The site-frequency spectrum associated with Ξ-coalescents

Jochen Blath[1], Mathias Christensen Cronjäger[2], Bjarki Eldon[1], Matthias Hammer[1]

August 28, 2015

Author affiliations:

1. TU Berlin, Institut für Mathematik
10623 Berlin, Germany

2. University of Oxford
OX1 2DL Oxford, UK

Running title: Frequency spectrum and Xi-coalescents

Keywords: site-frequency spectrum, Xi-coalescents, diploidy, Atlantic cod, simultaneous mergers

Corresponding author:
Bjarki Eldon
TU Berlin, Institut für Mathematik
Straße des 17. Juni 136
10623 Berlin, Germany
Email: eldon@math.tu-berlin.de
Phone: +49 30 314 25762
Fax: +49 30 314 21695

Abstract

We give recursions for the expected site-frequency spectrum associated with Xi-coalescents, that is exchangeable coalescents which admit *simultaneous multiple mergers* of ancestral lineages. Xi-coalescents arise, for example, in association with population models of skewed offspring distributions with diploidy, recurrent advantageous mutations, or strong bottlenecks. In contrast, Lambda-coalescents admit multiple mergers of lineages, but at most one such merger each time. Xi-coalescents, as well as Lambda-coalescents, can predict an excess of singletons, compared to the Kingman coalescent. We compare estimates of coalescent parameters when Xi-coalescent models are applied to data obtained from Lambda-coalescents, and vice versa. In general, Xi-coalescents predict fewer singletons than corresponding Lambda-coalescents, but higher count of mutations of 'size' larger than singletons. We analyse unfolded site-frequency spectra obtained for nuclear loci of the diploid Atlantic cod, and obtain different coalescent parameter estimates than previously obtained with Lambda-coalescents. Our results provide new inference tools, and suggest that for nuclear population genetic data from diploid or polyploid highly fecund populations who may have skewed offspring distributions, one should not apply Lambda-coalescents, but Xi-coalescents.

Introduction

The coalescent approach, ie. the idea of considering the (random) ancestral relations of alleles sampled from natual populations, has provided rich mathematical theory [cf. 4], and very useful inference tools (cf. eg. [14, 39] for reviews). Initiated by the Kingman coalescent [24, 26, 25], coalescent models now include the family of Lambda-(Λ-)coalescents [31, 32, 13], and Xi-(Ξ-)coalescents [35, 29, 33]. Ξ-coalescents admit *simultaneous multiple mergers* of ancestral lineages. Thus, in each merger event, distinct groups of ancestral lineages can merge at the same time, and each group can have *more* than two lineages. Λ-coalescents, in contrast, only allow one group - possibly containing more than two lineages - to merge each time. Thus, due to multiple mergers, the derivation and application of inference tools becomes harder as one moves from the Kingman coalescent to Lambda-coalescent models, and from Λ-coalescents to Ξ-coalescents. Ξ-coalescents can be obtained from diploid population models [6, 30, 11]. They also arise in models of repeated strong bottlenecks [8], and in models of selective sweeps [15, 16].

[20] obtained closed-form expressions for the expected site-frequency spectrum, as well as (co)-variances, when associated with the Kingman coalescent. [7] obtain recursions for expected values and (co)-variances when associated with Λ-coalescents. However, the complexity of the recursions means that (co)-variances, when associated with Λ-coalescents, can only be computed for small sample sizes. The expected values can be applied in distance statistics [21], as well as in an approximate likelihood approach [7, 18].

Multiple merger coalescents can be obtained from population models that admit large offspring numbers and skewed offspring distributions, characteristics associated with many marine populations [1, 19, 5, 9, 10, 22, 23, 34]. Indeed, [2] find much better fit (the sum of the squared distance between observed and expected values) between data on nuclear genes from Atlantic cod and Λ-coalescents than with the Kingman coalescent. Simulation results of [34] suggest that the site-frequency spectrum of (at least some) Ξ-coalescents is multimodal, a pattern observed in data on the nuclear *Ckma* gene in Atlantic cod [2]. Based on this evidence, a way to compute expected values of the site-frequency spectrum associated

with Ξ-coalescents should be a welcome and important addition to the set of inference tools for population genetics.

In this work, we obtain recursions for the expected site-frequency spectrum associated with general Ξ-coalescents, with an approach similar to the one applied by [7]. We compare estimates of coalescent parameters when applied to simulated data obtained under Λ-coalescents, and vice-versa. Since the recursions for the expected values are already fairly complex, and computationally intensive, we expect recursions for the (co)-variances to be even more so. The (co)-variances will therefore not be addressed.

We estimate coalescent parameters associated with 4-fold Xi-coalescents for the unfolded site-frequency spectrum of 3 nuclear loci [2] of the highly fecund Atlantic cod. Our simple method involves minimising the distance between observed and expected values, where the distance is not calibrated by the corresponding variance. Hence, it is not a formal test. Our estimates differ from previous estimates obtained with the use of Lambda-coalescents. The main biological implication of our results are that Xi-coalescents should be applied to nuclear data from highly fecund diploid (or polyploid) populations, and Lambda-coalescents to haploid data such as mitochondrial DNA.

The paper is structured as follows: First we give a precise mathematical description of the various coalescent models. We then state our main result on the expected site frequency spectrum of Ξ-coalescents, Theorem 2. This is followed by a discussion of some specific examples of Ξ-coalescents. Some numerical examples which illustrate the difference in the site-frequency spectrum between Lambda- and Xi-coalescents are then presented, followed by an application to nuclear Atlantic cod data. The proofs are collected in an appendix.

Theory

Coalescent models

We briefly review the basic coalescent models, namely the Kingman-, Λ-, and Ξ-coalescents. They all have in common to be continuous-time Markov chains, taking values in the space of

partitions of the natural numbers $\mathbb{N} := \{1, 2, \dots, \}$, whose restriction to the first n integers can be described as follows: Let $\mathcal{P}_n$ denote the space of partitions of $[n] := \{1, \dots, n\}$. We write π for a generic element of $\mathcal{P}_n$, and $\#\pi$ for the *size* of π, i.e. for the number of blocks $\pi_i \in \pi$. Thus for $\pi \in \mathcal{P}_n$ we have $\pi = \{\pi_1, \dots, \pi_{\#\pi}\}$ with $\#\pi \leq n$. If $\pi, \pi' \in \mathcal{P}_n$ with $\#\pi = m$, we write $\pi' \prec \pi$ if there exist $i, j \in [m]$ with $\pi' = \{\pi_\ell : \ell \in [m], \ell \notin \{i, j\}\} \cup \{\pi_i \cup \pi_j\}$, ie. π' is obtained from π by merging blocks π_i and π_j. If the transition rates of the continuous-time Markov chain $\{\Pi_t^{(K,n)}, t \geq 0\}$ with values in $\mathcal{P}_n$, and starting from state $\{\{1\}, \dots, \{n\}\}$ at time $t = 0$, are given by

$$q_{\pi,\pi'} = \begin{cases} 1 & \text{if } \pi' \prec \pi, \\ -\binom{m}{2} & \text{if } \pi' = \pi,\ \#\pi = m \geq 2, \\ 0 & \text{otherwise;} \end{cases} \tag{1}$$

we refer to $\{\Pi_t^{(K,n)}, t \geq 0\}$ as the Kingman-n-coalescent. The process is stopped at time $\inf\{t > 0 : \Pi_t^{(K,n)} = \{\{1, \dots, n\}\}\}$, ie. when the most recent common ancestor of the n lineages has been reached.

If $\pi, \pi' \in \mathcal{P}_n$ with $\#\pi = m$ and there exist indices $i_1, \dots, i_k \in [m]$ with $\pi' = \{\pi_\ell : \ell \in [m], \ell \notin \{i_1, \dots, i_k\}\} \cup \{\pi_{i_1} \cup \dots \cup \pi_{i_k}\}$, we write $\pi' \prec_{m,k} \pi$ and say that a k-merger has occurred, with $2 \leq k \leq m$. For a finite measure Λ on $[0, 1]$, define

$$\lambda_{m,k} := \int_0^1 x^{k-2}(1-x)^{m-k}\Lambda(dx), \tag{2}$$

$$\lambda_m := \int_0^1 [1 - (1-x)^m - mx(1-x)^{m-1}]x^{-2}\Lambda(dx) \tag{3}$$

if the integral in (3) exists, and

$$\lambda_m := \sum_{k=2}^{m} \binom{m}{k} \int_0^1 x^{k-2}(1-x)^{m-k}\Lambda(dx) \tag{4}$$

otherwise. A $\mathcal{P}_n$-valued continuous-time Markov chain $\{\Pi_t^{(\Lambda,n)}, t \geq 0\}$ with transition rates

$q_{\pi,\pi'}$ from π to π' given by

$$q_{\pi,\pi'} = \begin{cases} \lambda_{m,k} & \text{if } \pi' \prec_{m,k} \pi, \\ \lambda_m & \text{if } \pi' = \pi,\ \#\pi = m, \\ 0 & \text{otherwise;} \end{cases} \tag{5}$$

is referred to as a Λ-n-coalescent. The waiting time in state π is exponential with rate λ_m as in (3, 4).

Now we specify the transition rates for a Ξ-n-coalescent. We write $\pi' \prec_{m,\underline{k}} \pi$ if $\pi, \pi' \in \mathcal{P}_n$ with $\#\pi = m$ and there exist $r \in [\lfloor \frac{m}{2} \rfloor]$ groups of indices $i_1^{(j)}, \ldots, i_{k_j}^{(j)}$, $j = 1, \ldots, r$, such that

$$\pi' = \left\{ \pi_\ell : \ell \in [m], \ell \notin \{i_1^{(1)}, \ldots, i_{k_1}^{(1)}\} \cup \cdots \cup \{i_1^{(r)}, \ldots, i_{k_r}^{(r)}\} \right\} \cup \left\{ \pi_{i_1^{(1)}} \cup \cdots \cup \pi_{i_{k_1}^{(1)}} \right\} \cup \cdots \cup \left\{ \pi_{i_1^{(r)}} \cup \cdots \cup \pi_{i_{k_r}^{(r)}} \right\},$$

by which we denote a transition where blocks with indices $i_1^{(j)}, \ldots, i_{k_j}^{(j)}$ merge into a single block, for $j \in [r]$. Thus, a transition denoted by $\pi' \prec_{m,\underline{k}} \pi$ is a *simultaneous multiple merger*, where $k_j \geq 2$ blocks merge into a single block, and $r \in [\lfloor \frac{m}{2} \rfloor]$ such mergers occur simultaneously. The vector $\underline{k} = (k_1, \ldots, k_r)$ specifies the merger sizes, and we write $|\underline{k}| := \sum_{j=1}^r k_j$.

Let Δ denote the infinite simplex

$$\Delta := \left\{ \boldsymbol{x} = (x_1, x_2, \ldots) : x_1 \geq x_2 \geq \cdots \geq 0, \quad \sum_i x_i \leq 1 \right\}.$$

Let $\boldsymbol{x} \in \Delta$, $m \in \mathbb{N}$, $\underline{k} = (k_1, \ldots, k_r)$ with $k_i \geq 2$ and $|\underline{k}| = k_1 + \cdots + k_r \leq m$ the vector of the r merger sizes, and $s := m - |\underline{k}|$ the number of blocks unaffected by the given merger.

Define the functions $f(\boldsymbol{x}, m, \underline{k})$ and $g(\boldsymbol{x}, m, \underline{k})$, with $\boldsymbol{x} \in \Delta_{\mathbf{0}} := \Delta \setminus \{(0, 0, \ldots)\} = \Delta \setminus \{\mathbf{0}\}$,

$$\begin{aligned} f(\boldsymbol{x}, m, \underline{k}) &:= \frac{\sum_{\ell=0}^{s} \sum_{i_1 \neq \ldots \neq i_{r+\ell}} \binom{s}{\ell} x_{i_1}^{k_1} \cdots x_{i_r}^{k_r} x_{i_{r+1}} \cdots x_{i_{r+\ell}} \left(1 - \sum_j x_j\right)^{s-\ell}}{\sum_j x_j^2}, \\ g(\boldsymbol{x}, m) &:= \frac{1 - \sum_{\ell=0}^{m} \sum_{i_1 \neq \ldots \neq i_\ell} \binom{m}{\ell} x_{i_1} \cdots x_{i_\ell} \left(1 - \sum_j x_j\right)^{m-\ell}}{\sum_j x_j^2}. \end{aligned} \tag{6}$$

Let $\Xi_{\mathbf{0}}$ denote a finite measure on $\Delta_{\mathbf{0}}$, and write $\Xi := \Xi_{\mathbf{0}} + a\delta_{\{\mathbf{0}\}}$. Further, let $\mathcal{N}(m, \underline{k})$ denote the number of ways of arranging m items into r non-empty groups whose sizes are given by $\underline{k}$. With ℓ_j denoting the number of $k_1, \ldots, k_r$ equal to j, one checks that, with $s = m - |\underline{k}|$ [35],

$$\mathcal{N}(m, \underline{k}) = \binom{m}{k_1 \ldots k_r \; s} \frac{1}{\prod_{j=2}^{m} \ell_j!}. \tag{7}$$

Now define [35]

$$\lambda_{m,\underline{k}} := \int_{\Delta} f(\boldsymbol{x}, m, \underline{k}) \, \Xi_{\mathbf{0}}(d\boldsymbol{x}) + a \mathbb{1}_{(r=1, k_1=2)}, \tag{8}$$

$$\lambda_m := \int_{\Delta_{\mathbf{0}}} g(\boldsymbol{x}, m) \, \Xi_{\mathbf{0}}(d\boldsymbol{x}) + a \binom{m}{2} \tag{9}$$

if the integral in (9) exists, and

$$\lambda_m = \sum_{n=1}^{m-1} \sum_{\substack{k_1 \geq \ldots \geq k_r \geq 2 \\ m - |\underline{k}| + r = n}} \mathcal{N}(m, \underline{k}) \lambda_{m,\underline{k}} \tag{10}$$

otherwise. A continuous-time $\mathcal{P}_n$-valued Markov chain with transitions $q_{\pi,\pi'}$ given by [35]

$$q_{\pi,\pi'} = \begin{cases} \lambda_{m,\underline{k}} & \text{if } \pi' \prec_{m,\underline{k}} \pi, \\ -\lambda_m & \text{if } \pi' = \pi, \; \#\pi = m, \\ 0 & \text{otherwise;} \end{cases} \tag{11}$$

with $f(\boldsymbol{x}, m, \underline{k})$ and $g(\boldsymbol{x}, m)$ given by (6), is referred to as a Ξ-n-coalescent, and denoted by

$\{\Pi_t^{(\Xi,n)}, t \geq 0\}$. The waiting time in state π is exponential with rate λ_m as in (9, 10).

The site-frequency spectrum

The site-frequency spectrum is a simple summary statistic of the full DNA sequence data, but contains valuable information about variation among individuals. We assume the infinitely-many sites mutation model [40], in which mutations occur as independent Poisson processes on the branches of a given gene genealogy with rate $\theta/2$ for some constant $\theta > 0$, and no two mutations occur at the same site. The constant θ is determined by the ratio μ/c_N, where μ is the per-generation mutation rate, and c_N is the probability of two distinct individuals (gene copies) sharing a common ancestor in the previous generation. We refer to [18] for a discussion of the relation between mutation and timescales of different coalescent processes.

Given sample size m, we let $\xi_i^{(m)}$ denote the number of polymorphic sites at which one variant (the derived mutation) is observed in i copies. The collection

$$\boldsymbol{\xi}^{(n)} = \left(\xi_1^{(m)}, \ldots, \xi_{m-1}^{(m)}\right) \tag{12}$$

is known as the (unfolded) *site-frequency spectrum*. If information about ancestral states are unavailable, so that one does not know which variant is new, one considers the *folded* spectrum $\boldsymbol{\eta}^{(m)} = \left(\eta_1^{(m)}, \ldots, \eta_{\lfloor \frac{m}{2} \rfloor}^{(m)}\right)$ in which

$$\eta_i^{(m)} = \begin{cases} \xi_i^{(m)} + \xi_{m-i}^{(m)} & \text{if } i < \frac{m}{2}, \\ \xi_i^{(m)} & \text{if } i = \frac{m}{2}. \end{cases} \tag{13}$$

[20] obtains closed-form solutions for expected values and (co)-variances of the site-frequency spectrum associated with the Kingman coalescent $\Pi^{(K)}$. Indeed [20],

$$\mathbb{E}^{(K)}\left[\xi_i^{(m)}\right] = \frac{\theta}{i}, \quad i \in [m-1]. \tag{14}$$

Let $B_i^{(m)}$ denote the random total length of branches subtending $i \in [m-1]$ leaves. Result

(14) follows from [20]

$$\mathbb{E}^{(K)}\left[B_i^{(m)}\right] = \frac{2}{i}, \quad i \in [m-1], \tag{15}$$

and the infinitely-many sites mutation model.

Due to the multiple merger property of Λ- and Ξ-coalescents, closed-form expressions for $\mathbb{E}^{(\Lambda)}\left[\xi_i^{(m)}\right]$ and $\mathbb{E}^{(\Xi)}\left[\xi_i^{(m)}\right]$ are quite hard to obtain. A key quantity in computing $\mathbb{E}\left[B_i^{(m)}\right]$ for multiple merger coalescents is

$$p^{(m)}[k,i], \qquad i \in \{1,\ldots,m-k+1\},$$

which can be described as the probability that starting from m blocks, conditioned that there are at some point in time exactly k blocks, one of them, sampled uniformly at random, subtends $i \in [m-k+1]$ leaves. See Figure 1 for an illustration.

Figure 1: An illustration of an event with probability $p^{(n)}[k, b]$, with $k = 2$, and $2 \leq b = n-3$ for n leaves (shown as open circles). In the example shown, the process reaches $k = 2$ blocks in one transition, which is a simultaneous merger involving all of the b leaves (encircled) as shown. All the b leaves are subtended by the block represented by a black square. A 'level' refers to the values taken by the block-counting process; filled symbols represent ancestral blocks.

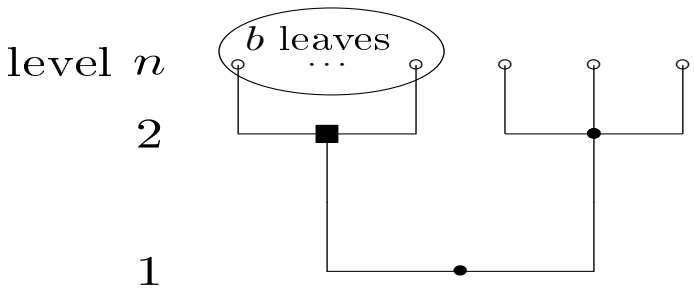

With $g(m,k)$ we denote the expected length of time during which we see $k \in \{2,\dots,m\}$ blocks, given that we started from $m \geq 2$ blocks. Given $p^{(m)}[k,i]$ and $g(m,k)$, $\mathbb{E}\left[B_i^{(m)}\right]$ can be computed as follows:

$$\mathbb{E}\left[B_i^{(m)}\right] = \sum_{k=2}^{m-i+1} p^{(m)}[k,i] \cdot k \cdot g(m,k), \tag{16}$$

where moreover $g(m,k)$ can be computed recursively. This was shown in [7] for Λ-coalescents but in fact holds for Ξ-coalescents as well. Hence, it suffices to obtain a recursion for $p^{(m)}[k,i]$. For the Λ-case, [7] obtain the recursion

$$\begin{aligned} p^{(m)}[k,i] = \sum_{n=k}^{m-1} p_{m,n} \frac{g(n,k)}{g(m,k)} \Bigg[& \frac{i-m+n}{n} p^{(n)}[k,i-m+n] \mathbb{1}_{(i>m-n)} \\ & + \frac{n-i}{n} p^{(n)}[k,i] \mathbb{1}_{(i<n)} \Bigg] \end{aligned} \tag{17}$$

in which $p_{m,n}$ is the probability of the block-counting process jumping from m to n blocks.

Before we turn to the expected site-frequency spectrum associated with Ξ-coalescents, we give the recursion to compute $g(m,k)$.

Lemma 1. *Let $p_{m,k}$ denote the probability of the block-counting process associated with a $\mathcal{P}_m$-valued Ξ-m-coalescent with transition rates (11) jumping from $m \geq 2$ to $k \in [m-1]$ blocks. For any $m > n \geq 1$, we have*

$$g(m,n) = \sum_{k=n}^{m-1} p_{m,k}\, g(k,n), \tag{18}$$

with the boundary condition

$$g(m,m) = \lambda_m^{-1} \tag{19}$$

for any $m \geq 2$, where $\lambda_m = -q_{\pi,\pi}$ with $\#\pi = m$, see (9, 10, 11).

A proof of Lemma 1 is given in the Appendix.

The expected site-frequency spectrum associated with Ξ-coalescents

Since formula (16) holds for any exchangeable coalescent, it suffices to obtain a recursion for $p^{(n)}[k,i]$ when associated with Ξ-coalescents. The recursion in Thm. (2) below for the quantity $p^{(m)}[k,i]$ is a key ingredient needed to compute $\mathbb{E}^{(\Xi)}\left[B_i^{(m)}\right]$. Before we state the result, we review our notation for partitions of positive integers. We need to consider partitions of integers, since the current active number of blocks can, in one transition, change from m to n in any number of ways in a Ξ-coalescent, if m and n are anything but small.

A *partition* of $n \in \mathbb{N}$ is a non-increasing sequence of positive integers whose sum is n. By $\tilde{n}$ we denote the set of all partitions of n, we denote by ν a generic element of $\tilde{n}$. By way of example,

$$\tilde{3} = \{(3), (2,1), (1,1,1)\}.$$

If $\nu \in \tilde{n}$, we write $|\nu| = n$. The *size* of a partition ν is defined as the length of the sequence, and is denoted by $\#\nu$. If $|\nu| = n$ with size $\#\nu = k \in [n]$, we write $|\nu| \overset{k}{=} n$. Thus, $|\nu| \overset{1}{=} n$ if and only if $\nu = (n)$; $|\nu| \overset{n}{=} n$ if and only if $\nu = (\underbrace{1, \ldots, 1}_{n \text{ times}})$.

Another way of representing an integer partition ν is by specifying how often each positive integer $i \in \mathbb{N}$ appears in ν (see [12] for details). Thus, we will also denote ν by $\langle \alpha_1, \alpha_2, \ldots \rangle$, where α_i denotes the number of times integer i appears in the given partition ν. A partition $\mu = \langle \beta_1, \ldots \rangle$ is a *sub-partition* of $\nu = \langle \alpha_1, \ldots \rangle$, denoted $\mu \subset \nu$, if and only if $\beta_i \leq \alpha_i$ for all i. For a set partition $\pi \in \mathcal{P}_n$, we define the *integer partition associated with* π, denoted $\pi^{\downarrow} \in \tilde{n}$, as the partition of n obtained by listing the block sizes of π in decreasing order. More detailed discussion of partitions of integers can be found eg. in [38].

The role of integer partitions in association with Ξ-coalescents should now be clear. We can enumerate all the possible ways the block counting process can jump from m to n active blocks by specifying the partitions of m (see eg. [12] for details). The elements of the sequence $\nu \in \tilde{m}$ specify the merger sizes, with the obvious exclusion of mergers of size 1. By way of example, integer partition $(3,2,1) \in \tilde{6}$ specifies a simultaneous merger of 3 blocks and 2 blocks, and one block remains unchanged, when we have 6 active blocks. By (11), any such

transition happens at rate $\lambda_{6,(3,2)}$.

More generally, given integer partition $|\nu| \stackrel{n}{=} m$ with $\nu = \langle \alpha_1, \alpha_2, \ldots \rangle$, put $r := n - \alpha_1$ for the number of elements of the sequence that are larger than 1, so that

$$\nu = (\nu_1, \nu_2, \ldots, \nu_r, \underbrace{1, \ldots, 1}_{\alpha_1 \text{ times}}).$$

Then $\nu_1 \geq \nu_2 \geq \cdots \geq \nu_r \geq 2$, and defining $\underline{k} := (\nu_1, \nu_2, \ldots, \nu_r)$, the corresponding transitions in which $\underline{k}$ specifies the sizes of the r mergers involved happen at rate $\lambda_{m,\underline{k}}$, see (11). Moreover, the probability $p_\nu^{(m)}$ of such a transition is given by [12, Lemma 2.2.2]

$$p_\nu^{(m)} = \frac{m!}{\prod_i \nu_i! \alpha_i!} \frac{\lambda_{m,\underline{k}}}{\lambda_m}, \tag{20}$$

where the factor

$$\frac{m!}{\prod_i \nu_i! \alpha_i!}$$

denotes the number of different ways of merging m blocks in r groups specified by $\underline{k}$ (see also (7)), and $\lambda_{m,\underline{k}}$, λ_m are given by (8, 9, 10).

Now we state our main theorem, which contains the recursion for $p^{(m)}[k,i]$ needed to compute $\mathbb{E}^{(\Xi)}\left[B_i^{(m)}\right]$. The theorem holds for all Xi-m-coalescents whose block-counting process visits every possible state with positive probability, which is true of all examples of Xi-coalescents that we consider. The assumption is not very restrictive, since it excludes only pathological cases like the star-shaped coalescent.

Theorem 2. *[12] Let $\{\Pi_t^{(\Xi,m)}, t \geq 0\}$ be a Ξ-m-coalescent with transition rates (11) such that the corresponding block-counting process hits every $k \in [m]$ with positive probability. Then, for $2 \leq k \leq m$ and $1 \leq i \leq m - k + 1$, we have*

$$p^{(m)}[k,i] = \sum_{n=k}^{m-1} \sum_{\substack{|\nu| \stackrel{n}{=} m \\ \nu = \langle \alpha_1, \ldots \rangle}} p_\nu^{(m)} \frac{g(n,k)}{g(m,k)} \sum_{j=1}^{i \wedge (n-k+1)} \sum_{\substack{\mu \subset \nu \\ |\mu| \stackrel{j}{=} i \\ \mu = \langle \beta_1, \ldots \rangle}} p^{(n)}[k,j] \frac{\prod_\ell \binom{\alpha_\ell}{\beta_\ell}}{\binom{n}{j}}, \tag{21}$$

with the boundary cases $p^{(m)}[m,i] = \mathbb{1}_{(i=1)}$.

A proof is provided in the Appendix.

Specific examples of Xi-coalescents

Out of the rich class of Xi-coalescents, several special cases have been identified either for biological / modeling relevance or mathematical tractability. The example most relevant for us is concerned with diploidy.

Haploid population models are probably the most common models in mathematical population genetics. Diploidy, and other forms of polyploidy, are, however, widely found in nature. Atlantic cod is diploid, and oysters show both tetraploidy and triploidy [cf. eg. 28]. In polyploid models which admit skewed offspring distribution, one should observe up to $M \geq 2$ simultaneous mergers, where M is some fixed number which reflects the level of polyploidy.

Indeed, a mathematical description of a Xi-coalescent which admits up to M simultaneous mergers is as follows. Take a finite measure Λ on $[0,1]$ (which would normally describe a Λ-coalescent). For convenience, let $F(dx) := \frac{\Lambda(dx)}{\Lambda([0,1])}$ be the corresponding normalized probability measure. Then, with $M \geq 2$, define the measure Ξ on the simplex Δ by

$$\Xi(d\mathbf{y}) := \frac{1}{M} \int_{[0,1]} \delta_{(\underbrace{\frac{x}{M}, \ldots, \frac{x}{M}}_{M \text{ times}}, 0, 0, \ldots)}(d\mathbf{y})\, F(dx). \tag{22}$$

The interpetation is this: If the normalized Lambda-measure F produces a multiple merger event, in which individual active ancestral lineages (blocks of the current partition) take part with probability $x \in (0,1]$, then the participating lineages are randomly grouped into M simultaneous mergers (each with probability $\frac{1}{M}$). Observe that if $F(dx) = \delta_0(x)dx$, then (22) becomes $\Xi(d\mathbf{y}) = \frac{1}{M}\int_{[0,1]} \delta_{(0,\ldots)}(d\mathbf{y})\delta_0(x)dx = \frac{1}{M}\delta_{(0,\ldots)}(d\mathbf{y})$, which corresponds to a Kingman-coalescent with time scaled by a factor $\frac{1}{M}$.

For a given merger of the Ξ-coalescent into $r \in [M]$ groups of sizes given by $\underline{k} = (k_1, \ldots, k_r)$, with $k_1 \geq \ldots \geq k_r \geq 2$ and $|\underline{k}| := k_1 + \cdots + k_r$, when n active ancestral

lineages are present, with $n - |\underline{k}|$ lineages unaffected by the given merger, the transition rates are given by

$$\lambda_{n;\underline{k}} = \sum_{\ell=0}^{(n-|\underline{k}|)\wedge(M-r)} \binom{n-|\underline{k}|}{\ell} (M)_{r+\ell} M^{-(|\underline{k}|+\ell)} \int_{[0,1]} x^{|\underline{k}|+\ell-2}(1-x)^{n-(|\underline{k}|+\ell)} F(dx). \tag{23}$$

The rates (23) depend on the choice of the probability measure F determined by the underlying population model. A proof of (23) is in the Appendix.

Xi-coalescents which admit at most $M = 4$ simultaneous mergers arise from diploid Cannings population models with skewed offspring distribution as shown in [30, 6], and are thus relevant for population genetics. In fact, if a haploid model (for example for mitochondrial DNA) is governed by a Λ-coalescent, the corresponding diploid model (concerning the core genome) might naturally lead to Xi-coalescents. Indeed, [6] derive a 4-fold Xi-coalescent from a diploid model, in which exactly one pair of diploid parents contribute diploid offspring in each reproduction event. Hence, since 4 parental chromosomes are involved in each event, one can observe up to 4 simultaneous mergers. This was also observed by [30] in association with a diploid population model, but under a more general reproduction law than considered by [6]. Xi-coalescents are also classified by [11] for a very general diploid exchangeable Cannings model in which arbitrary pairs of diploid parents contribute offspring in each reproduction event. This generalises the model by [30], in which each individual forms at most one parental pair in each generation. For a detailed classification of coalescent limits, see [30, 33, 29, 11].

In truly diploid models, as considered by [30, 6], selfing is excluded, which leads to a 'separation of timescales' phenomenon in the ancestral process , in which blocks which reside in the same diploid individual instantaneously 'disperse'; thus the configuration of blocks in diploid individuals becomes irrelevant in the ancestral process (see Cor. 4.3 in [30]).

A natural candidate for F may be the beta distribution with parameters $\vartheta > 0$ and $\gamma > 0$

(cf. e.g. [6]), with density

$$\frac{\Gamma(\vartheta+\gamma)}{\Gamma(\vartheta)\Gamma(\gamma)} x^{\vartheta-1}(1-x)^{\gamma-1}, \quad x\in[0,1].$$

In this case, the rate $\lambda_{n,\underline{k}}$ in (23) takes the form

$$\lambda_{n,\underline{k}} = \sum_{\ell=0}^{(n-|\underline{k}|)\wedge(4-r)} \binom{n-|\underline{k}|}{\ell} (4)_{r+\ell}\, 4^{-|\underline{k}|+\ell} \cdot \frac{B(|\underline{k}|+\ell+\vartheta-2, n+\gamma-(|\underline{k}|+\ell))}{B(\vartheta,\gamma)}. \tag{24}$$

A different choice is based on a model of Eldon and Wakeley [19], where

$$\Lambda(dx) = F(dx) = \frac{2}{2+\psi^2}\delta_0(dx) + \frac{\psi^2}{2+\psi^2}\delta_\psi(dx), \quad \psi\in[0,1].$$

In this case, the rates reduce to

$$\lambda_{n;\underline{k}} = \frac{1}{2+\psi^2} \sum_{l=0}^{(n-|\underline{k}|)\wedge(4-r)} \binom{n-|\underline{k}|}{\ell} (4)_{r+\ell}\, 4^{-|\underline{k}|+\ell} \cdot (1-\psi)^{n-(|\underline{k}|+\ell)} \psi^{|\underline{k}|+\ell} + \mathbb{1}_{(r=1,\, k_1=2)}\, \frac{1}{2}\frac{1}{2+\psi^2}. \tag{25}$$

In (25) the parameter ψ has a clear biological interpretation as the fraction of the diploid population replaced by the offspring of the reproducing parental pair in one generation. The interpretation of the parameters in (24) is perhaps less clear.

A multi-loci ancestral recombination graph in which simultaneous mergers are admitted are obtained by [6] in which the framework of [30] is borrowed. There, one can think of the reproduction model as a two-atom Lambda-measure, one atom at zero, and another at some point $\psi\in(0,1)$. If the atom at ψ has mass of order at most N^{-2}, the limit process admits simultaneous mergers. The order N^{-2} represents the order of the expected time which two gene copies need to coalesce when only 1 diploid offspring is produced in each reproduction event. In [6], complete dispersion of chromosomes also occurs, and the configuration of blocks among diploid individuals becomes irrelevant in the limit process.

Xi-coalescents can also be obtained from a population model where the population size varies substantially due to recurrent bottlenecks. This has been introduced and discussed

in [8], who obtain a randomly time-changed Kingman coalescent, which thus yields a Xi-coalescent.

Durrett and Schweinsberg [15, 16] show that a Xi-coalescent gives a good approximation [16, cf. Prop. 3.1] to the genealogy of a locus subject to recurrent beneficial mutations.

These examples suggest that Xi-coalescents form an important class of mathematical objects with which to study genetic diversity.

Numerical results

The result (21) in Thm. 2 holds for arbitrary number (up to $\lfloor m/2 \rfloor$) of simultaneous mergers. However, it is quite a challenge to compute $p^{(m)}[k,i]$ for large m due to the number of terms in the recursion (21). In our numerical examples, we restrict to Ξ-coalescents which admit simultaneous mergers in up to 4 groups. The number 4 represents the number of parental chromosomes involved in a large reproduction event, ie. when the number of diploid offspring of a given pair of diploid parents constitute a significant part of the total population. Such Ξ-coalescents can be shown to arise from *truly* diploid Cannings population models [30, 6], and are thus highly relevant for population genetics.

The recursion for $p^{(m)}[k,i]$ simplifies a bit when one restricts to Ξ-coalescents with at most $M \geq 2$ simultaneous mergers ($M = 1$ simply gives a Lambda-coalescent), as shown in Cor. 3. Before we state Cor. 3, we briefly review *ordered mergers*. Define $\mathbb{M} := \{2, 3, \ldots\}$, and let

$$\mathcal{M}(m,n) := \{\underline{k} = (k_1, \ldots, k_r) : r \in [M], k_i \in \mathbb{M}, k_1 \geq \ldots \geq k_r \geq 2, n = m - |\underline{k}| + r\} \quad (26)$$

denote the set of single and up to M simultaneous ordered mergers by which the block-counting process can jump from $m \geq 2$ to $n \in [m-1]$ blocks in $r \in [M]$ mergers. Thus, $\mathcal{M}(m,n)$ corresponds to the set of all integer partitions $|\nu| \stackrel{n}{=} m$ such that, if $\nu = \langle \gamma_1, \gamma_2, \ldots \rangle$, we have $\sum_{j \geq 2} \gamma_j = n - \gamma_1 \in [M]$. Indeed, for $1 \leq n < m$, we have a bijection

$$\{\nu = \langle \gamma_1, \gamma_2, \ldots \rangle : |\nu| \stackrel{n}{=} m, n - \gamma_1 \leq M\} \ni \nu \mapsto (\nu_1, \ldots, \nu_{r_\nu}) \in \mathcal{M}(m,n),$$

with $r_\nu := \max\{j : \nu_j \geq 2\} = n - \gamma_1$, for $\nu = \langle \gamma_1, \gamma_2, \ldots \rangle$. Obviously, the inverse bijection is given by

$$\mathcal{M}(m,n) \ni \underline{k} \mapsto \nu := (k_1, k_2, \ldots, k_r, \underbrace{1, \ldots, 1}_{n-r \text{ times}}).$$

Thus informally, ordered mergers are just integer partitions where all elements equal to one are omitted.

For ease of presentation, we will also write $\mathcal{M}(m,n) \ni \mu = (\alpha) \equiv (\alpha_2, \alpha_3, \ldots)$ where α_j denotes the number of occurrences of merger of size j in merger μ. By $\gamma = (\beta) \subset \mu = (\alpha)$ we denote a *submerger* γ of μ where $\beta_j \leq \alpha_j$ for all j, *including* the case $\beta_j = 0$. Finally, the *size* of $\mu = \underline{k} \in \mathcal{M}(m,n)$ is just the length of the sequence $\underline{k}$, i.e. $\#\mu = r$ for $\mu = \underline{k} = (k_1, \ldots, k_r)$. Note that if μ is the ordered merger corresponding to some integer partition $\nu = \langle \gamma_1, \gamma_2, \ldots \rangle$, $|\nu| \overset{n}{=} m$ as above, then $\#\nu = n = \#\mu + \gamma_1$.

Corollary 3. *Let $\{\Pi_t^{(\Xi,n)}, t \geq 0\}$ be a $\mathcal{P}_n$-valued Ξ-coalescent with transition rates (11) such that at most $M \geq 2$ simultaneous mergers are possible. Then*

$$\begin{aligned} p^{(m)}[k,i] = &\sum_{n=k}^{m-1} \frac{g(n,k)}{g(m,k)} \sum_{\substack{\mu \in \mathcal{M}(m,n) \\ \mu = (\alpha)}} p_\mu^{(m)} \\ &\cdot \sum_{j=1}^{i \wedge (n-k+1)} \sum_{\substack{\gamma \subset \mu \\ \gamma \in \mathcal{M}(i,j) \\ \gamma = (\beta)}} p^{(n)}[k,j] \binom{n}{j}^{-1} \prod_{\ell \geq 2} \binom{\alpha_\ell}{\beta_\ell} \binom{n - \#\mu}{j - \#\gamma}, \end{aligned} \tag{27}$$

with the boundary cases $p^{(m)}[m,i] = \mathbb{1}_{(i=1)}$.

Informally, the first sum in recursion (27) is over the number of blocks the block-counting process can jump to, given that it starts in m, and conditioned on it hits k. The second sum is over all (up to M simultaneous) mergers (μ) in which one can jump from m to n blocks, and the last sum is over all the ways mergers involving the i leaves can be nested within each given merger μ.

Corollary 3 is simply another way of representing $p^{(m)}[k,i]$ (21) (see Thm. 2) in terms of ordered mergers. The switch in focus to ordered mergers from integer partitions obliviates

the need to keep track of all the 1s. Also, since we restrict to Ξ-coalescents which admit at most M simultaneous mergers, the required mergers can be generated quite efficiently.

The computation of $p^{(m)}[k,i]$ can be checked by noting that for each fixed $m \geq 2$, and with $2 \leq k \leq m$, $\sum_i p^{(m)}[k,i] = 1$. One can also compute $\mathbb{E}^{(\Xi)}\left[B_1^{(m)}\right]$ with a simple recursion as follows. Let $\mathbb{E}^{(\Xi)}\left[B_1^{(m,i)}\right]$ denote the expected length of i external branches, given $m \geq 2$ active blocks (i of which are singleton blocks). Define, for all $m \in \mathbb{M}$,

$$\begin{aligned} \mathcal{C}_m :&= \{\underline{k} = (k_1, \ldots, k_r) : r \in [M], k_i \in \mathbb{M}, k_1 \geq \ldots \geq k_r \geq 2, |\underline{k}| \leq m\} \\ &= \bigcup_{n=1}^{m-1} \mathcal{M}(m,n) \end{aligned} \tag{28}$$

as the set of all ordered (up to M simultaneous) mergers given m active blocks, where $\mathcal{M}(m,n)$ was defined in (26).

Lemma 4. *Let $\Pi^{(\Xi,n)} \equiv \{\Pi_t^{(\Xi,n)}, t \geq 0\}$ be a $\mathcal{P}_n$-valued Ξ-coalescent with transition rates (11) such that at most M simultaneous mergers are possible. With λ_m given by (9, 10), let $\mathcal{C}_m$ be given by (28), and let $p_{\underline{k}}^{(m)}$ denote the probability of merger $\underline{k}$ when the number of active blocks is m. Then, with $|\underline{k}| = k_1 + \cdots + k_r$,*

$$\mathbb{E}^{(\Xi)}\left[B_1^{(m,i)}\right] = \frac{i}{\lambda_m} + \sum_{\underline{k} \in \mathcal{C}_m} p_{\underline{k}}^{(m)} \sum_{j=(i-m+|\underline{k}|)^+}^{i \wedge |\underline{k}|} \binom{m}{|\underline{k}|}^{-1} \binom{i}{j} \binom{m-i}{|\underline{k}|-j} \mathbb{E}^{(\Xi)}\left[B_1^{(m-|\underline{k}|+r, i-j)}\right] \tag{29}$$

with the boundary condition $\mathbb{E}^{(\Xi)}\left[B_1^{(m,0)}\right] = 0$.

Expected branch lengths $\mathbb{E}^{(m,\Pi)}[B_i]$

Under the infinitely many sites mutation model, the expected site-frequency spectrum is given by $\mathbb{E}^{(\Pi)}\left[\xi_i^{(m)}\right] = \theta\, \mathbb{E}^{(\Pi)}\left[B_i^{(m)}\right]$ where $\theta > 0$ is the appropriately scaled mutation rate. Hence, it suffices to consider $\mathbb{E}^{(\Pi)}\left[B_i^{(m)}\right]$ in a comparison of the site-frequency spectrum associated with different coalescent models. In Figure 2, we consider the 4-fold Ξ-coalescent,

when the measure F in (23) is associated with the beta-density, with $\alpha \in [1, 2)$ [36],

$$F(dx) = \frac{x^{1-\alpha}(1-x)^{\alpha-1}}{B(2-\alpha, \alpha)}dx, \quad x \in [0, 1]. \tag{30}$$

The range of α is the interval $[1, 2)$, since for $\alpha \in [1, 2)$, one obtains a Lambda-Beta coalescent from a supercritical branching population model [36]. We do not have a microscopic diploid population model which explicitly yields a Xi-Beta coalescent. However, the results of [30] indicate that such a process should exist. Also, the Lambda-Beta coalescent is one of the most studied examples of Lambda-coalescents. However, the existence of the Xi-Dirac coalescent was proved by [6]. The expected branch lengths associated with a Lambda-coalescent with F-measure (30), as well as the Kingman coalescent, are also shown in Figure 2 for comparison. We consider the normalised expected spectrum

$$\varphi_i^{(n,\Pi)} := \frac{\mathbb{E}^{(\Pi)}\left[B_i^{(n)}\right]}{\mathbb{E}^{(\Pi)}\left[B^{(n)}\right]}, \quad i \in [n-1]. \tag{31}$$

in which $\mathbb{E}^{(\Pi)}\left[B^{(n)}\right]$ is the expected total size of the genealogy $\left(B^{(n)} := B_1^{(n)} + \cdots + B_{n-1}^{(n)}\right)$. Define $R_i^{(n)} := \frac{B_i^{(n)}}{B^{(n)}}$. Let $\xi^{(n)} := \xi_1^{(n)} + \cdots + \xi_{n-1}^{(n)}$ denote the (random) total number of segregating sites, and define the normalised spectrum $\zeta_i^{(n)} := \frac{\xi_i^{(n)}}{\xi^{(n)}}$ (with $\zeta_i^{(n)} \equiv 0$ if $\xi^{(n)} = 0$). The reasons for our preference for $\varphi_i^{(n,\Pi)}$ over $\mathbb{E}^{(\Pi)}\left[\xi_i^{(n)}\right]$ are the following. The quantity $\varphi_i^{(n,\Pi)}$, which is an approximation of the expected normalised spectrum $\mathbb{E}^{(\Pi)}\left[\zeta_i^{(n)}\right]$ is, clearly, *not* a function of the mutation rate θ. The expected normalised spectrum $\mathbb{E}^{(\Pi)}\left[\zeta_i^{(n)}\right]$ is well approximated by $\mathbb{E}^{(\Pi)}\left[R_i^{(n)}\right]$, and is also quite robust to changes in mutation rate, if the mutation rate is not very small [18]. Since $\varphi_i^{(n,\Pi)}$ is a decent approximation of $\mathbb{E}^{(\Pi)}\left[R_i^{(n)}\right]$ (results not shown), $\varphi_i^{(n,\Pi)}$ is therefore a good approximation of $\mathbb{E}^{(\Pi)}\left[\zeta_i^{(n)}\right]$. One can therefore use $\varphi_i^{(n,\Pi)}$ to estimate coalescent parameters, for example by minimising sum-of-squares, without the need to jointly estimate the mutation rate.

As Figure 2 shows, the corresponding Lambda-Beta and Xi-Beta coalescents predict different patterns of the site-frequency spectrum, at least for $\alpha \in \left[1, \frac{3}{2}\right]$. Both processes can predict a significant excess of singletons relative to the Kingman coalescent.

Similar conclusions can be reached from Figure 3, in which the F measure is associated with the Dirac-measure $F(dx) = \delta_\psi(x)dx$ for some $\psi \in [0, 1]$. The Xi-Dirac-coalescent, for $\psi = 0.95$, displays a multimodal graph of $\varphi_i^{(n,\Pi)}$, but the relative 'height' of the modes is small ($< 1\%$ of the total expected length). As well as an excess of singleton polymorphisms, [2] observe multi-modality, or small 'bumps', in the site-frequency spectrum associated with the *Ckma* gene in Atlantic cod.

Figure 2: The relative expected lengths $\varphi_i^{(n,\Pi)}$ (31) for $n = 100$ and coalescent process Π as follows: `L(a)` denotes the Lambda-Beta coalescent with parameter value `a` (α) as shown, and `X(a)` denotes the Xi-Beta coalescent with parameter value `a` (α) as shown. In **A**, only the first 4 classes are shown; in **B**, the remaining classes.

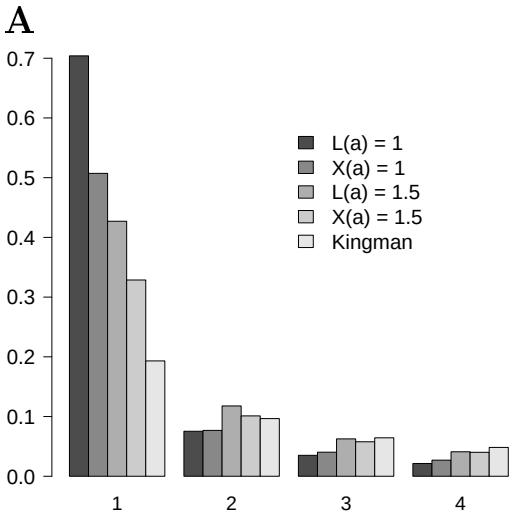

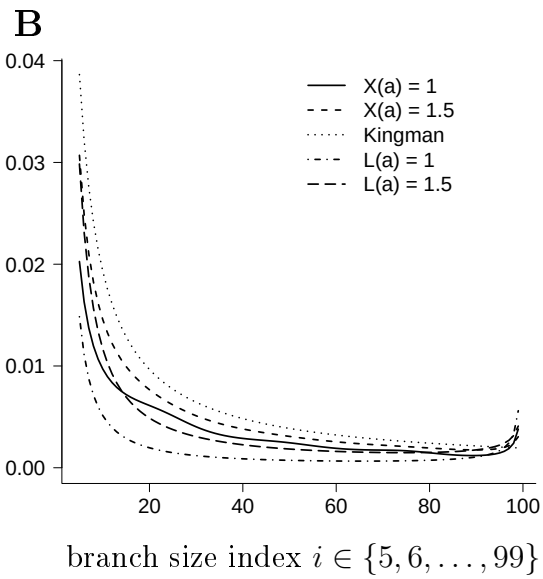

Figure 3: The relative expected lengths $\varphi_i^{(n,\Pi)}$ (31) for $n = 100$ and coalescent process Π as follows: `L(p)` denotes the Lambda-Dirac coalescent with parameter value `p` (ψ) as shown, and `X(p)` denotes the Xi-Dirac coalescent with parameter value `p` (ψ) as shown. In **A**, only the first 4 classes are shown; in **B**, the remaining classes.

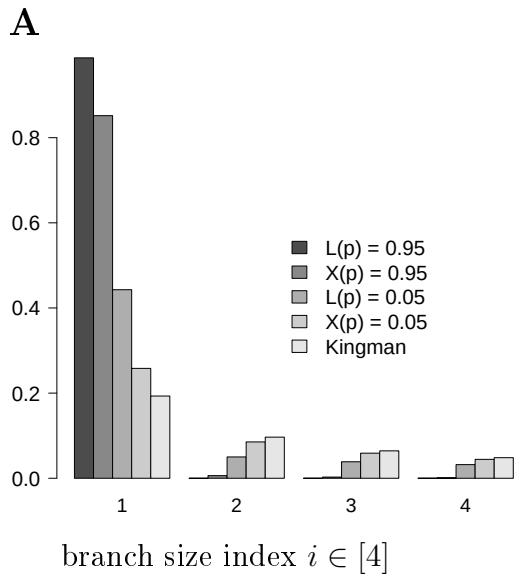

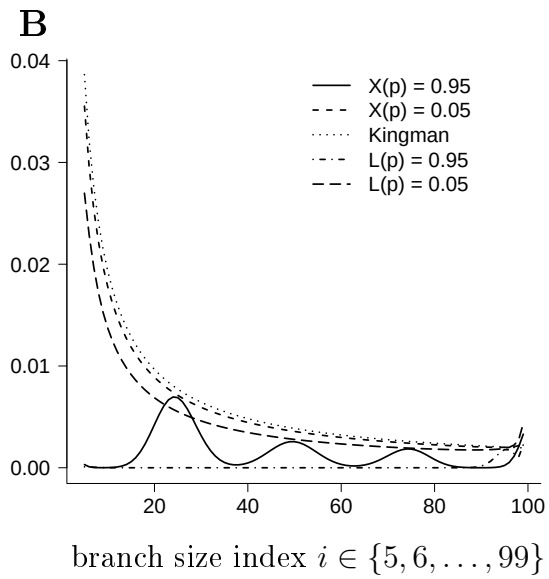

Coalescent parameter estimates

As Figures 2 and 3 indicate, one should be able to distinguish between corresponding Xi- and Lambda-coalescents from an observed site-frequency spectrum. In Table 1 we record parameter estimates obtained when the 'data' are branch lengths B_i simulated under either a 4-fold Xi-coalescent (23) or a Lambda-coalescent (5) with parameter values as shown. We consider an extensive comparison of different examples of the large class of Xi-coalescent a bit beyond the scope of the current work. If $0 \leq \vartheta < 1$, $\Xi(\vartheta)$ denotes a 4-fold Dirac-Xi coalescent, with F-measure $F(dx) = \delta_\psi(x)dx$ in (23), and $\Lambda(\vartheta)$ a Dirac-Lambda coalescent. If $\vartheta \in [1, 2)$, $\Xi(\vartheta)$ denotes a 4-fold Beta-Xi coalescent, and $\Lambda(\vartheta)$ a Beta-Lambda coalescent. Estimates of ϑ attributed to coalescent process Π_2 are obtained with an ℓ_2 norm applied to the normalised lengths $R_i := \frac{B_i}{B}$ drawn from coalescent process Π_1, and $\varphi_i^{(n,\Pi_2)}$,

$$\ell_2(\Pi_1, \Pi_2) = \sqrt{\sum_{i=1}^{n-1} \left(R_i - \varphi_i^{(n,\Pi_2)}\right)^2}. \tag{32}$$

If the B_i are drawn from a Xi-Dirac coalescent (Π_1), we estimate ϑ associated with a Lambda-Dirac coalescent (Π_2), and vice versa. If the B_i are drawn from a Xi-Beta coalescent (Π_1), we estimate ϑ associated with a Beta-Lambda coalescent (Π_2), and vice versa.

A more suitable distance statistic could be

$$d(\Pi_1, \Pi_2) = \sqrt{\sum_{i=1}^{n-1} \frac{\left(R_i - \mathbb{E}^{(n,\Pi_2)}[R_i]\right)^2}{\mathbb{V}^{(n,\Pi_2)}[R_i]}}, \tag{33}$$

where $\mathbb{V}^{(n,\Pi_2)}[R_i]$ denotes the variance of R_i computed with respect to Π_2. However, we can neither represent $\mathbb{V}^{(n,\Pi_2)}[R_i]$ nor $\mathbb{E}^{(n,\Pi_2)}[R_i]$ as simple functions of ϑ or n. In actual applications, one would replace R_i in (33) with $\chi_i := \frac{\xi_i}{\xi}$, the normalised site-frequency spectrum.

As Table 1 shows, a Lambda-Dirac coalescent underestimates ψ when the data are generated by a Xi-Dirac coalescent, and Lambda-Beta overestimates α when the data are generated by a Xi-Beta coalescent. When we switch the generation of data from Xi- to Lambda-

coalescents, we reach the opposite conclusions. A Xi-Dirac coalescent overestimates the parameter (ψ) when the data are generated by a Lambda-Dirac coalescent, and the Xi-Beta coalescent underestimates α when the data are generated by a Lambda-Beta coalescent.

The difference between the corresponding Xi- and Lambda-coalescents are further illustrated in Figure 4, in which the distance between the normalised expected spectra $\varphi_i^{(n,\Pi)}$ (31), with Π as shown, is quantified by the ℓ_2 norm (32). The graphs in Figure 4 show clearly that even when the parameters associated with the corresponding Xi- and Lambda-coalescents are the same, the difference in $\varphi_i^{(n,\Pi)}$ can be substantial (except of course when the Lambda-coalescent is the Kingman coalescent; which happens when $\alpha = 2$ or $\psi = 0$).

The difference in estimates between corresponding Lambda- and Xi-coalescents may be understood from the way the Xi-coalescent process is constructed. Indeed, given k blocks drawn from the associated Lambda-coalescent, we see one k-merger with probability 4^{1-k}, which quickly becomes small as k increases. A much more likely outcome is for the blocks to become (evenly) distributed into four groups. The effect on the genealogy of drawing a large number k is thus reduced in a Xi-coalescent relative to the corresponding Lambda-coalescent.

C code written for the computations is available at `http://page.math.tu-berlin.de/~eldon/programs.html`.

Table 1: Estimates of coalescent parameter ϑ, obtained by the use of the ℓ_2 norm (32), when the 'data' (branch lengths of a realized genealogy) are obtained by a Lambda- ($\Lambda(\vartheta)$) or a Xi-coalescent ($\Xi(\vartheta)$) as shown. In the top half, the 'data' is generated by a 4-fold Xi-coalescent with parameter value as shown, and a parameter estimate obtained for the corresponding Lambda-coalescent. In the bottom half, branch lengths are generated from a Lambda-coalescent with parameter values as shown, and parameter estimates obtained for the corresponding 4-fold Xi-coalescent. Parameter estimates (mean $\overline{\vartheta}$; standard deviation $\widehat{\vartheta}$) obtained for $n = 50$ leaves from 10^4 replicates. Estimates were obtained over the grids $\{0.0, 0.05, \ldots, 0.95\}$ and $\{1.0, 1.05, \ldots, 1.95\}$.

$\Pi(\vartheta)$	$\overline{\vartheta}$	$\widehat{\vartheta}$
$\Xi(0.05)$	0.02	0.029
$\Xi(0.95)$	0.30	0.225
$\Xi(1.0)$	1.44	0.257
$\Xi(1.5)$	1.70	0.208
$\Lambda(0.05)$	0.23	0.144
$\Lambda(0.95)$	0.95	0.035
$\Lambda(1.0)$	1.01	0.041
$\Lambda(1.5)$	1.20	0.256

Figure 4: The distance between $\varphi_i^{(n,\,\Xi\text{-Beta})}$ (31) and $\varphi_i^{(n,\,\Lambda\text{-Beta})}$ (**A**) as quantified by the ℓ_2 norm (32); the distance between $\varphi_i^{(n,\,\Xi\text{-Dirac})}$ and $\varphi_i^{(n,\,\Lambda\text{-Dirac})}$ (**B**) as quantified by the ℓ_2 norm. The number of leaves $n = 50$. Values were computed over the grid $\{1.0, 1.025, \ldots, 2\}$ when associated with the Beta-coalescent (**A**); and $\{0, 0.025, \ldots, 0.975\}$ when associated with the Dirac-coalescent (**B**).

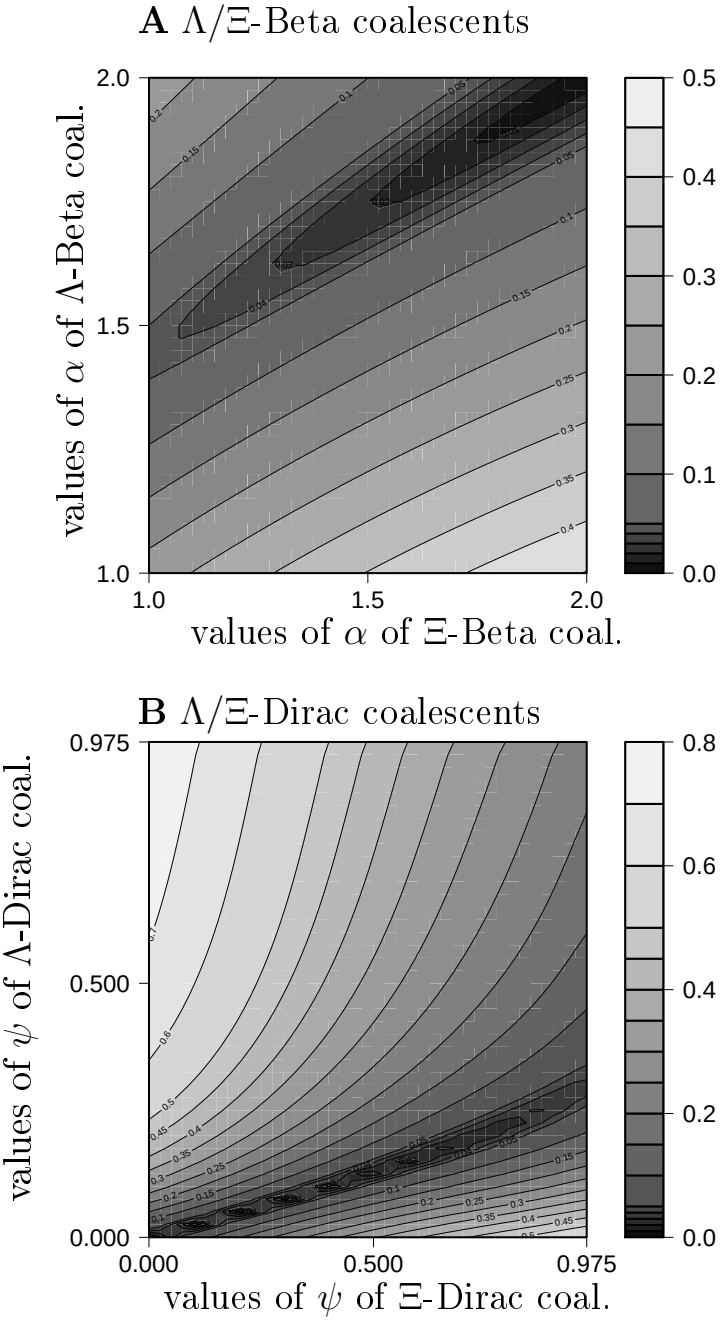

Application to Atlantic cod data

Atlantic cod is a diploid highly fecund marine organism, whose reproduction is potentially characterised by a skewed offspring distribution [1, 2]. Since Xi-coalescents can arise from diploid population models which admit skewed offspring distributions [30, 6], one should analyse population genetic data of nuclear loci in diploid highly fecund populations with Xi-coalescent models. Indeed, [2] obtain population genetic data at three nuclear loci from Atlantic cod. We use the ℓ_2-norm (32) to fit (see Table 2) the 4-fold Xi-Beta and Xi-Dirac coalescents to the unfolded site-frequency spectrum (USFS) of *Ckma*, *Myg*, and *HbA2* genes obtained by [2].

The USFS of the nuclear genes *Ckma*, *Myg*, and *HbA2* are all characterised by a high relative amount of singletons. Thus, singletons have the most weight in our estimate, in particular since we do not calibrate the difference between observed and expected values by the variance. The estimates of α associated with the Xi-Beta coalescent are therefore all at 1.0, which we attribute to the excessive amount of singletons. The excessive amount of singletons also increases the estimate of ψ associated with the Xi-Dirac coalescent. In particular, our Xi-based estimates of ψ are higher than the Lambda-based estimates obtained by [2] (see Table 2 in [2]). Possibly the Xi-coalescent assigns less mass to the external branches than the corresponding Lambda-coalescent for a given parameter value, but the exact shift in mass may vary between different Xi-coalescents. The Xi-Dirac coalescent is able to predict the excessive amount of singletons, the Xi-Beta coalescent much less so (Figures 5–6).

Our estimates of ψ for the combined data on *Ckma* is quite a bit smaller than for the partitioned data (into A and B alleles), and for the supposedly neutral loci *Myg* and *HbA2*. [2] also observe a similar pattern. Of the three loci, the Xi-coalescents give best fit to the *Ckma* data. In view of the modes in the right tail of the USFS for *Ckma*, [2] conclude that *Ckma* is under strong selection. Even though the Xi-Dirac coalescent does show multi-modal spectrum (Figure 3), the modes are small relative to the expected length of external branches, and do not quite explain the modes observed for *Ckma* (Figure 6).

Table 2: Parameter estimates $(\widehat{\alpha}, \widehat{\psi})$ obtained by minimising the ℓ_2-norm (32) for the unfolded site-frequency spectrum of Atlantic cod data [2]. Parameter α associated with the Xi-Beta coalescent was estimated over the grid $\{1.0, 0.05, \ldots, 1.95\}$, and ψ associated with the Xi-Dirac coalescent was estimated over the grid $\{0.0, 0.05, \ldots, 0.95\}$.

nuclear locus	sample size	$\widehat{\alpha}$	$\ell_2\left(\widehat{\alpha}\right)$	$\widehat{\psi}$	$\ell_2\left(\widehat{\psi}\right)$
Ckma	122	1.0	0.08	0.30	0.10
CkmaA	43	1.0	0.21	0.65	0.12
CkmaB	79	1.0	0.18	0.50	0.13
Myg	45	1.0	0.27	0.80	0.14
HbA2	114	1.0	0.34	0.80	0.13

Figure 5: Comparison of observed (χ_i; ○) and expected $\left(\varphi_i^{(n,\Pi)}; ■, ▲\right)$ (31) normalised site-frequency spectrum for *Myg* and *HbA2* [2] with 4-fold Xi-coalescent process as shown. The associated parameter values are given in Table 2. In **A**, the observed and expected values for the *Myg* gene are compared; in **B** for the *HbA2* gene.

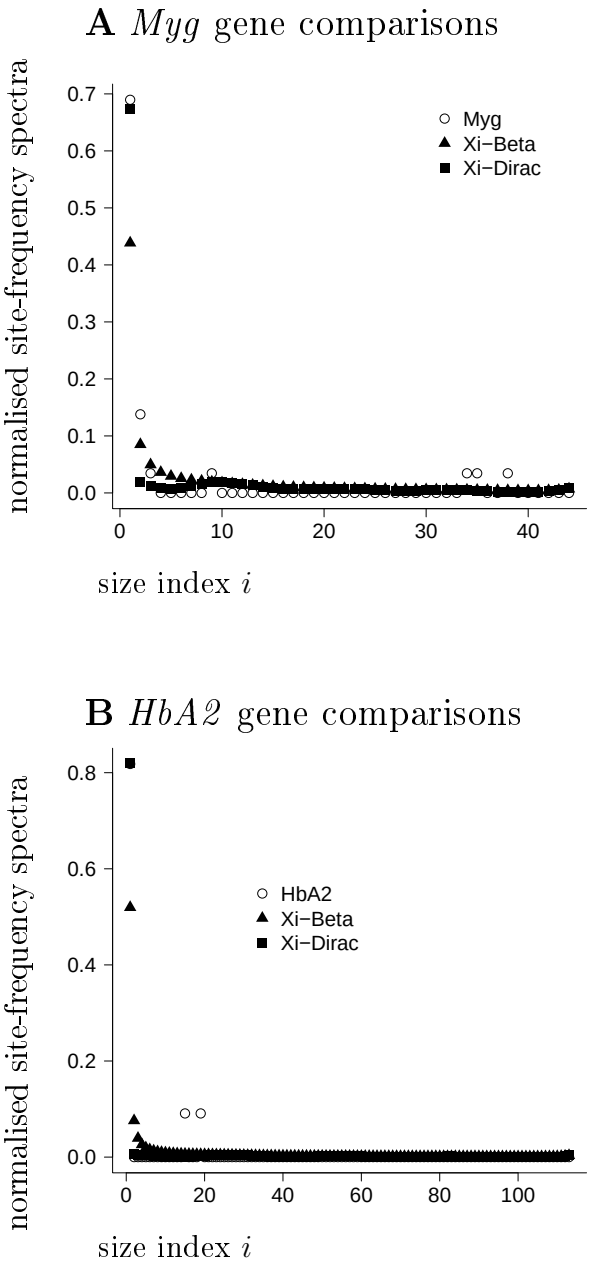

Figure 6: Comparison of observed (χ_i; ○) and expected $\left(\varphi_i^{(n,\Pi)}; ■, ▲\right)$ (31) normalised site-frequency spectrum for *Ckma* [2] with 4-fold Xi-coalescent process as shown. The associated parameter values are given in Table 2. In **A**, all the *CkmA* sequences are included; in **B**, only the *A* allele is considered; in **C**, only the *B* allele is considered.

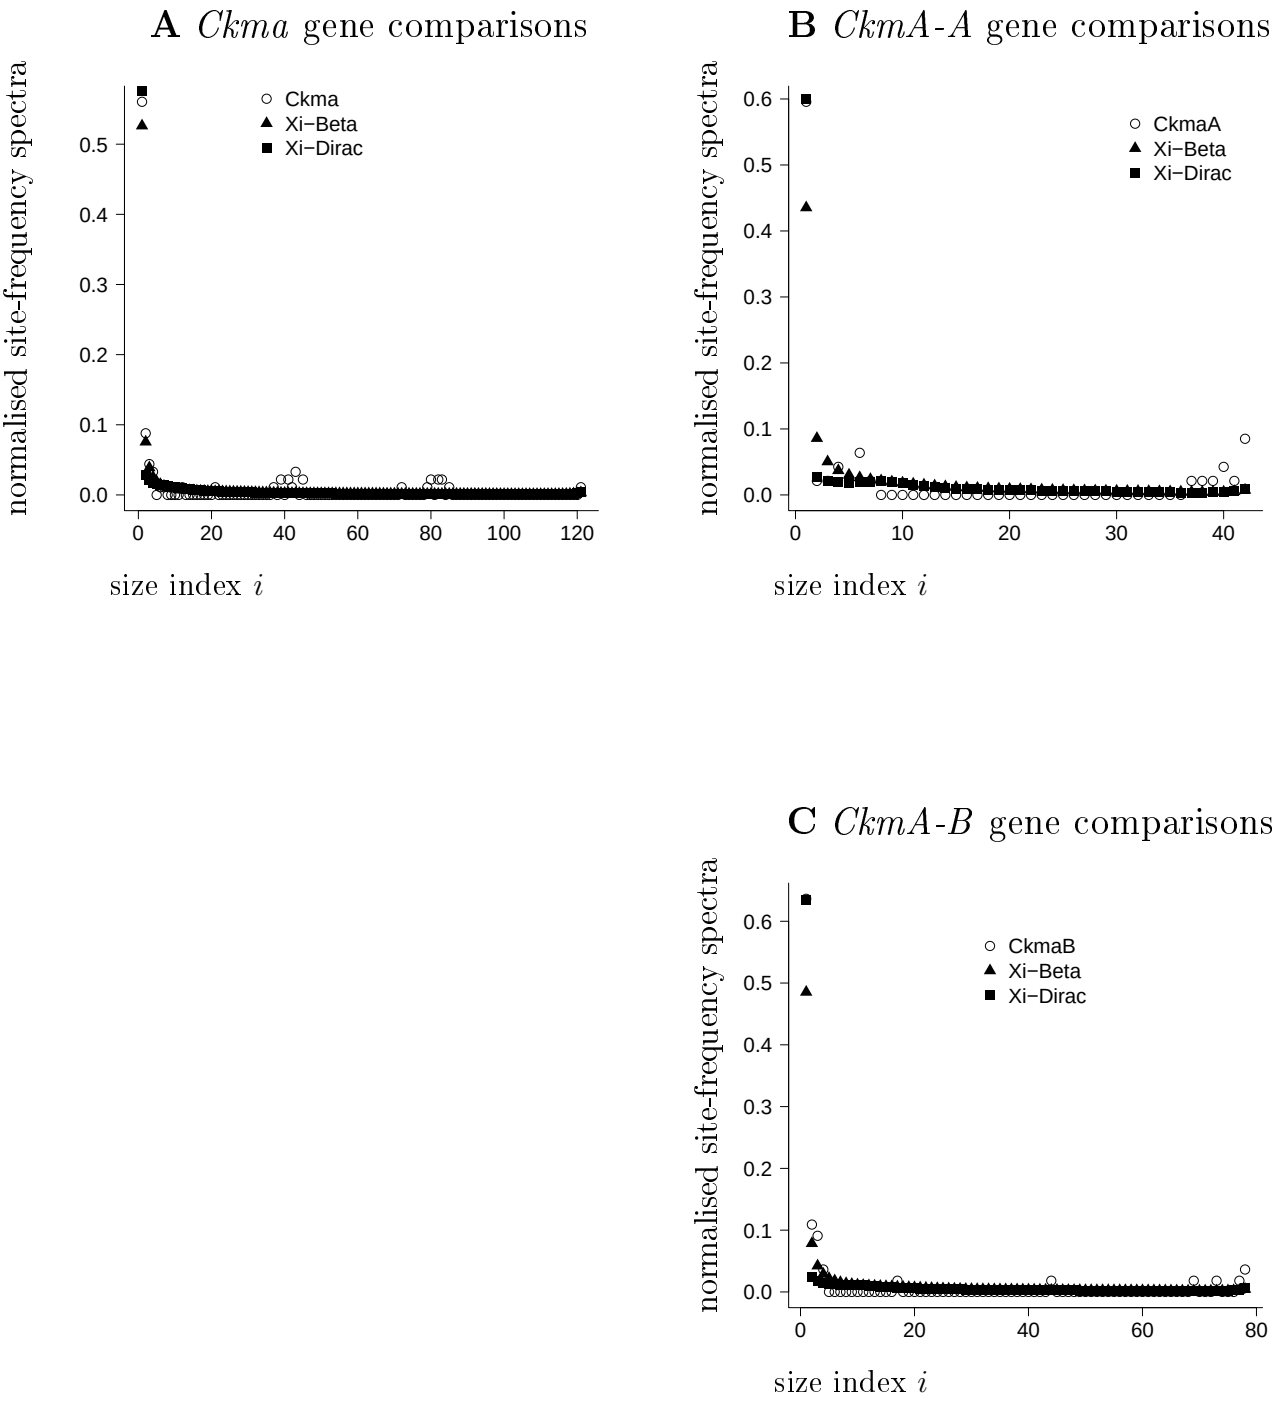

Discussion

We prove recursions for the expected site-frequency spectrum associated with Xi-coalescents which admit simultaneous multiple mergers of active ancestral lineages (blocks of the current partition). We give a class of Xi-coalescents which is 'driven' by a finite measure (a Lambda-measure) on the unit interval, which determines the law of the total number of active lineages which may merge each time. This class of Xi-coalescents can be applied to populations of arbitrary ploidy. We apply the recursions to compare estimates of coalescent parameters between Lambda- and Xi-coalescents. Finally, we estimate coalescent parameters associated with Xi-coalescents for Atlantic cod where the data are the unfolded site-frequency spectrum on nuclear loci.

The framework we develop will allow us to extend the recursions to more complicated frameworks, to, by way of example, populations structured into discrete subpopulations [17, 27], or possibly by considering some sort of continuous distribution in space [3]. However, one would extend the recursions to such more structured frameworks at the high risk of increasing their computational complexity.

Computing the full expected site-frequency spectrum for a 4-fold Xi-coalescent with sample size $n \geq 100$ takes unfortunately a bit of time (on the order of hours). We do not provide detailed analysis, but the time can be shortened by considering the lumped site-frequency spectrum, where one would collect all classes of size larger than some number $m \ll n$ into one class. How such lumping would affect the inference remains to be seen. However, one can analyse small samples ($n \leq 100$) with our recursions, as we provide an example of. Exact likelihood methods in the spirit of [5] are yet to be developed for Xi-coalescents, and will likely be computationally intensive.

Our simple method of minimising the deviation between observed and expected values of the unfolded site-frequency spectrum should not be applied in a formal test procedure, since we do not scale the deviations by the corresponding variance. We do not give recursions for the (co)-variances of the site-frequency spectrum associated with Xi-coalescents, as these will very likely be too computationally intensive to be useful. This has already been shown

to be the case for the much simpler Lambda-coalescents [7]. Our recursions provide a way to distinguish between Xi-coalescents and other demographic effects such as population growth, with the use of approximate likelihoods [18].

We obtain estimates of coalescent parameters associated with Xi-coalescents for data on nuclear loci of Atlantic cod [2]. Our estimates differ from previous estimates obtained with the use of Lambda-coalescents [2], due to the simultaneous merger characteristic of Xi-coalescents. Regardless of exact estimates, our results, coupled with those of [2], suggest that multiple merger coalescents might be the proper null model with which to analyse population genetic data on Atlantic cod, as well as other populations which may exhibit HFSOD. Our specific examples of Xi-coalescents may well represent an oversimplification of the actual mating schemes. However, one could use the framework of, say, [11], to develop new examples of Xi-coalescents for specific mating schemes, for example when one successful female produces offspring with many males.

Rigorous inference methods to distinguish the effects of high fecundity coupled with skewed offspring distribution (HFSOD) from selection are yet to be developed. The common notion is that selective sweeps lead to an excess of singletons. The main genetic signature of HFSOD is also an excess of singletons. The unfolded site-frequency spectrum of the *Ckma* gene [2] is trimodal, with an excess of singletons, and small modes of mutations of larger size. These smaller modes are not captured by the examples of Xi-coalescents that we apply. Durrett and Schweinsberg [37, 16] use a stick-breaking construction to obtain a good approximation ($\leq O(1/(\log(N))^2)$) to a selective sweep where N denotes the population size. [2] conclude that *Ckma* is under a form of balancing selection. Our examples of Xi-coalescents also give best fit to the *Ckma* gene of the three loci studied by [2], although our method does not constitute a formal test. We refer to [2] for more detailed discussion of the variation observed at *Ckma*, and the supposedly neutral nuclear loci *Myg* and *HbA2*. The congruence between our Xi-coalescent examples and *Ckma*, and between Xi-coalescents and selective sweeps studied by [37, 16], and the different site-frequency spectra predicted by Lambda- and Xi-coalescents, leads us to conclude that Xi-coalescents form an important

class of mathematical objects with which to study genetic diversity.

ACKNOWLEDGMENTS

JB and BE acknowledge support by Deutsche Forschungsgemeinschaft (DFG) grant BL 1105/3-1 as part of SPP Priority Programme 1590 'Probabilistic Structures in Evolution'.

References

[1] E Árnason. Mitochondrial cytochrome *b* variation in the high-fecundity Atlantic cod: trans-Atlantic clines and shallow gene genealogy. *Genetics*, 166:1871–1885, 2004.

[2] Einar Árnason and Katrín Halldórsdóttir. Nucleotide variation and balancing selection at the *Ckma* gene in Atlantic cod: analysis with multiple merger coalescent models. *PeerJ*, 3:e786, 2015.

[3] NH Barton, AM Etheridge, and A Véber. A new model for evolution in a spatial continuum. *Electron J Probab*, 15:162–216, 2010.

[4] N Berestycki. Recent progress in coalescent theory. *Ensaios Mathématicos*, 16:1–193, 2009.

[5] M Birkner and J Blath. Computing likelihoods for coalescents with multiple collisions in the infinitely many sites model. *J Math Biol*, 57:435–465, 2008.

[6] M Birkner, J Blath, and B Eldon. An ancestral recombination graph for diploid populations with skewed offspring distribution. *Genetics*, 193:255–290, 2013.

[7] M Birkner, J Blath, and B Eldon. Statistical properties of the site-frequency spectrum associated with Λ-coalescents. *Genetics*, 195:1037–1053, 2013.

[8] M Birkner, J Blath, M Möhle, M Steinrücken, and J Tams. A modified lookdown construction for the Xi-Fleming-Viot process with mutation and populations with recurrent bottlenecks. *ALEA Lat. Am. J. Probab. Math. Stat.*, 6:25–61, 2009.

[9] M Birkner, J Blath, and M Steinrücken. Importance sampling for Lambda-coalescents in the infinitely many sites model. *Theor Popul Biol*, 79:155–173, 2011.

[10] M Birkner, J Blath, and M Steinrücken. Analysis of DNA sequence variation within marine species using Beta-coalescents. *Theor Popul Biol*, 87:15–24, 2013.

[11] M Birkner, H Liu, and A Sturm. A note on coalescent results for diploid exchangeable population models. *Preprint*, 2015.

[12] M C Cronjäger. On the expected site-frequency spectrum associated with Ξ-coalescents. Master's thesis, TU Berlin, Berlin, 2014.

[13] P Donnelly and T G Kurtz. Particle representations for measure-valued population models. *Ann Probab*, 27:166–205, 1999.

[14] R Durrett. *Probability models for DNA sequence evolution*. Springer, New York, 2nd edition, 2008.

[15] R Durrett and J Schweinsberg. Approximating selective sweeps. *Theor Popul Biol*, 66:129–138, 2004.

[16] R Durrett and J Schweinsberg. A coalescent model for the effect of advantageous mutations on the genealogy of a population. *Stoch Proc Appl*, 115:1628–1657, 2005.

[17] B Eldon. Structured coalescent processes from a modified Moran model with large offspring numbers. *Theor Popul Biol*, 76:92–104, 2009.

[18] B Eldon, M Birkner, J Blath, and F Freund. Can the site-frequency spectrum distinguish exponential population growth from multiple-merger coalescents. *Genetics*, 199:841–856, 2015.

[19] B Eldon and J Wakeley. Coalescent processes when the distribution of offspring number among individuals is highly skewed. *Genetics*, 172:2621–2633, 2006.

[20] Y-X Fu. Statistical properties of segregating sites. *Theor Popul Biol*, 48:172–197, 1995.

[21] Y-X Fu. New statistical tests of neutrality for DNA samples from a population. *Genetics*, 143:557–570, 1996.

[22] D Hedgecock. Does variance in reproductive success limit effective population sizes of marine organisms? In A Beaumont, editor, *Genetics and evolution of Aquatic Organisms*, pages 1222–1344, London, 1994. Chapman and Hall.

[23] D Hedgecock and A I Pudovkin. Sweepstakes reproductive success in highly fecund marine fish and shellfish: a review and commentary. *Bull Marine Science*, 87:971–1002, 2011.

[24] J F C Kingman. The coalescent. *Stoch Proc Appl*, 13:235–248, 1982.

[25] J F C Kingman. Exchangeability and the evolution of large populations. In G Koch and F Spizzichino, editors, *Exchangeability in Probability and Statistics*, pages 97–112, Amsterdam, 1982. North-Holland.

[26] J F C Kingman. On the genealogy of large populations. *J App Probab*, 19A:27–43, 1982.

[27] V Limic and A Sturm. The spatial Λ-coalescent. *Electron J Probab*, 11:363–393, 2006.

[28] H McCombie, S Lapègue, F Cornette, C Ledu, and P Boudry. Chromosome loss in biparental progenies of tetraploid Pacific oysters *Crassostrea gigas*. *Aquaculture*, 247:97–105, 2005.

[29] M Möhle and S Sagitov. A classification of coalescent processes for haploid exchangeable population models. *Ann Probab*, 29:1547–1562, 2001.

[30] M Möhle and S Sagitov. Coalescent patterns in diploid exchangeable population models. *J Math Biol*, 47:337–352, 2003.

[31] J Pitman. Coalescents with multiple collisions. *Ann Probab*, 27:1870–1902, 1999.

[32] S Sagitov. The general coalescent with asynchronous mergers of ancestral lines. *J Appl Probab*, 36:1116–1125, 1999.

[33] S Sagitov. Convergence to the coalescent with simultaneous mergers. *J Appl Probab*, 40:839–854, 2003.

[34] O Sargsyan and J Wakeley. A coalescent process with simultaneous multiple mergers for approximating the gene genealogies of many marine organisms. *Theor Pop Biol*, 74:104–114, 2008.

[35] J Schweinsberg. Coalescents with simultaneous multiple collisions. *Electron J Probab*, 5:1–50, 2000.

[36] J Schweinsberg. Coalescent processes obtained from supercritical Galton-Watson processes. *Stoch Proc Appl*, 106:107–139, 2003.

[37] J Schweinsberg and R Durrett. Random partitions approximating the coalescence of lineages during a selective sweep. *Ann Appl Probab*, 1591–1651, 2005.

[38] R P Stanley. *Enumerative Combinatorics*, volume 2. Cambridge University Press, New York, 2001.

[39] J Wakeley. *Coalescent theory*. Roberts & Co, 2007.

[40] G A Watterson. On the number of segregating sites in genetical models without recombination. *Theor Pop Biol*, 7:256–276, 1975.

Appendix

Proof of recursion (21) for $p^{(n)}[k,b]$ (Thm. 2)

The requirement $p^{(n)}[n,1] = 1$ is obvious (therefore $p^{(n)}[n,k] = 0$ for all $k > 1$). From now on, we consider the case $1 < k < n$. By $\mathcal{P}_A$ we denote the set of all partitions of the set

A; and $[n] := \{1, 2, \ldots, n\}$ for all $n \in \mathbb{N} := \{1, 2, \ldots\}$. Let $\Pi \equiv \{\Pi_t, t \geq 0\}$ be a $\mathcal{P}_{\mathbb{N}}$-valued exchangeable coalescent, defined on the probability space $(\Omega, \mathcal{F}, \mathbb{P})$. By $\Pi^{(n)} \equiv \left\{\Pi_t^{(n)}, t \geq 0\right\}$ we denote the projection of Π onto $\mathcal{P}_{[n]}$, which will be associated with coalescent processes started from $n \geq 2$ leaves. Define

$$\tau_k^{(n)} := \inf\left(\left\{t > 0 | \#\Pi_t^{(n)} = k\right\} \cup \{\infty\}\right), \tag{34}$$

ie. the first time the block counting process associated with $\Pi^{(n)}$ hits state k, with $\tau_n^{(n)} = 0$. Let

$$\Omega_k^{(n)} := \{\tau_k^{(n)} < \infty\}.$$

Recall that we assume that $\mathbb{P}\left(\tau_k^{(n)} < \infty\right) > 0$ for each $k \in [n]$. Thus we can define the conditional law $\mathbb{P}_k^{(n)}(\cdot) := \mathbb{P}\left(\cdot | \Omega_k^{(n)}\right)$. On the conditional probability space $\left(\Omega_k^{(n)}, \mathcal{F}|_{\Omega_k^{(n)}}, \mathbb{P}_k^{(n)}\right)$, we sample uniformly at random a block π_0 from the blocks of $\Pi^{(n)}_{\tau_k^{(n)}}$, ie. $\pi_0(\omega) \in \Pi^{(n)}_{\tau_k^{(n)}}(\omega)$ for $\omega \in \Omega_k^{(n)}$. Then we can write

$$p^{(n)}[k, b] = \mathbb{P}_k^{(n)}\left(\#\pi_0 = b\right). \tag{35}$$

We define the first jump time τ of the block-counting process

$$\tau := \inf\left\{t > 0 : \#\Pi_t^{(n)} < n\right\}.$$

Now consider some $k \leq m < n$, and partition $|\nu| \stackrel{m}{=} n$, ie. an integer partition ν of n into m elements. Define

$$\Omega_k^{(n,\nu)} := \left\{\left(\Pi_\tau^{(n)}\right)^\downarrow = \nu\right\} \cap \Omega_k^{(n)}$$

(recall that $\pi^\downarrow$ denotes the integer partition associated to $\pi \in \mathcal{P}_{[n]}$ obtained by listing the block sizes of π in decreasing order). Then it is clear that we have the decomposition

$$\Omega_k^{(n)} = \bigcup_{m=k}^{n-1} \bigcup_{|\nu| \stackrel{m}{=} n} \Omega_k^{(n,\nu)}. \tag{36}$$

We define the conditional law $\mathbb{P}_k^{(n,\nu)}(\cdot) := \mathbb{P}\left(\cdot | \Omega_k^{(n,\nu)}\right)$.

For each $\omega \in \Omega_k^{(n,\nu)}$, we define the set of blocks

$$\pi^{(\nu,0)}(\omega) := \left\{\pi_i \in \Pi_\tau^{(n)}(\omega) : \pi_i \subseteq \pi_0, i \in [m]\right\}, \qquad \omega \in \Omega_k^{(n,\nu)}. \tag{37}$$

The set of blocks $\pi^{(\nu,0)}(\omega)$ contains all blocks of the partition $\Pi_\tau^{(n)}(\omega)$ that will eventually merge into the block $\pi_0(\omega)$.

For any integer subpartition $\varrho = \langle \beta_1, \beta_2, \ldots \rangle \subset \nu = \langle \alpha_1, \alpha_2, \ldots \rangle$, we need to be able to compute the probability $\mathbb{P}_k^{(n,\nu)}\left(\left(\pi^{(\nu,0)}\right)^\downarrow = \varrho\right)$ on $\Omega_k^{(n,\nu)}$. Indeed,

$$\begin{aligned}\mathbb{P}_k^{(n,\nu)}\left(\left(\pi^{(\nu,0)}\right)^\downarrow = \varrho\right) &= \mathbb{P}_k^{(n,\nu)}\left(\left(\pi^{(\nu,0)}\right)^\downarrow = \varrho \,|\, \#\pi^{(\nu,0)} = \#\varrho\right) \cdot \mathbb{P}_k^{(n,\nu)}\left(\#\pi^{(\nu,0)} = \#\varrho\right) \\ &= \mathbb{1}_{(\#\varrho \le \#\nu - k + 1)} \cdot p^{(\#\nu)}[k, \#\varrho] \frac{\prod_i \binom{\alpha_i}{\beta_i}}{\binom{\#\nu}{\#\varrho}}.\end{aligned} \tag{38}$$

The event $\left\{\left(\pi^{(\nu,0)}\right)^\downarrow = \varrho\right\}$ states that the sizes of the blocks of $\pi^{(\nu,0)}$ are given by the integer subpartition ϱ. Equation (38) follows from the exchangeability of Π, and the form of the probability density function of the multivariate hypergeometric distribution, which applies to the event of sampling β_i blocks from α_i for each i, for a total of $\#\varrho$ out of $\#\nu$. The condition $\#\varrho \le \#\nu - k + 1$ is required since we condition on the block counting process to hit k blocks.

Given (38) we can compute $\mathbb{P}_k^{(n,\nu)}(\#\pi_0 = b)$. Indeed, the decomposition

$$\{\#\pi_0 = b\} \cap \Omega_k^{(n,\nu)} = \bigcup_{\substack{\varrho \subset \nu \\ |\varrho| = b}} \left\{\left(\pi^{(\nu,0)}\right)^\downarrow = \varrho\right\}$$

(where the union is disjoint) together with (38) gives

$$\begin{aligned}\mathbb{P}_k^{(n,\nu)}\left(\#\pi_0 = b\right) &= \sum_{\substack{\varrho \subset \nu \\ |\varrho| = b}} \mathbb{P}_k^{(n,\nu)}\left((\pi^{(\nu,0)})^{\downarrow} = \varrho\right) \\ &= \sum_{\substack{\varrho \subset \nu \\ |\varrho| = b}} \mathbb{1}_{(\#\varrho \leq \#\nu - k + 1)} \cdot p^{(\#\nu)}[k, \#\varrho] \frac{\prod_i \binom{\alpha_i}{\beta_i}}{\binom{\#\nu}{\#\varrho}} \\ &= \sum_{c=1}^{b \wedge (\#\nu - k + 1)} \sum_{\substack{\varrho \subset \nu \\ |\varrho| \overset{c}{=} b}} p^{(\#\nu)}[k, c] \frac{\prod_i \binom{\alpha_i}{\beta_i}}{\binom{\#\nu}{c}}.\end{aligned} \tag{39}$$

Given Equation (39) one can now compute $p^{(n)}[k, b]$: By the decomposition (36), we obtain

$$\begin{aligned}p^{(n)}[k, b] &= \mathbb{P}_k^{(n)}\left(\#\pi_0 = b\right) \\ &= \sum_{m=k}^{n-1} \sum_{|\nu| \overset{m}{=} n} \mathbb{P}_k^{(n)}\left(\{\#\pi_0 = b\} \cap \Omega_k^{(n,\nu)}\right) \\ &= \sum_{m=k}^{n-1} \sum_{|\nu| \overset{m}{=} n} \mathbb{P}_k^{(n)}\left(\Omega_k^{(n,\nu)}\right) \mathbb{P}_k^{(n)}\left(\#\pi_0 = b \,|\, \Omega_k^{(n,\nu)}\right) \\ &= \sum_{m=k}^{n-1} \sum_{|\nu| \overset{m}{=} n} \mathbb{P}_k^{(n)}\left(\Omega_k^{(n,\nu)}\right) \mathbb{P}_k^{(n,\nu)}\left(\#\pi_0 = b\right).\end{aligned} \tag{40}$$

We apply Lemma 5 to $\mathbb{P}_k^{(n)}\left(\Omega_k^{(n,\nu)}\right)$ and Equation (39) to $\mathbb{P}_k^{(n,\nu)}\left(\#\pi_0 = b\right)$ to obtain

$$p^{(n)}[k, b] = \sum_{m=k}^{n-1} \sum_{|\nu| \overset{m}{=} n} p_\nu^{(n)} \frac{g(m,k)}{g(n,k)} \sum_{c=1}^{b \wedge (m-k+1)} \sum_{\substack{\varrho \subset \nu \\ |\varrho| \overset{c}{=} b}} p^{(m)}[k, c] \frac{\prod_i \binom{\alpha_i}{\beta_i}}{\binom{m}{c}}.$$

1 □

Proof of computation of $\mathbb{P}_k^{(n)}\left(\Omega_k^{(n,\nu)}\right)$ (Lemma 5)

Let $\Omega_k^{(n,\nu)}$, $\mathbb{P}_k^{(n)}(\cdot)$, $\tau_k^{(n)}$ (34) and τ be as defined in the proof for Thm. 2. Recall further that $p_\nu^{(n)} := \mathbb{P}\left((\Pi_\tau^{(n)})^\downarrow = \nu\right)$, see Equation (20).

By $\left\{Y_t^{(n)}; t \geq 0\right\}$, $Y_t^{(n)} := \#\Pi_t^{(n)}$ we denote the block-counting process associated with $\Pi^{(n)}$, starting from n. Let $g(n,k)$ denote the expected length of time that $Y_t^{(n)}$ spends in state $k \leq n$,

$$g(n,k) := \mathbb{E}\left[\int_0^\infty \mathbb{1}_{\left(Y_s^{(n)}=k\right)} ds\right]. \tag{41}$$

Clearly, $g(k,k) = \frac{1}{\lambda_k} = -q_{\pi,\pi}$ where $\#\pi = k$, see (11).

Lemma 5. *For $k \leq m < n$ and $|\nu| \overset{m}{=} n$, assuming $\mathbb{P}\left(\tau_k^{(n)} < \infty\right) > 0$, we have*

$$\mathbb{P}_k^{(n)}\left(\Omega_k^{(n,\nu)}\right) \equiv \mathbb{P}_k^{(n)}\left((\Pi_\tau^{(n)})^\downarrow = \nu\right) = p_\nu^{(n)} \frac{g(m,k)}{g(n,k)}, \tag{42}$$

where $g(\cdot,\cdot)$ is defined in (41).

Proof. Assuming $\mathbb{P}\left(\tau_k^{(n)} < \infty\right) > 0$, one obtains

$$\begin{aligned}
\mathbb{P}_k^{(n)}\left((\Pi_\tau^{(n)})^\downarrow = \nu\right) &= \frac{\mathbb{P}\left(\left\{\left(\Pi_\tau^{(n)}\right)^\downarrow = \nu\right\} \cap \left\{\tau_k^{(n)} < \infty\right\}\right)}{\mathbb{P}\left(\tau_k^{(n)} < \infty\right)} \\
&= \mathbb{P}\left((\Pi_\tau^{(n)})^\downarrow = \nu\right) \frac{\mathbb{P}\left(\tau_k^{(n)} < \infty \middle| \left(\Pi_\tau^{(n)}\right)^\downarrow = \nu\right)}{\mathbb{P}\left(\tau_k^{(n)} < \infty\right)} \\
&= p_\nu^{(n)} \frac{\mathbb{P}\left(\tau_k^{(m)} < \infty\right)}{\mathbb{P}\left(\tau_k^{(n)} < \infty\right)},
\end{aligned} \tag{43}$$

where we use $\#\nu = m$ and apply the strong Markov property of the process $\left(\Pi^{(n)}\right)^\downarrow$ at time τ to obtain the last equality.

Now we consider $\frac{g(m,k)}{g(n,k)}$, and obtain

$$\begin{aligned}\frac{g(m,k)}{g(n,k)} &= \frac{\mathbb{E}\left[\int_0^\infty \mathbb{1}_{\left(Y_s^{(m)}=k\right)} ds \mid \tau_k^{(m)} < \infty\right] \mathbb{P}\left(\tau_k^{(m)} < \infty\right)}{\mathbb{E}\left[\int_0^\infty \mathbb{1}_{\left(Y_s^{(n)}=k\right)} ds \mid \tau_k^{(n)} < \infty\right] \mathbb{P}\left(\tau_k^{(n)} < \infty\right)} \\ &= \frac{\mathbb{E}\left[\int_{\tau_k^{(m)}}^\infty \mathbb{1}_{\left(Y_s^{(m)}=k\right)} ds \mid \tau_k^{(m)} < \infty\right] \mathbb{P}\left(\tau_k^{(m)} < \infty\right)}{\mathbb{E}\left[\int_{\tau_k^{(n)}}^\infty \mathbb{1}_{\left(Y_s^{(n)}=k\right)} ds \mid \tau_k^{(n)} < \infty\right] \mathbb{P}\left(\tau_k^{(n)} < \infty\right)} \\ &= \frac{g(k,k)}{g(k,k)} \frac{\mathbb{P}\left(\tau_k^{(m)} < \infty\right)}{\mathbb{P}\left(\tau_k^{(n)} < \infty\right)},\end{aligned} \tag{44}$$

where we use the strong Markov property of the block-counting process in the last step. Now the statement follows from (43) and (44).

□

Proof of recursion (18) for $g(n,k)$ (Lemma 1)

Proof. For $m \geq 2$, let $\lambda_m := -q_{\pi,\pi}$ with $q_{\pi,\pi}$ given by (11), and $\#\pi = m$. Again write $Y_t^{(m)} := \#\Pi_t^{(m)}$ for the block-counting process starting from m.

Let $\tau := \inf\left\{t > 0 : Y_t^{(m)} < m\right\}$ denote the first jump time of the block-counting process, and let $p_{m,k} := \mathbb{P}\left(Y_\tau^{(m)} = k\right)$, $k \leq m-1$. For $n = m$, one obtains

$$\begin{aligned}g(m,m) &= \mathbb{E}\left[\int_0^\infty \mathbb{1}_{\left(Y_s^{(m)}=m\right)} ds\right] \\ &= \mathbb{E}\left[\int_0^\tau \mathbb{1}_{\left(Y_s^{(m)}=m\right)} ds\right] + 0 \\ &= \mathbb{E}[\tau] \\ &= \lambda_m^{-1},\end{aligned}$$

and (19) is established.

For $n < m$, we decompose according to the value of the block-counting process after the

first jump and obtain

$$\begin{aligned} g(m,n) &= \mathbb{E}\left[\int_0^\infty \mathbb{1}_{\left(Y_s^{(m)}=n\right)} ds\right] \\ &= \sum_{k=n}^{m-1} \mathbb{E}\left[\int_0^\infty \mathbb{1}_{\left(Y_s^{(m)}=n\right)} ds \middle| Y_\tau^{(m)} = k\right] \mathbb{P}\left(Y_\tau^{(m)} = k\right) \\ &= \sum_{k=n}^{m-1} \mathbb{E}\left[\int_\tau^\infty \mathbb{1}_{\left(Y_s^{(m)}=n\right)} ds \middle| Y_\tau^{(m)} = k\right] p_{m,k} \\ &= \sum_{k=n}^{m-1} \mathbb{E}\left[\int_0^\infty \mathbb{1}_{\left(Y_s^{(k)}=n\right)} ds\right] p_{m,k} \\ &= \sum_{k=n}^{m-1} p_{m,k} g(k,n), \end{aligned}$$

where we use the strong Markov property of Y in the second-to-last equality. Thus (18) is established.

Proof of Equation (23)

Fix $n \geq 2$, $\underline{k} := (k_1, \ldots, k_r)$ with $k_1 \geq \ldots \geq k_r \geq 2$, $r \in [M]$, and let $|\underline{k}| := k_1 + \cdots + k_r$, $s := n - |\underline{k}| = n - k_1 - \cdots - k_r$. Then we have to prove that

$$\lambda_{n;\underline{k}} = \sum_{\ell=0}^{s \wedge (M-r)} \binom{s}{\ell} (M)_{r+\ell} M^{-(|\underline{k}|+\ell)} \int_{[0,1]} x^{|\underline{k}|+\ell-2} (1-x)^{s-\ell} F(dx).$$

Let Δ denote the infinite simplex, and recall the function f from (6) used to describe the rates (11) of a Xi-coalescent. We define the map $h_{(n,\underline{k})} : \Delta \ni \xi \mapsto f(\xi, n, \underline{k}) \in \mathbb{R}$. We then have that the rate of a $\underline{k}$-merger is given by

$$\lambda_{n,\underline{k}} = \int_{\xi \in \Delta} h_{(n,\underline{k})}(\xi)\, \Xi(d\xi). \tag{45}$$

Since $\mathrm{supp}(\Xi) \subseteq \Delta_M := \{(\underbrace{\frac{x}{M}, \ldots, \frac{x}{M}}_{M \text{ times}}, 0, 0, \ldots) \in \Delta : x \in [0,1]\}$, we can rewrite the rate in (45) as

$$\lambda_{n,\underline{k}} = \frac{1}{M} \int_{x\in[0,1]} h_{(n,\underline{k})}(\underbrace{\tfrac{x}{M}, \ldots, \tfrac{x}{M}}_{M \text{ times}}, 0, 0, \ldots) F(dx). \tag{46}$$

Define $x_i := \frac{x}{M} \mathbb{1}_{(i \le M)}$. Then the integrand in (46) is given by

$$\frac{\sum_{\ell=0}^{s} \sum_{i_1 \neq \ldots \neq i_{r+\ell}} \binom{s}{\ell} x_{i_1}^{k_1} \cdots x_{i_r}^{k_r} x_{i_{r+1}} \cdots x_{i_{r+\ell}} \left(1 - \sum_j x_j\right)^{s-\ell}}{\sum_j x_j^2},$$

which is zero if $r + \ell > M$. For $r + \ell \le M$, one obtains

$$x_{i_1}^{k_1} \cdots x_{i_r}^{k_r} x_{i_{r+1}} \cdots x_{i_{r+\ell}} = \frac{x^{|\underline{k}|+\ell}}{M^{|\underline{k}|+\ell}}$$

for any choice of indices $i_1, \ldots, i_{r+\ell}$ which are all different and smaller than M. Therefore, with $(a)_m := a(a-1)\cdots(a-m+1)$ for $m \in \mathbb{N}$, $(a)_0 := 1$ denoting the falling factorial, we have

$$\sum_{i_1 \neq \ldots \neq i_{r+\ell}} x_{i_1}^{k_1} \cdots x_{i_r}^{k_r} x_{i_{r+1}} \cdots x_{i_{r+\ell}} = \mathbb{1}_{(r+\ell \le M)} \frac{(M)_{r+\ell}}{M^{|\underline{k}|+\ell}} x^{|\underline{k}|+\ell}.$$

We note further that

$$1 - \sum_j x_j = 1 - x, \qquad \sum_j x_j^2 = M \cdot \left(\frac{x}{M}\right)^2 = \frac{x^2}{M}.$$

Thus we get

$$h_{(n,\underline{k})}(\underbrace{\tfrac{x}{M}, \ldots, \tfrac{x}{M}}_{M \text{ times}}, 0, 0, \ldots) = \sum_{\ell=0}^{s \wedge (M-r)} \binom{s}{\ell} (1-x)^{s-\ell} x^{|\underline{k}|+\ell-2} (M)_{r+\ell} M^{1-|\underline{k}|-\ell},$$

and we can represent the rate $\lambda_{n,\underline{k}}$ as

$$\begin{aligned}\lambda_{n,\underline{k}} &= \frac{1}{M}\int_{x\in[0,1]} h_{(n,\underline{k})}(\underbrace{\tfrac{x}{M},\ldots,\tfrac{x}{M}}_{M \text{ times}},0,0,\ldots)F(dx)\\ &= \sum_{\ell=0}^{s\wedge(M-r)}\binom{s}{\ell}(M)_{r+\ell}M^{-(|\underline{k}|+\ell)}\int_{x\in[0,1]} x^{|\underline{k}|+\ell-2}(1-x)^{s-\ell}F(dx).\end{aligned}$$

This completes our proof. □