arXiv:1902.10834v2 [cs.NE] 24 Aug 2020

An evolutionary model that satisfies detailed balance

Jüri Lember*, Chris Watkins†

August 25, 2020

Abstract

We propose a class of evolution models that involves an arbitrary exchangeable process as the breeding process and different selection schemes. In those models, a new genome is born according to the breeding process, and after that a genome is removed according to the selection scheme that involves fitness. Thus, the population size remains constant. The process evolves according to a Markov chain, and, unlike in many other existing models, the stationary distribution – so called mutation-selection equilibrium – can easily found and studied. As a special case our model contains a (sub)class of Moran models. The behaviour of the stationary distribution when the population size increases is our main object of interest. Several phase-transition theorems are proved.

Keywords: Markov chain Monte Carlo, Dirichlet distribution, de Finetti theorem, weak convergence of probability measures, Moran model.

AMS classification: 60J10, 60B10

1 Introduction

We introduce a probability model of evolution that has three purposes: as an abstract model of biological evolution; as a class of efficiently implementable genetic algorithms that are easy to analyze theoretically; and as a link between genetic algorithms and Bayesian non-parametric MCMC methods. The stationary distribution of the model can be expressed in closed form for arbitrary fitness functions: this enables us to investigate the behaviour of the model for different population sizes, mutation rates, and fitness scalings. We find two phase transitions that occur for all fitness functions. Our approach is most applicable to evolution by sexual rather than asexual reproduction.

In any model of evolution under constant conditions, there is a temporal sequence of possibly overlapping populations. In the transition from each population to the next, one

*Institute of Mathematics and Statistics, University of Tartu, Estonia; jyril@ut.ee; +372 7375487
†Department of Computer Science, Royal Holloway, University of London, UK

or more new individuals are 'born', and the same number of individuals are removed from the population, or 'die'. Each new generation is a population that depends only upon the previous parent population, so that the sequence of populations is a Markov chain. In a model with mutation, the Markov chain is irreducible and has a unique stationary distribution that is also known as the mutation-selection equilibrium. Typically we wish to know what the stationary distribution is.

For Moran models with only a single locus, the stationary distribution can be characterised. But for models with more than one locus, and fitness that depends nonlinearly on the combination of alleles at the loci, the stationary distribution is notoriously hard to characterise. The reason is that in previous models of sexual reproduction with multiple loci, the Markov chain of populations is irreversible, so that there is no obvious method of finding the stationary distribution other than attempting to compute eigenvectors of the transition matrix directly, as done in [25], but these calculations are neither easy nor revealing for arbitrary fitness functions on multiple loci.

For a reversible Markov chain, the stationary distribution may be found by verifying detailed balance conditions. In our model, introduced in [26], we start by writing the mutation-selection equilibrium in a convenient closed form, and then exhibit MCMC kernels that implement reversible Markov chains with this stationary distribution, and for which the proposal and acceptance algorithms are plausible abstract models of breeding, mutation, and selection. Each generation starts with a population of n individuals. One new individual is 'bred' – that is, sampled conditionally on the existing population – to produce an expanded population of $n+1$ individuals; from this expanded population, one individual is selected to be discarded, leaving a new population of size n to start the next generation. This kind of evolution is typically modeled with a Moran model [18] and in subsections 2.1 and 3.4, we shall see that when the breeding process is a generalised Polya urn model, then our model can be indeed considered as a discrete time Moran model with mutation and selection. But our model is more general that just a Moran one, because we allow a more general breeding process. However, due to the connection with Moran models, the generalized Polya urn model (or, equivalently, Dirichlet-categorical process) is an important special case of breeding so that section 6 is devoted to that model.

In many evolutionary models such as [5], and in genetic algorithms such as [15], the reproductive fitness of a genome is modelled as the genome's rate of breeding, and not as its probability of death. In these models, in the breeding phase of each generation, fit genomes are chosen more frequently to breed than unfit genomes are: in these models, discarding individuals – or 'death' – is modelled as uniformly random deletion from the population. The second option is that the selection occurs at death: individuals are discarded (or 'die') with probability related to their fitness (termed viability selection): less-fit individuals are more likely to be discarded. For discrete time classical models such as the Wright-Fisher or Moran model, there is actually no big difference whether the selections occur at death or at reproduction (although formulas can be slightly different, see e.g. [20]) and in the literature both versions are present. In our model, breeding is conditional sampling from a very general exchangeable process and so it is quite natural that selection occurs at death. This is a significant design choice that turns out to greatly

simplify the analysis. Observe that selection at death it does not mean that all allele types have equal probability to be born. Our model might have breeding process as i.i.d. random variables with given distribution possibly very far form uniform, or the breeding process can be generalized Polya urn, where the history of evolution determines what types are more likely to be born.

As explained in detail in sections 2 and 3, careful choices in modelling breeding as conditional sampling, and in modelling fitness by stochastic rejection of the less fit, enable us to construct a reversible Markov chain of populations that is a Metropolis-Hastings process.

Mathematical formulation, overview of the main results and organization of the paper. Throughout the paper, we suppose there is a finite set $\mathcal{X} = \{1, \ldots, K\}$ of possible genomes. Our models apply both to the case of genomes with one locus with K possible alleles, and also to the case where genomes have more than one locus: for example, for genomes with L loci with two possible alleles at each, then $K = 2^L$. Throughout the paper, $\xi_1, \xi_2, \ldots$ denotes an exchangeable $\mathcal{X}$-valued stochastic process. For any population of n genomes $\mathbf{x} = (x_1, \ldots, x_n) \in \mathcal{X}^n$, we define

$$P_\xi(\mathbf{x}) := P_\xi(x_1, \ldots, x_n) := P(\xi_1 = x_1, \ldots, \xi_n = x_n) \tag{1.1}$$

By definition of exchangeability, we have that for any permutation σ,

$$P_\xi(x_1, \ldots, x_n) = P_\xi(x_{\sigma(1)}, \ldots, x_{\sigma(n)}).$$

Given a population of genomes $\mathbf{x} = (x_1, \ldots, x_n)$, we breed an $n+1$'th genome by sampling x_{n+1} conditionally:

$$x_{n+1} \sim P_\xi(\cdot \mid x_1, \ldots, x_n).$$

The 'fitness' of a genome x is denoted by $w(x) > 0$, where w is an arbitrary strictly positive function over $\mathcal{X}$. In the evolutionary models we propose, in each generation one individual is 'bred' by conditional sampling from the existing population, and then an individual is discarded in a fitness-biased way, so that less-fit individuals are more likely to be discarded. As explained in subsections 3.1, 3.2 and 3.3, there are many possible schemes to formalize this idea, and we shall show that all those schemes lead to one particular stationary distribution of populations that factorises into the form:

$$\underbrace{P_n(x_1, \ldots, x_n)}_{\text{stationary distribution}} = \frac{1}{Z_n} \underbrace{P_\xi(x_1, \ldots, x_n)}_{\text{breeding term}} \underbrace{w(x_1) \cdots w(x_n)}_{\text{fitness term}}. \tag{1.2}$$

The process is similar to, but not the same as non-parametric Bayesian MCMC, and we make this connection explicit in section 3.6.

The measure P_n is the main object of interest. Since ξ is exchangeable, there exists a prior measure π on the set of all probability measures on $\mathcal{X}$ (called simplex), denoted by

$$\mathcal{P} := \{(q_1, \ldots, q_k) : q_i \geq 0, \sum_i q_i = 1\},$$

such that for every $\mathbf{x} \in \mathcal{X}^n$

$$P_\xi(\mathbf{x}) = \int_{\mathcal{P}} \prod_{i=1}^{n} q(x_i)\pi(dq). \tag{1.3}$$

Here $q(k) = q_k$, in what follows, we shall use both notations (q_k and $q(k)$). Hence P_n is fully determined by prior measure π and fitness function w. In order to achieve more generality, we allow w and π to depend on population size n (thus writing w_n and π_n), and we ask: do the measures P_n converge (in some sense) to a limit? The sense of convergence needs to be specified, because P_n is defined on $\mathcal{X}^n$, so that the domain of the measure depends on n. Therefore we consider the array of random variables $(X_{n,1}, \ldots, X_{n,n}) \sim P_n$ and we ask: is there a stochastic process $X_1, X_2, \ldots$ so that for every m

$$(X_{1,n}, \ldots X_{m,n}) \Rightarrow (X_1, \ldots X_m)?$$

Another option to define the limit is based on the observation that the measure P_n is invariant with respect to the permutations (finitely exchangeable), and so it depends on the counts of vector $\mathbf{x}$, only. This, in turn, allows to define a counterpart of P_n – a probability measure Q_n on the simplex $\mathcal{P}$. The formal definition of Q_n is given in subsection 4.1. Since all measures Q_n are defined on the same domain (simplex $\mathcal{P}$), we can now ask whether the sequence Q_n converges in the classical sense of weak convergence of probability measures. Throughout the paper, these two approaches – measures P_n (random process) and measures Q_n – are handled in parallel.

Our first limit result, Theorem 5.1 considers the case when the prior measure π is independent of n, but arbitrary; the fitness function w_n depends on n in the following way (recall $\mathcal{X} = \{1, \ldots, K\}$):

$$w_n(k) = \exp\left[-\frac{\phi(k)}{n^\lambda}\right], \quad k = 1, \ldots, K \tag{1.4}$$

where $\lambda \geq 0$ and $0 \leq \phi(1) \leq \phi(2) \leq \cdots \leq \phi(K)$. The case $\lambda = 0$ corresponds to fitness that is constant in that it does not vary with n. Observe that for any n, the genotype 1 is the fittest. Theorem 5.1 shows that the following phase transition occurs with respect to λ:

1. when $\lambda \in [0, 1)$ and $\phi(1) < \phi(2)$, then the limit process $X_1, X_2, \ldots$ has one trajectory $1, 1, 1, \ldots$ a.s., and $Q_n \Rightarrow \delta_{q^*}$, where $q^* = (1, \ldots, 0)$. Thus, when $\lambda \in [0, 1)$ only the fittest genotype survives in the limit, no matter what the prior says; the influence of the prior vanishes and fitness prevails. The additional assumption $\phi(1) < \phi(2)$ means that $w_n(1) > w_n(k)$, $k = 2, \ldots, K$, and so it is quite natural for the result.

2. When $\lambda = 1$, Q_n converges to a nondegenerate distribution specified in Theorem 5.1 that depends on the prior π and on the function ϕ. When the prior π has density $\pi(q)$, then the limit distribution also has density proportional to

$$\pi(q) \exp[-\langle \phi, q \rangle]. \tag{1.5}$$

Then also the limit process $X_1, X_2, \ldots$ is exchangeable process with prior (1.5). non-degenerate.

3. When $\lambda > 1$, then the influence of fitness vanishes and only the prior matters – the limit process is equal to the breeding process and $Q_n \Rightarrow \pi$.

Section 6 considers Dirichlet priors

$$\pi_n = \mathrm{Dir}(n^{1-\lambda}\alpha_1, \ldots n^{1-\lambda}\alpha_K), \tag{1.6}$$

where $\alpha := (\alpha_1, \ldots, \alpha_K)$, $\alpha_k > 0$. Thus π_n admits a density on $\mathcal{P}$

$$\pi(q_1, \ldots, q_k) \propto \prod_{k=1}^{K} q_k^{n^{1-\lambda}\alpha_k}.$$

With a Dirichlet prior the breeding process is a (generalized) Polya urn process, also known as a Dirichlet-categorical process. This process is a natural choice because it admits a notion of 'mutation'. Conditional sampling from it is directly interpretable as a simplified model of sexual breeding with mutation, as explained in section 2. The fitness is previously defined as in (1.4) and the common parameter λ allows to interpolate between influence of fitness and prior. The case $\lambda = 1$ corresponds to the constant prior case. Our second main limit theorem, Theorem 6.1 shows that with respect to λ, phase transition occurs again:

1. when $\lambda = 0$, then the limit process $X_1, X_2, \ldots$ consists of i.i.d. random variables with certain distribution r^* (specified in Theorem 6.1) depending on ϕ and α, and $Q_n \Rightarrow \delta_{r^*}$.

2. when $\lambda \in (0, 1)$, then then the limit process $X_1, X_2, \ldots$ consists of i.i.d. random variables with another distribution q^* (specified in Theorem 6.1) also depending on ϕ and α, and $Q_n \Rightarrow \delta_{q^*}$;

3. when $\lambda = 1$ and $\phi(1) < \phi(2)$, then the limit is specified by Theorem 5.1. From (1.5) it follows that the limit distribution of allele proportions has density proportional to

$$\exp[-\langle \phi, q \rangle] \prod_{k=1}^{K} q_k^{\alpha_k - 1}. \tag{1.7}$$

The limits of type (1.7) are common for standard models in evolution theory, both for Wright-Fisher and Moran model (see, e.g. [6, 9, 10]). We shall discuss the connection with our and Moran model in subsections 2.1 and 3.4.

The phase transitions at $\lambda = 0$ and $\lambda = 1$ are sharp. These are the most important results of the paper.

The paper is organized as follows. In section 2 we give examples of exchangeable conditional sampling procedures that can be regarded as abstract models of breeding, and

section 3 gives examples of selection procedures that, together with any of the breeding procedures, will produce a reversible Markov chain of populations with the stationary distribution of (1.2). Subsection 3.5 establishes that the stationary distributions are invariant to multiplicative fitness noise; subsection 3.6 shows that special cases of this process are forms of Bayesian inference by MCMC. Next, section 4 establishes basic conditions on the convergence of the stationary distribution to a limit distribution as population size tends to infinity. Section 5 is devoted to Theorem 5.1 and in section 6 generalized Polya urn processes (Dirichlet priors) are considered. In this section, Theorem 6.1 is proved and the product of independent Polya urn processes are considered (subsection 6.2). We present computational experiments demonstrating our results in section 7. Finally we discuss implications of our results for genetic algorithms and evolutionary modelling in 8.

2 Breeding and mutation

We model breeding as Gibbs sampling [11] from an exchangeable distribution; exchangeable Gibbs sampling is a standard technique in statistical nonparametric MCMC methods, for example [21, 14], but in that context it is not of course regarded as a model of breeding. The property of exchangeability of P_ξ will be used in two ways: first, in section 3 we will use it to establish detailed balance for several selection procedures, which establishes that the stationary distribution of populations is indeed as given in (1.2). Second, in section 4 we use the de Finetti integral representation of ξ to establish limit properties of the stationary distribution as $n \to \infty$. We now give examples of conditional sampling that can be regarded as plausible models of breeding.

2.1 Dirichlet-Categorical Process

In the simplest case, each 'genome' consists of only one locus or 'gene' which can be one of K possible allele 'types', that we denote by $\{1, \ldots, K\}$. The exchangeable process ξ is the well known Polya urn model for a Dirichlet process with discrete base distribution. We recall the definition of this process. Let $\alpha = (\alpha_1, \ldots, \alpha_K)$, $\alpha_i > 0$ be the prior parameters of the base distribution; we write $|\alpha| := \alpha_1 + \cdots + \alpha_K$. Let $\xi_1, \xi_2, \ldots$ be a random process over $\{1, \ldots, K\}$. Let $n \geq 0$. Given $\mathbf{x} = x_1, \ldots, x_n$, we denote the number of k-s in the sequence by $n_k(\mathbf{x})$:

$$n_k = \sum_{i=1}^{n} I_{\{k\}}(x_i). \tag{2.1}$$

From the Polya urn, the next allele is drawn of type k with the following probability:

$$P_\xi(\xi_{n+1} = k \mid x_1, \ldots, x_n, \alpha) := \frac{n_k + \alpha_k}{n + |\alpha|}.$$

It follows that

$$P_\xi(x_1, \ldots, x_n \mid \alpha) = \frac{(\alpha_1)_{n_1} \cdots (\alpha_K)_{n_K}}{(|\alpha|)_n},$$

where for any $\alpha > 0$

$$(\alpha)_n := \alpha(\alpha+1)\cdots(\alpha+n-1).$$

The process $\xi_1, \xi_2, \ldots$ is infinitely exchangeable by inspection. Due to the exchangeability, by the formula above, it is clear that the random vector of counts has the distribution

$$P_\xi(n_1, \ldots, n_K) = \frac{n!}{n_1! \cdots n_K!} \frac{(\alpha_1)_{n_1} \cdots (\alpha_K)_{n_K}}{(|\alpha|)_n}. \tag{2.2}$$

By de Finetti's theorem

$$P_\xi(x_1, \ldots, x_n \mid \alpha) := \int_{\mathcal{P}} q(x_1) \cdots q(x_n) \pi(dq),$$

where the prior measure π is Dirichlet distribution, i.e. $\pi = \mathrm{Dir}(\alpha_1, \ldots, \alpha_K)$. This process is in a sense the central object of our study.

Concentration parameter, mutation rate and Moran models. The concentration parameter $|\alpha| = \alpha_1 + \cdots + \alpha_K$ may be viewed as determining a mutation rate that depends on n. With n balls in the sample (thus $n + |\alpha|$ balls in the urn) when a new ball is sampled, there is a probability $\frac{|\alpha|}{n+|\alpha|}$ that the ball will be sampled from the collection of 'prior' balls, rather than the actual balls in the urn. This probability is independent of the colors of the n actual balls present: since a new color may be introduced in this way, we regard this as analogous to a mutation. The mutation rate u is

$$u := \frac{|\alpha|}{|\alpha| + n} \quad \text{or equivalently} \quad |\alpha| = \frac{nu}{1-u}$$

To be more detailed, for any type $k = 1, \ldots, K$, the 'prior' ball of type k is sampled with probability

$$u_k := \frac{\alpha_k}{|\alpha| + n} \quad \text{or equivalently} \quad \alpha_i = \frac{nu_i}{1-u}. \tag{2.3}$$

Hence $u_1 + \cdots + u_K = u$ and u_k is the rate (probability) of mutating from any other type to k. In our evolutionary processes, one new genome is sampled, and one genome is then discarded at each generation, so that n remains constant. To define processes with the same mutation rate u but with different values of n, α must be adjusted to depend on n. Such single birth-death changes are often modeled via the well known Moran model, and it turns out that our model of breeding in some sense generalizes the drift-mutation Moran model. Let us explain that connection more closely. Recall that a standard (simple) K-allele Moran model without mutation and selection is as follows. We consider a population of n individuals and suppose that at time t, the number of type k individuals is n_k. An individual is randomly chosen to reproduce and another is independently chosen to die. When the type i is chosen to reproduce and type j is chosen to die, then the allele counts at time $t+1$ are $(n_1, \ldots, n_i + 1, \ldots, n_j - 1, \ldots n_K)$ and the probability of such a change

is $\frac{n_i n_j}{n^2}$. There are several ways to incorporate mutation into the Moran model, but in our model the probability that type i is born is

$$\frac{n_i + \alpha_i}{|\alpha| + n} = \frac{n_i}{n}\frac{n}{|\alpha| + n} + \frac{\alpha_i}{|\alpha| + n_i} = \frac{n_i}{n}(1-u) + u_i \tag{2.4}$$

In the sum above, the first part is the transition probability without mutation multiplied with the probability of no mutation, the second part is the probability that the selected type mutates to type i. When the probability of j-type to die is just $\frac{n_j}{n}$ we end up with a Moran model, where the probability that n_i becomes $n_i + 1$ and n_j becomes $n_j - 1$ is

$$\frac{n_i n_j}{n^2}(1-u) + \frac{n_j}{n}u_i. \tag{2.5}$$

It is easy to verify that a Markov chain with transition (2.5) satisfies the detailed balance equation with stationary distribution being (2.2), where α_k is defined as in (2.3). There are surely other ways to specify mutation in Moran model so that (2.2) is the stationary distribution. For $K = 2$, one such example is in [9]: let u_1 be the probability that type 2 mutates to 1 and u_2 be the probability that type 1 mutates to 2. Then the transition from (n_1, n_2) to $(n_1 + 1, n_2 - 1)$ has probability

$$\frac{n_1 n_2}{n^2}(1-u_2) + \frac{n_2^2}{n^2}u_1, \tag{2.6}$$

and detailed balance equation with (2.2) (with α_k being defined as in (2.3)) is easy to verify (this is (3.58) in [9]). Other examples of Moran models without selection having (2.2) as stationary distribution can be found in [24], where the authors consider a discrete time version of so-called two-allele decoupled Moran model introduced in [4] (see also [22]). In their model, the transition probability from n_i and n_j to $n_i + 1$ and $n_j - 1$ without selection is very close to (2.5):

$$\frac{n_i n_j}{n^2} + u_i \frac{n_j}{n} \tag{2.7}$$

and the the detailed balance equation with (2.2) is easy to to verify. In this case, $\alpha_i = nu_i$. A K-allele version of that model with transition matrix as in (2.7) is studied in [8] and the stationary distribution (14) in [8] is exactly (2.2). Hence our breeding model largely incorporates mutation-drift version of Moran models, but it is much general, because the Dirichlet distribution is only one possible choice of prior measure π in (1.3).

Note that in our model, mutations only occur at birth, and – importantly – the distribution of mutations does not depend on the frequencies of different colors that are currently in the population; mutations are distributed according to a fixed prior distribution determined by the frequencies of 'prior' balls in the urn.

In our model, we count as a mutation any ball which results from a draw of a 'prior' ball: this definition is consistent with an extension of our model to Dirichlet processes with continuous base distributions, which we intend to consider in future work.

2.2 Complex genomes: direct product of Dirichlet processes

More complex evolutionary models such as genetic algorithms require more complex genomes. Suppose that each genome is a vector of L genes, $x_i = (y_i^1, \ldots, y_i^L)$. Let ξ be a direct product of L independent exchangeable processes $\xi = (\xi^1, \ldots, \xi^L)$. Then

$$P_\xi(x_1, \ldots, x_n) := \prod_{j=1}^{L} P_{\xi^j}(y_1^j, \ldots, y_n^j) \tag{2.8}$$

which clearly is exchangeable as well. Exchangeable sampling from $P_\xi(\cdot \mid x_1, \ldots, x_n)$, where each x_i is a vector of discrete values, can be viewed as a model of sexual reproduction with the assumption of linkage equilibrium. A new vector $(y_{n+1}^1, \ldots, y_{n+1}^L)$ is sampled by, for $1 \leq j \leq L$, sampling the j-th component $y_{n+1}^j \sim P_{\xi^j}(\cdot \mid y_1^j, \ldots, y_n^j)$ independently from the rest of the components. In words, each new element of the vector x_{n+1} is either a copy of the corresponding element of a randomly chosen member of the existing population, or else a mutation. Instead of a new 'child' genome being constructed by random recombination of two parent genomes, it is instead a random recombination of all n existing genomes in the population, with mutations.

This method of constructing new genomes by 'n-way recombination' is a widely used approach in genetic algorithms, as used by [1, 2, 3] and others, and sexual reproduction with full linkage equilibrium is a standard simplified model of sexual reproduction in population genetics theory [5, 9].

The extension to Cartesian products of Dirichlet processes might appear rather simple because each component of a new genome is sampled independently of the others; however, this extension can lead to models of great complexity because the fitness function w, or equivalently ϕ, can be an arbitrary function on $\mathcal{X}^L$, so that the stationary distribution need not be a product distribution. In genetic language, the fitness function can have arbitrary epistasis.

Note that there are other discrete exchangeable distributions based on Dirichlet distributions that could also be used as P_ξ. A notable example is the discrete fragmentation-coagulation sequence process introduced in [7]; this was intended as statistical model for imputing phasing in genetic analysis, but it could also be used as a breeding distribution for our purposes.

3 Selection

Several MCMC sampling methods give the factorized stationary distribution given by equation (1.2) and at the same time are models of sexual reproduction that are as plausible as those used in evolutionary computation or simplified models in population genetics.

We suppose that each element $x \in \mathcal{X}$ has a strictly positive weight $w(x)$. In context, for brevity we will denote the weights $w(x_1), \ldots, w(x_n)$ as $w_1, \ldots, w_n$.

3.1 Single tournament selection

We suppose that when a new genome x_{n+1} is 'born' and added to the population, it competes to survive by having a tournament with another randomly selected member of the population, x_i say. The probability that x_{n+1} wins the tournament and ejects x_i from the population is $\frac{w_{n+1}}{w_{n+1}+w_i}$. This probability is always well defined since w is strictly positive. An equally valid tournament winning probability is that x_{n+1} wins with probability $\min\{1, \frac{w_{n+1}}{w_i}\}$. These two tournament winning probabilities are simply different formulations of the Metropolis-Hastings acceptance rule; the proof below establishes detailed balance for the first winning rule. The algorithm for performing one generation of breeding, mutation, and selection is:

1. Sample $x_{n+1} \sim P_\xi(\cdot \mid x_1, \ldots, x_n)$
2. Sample i randomly from $\{1, \ldots, n\}$
3. With probability $\frac{w_{n+1}}{w_i+w_{n+1}}$ replace x_i with x_{n+1} and discard x_i, otherwise discard x_{n+1}.

Let $\mathbf{x}$, $\mathbf{x}'$ be populations defined as

$$\mathbf{x} = x_1, \ldots, x_n \quad \mathbf{x}' = x_1, \ldots, x_{i-1}, x_{n+1}, x_{i+1}, \ldots, x_n$$

Recall the measure $P_n(\mathbf{x})$ defined in (1.2). We now show that this measure satisfies detailed balance. By exchangeability of ξ, we have:

$$P_\xi(\mathbf{x})P_\xi(x_{n+1} \mid \mathbf{x}) = P_\xi(x_1, \ldots, x_{n+1}) = P_\xi(\mathbf{x}')P_\xi(x_i \mid \mathbf{x}') \tag{3.1}$$

Note that x_{n+1} has tournament against x_i with probability $\frac{1}{n}$ and wins with probability $\frac{w_{n+1}}{w_i+w_{n+1}}$, so:

$$\begin{aligned}
P_n(\mathbf{x})P(\mathbf{x} \to \mathbf{x}') &= P_n(\mathbf{x}) \cdot P_\xi(x_{n+1} \mid \mathbf{x}) \cdot \frac{1}{n}\frac{w_{n+1}}{w_i + w_{n+1}} \\
&= \frac{1}{Z_n} P_\xi(\mathbf{x}) w_1 \cdots w_n \cdot P_\xi(x_{n+1} \mid \mathbf{x}) \cdot \frac{1}{n}\frac{w_{n+1}}{w_i + w_{n+1}} \\
&= \frac{1}{Z_n} P_\xi(x_1, \ldots, x_{n+1}) \cdot \frac{1}{n}\frac{w_1 \cdots w_{n+1}}{w_i + w_{n+1}} \\
&= \frac{1}{Z_n} P_\xi(\mathbf{x}')P_\xi(x_i \mid \mathbf{x}') \cdot \frac{1}{n}\frac{w_1 \cdots w_{n+1}}{w_i + w_{n+1}} \\
&= \frac{1}{Z_n} P_\xi(\mathbf{x}') w_1 \cdots w_{i-1} w_{n+1} w_{i+1} \cdots w_n \cdot P_\xi(x_i \mid \mathbf{x}') \cdot \frac{1}{n}\frac{w_i}{w_i + w_{n+1}} \\
&= P_n(\mathbf{x}')P(\mathbf{x}' \to \mathbf{x}).
\end{aligned}$$

3.2 Inverse fitness selection: limit of many tournaments

Suppose that many tournaments are fought, and each time the loser of the previous tournament fights another randomly chosen genome from the population. After many tournaments, and at a stopping time, the current loser is ejected. The current loser evolves according to a irreducible aperiodic Markov chain, and the limiting distribution of ejection is the stationary distribution of that chain:

$$P(\text{ eject } i\ ) = \frac{\frac{1}{w_i}}{\frac{1}{w_1} + \cdots + \frac{1}{w_{n+1}}}. \tag{3.2}$$

The algorithm for performing one generation of breeding and selection is then:

1. Sample $x_{n+1} \sim P_\xi(\cdot \mid x_1, \ldots, x_n)$
2. Sample i from $\{1, \ldots, n+1\}$ with probabilities proportional to $\{\frac{1}{w_1}, \ldots, \frac{1}{w_{n+1}}\}$
3. Discard x_i

This process too satisfies detailed balance. With $\mathbf{x}$ and $\mathbf{x}'$ defined as above, note that:

$$\begin{aligned}
P_n(\mathbf{x})P(\mathbf{x} \to \mathbf{x}') &= P_n(\mathbf{x}) \cdot P_\xi(x_{n+1} \mid \mathbf{x}) \cdot \frac{\frac{1}{w_i}}{\frac{1}{w_1} + \cdots + \frac{1}{w_{n+1}}} \\
&= \frac{1}{Z_n} P_\xi(\mathbf{x}) w_1 \cdots w_n \cdot P_\xi(x_{n+1} \mid \mathbf{x}) \cdot \frac{\frac{1}{w_i}}{\frac{1}{w_1} + \cdots + \frac{1}{w_{n+1}}} \\
&= \frac{1}{Z_n} P_\xi(x_1, \ldots, x_{n+1}) \frac{w_1 \cdots w_{n+1}}{w_{n+1}} \frac{\frac{1}{w_i}}{\frac{1}{w_1} + \cdots + \frac{1}{w_{n+1}}} \\
&= \frac{1}{Z_n} P_\xi(x_1, \ldots, x_{n+1}) \frac{w_1 \cdots w_{n+1}}{w_i} \frac{\frac{1}{w_{n+1}}}{\frac{1}{w_1} + \cdots + \frac{1}{w_{n+1}}} \\
&= P_n(\mathbf{x}') \cdot P_\xi(x_i \mid \mathbf{x}') \cdot \frac{\frac{1}{w_{n+1}}}{\frac{1}{w_1} + \cdots + \frac{1}{w_{n+1}}} \\
&= P_n(\mathbf{x}')P(\mathbf{x}' \to \mathbf{x}).
\end{aligned}$$

Note in passing that

$$\mathbf{E}\{\text{weight of rejected genome} \mid w_1, \ldots, w_{n+1}\} = \frac{1}{\frac{1}{w_1} + \cdots + \frac{1}{w_{n+1}}}. \tag{3.3}$$

That is, the expected weight of the rejected genome is the harmonic mean of the weights of the genomes in the current population, including the newly added genome x_{n+1}.

The case $n = 1$ reduces to Metropolis-Hastings: with a population of size 1, we have proposal distribution $P_\xi(x' \mid x)P_\xi(x) = P_\xi(x, x') = P_\xi(x \mid x')P_\xi(x')$ and acceptance probability of x' given x of $\frac{w(x')}{w(x)+w(x')}$, which gives a stationary distribution

$$P_1(x) = \frac{P_\xi(x)w(x)}{\sum_x P_\xi(x)w(x)}.$$

3.3 Exchangeable breeding of many offspring

Another MCMC process which is also interpretable as an evolutionary algorithm breeds some arbitrary number m of offspring to give a population of size $n + m$, and then from these selects n genomes to form the next generation. The algorithm for a single generation is as follows:

1. Pick a random number m of offspring to breed, and a random number t of tournaments to conduct. Both m and t should be independent of the current population $x_1, \ldots, x_n$
2. Breed $x_{n+1}, \ldots, x_{n+m}$ by sequential exchangeable sampling; that is, let $x_{n+i} \sim P_\xi(\cdot \mid x_1, \ldots, x_{n+i-1})$ for $1 \leq i \leq m$
3. Assign 'survival tickets' to each of $x_1, \ldots, x_n$; the newly bred genomes $x_{n+1}, \ldots, x_{n+m}$ have as yet no survival tickets.
4. Repeat t times:
   (a) Uniformly sample from the population a genome x_i which currently has a survival ticket, and a genome x_j which currently does not have a survival ticket.
   (b) Hold a tournament between x_i and x_j; x_i wins with probability $\frac{w(x_i)}{w(x_i)+w(x_j)}$.
   (c) The winner of the tournament gets the survival ticket; after the tournament, the winner has the ticket and the loser does not.
5. After t tournaments have taken place, the n genomes currently holding survival tickets are selected to be the new population $x'_1, \ldots, x'_n$; the genomes that are not holding tickets are discarded.

The arguments of section 3.1 can readily be extended to show that this algorithm has the same stationary distribution (1.2).

3.4 The selection schemes and Moran models

Recall that under the Dirichlet prior our breeding process can be considered as Moran model with birth and mutation probabilities given in (2.4). Including the single tournament and inverse fitness selections into the model, we end up with Moran model, where the transition probability that $n_i$becomes $n_i + 1$ and n_j becomes $n_j - 1$ are, respectively,

$$\big(\frac{n_i}{n}(1-u)+u_i\big)\frac{n_j}{n}\frac{w_i}{(w_i+w_j)}, \quad \big(\frac{n_i}{n}(1-u)+u_i\big)\frac{w_j^{-1}}{\sum_k n_k w_k^{-1} + w_i^{-1}}. \tag{3.4}$$

By the same arguments as in subsections 3.1 and 3.2, it is easy to verify that both transition matrices satisfy detailed balance equations with the stationary measure being (1.2):

$$P_n(n_1,\dots,n_K) = \frac{1}{Z_n}\frac{n!}{n_1!\cdots n_K!}\frac{(\alpha_1)_{n_1}\cdots(\alpha_K)_{n_K}}{(|\alpha|)_n} w_1^{n_1}\cdots w_K^{n_K}. \tag{3.5}$$

Also, when adding the same single tournament or inverse fitness selections to the two allele model (2.6), the detailed balance equations with (3.5) hold. The same holds for the decoupled Moran model (2.7): when adding single tournament or inverse fitness selections, the detailed balance equations with (3.5) hold. Of course, our selection schemes are far from the only possibilities for having (3.5) as stationary distribution. For example, when the mutation-drift Moran model (i.e. the probability of j-type to eject just being n_j/n, no selection) satisfies the detailed balance equation with (2.2) being the stationary distribution, then multiplying the transition probability that type i is chosen to reproduce and j to die simply by w_i (as it is often done in the literature, e.g. [4, 24, 6]) then the detailed balance equation with (3.5) as stationary distribution holds. Thus in single tournament selection, the term $\frac{w_i}{(w_i+w_j)}$ can be replaced by w_i and, again, detailed balance equations with (3.5) as stationary distribution hold (independently of drift-mutation probabilities, given that they satisfy detailed balance equation with (2.2)). To summarize: there are various different kinds of Moran models yielding the stationary distribution P_n with Π being the Dirichlet distribution

Typically in the evolution literature, for haploid models the weights w_k are in form $(1+s_k)$, in two allele models simply $(1+s)$ and 1, where s is either a positive or negative small number. In the large population limit, typically $s \cdot n \to \gamma$, where γ is constant. Hence, for large n, $(1+s) \approx (1+\frac{\gamma}{n}) \approx \exp[\frac{\gamma}{n}]$ and this justifies our choice of weights in (1.4). So, when in two allele model, $\phi(1) = 0$ and $\phi(2) = \gamma > 0$, thus $w_n(1) = 1$ and $w_n(2) = e^{-\gamma/n}$, then the limit measure (1.7) has density

$$\frac{1}{Z_n} e^{-\gamma q_2} q_1^{\alpha_1 - 1} q_2^{\alpha_2 - 1}$$

that (with negative γ) is a classical and well-known limit, sometimes called Wright's formula [4] (see also (9) in [24], Ch 5 in [9], (7.28) in [6], Ch 1 in [10], Ch 4 in [16]). We would like to stress that we obtain our limit without any diffusion approximation and our main theorems 5.1 and 6.1 generalize this result in various ways.

3.5 Invariance to fitness noise

In real life, survival depends not just on fitness but also on luck. Suppose that when each new 'genome' is 'born', it combines its own intrinsic fitness with an independent random amount of luck. It keeps this amount of luck, unchanged, throughout its life. Does individual luck at birth alter the stationary distribution?

To formalize the question, let $\psi_1, \psi_2, \ldots, \psi_n$ be discrete random variables that represent 'luck'. We consider the ψ_i as discrete random variables because the formal derivations become far simpler. The random variable ψ_i can be interpreted as the individual luck that multiplies the fitness of i-th individual. The log fitness function of i-th individual is $\phi(x) + \psi_i$.

$$P_n(\mathbf{x}|\psi_1, \ldots, \psi_n) = \frac{1}{Z'_n} P_\xi(\mathbf{x}) \exp\big[- \sum_{i=1}^n \big(\phi(x_i) + \psi_i\big)\big]$$

The joint probability

$$\begin{aligned} P_n(\mathbf{x}; \psi_1, \ldots, \psi_n) &= \frac{1}{Z'_n} P_\xi(\mathbf{x}) \exp\big[- \sum_{i=1}^n \big(\phi(x_i) + \psi_i\big)\big] P(\psi_1, \ldots, \psi_n) \\ &= \frac{1}{Z'_n} P_\xi(\mathbf{x}) \exp\big[- \sum_{i=1}^n \phi(x_i)\big] \exp\big[- \sum_i \psi_i\big] P(\psi_1, \ldots, \psi_n). \end{aligned}$$

We see that the sum factorizes, and so $Z'_n = Z_n \cdot Z_n^\psi$, where

$$Z_n = \sum_{\mathbf{x}} P_\xi(\mathbf{x}) \exp\big[- \sum_{i=1}^n \phi(x_i)\big], \quad Z_n^\psi = \sum_{\psi_1, \ldots, \psi_n} \exp\big[- \sum_i \psi_i\big] \prod_i P(\psi_1, \ldots, \psi_n)$$

Thus the joint measure factorizes

$$P_n(\mathbf{x}; \psi_1, \ldots, \psi_n) = P_n(\mathbf{x}) P_n^\psi(\psi_1, \ldots, \psi_n),$$

where

$$P_n^\psi(\psi_1, \ldots, \psi_n) = \frac{1}{Z_n^\psi} P(\psi_1, \ldots, \psi_n) \exp[- \sum_i \psi_i]$$

is a probability measure. Thus, after summing $\psi_1, \ldots, \psi_n$ out, we end with $P_n(\mathbf{x})$:

$$\sum_{\psi_1, \ldots, \psi_n} P_n(\mathbf{x}; \psi_1, \ldots, \psi_n) = P_n(\mathbf{x}).$$

It follows that the stationary distribution is unchanged by multiplicative noise that is independent of the genomes, for all population sizes. Also, considering $\psi_1, \ldots, \psi_n$ as the prior, we see that posterior measure is $P_n(\psi_1, \ldots, \psi_n)$ that is independent of $\mathbf{x}$.

3.6 Fitness as likelihood: a connection with non-parametric Bayesian MCMC

In Dirichlet Process mixture models, as described for example in [21, 23], items of data $d_1, \ldots, d_n$ are given, together with a likelihood function $l(x, d) = P(d \mid x)$, where $x \in \mathcal{X}$ is a discrete latent variable. The prior distribution over latent variables is given by the exchangeable process ξ, each item of data is associated with its corresponding latent variable, and the aim of MCMC fitting is to sample latent variables $x_1, \ldots, x_n$ from the distribution

$$P_n(x_1, \ldots, x_n) = \frac{1}{Z_n} P_\xi(x_1, \ldots, x_n) P(d_1 \mid x_1) \cdots P(d_n \mid x_n) \tag{3.6}$$

where Z_n is an appropriate normalizing constant. If we write

$$w_j(x_i) := P(d_j \mid x_i)$$

then this distribution (3.6) is produced by the MCMC algorithm of section 3.1 with the slight modification that the tournament between the new 'genome' x_{n+1} and the randomly selected x_j is a victory for x_{n+1} with probability $\frac{w_j(x_{n+1})}{w_j(x_j)+w_j(x_{n+1})}$.

One might construct an evolutionary "Just-So" story as follows. On a rock in the ocean, there are n niches, in each of which one member of the species ξ can live; the fitness of x living in niche j is $w_j(x)$. Evolution occurs when a new individual, bred by 'n-way recombination' of all n parents, then challenges the occupant of a randomly chosen niche by fighting a tournament, after which the victor survives and takes over the niche. This gives a (fanciful) evolutionary interpretation to Bayesian latent variable models with exchangeable priors.

3.7 Designing a reversible evolutionary model

Our larger aim is to develop a model of evolution that is sufficiently realistic to capture some of the computational power of natural evolution, but which is also simple and tractable for analysis. To ensure that our model satisfies detailed balance and has a stationary distribution that factorises into a breeding and a selection term, we have made the following simplifying assumptions in addition to the usual simplifications of population genetics or genetic algorithms:

Overlapping populations : We believe that overlapping populations are necessary for reversibility. If full replacement of the population is enforced at each generation, there can be do guarantee that the population at time t could be easily bred from the population at time $t + 1$. In MCMC, state changes typically occur through proposing changes that may or may not be accepted; in an overlapping generations model, if a proposed change is not accepted, we continue with the same population as before.

n-way recombination : in the product of Dirichlet processes breeding system, each new genome is bred from n parents rather than from two parents selected from the population. We conjecture that this is necessary for exact reversibility because, with long genomes and many mutations, if a child is bred from two parents, then the child will be more similar to each of its parents than to other individuals in the population, so that 'triples' of two parents and one child will be identifiable even in the stationary distribution. This breaks reversibility since the direction of time can be determined observing evolution in the stationary distribution.

Mutation as sampling : We consider mutation as sampling from a base distribution of possible alleles. This model of mutation is not as general as those found in biology, where mutation probabilities are not symmetric or reversible.

Fitness as lifetime : All members of the population 'breed' at the same rate, and differences in fitness affect only the expected lifetime of an individual. This clearly differs from many types of natural selection, but it is also well known that many organisms continue to produce offspring throughout their lives, so that their total reproductive success depends on their lifetime, as well as on other factors.

It is beyond the scope of this article to argue further whether our model successfully abstracts some essential computational aspects of evolution with sexual reproduction: we present it merely as an abstraction of sexual evolution which is significantly more tractable to analyse than other apparently simple models.

4 The measure P_n

The measure P_n in (1.2) is our main object of interest. In this section we show that the marginal distributions P_n over the set of genotypes converge as the population size $n \to \infty$; in the next section we characterize these limits.

Since ξ is exchangeable, by de Finetti's theorem there exists a prior measure π on the set of all probability measures on $\mathcal{X}$ (simplex) $\mathcal{P}$ such that for every $\mathbf{x} \in \mathcal{X}^n$, the measure $P_\xi(\mathbf{x})$ allows the representation (1.3). Hence we can write P_n as follows

$$P_n(\mathbf{x}) = \frac{1}{Z_n} \prod_{i=1}^{n} w(x_i) \int_{\mathcal{P}} \prod_{i=1}^{n} q(x_i)\pi(dq) = \frac{1}{Z_n} \int_{\mathcal{P}} \prod_{i=1}^{n} \big(q(x_i)w(x_i)\big)\pi(dq), \tag{4.1}$$

where Z_n is the normalizing constant. In order to analyze the measure, it is convenient to rewrite it as follows. First, let us introduce some notation

$$\langle q, w\rangle := \sum_{k=1}^{K} q(k)w(k), \quad r_q(k) := \frac{w(k)q(k)}{\langle q, w\rangle}, \quad k = 1, \dots, K.$$

Note that since $w(k) > 0$ for all k, r_q is correctly defined for every $q \in \mathcal{P}$. Thus $\langle q, w\rangle$ is the expected weight (under q-measure) and r_q is a probability measure on $\mathcal{P}$. Now

$$P_n(\mathbf{x}) = \frac{1}{Z_n}\int_{\mathcal{P}} \prod_{i=1}^{n} \big(q(x_i)w(x_i)\big)\pi(dq) = \frac{1}{Z_n}\int_{\mathcal{P}} \langle q, w\rangle^n \prod_{i=1}^{n} r_q(x_i)\pi(dq). \tag{4.2}$$

Since r_q is a probability measure, it holds that $\sum_{\mathbf{x}} \prod_{i=1}^{n} r_q(x_i) = 1$ and so the normalization constant for P_n is

$$Z_n = \sum_{\mathbf{x}\in\mathcal{X}^n} \int_{\mathcal{P}} \langle q, w\rangle^n \prod_{i=1}^{n} r_q(x_i)\pi(dq) = \int_{\mathcal{P}} \langle q, w\rangle^n \pi(dq).$$

Finally note that (4.2) can be rewritten more neatly by defining the measure

$$d\bar{\pi}_n := \frac{\langle q, w\rangle^n}{Z_n} d\pi. \tag{4.3}$$

With $\bar{\pi}_n$, we have

$$P_n(\mathbf{x}) = \int_{\mathcal{P}} \prod_{i=1}^{n} r_q(x_i)\bar{\pi}_n(dq) \quad \forall \quad \mathbf{x} \in \mathcal{X}^n. \tag{4.4}$$

From(4.4), it is easy to find all marginal distributions, namely for any $m = 1, \ldots, n$

$$P_n(x_1, \ldots, x_m) = \int_{\mathcal{P}} \prod_{i=1}^{m} r_q(x_i)\bar{\pi}_n(dq) = \frac{1}{Z_n}\int_{\mathcal{P}} \langle q, w\rangle^{n-m} \prod_{i=1}^{m} q(x_i)w(x_i)\pi(dq). \tag{4.5}$$

In particular, when $(X_1, \ldots, X_n) \sim P_n$, then

$$P(X_i = k) = \int_{\mathcal{P}} r_q(k)\bar{\pi}_n(dq) = \frac{1}{Z_n}\int \langle q, w\rangle^{n-1} w(k)q(k)\pi(dq)$$
$$P(X_i = k, X_j = l) = \int_{\mathcal{P}} r_q(k)r_q(l)\bar{\pi}_n(dq) = \frac{1}{Z_n}\int \langle q, w\rangle^{n-2} w(k)q(k)w(l)q(l)\pi(dq)$$

and so on. It is important to observe that $P_n(x_1, \ldots, x_m)$ depends on n.

The limit process. We have defined for every n the measure (4.4) that describes the genotype distribution of a n-element population. Now the natural question is: do these measures converge (in some sense) if the population size n grows? First we have to define the sense of convergence. Since every measure P_n is defined on different domain ($\mathcal{X}^n$), we cannot speak about standard (weak) convergence of measures. Instead, we ask about the existence of a limiting stochastic process. To explain the sense of convergence, consider

that we have defined a triangular array of random variables:

$$
\begin{aligned}
&X_{1,1} \sim P_1 \\
&(X_{2,1}, X_{2,2}) \sim P_2 \\
&(X_{3,1}, X_{3,2}, X_{3,3}) \sim P_3 \\
&\dots \\
&(X_{n,1}, X_{n,2}, \dots, X_{n,n}) \sim P_n \\
&\dots
\end{aligned}
$$

We also know that the joint distribution of the first m variables in every row depends on n. Therefore we ask: is there a stochastic process $X_1, X_2, \dots$ so that for every m the following convergence holds

$$(X_{1,n}, \dots X_{m,n}) \Rightarrow (X_1, \dots X_m)? \tag{4.6}$$

According to Kolmogorov's existence theorem, the existence of a stochastic process is equivalent to the existence of (finite dimensional) measures P^*_m on set $\mathcal{X}^m$, $m = 1, 2, \dots$ that satisfy the following consistency conditions: for every m and for every $(x_1, \dots, x_m) \in \mathcal{X}^m$, it holds that

$$\sum_{x_{m+1}} P^*_{m+1}(x_1, \dots, x_m, x_{m+1}) = P^*_m(x_1, \dots, x_m).$$

If we also want (4.6) to be true, then for every m and for every $(x_1, \dots, x_m) \in \mathcal{X}^m$ the following convergences must hold:

$$P_n(x_1, \dots, x_m) \to P^*(x_1, \dots, x_m), \quad \forall m, \quad \forall (x_1, \dots, x_m) \in \mathcal{X}^m. \tag{4.7}$$

We now present a general lemma that guarantees the convergence (4.7). To achieve the full generality, we let w also depend on n. Thus, we have weights w_n, and we define the measures $r_{q,n}$ as follows

$$r_{q,n}(k) := \frac{w_n(k)q(k)}{\langle q, w_n \rangle} \quad \forall k \in \mathcal{X}.$$

We start with the following observation, proven in appendix.

Claim 4.1 *If $w_n(k) \to w(ik)$ $\forall k \in \mathcal{X}$, and $r_{q,n}$ and r_q are defined with respect of w_n and w, respectively, then the following uniform convergence holds.*

$$\sup_{q \in \mathcal{P}} |r_{q,n}(k) - r_q(k)| \to 0. \tag{4.8}$$

In the following lemma, $\bar{\pi}_n$ are arbitrary probability measures on $\mathcal{P}$, not necessarily as in (4.3). The measure $\bar{\pi}_n$ define P_n as in (4.4).

Lemma 4.1 *Let $w_n(k) \to w(k)$ for every $k \in \mathcal{X}$. If there exists an probability measure $\bar{\pi}$ such that $\bar{\pi}_n \Rightarrow \bar{\pi}$, then for every m there exists a probability measure P^*_m on $\mathcal{X}^m$ so that (4.7) holds. Moreover, for every $(x_1, \ldots, x_m) \in \mathcal{X}^m$*

$$P^*_m(x_1, \ldots, x_m) = \int_{\mathcal{P}} \prod_{i=1}^{m} r_q(x_i)\bar{\pi}(dq), \quad \text{where} \quad r_q(k) = \frac{w(k)q(k)}{\langle w, q\rangle}, \quad k = 1, \ldots, K$$

and the measures P^*_m, $m = 1, 2, \ldots$ satisfy consistency conditions.

Proof. For every $x_1, \ldots, x_m$ from (4.8), it follows that

$$\sup_{q\in\mathcal{P}} |\prod_{i=1}^{m} r_{q,n}(x_i) - \prod_{i=1}^{m} r_q(x_i)| \to 0. \tag{4.9}$$

Since the functions

$$q \mapsto \prod_{i=1}^{m} r_{q,n}(x_i), \quad q \mapsto \prod_{i=1}^{m} r_q(x_i)$$

are bounded (by 1) continuous functions, from uniform convergence, it holds

$$P_n(x_1, \ldots, x_m) = \int_{\mathcal{P}} \prod_{i=1}^{m} r_{q,n}(x_i)\bar{\pi}_n(dq) \to \int_{\mathcal{P}} \prod_{i=1}^{m} r_q(x_i)\bar{\pi}(dq) =: P^*_m(x_1, \ldots, x_m).$$

Clearly P^*_m are probability measures. The consistency condition trivially holds, because

$$\sum_{m+1} P^*_{m+1}(x_1, \ldots, x_{m+1}) = \int_{\mathcal{P}} \sum_{x_{m+1}} \prod_{i=1}^{m+1} r_q(x_i)\bar{\pi}(dq) = \int_{\mathcal{P}} \prod_{i=1}^{m} r_q(x_i)\bar{\pi}(dq) = P^*_m(x_1, \ldots, x_m).$$

■

4.1 Frequencies: the measure Q_n

We now consider how to express the limit measure P^* in terms of a limiting measure on the simplex $\mathcal{P}$. Recall $n_k(\mathbf{x})$ defined in (2.1) and let

$$\mathbf{n}(\mathbf{x}) := (n_1, \ldots, n_K), \quad \text{where } n_k = n_k(\mathbf{x}), \text{ for } k = 1, \ldots, K$$

Since ξ is exchangeable, the probability $P_\xi(\mathbf{x})$ depends on the counts $\mathbf{n}(\mathbf{x})$ only, so

$$P_\xi(\mathbf{x}) =: g(n_1, \cdots, n_K).$$

We may now write:

$$g(n_1, \ldots, n_K) = P_\xi(\mathbf{x}) = \int_{\mathcal{P}} \prod_{i=1}^{n} q(x_i)\pi(dq) = \int_{\mathcal{P}} \prod_{k=1}^{K} q(k)^{n_k}\pi(dq). \tag{4.10}$$

In what follows, let us denote

$$\mathbb{N}_n := \{(n_1, \dots, n_K) : \sum_k n_k = n\}.$$

Observe that

$$\sum_{(n_1,\dots,n_k)\in\mathbb{N}_n} \frac{n!}{n_1!\cdots n_K!} g(n_1, \cdots, n_K) = 1.$$

Therefore, the measure P_n can be defined on the set $\mathbb{N}_n$ as follows:

$$P_n(\mathbf{n}) := \frac{1}{Z_n}\frac{n!}{n_1!\cdots n_K!} g(n_1, \dots, n_K) \prod_{k=1}^{K} \big(w_n(k)\big)^{n_k} = \frac{n!}{n_1!\cdots n_K!} \int_{\mathcal{P}} \prod_{k=1}^{K} \big(r_{q,n}(k)\big)^{n_k} \bar{\pi}_n(dq). \tag{4.11}$$

Considering the frequencies instead of counts, we can define the corresponding measure on the simplex $\mathcal{P}$. Let us denote that measure as Q_n, so that with $\mathbf{n}/n := (n_1/n, \dots n_k/n)$

$$Q_n\left(\frac{\mathbf{n}}{n}\right) := P_n(\mathbf{n}), \quad \forall \mathbf{n} \in \mathbb{N}_n. \tag{4.12}$$

Thus Q_n is a discrete measure

$$Q_n = \sum_{\mathbf{n}\in\mathbb{N}_n} P_n(\mathbf{n}) \delta_{\frac{\mathbf{n}}{n}}.$$

The advantage of Q_n over the measure P_n on $\mathbb{N}_n$ is that for any n, Q_n is defined on the same domain $\mathcal{P}$, and so one can speak about the weak convergence of Q_n. Essentially, obviously, the measure P_n on $\mathcal{X}^n$, the measure P_n on $\mathbb{N}_n$ and Q_n on $\mathcal{P}$ are all the same, just the domains are different.

Since the measures Q_n are defined on the same space (simplex), it is now natural to ask, whether there exists a probability measure Q^* so that $Q_n \Rightarrow Q^*$? It turns out the if the assumption of Lemma 4.1 holds, i.e. $\pi_n \Rightarrow \bar{\pi}$ and $w_n \to w$ (pointwise), then the limit measure is actually $\bar{\pi} r^{-1}$, where

$$r : \mathcal{P} \mapsto \mathcal{P}, \quad r(q) = r_q$$

and r is defined with respect to limit weight function w. Thus for a measurable $E \subset \mathcal{P}$,

$$\bar{\pi} r^{-1}(E) = \bar{\pi}\big(r^{-1}(E)\big).$$

For example, if $\bar{\pi} = \delta_{q^*}$ (the measure is concentrated on one point), then

$$\bar{\pi} r^{-1} = \delta_{r(q^*)},$$

because

$$\delta_{q^*} r^{-1}(E) = 1 \quad \Leftrightarrow \quad q^* \in r^{-1}(E) \quad \Leftrightarrow \quad r(q^*) \in E.$$

The following lemma is the counterpart of Lemma 4.1. Again, $\bar{\pi}_n$ is an arbitrary sequence of probability measures on $\mathcal{P}$, P_n are defined via $\bar{\pi}_n$ by (4.11) and Q_n via P_n as in (4.12).

Lemma 4.2 *Let $w_n(k) \to w(k)$ for every $k \in \mathcal{X}$. If there exists a probability measure $\bar{\pi}$ such that $\bar{\pi}_n \Rightarrow \bar{\pi}$, then $Q_n \Rightarrow \bar{\pi} r^{-1}$.*

Proof. Let $f : \mathcal{P} \to \mathbb{R}$ be a (K-variable) bounded continuous function. By definition of the weak convergence, it suffices to show that

$$\int f(q) Q_n(dq) \to \int f(q) \bar{\pi} r^{-1}(dq) = \int f(r_q) \bar{\pi}(dq), \tag{4.13}$$

where the last equality holds by the change of variable formula. Note that

$$\begin{aligned}
\int f(q) Q_n(dq) &= \sum_{\mathbf{n} \in \mathbb{N}_n} f\Big(\frac{\mathbf{n}}{n}\Big) P_n(\mathbf{n}) = \sum_{\mathbf{n} \in \mathbb{N}_n} f\Big(\frac{\mathbf{n}}{n}\Big) \frac{n!}{n_1! \cdots n_K!} \int \prod_{k=1}^{K} \big(r_{q,n}(k)\big)^{n_k} \bar{\pi}_n(dq) \\
&= \int \Big(\sum_{\mathbf{n} \in \mathbb{N}_n} f\Big(\frac{\mathbf{n}}{n}\Big) \frac{n!}{n_1! \cdots n_K!} \prod_{k=1}^{K} \big(r_{q,n}(k)\big)^{n_k} \Big) \bar{\pi}_n(dq) \\
&= \int f_n(r_{q,n}(1), \ldots, r_{q,n}(K)) \bar{\pi}_n(dq) = \int f_n(r_{q,n}) \bar{\pi}_n(dq),
\end{aligned}$$

where

$$f_n(r_{q,n}) :=: f_n(r_{q,n}(1), \ldots, r_{q,n}(K)) := \sum_{(n_1, \ldots, n_K) \in \mathbb{N}_n :} f\Big(\frac{n_1}{n}, \ldots, \frac{n_1}{n}\Big) \frac{n!}{n_1! \cdots n_K!} \prod_{k=1}^{K} \big(r_{q,n}(k)\big)^{n_k}$$

is the Bernstein polynomial evaluated at $r_{q,n} = (r_{q,n}(1), \ldots, r_{q,n}(K))$. It is easy to see and well known that for any vector $r \in \mathcal{P}$, $f_n(r) \to f(r)$, moreover, the convergence is uniform over $\mathcal{P}$:

$$\sup_{r \in \mathcal{P}} \Big| f_n(r) - f(r) \Big| \to 0. \tag{4.14}$$

Since for every q, $r_{q,n}$ is a probability vector, then

$$\frac{n!}{n_1! \cdots n_K!} \prod_{k=1}^{K} \big(r_{q,n}(k)\big)^{n_k} \leq 1,$$

and since f is bounded, we see that for every n, $q \mapsto f_n(r_{q,n}) =: b_n(q)$ is a bounded continuous function. Also the function $q \mapsto f(r_q) =: b(q)$ is a bounded continuous function. Then

$$\sup_q |b_n(q) - b(q)| \leq \sup_q |f_n(r_{q,n}) - f(r_{q,n})| + \sup_q |f(r_{q,n}) - f(r_q)|.$$

By (4.8), $\sup_q |r_{q,n}(k) - r_q(k)| \to 0$ for every $k \in \mathcal{X}$. Then also $\sup_q \|r_{q,n} - r_q\| \to 0$. A continuous function on compact space is uniformly continuous, so

$$\sup_q |f(r_{q,n}) - f(r_q)| \to 0.$$

By (4.14),

$$\sup_q |f_n(r_{q,n}) - f(r_{q,n})| \leq \sup_r |f_n(r) - f(r)| \to 0.$$

Therefore, we have shown that $\sup_q |b_n(q) - b(q)|$ implying that

$$\int f(q)Q_n(dq) = \int b_n(q)\bar{\pi}_n(dq) \to \int b(q)\bar{\pi}(dq) = \int f(r_q)\bar{\pi}(dq).$$

■

5 P^* and Q^* in the large population limit

In what follows, let us rewrite fitnesses in terms of $\phi(k) := -\ln(w(k))$, so that for any genome $k \in \mathcal{X}$, $w(k) = \exp(-\phi(k))$. Moreover, in order to increase the influence of prior, we let weights w_n depend on n as defined in (1.4), i.e. $w(k) = \exp(-\phi(k)/n^\lambda)$. Clearly, for every k, $w_n(k) \to w(k)$ and λ controls the speed of that convergence. When $\lambda > 0$, then $w(k) = 1$ implying that in this case the mapping r is identity, i.e. for every q, $r_q = q$.

Let us return to our original $\bar{\pi}_n$, defined as in (4.3) with w_n:

$$\bar{\pi}_n(E) = \int_E \frac{\langle q, w_n\rangle^n}{Z_n}\pi(dq), \quad Z_n = \int \langle q, w\rangle^n \pi(dq). \tag{5.1}$$

In this section we consider the case where the prior measure π is independent of n, and the support of π is the whole simplex $\mathcal{P}$. When $w(1) > w(2)$, the function $q \mapsto \langle q, w\rangle$ has unique maximizer $q^* := (1, 0, \ldots, 0)$. The following theorem states that under $w(1) > w(2)$ that phase transition occurs: when $0 \leq \lambda < 1$, then $\bar{\pi}_n \Rightarrow \delta_{q^*}$ and $P_n(x_1, \ldots, x_m) \to P^*(x_1, \ldots, x_m)$, where

$$P^*(x_1, \ldots, x_m) = \prod_{i=1}^m r_{q^*}(x_i) = \prod_{i=1}^m q^*(x_i) = \begin{cases} 1, & \text{if for every } i,\ x_i = 1\ ; \\ 0, & \text{else.} \end{cases},$$

because in both cases (i.e. $\lambda = 0$ and $\lambda \in (0, 1)$), it holds that $r_{q^*} = q^*$. Thus the limit process $X_1, X_2, \ldots$ has only one realization: $1, 1, \ldots$. When $\lambda = 1$, then $\bar{\pi}_n$ as well Q_n converges to a nondegenerate distribution, when $\lambda > 1$, then $\bar{\pi}_n \Rightarrow \pi$, $Q_n \Rightarrow \pi$ and the measure P^* is the law of birth process ξ.

Theorem 5.1 *Let the fitness function be defined as in (1.4) and assume that the support of the prior π is $\mathcal{P}$. Then the following convergences hold:*

1) *If $\lambda \in [0, 1)$ and $\phi(1) < \phi(2)$, then $\bar{\pi}_n \Rightarrow \delta_{q^*}$, $Q_n \Rightarrow \delta_{q^*}$ and (4.7) holds with*

$$P^*(x_1, \ldots, x_m) = \prod_{i=1}^m q^*(x_i), \text{ where } q^* = (1, 0, \ldots, 0).$$

2) *If $\lambda = 1$, then $\bar{\pi}_n \Rightarrow \bar{\pi}$, $Q_n \Rightarrow \bar{\pi}$ and (4.7) holds with*

$$P^*(x_1, \ldots, x_m) = \int \prod_{i=1}^{m} q(x_i)\bar{\pi}(dq),$$

where for every $E \subset \mathcal{P}$,

$$\bar{\pi}(E) = \frac{1}{Z}\int_E \exp[-\langle \phi, q\rangle]\pi(dq), \quad Z = \int \exp[-\langle \phi, q\rangle]\pi(dq).$$

3) *If $\lambda > 1$, then $\bar{\pi}_n \Rightarrow \pi$, $Q_n \Rightarrow \pi$ and (4.7) holds with*

$$P^*(x_1, \ldots, x_m) = P_\xi(x_1, \ldots, x_m).$$

5.1 Proof of Theorem 5.1

Before proving the theorem, let us state a very useful preliminary result. Recall that simplex $\mathcal{P}$ is a compact set. Let $f_n, f : \mathcal{P} \to \mathbb{R}^+$ be continuous, hence bounded measurable functions so that $f_n \to f$ uniformly and let $m_n \to \infty$ be an increasing sequence. We are given a measure π on $\mathcal{P}$ having full support (i.e. the support of π is $\mathcal{P}$), and we are interested in the asymptotic behavior of the measure ν_n, where

$$\nu_n(E) := \int_E h_n(q)\pi(dq), \quad h_n(q) := \frac{f_n^{m_n}(q)}{\int f_n^{m_n}(q)\pi(dq)} = \Big(\frac{f_n(q)}{\|f_n\|_{m_n}}\Big)^{m_n}.$$

Here $\|f\|_m$ stands for L_m-norm and we assume that $\int f_n^{m_n} d\pi < \infty$ for every n. If π is a finite measure, then the conditions automatically holds due to the boundedness of f_n. In what follows, let

$$\mathcal{P}^* := \{q \in \mathcal{P} : f(q) = \|f\|_\infty\}, \quad \mathcal{P}^*_\delta := \{q \in \mathcal{P} : f(q) > \|f\|_\infty - \delta\}. \tag{5.2}$$

Here $\|f\|_\infty$ is the essential supremum of f with respect to the π-measure. If f is continuous, then $\|f\|_\infty = \sup_q f(q)$ (because π has full support). The proof of the following Proposition 5.1 is given in appendix.

Proposition 5.1 *Let $f_n \to f$ uniformly and let π be a finite measure on $\mathcal{P}$ having full support. Then for every $\delta > 0$, $\nu_n\big(\mathcal{P}^*_\delta\big) \to 1$. If $\mathcal{P}^* = \{q^*\}$, then $\nu_n \Rightarrow \delta_{q^*}$.*

Besides Proposition 5.1, the proof of Theorem 5.1 is based on the following well-known observation: when $m \to \infty$, then

$$\sup_{q\in\mathcal{P}} \Big|\big\langle \exp[-\frac{\phi}{m}], q\big\rangle^m - \exp[-\langle \phi, q\rangle]\Big| \to 0. \tag{5.3}$$

Proof. (Theorem 5.1)

1) For $\lambda = 0$, take $f_n(q) = f(q) = \langle w, q\rangle$, $m_n = n$. By $w(1) > w(2) \geq \cdots \geq w(K)$, we have $\mathcal{P}^* = \{q^*\}$, from Proposition 5.1, it follows that $\bar{\pi}_n \Rightarrow \delta_{q^*}$. Since for any weight w, $r_{q^*}(k) = q^*(k)$, from Lemma 4.1, it follows that

$$P^*(x_1, \ldots, x_m) = \prod_{i=1}^{m} r_{q^*}(x_i) = \prod_{i=1}^{m} q^*(x_i).$$

Since $\bar{\pi}_n r^{-1}(q^*) = q^*$, from Lemma 4.2, it follows that $Q_n \Rightarrow \delta_{q^*}$. So, for $\lambda = 0$, the statement is proven and we now consider the case $\lambda \in (0, 1)$. Let

$$f_n(q) := \langle \exp[-\frac{\phi}{n^\lambda}], q\rangle^{n^\lambda}, \quad f(q) := \exp[-\langle \phi, q\rangle].$$

By (5.3), $\|f_n - f\|_\infty \to 0$. Since $\lambda \in (0, 1)$, take $m_n = n^{1-\lambda}$. Then

$$h_n(q) := \frac{\langle \exp[-\frac{\phi}{n^\lambda}], q\rangle^n}{Z_n}$$

is the density of $\bar{\pi}_n$ with respect to π. Since f is continuous, the set $\mathcal{P}^*$ in (5.2) is

$$\mathcal{P}^* = \arg\max_{q \in \mathcal{P}} f(q) = \arg\min_{q \in \mathcal{P}} \langle \phi, q\rangle = \{(1, 0, \ldots, 0)\} = \{q^*\}.$$

By Proposition 5.1, $\bar{\pi}_n \Rightarrow \delta_{q^*}$. As in the case of $\lambda = 0$, it follows that $Q_n \Rightarrow \delta_{q^*}$ and $P^*(x_1, \ldots, x_m) = 1$ if and only if $x_1 = \cdots = x_m = 1$.

2) Since for any q and any n, it holds

$$\langle \exp[-\frac{\phi}{n}], q\rangle^n \leq e^{-\phi(1)} \leq 1,$$

we obtain from (5.3) and bounded convergence that for any measurable E

$$\int_E \langle \exp[-\frac{\phi}{n}], q\rangle^n \pi(dq) \to \int_E \exp[-\langle \phi, q\rangle]\pi(dq). \tag{5.4}$$

Recall

$$\bar{\pi}(E) = \frac{1}{Z}\int_E \exp[-\langle \phi, q\rangle]\pi(dq), \quad \text{where} \quad Z = \int \exp[-\langle \phi, q\rangle]\pi(dq).$$

Therefore, from (5.4), it follows that when $\lambda = 1$, we have for every measurable E

$$\bar{\pi}_n(E) \to \bar{\pi}(E),$$

meaning that $\bar{\pi}_n \Rightarrow \bar{\pi}$ (even in a stronger sense). Since $r_q = q$, from Lemma 4.1, it follows that the limits of $P_n(x_1, \ldots, x_m)$ are

$$P^*(x_1, \ldots, x_m) = \int \prod_{i=1}^{m} q(x_i)\bar{\pi}(dq) = \frac{1}{Z}\int \prod_{i=1}^{m} q(x_i) \exp[-\langle \phi, q\rangle]\pi(dq).$$

Since r is identity function, by Lemma 4.2 the limit measure of frequencies is $\bar{\pi}$, i.e. $Q_n \Rightarrow \bar{\pi}$.

3) Since for any q,

$$\langle \exp[-\frac{\phi}{n^\lambda}], q\rangle^n \to 1$$

by dominated convergence, again, for any measurable E

$$\int_E \langle \exp[-\frac{\phi}{n}], q\rangle^n \pi(dq) \to \pi(E).$$

Therefore $\bar{\pi}_n \Rightarrow \pi$. By Lemma 4.1, the limits of $P_n(x_1, \ldots, x_m)$ are

$$P^*(x_1, \ldots, x_m) = \int \prod_{i=1}^m q(x_i)\pi(dq),$$

so that the limit process is ξ. The convergence $Q_n \Rightarrow \pi$ follows from Lemma 4.2.

■

We have seen that the critical case $\lambda = 1$ is the only case where the prior and fitnesses both determine the limit measure. In this case, the limit process $X_1, X_2, \ldots$, governed by P^* has marginals

$$P(X_i = k) = P(X_1 = k) = \int q(k)\bar{\pi}(dq) = \frac{1}{Z}\int e^{-\langle \phi, q\rangle} q(k)\pi(dq),$$

$$P(X_i = k, X_j = l) = P(X_1 = k, X_2 = l) = \frac{1}{Z}\int e^{-\langle \phi, q\rangle} q(k)q(l)\pi(dq).$$

It is also interesting to point out that in the critical case $\lambda = 1$, the measure $\bar{\pi}$ satisfies

$$\bar{\pi} = \arg\min_{\pi' \in E} D(\pi \| \pi'),$$

where E is a set of probability measures on $\mathcal{P}$, namely $E := \{\pi' : \int \langle \phi, q\rangle \pi'(dq) \geq c\}$, D stands for Kullback-Leibler divergence and $c > 0$ is a constant.

6 Dirichlet prior

Also in the current section we consider the weights $w_n(k)$ as in (1.4), where $\lambda \in [0, 1]$. We already know that in the case of constant priors, the case $\lambda < 1$ means that the fitnesses will prevail over the prior. Therefore, it is meaningful to consider non-constant priors so that the influence of prior increases with suitable rate. This motivates us to consider Dirichlet priors (1.6), i.e. $\pi_n = \mathrm{Dir}(n^{1-\lambda}\alpha_1, \ldots n^{1-\lambda}\alpha_K)$, where $\alpha := (\alpha_1, \ldots, \alpha_K)$, $\alpha_k > 0$ and $|\alpha| := \sum_k \alpha_k$. The constant $\sum_i n^{1-\lambda}\alpha_k = |\alpha| n^{1-\lambda}$ is the so called *concentration* or *precision* parameter, the bigger that parameter, the more prior is concentrated over its expectation $(\alpha_1/|\alpha|, \ldots, \alpha_K/|\alpha|)$. Increasing the concentration parameter increases the influence of prior, and now it is clear the the smaller is λ, the bigger must be the prior influence. This justifies the choice of $n^{1-\lambda}$. The case $\lambda = 1$ corresponds to already studied case of constant priors, therefore we now consider the case $\lambda \in [0, 1)$. The following theorem shows that the phase transition occurs again.

Theorem 6.1 *Let the fitness function be defined be defined as in (1.4) and the prior π_n as in (1.6). Let $\bar{\pi}_n$ be defined as in (5.1) with π_n instead of π. Then the following convergences hold:*

1) *If $\lambda = 0$, then $\bar{\pi}_n \Rightarrow \delta_{q^*}$, where q^* is the unique maximizer of the following function*

$$\ln\langle e^{-\phi}, q\rangle + \sum_{k=1}^{K} \alpha_k \ln q(k). \tag{6.1}$$

Then $Q_n \Rightarrow \delta_{r^*}$, where $r^* = r_{q^*}$, so that $r^*(k) \propto q^*(k)w(k)$, and (4.7) holds with

$$P^*(x_1, \ldots, x_m) = \prod_{i=1}^{m} r^*(x_i).$$

2) *If $\lambda \in (0, 1)$, then $\bar{\pi}_n \Rightarrow \delta_{q^*}$, where q^* is the unique maximizer of the following function:*

$$-\langle \phi, q\rangle + \sum_{k=1}^{K} \alpha_k \ln q(k) \tag{6.2}$$

Then $Q_n \Rightarrow \delta_{q^*}$ and (4.7) holds with

$$P^*(x_1, \ldots, x_m) = \prod_{i=1}^{m} q^*(x_i).$$

Let us start with proving the uniqueness of the solutions of (6.1) and (6.2). The proof of the following lemma is in appendix.

Lemma 6.1

1) *The function (6.1) has an unique maximizer q^*, where*

$$q^*(k) = \frac{\alpha_k}{(1 + |\alpha|) - \frac{w(k)}{\theta}}, \quad k = 1, \ldots, K \tag{6.3}$$

where $\theta > 0$ is a parameter satisfying $\theta = \langle w, q^*\rangle$.

2) *The function (6.2) has an unique maximizer q^*, where*

$$q^*(k) = \frac{\alpha_k}{\phi(k) + |\alpha| - \theta}, \quad k = 1, \ldots, K, \tag{6.4}$$

where $\theta > 0$ is the parameter satisfying $\theta = \langle \phi, q^*\rangle$.

Proof of theorem 6.1.

1) In the case $\lambda = 0$, the measure $\bar{\pi}_n$ has the following density with respect to the Lebesgue measure:

$$\bar{\pi}_n(q) = \frac{1}{Z_n}\langle e^{-\phi}, q\rangle^n \cdot \frac{1}{B(n\alpha)}\prod_{k=1}^{K}(q(k))^{n(\alpha_k - \frac{1}{n})} = \frac{f_n(q)^n}{Z'_n},$$

where

$$f_n(q) := \langle e^{-\phi}, q\rangle \prod_{k=1}^{K}(q(k))^{(\alpha_k - \frac{1}{n})}, \quad Z'_n := \int f_n^n(q)dq.$$

Clearly for every q, $f_n(q) \to f(q)$, where

$$f(q) = \langle e^{-\phi}, q\rangle \prod_{k=1}^{K}(q(k))^{\alpha_k}.$$

It is not hard to see that the convergence is uniform, i.e. $\sup_q |f_n(q) - f(q)| \to 0$. By **1)** of Lemma 6.1, the function f has unique maximizer q^* (6.3), i.e. $\mathcal{P}^* = \{q^*\}$. Now apply Proposition 5.1 with π being the Lebesgue measure on $\mathcal{P}$ (hence π is finite) and $m_n = n$ so that

$$\nu_n(E) = \int_E h_n dq = \frac{1}{Z'_n}\int_E f_n(q)^n dq = \bar{\pi}_n(E).$$

Since all assumptions are fulfilled, we have $\bar{\pi}_n \Rightarrow \delta_{q^*}$. Since $w_n = w$, by Lemma 4.1 the limit process has finite-dimensional distributions

$$P^*(x_1, \ldots, x_m) = \prod_{i=1}^{m} r^*(x_i), \quad \text{where} \quad r^* = r_{q^*}$$

so that the limit process P^* corresponds to a i.i.d. sequence $X_1, X_2, \ldots$ with $X_1 \sim r^*$. According to Lemma 4.2, the frequencies Q_n converge weakly to the measure r^* and this is also quite obvious by SLLN.

2) The proof is similar: $\bar{\pi}_n$ has density (with respect to Lebesgue measure)

$$\frac{f_n(q)^{m_n}}{Z'_n}, \quad \text{where} \quad f_n(q) = \langle e^{-\frac{\phi}{n^\lambda}}, q\rangle^{n^\lambda} \prod_{k=1}^{K}(q(k))^{(\alpha_k - \frac{1}{n^{1-\lambda}})}, \quad m_n = n^{(1-\lambda)}.$$

Since

$$\langle e^{-\frac{\phi}{n^\lambda}}, q\rangle^{n^\lambda} \to e^{\langle \phi, q\rangle},$$

uniformly over q, we have that sequence f_n converges uniformly to

$$f(q) = e^{-\langle \phi, q\rangle} \prod_{k=1}^{K}(q(k))^{\alpha_k}.$$

By **2)** of Lemma 6.1, the function f has unique maximizer q^* (6.4). As in the case **1)**, it is easy to see that the assumptions of Proposition 5.1 are fulfilled with π being Lebesgue measure on $\mathcal{P}$, $m_n = n^{1-\lambda}$ and so $\bar{\pi}_n \Rightarrow \delta_{q^*}$. In the present case, for every $k = 1, \ldots, K$, $w_n(k) \to 1$ and so by Lemma 4.1, the limit process P^* is i.i.d. process with distribution q^* (because r is identity function). According to Lemma 4.2, the frequencies Q_n converge weakly to the measure δ_{q^*}.

6.1 Relation between $\lambda = 0$ and $\lambda \in (0, 1)$

From (5.3), it follows:

$$\sup_{q \in \mathcal{P}} \left| m \ln\langle \exp[-\frac{\phi}{m}], q\rangle + \langle \phi, q\rangle] \right| \to 0 \tag{6.5}$$

so that with

$$f_m(q) := \ln\langle \exp[-\frac{\phi}{m}], q\rangle + \sum_k \frac{\alpha_k}{m} \ln q(k), \quad f(q) = -\langle \phi, q\rangle + \sum_k \alpha_k \ln q(k),$$

we have

$$\sup_{q \in \mathcal{P}} |m f_m(q) - f(q)| \to 0.$$

Since $f(q)$ is as in (6.2), it has the unique maximizer q^* given in (6.3). On the other hand, the maximizer of $m f_m(q)$ is the same as the maximizer of $f_m(q)$, which corresponds to (6.1) where ϕ is replaced by ϕ/m and α is replaced by α/m. Let this unique maximizer be q^*_m Since the functions $m f_m(\cdot)$ and $f(\cdot)$ are continuous, uniformly convergent and having unique maximum, it follows that $q^*_m \to q^*$ (in usual sense, because $\mathcal{P}$ is compact). Thus, we have proven the following proposition.

Proposition 6.1 *Let*

$$q^*_m = \arg\max_q \Big(\ln\langle \exp[-\frac{\phi}{m}], q\rangle + \sum_k \frac{\alpha_k}{m} \ln q(k) \Big)$$

and let q^* be the maximizer of (6.2). Let r^*_m be the corresponding r measure, i.e. $r^*_m(k) \propto q^m(k) \exp[-\frac{\phi(k)}{m}]$. Then $q^*_m \to q^*$ and $r^*_m \to q^*$.

6.2 Product of Dirichlet priors

Recall the setup in Subsection 2.2. The set of genomes is now $\mathcal{X}^L = \overbrace{\mathcal{X} \times \cdots \times \mathcal{X}}^{L}$ and the breeding process $\xi = (\xi^1, \ldots, \xi^L)$, where ξ^l are independent exchangeable processes. We now assume that the prior of ξ^l is $\pi^l = \mathrm{Dir}(\alpha^l)$, where $\alpha^l = (a^l_1, \ldots, \alpha^l_K)$, $l = 1, \ldots, L$. In this model, L different Polya urns are run independently. Let $\mathcal{P}^L$ be the set of L-fold product measures:

$$\mathcal{P}^L := \{q^1 \times \cdots \times q^L : q^j \in \mathcal{P}\},$$

where $\mathcal{P}$, as previously, stands for the $(K-1)$-dimensional simplex. Observe that $\mathcal{P}^L$ is a compact subset of the set of all possible probability measures on $\mathcal{X}^L$. Since the components of ξ are independent, the prior π of ξ is the product of Dirichlet measures $\pi = \pi^1 \times \cdots \times \pi^L$. This means that the support of π is $\mathcal{P}^L$ and for every element $q = q^1 \times \cdots \times q^L \in \mathcal{P}^L$, the density is (with slight abuse of notation, π stands for the measure as well as for its density)

$$\pi(q) = \prod_{l=1}^{L} \pi^l(q^l) = \frac{1}{B} \prod_{l=1}^{L} \prod_{k=1}^{K} \left(q(k)^l\right)^{\alpha_k^l - 1}, \quad B := \prod_{l=1}^{L} B(\alpha^l).$$

The function ϕ is now defined on the set $\mathcal{X}^L$, and so for any $q \in \mathcal{P}^L$,

$$\langle \phi, q \rangle = \sum_{(k_1,\ldots,k_L) \in \mathcal{X}^L} \phi(k_1, \ldots, k_L) q^1(k_1) \cdots q^L(k_L)$$

and $\langle e^{-\phi}, q \rangle$ is defined similarly. When $\lambda = 0$, the measure $\bar{\pi}_n$ has density $f_n(q)^n / Z'_n$, where

$$f_n(q) = \langle e^{-\phi}, q \rangle \prod_{l=1}^{L} \prod_{k=1}^{K} (q^l(k))^{(\alpha_k^l - \frac{1}{n})}, \quad Z'_n := \int f_n^n(q) dq.$$

Clearly $f_n(q)$ converges uniformly to

$$f(q) := \langle e^{-\phi}, q \rangle \prod_{l=1}^{L} \prod_{k=1}^{K} (q^l(k))^{(\alpha_k^l)}. \tag{6.6}$$

Similarly, when $\lambda \in (0,1)$ the measure $\bar{\pi}_n$ has density $f_n(q)^{m_n} / Z'_n$, where $m_n = n^{1-\lambda}$,

$$f_n(q) = \langle e^{\frac{-\phi}{n^\lambda}}, q \rangle^{n^\lambda} \prod_{l=1}^{L} \prod_{k=1}^{K} (q^l(k))^{(\alpha_k^l - \frac{1}{n^{1-\lambda}})}, \quad Z'_n := \int f_n^{m_n}(q) dq.$$

Again, $f_n(q)$ converges uniformly to

$$f(q) := e^{\langle -\phi, q \rangle} \prod_{l=1}^{L} \prod_{k=1}^{K} (q^l(k))^{(\alpha_k^l)}. \tag{6.7}$$

When (6.6) (resp. (6.7)) have unique maximizer q^*, then the statements of Theorem 6.1 hold (the proof is the same):

1) Suppose $\lambda = 0$ and (6.6) has unique maximizer $q^* = q^1 \times \cdots \times q^L$. Then $\bar{\pi}_n \Rightarrow \delta_{q^*}$, and $Q_n \Rightarrow \delta_{r^*}$, where

$$r^*(k_1, \ldots, k_L) \propto w(k_1, \ldots, k_L) q^1(k_1) \cdots q^L(k_L),$$

and $w(k_1, \ldots, k_L) = \exp[-\phi(k_1, \ldots, k_L)]$. Observe that the measure r^* is not necessarily a product measure. Then also (4.7) holds with

$$P^*(x_1, \ldots, x_m) = \prod_{i=1}^{m} r^*(x_i), \quad x_i \in \mathcal{X}^L.$$

2) Suppose $\lambda \in (0,1)$ and (6.7) has unique maximizer $q^* = q^1 \times \cdots \times q^L$. Then $\bar{\pi}_n \Rightarrow \delta_{q^*}$, $Q_n \Rightarrow \delta_{q^*}$ and (4.7) holds with

$$P^*(x_1, \ldots, x_m) = \prod_{i=1}^{m} q^*(x_i), \quad x_i \in \mathcal{X}^L.$$

In the case $L > 1$, the maximizer of (6.6) and (6.7) is not always unique. Whether it is unique or not depends on ϕ and vectors α^l. Indeed, maximizing (6.6) is equivalent to minimizing

$$-\ln\left(\langle e^{-\phi}, q\rangle\right) - \sum_{l=1}^{L}\sum_{k=1}^{K} \alpha_k^l \ln(q^l(k)), \tag{6.8}$$

and $-\sum_{l=1}^{L}\sum_{k=1}^{K} \alpha_k^l \ln(q^l(k))$ is always a convex function. So, when the parameters α_k^l are big enough, then the whole function (6.8) becomes convex. The same argument holds for (6.7). We shall present some sufficient conditions for convexity of (6.6) and (6.7) for the case $K = L = 2$ below.

Recall: when a positive continuous function $f(q)$ has one maximizer q^*, then for any sequence $m_n \to \infty$, the measures ν_n with densities proportional to $f^{m_n}(q)$ converge weakly to δ_{q^*}. When the function has, say, two maximizers, q_1^* and q_2^*, then by Proposition 5.1, for all disjoint open balls B_1 and B_2 so that $q_i^* \in B_i$, it still holds that $\nu_n(B_1) + \nu_n(B_2) \to 1$. Thus, when the measures ν_n are weakly convergent, then the limit measure is concentrated on $\{q_1^*, q_2^*\}$ so that the limit measure must be $p\delta_{q_1^*} + (1-p)\delta_{q_2^*}$, for some $p \in [0,1]$. In this case $\nu_n(B_1) \to p$. However, the function f might be so that the limits $\nu_n(B_i)$ do not exist. And even if they do exist (i.e. the measures ν_n are weakly convergent), the limits p and $1-p$ might be arbitrary real numbers, and hard to determine. Therefore the following theorem adapted from [13] (Theorem 5.7 and a remark after it) might be very useful.

Theorem 6.2 *Suppose $K \subset \mathbb{R}^k$ is a compact non-empty subset and let $g : K \to [0, \infty)$ be a twice continuously differentiable function with finitely many minimum points $\{a_1, \ldots, a_r\}$ all located in the interior of K. Let, for every $i = 1, \ldots r$ the Hessian of g at a_i be positive definite. Given any increasing sequence m_n, define the sequence of measures*

$$\nu_n(E) := \frac{1}{Z_n}\int_E \exp[-m_n \cdot g(x)]dx, \quad Z_n := \int \exp[-m_n \cdot g(x)]dx.$$

Then $\nu_n \Rightarrow \nu$, where

$$\nu = \sum_{i=1}^{r} p_i \delta_{a_i}, \qquad \text{with } p_i \propto \frac{1}{\sqrt{\det H(a_i)}}$$

and $\det H(a_i)$ is a determinant of Hessian evaluated at a_i.

To apply the theorem in our case, let us first note that any $K-1$ dimensional simplex can be considered as a $K-1$-dimensional non-empty compact set

$$\mathcal{P}_K = \{(q_1, \ldots, q_{K-1} : q_k \geq 0, \sum_{k}^{K-1} q_k \leq 1\}.$$

Therefore, our search space $\mathcal{P}^L$ can be considered as a subset in $\mathbb{R}^{L(K-1)}$. This subset has non-empty interior. Clearly any solution of (6.6) and (6.7) has all components strictly positive so that all maximizers of maximizer of (6.6) are interior points and the same holds for (6.7). We take $g(q) = -\ln f(q)$, where f is as in (6.6) or (6.7). Thus, for any m, $\exp[-mg(y)] = f^m(q)$ so that the measure defined in the statement of theorem is

$$\nu_n(E) \propto \int_E f^{m_n}(q)dq.$$

However, even when the measures ν_n converge weakly to a limit, it does not automatically follow that the measures $\bar{\pi}_n$ converge to the same limit even if f_n converges to f uniformly. This convergence might depend on the speed of the uniform convergence, and we leave it for the further studies and proceed with an example instead.

6.2.1 The case $K = L = 2$

Let us analyze more closely the case $K = 2$ and $L = 2$. Denote $q^1(1) =: z_1$ and $q^2(1) =: z_2$. Also denote $\alpha_k^1 = \alpha_k$ and $\alpha_k^2 = \beta_k$. The function (6.8) is

$$g(z_1, z_2) = -\ln\big(\langle w, z\rangle\big) - \alpha_1 \ln z_1 - \alpha_2 \ln(1 - z_1) - \beta_1 \ln z_2 - \beta_2 \ln(1 - z_2), \tag{6.9}$$

where

$$\langle w, z\rangle = w(1,1)z_1z_2 + w(1,2)z_1(1 - z_2) + w(2,1)(1 - z_1)z_2 + w(2,2)(1 - z_1)(1 - z_2).$$

Thus with $w^* := w(1,1) - w(1,2) - w(2,1) + w(2,2)$ and

$$\begin{aligned} \theta_1^1 &:= w(1,1)z_2 + w(1,2)(1 - z_2), \quad & \theta_2^1 &= w(2,1)z_2 + w(2,2)(1 - z_2) \\ \theta_1^2 &:= w(1,1)z_1 + w(2,1)(1 - z_1), \quad & \theta_2^2 &:= w(1,2)z_1 + w(2,2)(1 - z_1). \end{aligned}$$

we obtain the Hessian

$$\begin{pmatrix} \frac{\partial^2 g}{\partial z_1^2} & \frac{\partial^2 g}{\partial z_1 \partial z_2} \\ \frac{\partial^2 g}{\partial z_2 \partial z_1} & \frac{\partial^2 g}{\partial z_2^2} \end{pmatrix} = \begin{pmatrix} \frac{\alpha_1}{z_1^2} + \frac{\alpha_2}{(1-z_1)^2} + \frac{(\theta_1^1-\theta_2^1)^2}{\langle w,z\rangle^2} & \frac{(\theta_1^1-\theta_2^1)(\theta_1^2-\theta_2^2)-w^*\langle w,z\rangle}{\langle w,z\rangle^2} \\ \frac{(\theta_1^1-\theta_2^1)(\theta_1^2-\theta_2^2)-w^*\langle w,z\rangle}{\langle w,z\rangle^2} & \frac{\beta_1}{z_2^2} + \frac{\beta_2}{(1-z_2)^2} + \frac{(\theta_1^2-\theta_2^2)^2}{\langle w,z\rangle^2} \end{pmatrix}.$$

The elements in the main diagonal are strictly positive and therefore the matrix is positive definite if the determinant is positive, i.e.

$$\Big(\frac{\alpha_1}{z_1^2} + \frac{\alpha_2}{(1 - z_1)^2} + \frac{(\theta_1^1 - \theta_2^1)^2}{\langle w, z\rangle^2}\Big)\Big(\frac{\beta_1}{z_2^2} + \frac{\beta_2}{(1 - z_2)^2} + \frac{(\theta_1^2 - \theta_2^2)^2}{\langle w, z\rangle^2}\Big) > \Big(\frac{(\theta_1^1 - \theta_2^1)(\theta_1^2 - \theta_2^2)}{\langle w, z\rangle^2} - \frac{w^*}{\langle w, z\rangle}\Big)^2.$$

Observe: $\langle w, z\rangle \geq \min_{i,j} w(i,j) > 0$ and so

$$\Big|\frac{w^*}{\langle w, z\rangle}\Big| \leq \frac{|w^*|}{\min_{i,j} w(i,j)}.$$

On the other hand, for any $(z_1, z_2) \in [0,1] \times [0,1]$,

$$\frac{\alpha_1}{z_1^2} + \frac{\alpha_2}{(1-z_1)^2} \geq (\alpha_1^{\frac{1}{3}} + \alpha_2^{\frac{1}{3}})^3, \quad \frac{\beta_1}{z_2^2} + \frac{\beta_2}{(1-z_2)^2} \geq (\beta_1^{\frac{1}{3}} + \beta_2^{\frac{1}{3}})^3.$$

Therefore, it is not hard to see that when

$$(\alpha_1^{\frac{1}{3}} + \alpha_2^{\frac{1}{3}})^3 (\beta_1^{\frac{1}{3}} + \beta_2^{\frac{1}{3}})^3 > \big(\frac{w^*}{\min_{i,j} w(i,j)}\big)^2, \tag{6.10}$$

then the Hessian is always positive definite, i.e. the function g in (6.9) is strictly convex, and the minimum unique. In particular, the condition holds if $w^* = 0$. In particular, if $\alpha_1 = \beta_1$, $\alpha_2 = \beta_2$ and $w(1,2) = w(2,1)$, then under (6.10) the unique solution z_1, z_2 satisfies $z_1 = z_2$ (by symmetry). Indeed, in symmetric case it holds: if (z_1, z_2) is a solution, then so must be (z_2, z_1), and if the solution is unique, then it must be that $z_1 = z_2$. If (6.10) fails, then the function might be non-convex, and for small α_k and β_k values it is (given $w^* \neq 0$), but evaluated at the minimums, the Hessian might still be positive definite and so Theorem 6.2 might apply.

Similarly, for (6.7)

$$g(z) = -\ln f(z) = \langle \phi, z \rangle - \alpha_1^1 \ln z_1 - \alpha_2^1 \ln(1-z_1) - \alpha_1^2 \ln z_2 - \alpha_2^2 \ln(1-z_2) \tag{6.11}$$

and so the Hessian is

$$\begin{pmatrix} \frac{\partial^2 g}{\partial z_1^2} & \frac{\partial^2 g}{\partial z_1 \partial z_2} \\ \frac{\partial^2 g}{\partial z_2 \partial z_1} & \frac{\partial^2 g}{\partial z_2^2} \end{pmatrix} = \begin{pmatrix} \frac{\alpha_1}{z_1^2} + \frac{\alpha_2}{(1-z_1)^2} & \phi^* \\ \phi^* & \frac{\beta}{z_2^2} + \frac{\beta_2}{(1-z_2)^2} \end{pmatrix},$$

where

$$\phi^* := \phi(1,1) - \phi(1,2) - \phi(2,1) + \phi(2,2).$$

Since the elements of the main diagonal are are strictly positive, the matrix is positive definite if and only if the determinant is positive:

$$\big(\frac{\alpha_1}{z_1^2} + \frac{\alpha_2}{(1-z_1)^2}\big)\big(\frac{\beta_1}{z_2^2} + \frac{\beta_2}{(1-z_2)^2}\big) > (\phi^*)^2. \tag{6.12}$$

Thus, when the following inequality holds

$$(\alpha_1^{\frac{1}{3}} + \alpha_2^{\frac{1}{3}})^3 (\beta_1^{\frac{1}{3}} + \beta_2^{\frac{1}{3}})^3 > (\phi^*)^2, \tag{6.13}$$

then the Hessian is always positive definite and the function (6.11) strictly convex implying that the minimum is unique. If the Hessian is not always unique, but (6.12) holds for minimums, then Theorem 6.2 applies. Again, when $\alpha = \beta$ and $\phi(1,2) = \phi(2,1)$, then under (6.13) the unique minimum is such that $z_1 = z_2$. For example, when $\alpha = \beta = (2,2)$ and $\phi(1,1) = 1, \phi(1,2) = \phi(2,1) = 2, \phi(2,2) = 3$, then $\phi^* = 0$ and, therefore, (6.13) holds. It means the minimum is unique, $z_1 = z_2$, and one can verify that

$$z_1 = z_2 = \frac{\sqrt{17} - 3}{2} \approx 0.561.$$

So the unique limit distribution in this case is $q \times q$, where $q = (z, 1-z)$. But when $\alpha = (2,3)$, $\beta = (3,2)$ and ϕ is as previously, then the minimum is again unique but since $\alpha \neq \beta$, we now have $z_1 \neq z_2$: $z_1 = \sqrt{6}-2 \approx 0.45$, $\quad z_2 = \sqrt{7}-2 \approx 0.645$. This means: the unique limit distribution is $q^1 \times q^2$, where $q^1 = (\sqrt{6}-2, 3-\sqrt{6})$, $q^2 = (\sqrt{7}-2, 3-\sqrt{7})$.

When $\alpha = \beta = (2,2)$, $\phi(1,1) = 4, \phi(1,2) = \phi(2,1) = 2, \phi(2,2) = 4$, then $\phi^* = 4$, but (6.13) still holds and therefore there is unique minimum: $z_1 = 0.5, z_2 = 0.5$. However, when the ϕ is as previously, but $\alpha = \beta = (0.25, 0.25)$, then (6.13) fails. It turns out that now the function is not convex and there are two minima:

$$(z_1 = \frac{2-\sqrt{2}}{4}, z_2 = \frac{2+\sqrt{2}}{4}), \quad (z_1 = \frac{\sqrt{2}+2}{4}, z_2 = \frac{\sqrt{2}-2}{4})$$

Observe that $z_2 = 1 - z_1$. Thus the limit measures are $q^1 \times q^2$ and $q^2 \times q^1$, where $q^1 = (z_1, z_2)$ and $q^2 = (z_2, z_1)$. These two product measures are different. Finally observe that in both cases (6.12) holds, so that by Theorem 6.2,

$$\nu_n \Rightarrow \frac{1}{2}\delta_{q^1 \times q^2} + \frac{1}{2}\delta_{q^2 \times q^1}.$$

The function

$$f_n(q) = \langle e^{-\frac{\phi}{n^\lambda}}, q\rangle^{n^\lambda} \prod_{l=1}^{2}\prod_{k=1}^{2}(q(k)^l)^{0.25-\frac{1}{n^{1-\lambda}}}$$

is symmetric, i.e. $f_n(z_1, z_2) = f_n(z_2, z_1)$ and then $\bar{\pi}_n \Rightarrow \frac{1}{2}\delta_{q^1 \times q^2} + \frac{1}{2}\delta_{q^2 \times q^1}$. Since now $r_q = q$, by Lemma 4.1, (4.7) holds, with

$$P^*(x_1, \ldots, x_m) = \frac{1}{2}\prod_{i=1}^{m} q^1 \times q^2(x_i) + \frac{1}{2}\prod_{i=1}^{m} q^2 \times q^1(x_i), \quad x_i \in \mathcal{X}^L.$$

6.3 An application: exact calculation of stationary distributions for epistatic fitness

Exchangeable sampling with a product of Dirichlet priors is an abstract model of sexual reproduction, with the constraint that the population is in linkage equilibrium. For a considerable class of fitness functions it is possible to compute the stationary distribution exactly, so that the model is potentially a tool for the investigation of the properties of genetic architectures. This line of research is beyond the scope of this paper, but we give one example of computation of the stationary distribution for a non-trivial fitness function in a multi-locus system.

Consider a the case where $\mathcal{X} = \{1,2\}^L$, and the fitness of a genome $x = (x^1, \ldots, x^L)$ is defined as

$$w(x^1, \ldots, x^L) = \begin{cases} 1+s & \text{if } x^1 = \cdots = x^L = 1 \\ 1 & \text{otherwise} \end{cases}$$

In other words, if a genome $(x^1, \ldots, x^L)$ is 'perfect', consisting only of the preferred alleles (represented as 1s), then it has a fitness advantage of s, but if even a single element is

imperfect, then there is no fitness advantage. It seems plausible that there are that there are combinations of well-separated alleles in biology that have a fitness advantage only if all alleles in the set are of the correct type; this is a simple model of such a situation.

To keep the notation simple, we assume that all independent components of breeding process ξ^l, $l = 1, \ldots, L$ have the same Dirichlet prior with parameter $\alpha = (\alpha_1, \ldots, \alpha_K)$. We shall denote the law of the breeding process by $P_\xi(\cdot|L, \alpha)$. When $(\xi_1, \ldots, \xi_n) \sim P_\xi(\cdot|L, \alpha)$, then the probability that the number of prefect genomes (vectors ξ_i consisting of ones, only) equals to n_c is

$$P_\xi(n_c, n - n_c; L, \alpha) := P\left(\sum_{i=1}^{n} I_{1+s}(w(\xi_i)) = n_c\right).$$

In words, $P_\xi(n_c, n - n_c; L, \alpha)$ denotes the probability of n_c perfect genomes and $n - n_c$ imperfect genomes in a population of n genomes, each of length L, sampled from a product of L Polya urn processes, each with prior parameter α.

We calculate $P_\xi(n_c, n - n_c; L, \alpha)$ using recursion on L. When $L = 1$, n_c is simply the number of 1s in the population of n, so

$$P_\xi(n_c, n - n_c; 1, \alpha) = \binom{n}{n_c} \frac{(\alpha_1)_{n_c}(\alpha_2)_{n-n_c}}{(|\alpha|)_n}.$$

It is not hard to show, using independence of $\xi^1, \ldots, \xi^L$ and exchangeability, that

$$P_\xi(n_c, n - n_c; L, \alpha) = \sum_{m=n_c}^{N} P_\xi(m, n - m; L - 1, \alpha) P_\xi(n_c, m - n_c; 1, \alpha)$$

Finally, the stationary distribution of the number of perfect genomes in a population of n is:

$$P_n(n_c, n - n_c; L, \alpha) = \frac{1}{Z(n, L, \alpha)} P_\xi(n_c, n - n_c; 1, \alpha)(1 + s)^{n_c} \tag{6.14}$$

where the normalising constant is

$$Z(n, L, \alpha) = \sum_{n_c=0}^{n} P_\xi(n_c, n - n_c; 1, \alpha)(1 + s)^{n_c}.$$

Thus in our model the stationary distribution of the number of perfect genomes is simple to derive without making any approximations. In general, because the processes $\xi^1, \ldots, \xi^L$ are independent, and as a result of the factorisation of the stationary distribution in equation 1.2, the stationary distribution for a fitness function w can be derived exactly using the sum-product algorithm if w can be represented as a tractable factor graph [17] over the L urn distributions. This introduces a new method for studying the effect of epistatic fitness on population structure.

7 Experiments

Let $K = 2$, $\phi(1) = 0$, $\phi(2) = \ln 6$, $\alpha_1 = 0.3$, $\alpha_2 = 0.7$. Let us find the limit measures q^* as in Theorem 6.1 in the following cases: $\lambda = 0$, $\lambda \in (0, 1)$ and $\lambda = 1$.

Case $\lambda = 0$: Then, as it can be easily checked by verifying (6.3) that $q^* = (3/5, 2/5)$. Since $\theta = \langle q^*, w\rangle = 2/3$, the measure r^* is as follows: $r_1^* = w(1)q^*(1)/\theta = 9/10$ and $r^*(2) = w(2)q^*(2)/\theta = 1/10$. Therefore the limit process governed by P^* is an i.i.d. process with measure r^*, and so due to the weight function, the average proportion of the first genotype has increased from 0.3 (according to the prior) to 0.9. Figure 1 illustrates the convergence.

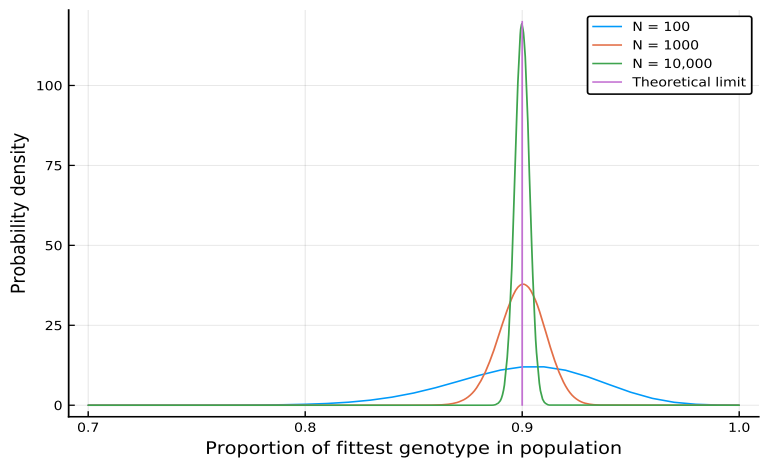

Figure 1: Density histogram of the fraction of the first type in the population, for a Dirichlet prior $\alpha = (0.3, 0.7)$, fitness $\phi = (0, \ln 6)$, and $\lambda = 0$, and three population sizes 10^2, 10^3, and 10^4. The histograms were constructed by recording the fraction of the first type in the population over 10^8 MCMC samples according to the process described in section 3.2

.

Case $\lambda \in (0, 1)$: The solution of (6.4) is

$$q^*(1) = \frac{\ln(6) - 1 + \sqrt{(\ln(6) - 1)^2 + 1.2\ln(6)}}{2\ln(6)} \approx 0.686. \tag{7.1}$$

Now $r^* = q^*$, thus we see that the limit process governed by P^* is an i.i.d. process with measure q^*, and so due to the weight function, the proportion of first genotype has increased from 0.3 (according to the prior) to 0.689. The increase is smaller than in the previous case. Figures 2, 3 and 4 illustrates the convergence for $\lambda = 0.25, 0.5, 0.75$, respectively. We see that although the limit is the same, the speed of convergence depends very much on λ. Let $q_m^* = (q_m(1), q_m(2))$ be the maximizer of

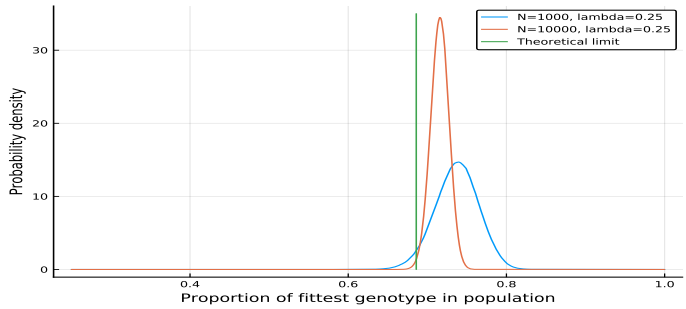

Figure 2: Case $\lambda = 0.25$: density histogram of the fraction of the first type in the population, for a Dirichlet prior $\alpha = (0.3, 0.7)$, fitness $\phi = (0, \ln 6)$, and $\lambda = 0.25$, and population sizes 10^3 and 10^4. The histograms were constructed by recording the fraction of the fitter allele in the population over 10^8 MCMC samples according to the process described in section 3.2

.

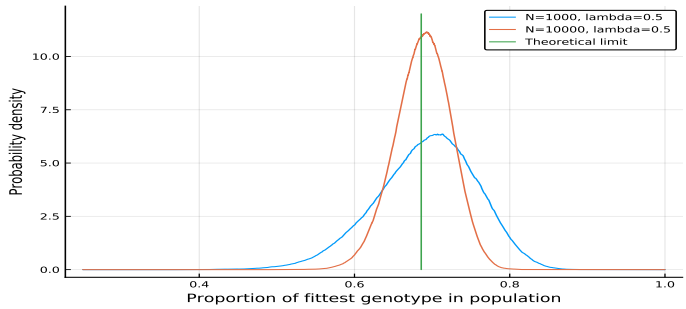

Figure 3: Case $\lambda = 0.5$: density histogram of the fraction of the first type in the population, for a Dirichlet prior $\alpha = (0.3, 0.7)$, fitness $\phi = (0, \ln 6)$, and population sizes 10^3 and 10^4. The histograms were constructed by recording the fraction of the fitter allele in the population over 10^8 MCMC samples according to the process described in section 3.2

.

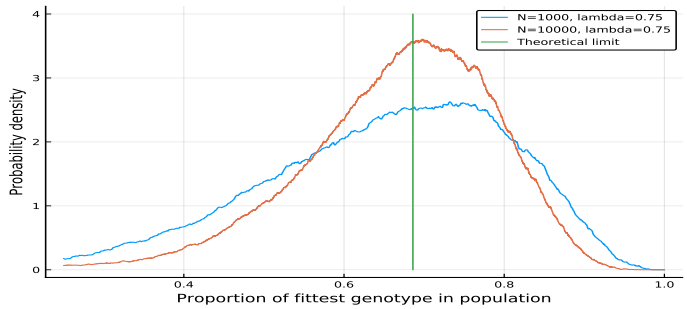

Figure 4: Case $\lambda = 0.75$: density histogram of the fraction of the first type in the population, for a Dirichlet prior $\alpha = (0.3, 0.7)$, fitness $\phi = (0, \ln 6)$, and population sizes 10^3 and 10^4. The histograms were constructed by recording the fraction of the fitter allele in the population over 10^{10} MCMC samples according to the process described in section 3.2

.

$$\ln\langle \exp[-\frac{\phi}{m}], q\rangle + \sum_k \frac{\alpha_k}{m} \ln q(k).$$

In our example $q_1(1) = 3/5$ (the case $\lambda = 0$) and

$$q_m(1) = \frac{b_m + \sqrt{b_m^2 + \frac{12}{10}(1+m)\big(1-(1/6)^{\frac{1}{m}}\big)(1/6)^{\frac{1}{m}}}}{2(1+m)\big(1-(1/6)^{\frac{1}{m}}\big)}, \text{ where } b_m = (m{+}0.3){-}(1/6)^{\frac{1}{m}}(1.3{+}m).$$

According to Proposition 6.1, in the process $m \to \infty$, $q_m(1)$ tends to $q^*(1)$ from (7.1), and one can easily verify that it is really so.

Case $\lambda = 1$: According to **2)** of Theorem 5.1, the limit measure $\bar{\pi}$ has density with respect to the Lebesgue'i measure on $[0, 1]$:

$$\bar{\pi}(q) = \frac{1}{Z} \exp[-\ln 6 \cdot q](1-q)^{0.3-1} q^{0.7-1},$$

where Z is the valuer of the moment generating function of Beta(0.7,0.3)-distributed random variable evaluated at $\ln 6$. In the density above, q stands for $q(2)$. Therefore, the limit proportion of the second genotype in the stochastic process governed by P^* is a random variable, its distribution has density $\bar{\pi}(q)$ as stated above. Figure 5 illustrates the convergence.

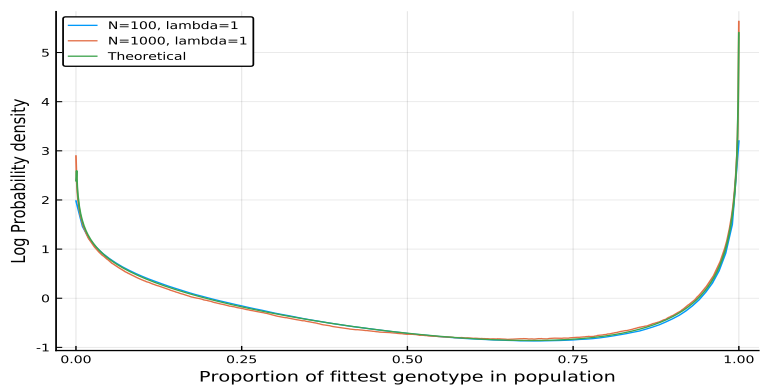

Figure 5: Case $\lambda = 1$: Density histogram (log scale) of the fraction of the first type in the population, for a Dirichlet prior $\alpha = (0.3, 0.7)$, fitness $\phi = (0, \ln 6)$, and population sizes 10^2 and 10^3. Note that because the concentration parameter is kept constant, the limit distribution is a reweighted β distribution, with infinities at 0 and 1: when $\lambda = 0$, the prior is the same for all N, and the log-fitness term is reduced by a factor of N, so that P_n never concentrates. The histograms were constructed by recording the fraction of the fitter allele in the population over 10^9 and 10^{10} MCMC samples respectively according to the process described in section 3.2

.

8 Conclusions

We have constructed MCMC algorithms that are similar to existing genetic algorithms. The 'breeding' consists of sampling from exchangeable distributions based on the Dirichlet distribution, and the 'selection' is essentially Metropolis-Hastings. The sequence of populations forms a reversible Markov chain that satisfies detailed balance conditions. We have exhibited two possible sampling distributions: more elaborate exchangeable sampling distributions are possible. The entire MCMC procedure is a population generalisation of Metropolis-Hastings. As far as we are aware, this is the first

implementable computational model of sexual reproduction that exactly satisfies detailed balance, and for which the stationary distribution can be written in closed form for arbitrary fitness functions.

We also explored some properties of the stationary distribution, and showed that for any fitness function there are three non-trivial limiting distributions for large population sizes, with two phase transitions. This is a first step towards a more general understanding of the interaction of the population size, fitness scaling, and mutation rate in genetic algorithms and evolutionary models.

Formulating a genetic model as a MCMC procedure opens a new research direction in using the many techniques developed in MCMC to achieve faster convergence to the stationary distribution using different MCMC kernels.

We have shown that the stationary distribution is unaffected by multiplicative noise in fitness evaluations. This has been suggested by, for example, [19], but our techniques allow a proof of this effect.

Finally there is a more general conclusion from our analysis. For many years, since [15] and [12], a widely suggested folk-motivation for genetic algorithms has been that because they are inspired by natural biological evolution, and because evolution has produced the variety of life on earth, genetic algorithms should be in some sense generally effective. Our analysis makes it clear that genetic algorithms are more closely related to conventional MCMC methods for non-parametric Bayesian inference than has previously been recognised.

Appendix

Proof of claim 4.1 **Proof.** First note that

$$\sup_{q\in\mathcal{P}} |\langle w_n, q\rangle - \langle w, q\rangle| \leq \|w_n - w\|\|q\| \leq \|w_n - w\| \to 0.$$

Now use the fact that if f_n, f, g_n, g are nonnegative functions such that $\sup_x |f_n(x) - f(x)| \to 0$, $\sup_x |g_n(x) - g(x)| \to 0$, $\inf_x g(x) = g_* > 0$ and f, g are bounded above, then

with $g^* = \sup_x g(x)$ and $f^* = \sup_x f(x)$

$$\begin{aligned}\sup_x \Big|\frac{f_n(x)}{f(x)} - \frac{g_n(x)}{g(x)}\Big| &= \sup_x \Big|\frac{f_n(x)g(x) - g_n(x)f(x)}{g_n(x)g(x)}\Big| \\ &\le \sup_x \Big|\frac{(f_n(x) - f(x))g(x)}{g_n(x)g(x)}\Big| + \sup_x \Big|\frac{(g_n(x) - g(x))f(x)}{g_n(x)g(x)}\Big| \\ &\le \sup_x \Big|\frac{(f_n(x) - f(x))g^*}{g_n(x)g(x)}\Big| + \sup_x \Big|\frac{(g_n(x) - g(x))f^*}{g_n(x)g(x)}\Big| \to 0,\end{aligned}$$

because for n big enough $g_n(x)g(x) \ge \frac{g_*^2}{2}$ for every x. Take $x = q$, $f_n(q) = w_n(k)q(k)$, $f(q) = w(k)q(k)$, $g_n(q) = \langle w_n, q\rangle$ and $g(q) = \langle w, q\rangle$. Then $g_* = w(K) > 0$, $f^* = g^* = w(1)$ and so (4.8) follows. ■

8.1 Proof of Proposition 5.1

Recall that f_n and f are continuous and bounded functions on $\mathcal{P}$ so that $\|f_n\|_\infty < \infty$ and $\|f\|_\infty < \infty$. By assumption, π is a finite measure. Since f_n converges to f uniformly, it follows that $\|f_n - f\|_\infty \to 0$ and so $\|f_n\|_\infty \to \|f\|_\infty$. For every m,

$$|\|f_n\|_m - \|f\|_m| \le \|f_n - f\|_m \le \pi(\mathcal{P})^{\frac{1}{m}}\|f_n - f\|_\infty \to 0.$$

Since $\|f\|_{m_n} \to \|f\|_\infty$, we have

$$\begin{aligned}\big|\|f_n\|_{m_n} - \|f\|_\infty\big| &\le \big|\|f_n\|_{m_n} - \|f\|_{m_n}\big| + \big|\|f\|_{m_n} - |f\|_\infty\big| \\ &\le \|f_n - f\|_{m_n} + \big|\|f\|_{m_n} - |f\|_\infty\big| \le \pi(\mathcal{P})^{\frac{1}{m_n}}\|f_n - f\|_\infty + \big|\|f\|_{m_n} - |f\|_\infty\big| \to 0.\end{aligned}$$

Now fix $\delta > 0$. Since $\mathcal{P}_\delta^* := \{q : f(q) > \|f\|_\infty - \delta\}$, we have $\mathcal{P} - \mathcal{P}_\delta^* = \{q : f(q) \le \|f\|_\infty - \delta\}$. Define $\delta' := \delta/\|f\|_\infty$. Then

$$\begin{aligned}\sup_{q\in\mathcal{P}-\mathcal{P}_\delta^*} \frac{f_n(q)}{\|f_n\|_{m_n}} &= \sup_{q\in\mathcal{P}-\mathcal{P}_\delta^*} \frac{f(q) + (f_n(q) - f(q))}{\|f_n\|_{m_n}} = \sup_{q\in\mathcal{P}-\mathcal{P}_\delta^*} \frac{f(q) + (f_n(q) - f(q))}{\|f\|_\infty}\frac{\|f\|_\infty}{\|f_n\|_{m_n}} \\ &\le \sup_{q\in\mathcal{P}-\mathcal{P}_\delta^*} \frac{f(q)}{\|f\|_\infty}\frac{\|f\|_\infty}{\|f_n\|_{m_n}} + \frac{\|f_n - f\|_\infty}{\|f_n\|_{m_n}} \le 1 - \frac{\delta'}{2},\end{aligned}$$

provided n is big enough. Thus,

$$\sup_{q\in\mathcal{P}-\mathcal{P}_\delta^*} h_n(q) \le \Big(1 - \frac{\delta'}{2}\Big)^{m_n} \to 0,$$

so that $\nu_n(\mathcal{P}_\delta^*) \to 1$. We now argue that when $\mathcal{P}^* = \{q^*\}$, then for any $\epsilon > 0$ there exists $\delta > 0$ so that

$$\mathcal{P}_\delta^* \subset B(q^*, \epsilon), \tag{8.1}$$

where $B(q^*, \epsilon)$ is a ball in Euclidean sense. If, for an $\epsilon > 0$, such a $\delta > 0$ would not exists, then there would exist a sequence $q_n \to q$ such that $f(q_n) \nearrow f(q)$, but $\|q_n - q\| \ge \epsilon$. Since $\mathcal{P}$ is compact, along a subsequence $q_{n'} \to q$ and by continuity $f(q_{n'}) \to f(q)$. On the other hand $\|q - q^*\| > \epsilon$ and that would contradict the uniqueness of q^*. Therefore (8.1) holds and so for any $\epsilon > 0$, it holds that $\nu_n\big(B(q^*, \epsilon)\big) \to 1$. From the definition of the weak convergence, it now follows that $\nu_n \Rightarrow \delta_{q^*}$.

8.2 Proof of Lemma 6.1

1) To find

$$q^* = \arg\max_{q\in\mathcal{P}} \; [\ln\langle e^{-\phi}, q\rangle + \sum_k \alpha_k \ln q(k)], \tag{8.2}$$

we define Lagrangian

$$L(q,\beta) = \ln\langle e^{-\phi}, q\rangle + \sum_k \alpha_k \ln q(k) - \beta(\sum_k q(k) - 1)$$

(here β is a scalar) and maximize $L(q,\beta)$ over $q > 0$ (all entries are positive). Taking partial derivatives with respect to $q(k)$, we have

$$\frac{e^{-\phi(k)}}{\langle e^{-\phi}, q\rangle} + \frac{\alpha_k}{q(k)} = \beta, \quad \Rightarrow \quad \frac{e^{-\phi(k)} q(k)}{\langle e^{-\phi}, q\rangle} + \alpha_k = q(k)\beta \quad \forall k.$$

With $|\alpha| = \sum_k \alpha_k$, we have thus $\beta = 1 + |\alpha|$ and so the solution q^* satisfies the set of equalities

$$q^*(k) = \frac{1}{1+|\alpha|}\Big(\frac{e^{-\phi(k)} q^*_k}{\langle e^{-\phi}, q^*\rangle} + \alpha_k\Big), \quad \forall k. \tag{8.3}$$

Now with $w(k) = e^{-\phi(k)}$ define parameter $\theta := \langle w, q^*\rangle$ and rewrite (8.3) as follows

$$q^*(k) = \frac{\alpha_k}{(1+|\alpha|) - \frac{w(k)}{\theta}} \quad k = 1, \ldots, K. \tag{8.4}$$

We see that amongst the probability vectors satisfying $\langle w, q^*\rangle = \theta$, the solution is unique. Since $\alpha_k > 0$ for every k, it is easy to see that there is only one parameter θ such that the right hand side of (8.4) would be a probability measure: if $\theta' > \theta$, then for every k, we have

$$\frac{\alpha_k}{(1+|\alpha|) - \frac{w(k)}{\theta}} > \frac{\alpha_k}{(1+|\alpha|) - \frac{w(k)}{\theta'}}.$$

Therefore a solution of (8.2) is unique vector q^* given by (8.4), where $\theta = \langle w, q^*\rangle$.

2) To find

$$q^* = \arg\max_{q\in\mathcal{P}} [-\langle \phi, q\rangle + \sum_k \alpha_k \ln q(k)], \tag{8.5}$$

we define Lagrangian

$$L(q,\beta) = -\langle \phi, q\rangle + \sum_k \alpha_k \ln q(k) - \beta(\sum_k q(k) - 1).$$

Partial derivatives with respect to $q(k)$ give us the equalities

$$-\phi(k) + \frac{\alpha_k}{q(k)} = \beta \quad \forall k \quad \Rightarrow \quad -\langle \phi, q\rangle + |\alpha| = \beta.$$

Therefore, the inequalities for $q^*(k)$ are

$$q^*(k) = \frac{\phi(k)q^*(k) - \alpha_k}{\langle \phi, q^* \rangle - |\alpha|} = \frac{\phi(k)q^*(k) - \alpha_k}{\theta - |\alpha|}, \quad \theta := \langle \phi, q^* \rangle. \tag{8.6}$$

After rewriting (8.6), we obtain

$$q^*(k) = \frac{\alpha_k}{\phi(k) + |\alpha| - \theta}, \quad k = 1, \ldots, K,$$

Thus, there cannot be two solutions having the same θ. As in the case **1)**, it is easy to see that when $\alpha_k > 0$ there is only one θ so that (8.2) sums up to one. Therefore, the solution to the problem (8.5) is unique. Note that the solution is independent of λ.

Acknowledgments

Research by Jüri Lember is supported by Estonian Institutional research funding IUT34-5 and PRG 865. Research by Chris Watkins is supported by grant number (FQXi Grant number FQXi-RFP-IPW-1913) from the Foundational Questions Institute and Fetzer Franklin Fund, a donor advised fund of Silicon Valley Community Foundation.

References

[1] S. Baluja and R. Caruana. Removing the genetics from the standard genetic algorithm. In *ICML*, pages 38–46. Morgan Kauffman Publishers, Inc., 1995.

[2] Shumeet Baluja. Genetic algorithms and explicit search statistics. In *Advances in Neural Information Processing Systems*, pages 319–325, 1997.

[3] E.B. Baum, D. Boneh, and C. Garrett. Where genetic algorithms excel. *Evolutionary Computation*, 9(1):93–124, 2001.

[4] Robert Bialowons and Ellen Baake. Ancestral processes with selection: Branching and moran models. *Banach Center Publications*, 80, 2008.

[5] James Franklin Crow and Motoo Kimura. *An introduction to population genetics theory.* New York, Evanston and London: Harper & Row, Publishers, 1970.

[6] Richard Durrett. *Probability models for DNA sequence evolution.* Springer Science & Business Media, 2008.

[7] Lloyd Elliott and Yee Whye Teh. Scalable imputation of genetic data with a discrete fragmentation-coagulation process. In *Advances in Neural Information Processing Systems 25*, pages 2861–2869, 2012.

[8] Alison M Etheridge and Robert C Griffiths. A coalescent dual process in a moran model with genic selection. *Theoretical population biology*, 75(4):320–330, 2009.

[9] Warren J Ewens. *Mathematical population genetics. 2nd.* Springer New York:, 2004.

[10] Shui Feng. *The Poisson-Dirichlet distribution and related topics: models and asymptotic behaviors.* Springer Science & Business Media, 2010.

[11] Stuart Geman and Donald Geman. Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, (6):721–741, 1984.

[12] D.E. Goldberg. *Genetic algorithms in search, optimization, and machine learning.* Addison-wesley, 1989.

[13] Heikki Haario and Eero Saksman. Simulated annealing process in general state space. *Advances in Applied Probability*, 23(4):866–893, 1991.

[14] Nils Lid Hjort, Chris Holmes, Peter Muller, and Stephen Walker. *Bayesian nonparametrics.* Number 28. Cambridge University Press, 2010.

[15] J.H. Holland. *Adaptation in natural and artificial systems.* University of Michigan press, 1975.

[16] John Frank Charles Kingman. *Mathematics of genetic diversity*, volume 34. SIAM, 1980.

[17] Frank R Kschischang, Brendan J Frey, and H-A Loeliger. Factor graphs and the sum-product algorithm. *IEEE Transactions on information theory*, 47(2):498–519, 2001.

[18] P.A.P. Moran. *The statistical processes of evolutionary theory.* Clarendon Press; Oxford University Press., 1962.

[19] Gregory Morse and Kenneth O Stanley. Simple evolutionary optimization can rival stochastic gradient descent in neural networks. In *Proceedings of the Genetic and Evolutionary Computation Conference 2016*, pages 477–484. ACM, 2016.

[20] Christina A Muirhead and John Wakeley. Modeling multiallelic selection using a moran model. *Genetics*, 182(4):1141–1157, 2009.

[21] R.M. Neal. Markov chain sampling methods for Dirichlet process mixture models. *Journal of computational and graphical statistics*, pages 249–265, 2000.

[22] Dominik Schrempf and Asger Hobolth. An alternative derivation of the stationary distribution of the multivariate neutral wright–fisher model for low mutation rates with a view to mutation rate estimation from site frequency data. *Theoretical population biology*, 114:88–94, 2017.

[23] Yee Whye Teh and Michael I Jordan. Hierarchical Bayesian nonparametric models with applications. In *Bayesian nonparametrics*, pages 158–207. Cambridge University Press, 2010.

[24] Claus Vogl and Florian Clemente. The allele-frequency spectrum in a decoupled moran model with mutation, drift, and directional selection, assuming small mutation rates. *Theoretical population biology*, 81(3):197–209, 2012.

[25] M.D. Vose. *The simple genetic algorithm: foundations and theory.* The MIT Press, 1999.

[26] Chris Watkins and Yvonne Buttkewitz. Sex as Gibbs sampling: a probability model of evolution. *arXiv preprint arXiv:1402.2704*, 2014.