMS, H. 2015 Expanded metabolic versatility of ubiquitous nitrite-oxidizing bacteria from the genus *Nitrospira*. *Proceedings of the National Academy of Sciences* **112** (36), 11371–11376.
- KUO, L. & MALICK, B. 1998 Variable selection for regression models. *Sankhy: The Indian Journal of Statistics, Series B (1960-2002)* **60** (1), 65–81.
- LAFONTAINE, D. L. & TOLLERVEY, D. 2001 The function and synthesis of ribosomes. *Nature Reviews Molecular Cell Biology* **2**, 514–520.
- LARTILLOT, N., BLANQUART, S. & LEPAGE, T. 2004 *PhyloBayes 3.3: A Bayesian software for phylogenetic reconstruction and molecular dating using mixture models*.
- LARTILLOT, N. & PHILIPPE, H. 2006 Computing Bayes factors using thermodynamic integration. *Systematic Biology* **55**, 195–207.
- LEE, K. H., COULL, B., MOSCICKI, A.-B., PASTER, B. & STARR, J. 2018 Bayesian variable selection for multivariate zero-inflated models: Application to microbiome count data. *Biostatistics* .
- LEI, G., BOYS, R. J., GILLESPIE, C. S., GREENALL, A. & WILKINSON, D. J. 2011 Bayesian inference for sparse VAR (1) models, with application to time course microarray data. *Journal of Biometrics & Biostatistics* **2** (5).
- LI, H., SHAO, R., SONG, N., SONG, F., JIANG, P., LI, Z. & CAI, W. 2015 Higher-level phylogeny of paraneopteran insects inferred from mitochondrial genome sequences. *Scientific Reports* **5**, 8527.
- LIND, P. A. & ANDERSSON, D. I. 2008 Whole-genome mutational biases in bacteria. *Proceedings of the National Academy of Sciences of the United States of America* **105** (46), 17878–17883.
- LITTLE, R. J. & RUBIN, D. B. 2019 *Statistical Analysis with Missing Data (3rd ed.)*. Hoboken, NJ, USA: Wiley.
- LOTKA, A. J. 1926 Elements of physical biology. *Science Progress in the Twentieth Century (1919-1933)* **21** (82), 341–343.
- MAPLESOFT 2020 Maple.

- MCILROY, S. J., SAUNDERS, A. M., ALBERTSEN, M., NIERYCHLO, M., MCILROY, B., HANSEN, A. A., ... & NIELSEN, P. H. 2015 MiDAS: The field guide to the microbes of activated sludge. *Database* **2015**, bav062.
- MCMURDIE, P. J. & HOLMES, S. 2014 Waste not, want not: Why rarefying microbiome data is inadmissible. *PLOS Computational Biology* **10** (4), 1–12.
- MEHRANI, M., SOBOTKA, D., KOWAL, P., CIESIELSKI, S. & MAKINIA, J. 2020 The occurrence and role of Nitrospira in nitrogen removal systems. *Bioresource Technology* **303**, 122936.
- MITCHELL, T. J. & BEAUCHAMP, J. J. 1988 Bayesian variable selection in linear regression. *Journal of the American Statistical Association* **83** (404), 1023–1032.
- MORGAN, C. C., FOSTER, P. G., WEBB, A. E., PISANI, D. & O’CONNELL, J. O. M. M. J. 2013 Heterogeneous models place the root of the placental mammal phylogeny. *Molecular Biology and Evolution* **30** (9), 2145–2156.
- MOUNIER, J., MONNET, C., VALLAEYS, T., ARDITI, R., SARTHOU, A., HÉLIAS, A. & IRLINGER, F. 2008 Microbial interactions within a cheese microbial community. *Applied and Environmental Microbiology* **74** (1), 172–181.
- NEWTON, M. A. & RAFTERY, A. E. 1994 Approximate Bayesian inference by the weighted likelihood bootstrap (with discussion). *Journal of the Royal Statistical Society: Series B* **56**, 3–48.
- NYE, T. M. 2020 Random walks and Brownian motion on cubical complexes. *Stochastic Processes and their Applications* **130** (4), 2185 – 2199.
- OAKS, J. R., COBB, K. A., MININ, V. N. & LEACHÉ, A. D. 2019 Marginal likelihoods in phylogenetics: A review of methods and applications. *Systematic Biology* .
- OCHOA-HERRERA, V., BANIHANI, Q., LEN, G., KHATRI, C., FIELD, J. A. & SIERRA-ALVAREZ, R. 2009 Toxicity of fluoride to microorganisms in biological wastewater treatment systems. *Water Research* **43** (13), 3177–3186.
- O’HARA, R. & SILLANPÄÄ, M. 2009 A review of Bayesian variable selection methods: What, how and which. *Bayesian Analysis* **4** (1), 85–118.
- PAULSON, J. N., STINE, O. C., BRAVO, H. C. & POP, M. 2013 Differential abundance analysis for microbial marker-gene surveys. *Nature Methods* **10**, 1200–1203.
- PHILLIPS, M. J., GIBB, G. C., CRIMP, E. A. & PENNY, D. 2009 Tinamous and Moa flock together: Mitochondrial genome sequence analysis reveals independent losses of flight among ratites. *Systematic Biology* **59** (1), 90–107.

- PHILLIPS, M. J., MCLENACHAN, P. A., DOWN, C., GIBB, G. C. & PENNY, D. 2006 Combined mitochondrial and nuclear DNA sequences resolve the interrelations of the major Australasian marsupial radiations. *Systematic Biology* **55** (1), 122–137.
- PIIRONEN, J. & VEHTARI, A. 2017 Sparsity information and regularization in the horse-shoe and other shrinkages priors. *Electronic Journal of Statistics* **11**, 5018–5051.
- PUJALTE, M. J., LUCENA, T., RUVIRA, M. A., ARAHAL, D. R. & MACIÁN, M. C. 2014 The family rhodobacteraceae. In *The Prokaryotes: Alphaproteobacteria and Betaproteobacteria*, pp. 439–512. Springer Berlin Heidelberg.
- QUINCE, C., WALKER, A. W., SIMPSON, J. T., LOMAN, N. J. & SEGATA, N. 2017 Shotgun metagenomics, from sampling to analysis. *Nature Biotechnology* **35**, 833–844.
- RAMSAY, J. O., HOOKER, G., CAMPBELL, D. & CAO, J. 2007 Parameter estimation for differential equations: a generalized smoothing approach. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **69** (5), 741–796.
- RIBOSOMAL DATABASE PROJECT 2016 RDP Release 11 – Sequence Analysis Tools. <https://rdp.cme.msu.edu/>, accessed: 2020-04-29.
- ROMIGUIER, J. & ROUX, C. 2017 Analytical biases associated with GC-content in molecular evolution. *Frontiers in Genetics* **8**, 16.
- ROSENBERG, M. S., SUBRAMANIAN, S. & KUMAR, S. 2003 Patterns of transitional mutation biases within and among mammalian genomes. *Molecular Biology and Evolution* **20** (6), 988–993.
- RUBIN, D. B. 1976 Inference and missing data. *Biometrika* **63** (3), 581–592.
- SCHWARZ, G. 1978 Estimating the Dimension of a Model. *The Annals of Statistics* **6** (2), 461 – 464.
- SHCHEGOLKOVA, N. M., KRASNOV, G. S., BELOVA, A. A., DMITRIEV, A. A., KHARITONOV, S. L., KLIMINA, K. M., MELNIKOVA, N. V. & KUDRYAVTSEVA, A. V. 2016 Microbial community structure of activated sludge in treatment plants with different wastewater compositions. *Frontiers in Microbiology* **7**, 90.
- SHUMWAY, R. H. & STOFFER, D. S. 2017 *Time Series Analysis and Its Applications with R Examples*, 4th edn. Springer.
- SPIEGELHALTER, D. J., BEST, N. G., CARLIN, B. P. & VAN DER LINDE, A. 2002 Bayesian measures of model complexity and fit. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **64** (4), 583–639.

- SQUARTINI, F. & ARNDT, P. F. 2008 Quantifying the stationarity and time reversibility of the nucleotide substitution process. *Molecular Biology and Evolution* **25**, 2525–2535.
- STAN DEVELOPMENT TEAM 2020 RStan: the R interface to Stan. R package version 2.21.2.
- STAN DEVELOPMENT TEAM 2021 Stan modeling language users guide and reference manual, version 2.26.
- STEEL, M. & MCKENZIE, A. 2001 Properties of phylogenetic trees generated by Yule-type speciation models. *Mathematical Biosciences* **170**, 91–112.
- STEIN, R. R., BUCCI, V., TOUSSAINT, N. C., BUFFIE, C. G., RÄTSCH, G., PAMER, E. G., SANDER, C. & XAVIER, J. B. 2013 Ecological modeling from time-series inference: Insight into dynamics and stability of intestinal microbiota. *PLOS Computational Biology* **9** (12), 1–11.
- SUMNER, J. G., FERNÁNDEZ-SÁNCHEZ, J. & JARVIS, P. D. 2012a Lie Markov models. *Journal of Theoretical Biology* **298**, 16–31.
- SUMNER, J. G., JARVIS, P. D., FERNÁNDEZ-SÁNCHEZ, J., KAINE, B. T., WOODHAMS, M. D. & HOLLAND, B. R. 2012b Is the general time-reversible model bad for molecular phylogenetics? *Systematic Biology* **61** (6), 1069–1074.
- TARRÍO, R., RODRÍGUEZ-TRELLES, F. & AYALAA, F. J. 2000 Tree rooting with outgroups when they differ in their nucleotide composition from the ingroup: The *Drosophila saltans* and *willistoni* groups, a case study. *Molecular Phylogenetics and Evolution* **16** (3), 344–349.
- TAVARÉ, S. 1986 Some probabilistic and statistical problems in the analysis of DNA sequences. *Lectures on Mathematics in the Life Sciences* **17**, 57–86.
- VEHTARI, A., GELMAN, A., SIMPSON, D., CARPENTER, B. & BRKNER, P.-C. 2020 Rank-normalization, folding, and localization: An improved \hat{R} for assessing convergence of MCMC. *Bayesian Analysis* .
- VOLTERRA, V. 1926 Fluctuations in the abundance of a species considered mathematically. *Nature* **118**, 558–560.
- WAGNER, M., RATH, G., KOOPS, H.-P., FLOOD, J. & AMANN, R. 1996 In situ analysis of nitrifying bacteria in sewage treatment plants. *Water Science and Technology* **34** (1-2), 237–244.

- WANG, Q., GARRITY, G. M., TIEDJE, J. M. & COLE, J. R. 2007 Naïve Bayesian classifier for rapid assignment of rRNA sequences into the new bacterial taxonomy. *Applied and Environmental Microbiology* **73**, 5261–5267.
- WEISS, S., XU, Z. Z., PEDDADA, S., AMIR, A., BITTINGER, K., GONZALEZ, A., ... & KNIGHT, R. 2017 Normalization and microbial differential abundance strategies depend upon data characteristics. *Microbiome* **5**, 27.
- WETZEL, R. G. 2001 12 - The nitrogen cycle. In *Limnology*, 3rd edn. (ed. R. G. Wetzel), pp. 205–237. San Diego: Academic Press.
- WILKINSON, D. J. 2011 Chemical and biochemical kinetics. In *Stochastic Modelling for Systems Biology*, 2nd edn., pp. 171–202. CRC Press LLC.
- WILLIAMS, K. P., SOBRAL, B. W. & DICKERMAN, A. W. 2007 A robust species tree for the Alphaproteobacteria. *Journal of Bacteriology* pp. 4578–4586.
- WILLIAMS, T. A., HEAPS, S. E., CHERLIN, S., NYE, T. M. W., BOYS, R. J. & EMBLEY, T. M. 2015 New substitution models for rooting phylogenetic trees. *Philosophical Transactions of the Royal Society B: Biological Sciences* **370** (1678).
- WOESE, C. R. & FOX, G. E. 1977 Phylogenetic structure of the prokaryotic domain: the primary kingdoms. *Proceedings of the National Academy of Sciences of the United States of America* **74**, 5088–5090.
- WOODHAMS, M. D., FERNÁNDEZ-SÁNCHEZ, J. & SUMNER, J. G. 2015 A new hierarchy of phylogenetic models consistent with heterogeneous substitution rates. *Systematic Biology* **64**, 638–650.
- XIA, Y., SUN, J. & CHEN, D.-G. 2018 *Statistical Analysis of Microbiome Data with R*. Singapore: Springer Singapore.
- XIE, W., LEWIS, P. O., FAN, Y., KUO, L. & CHEN, M.-H. 2011 Improving marginal likelihood estimation for Bayesian phylogenetic model selection. *Systematic Biology* **60**, 150–160.
- YANG, Z. 1993 Maximum-likelihood estimation of phylogeny from dna sequences when substitution rates differ over sites. *Molecular Biology and Evolution* **10**, 1396–1401.
- YANG, Z. 1994 Maximum likelihood phylogenetic estimation from DNA sequences with variable rates over sites: Approximate methods. *Journal of Molecular Evolution* **39**, 306–314.

- YANG, Z. 1996 Among-site rate variation and its impact on phylogenetic analyses. *Trends in Ecology & Evolution* **11** (9), 367–372.
- YANG, Z. 2006 *Computational Molecular Evolution*. Oxford, UK: Oxford University Press.
- YANG, Z. & ROBERTS, D. 1995 On the use of nucleic acid sequences to infer early branchings in the tree of life. *Molecular Biology and Evolution* **12** (3), 451–458.
- YIN, J., MAO, X. & WU, F. 2009 Generalized stochastic delay Lotka-Volterra systems. *Stochastic Models* **25** (3), 436–454.
- ZWICKL, D. J. & HOLDER, M. T. 2004 Model parameterization, prior distributions, and the general time-reversible model in Bayesian phylogenetics. *Systematic Biology* **53** (6), 877–888.