

Research



Cite this article: McAvoy A, Adlam B, Allen B, Nowak MA. 2018 Stationary frequencies and mixing times for neutral drift processes with spatial structure. *Proc. R. Soc. A* **474**: 20180238. <http://dx.doi.org/10.1098/rspa.2018.0238>

Received: 9 April 2018

Accepted: 25 September 2018

Subject Areas:

applied mathematics, mathematical modelling

Keywords:

evolution, mixing time, neutral drift, structured population

Author for correspondence:

Martin A. Nowak

e-mail: martin_nowak@harvard.edu

Stationary frequencies and mixing times for neutral drift processes with spatial structure

Alex McAvoy¹, Ben Adlam^{1,2}, Benjamin Allen^{1,5},
Martin A. Nowak^{1,3,4}

¹Program for Evolutionary Dynamics, ²School of Engineering and Applied Sciences, ³Department of Organismic and Evolutionary Biology, and ⁴Department of Mathematics, Harvard University, Cambridge, MA 02138, USA

⁵Department of Mathematics, Emmanuel College, Boston, MA 02115, USA

AM, 0000-0002-9110-4635; BA, 0000-0002-9746-7613

We study a general setting of neutral evolution in which the population is of finite, constant size and can have spatial structure. Mutation leads to different genetic types (traits), which can be discrete or continuous. Under minimal assumptions, we show that the marginal trait distributions of the evolutionary process, which specify the probability that any given individual has a certain trait, all converge to the stationary distribution of the mutation process. In particular, the stationary frequencies of traits in the population are independent of its size, spatial structure and evolutionary update rule, and these frequencies can be calculated by evaluating a simple stochastic process describing a population of size one (i.e. the mutation process itself). We conclude by analysing mixing times, which characterize rates of convergence of the mutation process along the lineages, in terms of demographic variables of the evolutionary process.

1. Introduction

At the heart of evolutionary theory lies the question of how individual-level properties affect the long-run composition of a population. Microscopic quantities including mutation rates, selective differences between competing types, and number of interaction partners can

have profound effects on the evolutionary success of genetic types (traits) [1–12], but the precise relationship between local and global properties is often difficult to quantify. Neutral drift, which involves no selective differences between the traits, provides a natural setting in which to study this relationship [13–17]. Neutral population dynamics are also important to the study of populations with selection because the latter can be viewed as a perturbation of the former provided selection is sufficiently weak [18–28]. The purpose of this article is to study the long-term distribution of traits arising along the lineages of a neutral population, as a function of the population's size, structure, evolutionary update rule and underlying mutation process.

To motivate the discussion, let us first consider a population of size N of haploid organisms reproducing asexually. Each organism has a type, which might indicate a particular trait. We assume that all types are equivalent from a reproductive standpoint. Let S be the set of all possible types, which we initially assume is finite (or at least discrete). In the Moran process [29], which is a simple example of a finite-population evolutionary process, an individual is first chosen to reproduce. The offspring then replaces another individual chosen for death. When an individual of type s reproduces in the population, mutation is possible, and the offspring acquires type s' with probability $M_{s,s'} \in [0, 1]$. This procedure repeats ad infinitum, and the resulting process can be used as a model to study evolution's effects on the long-run composition of the population.

Mutations can be represented by a stochastic matrix, M , which gives rise to a Markov chain on the set of possible types, S . The resulting stochastic process is called 'mutation process' [30]. In the mutation process, a transition from s to s' occurs with probability $M_{s,s'}$. When evolution is neutral, there is a meaningful comparison between the evolutionary process and the mutation process since then neither one involves different types reproducing at different rates. In effect, the mutation process is just an evolutionary process in a population of size $N = 1$. A natural question to ask is whether population size, structure or update rule affect average trait frequency relative to the mutation process. That is, is the average frequency of type s in a population of any size, $N > 1$, different from that of a population of size $N = 1$?

Suppose that there are n possible types and that a mutation involves switching to a new type uniformly at random. Specifically, with probability u , a mutation to a random type occurs, and with probability $1 - u$, the offspring acquires the parental type. This kind of mutation, which is common in evolutionary game theory, is symmetric in the sense that it acts on all competing types in the same way. With mutations of this form, all types are all equally abundant in the stationary distribution [19,20]. Stated differently, in a neutral evolutionary process with symmetric mutations among n types in a population of size N , the average frequency of each type ($1/n$) is the same as one would find in a population of size $N = 1$. We extend this result to any mutation process, showing that the average frequency of a type in a population (actually, even along any lineage) of any size is the same as that in a population of size $N = 1$. Consequently, dispersal patterns and population size have no effect on the average frequencies of the traits.

In some ways, this result is not surprising, and special cases have indeed been known for some time [31]. When there are no selective differences between the traits, there is evidently nothing in the population that drives one trait to a higher average frequency in a population than in the original mutation process. However, this reasoning does not constitute a proof. Given the increasing interest in using neutral frequency as a baseline measure to understand strategy selection, we provide a general proof here. An interesting consequence of the proof is that it allows one to quantify the rate of convergence of the trait distribution in terms of simple demographic variables of the process. This rate, which we characterize in terms of mixing times, exhibits much more interesting behaviour than the limit itself (which is not affected by demographic components).

Our results apply to not only any population structure but also to any trait space, S . In other words, a player's trait could be an element of a finite set, a denumerable (countably infinite) set, or even an uncountably infinite set. Finite sets, even those consisting of just two types, have traditionally been the focus of many models in evolutionary game theory [32]. Other frameworks, such as adaptive dynamics [33–40] and even some within evolutionary game theory [41–50],

allow for continuous trait spaces. Accordingly, we make no assumption that there are only finitely many traits.

2. Neutral evolution

Before treating stationary trait frequencies in neutral evolution, we first turn to a generic way of describing neutral evolution itself. Allen & Tarnita [26] provide a method for dealing with evolution in finite populations of fixed size and structure. This general framework applies also to processes with selective differences between the types, but here we recall just the portion that applies to neutral processes. For simplicity, we introduce this framework in the context of haploid, asexually reproducing populations, but we also discuss how it can be extended to non-haploid populations with sexual reproduction.

The atom of the mathematical framework we consider is a ‘replacement event’ [26]. A replacement event is a pair, (R, α) , where $R \subseteq \{1, \dots, N\}$ is the set of individuals (or locations) to be replaced, and $\alpha : R \rightarrow \{1, \dots, N\}$ is a map where $\alpha(i)$ indicates that individual i is replaced by the offspring of individual $\alpha(i)$. A neutral evolutionary process may then be described by a probability distribution over all possible replacement events, $\{p_{(R, \alpha)}\}_{(R, \alpha)}$. The probabilities $\{p_{(R, \alpha)}\}_{(R, \alpha)}$ encode both the population structure and the demographic update rule (figure 1). This process is neutral because the distribution over replacement events is the same at each time step, regardless of the state [51]. The population structure and update rule are fixed, in that the replacement events at different points in time are chosen independently from the same distribution. While this framework describes a broad variety of neutral evolutionary processes, there is one further ‘unity’ assumption that must be made, which we recall in slightly modified form from [26]:

(unity) for each $i = 1, \dots, N$, there exists a sequence of replacement events, $\{(R_k, \alpha_k)\}_{k=1}^\ell$, such that (i) $p_{(R_k, \alpha_k)} > 0$ for each $k = 1, \dots, \ell$ and (ii) for each j , we have $\tilde{\alpha}_1 \circ \dots \circ \tilde{\alpha}_\ell(j) = i$, where the mapping $\tilde{\alpha}_k : \{1, \dots, N\} \rightarrow \{1, \dots, N\}$ is defined by $\tilde{\alpha}_k(m) = \alpha_k(m)$ for $m \in R_k$ and $\tilde{\alpha}_k(m) = m$ for $m \notin R_k$.

This assumption formalizes the notion that the population evolves as a coherent unit, and that every individual can produce a lineage that takes over the entire population. Importantly, it does not imply that the population structure is trivial. Many interesting population structures, including heterogeneous graphs, sets and subdivided populations with migration, satisfy the unity condition; we refer the reader to [26] for further examples. In example 2.2, we give an example of a population that does not satisfy the unity condition.

At each point in time, every individual, $i = 1, \dots, N$, has a genetic type, $s_i \in S$. For the moment, we assume that S is finite, but in the next section this assumption is relaxed. The current population state is given by the vector $\mathbf{s} = (s_1, \dots, s_N) \in S^N$. Let M be the transition matrix for a mutation process on S . Transitions from a given population state $\mathbf{s} \in S^N$ occur as follows: first a replacement event (R, α) is chosen from the distribution $\{p_{(R, \alpha)}\}_{(R, \alpha)}$. The probability that individual $i \in R$ is of type s after this update is $M_{s, \alpha(i), s}$. If $i \notin R$, the probability that individual i is of type s is $\delta_{s_i, s}$ (that is, individuals who are not replaced retain their type). This process defines a Markov chain on S^N , which we denote by $\{\mathbf{S}(T)\}_{T=0}^\infty = \{(S_1(T), \dots, S_N(T))\}_{T=0}^\infty$. Note that we use t to indicate time in the underlying mutation process, $\{X(t)\}_{t=0}^\infty$, and T to denote time (i.e. number of update steps) in the evolutionary process, $\{\mathbf{S}(T)\}_{T=0}^\infty$; both t and T are discrete. The number of updates in the mutation process along a lineage is bounded from above by the number of update steps in the overall process, but in general these two measures of time need not necessarily coincide.

The distribution $\{p_{(R, \alpha)}\}_{(R, \alpha)}$ gives rise to several useful demographic variables [26], namely

$$e_{ij} := \sum_{\substack{(R, \alpha) \\ j \in R, \alpha(j)=i}} p_{(R, \alpha)}; \quad b_i := \sum_{j=1}^N e_{ij}; \quad d_i := \sum_{j=1}^N e_{ji}, \quad (2.1)$$

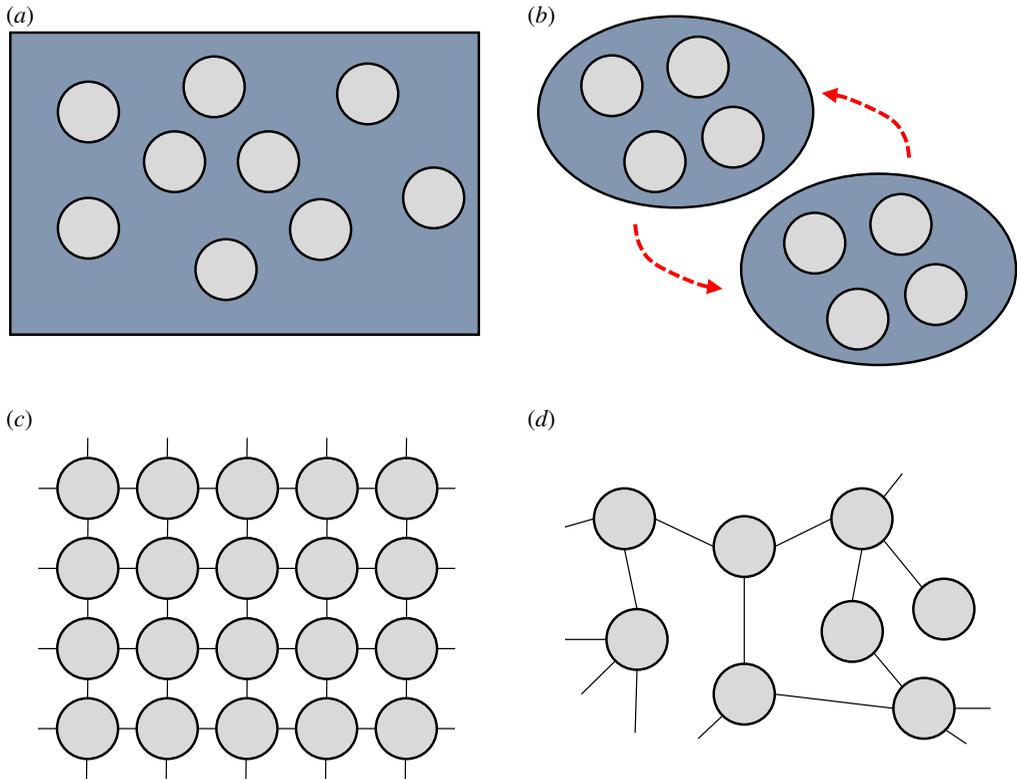


Figure 1. Four examples of evolving populations, ordered by increasing complexity of the spatial structure. In (a), the population is unstructured. (b) Depicts a subdivided (or ‘deme-structured’) population that consists of unstructured subpopulations (blue) with migration between them (red arrows). (c) A regular grid (square lattice) in which all players have exactly four neighbours. (d) A heterogeneous graph-structured population. Each individual resides on the vertex of a graph, and links indicate who is a neighbour of whom. The number of neighbours can vary from individual to individual, which results in structural asymmetries. In general, if the graph indicates an offspring-dispersal structure, then $p_{(R,\alpha)} > 0$ only if $\alpha(i)$ is a neighbour of i whenever $i \in R$. (Online version in colour.)

where e_{ij} is the probability that i transmits its offspring to j ; b_i is the birth rate of i ; and d_i is the death rate of i .

The way in which mutations are incorporated is essentially the same as they are in [26], except that here M and S can be arbitrary and mutations need not be symmetric with respect to the traits.

Theorem 2.1. *If the unity condition holds, then $\{S(T)\}_{T=0}^{\infty}$ is ergodic whenever $\{X(t)\}_{t=0}^{\infty}$ is ergodic.*

Proof. Consider two states, $\mathbf{s}, \mathbf{s}' \in S^N$. If $\{X(t)\}_{t=0}^{\infty}$ is ergodic, then there exists m_0 such that $(M^m)_{r,s} > 0$ whenever $r, s \in S$ and $m \geq m_0$. If $m_0 = 1$, then T is positive. Otherwise, by the unity condition, we can find an ordered sequence, $(R_1, \alpha_1), \dots, (R_{m_0}, \alpha_{m_0})$ with $p_{(R_k, \alpha_k)} > 0$ for each k , together with a collection $(i_0, i_1, \dots, i_{m_0})$ such that $\alpha_k(i_{k-1}) = i_k$ for each $k = 1, \dots, m_0$. After starting in state \mathbf{s} , and given this sequence of replacement events, the probability that player i_0 has type s is $M_{s_{i_0}, s}^{m_0} > 0$. By the unity condition, i_0 can propagate its offspring to all other nodes in a finite number of steps; together with the fact that M^m is positive whenever $m \geq m_0$, we see that there is a positive probability of reaching state \mathbf{s}' . Thus, $\{S(T)\}_{T=0}^{\infty}$ is irreducible. Aperiodicity of $\{S(T)\}_{T=0}^{\infty}$ follows from essentially the same argument because, after a similarly chosen sequence of replacements, the probability of staying in the state in which the chain started is positive. ■

If the unity condition is not satisfied, then theorem 2.1 need not necessarily hold, as the following example illustrates:

Example 2.2 (Line graph). As an example of a process that does not satisfy the unity condition, consider a population arranged on a line. If reproduction and replacement flow in only one direction, then player ‘1,’ i.e. the player at the ‘beginning’ of the line, is never replaced. Therefore, for each $k \in S$, there exists a stationary distribution, $\mu^{(k)}$, for $\{\mathbf{S}(T)\}_{T=0}^{\infty}$ with $\mu^{(k)}(s_1 = k) = 1$. In particular, there is more than one stationary distribution even when $\{X(t)\}_{t=0}^{\infty}$ is ergodic, so theorem 2.1 does not hold in this case.

When the unity condition holds and $\{X(t)\}_{t=0}^{\infty}$ is ergodic, theorem 2.1 implies that $\{\mathbf{S}(T)\}_{T=0}^{\infty}$ has a unique stationary distribution, μ . The next section establishes the simple fact that, if π is the stationary distribution of the mutation process, M , then $\mu(s_i = k) = \pi_k$ for every $k \in S$. When $\{\mathbf{S}(T)\}_{T=0}^{\infty}$ has more than one stationary distribution, it need not be the case that every such distribution satisfies $\mu(s_i = k) = \pi_k$ for every $k \in S$, which we illustrate using an irreducible mutation chain that is not ergodic:

Example 2.3 (Mutation process is irreducible but not ergodic). Suppose that $S = \{1, 2\}$ and let M be the matrix

$$M := \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}. \quad (2.2)$$

$\{X(t)\}_{t=0}^{\infty}$ is irreducible but periodic, with period 2, and its (unique) stationary distribution is $(1/2, 1/2)$. Consider a birth–death process on a star graph. Each player is chosen uniformly-at-random to reproduce. If a peripheral player reproduces, their offspring is subjected to the mutation operator and propagated to the central node. If the player at the central node reproduces, then this player propagates an offspring to every peripheral location (and, again, each offspring is subjected to the mutation operator). It is easy to see that the state with type 1 at the central node and type 2 at all of the peripheral nodes is stationary (figure 2). For this stationary distribution, $(1/N) \sum_{i=1}^N \mu(s_i = 1) = 1/N$, which is not equal to $\pi_1 = 1/2$ when $N > 2$.

The following result is the main tool we need to show that population size and structure do not influence stationary trait frequencies. The essence of this result is that for any $t \geq 0$, if one looks sufficiently far into the future, then every individual in the population will have a lineage of length at least t :

Lemma 2.4. Let B_i^T be the number of birth events after T updates in the lineage leading to individual i (a random variable). Then, for any $i = 1, \dots, N$ and any $t \geq 0$, we have $\lim_{T \rightarrow \infty} \mathbf{P}[B_i^T \geq t] = 1$.

Proof. The probability that i is replaced in any given update step (i.e. the death rate of i) is d_i (equation (2.1)), which is the same at every update since the process is neutral. Let $d_* := \min_{1 \leq i \leq N} d_i$ and $d^* := \max_{1 \leq i \leq N} d_i$. Since an individual is replaced with probability at most d^* and not replaced with probability at most $1 - d_*$ in each step, the probability that there are exactly t birth events by time T on the lineage leading to i satisfies

$$\mathbf{P}[B_i^T = t] \leq \binom{T}{t} (1 - d_*)^{T-t} (d^*)^t. \quad (2.3)$$

Therefore, provided $d_* > 0$, we see that $\lim_{T \rightarrow \infty} \mathbf{P}[B_i^T \geq t] = 1$ for any $t \geq 0$. ■

In the next section, we use this simple fact to derive the long-run trait frequencies.

3. Marginal distributions and stationary frequencies

Although the assumption that S is finite is reasonable in many cases, there are also scenarios in which one would like to consider continuous trait spaces. To capture a general notion of a trait, we let S be a measurable space, which contains as special cases finite, denumerable and uncountably infinite trait spaces.

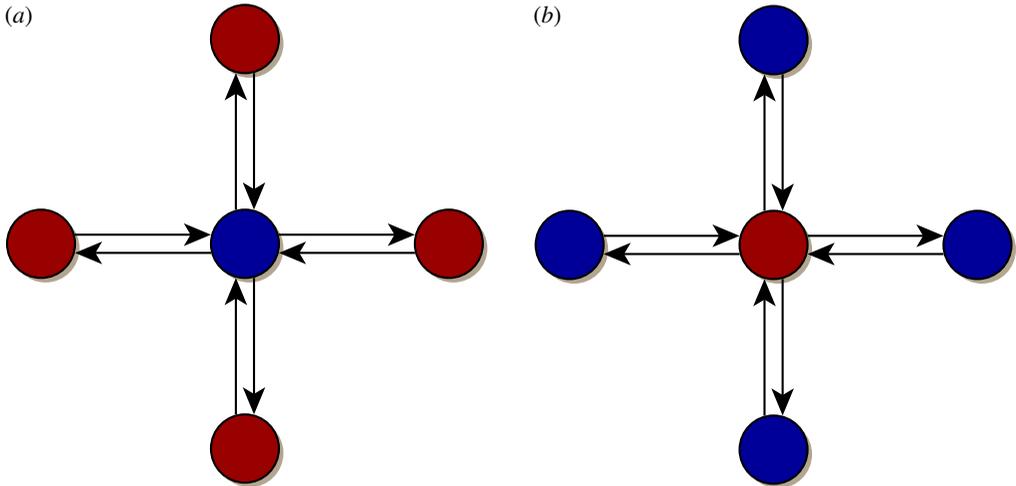


Figure 2. A birth–death process on a star graph of size $N = 5$. When a player reproduces, the offspring has the opposite type of the parent (i.e. if the parent is blue, the child is red; if the parent is red, the child is blue). When a peripheral individual reproduces, the central individual dies and is replaced by the offspring. If the central individual reproduces, then all four peripheral players die and are replaced by offspring of the central player. Panels (a) and (b) give stationary states, which are unaffected by the evolutionary process. The frequency of blue is $1/5$ in (a) and $4/5$ in (b), while the frequency of blue in the stationary distribution of the mutation process (i.e. in a population of size $N = 1$) is $1/2$. Therefore, even though $\{X(t)\}_{t=0}^{\infty}$ is irreducible, the evolutionary process has multiple stationary distributions due to the periodicity of M , and these stationary distributions need not all exhibit the same trait frequencies as the original mutation process. However, there does exist a stationary distribution for the birth–death process with trait frequency $1/2$ for each of blue and red, namely the distribution that assigns probability $1/2$ to state (a) and $1/2$ to state (b). (Online version in colour.)

In what follows, we denote by \mathbf{P}_ν and \mathbf{E}_ν the distribution and expectation, respectively, of a random variable that depends on another distribution, ν . Such is the case when ν is the initial distribution of a Markov chain and the random variable under consideration is the time- t state of the chain. An element $s \in S$ in place of ν (e.g. \mathbf{P}_s or \mathbf{E}_s) indicates that the starting point of the Markov chain is at s .

When dealing with a general trait space, we can no longer necessarily represent a mutation chain by a transition matrix. Instead, such a mutation process is described by a transition kernel. Let $\mathcal{F}(S)$ be a σ -algebra of subsets on S and denote by $\Delta(S)$ the space of probability measures on S . A mutation process on S , $\{X(t)\}_{t=0}^{\infty}$, is then defined by a Markov kernel, $\kappa : S \rightarrow \Delta(S)$, where for $s \in S$ and $E \in \mathcal{F}(S)$, $\kappa(s, E)$ is the probability that the chain is in E after being in state s . If S is finite and the transition matrix for the mutation process is M , then $M_{s,s'} = \kappa(s, \{s'\})$ for each $s, s' \in S$. To extend the notion of ergodicity to a Markov chain on a general state space, one needs the notion of a Harris chain, which we recall from [52]:

Definition 3.1. A Markov chain, $\{X(t)\}_{t=0}^{\infty}$, on S with kernel κ is a Harris chain if there exist $A, B \in \mathcal{F}(S)$, $\varepsilon > 0$, a function $q : A \times B \rightarrow \mathbb{R}$ with $q(s, s') \geq \varepsilon$ for every $s \in A$ and $s' \in B$, and $\rho \in \Delta(B)$ such that

- (i) $\mathbf{P}_s[\tau_A < \infty] > 0$ for every $s \in S$, where $\tau_A = \inf\{t \in \{0, 1, 2, \dots\} \mid X(t) \in A\}$;
- (ii) $\kappa(s, C) \geq \int_{s' \in C} q(s, s') d\rho(s')$ for every $s \in A$ and $C \in \mathcal{F}(B)$.

Furthermore, recurrence and aperiodicity for Harris chains are defined as follows [53]:

Definition 3.2. A Harris chain on S , $\{X(t)\}_{t=0}^{\infty}$, is recurrent if $\mathbf{P}_s[\tau_A < \infty] = 1$ for every $s \in A$. A recurrent Harris chain is aperiodic if for every $s \in S$, there exists t_0 such that $\mathbf{P}_s[X(t) \in A] > 0$ whenever $t \geq t_0$.

We refer to a recurrent, aperiodic Harris chain as ergodic. The key result we need is the following, which can be found in [53]: if κ defines an ergodic Harris chain with stationary distribution $\pi \in \Delta(S)$, then

$$\lim_{t \rightarrow \infty} \sup_{E \in \mathcal{F}(S)} |\kappa^t(s, E) - \pi(E)| = 0 \quad (3.1)$$

whenever $s \in S$ satisfies $\mathbf{P}_s[R_A < \infty] = 1$, where $R_A = \inf\{t \in \{1, 2, \dots\} \mid X(t) \in A\}$. In other words, the t -step transition kernel when starting from s , $\kappa^t(s, -)$, converges in total variation to π as $t \rightarrow \infty$.

Our main result with respect to marginal trait distributions can be stated succinctly as follows:

Theorem 3.3. $\lim_{T \rightarrow \infty} \mathbf{P}_{\mu_0}[S_i(T) \in E] = \pi(E)$ for every $i = 1, \dots, N$; initial distribution, $\mu_0 \in \Delta(S^N)$; and $E \in \mathcal{F}(S)$.

Proof. Since $\{X(t)\}_{t=0}^\infty$ is ergodic, for each $s \in S$ and $\varepsilon > 0$, there exists $t_0 \geq 0$ such that whenever $t \geq t_0$,

$$\sup_{E \in \mathcal{F}(S)} |\mathbf{P}_s[X(t) \in E] - \pi(E)| < \varepsilon. \quad (3.2)$$

Thus, for each $E \in \mathcal{F}(S)$ and $t \geq t_0$, we have $\pi(E) - \varepsilon < \mathbf{P}_s[X(t) \in E] < \pi(E) + \varepsilon$, which gives

$$(\pi(E) - \varepsilon)\mathbf{P}[B_i^T \geq t_0] < \sum_{t=t_0}^T \mathbf{P}_s[X(t) \in E]\mathbf{P}[B_i^T = t] < (\pi(E) + \varepsilon)\mathbf{P}[B_i^T \geq t_0]. \quad (3.3)$$

Since $\lim_{T \rightarrow \infty} \mathbf{P}[B_i^T \geq t_0] = 1$ for any $t_0 \geq 0$, we have

$$\begin{aligned} \lim_{T \rightarrow \infty} \sum_{t=0}^T \mathbf{P}_s[X(t) \in E]\mathbf{P}[B_i^T = t] &= \sum_{t=0}^{t_0-1} \mathbf{P}_s[X(t) \in E] \lim_{T \rightarrow \infty} \mathbf{P}[B_i^T = t] \\ &+ \lim_{T \rightarrow \infty} \sum_{t=t_0}^T \mathbf{P}_s[X(t) \in E]\mathbf{P}[B_i^T = t] \\ &= \lim_{T \rightarrow \infty} \sum_{t=t_0}^T \mathbf{P}_s[X(t) \in E]\mathbf{P}[B_i^T = t] \\ &\in (\pi(E) - \varepsilon, \pi(E) + \varepsilon). \end{aligned} \quad (3.4)$$

Since equation (3.4) holds for every $s \in S$, and since $\varepsilon > 0$ was arbitrary, it follows that $\lim_{T \rightarrow \infty} \mathbf{P}_{\mu_0}[S_i(T) \in E] = \pi(E)$ for every $\mu_0 \in \Delta(S^N)$. ■

Let δ_s denote the Dirac measure on S centred at $s \in S$. The frequency of players whose type lies in $E \in \mathcal{F}(S)$ at time T is

$$\mathbf{E}_{\mu_0} \left[\frac{1}{N} \sum_{i=1}^N \delta_{S_i(T)}(E) \right] = \frac{1}{N} \sum_{i=1}^N \mathbf{P}_{\mu_0}[S_i(T) \in E], \quad (3.5)$$

which approaches $\pi(E)$ as $t \rightarrow \infty$ by theorem 3.3; therefore, π gives the long-term trait frequencies. In the setting of theorem 2.1, if $\{\mathbf{S}(T)\}_{T=0}^\infty$ has a unique stationary distribution, μ , then $\mu(s_i = k) = \pi_k$ for every $i = 1, \dots, N$ and $k \in S$, i.e. the marginal distribution of each individual is exactly π .

Although our focus has been on haploid individuals, we note that theorem 3.3 can be extended to populations with diploid, haplodiploid or polyploid genetics, with either sexual or asexual reproduction (or a combination of both) as long as (i) only a single genetic locus is considered, (ii) there is no recombination within this locus and (iii) the theorem is understood to apply at the level of alleles, rather than allele combinations or traits. The idea of this extension is to treat the alleles at this locus as asexual replicators, and let $i = 1, \dots, N$ index *genetic sites* rather than individuals [54]. These genetic sites are distributed among individuals, so that each individual has a number of genetic sites equal to its ploidy. Each genetic site i contains a single allele $s_i \in S$, where S represents the set of possible alleles at this locus. During transitions, the alleles

in a subset R of genetic sites are replaced by (possibly mutated) copies of the alleles in other genetic sites, as determined by the chosen replacement event, (R, α) . The probability distribution over replacement events, $\{p_{(R, \alpha)}\}_{(R, \alpha)}$, encodes all necessary information about ploidy, sexes and mating.

Formally, this gene's-eye framework is mathematically equivalent to the framework based on haploid individuals [54]. All of our results therefore carry over to sexually reproducing populations without any additional mathematical assumptions. In applying these results to sexually reproducing populations, there is, however, an implicit *biological* assumption that there is no recombination *within the locus in question*. That is, each allele in a new offspring is a copy of a single allele in one parent, not a mixture of two (or more) alleles, which is reasonable if the locus represented by S is small enough that linkage within the locus can be assumed complete. We also emphasize that theorem 3.3 characterizes the limiting frequencies of alleles but not of allele combinations in individuals. For example, in a diploid population with alleles A and a , theorem 3.3 characterizes the limiting frequencies of A and a , but not of AA , Aa and aa . If different allele combinations correspond to different traits, theorem 3.3 does not characterize the trait distribution in this case. The limiting frequencies of allele combinations in individuals depend on the population structure, i.e. on the replacement rule.

Even if the mutation process does not have a unique stationary distribution (such as when there are multiple absorbing states), the proof of theorem 3.3 can still be used to derive the marginal distributions of every individual when the starting condition is monomorphic. Suppose that every individual in the population initially has trait $s \in S$. There will eventually be enough births along every lineage for the mutation process to reach its limiting distribution (provided the limit exists). Moreover, since the lineage leading to i at time T can start at any $j \in \{1, \dots, N\}$ at time 0, we know the initial condition of this chain, s , provided the population starts out monomorphic. We state this observation as a proposition:

Proposition 3.4. *If $s \in S$, $E \in \mathcal{F}(S)$, and $i \in \{1, \dots, N\}$, then, provided the limits exist,*

$$\lim_{T \rightarrow \infty} P_{(s, \dots, s)}[S_i(T) \in E] = \lim_{t \rightarrow \infty} P_s[X(t) \in E]. \quad (3.6)$$

Example 3.5 (Regeneration process). Suppose that each individual in a population is a (bacterial) cell dividing at a constant rate, and mutations occur during cell division. We denote by A_0 the 'wild-type' cells, which mutate to A_1 cells with probability w . A_1 cells contain the starting condition for the search process (e.g. the duplicated gene). For $k > 0$, mutation with probability u leads from an A_k cell to an A_{k+1} cell. These steps are the forward mutations in the search process toward the new function. The search is lost with probability v , i.e. each A_k cell (with $k > 0$) mutates back to an A_0 cell with probability v . Here v can represent the rate of deletion events, nonsense mutations or any missense mutation that leads away from the target because then the search is essentially lost. It is natural to assume that $v > u$, meaning that at each step, it is more likely that the search is lost than that a mutation is made in the direction of the target. This mutation scheme, which was introduced by Knoll & Nowak [55], is known as the 'regeneration process'.

Consider a Wright–Fisher process in a population of size N . Generations are non-overlapping, and at each time step the current generation is sampled (with replacement) from the previous generation [56–58]. In other words, every individual in the current generation is the offspring of individual i in the previous generation with probability $1/N$ for $i = 1, \dots, N$. The total population size is strictly constant, and all cells have the same reproductive rate. In other words, evolution is neutral. On each birth event, the offspring mutates from the parent's type according to the regeneration process. We are interested in the stationary distribution of this stochastic process, and, in particular, the probability that a randomly drawn cell is of type A_k . Denote by π_k this probability and suppose first that $N = 1$. By the definition of a stationary distribution and the

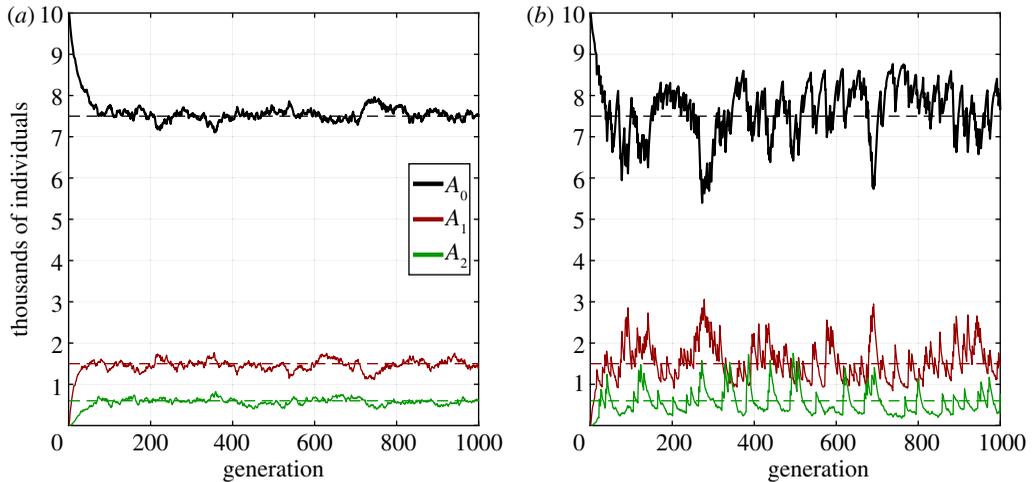


Figure 3. The frequency of three cell types in a population of size $N = 10^4$. Mutations in both panels are governed by the regeneration process with $m = 10^3$ and mutation rates $w = 0.01$, $u = 0.02$ and $v = 0.03$. The mean abundance of each type in this mutation process is indicated by a dashed line. In (a), the population is updated according to a Wright–Fisher rule. (b) illustrates a modified (and asymmetric) Wright–Fisher rule in which there exists a single cell in every generation that is more likely to reproduce than the others (independent of cell type). In (a), each offspring chooses a parent uniformly-at-random from the previous generation (i.e. each with probability $1/N = 10^{-4}$). In (b), one marked individual is chosen as the parent with probability 0.05 and each of the $N - 1$ remaining individuals is chosen with probability $(1 - 0.05)/(N - 1) \ll 0.05$. Despite the heterogeneity in (b), the mean trait abundances are the same; only the fluctuations around these means change. (Online version in colour.)

mutational scheme for the regeneration process, we have the recurrence relations

$$\pi_0 = \pi_0(1 - w) + (1 - \pi_0)v; \quad (3.7a)$$

$$\pi_1 = \pi_0 w + \pi_1(1 - u - v); \quad (3.7b)$$

$$\pi_k = \pi_{k-1}u + \pi_k(1 - u - v); \quad (1 < k < m) \quad (3.7c)$$

and

$$\pi_m = \pi_{m-1}u + \pi_m(1 - v). \quad (3.7d)$$

The solution to this system of equations is given explicitly by

$$\pi_0 = \frac{v}{v + w}; \quad (3.8a)$$

$$\pi_k = \left(\frac{u}{u + v}\right)^{k-1} \left(\frac{w}{u + v}\right) \left(\frac{v}{v + w}\right); \quad (0 < k < m) \quad (3.8b)$$

and

$$\pi_m = \left(\frac{u}{v}\right) \left(\frac{u}{u + v}\right)^{m-2} \left(\frac{w}{u + v}\right) \left(\frac{v}{v + w}\right). \quad (3.8c)$$

For simplicity, we assume that u and v do not depend on A_k . However, one can similarly calculate π when u_k is the probability of going from A_k to A_{k+1} and v_k is the probability of going from A_k to A_0 .

Determining the equilibrium frequency of A_k in a population of any size, N , is actually quite simple: it is still just π_k and can be calculated by considering a stochastic process that follows a single lineage. That is, the frequencies of the cell types are the same for a population of size $N = 1$ as it is for a population of any size. Furthermore, the probability that individual i has type A_k is π_k (which is independent of i). Figure 3 shows that heterogeneity induced by population structure can, however, strongly influence a type’s fluctuation around its mean frequency.

Example 3.6 (One-dimensional Markov chain). Another simple example is when the mutation process, defined by transition matrix M , is one-dimensional. In this case, the state space is $S = \{0, 1, \dots, m\}$ for some $m \geq 0$, and transitions from a given state are at most one step in either direction. In other words, $M_{s,s'} = 0$ whenever $|s - s'| > 1$. One can easily check that the stationary distribution for M is given by

$$\pi_s = \frac{\prod_{k=0}^{s-1} (M_{k,k+1}/M_{k+1,k})}{1 + \sum_{s=1}^m \prod_{k=0}^{s-1} (M_{k,k+1}/M_{k+1,k})}. \quad (3.9)$$

The above equation (3.9) gives a simple expression for the long-run probability of finding an individual in type s , both in the mutation process itself and in any evolving population in which mutations are neutral and governed by M .

4. Intra-lineage mixing times

So far, we have focused on neutral trait frequencies, which are not affected by the distribution over replacement events, $\{p_{(R,\alpha)}\}_{(R,\alpha)}$. We now turn to rate of convergence to these equilibrium frequencies along the lineages, characterized in terms of mixing times. The proof of theorem 3.3 is based on the fact that, eventually, all lineages will contain sufficiently many birth events for the mutation process to mix. Here, we consider the amount of time—measured in update steps—for this mixing to occur along a given lineage. Since both the mutation processes and evolutionary update rules we consider are discrete in time, all references to time shall indicate non-negative integers.

Let $\pi \in \Delta(S)$ be the stationary distribution for the mutation process. For $t \in \{1, 2, \dots\}$, let

$$\psi(t) := \sup_{s \in S} \sup_{E \in \mathcal{F}(S)} |\mathbf{P}_s[X(t) \in E] - \pi(E)|. \quad (4.1)$$

The ε -mixing time of the mutation process is then $\tau(\varepsilon) := \min\{t \in \{1, 2, \dots\} : \psi(t) < \varepsilon\}$. (We use the symbol ψ to denote this distance rather than the usual d in order to avoid confusion with death rates.)

At the population level, we can consider an analogue of mixing time along lineages. Recall that B_i^T is the number of birth events in the lineage leading to individual i after T update steps. Since the marginal distributions all converge to π , the distance between i 's trait at time $T \in \{1, 2, \dots\}$ and the stationary distribution is

$$\Psi_i(T) = \sup_{s \in S^N} \sup_{E \in \mathcal{F}(S)} |\mathbf{P}_s[S_i(T) \in E] - \pi(E)|. \quad (4.2)$$

The analogue of $\tau(\varepsilon)$ in this case is $\mathcal{T}_i(\varepsilon) := \min\{T \in \{1, 2, \dots\} : \Psi_i(T) < \varepsilon\}$, which is the ε -mixing time of the process along the lineage leading to individual i . We have the trivial lower bound $\mathcal{T}_i(\varepsilon) \geq \tau(\varepsilon)$ for all $i = 1, \dots, N$.

Example 4.1 (Non-overlapping generations). If generations do not overlap, then $\mathbf{P}[B_i^T = T] = 1$. Trivially, then, we have $\Psi_i(T) = \psi(T)$ and $\mathcal{T}_i(\varepsilon) = \tau(\varepsilon)$ for every $i = 1, \dots, N$.

Suppose, for instance, that the death rate is constant and equal to $d \in (0, 1)$, so that each individual is updated every $1/d$ update steps (on average). Since all individuals have the same probability of being replaced, $\mathcal{T}_i(\varepsilon)$ is independent of i , and we call this mixing time simply $\mathcal{T}(\varepsilon)$. If τ is the mixing time of the mutation chain, then a natural guess for \mathcal{T} is simply τ/d because, on average, the lineage experiences a mutation (i.e. a step in the mutation chain) every $1/d$ updates. However, it turns out that τ/d is generally a bad approximation of \mathcal{T} because it does not take into account enough information about the distribution of lineage length. In fact, one cannot even use τ/d to establish a general upper or lower bound on \mathcal{T}_i . Figure 4 illustrates this comparison for death–birth updating with $d = 1/N$, where N is the population size.

Instead, we can bound $\Psi_i(T)$ by the average of $\psi(B_i^T)$, which we state as a simple lemma:

Lemma 4.2. $\Psi_i(T) \leq E[\psi(B_i^T)]$ for every $i = 1, \dots, N$ and $T \in \{1, 2, \dots\}$.

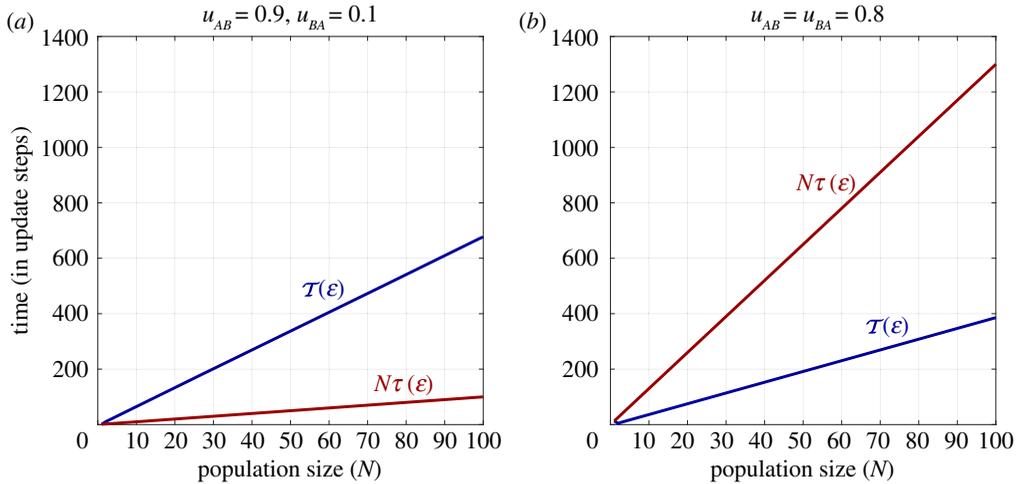


Figure 4. Intra-lineage mixing time, $\mathcal{T}(\varepsilon)$, relative to the mixing time of the underlying mutation process, $\tau(\varepsilon)$, for death–birth updating (shown here for $\varepsilon = 10^{-3}$). Every individual has one of two types, A or B , and upon reproduction A mutates to B with probability u_{AB} and B mutates to A with probability u_{BA} . The stationary distribution of this process puts probability $u_{BA}/(u_{AB} + u_{BA})$ on state A and probability $u_{AB}/(u_{AB} + u_{BA})$ on state B . Since the death rate is constant and equal to $1/N$ (regardless of the population’s spatial structure), each individual is updated every N steps (on average). However, $\mathcal{T}(\varepsilon)$ differs significantly from the rescaled mixing time of the mutation process, $N\tau(\varepsilon)$. In fact, $N\tau(\varepsilon)$ neither a general upper nor lower bound on $\mathcal{T}(\varepsilon)$, which can be seen by simply varying u_{AB} and u_{BA} . (Online version in colour.)

Proof. Fix $i = 1, \dots, N$ and $T \in \{1, 2, \dots\}$. Straightforward manipulations give

$$\begin{aligned}
 \Psi_i(T) &= \sup_{s \in \mathcal{S}^N} \sup_{E \in \mathcal{F}(S)} |\mathbf{P}_s[S_i(T) \in E] - \pi(E)| \\
 &\leq \sum_{t=0}^T \sup_{s \in \mathcal{S}} \sup_{E \in \mathcal{F}(S)} |\mathbf{P}_s[X(t) \in E] - \pi(E)| \mathbf{P}[B_i^T = t] \\
 &= \mathbf{E} \left[\sup_{s \in \mathcal{S}} \sup_{E \in \mathcal{F}(S)} |\mathbf{P}_s[X(B_i^T) \in E] - \pi(E)| \right] \\
 &= \mathbf{E}[\psi(B_i^T)]
 \end{aligned} \tag{4.3}$$

as desired. ■

To calculate $\mathbf{E}[\psi(B_i^T)]$, we note that $\mathbf{P}[B_i^T = t]$ satisfies the multivariate recurrence relation

$$\mathbf{P}[B_i^T = t] = \mathbf{P}[B_i^{T-1} = t](1 - d_i) + \sum_{j=1}^N \mathbf{P}[B_j^{T-1} = t-1] e_{ji}, \tag{4.4}$$

with boundary conditions $\mathbf{P}[B_i^T = 0] = (1 - d_i)^T$ for $i = 1, \dots, N$. Explicitly, for $1 < t \leq T$, we have

$$\mathbf{P}[B_i^T = t] = \sum_{\substack{k_0, \dots, k_t \geq 1 \\ k_0 + \dots + k_t = T}} (1 - d_i)^{k_t - 1} \sum_{j_{t-1}=1}^N (1 - d_{j_{t-1}})^{k_{t-1} - 1} e_{j_{t-1}, j_t} \cdots \sum_{j_0=1}^N (1 - d_{j_0})^{k_0 - 1} e_{j_0, j_1}. \tag{4.5}$$

While the general expression for $\mathbf{P}[B_i^T = t]$ (equation (4.5)) can appear complicated (depending on the replacement rule, $\{p_{(R, \alpha)}\}_{(R, \alpha)}$, and the resulting demographic variables), we can further simplify the bound given by lemma 4.2 if some additional properties hold. We consider two cases in which one can be more explicit.

(a) Finite, ergodic mutation chains

If S is finite and $\{X(t)\}_{t=0}^{\infty}$ is an irreducible and aperiodic, then there exist $C > 0$ and $\alpha \in (0, 1)$ such that $\psi(t) \leq C\alpha^t$ for every $t \geq 1$ [52,59]. From lemma 4.2,

$$\Psi_i(T) \leq CE[\alpha^{B_i^T}]. \quad (4.6)$$

Example 4.3 (Constant death rate). If the death rate is constant, meaning there exists d with $d_i = d$ for $i = 1, \dots, N$, then $\mathbf{P}[B_i^T = t] = \binom{T}{t}(1-d)^{T-t}d^t$. From equation (4.6), we have the upper bound

$$\Psi_i(T) \leq C(1-d(1-\alpha))^T. \quad (4.7)$$

Under death–birth updating, for example, the death rate is constant and equal to $1/N$. The death rate is also clearly constant when generations are non-overlapping, and in this case $1-d(1-\alpha) = \alpha$, which gives back the same upper bound of $C\alpha^T$. On the other hand, when $d < 1$ we have $1-d(1-\alpha) > \alpha$.

(b) Reversible mutation chains and spectral theory

Suppose that the mutation process, M , is reversible with respect to π , meaning $\pi_s M_{s,s'} = \pi_{s'} M_{s',s}$ for every $s, s' \in S$. The matrix Λ defined by $\Lambda_{s,s'} = \pi_s^{1/2} M_{s,s'} \pi_{s'}^{-1/2}$ is then symmetric with all of its eigenvalues real, $1 = \lambda_1 > \lambda_2 \geq \dots \geq \lambda_{|S|} \geq -1$. These eigenvalues are the same as those of M since M and Λ represent the same linear transformation. Let $\pi_* := \min_{1 \leq i \leq |S|} \pi_{s_i}$. By standard results on mixing times in reversible Markov chains [60], we have

$$\psi(T) \leq \frac{1}{2} \sqrt{\frac{1-\pi_*}{\pi_*}} \max\{|\lambda_2|, |\lambda_{|S|}|\}^T. \quad (4.8)$$

To see how the evolutionary process affects mixing along its lineages, we can again consider the case of a constant death rate for simplicity. If d is the death rate for every individual, i , then the mutation process along any lineage is a Markov chain with transition matrix $M(d) := (1-d)I + dM$. Moreover, $M(d)$ is reversible with respect to π since M is. Letting $\Lambda(d) = (1-d)I + d\Lambda$, we see that if $\Lambda v = \lambda v$, then $\Lambda(d)v = (1-d(1-\lambda))v$. Since the map $\lambda \mapsto 1-d(1-\lambda)$ sends eigenvalues of Λ to eigenvalues of $\Lambda(d)$, and since $\max_{2 \leq i \leq |S|} |1-d(1-\lambda_i)| = \max\{|1-d(1-\lambda_2)|, |1-d(1-\lambda_{|S|})|\}$, we have

$$\Psi(T) \leq \frac{1}{2} \sqrt{\frac{1-\pi_*}{\pi_*}} \max\{|1-d(1-\lambda_2)|, |1-d(1-\lambda_{|S|})|\}^T. \quad (4.9)$$

Figure 5 gives an example of this mixing time bound for the reversible chain with $|S| = 3$ and transition matrix

$$M = \begin{pmatrix} 0.9429 & 0.0284 & 0.0287 \\ 0.0250 & 0.9638 & 0.0112 \\ 0.0360 & 0.0159 & 0.9481 \end{pmatrix}. \quad (4.10)$$

On the other hand, if the death rate depends on i , we cannot necessarily transform the mutation process into one whose transition matrix is $M(d) = (1-d)I + dM$ for some d . Instead, starting at location i , one needs to keep track of where in the population each ancestor arises. Since M is reversible, its time-reversal, \tilde{M} , is again just M . Consider the Markov chain on $\{1, \dots, N\} \times S$, $\{Y(T)\}_{T=0}^{\infty}$, with transitions given by

$$P_{(i,s),(j,s')} = \begin{cases} e_{ji} M_{s,s'} & i \neq j \text{ or } s \neq s', \\ e_{ii} M_{s,s} + 1 - d_i & i = j \text{ and } s = s'. \end{cases} \quad (4.11)$$

Starting at location i , one can then ask how many steps is required for the marginal trait distribution to be close to π (in total variation). From the definition of $\Psi_i(T)$, it is easily seen

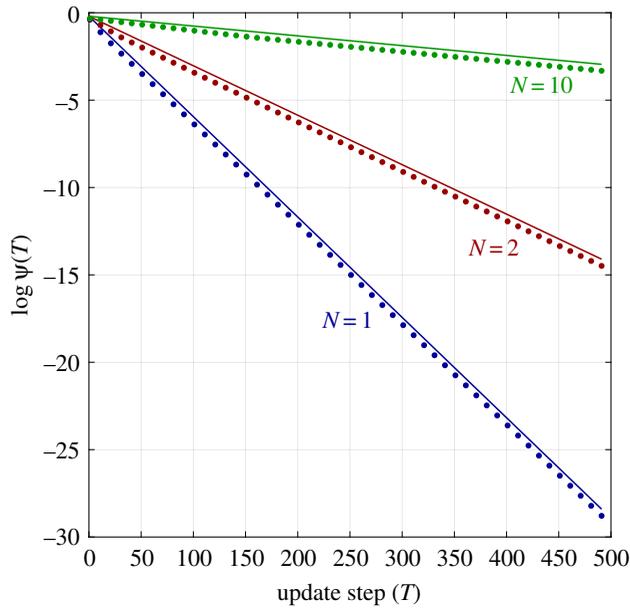


Figure 5. The distance between individual i 's trait and the stationary distribution after T update steps, $\Psi(T)$, under death–birth updating. All of the marginal distributions converge to π as $T \rightarrow \infty$, but the rates at which they do so depend on the population size, structure and update rule. A population of size $N = 1$ is the same as the underlying mutation process itself, and in this case $\Psi(T) = \psi(T) \leq C\alpha^T$ for some $C > 0$ and $\alpha \in (0, 1)$. The mutation process depicted here is reversible with $|S| = 3$ traits and transition matrix equation (4.10), so we can set $C = \frac{1}{2}\sqrt{(1 - \pi_*)/\pi_*}$ and $\alpha = \max\{|\lambda_2|, |\lambda_3|\}$ (see equation (4.8)). If the population has size N , then $\Psi(T) \leq C(1 - (1/N)(1 - \alpha))^T$ since death rate for all individuals is $1/N$ under death–birth updating. Since $\Psi(T)$ converges rapidly to 0 in all cases, we show both the predicted upper bounds (solid lines) and the actual values of $\Psi(T)$ (dots) on a logarithmic scale. (Online version in colour.)

that

$$\Psi_i(T) = \sup_{s \in S} \sup_{E \in \mathcal{F}(S)} |\mathbf{P}_{(i,s)}[Y(T) \in \{1, \dots, N\} \times E] - \pi(E)|. \quad (4.12)$$

As a result, the mixing time of the marginal trait distribution of $\{Y(T)\}_{T=0}^\infty$ starting from i is $\mathcal{T}_i(\epsilon)$.

(c) Mutations induced by group actions

One way to represent mutation is as a group action. Let S be a finite set representing a collection of heritable traits, and let G be a group (in the mathematical sense (e.g. [61])) representing the possible effects of mutation. Suppose that G acts on S , meaning there exists a map $G \times S \rightarrow S$ sending (g, s) to $gs \in S$ that satisfies the following properties: (i) if e is the identity element of G , then $es = s$ for every $s \in S$ and (ii) if $g, h \in G$ and $s \in S$, then $(gh)s = g(hs)$. A special case is when $S = G$ and the action is the same as the binary operation in G .

For example, G might be a group of three-dimensional rotations (a subgroup of $SO(3)$), acting on a set S of possible orientations of some aspect of an organism's morphology. Or G might be the group of permutations on n elements, \mathfrak{S}_n , acting on the set S of possible arrangements of n transposable elements in a genome. An especially relevant example is $G = S = (\mathbb{Z}/2\mathbb{Z})^m$; here S can be understood as the group of binary strings (idealized genomes) of length m .

If $\nu \in \Delta(G)$ is a probability distribution on G , then one can define a Markov chain on S with transitions

$$M_{s,s'} := \sum_{\substack{g \in G \\ gs=s'}} \nu(g). \quad (4.13)$$

In order to ensure that this chain has a unique stationary distribution, we need to assume that the group action is transitive; that is, for any $s, s' \in S$, there exists $g \in G$ with $gs = s'$. For practical reasons, we make at least one of two other assumptions on the group action when considering the mixing times of M :

(i) Free group actions

If the action of G on S is free, then $gs = hs$ implies that $g = h$. If the action of G is on itself via its binary operation (in which case equation (4.13) defines a random walk on G), then the action is free because every element of G has an inverse. In other words, the only element of G that sends $g \in G$ to itself is the identity element. A free action generalizes this property to actions of G on another set, S .

For a free action, $M_{s,gs} := \nu(g)$ for every $s \in S$ and $g \in G$. The distribution π defined by $\pi_g = |G|^{-1}$ for $g \in G$ (i.e. the uniform distribution on G) is the stationary distribution for both M and its time-reversal, \tilde{M} , which satisfies $\tilde{M}_{s,gs} = \nu(g^{-1})$. A standard result on mixing times for random walks on groups (see [59], Lemma 4.13 and Corollary 4.14) is then easily seen to imply that M and \tilde{M} have the same mixing time.

(ii) Abelian group actions

Let G be an abelian group, meaning $gh = hg$ for every $g, h \in G$. An example of an abelian group is the set of binary strings of length m , $(\mathbb{Z}/2\mathbb{Z})^m$, with the operation of addition. (The group of permutations on $n > 2$ letters, \mathfrak{S}_n , is an example of a group that is *not* abelian.) If $s \in S$ and $g, h \in G$ satisfy $gs = hs$, then $g^{-1}s = h^{-1}s$. The time-reversal of M then satisfies

$$\tilde{M}_{s,s'} = \sum_{\substack{g \in G \\ gs=s'}} \nu(g^{-1}). \quad (4.14)$$

Again, both M and \tilde{M} have the uniform stationary distribution, and straightforward modifications of the arguments presented in [59, Lemma 4.13] show that M and \tilde{M} have the same mixing time.

If the group action is transitive and either free or abelian, it follows that one can obtain $T_i(\epsilon)$ from the time-reversed Markov chain on $\{1, \dots, N\} \times S$, $\{Y(T)\}_{T=0}^\infty$, whose transitions are defined by

$$P_{(i,s),(j,s')} = \begin{cases} e_{ji}\tilde{M}_{s,s'} & i \neq j \text{ or } s \neq s', \\ e_{ii}\tilde{M}_{s,s} + 1 - d_i & i = j \text{ and } s = s'. \end{cases} \quad (4.15)$$

In particular, equation (4.12) holds even though the mutation chain, $\{X(t)\}_{t=0}^\infty$, is not necessarily reversible.

Remark 4.4. Although one can define a time-reversal of any mutation chain (not just one on a group or one that is reversible), the mixing time of the reversed chain need not coincide with the original chain. Therefore, the utility of reversing a chain and studying it along lineages is limited, as least for the purpose of analysing mixing times.

5. Discussion

The result on stationary trait frequencies is well established in the special case of two competing types on a homogeneous population structure [62,63]. Similarly, in population genetics, it has been noted that the so-called ‘common ancestor’ process is the same as the mutation process when there is no selection [30,31,64]. Comparisons between neutrally evolving populations and their mutation processes have also been used to show that evolution favours types that are robust against mutation [65]. Given the coupling of birth and mutation, how neutral evolution affects trait frequencies (allele frequencies, in the sexually reproducing case) is a natural question to ask. Under mild assumptions, we have seen that neutral evolution results in the same trait frequencies as the mutation process.

This method can also be used to characterize convergence to the stationary distribution along the lineages. Since each lineage can contain at most one birth event per update step, the lineages mix at least as slowly as does the original mutation process. In general, if $\Psi_i(T)$ is the distance between the stationary distribution and the lineage leading to individual i , then $\Psi_i(T) \leq \mathbb{E}[\psi(B_i^T)]$, where $\psi(t)$ is the distance between the mutation process and the stationary distribution after t steps and B_i^T is a random variable giving the number of birth events along the lineage leading to i at time T . The distribution of B_i^T can be given explicitly in terms of the demographic variables of the update rule (equation (4.5)), which gives a bound on $\Psi_i(T)$ that is straightforward to calculate (although the expression itself is not necessarily succinct).

Matrix games, as well as the regeneration process, have finitely many possible types, but the main result on long-term trait frequency holds for mutation processes with continuous state spaces as well. A genetic type of this form could be an element of an interval such as $[0, 1]$, for example, representing partial expression of a trait or the tendency to be of a particular binary type. Mutations can also be supported on uncountably infinitely many points in the trait space, one example being when an individual with type $x \in [0, 1]$ mutates to a nearby type with probability determined by a Gaussian distribution centred at x [66,67]. On a generic state space, the technical condition we require is that the mutation process be an ergodic Harris chain [52].

Our focus here has been on evolving populations of fixed size and structure. These assumptions are not strictly necessary, but they do make the notation more convenient since an evolutionary process can then be described by a probability distribution over simple replacement events [26]. Extinction also becomes a possibility when the population size fluctuates, which can make analysing such a process more complicated [68–73]. Nonetheless, the intuition from our analysis here carries over to other situations: as long as an individual arises on a lineage with sufficiently many prior birth events, the mutation process along this lineage will converge to its stationary distribution, π . If all individuals in the population have this property sufficiently far into the future, then π must give the long-term average trait frequencies in the population.

The study of trait dynamics along lineages leads to an interesting question: how does selection change the stationary trait distribution of individual i ? This question, of course, is not new and has been studied in constant-fitness models over the past two decades via the so-called ‘ancestral selection graph’ [74–76], which is a version of Kingman’s coalescent [77] that allows for selective differences between the types. In populations with heterogeneous structure, selection can affect the marginal trait distributions in different ways at different locations, and there are still many open questions in this area, particularly with respect to frequency-dependent selection (games). For example, in an evolutionary game, the standard method of measuring the influence of selection is through its effects on average trait frequency in the population [21]. However, this type of averaging over the population does not provide a description of the locations at which a given trait is more likely to be found in a population, and it says nothing about selection’s effects on mixing times.

Here, we have studied the individual lineages generated by neutral evolution. The trait distributions along those lineages are given by that of the mutation process. Therefore, the trait frequencies do not depend on population size, structure, and update rule. The mixing times of neutral evolution, however, do depend on these demographic components. Selection can perturb neutral evolutionary dynamics in complicated ways, and how it changes both marginal distributions and mixing times, especially within the realm of evolutionary game theory, is a more complicated question but nonetheless a natural extension of the one considered here.

Data accessibility. This article does not contain any additional data.

Authors’ contributions. All authors designed research, performed research, and wrote the paper.

Competing interests. We declare we have no competing interests.

Funding. A.M. and M.A.N. are supported by the Office of Naval Research, grant no. N00014-16-1-2914. B.A. is supported by the National Science Foundation, award no. 1715315.

Acknowledgements. The authors are grateful to Christoph Hauert, John Wakeley, and the referees for helpful comments.

References

1. Wright S. 1931 Evolution in Mendelian populations. *Genetics* **16**, 97–159.
2. Maruyama T. 1970 Effective number of alleles in a subdivided population. *Theor. Popul. Biol.* **1**, 273–306. (doi:10.1016/0040-5809(70)90047-x)
3. Maruyama T. 1970 On the fixation probability of mutant genes in a subdivided population. *Genet. Res.* **15**, 221–225. (doi:10.1017/s0016672300001543)
4. Maruyama T. 1974 A simple proof that certain quantities are independent of the geographical structure of population. *Theor. Popul. Biol.* **5**, 148–154. (doi:10.1016/0040-5809(74)90037-9)
5. Nagylaki T. 1980 The strong-migration limit in geographically structured populations. *J. Math. Biol.* **9**, 101–114. (doi:10.1007/bf00275916)
6. Nagylaki T. 1992 *Introduction to theoretical population genetics*. Berlin, Germany: Springer.
7. Lieberman E, Hauert C, Nowak MA. 2005 Evolutionary dynamics on graphs. *Nature* **433**, 312–316. (doi:10.1038/nature03204)
8. Ohtsuki H, Hauert C, Lieberman E, Nowak MA. 2006 A simple rule for the evolution of cooperation on graphs and social networks. *Nature* **441**, 502–505. (doi:10.1038/nature04605)
9. Tarnita CE, Antal T, Ohtsuki H, Nowak MA. 2009 Evolutionary dynamics in set structured populations. *Proc. Natl Acad. Sci. USA* **106**, 8601–8604. (doi:10.1073/pnas.0903019106)
10. Traulsen A, Hauert C, De Silva H, Nowak MA, Sigmund K. 2009 Exploration dynamics in evolutionary games. *Proc. Natl Acad. Sci. USA* **106**, 709–712. (doi:10.1073/pnas.0808450106)
11. Loewe L, Hill WG. 2010 The population genetics of mutations: good, bad and indifferent. *Phil. Trans. R. Soc. B* **365**, 1153–1167. (doi:10.1098/rstb.2009.0317)
12. Hauert C, Imhof L. 2012 Evolutionary games in deme structured, finite populations. *J. Theor. Biol.* **299**, 106–112. (doi:10.1016/j.jtbi.2011.06.010)
13. Kimura M. 1983 *The neutral theory of molecular evolution*. Cambridge, UK: Cambridge University Press.
14. Lynch M, Hill WG. 1986 Phenotypic evolution by neutral mutation. *Evolution* **40**, 915–935. (doi:10.2307/2408753)
15. Ochman H. 2003 Neutral mutations and neutral substitutions in bacterial genomes. *Mol. Biol. Evol.* **20**, 2091–2096. (doi:10.1093/molbev/msg229)
16. Nei M. 2005 Selectionism and neutralism in molecular evolution. *Mol. Biol. Evol.* **22**, 2318–2342. (doi:10.1093/molbev/msi242)
17. Bloom JD, Romero PA, Lu Z, Arnold FH. 2007 Neutral genetic drift can alter promiscuous protein functions, potentially aiding functional evolution. *Biol. Direct* **2**, 17. (doi:10.1186/1745-6150-2-17)
18. Antal T, Ohtsuki H, Wakeley J, Taylor PD, Nowak MA. 2009 Evolution of cooperation by phenotypic similarity. *Proc. Natl Acad. Sci. USA* **106**, 8597–8600. (doi:10.1073/pnas.0902528106)
19. Antal T, Traulsen A, Ohtsuki H, Tarnita CE, Nowak MA. 2009 Mutation-selection equilibrium in games with multiple strategies. *J. Theor. Biol.* **258**, 614–622. (doi:10.1016/j.jtbi.2009.02.010)
20. Tarnita CE, Ohtsuki H, Antal T, Fu F, Nowak MA. 2009 Strategy selection in structured populations. *J. Theor. Biol.* **259**, 570–581. (doi:10.1016/j.jtbi.2009.03.035)
21. Nowak MA, Tarnita CE, Antal T. 2009 Evolutionary dynamics in structured populations. *Phil. Trans. R. Soc. B* **365**, 19–30. (doi:10.1098/rstb.2009.0215)
22. Tarnita CE, Wage N, Nowak MA. 2011 Multiple strategies in structured populations. *Proc. Natl Acad. Sci. USA* **108**, 2334–2337. (doi:10.1073/pnas.1016008108)
23. Gokhale CS, Traulsen A. 2011 Strategy abundance in evolutionary many-player games with multiple strategies. *J. Theor. Biol.* **283**, 180–191. (doi:10.1016/j.jtbi.2011.05.031)
24. Wu B, García J, Hauert C, Traulsen A. 2013 Extrapolating weak selection in evolutionary games. *PLoS Comput. Biol.* **9**, e1003381. (doi:10.1371/journal.pcbi.1003381)
25. Wu B, Traulsen A, Gokhale CS. 2013 Dynamic properties of evolutionary multi-player games in finite populations. *Games* **4**, 182–199. (doi:10.3390/g4020182)
26. Allen B, Tarnita CE. 2014 Measures of success in a class of evolutionary models with fixed population size and structure. *J. Math. Biol.* **68**, 109–143. (doi:10.1007/s00285-012-0622-x)
27. Catalán P, Seoane JM, Sanjuán MAF. 2015 Mutation-selection equilibrium in finite populations playing a Hawk–Dove game. *Commun. Nonlinear Sci. Numer. Simul.* **25**, 66–73. (doi:10.1016/j.cnsns.2015.01.012)
28. Zhang Y, Liu A, Sun C. 2016 Impact of migration on the multi-strategy selection in finite group-structured populations. *Sci. Rep.* **6**, 35114. (doi:10.1038/srep35114)

29. Moran PAP. 1958 Random processes in genetics. *Math. Proc. Cambridge Philos. Soc.* **54**, 60–71. (doi:10.1017/s0305004100033193)
30. Taylor J. 2007 The common ancestor process for a Wright-Fisher diffusion. *Elect. J. Probab.* **12**, 808–847. (doi:10.1214/ejp.v12-418)
31. Birky CW, Walsh JB. 1988 Effects of linkage on rates of molecular evolution. *Proc. Natl Acad. Sci. USA* **85**, 6414–6418. (doi:10.1073/pnas.85.17.6414)
32. Sigmund K. 2010 *The calculus of selfishness*. Princeton, NJ: Princeton University Press.
33. Nowak M, Sigmund K. 1990 The evolution of stochastic strategies in the Prisoner's dilemma. *Acta Appl. Math.* **20**, 247–265. (doi:10.1007/bf00049570)
34. Geritz SAH, Kisdi É, Meszéna G, Metz JAJ. 1998 Evolutionarily singular strategies and the adaptive growth and branching of the evolutionary tree. *Evol. Ecol.* **12**, 35–57. (doi:10.1023/a:1006554906681)
35. Dieckmann U, Doebeli M. 1999 On the origin of species by sympatric speciation. *Nature* **400**, 354–357. (doi:10.1038/22521)
36. Killingback T, Doebeli M, Knowlton N. 1999 Variable investment, the Continuous Prisoner's Dilemma, and the origin of cooperation. *Proc. R. Soc. Lond. B* **266**, 1723–1728. (doi:10.1098/rspb.1999.0838)
37. Killingback T, Doebeli M. 2002 The Continuous Prisoner's Dilemma and the evolution of cooperation through reciprocal altruism with variable investment. *Am. Nat.* **160**, 421–438. (doi:10.1086/342070)
38. Dieckmann U, Metz JAJ. 2006 Surprising evolutionary predictions from enhanced ecological realism. *Theor. Popul. Biol.* **69**, 263–281. (doi:10.1016/j.tpb.2005.12.001)
39. Imhof LA, Nowak MA. 2009 Stochastic evolutionary dynamics of direct reciprocity. *Proc. R. Soc. B* **277**, 463–468. (doi:10.1098/rspb.2009.1171)
40. Doebeli M. 2011 *Adaptive diversification (MPB-48)*. Princeton, NJ: Princeton University Press.
41. Nowak M. 1990 Stochastic strategies in the Prisoner's Dilemma. *Theor. Popul. Biol.* **38**, 93–112. (doi:10.1016/0040-5809(90)90005-G)
42. Nowak M, Sigmund K. 1993 A strategy of win-stay, lose-shift that outperforms tit-for-tat in the Prisoner's Dilemma game. *Nature* **364**, 56–58. (doi:10.1038/364056a0)
43. Roberts G, Sherratt TN. 1998 Development of cooperative relationships through increasing investment. *Nature* **394**, 175–179. (doi:10.1038/28160)
44. Wahl LM, Nowak MA. 1999 The continuous Prisoner's dilemma. I. Linear reactive strategies. *J. Theor. Biol.* **200**, 307–321. (doi:10.1006/jtbi.1999.0996)
45. Wahl LM, Nowak MA. 1999 The continuous Prisoner's dilemma. II. Linear reactive strategies with noise. *J. Theor. Biol.* **200**, 323–338. (doi:10.1006/jtbi.1999.0997)
46. Oechssler J, Riedel F. 2001 Evolutionary dynamics on infinite strategy spaces. *Econ. Theory* **17**, 141–162. (doi:10.1007/pl00004092)
47. van Veelen M, Spreij P. 2008 Evolution in games with a continuous action space. *Econ. Theory* **39**, 355–376. (doi:10.1007/s00199-008-0338-8)
48. Cleveland J, Ackleh AS. 2013 Evolutionary game theory on measure spaces: well-posedness. *Nonlinear Anal. Real World Appl.* **14**, 785–797. (doi:10.1016/j.nonrwa.2012.08.002)
49. Cheung M-W. 2014 Pairwise comparison dynamics for games with continuous strategy space. *J. Econ. Theory* **153**, 344–375. (doi:10.1016/j.jet.2014.07.001)
50. Cheung M-W. 2016 Imitative dynamics for games with continuous strategy space. *Games Econ. Behav.* **99**, 206–223. (doi:10.1016/j.geb.2016.08.003)
51. Allen B, Sample C, Dementieva Y, Medeiros RC, Paoletti C, Nowak MA. 2015 The molecular clock of neutral evolution can be accelerated or slowed by asymmetric spatial structure. *PLoS Comput. Biol.* **11**, e1004108. (doi:10.1371/journal.pcbi.1004108)
52. Durrett R. 2009 *Probability: theory and examples*. Princeton, NJ: Princeton University Press.
53. Huber ML. 2016 *Perfect simulation*. London, UK: Chapman and Hall.
54. Allen B, McAvoy A. 2018 A mathematical formalism for natural selection with arbitrary spatial and genetic structure. (<http://arxiv.org/abs/1806.04717>)
55. Knoll AH, Nowak MA. 2017 The timetable of evolution. *Sci. Adv.* **3**, e1603076. (doi:10.1126/sciadv.1603076)
56. Ewens WJ. 2004 *Mathematical population genetics*. New York, NY: Springer.
57. Imhof LA, Nowak MA. 2006 Evolutionary game dynamics in a Wright-Fisher process. *J. Math. Biol.* **52**, 667–681. (doi:10.1007/s00285-005-0369-8)
58. Der R, Epstein CL, Plotkin JB. 2011 Generalized population models and the nature of genetic drift. *Theor. Popul. Biol.* **80**, 80–99. (doi:10.1016/j.tpb.2011.06.004)

59. Levin D, Peres Y, Wilmer E. 2008 *Markov chains and mixing times*. Providence, RI: American Mathematical Society.
60. Montenegro R, Tetali P. 2005 Mathematical aspects of mixing times in Markov chains. *Found. Trends Theoret. Comput. Sci.* **1**, 237–354. (doi:10.1561/0400000003)
61. Knapp AW. 2006 *Basic algebra*. Boston, MA: Birkhäuser.
62. Taylor PD, Day T, Wild G. 2007 From inclusive fitness to fixation probability in homogeneous structured populations. *J. Theor. Biol.* **249**, 101–110. (doi:10.1016/j.jtbi.2007.07.006)
63. Débarre F. 2017 Fidelity of parent-offspring transmission and the evolution of social behavior in structured populations. *J. Theor. Biol.* **420**, 26–35. (doi:10.1016/j.jtbi.2017.02.027)
64. Fearnhead P. 2002 The common ancestor at a nonneutral locus. *J. Appl. Prob.* **39**, 38–54. (doi:10.1017/s0021900200021495)
65. van Nimwegen E, Crutchfield JP, Huynen M. 1999 Neutral evolution of mutational robustness. *Proc. Natl Acad. Sci. USA* **96**, 9716–9720. (doi:10.1073/pnas.96.17.9716)
66. Doebeli M, Hauert C, Killingback T. 2004 The evolutionary origin of cooperators and defectors. *Science* **306**, 859–862. (doi:10.1126/science.1101456)
67. Wakano JY, Iwasa Y. 2012 Evolutionary branching in a finite population: deterministic branching versus stochastic branching. *Genetics* **193**, 229–241. (doi:10.1534/genetics.112.144980)
68. Haccou P, Jagers P, Vatutin VA. 2005 *Branching processes: variation growth and extinction of populations*. Cambridge, UK: Cambridge University Press.
69. Lambert A. 2006 Probability of fixation under weak selection: a branching process unifying approach. *Theor. Popul. Biol.* **69**, 419–441. (doi:10.1016/j.tpb.2006.01.002)
70. Faure M, Schreiber SJ. 2014 Quasi-stationary distributions for randomly perturbed dynamical systems. *Ann. Appl. Probab.* **24**, 553–598. (doi:10.1214/13-aap923)
71. Hamza K, Jagers P, Klebaner FC. 2015 On the establishment, persistence, and inevitable extinction of populations. *J. Math. Biol.* **72**, 797–820. (doi:10.1007/s00285-015-0903-2)
72. Huang W, Hauert C, Traulsen A. 2015 Stochastic game dynamics under demographic fluctuations. *Proc. Natl Acad. Sci. USA* **112**, 9064–9069. (doi:10.1073/pnas.1418745112)
73. McAvoy A, Fraiman N, Hauert C, Wakeley J, Nowak MA. 2018 Public goods games in populations with fluctuating size. *Theor. Popul. Biol.* **121**, 72–84. (doi:10.1016/j.tpb.2018.01.004)
74. Krone SM, Neuhauser C. 1997 Ancestral processes with selection. *Theor. Popul. Biol.* **51**, 210–237. (doi:10.1006/tpbi.1997.1299)
75. Neuhauser C. 1999 The ancestral graph and gene genealogy under frequency-dependent selection. *Theor. Popul. Biol.* **56**, 203–214. (doi:10.1006/tpbi.1999.1412)
76. Slade PF, Wakeley J. 2005 The structured ancestral selection graph and the many-demes limit. *Genetics* **169**, 1117–1131. (doi:10.1534/genetics.104.032276)
77. Kingman JFC. 1982 On the genealogy of large populations. *J. Appl. Probab.* **19**, 27–43. (doi:10.2307/3213548)