

¹ Inference of multiple mergers while dating a pathogen phylogeny

² David Helekal¹, Jere Koskela², Xavier Didelot^{3,*}

³ ¹ Centre for Doctoral Training in Mathematics for Real-World Systems, University of Warwick, UK

⁴ ² School of Mathematics, Statistics and Physics, Newcastle University, UK

⁵ ³ School of Life Sciences and Department of Statistics, University of Warwick, UK

⁶ * Corresponding author. Tel: 0044 (0)2476 572827. Email: xavier.didelot@warwick.ac.uk

⁷ Running title: Inference of multiple mergers in phylogenies

⁸ Keywords: pathogen phylogenetics; coalescent model; multiple mergers; phylogenetic dating

9 ABSTRACT

10 The vast majority of pathogen phylogenetic studies do not consider the possibility of multiple merger
11 events being present, where a single node of the tree leads to more than two descendent branches.
12 These events are however likely to occur when studying a relatively small population or if there is
13 high variability in the reproductive chances. Here we consider the problem of detecting the presence of
14 multiple mergers in the context of dating a phylogeny, that is determining the date of each of the nodes.
15 We use the Lambda-coalescent theory as a modelling framework and show how Bayesian inference can
16 be efficiently performed using a Billera-Holmes-Vogtmann space embedding and a customised Markov
17 Chain Monte Carlo sampling scheme. We applied this new analysis methodology to a large number
18 of simulated datasets to show that it is possible to infer if and when multiple merger events occurred,
19 and that the phylogenetic dating is improved as a result of taking this information into account. We
20 also analysed real datasets of *Vibrio cholerae* and *Mycobacterium tuberculosis* to demonstrate the
21 relevance of our approach to real pathogen evolutionary epidemiology. We have implemented our new
22 methodology in a R package which is freely available at <https://github.com/dhelekal/MMCTime>.

23 INTRODUCTION

24 Dated phylogenies have risen to prominence in many research areas of the life sciences, from the study
25 of evolutionary histories of higher organisms, genomic epidemiology of infectious disease, through
26 to understanding diversity of microbial organisms. Most existing approaches to reconstructing and
27 analysing dated phylogenies are restricted to binary trees, where each internal node has exactly two
28 descendent branches. Indeed, Kingman's coalescent (Kingman 1982) on a continuous real time scale
29 (Drummond et al. 2002) is the most popular framework for modelling dated phylogenies of measurably
30 evolving populations (Drummond et al. 2003; Biek et al. 2015), and in this model only two lineages
31 may merge into the same ancestor at once. However, this model is only applicable if both the sample
32 size and typical family sizes are small in comparison to the effective population size, which can be
33 orders of magnitude smaller than the census population size for example due to heterogeneity of the
34 reproduction success (Charlesworth 2009).

35 In contrast with the standard Kingman's coalescent, Lambda-coalescent models, also known as
36 multiple merger coalescents, can be used to describe dated phylogenies where more than two
37 lineages may coalesce into the same ancestor at once (Pitman 1999; Sagitov 1999; Donnelly and
38 Kurtz 1999). Multiple merger events can be the result of various biological phenomena of interest,
39 such as populations undergoing rapid adaptation (Neher and Hallatschek 2013; Desai et al. 2013),
40 superspreading (Hoscheit and Pybus 2019; Lemieux et al. 2020; Gómez-Carballa et al. 2021) or some
41 other form of sweepstakes reproduction (Menardo et al. 2021; Árnason et al. 2023). In particular,
42 the Beta-coalescent (Schweinsberg 2003) is a specific type of Lambda-coalescent that has been used to
43 explain the shallow genealogies observed in cod (Birkner et al. 2013), to study pathogen superspreading
44 (Hoscheit and Pybus 2019) and to characterise *Mycobacterium tuberculosis* outbreak genealogies
45 (Menardo et al. 2021).

46 Let us take as our starting point an unrooted, undated tree produced by a maximum likelihood tree
47 reconstruction software such as RAxML (Stamatakis 2014), IQ-TREE (Minh et al. 2020) or PhyML
48 (Guindon et al. 2010). Such a tree may contain polytomies, where a node leads to more than two
49 branches. This can be either because of a multiple merger event, or because of at least one branch
50 covering a short interval of time, so that no substitution occurred as expected under any molecular
51 clock model (Bromham and Penny 2003). Multiple heuristic approaches have been developed for either

52 breaking up or collapsing polytomies in undated phylogenies (Kuhn et al. 2011; Lin et al. 2011; Lewis
53 et al. 2005). Here instead we use a Lambda-coalescent framework to infer which polytomies are caused
54 by multiple merger events and which are caused by a lack of phylogenetic signal. We do so in the
55 context of dating the tree, which means to use it as well as the dates of the leaves in order to produce
56 a dated phylogeny (To et al. 2016; Volz and Frost 2017; Didelot et al. 2018). The dated phylogeny
57 can then be used for a broad range of epidemiological investigations (Didelot and Parkhill 2022). To
58 reconstruct this dated phylogeny, we must infer the root position, ancestral node times, as well as
59 parameters associated with the clock and genealogical models. We must distinguish which polytomies
60 are consistent with multiple mergers, and which are better explained by quick binary branching and
61 therefore should be resolved. In the latter case we must also estimate the branching order within
62 the polytomies returned by the maximum likelihood estimation software. This is important as the
63 branching order within the polytomies is random and likely inconsistent with the temporal structure
64 of the tree, as previously noted (Sagulenko et al. 2018).

65 To achieve these aims several issues must be addressed. We need to choose a set of prior models for
66 the latent genealogies which take into account biological realism and statistical tractability. We focus
67 on the Beta-coalescent (Schweinsberg 2003) and an extension of it described by Eldon and Stephan
68 (2023) in which the Beta-coalescent and Kingman's coalescent are mixed together, combining low-
69 variance family size reproduction with occasional high-variance sweepstakes. We also need to specify
70 a molecular clock model to establish the relationship between dated and undated phylogenies, and for
71 this we use the Additive Relaxed Clock (ARC) model (Didelot et al. 2021). Next, we need to specify a
72 computational scheme for representing multiple merger trees where a single node may have an arbitrary
73 number of descendants. This representation needs to enable efficient computation of likelihoods and to
74 be statistically efficient. To this end we use the Billera-Holmes-Vogtmann (BHV) space (Billera et al.
75 2001) augmented with a spike-and-slab construction (George and McCulloch 1993). Finally we design
76 a Markov Chain Monte-Carlo (MCMC) sampling scheme targeting the posterior in order to perform
77 Bayesian inference.

78 METHODS

79 Lambda-coalescent

80 Lambda-coalescents are a class of stochastic genealogical processes (Pitman 1999; Sagitov 1999;
81 Donnelly and Kurtz 1999) that generalise the popular Kingman's coalescent (Kingman 1982) to a
82 setting where more than two lineages may merge into the same parent, i.e. they permit multiple
83 mergers. These processes commonly describe genealogies arising from various forwards in time models
84 in population genetics, typically in scenarios where there is a high variability in the number of surviving
85 offspring, or when selection and recombination are taken into account (Berestycki 2009; Tellier and
86 Lemaire 2014). Examples of such scenarios include heavy-tailed offspring distributions (Schweinsberg
87 2003; Matuszewski et al. 2018), recurrent selective sweeps in presence of recombination (Durrett and
88 Schweinsberg 2005), rapidly adapting populations (Neher and Hallatschek 2013; Desai et al. 2013), as
89 well as strong purifying selection (Cvijović et al. 2018).

90 Lambda-coalescents are uniquely specified by a finite measure Λ on $[0, 1]$ that governs the merger
91 sizes of the process. Intuitively, one can think of this relationship as sampling a probability $p \sim \Lambda$
92 proportional to the density of Λ , selecting with probability p each of the lineages currently in the
93 process, and merging all selected lineages. The instantaneous block merger rates $\lambda_{b,k}$, that is the rate
94 at which every $k \geq 2$ lineages merges into one parent when there are b lineages in total, are then given
95 by

$$\lambda_{b,k} = \int_0^1 p^{k-2} (1-p)^{b-k} \Lambda(dp). \quad (1)$$

96 The factor of p^{-2} arises from the fact that at least two lineages must participate in order for a merger
97 to happen. When $\Lambda = \delta_0$, that is a point mass at zero, we can see that

$$\lambda_{b,k} = \begin{cases} 1, & \text{if } k = 2 \\ 0, & \text{otherwise} \end{cases} \quad (2)$$

98 In other words the resulting Lambda-coalescent is Kingman's coalescent in which only two lineages
99 may merge.

100 **Beta-coalescent**

101 The approach presented will mostly focus on the Beta-coalescent and an extension of it. The rationale
102 for this is two-fold. First, the Beta-coalescent is relatively well studied, admits the frequently used
103 Kingman's coalescent as a special case, and the instantaneous block merger rates are available in
104 closed form. Second, it arises from models in which the variance of the offspring distribution can be
105 very high, and has been considered in the context of pathogen populations before, see for example
106 (Hoscheit and Pybus 2019). As the name suggests, in the case of the Beta-coalescent the measure Λ
107 is simply the Beta distribution. Usually in the context of Lambda-coalescents, the Beta distribution
108 is parameterised as (Schweinsberg 2003)

$$\Lambda = \text{Beta}(2 - \alpha, \alpha) \quad \alpha \in [1, 2] \quad (3)$$

109 The reason for this is a connection to a model of populations with skewed offspring distributions in
110 which α governs the tail behaviour of the offspring distribution. The forwards-in-time model in the
111 derivation of (Schweinsberg 2003) follows dynamics of a supercritical Galton–Watson process where in
112 successive non-overlapping generations each of the N individuals produces ν_i offspring i.i.d. according
113 to a distribution with the tail index $k^{-\alpha}$. This implies that the offspring distribution has infinite
114 variance if $1 < \alpha < 2$ and infinite mean if $\alpha = 1$. Offspring are then randomly killed in order to
115 keep the population size constant and equal to N . Schweinsberg (2003) showed that the genealogies
116 arising for this process converge to Kingman's coalescent for values of $\alpha \geq 2$ and to the Beta-coalescent
117 parameterised as in Equation 3 for values of $\alpha \in (2, 1]$, under suitable time-rescalings and as $N \rightarrow \infty$.
118 In plain words if the distribution of the number of offspring produced by any of the individuals is
119 sufficiently skewed, and this situation arises frequently enough, every once a while an individual may
120 get lucky and produce enough offspring to replace a non-negligible fraction of the population. Therefore
121 the probability that multiple individuals find the same ancestor at once does not vanish in the large
122 population size limit. The parameter α relates to how skewed the offspring distribution is. The limiting
123 case of $\alpha = 1$ corresponds to the Bolthausen-Sznitman coalescent (Bertoin and Le Gall 2000), which
124 has been shown to arise in scenarios corresponding to rapid adaptation and clonal interference (Desai
125 et al. 2013; Neher and Hallatschek 2013; Schweinsberg 2017). On the other hand, in the limiting case
126 of $\alpha = 2$, the Beta distribution collapses into an atom at 0 and thus the resulting coalescent is the
127 Kingman's coalescent (Schweinsberg 2003).

128 There are several limitations of the Beta-coalescent. First amongst these is the assumption that every
129 individual may produce a large number of offspring in every generation. Often it is easier to imagine
130 that such large reproductive events may only occur if correct circumstances are met. Related to this is
131 the second limitation. The Beta-coalescent has a time scale of $N^{1-\alpha}$ if $\alpha \in (1, 2]$ and $\log(N)^{-1}$ if $\alpha = 1$.
132 For many populations, this implies mutation rates or population sizes orders of magnitude higher than
133 what would be biologically realistic. This is especially the case if α is close to 1 as previously noted
134 (Eldon and Stephan 2023). This may however not be an issue for moderately large values of α in the
135 case of within-host evolution where the population sizes are going to be very large.

136 **Extended Beta-coalescent**

137 We now introduce an extension of the Beta-coalescent which we will refer to as the extended Beta-
138 coalescent. This modification is a mixture of the Beta-coalescent and Kingman's coalescent. This is
139 achieved by defining the measure Λ characterising this process as

$$\Lambda = \delta_0 + c\text{Beta}(2 - \alpha, \alpha) \quad c \in [0, \infty), \alpha \in [1, 2] \quad (4)$$

140 The reasons for introducing this are two-fold. For one it will provide us with a convenient example of
141 a Lambda-coalescent of a form shared by for example the Durrett–Schweinsberg coalescent (Durrett
142 and Schweinsberg 2005) arising from selective sweeps. The second reason is based on the modifications
143 to random sweepstakes reproduction presented in (Eldon and Stephan 2023). They considered a
144 modification to the construction in (Schweinsberg 2003) to address the problematic assumption of
145 very frequent large family sizes. In this modification, in each generation a coin with probability ϵ is
146 flipped. On success with probability ϵ each individual produces offspring according to an offspring
147 distribution with a “small” $\alpha \in [1, 2)$ and with probability $1 - \epsilon$ according to an offspring distribution
148 $\alpha \geq 2$. For a suitable choice of $\epsilon = \epsilon_N \rightarrow 0$ as $N \rightarrow \infty$, the resulting coalescent process is exactly the
149 Lambda-coalescent specified by Equation 4.

150 **Selection of priors**

151 For the Beta-coalescent we parametrise the measure characterising the genealogical prior as:

$$\Lambda = \nu \text{Beta}(1 - \alpha^*, 1 + \alpha^*) \quad (5)$$

152 ν and α^* are unknown parameters we wish to infer. $\nu \in \mathbb{R}_+$ is the time scale of the process and
 153 $\alpha^* \in [0, 1]$ controls the Beta distribution governing the merger sizes. α^* relates to the original
 154 parameter α from Equation 3 (Schweinsberg 2003) as $\alpha = \alpha^* + 1$. We equip these parameters with
 155 the following prior distributions:

$$\begin{aligned} \log \nu &\sim \text{Normal}(0, 4) \\ \alpha^* &\sim \text{Beta}(3, 1) \end{aligned} \quad (6)$$

156 We caution that it is not straightforward to interpret ν as the timescale of the Beta-coalescent is
 157 $1/N^{1-\alpha}$ and therefore ν only corresponds to the usual effective population size if $\alpha = 2$.

158 For the extended Beta-coalescent we parameterise the genealogical prior as:

$$\Lambda = \nu((1 - \phi)\delta_0 + \phi \text{Beta}(1 - \alpha^*, 1 + \alpha^*)) \quad (7)$$

159 $\nu \in \mathbb{R}_+$ corresponds to the process rate, $\phi \in [0, 1]$ is the mixing proportion between the Kingman
 160 component and the Beta component, and $\alpha^* \in [0, 1]$ once again controls the Beta distribution and
 161 therefore the size of multiple mergers. We equip these parameters with the following prior distribution:

$$\begin{aligned} \log \nu &\sim \text{Normal}(0, 4) \\ \phi &\sim \text{Beta}(1, 3) \\ \alpha^* &\sim \text{Beta}(1, 2) \end{aligned} \quad (8)$$

162 Finally, for the molecular clock model we use the Additive Relaxed Clock (ARC) (Didelot et al. 2021)
 163 in which a branch of length l carries a number of substitutions x distributed as:

$$x \sim \text{NegBin}\left(\frac{\mu l}{\omega}, \frac{\omega}{1 + \omega}\right) \quad (9)$$

164 where μ represents the mean clock rate and ω controls the amount of relaxation relative to a strict
165 clock model. We use the following prior distribution:

$$\begin{aligned}\mu &\sim \text{Gamma}(2, 8) \\ \omega &\sim \text{Normal}(0, 2) \mathbb{1}_{[0, \infty)}\end{aligned}\tag{10}$$

166 **Billera-Holmes-Vogtmann space embedding**

167 In order to efficiently sample genealogies that admit multiple mergers we leverage two concepts. First
168 we embed binary trees in a space of phylogenetic trees where coordinates correspond to branch lengths,
169 specifically the Billera-Holmes-Vogtmann (BHV) space. Second we use a spike-and-slab construction
170 to put positive mass on the set of trees where at least one branch length is shrunk to exactly zero and
171 identify these trees with multiple merger trees obtained by collapsing all branches with lengths equal
172 to exactly zero.

173 BHV space is a metric space introduced in the study of phylogenetic tree geometry (Billera et al.
174 2001). For a fixed number of tips n the BHV space is constructed from a set of $(2n - 3)!!$ orthants
175 in $(0, \infty)^{n-2}$ each corresponding to a particular topology of a rooted binary tip labeled n -tree. Each
176 coordinate within the $n - 2$ dimensional orthant corresponds to the length of an interior branch in
177 the given tree topology. At any of the zero boundaries, the binary topology degenerates into a tree
178 that has branching greater than two. The orthants are “glued” together at boundaries corresponding
179 to the same topology. For example if only one coordinate approaches zero there is a junction of three
180 different $n - 2$ dimensional orthants corresponding to an $n - 3$ dimensional orthant face. For existence
181 of centroids, proof of constant negative curvature and properties of geodesics refer to (Billera et al.
182 2001). We will base our sampling scheme construction around the BHV space. The BHV space has
183 been used previously in phylogenetic reconstruction for example to construct an embedding amenable
184 to sampling binary trees using discrete Hamiltonian Monte Carlo (Dinh et al. 2017). Typically such
185 approaches simply select the base measure to be Lebesgue within each of the $n - 2$ dimensional orthants
186 corresponding to binary topologies, and each $n - 2$ dimensional orthant is assigned equal probability.
187 The lower dimensional orthants are then null sets within this measure space. However such an approach
188 is not appropriate for inference with multiple merger trees as any set consisting purely of these would

189 be given a zero probability. One option would be to put mass on lower dimensional orthants, and set
190 up a trans-dimensional sampling scheme using reversible jump MCMC (Green 1995). However this
191 would be unlikely to work efficiently as designing moves that remove or add more than one branch at
192 a time and don't suffer from a rapidly diminishing acceptance ratio would be challenging, and would
193 not take advantage of the natural geometry of the space.

194 We will therefore augment the BHV space using a spike-and-slab construction (George and McCulloch
195 1993). Denote by \mathcal{T} the set of all rooted, labeled, metric n -trees. An n -tree is said to be metric if all
196 of its branch lengths are strictly greater than 0. Let \mathbb{T} denote the set of all labeled, rooted, binary
197 n -tree topologies. Denote the closed $n - 2$ dimensional orthant of the BHV space corresponding to a
198 particular binary n -tree topology as V^τ for $\tau \in \mathbb{T}$. We identify points in the closed $n - 2$ dimensional
199 orthants of BHV space by the tuple (X, Q, τ) where $X \in [0, \infty)^{n-2}$ denotes the location within an
200 orthant, $Q \in \{0, 1\}^{n-2}$ is a vector of indicators where $q_i = 1$ if and only if $x_i = 0$ and $\tau \in \mathbb{T}$ denotes
201 the orthant index, i.e. the corresponding binary topology. We next identify all points on the boundary
202 ∂V^τ , that is all points with at least one coordinate equal to zero with the corresponding k -ary metric
203 n -tree by collapsing each and every branch for which $q_i = 1$, i.e. that is of length exactly zero. We
204 now define the base measure to be

$$\nu(dx \times d\tau) = \delta_\tau(d\tau) \otimes \mu(dx) \quad (11)$$

205 That is we assign uniform mass to each of the $n - 2$ dimensional orthants and within each orthant we
206 have

$$\mu = (\delta_0 + \mu_0)^{n-2} \quad (12)$$

207 Where δ_0 is an atom at zero and μ_0 is the Lebesgue measure on $[0, \infty)$. This construction assigns
208 positive probability to sets of binary trees with one or more branches with length exactly zero. By
209 identifying such trees with a (metric) multiple merger tree we can see that it therefore puts a positive
210 probability on trees with multiple mergers.

²¹¹ **Parametrisation of the genealogy**

Having outlined the construction above we now give an explicit parametrisation of the genealogy that will allow us to construct an MCMC scheme as well as enable the computation of necessary quantities. As an input we assume we are given a rooted binary phylogeny with n tips, with branch lengths corresponding to the estimated number of substitutions along a branch. This phylogeny is assumed to be a point estimate obtained by ML phylogenetic software. The root position may be assumed to be known a-priori, or to be estimated, in which case the initial rooting is assumed to be chosen arbitrarily. In practice the estimated number of substitutions along a branch may not be an integer even though it is likely to be very close to one when the number of substitutions per site is low. The clock models used require an integer number of mutations. Therefore all branch lengths are coerced to integer values by rounding. Based on the undated input phylogeny we can construct a rooted binary genealogy denoted as τ as follows. We define the following notations:

$\{1, 2, \dots, 2n - 1\}$	Set of all nodes
$\mathcal{S} := \{1, 2, \dots, n\}$	Set of tips
$\mathcal{I} := \{n + 1, n + 2, \dots, 2n - 1\}$	Set of internal nodes
$\{2n - 1\}$	Root node
$\text{pa}_\tau(i), i \in \{1, \dots, 2n - 2\}$	Parent of node i
$m_\tau(i), i \in \{1, \dots, 2n - 2\}$	Number of mutations on edge above node i
$s_\tau(i), i \in \mathcal{S}$	Sampling date of tip i
$j \prec_\tau i$	Node j is a descendant of node i
$\text{desc}_\tau(i) := \{j \in \mathcal{S} : j \prec_\tau i\}$	Set of leaves descendant of node i
$\mathcal{C}_\tau := \{c_l, c_r\}$	Children of the root node

²¹² Having defined necessary notation for characterising the rooted tree topology we aim to define the
²¹³ parametrise the internal node heights in the rooted genealogy. This corresponds to the within orthant
²¹⁴ parametrisation. We now proceed to define the collection of node times $\mathbf{t} := \{t_i\}_{1 \leq i \leq 2n-1}$. A subset
²¹⁵ of these variables corresponds to the internal node times, that is the free parameters of the model. We
²¹⁶ denote these as $\mathbf{t}^I := \{t_i \in \mathbf{t} : i \in \mathcal{I}\}$. The remaining first n variables correspond to sampling times

217 and these are considered as an input. Denote these set of sampling times by $\mathbf{t}^S := \{t_i \in \mathbf{t} : i \in S\}$. In
 218 practice, the vector \mathbf{t} needs to satisfy a complex set of constraints, that is

$$i \prec_{\tau} j \implies \begin{cases} t_i \leq t_j & i > n \\ t_i < t_j & i \leq n \end{cases} \quad (13)$$

219 Where $i \prec_{\tau} j$ denotes the partial order induced by the tree topology, that is $i \prec_{\tau} j$ iff i is a descendant
 220 of j .

221 The complex set of constraints that the node times have to satisfy would make defining a Metropolis-
 222 Hastings move capable of targeting the boundaries complicated. Hence we define the vector of positive
 223 height variables $\chi \in \mathbb{R}_+^{n-2}$ as well as the tree height variable $H \in \mathbb{R}$. Once again for convenience we
 224 will abuse notation and denote by $\chi(i)$ the component of χ associated with node i . This is in contrast
 225 to χ_i which denotes the i -th component of the vector χ . To define the transformation for χ to \mathbf{t} we
 226 first need to introduce the set of lower bounds b_{τ}

$$b_{\tau}(i) = \max_{j: \text{desc}_{\tau}(i)} s_{\tau}(j), \quad i \in [2n-1] \quad (14)$$

227 We now map χ, H to \mathbf{t} via the mapping $g : (H, \chi) \mapsto \mathbf{t}$. The first part of the mapping parametrises
 228 the height of the root node by transforming H

$$t_{2n-1} = e^H + b_{\tau}(2n-1) \quad (15)$$

229 The second part of mapping parametrises the height of internal non-root nodes by transforming χ

$$t_i = (t_{\text{pa}_{\tau}(i)} - b_{\tau}(i))e^{-\chi(i)} + b_{\tau}(i) \quad (16)$$

230 Note that $\chi(i) = 0$ implies that the length of the branch above node i scaled in time units is also 0.
 231 An illustration of how such parametrisation looks in practice can be found in Figure 1.




Figure 1: Example of tree parametrisation. In this case the bounds are $b(6) = s_1, b(7) = s_2, b(8) = s_3, b(9) = s_3$. The components χ are mapped as $\chi(6) = \chi_1, \chi(7) = \chi_2, \chi(8) = \chi_3$.

232 This parametrisation is reminiscent of the ratio transform of (Ji et al. 2021), with the key difference
 233 that we express non-root internal node heights in terms of a transform of the distance from the parent
 234 constrained so that the age of the child is strictly less than the oldest descendant tip as opposed to
 235 a transform of the ratio of the remaining height. This is because it is necessary in the construction
 236 presented here for the boundary corresponding to 0 branch length to be accessible. Finally we introduce
 237 the vector of indicators that determine whether a coordinate is allocated to the boundary:

$$\begin{aligned} \mathbf{q} &\in \{0, 1\}^{n-1} \\ q(i) = 1 &\Leftrightarrow \chi(i) = 0 \end{aligned} \tag{17}$$

²³⁸ We can now parametrise a genealogy by the tuple $\mathbf{Y} := (H, \chi, \mathbf{q}, \tau; \mathbf{s})$ where τ specifies the binary
²³⁹ topology and hence a full dimensional orthant of the BHV space, and H, χ, \mathbf{q} characterise the position
²⁴⁰ within the orthant conditional on the tip date constraints \mathbf{s} .

²⁴¹ **Lambda-coalescent likelihood**

²⁴² In order to compute the coalescent likelihood of a genealogy, we first need to convert the (possibly
²⁴³ non-metric) binary tree embedding to the respective metric multiple merger tree. To achieve this first
²⁴⁴ denote the set of all internal nodes descending from node i as

$$r(i) := \{j \in [2n - 1] : j \prec_{\tau} i\} \quad (18)$$

²⁴⁵ Next defining the set of internal descendants of i coincident with i as

$$z(i) := \{j \in r(i) : t(i) = t(j)\} \quad (19)$$

²⁴⁶ We can see that for any internal node i such that $q(i) = 0$, $|z(i)|$ corresponds to the merger size
²⁴⁷ minus 2. Using these to definitions we can compute the likelihood of the binary tree embedding
²⁴⁸ $\mathbf{Y} := (H, \chi, \mathbf{q}, \tau; \mathbf{s})$ under the Lambda-coalescent Λ^{θ} with block merger rates $\lambda_{b,k}^{\theta}$ as:

$$\begin{aligned} f_{\theta}^{\Lambda}(\mathbf{Y}) &= \prod_{\substack{i \in I: \\ q_i = 0}} \lambda_{A(t_i), |z(i)|+2}^{\theta} \\ &\times \prod_{\substack{1 < i \leq 2n-1: \\ q_i = 0}} e^{-R(A(t_i))(t_i - \max\{t_j : t_j < t_i\})} \\ &\times J_g \Big|_{\mathbf{Y}} \\ &\times \frac{1}{Z(\mathbf{Y})} \end{aligned} \quad (20)$$

²⁴⁹ Where $R : \mathbb{N} \mapsto \mathbb{R}_+$ denotes the total coalescent rate of the process

$$R(b) := \sum_{k=1}^b \binom{b}{k} \lambda_{b,k}^{\theta} \quad (21)$$

250 and $A : \mathbb{R}_+ \mapsto \mathbb{Z}_+$ denotes the lineages through time function at time t , defined as

$$A(t) := \sum_{i \in \mathcal{S}} \mathbb{1}_{t > t_i}(t_i) - \sum_{i \in \mathcal{I}} \mathbb{1}_0(q_i) \mathbb{1}_{t > t_i}(t_i) (|z(i)| + 1) \quad (22)$$

251 This is also known as the block counting process (Kukla and Möhle 2018). As the genealogical prior
 252 is expressed in terms of node heights \mathbf{t} , we must account for the transformation from χ to \mathbf{t} in the
 253 density. To do so we require the Jacobian for the mapping $g : (H, \chi) \mapsto \mathbf{t}$. This is straightforward to
 254 compute as the matrix of first order partials has a diagonal structure and hence is equal to

$$J_g \Big|_{\mathbf{Y}} = e^H \prod_{i \in \mathcal{I}} \left[\mathbb{1}_0(q(i)) (t_{\text{pa}_{\mathbf{r}}(i)} - s_{\mathbf{r}}(i)) e^{-\chi(i)} + \mathbb{1}_1(q(i)) \right] \quad (23)$$

255 Note that as the density of all $\chi_i = 0$ is with respect to an atomic measure these do not play a role within
 256 the relevant Jacobian adjustment. Finally we note that the embedding of a given multimerger tree as
 257 a binary tree is not unique. Therefore, in order for the density of a given tree to be proportional to the
 258 density given by the Lambda-coalescent density, we need to re-weight the density of the embedding to
 259 account for the overcounting. In general given a multiple merger of size m there are $(2m - 3)!!$ ways
 260 to resolve this as a sequence of binary mergers. This is the number of rooted labeled binary trees with
 261 m tips (Billera et al. 2001). Therefore the adjustment for the embedding Y is:

$$Z(\mathbf{Y}) = \prod_{i \in \mathcal{I}} (2|z(i)| + 4 - 3)!! \quad (24)$$

262 Branch likelihood

263 The likelihood is based on the ARC model (Didelot et al. 2021). Each branch in the input ML phylogeny
 264 is scaled in estimated number of substitutions along that branch. This number of substitutions is
 265 distributed according to Equation 9. The number of substitutions returned by ML software will in
 266 general not be an integer, although in many cases will be close to an integer value. As the ARC model
 267 requires integer valued number of substitutions all values are rounded to the nearest integer. Define the
 268 length of a branch above node i as $l(i) = t_{\text{pa}_{\mathbf{r}}(i)} - t_i$. Given the current genealogy state parameterised
 269 by \mathbf{Y} , and the clock parameters μ and ω , the likelihood for all branches not incident to the root can

270 be computed as:

$$\prod_{\substack{i \in I: \\ i \notin \mathcal{C}_\tau \cup \{2n-1\}}} \text{NegBin} \left(m_\tau(i); \frac{\mu l(i)}{\omega}, \frac{\omega}{1+\omega} \right) \quad (25)$$

271 This has to be multiplied by the likelihood of the branches incident to the root. This depends on
 272 whether the root is considered fixed, or if it is unknown. If the root is fixed, the branches incident to
 273 root are treated as any other branch and their contribution is

$$\prod_{i \in \mathcal{C}_\tau} \text{NegBin} \left(m_\tau(i); \frac{\mu l(i)}{\omega}, \frac{\omega}{1+\omega} \right) \quad (26)$$

274 If the root position is considered to be unknown then the position of the root on the branch is
 275 marginalised out and the contribution becomes

$$\text{NegBin} \left(m_\tau(c_l) + m_\tau(c_r); \frac{\mu(l(c_l) + l(c_r))}{\omega}, \frac{\omega}{1+\omega} \right) \quad (27)$$

276 This follows from the additive nature of the ARC clock model (Didelot et al. 2021).

277 MCMC scheme

278 We use four types of moves to sample parameters characterising Y . These moves can be categorised
 279 into three families. The first family consists of moves covering transitions that update the position
 280 within a full dimensional orthant including the boundary. The moves within the first family update
 281 branch lengths, and therefore merger sizes and internal node heights. The second family consists of
 282 moves for proposing transitions between full dimensional orthants. The moves within the second family
 283 update the root position and branching order within polytomies. The third family consists of a single
 284 move that updates the parameters of the observation model and genealogy model. The scheme consists
 285 of a single sweep through all three families, selecting one move from each family uniformly at random.

286 Orthant interior move

287 The orthant interior move is a random walk Metropolis (RWM) (Metropolis et al. 1953) move restricted
 288 to updating a subset of the vector χ corresponding to those coordinates which are currently not

²⁸⁹ restricted to the boundary, and thus a part of a multiple merger

$$k_{\text{RWM}}(\chi'_i \mid \chi_i, q_i)_{\sigma} = \mathbb{1}_0(q_i)\delta_0 + |\mathcal{N}|(\chi'_i - \chi_i; \sigma_i), \quad (28)$$

²⁹⁰ where $|\mathcal{N}|(\cdot; \sigma)$ denotes the density of the modulus of a normal random variable with variance σ^2 . As
²⁹¹ expected for RWM the proposal ratio equals to one.

²⁹² **Orthant boundary move**

²⁹³ The orthant boundary move is responsible for transitioning between binary topologies and multiple
²⁹⁴ mergers. To do so the move needs to propose transitions between the boundary and orthant interior.
²⁹⁵ The main challenge with designing this move is that due to the structure of the likelihood of the
²⁹⁶ extended Beta-coalescent there is a relatively sharp ridge between a 2-merger and a 3-merger when the
²⁹⁷ number of active lineages is large, i.e. $b \gg 3$. Therefore the move must be able to propose transitions
²⁹⁸ that move several coordinates of χ to 0 or back at once. The first step consists of sampling an internal
²⁹⁹ node index $i \in \mathcal{I} \setminus \{2n - 1\}$ at random with probability

$$P[i = x \mid \chi] \propto w_1(\chi(x)) \quad (29)$$

³⁰⁰ If the coordinate is allocated to the interior (therefore to the slab component of the base measure), i.e.
³⁰¹ if $q(i) = 0, \chi(i) > 0$ then the move begins by proposing $\chi'(i) = 0$. Denoting the parent node above
³⁰² $j = \text{pa}(i)$ the move continues upwards and proposing to shrink $\chi(j)$ to 0 with probability $w_2(\chi(x))$
³⁰³ as long as it is accessible to this move, i.e. $q(\text{pa}(j)) = 0$ and $j \neq 2n - 1$. This process repeats until a
³⁰⁴ coordinate fails to shrink either due to it not being accessible or due to the coin flip failing.
³⁰⁵ If the coordinate initially selected is allocated to the boundary then there are two options. The move
³⁰⁶ either expands that coordinate, proposing $\chi'(i) > 0$. Alternatively the move expands that coordinate
³⁰⁷ and then attempts to expand the coordinate above it if it is allocated to the boundary. If the coordinate
³⁰⁸ above it is allocated to the boundary it is expanded and this procedure repeats, terminating when either
³⁰⁹ the root is reached or a non-zero coordinate is reached. The expanded coordinates are sampled from
³¹⁰ a proposal distribution d , i.e. $\chi' \sim d$. Whether the move expands coordinates recursively or not
³¹¹ is decided uniformly at random. Denote the sequence of nodes that have had their corresponding

³¹² coordinates modified by v_1, v_2, \dots, v_m , where v_1 is the first node modified by the move. If the initial
³¹³ coordinate chosen was allocated to the interior the proposal ratio is equal to

$$a(\chi', \chi) = \frac{w_1(\chi'(v_1)) \sum_{j \in I \setminus \{2n-1\}} w_1(\chi(j))}{w_1(\chi(v_1)) \sum_{j \in I \setminus \{2n-1\}} w_1(\chi'(j))} d(\chi(v_1))^{1-2q(v_1)} \\ \times \left(\mathbb{1}_{x \geq 2}(n) \prod_{v_j: 2 \leq j \leq n} \left(\frac{d(\chi(v_j))}{w_2(\chi(v_j))} \right)^{1-2q(v_1)} + \mathbb{1}_1(n) \right) \\ \times \left[\frac{1}{2} (1 + \mathbb{1}_1(n)) \right]^{1-2q(v_1)} \\ \times \left(\frac{1}{P_s} \right)^{1-2q(v_1)} \quad (30)$$

³¹⁴ Note that $1 - 2q(v_1) = 1$ if the move is shrinking coordinates towards the boundary and -1 if it is
³¹⁵ expanding coordinates. Therefore it determines the direction of the move. The first term corresponds
³¹⁶ to the probability of selecting the same starting node, the second term corresponds to the likelihood
³¹⁷ of the coordinate transformation for subsequent nodes, the third accounts for the possibility of the
³¹⁸ reverse move being chosen, and the fourth term corresponds accounts for the stopping probability of
³¹⁹ the recursion, P_s which is equal to

$$P_s = \begin{cases} 1 - w_2(\text{pa}(v_n)) & \text{if } \text{pa}(v_n) \text{ is accessible} \\ 1 & \text{otherwise} \end{cases} \quad (31)$$

³²⁰ Crucially setting $w_2(\chi)$ to be equal to the density of $d(\chi)$ leads to the second term cancelling to one.
³²¹ In practice we use

$$d(x) = \frac{\sqrt{2}}{b\sqrt{\pi}} \exp -\frac{x^2}{2b^2} \mathbb{1}_{x \geq 0} \quad (32)$$

$$w_2(x) = w_1(x) = \exp -\frac{x^2}{2a^2}$$

³²² Setting $a = b = \sqrt{2/\pi}$ leads to $d(x) = w_1(x) = w_2(x)$.

³²³ Root NNI move

³²⁴ The root NNI move updates the root position, and is used if the root position is considered unknown.
³²⁵ It is a version of the nearest neighbour interchange (NNI) commonly used in phylogenetic inference
³²⁶ (Yang and Rannala 1997). This moves first proceeds in selecting one of the child nodes of the root as a

327 pivot. The move then moves the root to one of the descendant edges of the pivot. The pivot is chosen
 328 with equal probability from both root descendant nodes at random with the exception of two special
 329 cases. If one of the descendant nodes is a tip node the other descendant of the root is selected as a
 330 pivot. This is because NNI move is undefined for a tip node selected as a pivot. The other special case
 331 is if one of the root descendants has the branch above it collapsed as a part of a multiple merger but
 332 the other does not. In this case the descendant with the edge collapsed is chosen as the pivot. This is
 333 to prevent the move from wasting computational effort as this way the node with edge collapsed stays
 334 adjacent to the root. Otherwise it may become adjacent to an edge that cannot support a multiple
 335 merger leading to 0 likelihood. With the pivot selected the branch to move the root to is sampled with
 336 equal probability from two edges that descend from the pivot. We denote the pivot by $p \in \mathcal{C}_\tau$ and its
 337 probability mass function under the topology τ

$$P_\tau[p = i] = \begin{cases} \frac{1}{2}, & \text{if } (c_l > n) \wedge (c_r > n) \wedge (q(c_l) = q(c_r)) \\ \mathbb{1}_{c_l}(i), & \text{if } (c_r \leq n) \vee ((q(c_l) = 1) \wedge (q(c_r) = 0)) \\ \mathbb{1}_{c_r}(i), & \text{if } (c_l \leq n) \vee ((q(c_r) = 1) \wedge (q(c_l) = 0)) \end{cases} \quad (33)$$

338 With p fixed, sample j from the two descendants of p uniformly at random. Based on this generate a
 339 new rooted tree topology τ' , which is the same as τ except where:

$$\text{pa}_{\tau'}(i) = \begin{cases} p, & \text{if } i \in \mathcal{C}_\tau \wedge i \neq p \\ 2n - 2, & \text{if } i = j \\ \text{pa}_\tau(i), & \text{otherwise} \end{cases} \quad (34)$$

340 The mutations above the affected nodes are then adjusted accordingly:

$$m_{\tau'}(i) = \begin{cases} m_\tau(c_r) + m_\tau(c_l), & \text{if } i \in \mathcal{C}_\tau \wedge i \neq p \\ \lfloor m_\tau(j)s \rfloor, & \text{if } i = j \\ \lfloor m_\tau(j)(1-s) \rfloor, & \text{if } i = p \\ \text{pa}_\tau(i), & \text{otherwise} \end{cases} \quad (35)$$

³⁴¹ For some arbitrarily chosen $s \in [0, 1]$.

³⁴² The proposal ratio is

$$a(\boldsymbol{\tau}', \boldsymbol{\tau}) = \frac{P_{\boldsymbol{\tau}'}[p = i]}{P_{\boldsymbol{\tau}}[p = i]} \quad (36)$$

³⁴³ **Polytomy neighbour interchange move**

³⁴⁴ The polytomy neighbour interchange move changes the binary topology within polytomies as this
³⁴⁵ topology is randomly resolved (if at all) by the ML estimation program. It does so by first sampling a
³⁴⁶ pivot node p from all nodes for which the edge above contains 0 mutations. With the pivot selected a
³⁴⁷ node j is chosen uniformly at random from all the descendants of p that are adjacent to the polytomy.
³⁴⁸ That is from all descendant nodes such that the parent of that node is 0 mutations away from p .
³⁴⁹ Denote this set of nodes $A_{\boldsymbol{\tau}}$. Next a node k is selected from all nodes in the outgroup relative to p
³⁵⁰ such that the distance of the parent of those nodes to the parent of p is 0 mutations. Denote this set of
³⁵¹ nodes $B_{\boldsymbol{\tau}}$. Based on this generate a new rooted tree topology $\boldsymbol{\tau}'$, which is the same as $\boldsymbol{\tau}$ except where:

$$\text{pa}_{\boldsymbol{\tau}'}(i) = \begin{cases} \text{pa}_{\boldsymbol{\tau}}(j), & \text{if } i = k \\ \text{pa}_{\boldsymbol{\tau}}(k), & \text{if } i = j \\ \text{pa}_{\boldsymbol{\tau}}(i), & \text{otherwise} \end{cases} \quad (37)$$

³⁵² The reverse move consists of selecting the same pivot and then the corresponding descendants. Denote
³⁵³ the sibling of the pivot p with s_p . The proposal ratio is

$$a(\boldsymbol{\tau}', \boldsymbol{\tau}) = \frac{|A_{\boldsymbol{\tau}'}|^{-1} |B_{\boldsymbol{\tau}'}|^{-1} + (1 - \mathbb{1}_{s_p}(k))(|A_{\boldsymbol{\tau}'}| + 1)^{-1} (|B_{\boldsymbol{\tau}'}| - 1)^{-1}}{|A_{\boldsymbol{\tau}}|^{-1} |B_{\boldsymbol{\tau}}|^{-1} + (1 - \mathbb{1}_{s_p}(k))(|A_{\boldsymbol{\tau}}| + 1)^{-1} (|B_{\boldsymbol{\tau}}| - 1)^{-1}} \quad (38)$$

³⁵⁴ **Tuning the sampler**

³⁵⁵ In order to tune the sampler a preconditioning matrix of individual parameter variances is estimated,
³⁵⁶ along with a scale factor multiplying this matrix. The RWM moves updating the parameters and node
³⁵⁷ height variables are then scaled accordingly. This is done in a sequence of six steps. Let T denote
³⁵⁸ the thinning factor. The first two steps consists of running the sampler with orthant boundary moves

359 disabled. The first step consists of burn-in for $100 \times T$ iterations. The second step then estimates the
360 χ preconditioning matrix for $100 \times T$ iterations. For the following steps all moves are enabled. The
361 third step consists of burn-in for $100 \times T$ steps. The fourth step estimates the preconditioning matrix
362 for the parameter moves for $100 \times T$ steps. The fifth and sixth steps estimate the step scaling for
363 $50 \times T$ steps. All previous iterations are then discarded.

364 RESULTS

365 Implementation

366 We implemented the approach as described in an new R package titled *MMCTime*, which is available
367 at
368 <https://github.com/dhelekal/MMCTime>. The package uses *ape* (Paradis and Schliep 2019) as a
369 backend for handling phylogenies. *bayesplot* is used for handling MCMC diagnostic visualisations,
370 (Gabry et al. 2019) and *ggtree* (Yu et al. 2017) is used for visualising phylogenies. The package
371 *posterior* (Vehtari et al. 2021) is used for computing MCMC diagnostics.

372 Illustration on simulated phylogenies

373 All simulations follow the same protocol. First a dated phylogeny is simulated from a genealogical
374 model, conditional on appropriate parameter values and tip sample times. Then the expected number
375 of substitutions for each branch is sampled using the ARC model (Didelot et al. 2021). Then Seq-Gen
376 (Rambaut and Grass 1997) is used to generate sequences of a given length under the HKY model
377 (Hasegawa et al. 1985). For all simulation experiments the length was set to 10000bp, except if
378 otherwise mentioned. Finally, a ML phylogenetic tree is reconstructed using IQ-TREE (Minh et al.
379 2020). This, along with the tip dates, serves as a starting point for the analysis. For all simulation
380 benchmarks, the root position was assumed to be unknown and to be inferred.

381 A first dated phylogeny was simulated under the Beta-coalescent (Equation 5) with parameters
382 $\{\nu = 1/12, \alpha^* = 1/2\}$ as shown in Figure 2A. The ARC clock model with parameters $\{\mu = 1, \omega = 1\}$

383 was applied, and the input ML phylogeny is shown in Figure S1. Inference was performed under the
384 Beta-coalescent model using four chains, sampling every 2000 iterations for a total of 1000 samples
385 retained per chain. Assessing mixing and qualities of estimates is challenging in this setting as the
386 topology changes due to the uncertain branching on polytomies. Furthermore standard metrics like
387 the Robinson-Foulds distance (Robinson and Foulds 1981) are inappropriate for dated trees and not
388 applicable to multiple merger trees. To circumvent this we compute effective sample sizes (ESS) and
389 \hat{r} estimates for the genealogical parameters as well as the clock parameters along with the tree height
390 and the two following summaries: the number of multiple mergers and the maximum merger size. The
391 \hat{r} and ESS estimates are computed using technique in (Vehtari et al. 2021) as implemented in the
392 R package **posterior**. This confirmed that the MCMC had converged and mixed as expected. The
393 MCMC traces are shown in Figure S2 and the inferred parameters in Figure S3, with all inferred ranges
394 covering the correct values. Nine posterior samples of the dated phylogeny are shown in Figure S4. To
395 summarise the full posterior sample of dated phylogenies, we use a modified version of the DensiTree
396 representation (Bouckaert 2010) as shown in Figure 2B. Comparison of the simulated (Figure 2A) and
397 inferred (Figure 2B) phylogenies demonstrate the accuracy of the inference, including the identification
398 of which nodes are likely to be multiple merger events.

399 Next a dated phylogeny was simulated under the extended Beta-coalescent (Equation 7) with
400 parameters $\{\nu = 1/12, \alpha^* = 1/5, \phi = 2/5\}$ (Figure 3A). The same analysis as above was performed,
401 except that the extended Beta-coalescent model was used for inference. The ML tree is shown in Figure
402 S5, the MCMC traces in Figure S6, the parameters in Figure S7, nine posterior sampled phylogenies in
403 Figure S8 and the posterior phylogeny summary in Figure 3B. Once again we find that the parameters
404 and phylogeny are inferred satisfactorily. There were only a few multiple merger events in the simulated
405 tree, most of which behaved as a Kingman's coalescent tree. This represents a good illustration of
406 what can be achieved with the extended Beta-coalescent, and would be very unlikely to happen under
407 the Beta-coalescent model.

408 Benchmark under the Beta-coalescent

409 To benchmark the performance of inference under the Beta-coalescent, we considered three ARC clocks
410 with the following parameters: Clock 1 has $\mu = 1.5$ and $\omega = 0.5$, Clock 2 has $\mu = 3$ and $\omega = 1.0$,




Figure 2: Simulation and inference under the Beta-coalescent model. (A) The simulated phylogeny. (B) A qualitative summary of the posterior, showing locations possible multiple mergers and uncertainty in polytomy topology. Clades appearing in over 50% of posterior samples are indicated with black dots fixed at median height, and grey bars overlayed indicated the 95% posterior credible interval for the height of these nodes.




Figure 3: Simulation and inference under the extended Beta-coalescent. (A) The simulated phylogeny. (B) A qualitative summary of the posterior, showing locations possible multiple mergers and uncertainty in polytomy topology. Clades appearing in over 50% of posterior samples are indicated with black dots fixed at median height, and grey bars overlayed indicated the 95% posterior credible interval for the height of these nodes.

411 Clock 3 has $\mu = 6$ and $\omega = 2$. Figure S9 shows the distributions of the number of substitutions per
412 site under each of these clock models. For each clock model, Figure 4 shows the results of parameter
413 inference under the Beta-coalescent for 150 datasets generated with $\nu = 0.1$ and values of α^* increasing
414 between 0 and 1. In every case the parameters are correctly inferred, except for α^* which is slightly
415 overestimated when the correct value was lower than 0.1. Low values of α^* lead to trees with a high
416 probability of large multiple merger events. To avoid this biologically implausible scenario we used a
417 Beta(3,1) prior for inference, which only has cumulative probability 0.001 for $\alpha^* \in [0, 0.1]$. The use of
418 this prior explains the slight overestimation of α^* when the correct value was low.

419 Figure S10 shows the number of nodes in the inferred phylogenies minus the true number of nodes in
420 the simulated phylogenies. Positive values indicate an excess of nodes and therefore an underestimation
421 of the number of multiple merger events. Negative values indicate a lack of nodes and therefore an
422 underestimation of the number of multiple merger events. Most intervals cover the correct value of
423 zero. There is a slight tendency to overestimate the number of nodes overall, which is again driven by
424 our use of a conservative prior for α^* .

425 Benchmark under the extended Beta-coalescent

426 We performed a similar benchmark for the inference under the extended Beta-coalescent. Four
427 scenarios were considered with $\phi \in \{0, 0.5, 0.75, 1\}$. Each scenario consists of 48 genealogies. Within a
428 scenario the values of α^* were linearly varied from $\alpha^* = 0.01$ to $\alpha^* = 0.75$. The upper limit was chosen
429 to 0.75 since as $\alpha^* \rightarrow 1$ the process behaves like Kingman's coalescent irrespective of the value of ϕ ,
430 rendering the scenario meaningless. As for the previous benchmark the genealogies were then used to
431 generate three datasets each with different clock parameters. Figure S11 shows the distributions of
432 the number of substitutions per site under each of the three clock models.

433 The inferred values of the parameters ϕ and α^* are shown in Figures S12 and S13, respectively.
434 The scenario $\phi = 0$ corresponds to the Kingman's coalescent. In this scenario the parameter ϕ was
435 consistently estimated to be low as expected, and α^* could not be estimated (i.e. the posterior was
436 approximately equal to the prior) since in this scenario this parameter does not play a role. In the
437 scenarios where the Kingman's coalescent and Beta-coalescent were mixed with $\phi = 0.5$ and $\phi = 0.75$




Figure 4: Simulation benchmark under the Beta-coalescent. Posterior summaries for analysis of limitations of the likelihood approximation. Red lines indicate ground truth. Vertical bars represent 95% posterior credible intervals, and points represent the median. Red crosses indicate insufficient mixing in the corresponding.

438 the parameter ϕ was usually underestimated to the point that the inferred values of α^* did not follow
439 the correct values. However, in the scenario $\phi = 1$, which corresponds to the pure Beta-coalescent, it
440 was possible to infer the values of ϕ and α^* as long as α^* was not too high. When α^* is high the Beta-
441 coalescent component of the extended Beta-coalescent prior behaves like the Kingman's coalescent
442 component, so that the mixing proportion ϕ does not have much effect on the data.

443 Thus the extended Beta-coalescent suffers from identifiability issues on the parameters ϕ and α^* in
444 the part of the parameter space where the model reduces to the Kingman's coalescent, namely when ϕ
445 is low and/or α^* is high. This does not affect the estimates of the remaining parameters though. The
446 parameters ν , μ and ω are shown in Figures S14, S15 and S16, respectively, and are all well estimated.
447 Figures S17 and S18 show that the time to the most recent common ancestor and number of nodes in
448 the tree are also estimated around their correct values. Note that in the scenario with $\phi = 0$ the correct
449 tree is completely binary and so the number of nodes can only be underestimated. Finally, Figure S19
450 shows the estimated probabilities that a tree sampled from the posterior contains a multiple merger,
451 which increases as expected as ϕ increases.

452 **High mutation rate limitation**

453 A difficulty arises when the mutation rate per site is too high. In this case the probability of reversal
454 or homoplasious mutation increases, such that the maximum likelihood estimated branch lengths of
455 the input phylogeny become unlikely to be exactly zero even when a multiple merger event occurred.
456 This is related to the branch saturation observed in dating methods that use a maximum likelihood
457 tree as input, such as LSD (To et al. 2016) or BactDating (Didelot et al. 2018). In order to gain an
458 understanding of when this phenomenon becomes problematic, we repeated the simulation benchmark
459 under the Beta-coalescent but with a genome length of 1000bp (i.e. 10 times less than previously) and
460 mutation rates doubled for each clock model, i.e. $\mu \in \{3, 6, 12\}$. Figure 5 shows the results for this
461 analysis. There is a clear and consistent overestimation of α^* when the correct value of this parameter
462 was low. This corresponds to bias against configurations of the Beta-coalescent that produce larger
463 multiple mergers. This was accompanied by underestimation of the process timescale parameter ν .
464 This bias worsens as the mutation rate increases and thus the expected number of substitutions per site
465 increases. This result shows that the approach presented here is not appropriate when the number of

466 substitutions per site is too high, specifically in the order of $\mu \approx 10$ per year. Since the simulated trees
467 had sums of branch lengths in excess of 100 years (eg Figure 4A), and the number of sites simulated
468 was 1000bp, this corresponds to an expected number of substitutions greater than one for each site,
469 which would not be an issue in practice for the applications envisaged here.

470 Case Study: Spread of *Vibrio cholerae* in Argentina

471 A recent study compared genome sequences of *Vibrio cholerae*, the causative agent of cholera, sampled
472 from Argentina and neighboring countries between 1992 and 2000 in order to characterise its population
473 structure (Dorman et al. 2020). We selected from the previously published phylogeny the genomes that
474 had been isolated in Argentina and for which the isolation date was known, resulting in a phylogeny
475 containing 411 leaves as shown in Figure S20. We applied inference under the extended Beta-coalescent
476 model, which produced the traces shown in Figure S21 and the parameter estimates shown in Figure
477 S22. The rooting of the tree was fixed using an outgroup. Nine samples from the posterior phylogeny
478 are shown in Figure S23. A phylogenetic posterior sample is summarised as a DensiTree in Figure 6A.
479 This contains several large well supported multiple merger events, consistent with the high estimate
480 of ϕ and low estimate of α^* .

481 For comparison purposes, we also performed inference under the pure Kingman's coalescent model,
482 and nine samples from the posterior phylogeny are shown in Figure S24. As we can see in Figure
483 6B the clock rate estimated for the *Vibrio cholerae* genealogy under Kingman's coalescent is much
484 higher than the one under the extended Beta-coalescent. Furthermore the relaxation parameter ω is
485 higher when using the Kingman's coalescent, indicating that evolution is less clock-like. The estimated
486 mutation rate under the extended Beta-coalescent falls into a posterior 95% credible interval of [1.90
487 - 2.84] mutations per genome per year. This is in good agreement with previous estimates of the
488 *V. cholerae* clock rate based on sparsely sampled worldwide collections of genomes (Motreja et al.
489 2011; Didelot et al. 2015). In contrast the substitution rate estimated under Kingman's coalescent
490 is higher with credible interval [4.98 - 6.64] mutations per genome per year, which is inconsistent
491 with previous estimates. Consequently, the time to the most recent common ancestor for the whole
492 Argentinian dataset is underestimated when using Kingman's coalescent as opposed to the extended
493 Beta-coalescent (Figure 6B).




Figure 5: Posterior summaries for analysis of limitations of the likelihood approximation. Red lines indicate ground truth. Vertical bars represent 95% posterior credible intervals, and points represent the median. Red crosses indicate insufficient mixing in the corresponding run.




Figure 6: Analysis of the *Vibrio cholerae* dataset. (A) Qualitative summaries of the posterior inferred under the extended Beta-coalescent and Kingman's coalescent, showing locations possible multiple mergers and uncertainty in polytomy topology. Note that the tip ordering is not identical between the two summaries. Clades appearing in over 50% of posterior samples are indicated with black dots fixed at median height, and grey bars overlayed indicated the 95% posterior credible interval for the height of these nodes. (B) A comparison of estimated parameters under the extended Beta-coalescent and Kingman's coalescent models.

⁴⁹⁴ **Case Study: *Mycobacterium tuberculosis* outbreak phylogenies**

⁴⁹⁵ The importance of multiple merger genealogies to study tuberculosis outbreaks has been recently
⁴⁹⁶ demonstrated (Menardo et al. 2021) using data from eleven previously published outbreaks.
⁴⁹⁷ We selected three of these for reanalysis, labelled *Bainomugisa2018* (Bainomugisa et al. 2018),
⁴⁹⁸ *Eldholm2015* (Eldholm et al. 2015) and *Lee2015* (Lee et al. 2015). These three datasets were selected
⁴⁹⁹ because they had more than 90% probability of the model selected being a Beta-coalescent in the
⁵⁰⁰ previous analysis (Menardo et al. 2021) and their phylogenies were not being excessively large. Analysis
⁵⁰¹ was performed for each of the three phylogenies under three models: the extended Beta-coalescent, the
⁵⁰² Beta-coalescent and the Kingman's coalescent. The three input trees are shown in Figure S25. The
⁵⁰³ rooting of the trees was fixed using outgroup rooting.

⁵⁰⁴ As can be seen in Figure 7, analysis under Kingman's coalescent leads to a higher clock relaxation
⁵⁰⁵ parameter ω value, suggesting that this model is less appropriate. This is also shown by the fact that α^*
⁵⁰⁶ was always inferred much smaller than one. The effect is most pronounced for the *Eldhom2015* dataset.
⁵⁰⁷ In the extended Beta-coalescent the parameter ϕ was estimated to be very close to one for this dataset,
⁵⁰⁸ in which case it becomes approximately equivalent to the Beta-coalescent. A qualitative summary of
⁵⁰⁹ the genealogies inferred for this dataset can be seen in Figure 8. Nine realisations of the dated
⁵¹⁰ genealogies are shown for the Beta-coalescent, extended Beta-coalescent and Kingman's coalescent
⁵¹¹ in Figures S26, S27 and S28, respectively. As expected, the results under the Beta-coalescent and
⁵¹² extended Beta-coalescent are very similar, including evidence for several large multiple merger events.
⁵¹³ These results are in good agreement with the previous analysis of these three datasets (Menardo et al.
⁵¹⁴ 2021), and highlight the importance of considering multiple mergers when analysing phylogenetic data
⁵¹⁵ from tuberculosis outbreaks.

⁵¹⁶ **DISCUSSION**

⁵¹⁷ We have presented an approach to reconstructing dated phylogenies with multiple mergers under
⁵¹⁸ Lambda-coalescent models. To our knowledge this is the first such approach that scales to real world
⁵¹⁹ phylogeny sizes, explicitly reconstructs the underlying multiple merger genealogy, and does not rely




Figure 7: A comparison of parameter marginals estimated under different models for each dataset.




Figure 8: A comparison of three qualitative summaries of posteriors inferred for the the underlying genealogies for the *Eldhom2015* dataset when timed under different coalescent priors. Note that the tip ordering is not exactly the same between the different summaries.

520 on likelihood-free approaches such as Approximate Bayesian Computation. Our focus has been on
521 the implementation of the methodology presented, extensive benchmarks, and applications to real
522 world examples from pathogen phylogenetics. On the other hand we have not yet addressed how
523 this reconstruction can be used to further study of pathogen population genetics. Some implications
524 are relatively straightforward. For example our method could be used as a starting point to do
525 multifurcating skyline plot analysis (Ho and Shapiro 2011). We also envisage that it will lead to other
526 forms of studies becoming possible, for example using genomic data to learn about superspreading,
527 outbreaks, and impacts of selection.

528 Other Lambda-coalescents models that the ones we used have been studied or derived. While this study
529 was only focused on the Beta-coalescent and its mixture with a Kingman's coalescent, it is worthwhile
530 to mention some alternatives. In principle there is no reason why the presented methodology would not
531 be applicable to them. A first alternative class of Lambda-coalescent is the extinction-recolonisation
532 Dirac coalescent of (Eldon and Wakeley 2006). In the corresponding forwards in time model all large
533 reproductive events replace a fixed proportion of the population. Each time a potential multiple merger
534 event occurs, a biased coin with probability of heads $p \in (0, 1]$ gets flipped for every extant lineage
535 in the process. Lineages whose coin shows heads all merge into a common ancestor. This model is
536 noteworthy primarily because it was amongst the first to be derived, is simple, and well studied, but
537 it may not be the most biologically plausible.

538 Another alternative class of Lambda-coalescent is the Durrett-Schweinsberg (DS) coalescent (Durrett
539 and Schweinsberg 2005). This model describes populations undergoing successive hard selective sweeps
540 throughout the genome, and in particular, the hitchhiking effect of those sweeps on a fixed, neutral
541 site. The sweeps are modelled as points in a Poisson process of fixed rate. During a sweep, some
542 ancestral lineages carrying the neutral site of interest can escape the sweep by recombining, while
543 those lineages which don't recombine will merge to a common ancestor which initiated the sweep.
544 Between individual sweeps the population follows neutral Moran type dynamics. This class of model
545 has recently been found to describe the genetic diversity in cod populations (Árnason et al. 2023). In
546 general, the measure Λ associated with this class of coalescents takes the form of $\Lambda = \delta_0 + \Lambda_0$ where
547 δ_0 is an atom at 0 responsible for Kingman-like mergers between sweeps and Λ_0 is a finite measure
548 on $[0, 1]$ without an atom at 0, which drives multiple mergers due to selective sweeps. This model is
549 not directly applicable to pathogen populations, primarily due to traditional recombination being less

550 frequent in bacterial and viral pathogens. An adaptation of this model to bacterial pathogens may be
551 possible but is outside of the scope of this study.

552 Future studies are needed to investigate what type of Lambda-coalescents best describe pathogen
553 dynamics. The approach presented here is an approximate, albeit explicit approach to Bayesian
554 inference of multiple merger genealogies. It has limitations: for instance, it is not appropriate for
555 studying genealogies spanning geological timescales, as was demonstrated by worsening bias as the
556 number of substitutions per site becomes high (Figure 5). The possibility of extending the inference
557 under Lambda-coalescents to the fully Bayesian setting, incorporating uncertainty about the phylogeny
558 and relying on the phylogenetic likelihood, remains an open problem. Extending some aspects of
559 the parametrisation and construction of multiple merger genealogies presented in this work to the
560 aforementioned setting is straightforward. However, we anticipate that the parametrisation presented
561 here may not be computationally efficient when used in such a setting.

562 Finally there is the question of extending the approach presented here to, for example, joint estimation
563 of varying effective population size (Ho and Shapiro 2011). This is relevant both for the sake of the past
564 effective population size being an interesting quantity, as well as being relevant from the perspective of
565 statistical robustness. As Menardo et al. (2021) noted, population expansion can be misidentified as
566 a Lambda-coalescent if population growth is not properly accounted for. In order to do this there are
567 separate questions that have to be answered. Firstly, the impact of varying effective population size on
568 the genealogy depends on the forwards-in-time model. Therefore it is first necessary to decide which
569 Lambda-coalescent is the correct one to use for a given scenario. Secondly, adding a non-parametric
570 model for the effective population size will increase the complexity of the inference problem. More
571 efficient MCMC schemes, or other inference tools, would therefore need to be investigated. This might
572 include non-reversible samplers with boundary conditions such as (Bierkens et al. 2023) or Hamiltonian
573 Monte-Carlo methods (Dinh et al. 2017). Note that some phenomena might be indistinguishable
574 from multiple mergers, such as population structure. For example a sufficiently fast expansion of a
575 subpopulation that shares identity by descent (?) will likely lead to multiple mergers in the genealogy.

576 ACKNOWLEDGEMENTS

577 DH was supported by the UK Engineering and Physical Sciences Research Council (EPSRC) grant
578 EP/S022244/1 for the EPSRC Centre for Doctoral Training in Mathematics for Real-World Systems
579 II. JK was supported by EPSRC research grant EP/V049208/1. XD acknowledges funding from the
580 National Institute for Health and Care Research (NIHR) Health Protection Research Unit in Genomics
581 and Enabling Data (NIHR200892).

582 References

583 Bainomugisa, A., E. Lavu, S. Hiashiri, S. Majumdar, A. Honjepari, R. Moke, P. Dakulala, G. A.
584 Hill-Cawthorne, S. Pandey, B. J. Marais, et al. 2018. Multi-clonal evolution of multi-drug-
585 resistant/extensively drug-resistant *Mycobacterium tuberculosis* in a high-prevalence setting of
586 Papua New Guinea for over three decades. *Microbial genomics* 4.

587 Berestycki, N. 2009. Recent progress in coalescent theory.

588 Bertoin, J. and J.-F. Le Gall. 2000. The Bolthausen–Sznitman coalescent and the genealogy of
589 continuous-state branching processes. *Probability Theory and Related Fields* 117:249–266.

590 Biek, R., O. G. Pybus, J. O. Lloyd-Smith, and X. Didelot. 2015. Measurably Evolving Pathogens in
591 the Genomic Era. *Trends in Ecology & Evolution* 30:306–313 publisher: Elsevier Ltd.

592 Bierkens, J., S. Grazzi, F. v. d. Meulen, and M. Schauer. 2023. Sticky PDMP samplers for sparse and
593 local inference problems. *Statistics and Computing* 33:8.

594 Billera, L. J., S. P. Holmes, and K. Vogtmann. 2001. Geometry of the Space of Phylogenetic Trees.
595 *Advances in Applied Mathematics* 27:733–767.

596 Birkner, M., J. Blath, and B. Eldon. 2013. Statistical Properties of the Site-Frequency Spectrum
597 Associated with Λ -Coalescents. *Genetics* 195:1037–1053.

598 Bouckaert, R. R. 2010. DensiTree: making sense of sets of phylogenetic trees. *Bioinformatics* 26:1372–
599 1373.

600 Bromham, L. and D. Penny. 2003. The modern molecular clock. *Nature Reviews Genetics* 4:216–224.

601 Charlesworth, B. 2009. Fundamental concepts in genetics: Effective population size and patterns of
602 molecular evolution and variation. *Nature Reviews Genetics* 10:195–205.

603 Cvijović, I., B. H. Good, and M. M. Desai. 2018. The Effect of Strong Purifying Selection on Genetic
604 Diversity. *Genetics* 209:1235–1278.

605 Desai, M. M., A. M. Walczak, and D. S. Fisher. 2013. Genetic Diversity and the Structure of
606 Genealogies in Rapidly Adapting Populations. *Genetics* 193:565–585.

607 Didelot, X., N. J. Croucher, S. D. Bentley, S. R. Harris, and D. J. Wilson. 2018. Bayesian inference of
608 ancestral dates on bacterial phylogenetic trees. *Nucleic Acids Res.* 46:e134–e134.

609 Didelot, X., B. Pang, Z. Zhou, A. McCann, P. Ni, D. Li, M. Achtman, and B. Kan. 2015. The Role of
610 China in the Global Spread of the Current Cholera Pandemic. *PLoS Genetics* 11:e1005072.

611 Didelot, X. and J. Parkhill. 2022. A Scalable Analytical Approach from Bacterial Genomes to
612 Epidemiology. *Philosophical Transactions of the Royal Society B: Biological Sciences* 377:20210246
613 publisher: Cold Spring Harbor Laboratory.

614 Didelot, X., I. Siveroni, and E. M. Volz. 2021. Additive uncorrelated relaxed clock models for the
615 dating of genomic epidemiology phylogenies. *Mol. Biol. Evol.* 38:307–317.

616 Dinh, V., A. Bilge, C. Zhang, and F. A. M. Iv. 2017. Probabilistic Path Hamiltonian Monte Carlo.
617 Pages 1009–1018 in *Proceedings of the 34th International Conference on Machine Learning* PMLR
618 iSSN: 2640-3498.

619 Donnelly, P. and T. G. Kurtz. 1999. Particle Representations for Measure-Valued Population Models.
620 The *Annals of Probability* 27:166–205 publisher: Institute of Mathematical Statistics.

621 Dorman, M. J., D. Domman, T. Poklepovich, C. Tolley, G. Zolezzi, L. Kane, M. R. Viñas,
622 M. Panagópulo, M. Moroni, N. Binsztein, M. I. Caffer, S. Clare, G. Dougan, G. P. C.
623 Salmond, J. Parkhill, J. Campos, and N. R. Thomson. 2020. Genomics of the Argentinian cholera
624 epidemic elucidate the contrasting dynamics of epidemic and endemic *Vibrio cholerae*. *Nature
625 Communications* 11:4918.

626 Drummond, A. J., G. K. Nicholls, A. G. Rodrigo, and W. Solomon. 2002. Estimating mutation
627 parameters, population history and genealogy simultaneously from temporally spaced sequence data.
628 *Genetics* 161:1307–1320.

629 Drummond, A. J., O. G. Pybus, A. Rambaut, R. Forsberg, and A. G. Rodrigo. 2003. Measurably
630 Evolving Populations. *Trends in Ecology and Evolution* 18:481–488.

631 Durrett, R. and J. Schweinsberg. 2005. A coalescent model for the effect of advantageous mutations
632 on the genealogy of a population. *Stochastic Processes and their Applications* 115:1628–1657.

633 Eldholm, V., J. Monteserin, A. Rieux, B. Lopez, B. Sobkowiak, V. Ritacco, and F. Balloux. 2015. Four
634 decades of transmission of a multidrug-resistant *Mycobacterium tuberculosis* outbreak strain. *Nat.*
635 *Commun.* 6:7119 publisher: Nature Publishing Group.

636 Eldon, B. and W. Stephan. 2023. Sweepstakes reproduction facilitates rapid adaptation in highly
637 fecund populations. *Molecular Ecology* n/a *preprint*:
638 <https://onlinelibrary.wiley.com/doi/pdf/10.1111/mec.16903>.

639 Eldon, B. and J. Wakeley. 2006. Coalescent Processes When the Distribution of Offspring Number
640 Among Individuals Is Highly Skewed. *Genetics* 172:2621–2633.

641 Gabry, J., D. Simpson, A. Vehtari, M. Betancourt, and A. Gelman. 2019. Visualization in Bayesian
642 workflow. *Journal of the Royal Statistical Society Series A: Statistics in Society* 182:389–402.

643 George, E. I. and R. E. McCulloch. 1993. Variable Selection via Gibbs Sampling. *Journal*
644 *of the American Statistical Association* 88:881–889 publisher: Taylor & Francis *preprint*:
645 <https://www.tandfonline.com/doi/pdf/10.1080/01621459.1993.10476353>.

646 Green, P. J. 1995. Reversible Jump Markov Chain Monte Carlo Computation and Bayesian Model
647 Determination. *Biometrika* 82:711–732.

648 Guindon, S., J.-F. Dufayard, V. Lefort, M. Anisimova, W. Hordijk, and O. Gascuel. 2010. New
649 Algorithms and Methods to Estimate Maximum-Likelihood Phylogenies: Assessing the Performance
650 of PhyML 3.0. *Systematic biology* 59:307–21.

651 Gómez-Carballa, A., J. Pardo-Seco, X. Bello, F. Martínón-Torres, and A. Salas. 2021. Superspreading
652 in the emergence of COVID-19 variants. *Trends in Genetics* 37:1069–1080.

653 Hasegawa, M., H. Kishino, and T. Yano. 1985. Dating of the Human-Ape Splitting by a Molecular
654 Clock of Mitochondrial DNA. *Journal of molecular evolution* 22:160–174.

655 Helekal, D., A. Ledda, E. Volz, D. Wyllie, and X. Didelot. 2021. Bayesian Inference of Clonal
656 Expansions in a Dated Phylogeny. *Systematic Biology* Page syab095.

657 Ho, S. Y. W. and B. Shapiro. 2011. Skyline-plot methods for estimating demographic history from
658 nucleotide sequences. *Mol. Ecol. Resour.* 11:423–434 ISBN: 1755-0998 (Electronic)\backslashbackslashr1755-
659 098X (Linking).

660 Hoscheit, P. and O. G. Pybus. 2019. The multifurcating skyline plot. *Virus Evol.* 5:1–10 ISBN:
661 0000000280013.

662 Ji, X., A. A. Fisher, S. Su, J. L. Thorne, B. Potter, P. Lemey, G. Baele, and M. A. Suchard. 2021.
663 Scalable Bayesian divergence time estimation with ratio transformations. *ArXiv:2110.13298* [q-bio,
664 stat].

665 Kingman, J. F. C. 1982. The coalescent. *Stochastic Processes and their Applications* 13:235–248.

666 Kuhn, T. S., A. O. Mooers, and G. H. Thomas. 2011.
667 A simple polytomy resolver for dated phylogenies. *Methods in Ecology and Evolution* 2:427–436
668 eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.2041-210X.2011.00103.x>.

669 Kukla, J. and M. Möhle. 2018. On the block counting process and the fixation line of the Bolthausen–
670 Sznitman coalescent. *Stochastic Processes and their Applications* 128:939–962.

671 Lee, R. S., N. Radomski, J.-F. Proulx, I. Levade, B. J. Shapiro, F. McIntosh, H. Soualhine, D. Menzies,
672 and M. A. Behr. 2015. Population genomics of *Mycobacterium tuberculosis* in the Inuit. *Proceedings
673 of the National Academy of Sciences* 112:13609–13614.

674 Lemieux, J. E., K. J. Siddle, B. M. Shaw, C. Loreth, S. F. Schaffner, A. Gladden-Young, G. Adams,
675 T. Fink, C. H. Tomkins-Tinch, L. A. Krasilnikova, K. C. DeRuff, M. Rudy, M. R. Bauer, K. A.
676 Lagerborg, E. Normandin, S. B. Chapman, S. K. Reilly, M. N. Anahtar, A. E. Lin, A. Carter,
677 C. Myhrvold, M. E. Kemball, S. Chaluvadi, C. Cusick, K. Flowers, A. Neumann, F. Cerrato,
678 M. Farhat, D. Slater, J. B. Harris, J. Branda, D. Hooper, J. M. Gaeta, T. P. Baggett, J. O’Connell,
679 A. Gnirke, T. D. Lieberman, A. Philippakis, M. Burns, C. M. Brown, J. Luban, E. T. Ryan, S. E.
680 Turbett, R. C. LaRocque, W. P. Hanage, G. R. Gallagher, L. C. Madoff, S. Smole, V. M. Pierce,
681 E. Rosenberg, P. C. Sabeti, D. J. Park, and B. L. MacInnis. 2020. Phylogenetic analysis of SARS-
682 CoV-2 in the Boston area highlights the role of recurrent importation and superspreading events.
683 preprint *Epidemiology*.

684 Lewis, P. O., M. T. Holder, and K. E. Holsinger. 2005. Polytomies and Bayesian Phylogenetic Inference.
685 *Systematic Biology* 54:241–253.

686 Lin, G. N., C. Zhang, and D. Xu. 2011. Polytomy identification in microbial phylogenetic
687 reconstruction. *BMC Systems Biology* 5:S2.

688 Matuszewski, S., M. E. Hildebrandt, G. Achaz, and J. D. Jensen. 2018. Coalescent Processes with
689 Skewed Offspring Distributions and Nonequilibrium Demography. *Genetics* 208:323–338.

690 Menardo, F., S. Gagneux, and F. Freund. 2021. Multiple Merger Genealogies in Outbreaks of
691 *Mycobacterium tuberculosis*. *Molecular Biology and Evolution* 38:290–306.

692 Metropolis, N., A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller. 1953. Equation of
693 State Calculations by Fast Computing Machines. *Journal of Chemical Physics* 21:1087–1092 aDS
694 Bibcode: 1953JChPh..21.1087M.

695 Minh, B. Q., H. A. Schmidt, O. Chernomor, D. Schrempf, M. D. Woodhams, A. von Haeseler, and
696 R. Lanfear. 2020. IQ-TREE 2: New Models and Efficient Methods for Phylogenetic Inference in the
697 Genomic Era. *Molecular Biology and Evolution* 37:1530–1534.

698 Mutreja, A., D. W. Kim, N. R. Thomson, T. R. Connor, J. H. Lee, S. Kariuki, N. J. Croucher, S. Y.
699 Choi, S. R. Harris, M. Lebans, S. K. Niyogi, E. J. Kim, T. Ramamurthy, J. Chun, J. L. N. Wood,
700 J. D. Clemens, C. Czerkinsky, G. B. Nair, J. Holmgren, J. Parkhill, and G. Dougan. 2011. Evidence
701 for several waves of global transmission in the seventh cholera pandemic. *Nature* 477:462–465.

702 Neher, R. A. and O. Hallatschek. 2013. Genealogies of rapidly adapting populations. *Proceedings of
703 the National Academy of Sciences* 110:437–442 publisher: *Proceedings of the National Academy of
704 Sciences*.

705 Paradis, E. and K. Schliep. 2019. ape 5.0: an environment for modern phylogenetics and evolutionary
706 analyses in R. *Bioinformatics* 35:526–528 publisher: Oxford University Press.

707 Pitman, J. 1999. Coalescents With Multiple Collisions. *The Annals of Probability* 27:1870–1902
708 publisher: Institute of Mathematical Statistics.

709 Rambaut, A. and N. C. Grass. 1997. Seq-Gen: an application for the Monte Carlo simulation of DNA
710 sequence evolution along phylogenetic trees. *Bioinformatics* 13:235–238.

711 Robinson, D. F. and L. R. Foulds. 1981. Comparison of phylogenetic trees. *Mathematical Biosciences*
712 53:131–147.

713 Sagitov, S. 1999. The general coalescent with asynchronous mergers of ancestral lines. *Journal of*
714 *Applied Probability* 36:1116–1125 publisher: Cambridge University Press.

715 Sagulenko, P., V. Puller, and R. A. Neher. 2018. TreeTime: Maximum likelihood phylodynamic
716 analysis. *Virus Evol.* 4:vex042.

717 Schweinsberg, J. 2003. Coalescent processes obtained from supercritical Galton–Watson processes.
718 *Stochastic Processes and their Applications* 106:107–139.

719 Schweinsberg, J. 2017. Rigorous results for a population model with selection II: genealogy of the
720 population. *Electronic Journal of Probability* 22:1–54 publisher: Institute of Mathematical Statistics
721 and Bernoulli Society.

722 Stamatakis, A. 2014. RAxML version 8: a tool for phylogenetic analysis and post-analysis of large
723 phylogenies. *Bioinformatics* 30:1312–1313.

724 Tellier, A. and C. Lemaire. 2014. Coalescence 2.0: a multiple branching of recent theoretical
725 developments and their applications. *Molecular ecology* 23:2637–2652.

726 To, T.-H., M. Jung, S. Lycett, and O. Gascuel. 2016. Fast Dating Using Least-Squares Criteria and
727 Algorithms. *Systematic Biology* 65:82–97.

728 Vehtari, A., A. Gelman, D. Simpson, B. Carpenter, and P. C. Burkner. 2021. Rank-Normalization,
729 Folding, and Localization: An Improved R hat for Assessing Convergence of MCMC. *Bayesian Anal.*
730 16:667–718 _eprint: arXiv:1903.08008v5.

731 Volz, E. M. and S. D. W. Frost. 2017. Scalable Relaxed Clock Phylogenetic Dating. *Virus Evolution*
732 3:vex025.

733 Yang, Z. and B. Rannala. 1997. Bayesian Phylogenetic Inference using DNA Sequences: A Markov
734 Chain Monte Carlo Method. *Molecular Biology and Evolution* 14:717–724 publisher: SMBE.

735 Yu, G., D. K. Smith, H. Zhu, Y. Guan, and T. T. Y. Lam. 2017. Ggtree: an R Package for Visualization
736 and Annotation of Phylogenetic Trees With Their Covariates and Other Associated Data. *Methods*
737 *Ecol. Evol.* 8:28–36.

738 Árnason, E., J. Koskela, K. Halldórsdóttir, and B. Eldon. 2023. Sweepstakes reproductive success
739 via pervasive and recurrent selective sweeps. *eLife* 12:e80781 publisher: eLife Sciences Publications,
740 Ltd.