

Mixing Time of Markov Chains, Dynamical Systems and Evolution

Ioannis Panageas^{*1} and Nisheeth K. Vishnoi²

1 Georgia Institute of Technology, Atlanta, USA
ioannis@gatech.edu

2 École Polytechnique Fédérale de Lausanne (EPFL), Lausanne, Switzerland
nisheeth.vishnoi@epfl.ch

Abstract

In this paper we study the mixing time of evolutionary Markov chains over populations of a fixed size (N) in which each individual can be one of m types. These Markov chains have the property that they are guided by a dynamical system from the m -dimensional probability simplex to itself. Roughly, given the current state of the Markov chain, which can be viewed as a probability distribution over the m types, the next state is generated by applying this dynamical system to this distribution, and then sampling from it N times. Many processes in nature, from biology to sociology, are evolutionary and such chains can be used to model them. In this study, the mixing time is of particular interest as it determines the speed of evolution and whether the statistics of the steady state can be efficiently computed. In a recent result [Panageas, Srivastava, Vishnoi, Soda, 2016], it was suggested that the mixing time of such Markov chains is connected to the geometry of this guiding dynamical system. In particular, when the dynamical system has a fixed point which is a global attractor, then the mixing is fast. The limit sets of dynamical systems, however, can exhibit more complex behavior: they could have multiple fixed points that are not necessarily stable, periodic orbits, or even chaos. Such behavior arises in important evolutionary settings such as the dynamics of sexual evolution and that of grammar acquisition. In this paper we prove that the geometry of the dynamical system can also give tight mixing time bounds when the dynamical system has multiple fixed points and periodic orbits. We show that the mixing time continues to remain small in the presence of several unstable fixed points and is exponential in N when there are two or more stable fixed points. As a consequence of our results, we obtain a phase transition result for the mixing time of the sexual/grammar model mentioned above. We arrive at the conclusion that in the interesting parameter regime for these models, i.e., when there are multiple stable fixed points, the mixing is slow. Our techniques strengthen the connections between Markov chains and dynamical systems and we expect that the tools developed in this paper should have a wider applicability.

1998 ACM Subject Classification G.3 Probability and Statistics, F.2.2 Nonnumerical Algorithms and Problems

Keywords and phrases Markov chains, Mixing time, Dynamical Systems, Evolutionary dynamics, Language evolution

Digital Object Identifier 10.4230/LIPIcs.ICALP.2016.63

* Ioannis Panageas was supported by NSF EAGER award grants CCF-1415496 and CCF-1415498. Part of this work was done while he was visiting Simons Institute for the Theory of Computing at Berkeley.



© Ioannis Panageas and Nisheeth K. Vishnoi;
licensed under Creative Commons License CC-BY

43rd International Colloquium on Automata, Languages, and Programming (ICALP 2016).

Editors: Ioannis Chatzigiannakis, Michael Mitzenmacher, Yuval Rabani, and Davide Sangiorgi;
Article No. 63; pp. 63:1–63:14



Leibniz International Proceedings in Informatics

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany



1 Introduction

Evolutionary Markov chains and mixing time

In this paper we study Markov chains that arise in the context of evolution and which have also been used to model a wide variety of social, economical and cultural phenomena, see [18]. Typically, in such Markov chains, each state consists of a population of size N where each individual is of one of m types. Thus, the state space Ω has size $\binom{N+m-1}{m-1}$. At a very high level, in each iteration, the different types in the current generation reproduce according to their *fitnesses*, the reproduction could be *asexual* or *sexual* and have *mutations* that transform one type into another. This gives rise to an intermediate population that is subjected to the force of *selection*; a sample of size N is selected giving us the new generation. The specific way in which each of the *reproduction*, *mutation* and *selection* steps happen determine the transition matrix of the corresponding Markov chain. The size of the population (N), the number of types (m), the fitness of each type ($\{a_i \geq 0 : i \in [m]\}$), and the probabilities of mutation of one type to another ($\{Q_{ij} \geq 0 : i, j \in [m]\}$) are the parameters of the model. If we make the natural assumption that all the fitnesses are strictly positive and there is a non-zero probability of mutating from any type to the other, $Q_{ij} > 0$ for all $i, j \in [m]$, then the underlying chain is ergodic and has a unique steady state.

Most questions in evolution reduce to understanding the statistical properties of the steady state of an evolutionary Markov chain and how it changes with its parameters. However, in general, there seems to be no way to compute the desired statistical properties other than to sample from (close to) the steady state distribution by running the Markov chain for sufficiently long [5]. In the chains of interest, while there is an efficient way to sample the next state given the current state, typically, the state space is huge¹ This is captured by the notion of its *mixing time*. The mixing time of a Markov chain, t_{mix} , is defined to be the smallest time t such that for all $x \in \Omega$, the distribution of the Markov chain starting at x after t -time steps is within an ℓ_1 -distance of $1/4$ of the steady state.² Apart from dictating the computational feasibility of sampling procedures, the mixing time also gives us the number of generations required to reach a steady state; an important consideration for validating evolutionary models [24, 5]. However, despite the importance of understanding when an evolutionary Markov chain mixes fast (i.e., is significantly smaller than the size of the state space), until recently, there has been a lack of rigorous mixing time bounds for the full range of evolutionary parameters in even in the simplest of stochastic evolutionary models; see [7, 9, 8] for results under restricted assumptions and [5, 25] for an extended discussion on mixing time bounds in evolutionary Markov chains.

The expected motion of a Markov chain

In a recent result [20], a new approach for bounding the mixing time of such Markov chains was suggested. Towards this, it is convenient to think of each state of an evolutionary Markov chain as a vector which captures the fraction of each type in the current population. Thus, each state is a point in the m -dimensional probability simplex Δ_m ,³ and we can think of $\Omega \subseteq \Delta_m$. If $\mathbf{X}^{(t)}$ is the current state, then we define the *expected motion* of the chain at $\mathbf{X}^{(t)}$

¹ For example, even when $m = 40$ and the population is of size 10,000, the number of states is more than 2^{300} , i.e., more than the number of atoms in the universe!

² It is well-known that if one is willing to pay an additional factor of $\log 1/\varepsilon$, one can bring down the error from $1/4$ to ε for any $\varepsilon > 0$; see [14].

³ The probability simplex Δ_m is defined to be $\{p \in \mathbb{R}^m : p_i \geq 0 \forall i, \sum_i p_i = 1\}$.

to be the function

$$f(\mathbf{X}^{(t)}) \doteq \mathbb{E} \left[\mathbf{X}^{(t+1)} | \mathbf{X}^{(t)} \right]$$

where the expectation is over one step of the chain. Notice that while the domain of f is Ω , its range could be a larger subset of Δ_m . *What can the expected motion of a Markov chain tell us about the mixing time of a Markov chain?* Of course, without imposing additional structure on the Markov chain, we do not expect a very interesting answer. However, [20] suggested that, the expected motion can be helpful in establishing mixing time bounds, at least in the context of evolutionary dynamics. The first observation is that, while in the case of general Markov chains, the expected motion function is only defined at a subset of Δ_m , in the case of evolutionary Markov chains, the expected motion turns out to be a *dynamical system*; defined on *all* points of Δ_m . Further, the Markov chain can be recovered from the dynamical system: it can be shown that given a state $\mathbf{X}^{(t)}$ of the Markov chain, one can generate $\mathbf{X}^{(t+1)}$ *equivalently* by computing the probability distribution $f(\mathbf{X}^{(t)})$ and taking N i.i.d. samples from it. Subsequently, their main result is to prove that if this dynamical system has a *unique stable fixed point* and also all the trajectories converge to this point, then the evolutionary Markov chain mixes rapidly. Roughly, this is achieved by using the geometry of the dynamical system around this unique fixed point to construct a *contractive coupling*. As an application, this enabled them to establish rapid mixing for evolutionary Markov chains in which the reproduction is *asexual*.

What if the limit sets of the expected motion are complex: multiple fixed points – some stable and some unstable, or even periodic orbits? Not only are these natural mathematical questions given the previous work, such behavior arises in several important evolutionary settings; e.g., in the case when the reproduction is sexual (see [3, 17] and Chapter 20 in [11]) and an equivalent model for how children acquire grammar [19, 12]. While we describe these models later, we note that, as one changes the parameters of the model, the limit sets of the expected motion can exhibit the kind of complex behavior mentioned above and a finer understanding of how they influence the mixing time is desired.

Our contribution

In this paper we introduce prove that the geometry of the dynamical system can also give tight mixing time bounds when the dynamical system has multiple fixed points and periodic orbits. This completes the picture left open by the previous work. Recall that [20] proved that when there is a unique stable fixed point, then the mixing time is about $O(\log N)$ when N is large compared to the parameters of the model. We complement their result by proving the following mixing time bounds which depend on the structure of the limit sets of the expected motion:

- One stable fixed point and multiple unstable fixed points – the mixing time is $O(\log N)$, see Theorem 6.
- Multiple stable fixed points – the mixing time is $e^{\Omega(N)}$, see Theorem 7.
- Periodic orbits – the mixing time is $e^{\Omega(N)}$, see Theorem 8.

Thus, we can prove that despite the presence of unstable fixed points the mixing time continues to remain small. On the other hand, if there are two or more stable fixed points, the mixing time can undergo a phase transition and become exponential in N .

As an application, we characterize the mixing time of the dynamics of *grammar acquisition* (or, as explained later, sexual evolution). This Markov chain attempts to model a fascinating and important problem in linguistics; to understand the mechanism by which a child acquires

the capacity to comprehend a language and effectively communicate [16, 10]. Here, a parameter of interest is the *mutation rate* τ which is to be thought of as quantifying the *error of learning*; see Section 2.1. Corresponding to this, the probabilities of mutation $Q_{ij} = \tau$ for all $i \neq j$ and $Q_{ii} = 1 - (m-1)\tau$. We first prove that there is a critical value where the expected motion dynamical system goes through a bifurcation from multiple stable fixed points to one stable fixed point. Our main results then imply that for $\tau < \tau_c$ the mixing time is exponential in N and for $\tau > \tau_c$ it is $O(\log N)$, see Theorem 9. Thus, we arrive at the conclusion that, in the interesting parameter regime for an important and natural dynamics, i.e., when there is a stable fixed point other than the uniform one, the mixing is very slow.

Technically, there have been several influential works in the probability literature that use dynamical systems to analyze stochastic processes, see for example [22, 2, 26, 15]. While the techniques used in these results bear some similarity to ours, to the best of our knowledge, ours is the first paper which studies the question of how the *mixing time* of a Markov chain behaves as a function of the guiding dynamical system formally.

Organization of the paper

The rest of the paper is organized as follows. In Section 2 we present the formal statement of our main theorems and the model of grammar acquisition/sexual evolution. In Section 4, we present an overview of the proofs of our main theorem. Due to space constraints, details of the proofs have been omitted from this version and appear in the full version of this paper.

2 Formal statement of our results

In this section we present formal statements of our main results. We begin by introducing the required notation and preliminaries.

Notation

We use boldface letters, e.g., \mathbf{x} , to denote column vectors (points), and denote a vector's i^{th} coordinate by x_i . We use \mathbf{X} and \mathbf{Y} (often with time superscripts and coordinate subscripts as appropriate) to denote random vectors. For a function $f : \Delta_m \rightarrow \Delta_m$, by f^n we denote the composition of f with itself n times, namely $\underbrace{f \circ f \circ \dots \circ f}_{n \text{ times}}$. We use $J_f[\mathbf{x}]$ to denote the

Jacobian matrix of f at the point \mathbf{x} . When the function f is clear from the context, we omit the subscript and simply denote it by $J[\mathbf{x}]$. Similarly, we sometimes use $J^n[\mathbf{x}]$ to denote the Jacobian of f^n at \mathbf{x} . We denote by $\text{sp}(A)$ the spectral radius of a matrix A and by $(A\mathbf{x})_i$ the sum $\sum_j A_{ij}x_j$.

Dynamical Systems

Let $\mathbf{x}^{(t+1)} = f(\mathbf{x}^{(t)})$ be a *discrete time* dynamical system with update rule $f : \Delta_m \rightarrow \Delta_m$. The point \mathbf{z} is called a *fixed point* of f if $f(\mathbf{z}) = \mathbf{z}$. We call a fixed point \mathbf{z} *stable* if, for the Jacobian $J[\mathbf{z}]$ of f , it holds that $\text{sp}(J[\mathbf{z}]) < \rho < 1$. A sequence $(f^t(\mathbf{x}^{(0)}))_{t \in \mathbb{N}}$ is called a *trajectory* of the dynamics with $x^{(0)}$ as starting point. A common technique to show that a dynamical system converges to a fixed point is to construct a function $P : \Delta_m \rightarrow \mathbb{R}$ such that $P(f(\mathbf{x})) > P(\mathbf{x})$ unless \mathbf{x} is a fixed point. We call P a *potential* function. One of our results deals with dynamical systems that have stable periodic orbits.

► **Definition 1.** $C = \{\mathbf{x}_1, \dots, \mathbf{x}_k\}$ is called a *periodic orbit* of size k if $\mathbf{x}_{i+1} = f(\mathbf{x}_i)$ for $1 \leq i \leq k-1$ and $f(\mathbf{x}_k) = \mathbf{x}_1$. If $\text{sp}(J_{f^k}[\mathbf{x}_1]) < \rho < 1$, we call C a *stable periodic orbit* (we also use the terminology *stable limit cycle*).

► **Remark.** Since $f : \Delta_m \rightarrow \Delta_m$ and hence $\sum_i f_i(\mathbf{x}) = 1$ for all $\mathbf{x} \in \Delta_m$, if we define $h_i(\mathbf{x}) = \frac{f_i(\mathbf{x})}{\sum_i f_i(\mathbf{x})}$ so that $h(\mathbf{x}) = f(\mathbf{x})$ for all $\mathbf{x} \in \Delta_m$, we get that $\sum_i \frac{\partial h_i(\mathbf{x})}{\partial x_j} = 0$ for all $j \in [m]$. This means without loss of generality we can assume that the Jacobian $J[\mathbf{x}]$ of f has $\mathbf{1}^\top$ (the all-ones vector) as a left eigenvector with eigenvalue 0.

The definition below quantifies the instability of a fixed point as is standard in the literature. Essentially, an α unstable fixed point is repelling in any direction.

► **Definition 2.** Let \mathbf{z} be a fixed point of a dynamical system f . The point \mathbf{z} is called α -*unstable* if $|\lambda_{\min}(J[\mathbf{z}])| > \alpha > 1$ where λ_{\min} corresponds to the minimum eigenvalue of the Jacobian of f at the fixed point \mathbf{z} , excluding the eigenvalue 0 that corresponds to the left eigenvector $\mathbf{1}^\top$.

Stochastic Evolution

► **Definition 3.** Given an $f : \Delta_m \rightarrow \Delta_m$ which is smooth,⁴ and a population parameter N , we define a Markov chain called the *stochastic evolution guided by f* as follows. The state at time t is a probability vector $\mathbf{X}^{(t)} \in \Delta_m$. The state $\mathbf{X}^{(t+1)}$ is then obtained in the following manner. Define $\mathbf{Y}^{(t)} = f(\mathbf{X}^{(t)})$. Obtain N independent samples from the probability distribution $\mathbf{Y}^{(t)}$, and denote by $\mathbf{Z}^{(t)}$ the resulting counting vector over $[m]$. Then

$$\mathbf{X}^{(t+1)} := \frac{1}{N} \mathbf{Z}^{(t)} \quad \text{and therefore } \mathbb{E}[\mathbf{X}^{(t+1)} | \mathbf{X}^{(t)}] = f(\mathbf{X}^{(t)}).$$

We call f the *expected motion* of the stochastic evolution.

► **Definition 4 (Smooth contractive evolution).** A function $f : \Delta_m \rightarrow \Delta_m$ is said to be a *smooth contractive evolution* if it is smooth⁴, has a *unique* fixed point \mathbf{z} in the interior of Δ_m , this unique point is *stable*, and, for every $\varepsilon > 0$, there exists an ℓ such that for any $\mathbf{x} \in \Delta_m$, it holds that $\|f^\ell(\mathbf{x}) - \mathbf{z}\|_1 < \varepsilon$ (i.e., f converges to the fixed point).

The main result in [20] was Theorem 5 below. This theorem gives a bound on the mixing time of a stochastic evolution guided by a function f that satisfies Definition 4.

► **Theorem 5 (Main theorem in [20]).** *Let f be a smooth contractive evolution, and let \mathcal{M} be the stochastic evolution guided by f on a population of size N . Then, the mixing time of \mathcal{M} is $O(\log N)$.*

Our Results

Given a dynamical system f , one of the main questions that one can ask is does it converge, and if so, how fast. In general, if the behavior of a system is non-chaotic, we expect the system to reach some steady state (e.g., a fixed point or periodic orbit). This steady state might be some (local) optimum solution to a non-linear optimization problem. Therefore, it is important to understand what traits make a dynamical system converge fast. The

⁴ For our purposes, we call a function f is *smooth* if it is twice differentiable in the relative interior of Δ_m with bounded second derivative.

existence of many fixed points which are unstable can slow down the speed of convergence of a dynamical system. In the case of the stochastic evolution guided by f , one would expect the existence of multiple unstable fixed points to similarly slow down the mixing time. Nevertheless, our Theorem 6 shows rapid mixing in the presence of α -unstable fixed points. Additionally, we change the assumption *convergence to the fixed point* in 5 to the assumption that for all $\mathbf{x} \in \Delta_m$ the limit $\lim_{t \rightarrow \infty} f^t(\mathbf{x})$ exists and is equal to some fixed point \mathbf{z} , i.e., as in 5, there are no limit cycles.

► **Theorem 6.** *Let $f : \Delta_m \rightarrow \Delta_m$ be twice differentiable in the interior of Δ_m with bounded second derivative. Assume that $f(\mathbf{x})$ has a finite number of fixed points $\mathbf{z}_0, \dots, \mathbf{z}_l$ in the interior, where \mathbf{z}_0 is a stable fixed point, i.e., $\text{sp}(J[\mathbf{z}_0]) < \rho < 1$ and $\mathbf{z}_1, \dots, \mathbf{z}_l$ are α -unstable fixed points ($\alpha > 1$). Furthermore, assume that $\lim_{t \rightarrow \infty} f^t(\mathbf{x})$ exists for all $\mathbf{x} \in \Delta_m$. Then, the stochastic evolution guided by f has mixing time $O(\log N)$.*

In our second result, we allow f to have multiple stable fixed points (in addition to any number of unstable fixed points). For this setting, we prove that the stochastic evolution guided by f has mixing time $e^{\Omega(N)}$. Our phase transition result on a linguistic/sexual evolution model discussed in Section 2.1 relies crucially on Theorem 7.

► **Theorem 7.** *Let $f : \Delta_m \rightarrow \Delta_m$ be continuously differentiable in the interior of Δ_m . Assume that $f(\mathbf{x})$ has at least two stable fixed points in the interior $\mathbf{z}_1, \dots, \mathbf{z}_l$, i.e., $\text{sp}(J[\mathbf{z}_i]) < \rho_i < 1$ for $i = 1, 2, \dots, l$. Then, the stochastic evolution guided by f has mixing time $e^{\Omega(N)}$.*

Finally, we allow f to have a stable limit cycle. We prove that in this setting the stochastic evolution guided by f has mixing time $e^{\Omega(N)}$. This result seems important for evolutionary dynamics as periodic orbits often appear [23, 21].

► **Theorem 8.** *Let $f : \Delta_m \rightarrow \Delta_m$ be continuously differentiable in the interior of Δ_m . Assume that $f(\mathbf{x})$ has a stable limit cycle with points $\mathbf{w}_1, \dots, \mathbf{w}_s$ of size $s \geq 2$ in the sense that $\text{sp}(\prod_{i=1}^s J[\mathbf{w}_{s-i+1}]) < \rho < 1$. Then the stochastic evolution guided by f has mixing time $e^{\Omega(N)}$.*

2.1 Dynamics of grammar acquisition and sexual evolution

We begin by describing the evolutionary processes for grammar acquisition and sexual evolution. As we will explain, the two turn out to be identical and hence we primarily focus on the model for grammar acquisition in the remainder of the paper.

The starting point of the model is Chomsky's *Universal Grammar* theory [4].⁵ In his theory, language learning *is facilitated by a predisposition that our brains have for certain structures of language*. This universal grammar (UG) is believed to be innate and embedded in the neuronal circuitry. Based on this theory, an influential model for how children acquire grammar was given by appealing to evolutionary dynamics for infinite and finite populations respectively in [19] and [12]. We first describe the infinite population model, which is a dynamical system that guides the stochastic, finite population model. Each individual speaks exactly one of the m grammars from the set of inherited UGs $\{G_1, \dots, G_m\}$; denote by x_i the fraction of the population using G_i . The model associates a *fitness* to every individual on the basis of the grammar she *and others* use. Let A_{ij} be the probability that a person who speaks grammar j understands a randomly chosen sentence spoken by an individual

⁵ Like any important problem in the sciences, Chomsky's theory is not uncontroversial; see [10] for an in-depth discussion.

using grammar i . This can be viewed as the fraction of sentences according to grammar i that are also valid according to grammar j . Clearly, $A_{ii} = 1$. The pairwise *compatibility* between two individuals speaking grammars i and j is $B_{ij} := \frac{A_{ij} + A_{ji}}{2}$, and the fitness of an individual using G_i is $f_i := \sum_{j=1}^m x_j B_{ij}$, i.e., the probability that such an individual is able to meaningfully communicate with a randomly selected member of the population.

In the reproduction phase each individual produces a number of offsprings proportional to her fitness. Each child speaks one grammar, but the exact learning model can vary and allows for the child to incorrectly learn the grammar of her parent. We define the matrix Q where the entry Q_{ij} denotes the probability that the child of an individual using grammar i learns grammar j (i.e. Q is column stochastic matrix); once a child learns a grammar it is fixed and she does not later use a different grammar. Thus, the frequency x'_i of the individuals that use grammar G_i in the next generation will be

$$x'_i = g_i(\mathbf{x}) := \sum_{j=1}^m \frac{Q_{ji} x_j (B\mathbf{x})_j}{\mathbf{x}^\top B\mathbf{x}}$$

(with $g : \Delta_m \mapsto \Delta_m$ encoding the update rule). Nowak et al. [19] study the symmetric case, i.e., $B_{ij} = b$ and $Q_{ij} = \tau \in (0, 1/m]$ for all $i \neq j$ and observe a threshold: When τ , which can be thought of as quantifying the *error of learning or mutation*, is above a critical value, the only stable fixed point is the uniform distribution (all $1/m$) and below it, there are multiple stable fixed points.

Finite population models can be derived from the linguistic dynamics in a standard way. We describe the Wright-Fisher finite population model for the linguistic dynamics. The population size remains N at all times and the generations are non-overlapping. The current state of the population is described by the frequency vector $\mathbf{X}^{(t)}$ at time t which is a random vector in Δ_m and notice also that the population that uses G_i is $NX_i^{(t)}$. How does one generate $\mathbf{X}^{(t+1)}$? To do this, in the replication (R) stage, one first replaces the individuals that speak grammar G_i in the current population by $NX_i^{(t)}(B(N\mathbf{X}^{(t)}))_i$ and the total population has size $N^2\mathbf{X}^{(t)\top}B\mathbf{X}^{(t)}$.⁶ In the selection (S) stage, one selects N individuals from this population by sampling independently with replacement. Since the evolution is error prone, in the mutation (M) stage, the grammar of each individual in this intermediate population is mutated independently at random according to the matrix Q to obtain frequency vector $\mathbf{X}^{(t+1)}$. Given these rules, note that

$$\mathbb{E}[\mathbf{X}^{(t+1)} | \mathbf{X}^{(t)}] = g(\mathbf{X}^{(t)}).$$

In other words, in expectation, fixing $\mathbf{X}^{(t)}$, the next generation's frequency vector $\mathbf{X}^{(t+1)}$ is exactly $g(\mathbf{X}^{(t)})$, where g is the linguistic dynamics. Of course, this holds only for one step of the process. This process is a Markov chain with state space $\{(y_1, \dots, y_m) : y_i \in \mathbb{N}, \sum_i y_i = N\}$ of size $\binom{N+m-1}{m-1}$. If $Q > 0$ then it is ergodic (i.e., it is irreducible and aperiodic) and thus has a unique stationary distribution. In our analysis, we consider the symmetric case as in Nowak et al. [19], i.e., $B_{ij} = b$ and $Q_{ij} = \tau \in (0, 1/m]$ for all $i \neq j$.

Note that the linguistics model described above can also be seen as a (finite population) sexual evolution model: Assume there are N individuals and m types. Let $\mathbf{Y}^{(t)}$ be a vector of frequencies at time t , where $Y_i^{(t)}$ denotes the fraction of individuals of type i . Let F be a fitness matrix where F_{ij} corresponds to the number of offspring of type i , if an individual of

⁶ Here we assume that B_{ij} is an positive integer and thus $N^2X_i^{(t)}(B\mathbf{X}^{(t)})_i$ is an integer since the individuals are whole entities; this can be achieved by scaling and is without loss of generality.

type i chooses to mate with an individual of type j (assume $F_{ij} \in \mathbb{N}$). At every generation, each individual mates with every other individual. It is not hard to show that the number of offspring after the matings will be $N^2(\mathbf{Y}^{(t)\top} F \mathbf{Y}^{(t)})$ and there will be $N^2 \mathbf{Y}_i^{(t)}(F \mathbf{Y}^{(t)})_i$ individuals of type i . After the reproduction step, we select N individuals at random with replacement, i.e., we sample an individual of type i with probability $\frac{\mathbf{Y}_i^{(t)}(F \mathbf{Y}^{(t)})_i}{\mathbf{Y}^{(t)\top} F \mathbf{Y}^{(t)}}$. Finally in the mutation step, every individual of type i mutates with probability τ (*mutation parameter*) to some type j . Let $F_{ii} = A$, $F_{ij} = B$ for all $i \neq j$ with $A > B$ (this is called homozygote advantage) and set $b = \frac{B}{A} < 1$. It is self-evident that this sexual evolution model is identical with the (finite population) linguistic model described above since both end up having the same reproduction, selection and mutation rule. It holds that $E[\mathbf{X}^{(t+1)} | \mathbf{X}^{(t)}] = g(\mathbf{X}^{(t)})^7$ with

$$g_i(\mathbf{x}) = (1 - (m-1)\tau) \frac{N^2 x_i (B \mathbf{x})_i}{N^2 (\mathbf{x}^\top B \mathbf{x})} + \sum_{j \neq i} \tau \frac{N^2 x_j (B \mathbf{x})_j}{N^2 (\mathbf{x}^\top B \mathbf{x})} = (1 - m\tau) \frac{x_i (B \mathbf{x})_i}{(\mathbf{x}^\top B \mathbf{x})} + \tau$$

where $B_{ii} = 1$, $B_{ij} = b$ with $i \neq j$.⁸ For the Markov chains described above (symmetric case) we can prove the following phase transition result.

► **Theorem 9.** *There is a critical value τ_c of the error in learning/mutation parameter τ such that the mixing time is: (i) $\exp(\Omega(N))$ for $0 < \tau < \tau_c$ and (ii) $O(\log N)$ for $\tau > \tau_c$ where N is the size of the population.*

The theorem below will be used to prove the rapid mixing result for the finite linguistic model when $\tau > \tau_c$. It is used to construct a potential function and show that the deterministic dynamics g converges to fixed points.

► **Theorem 10 (Baum and Eagon Inequality [1]).** *Let $P(\mathbf{x}) = P(\{x_{ij}\})$ be a polynomial with nonnegative coefficients homogeneous of degree d in its variables $\{x_{ij}\}$. Let $\mathbf{x} = \{x_{ij}\}$ be any point of the domain $D : x_{ij} \geq 0, \sum_{j=1}^{q_i} x_{ij} = 1, i = 1, \dots, p, j = 1, \dots, q_i$. For $\mathbf{x} = \{x_{ij}\} \in D$, let $\Xi(\mathbf{x}) = \Xi\{x_{ij}\}$ denote the point of D whose i, j -th coordinate is*

$$\Xi(\mathbf{x})_{ij} = \left(x_{ij} \frac{\partial P}{\partial x_{ij}} \Big|_{(\mathbf{x})} \right) \cdot \left(\sum_{j=1}^{q_i} x_{ij} \frac{\partial P}{\partial x_{ij}} \Big|_{(\mathbf{x})} \right)^{-1}.$$

Then $P(\Xi(\mathbf{x})) > P(\mathbf{x})$ unless $\Xi(\mathbf{x}) = \mathbf{x}$.

3 Preliminaries

Couplings and Mixing Times

Let $\mathbf{p}, \mathbf{q} \in \Delta_m$ be two probability distributions on m objects. A *coupling* \mathcal{C} of \mathbf{p} and \mathbf{q} is a distribution on ordered pairs in $[m] \times [m]$, such that its marginal distribution on the first coordinate is equal to \mathbf{p} and that on the second coordinate is equal to \mathbf{q} . Couplings allow a very useful dual characterization of the total variation distance, as stated in the following well known lemma.

► **Lemma 11 (Coupling lemma [14]).** *Let $\mathbf{p}, \mathbf{q} \in \Delta_m$ be two probability distributions on m objects. Then,*

$$\|\mathbf{p} - \mathbf{q}\|_{\text{TV}} = \frac{1}{2} \|\mathbf{p} - \mathbf{q}\|_1 = \min_{\mathcal{C}} \mathbb{P}_{(A,B) \sim \mathcal{C}} [A \neq B],$$

where the minimum is taken over all valid couplings \mathcal{C} of \mathbf{p} and \mathbf{q} .

⁷ We use same notation for the update rule as before, i.e. g because it turns out to be the same function.

⁸ Observe that this rule is invariant under scaling of fitness matrix B .

► **Definition 12** (Mixing time [14]). Let \mathcal{M} be an ergodic Markov chain on a finite state space Ω with stationary distribution π . Then, the *mixing time* $t_{\text{mix}}(\varepsilon)$ is defined as the smallest time such that for any starting state $\mathbf{X}^{(0)}$, the distribution of the state $\mathbf{X}^{(t)}$ at time t is within total variation distance ε of π . The term mixing time is also used for $t_{\text{mix}}(\varepsilon)$ for a *fixed* values of $\varepsilon < 1/2$.

A well-known technique for obtaining upper bounds on mixing times is to use the Coupling Lemma above. Suppose $\mathbf{X}^{(t)}$ and $\mathbf{Y}^{(t)}$ are two evolutions of an ergodic chain \mathcal{M} such that their evolutions are coupled according to some coupling \mathcal{C} . Let T be the smallest time such that $\mathbf{X}^{(T)} = \mathbf{Y}^{(T)}$. If it can be shown that $\mathbb{P}[T > t] \leq 1/4$ for every pair of starting states $(\mathbf{X}^{(0)}, \mathbf{Y}^{(0)})$, then it follows that $t_{\text{mix}} := t_{\text{mix}}(1/4) \leq t$.

Operators, Norms

The following theorem, stated here only in the special case of the $1 \rightarrow 1$ norm, relates the spectral radius with other matrix norms.

► **Theorem 13** (Gelfand's formula, specialized to the $1 \rightarrow 1$ norm [13]). *For any square matrix A , we have*

$$\text{sp}(A) = \lim_{\ell \rightarrow \infty} \|A^\ell\|_{1 \rightarrow 1}^{1/\ell}.$$

Taylor Theorem (First order Remainder)

► **Theorem 14.** *Let $f : \mathbb{R}^m \rightarrow \mathbb{R}$ be differentiable and $\mathbf{x}, \mathbf{y} \in \mathbb{R}^m$. Then there exists some ξ in the line segment from \mathbf{x} to \mathbf{y} such that $f(\mathbf{y}) = f(\mathbf{x}) + \nabla f(\xi)(\mathbf{y} - \mathbf{x})$.*

Concentration

We also mention some standard Chernoff-Hoeffding type bounds that will be used in our later arguments.

► **Theorem 15** (Chernoff-Hoeffding bounds [6]). *Let Z_1, Z_2, \dots, Z_N be i.i.d. Bernoulli random variables with mean μ . We then have for all $\varepsilon > 0$,*

$$\mathbb{P} \left[\left| \frac{1}{N} \sum_{i=1}^N Z_i - \mu \right| > \varepsilon \right] \leq 2 \exp(-2N\varepsilon^2).$$

4 Overview of proofs

We begin by explaining the proof technique of Theorem 5 in [20]. In order to prove a bound on the mixing time, the authors constructed a coupling that contracts the distance between two chains. This contraction does not happen at every step, rather at every k steps where k is some constant and depends on the function f . Essentially, it is shown that given two chains $\mathbf{X}^{(t)}, \mathbf{Y}^{(t)}$ that are close to the unique fixed point \mathbf{z} of f , it holds that

$$\left\| \mathbf{X}^{(t+1)} - \mathbf{Y}^{(t+1)} \right\|_1 \approx \left\| J[\mathbf{z}](\mathbf{X}^{(t)} - \mathbf{Y}^{(t)}) \right\|_1 \leq \|J[\mathbf{z}]\|_1 \left\| \mathbf{X}^{(t)} - \mathbf{Y}^{(t)} \right\|_1$$

with high probability due to Chernoff bounds. Thus, the ℓ_1 norm of the Jacobian captures the contraction if it indeed exists. However it might be the case that $\|J[\mathbf{z}]\|_1 > 1$. On the positive side, using Gelfand's Theorem they were able to show a k -step contraction, since

$$\left\| J^k[\mathbf{z}] \right\|_1 \approx (\text{sp}(J[\mathbf{z}]))^k < \rho^k < 1$$

for some $k \in \mathbb{N}$. Our proofs also use the idea of Gelfand's formula to show contraction/expansion (in Theorems 7 and 6 respectively) and also make use of Theorem 5). Nevertheless, there are important technical barriers that need to be crossed in order to prove our results as explained below.

4.1 Overview of Theorem 6

The main difficulty to prove this theorem is the existence of multiple unstable fixed points in the simplex from which the Markov chain should get away fast. As before, we study the time T required for two stochastic evolutions with arbitrary initial states $\mathbf{X}^{(0)}$ and $\mathbf{Y}^{(0)}$, guided by some function f , to collide. By the conditions of Theorem 6, function f has a unique stable fixed point \mathbf{z}_0 with

$$\text{sp}(J[\mathbf{z}_0]) < \rho < 1.$$

Additionally, it has α -unstable fixed points. Moreover, for all starting points $\mathbf{x}_0 \in \Delta_m$, the sequence $(f^t(\mathbf{x}_0))_{t \in \mathbb{N}}$ has a limit. We can show that there exists constant c_0 such that $\mathbb{P}[T > c_0 \log N] \leq \frac{1}{4}$, from which it follows that $t_{\text{mix}}(1/4) \leq c_0 \log N$. In order to show collision after $O(\log N)$ steps, it suffices first to run each chain *independently* for $O(\log N)$ steps. We first show that with probability $\Theta(1)$, each chain will reach $B(\mathbf{z}_0, \frac{1}{N^{1-\varepsilon}})$ after at most $O(\log N)$ steps, for some $\varepsilon > 0$.⁹ As long as this is true, the coupling constructed in [20] can be used to show collision (see Section 3 for the definition of a coupling). To explain why our claim holds, we break the proof into three parts.

(a) First, it is shown that as long as the state of the Markov chain is within $o\left(\frac{\log^{2/3} N}{\sqrt{N}}\right)$ in ℓ_1 distance from some α -unstable fixed point \mathbf{w} , then, with probability $\Theta(1)$, it reaches distance $\Omega\left(\frac{\log^{2/3} N}{\sqrt{N}}\right)$ after $O(\log N)$ steps. Step (a) has the technical difficulty that as long as a chain starts from a $o\left(\frac{1}{\sqrt{N}}\right)$ distance from an unstable fixed point, the variance of the process dominates the expansion due to the fact the fixed point is unstable.

(b) Assuming (a), we show that with probability $1 - \frac{1}{\text{poly}(N)}$ the Markov chain reaches distance $\Theta(1)$ from any unstable fixed point after $O(\log N)$ steps.

(c) Finally, if the Markov chain has $\Theta(1)$ distance from any unstable fixed point (the fixed points have pairwise ℓ_1 distance independent of N , i.e., they are “well separated”), it will reach some $\frac{1}{N^{1-\varepsilon}}$ -neighborhood of the stable fixed point \mathbf{z}_0 exponentially fast (i.e., after $O(\log N)$ steps). For showing (a) and (b), we must prove an expansion argument for $\|f^t(\mathbf{x}) - \mathbf{w}\|_1$ as t increases, where \mathbf{w} is an α -unstable fixed point and also taking care of the random perturbations due to the stochastic evolution. Ideally what we want (but is not true) is the following to hold:

$$\|f^{t+1}(\mathbf{x}) - \mathbf{w}\|_1 \geq \alpha \|f^t(\mathbf{x}) - \mathbf{w}\|_1,$$

i.e., one step expansion. The first important fact is that f^{-1} is well-defined in a small neighborhood of \mathbf{w} due to the Inverse Function Theorem, and it also holds that

$$\|f^t(\mathbf{x}) - \mathbf{w}\|_1 \approx \|J^{-1}[\mathbf{w}](f^{t+1}(\mathbf{x}) - \mathbf{w})\|_1 \leq \|J^{-1}[\mathbf{w}]\|_1 \|f^{t+1}(\mathbf{x}) - \mathbf{w}\|_1,$$

where \mathbf{x} is in some neighborhood of \mathbf{w} and $J^{-1}[\mathbf{w}]$ is the *pseudoinverse* of $J[\mathbf{w}]$ (see the remark in Section 2). However even if \mathbf{w} is α -unstable and $\text{sp}(J^{-1}[\mathbf{w}]) < \frac{1}{\alpha}$, it can hold that

⁹ $B(\mathbf{x}, r)$ denotes the open ball with center \mathbf{x} and radius r in ℓ_1 , which we call an r -neighborhood of \mathbf{x} .

$\|J^{-1}[\mathbf{w}]\|_1 > 1$. At this point, we use Gelfand's formula (Theorem 13) as in the proof of [20]. Since $\lim_{t \rightarrow \infty} (\|A^t\|_1)^{1/t} \rightarrow \text{sp}(A)$, for all $\varepsilon > 0$, there exists a k_0 such that for all $k \geq k_0$ we have

$$\left| \|A^k\|_1 - (\text{sp}(A))^k \right| < \varepsilon.$$

We use this important theorem to show that for small $\varepsilon > 0$, there exists a k such that

$$\|f^t(\mathbf{x}) - \mathbf{w}\|_1 \approx \|(J^{-1}[\mathbf{w}])^k(f^{t+k}(\mathbf{x}) - \mathbf{w})\|_1 \leq \frac{1}{\alpha^k} \|f^{t+k}(\mathbf{x}) - \mathbf{w}\|_1,$$

where we used the fact that

$$\|(J^{-1}[\mathbf{w}])^k\|_1 < (\text{sp}(J^{-1}[\mathbf{w}]))^k - \varepsilon \leq \frac{1}{\alpha^k}.$$

By taking advantage of the continuity of the $J^{-1}[\mathbf{x}]$ around the unstable fixed point \mathbf{w} , we can show expansion for every k steps of the dynamical system. It remains to show for (a) and (b) how one can handle the perturbations due to the randomness of the stochastic evolution. In particular, if $\|\mathbf{X}^{(0)} - \mathbf{w}\|_1$ is $o\left(\frac{1}{\sqrt{N}}\right)$, even with the expansion we have from the deterministic dynamics (as discussed above), variance dominates. We examine case (b) first, which is relatively easy (the *drift* dominates at this step). Due to Chernoff bounds, the difference $\|\mathbf{X}^{(t+k)} - \mathbf{w}\|_1 - \|f^k(\mathbf{X}^{(t)}) - \mathbf{w}\|_1$ is $O\left(\sqrt{\frac{\log N}{N}}\right)$ (this captures the deviation on running the stochastic evolution for k steps vs running the deterministic dynamics for k steps, both starting from $\mathbf{X}^{(t)}$) with probability $1 - \frac{1}{\text{poly}(N)}$. Since $\|\mathbf{X}^{(t)} - \mathbf{w}\|_1$ is $\Omega\left(\frac{\log^{2/3} N}{\sqrt{N}}\right)$, then

$$\|\mathbf{X}^{(t+k)} - \mathbf{w}\|_1 \geq (\alpha^k - o_N(1)) \|\mathbf{X}^{(t)} - \mathbf{w}\|_1.$$

For (a), first we show that with probability $\Theta(1)$, after one step the Markov chain has distance $\Omega\left(\frac{1}{\sqrt{N}}\right)$ of \mathbf{w} . This claim just uses properties of the multinomial distribution. After reaching distance $\Omega\left(\frac{1}{\sqrt{N}}\right)$, we can use again the idea of expansion and being careful with the variance and we can show expansion with probability at least $\frac{1}{2}$, every k steps. Then we can show that with probability at least $\frac{1}{\log^{2/3} N}$, distance $\frac{\log^{2/3} N}{\sqrt{N}}$ is reached after $O(\log \log N)$ steps and basically we finish with (b). For (c), we use a couple of technical lemmas from [20], we explain in words below: Let Δ be some compact subset of Δ_m , where we have excluded all the α -unstable fixed points along with some open ball around each unstable fixed point of constant radius. We can show that given that the initial state of the Markov chain belongs to Δ , it reaches a $B(\mathbf{z}_0, \frac{1}{N^{1-\varepsilon}})$ for some $\varepsilon > 0$ as long as the dynamical system converges for all starting points in Δ (and it should converge to the stable fixed point \mathbf{z}_0). We have roughly that the dynamical system converges exponentially fast for every starting point in B to the stable fixed point \mathbf{z}_0 and that with probability $1 - \frac{1}{\text{poly}(n)}$ two arbitrary chains independently will reach a $\frac{1}{N^\varepsilon}$ neighborhood of the stable fixed point \mathbf{z}_0 . Therefore by (a), (b), (c) and the coupling from [20], we conclude the proof of Theorem 6.

4.2 Overview of Theorems 7 and 8

To prove Theorem 8, we make use of Theorem 7, i.e., we reduce the case of the stable limit cycle to the case of multiple stable fixed points. If s is the length of the limit cycle, roughly the bound $e^{\Omega(N)}$ on the mixing time loses a factor $\frac{1}{s}$ compared to the case of multiple stable

fixed points. We now present the ideas behind the proof of Theorem 7. First as explained above, we can show contraction after k steps (for some constant k) for the deterministic dynamics around a stable fixed point \mathbf{z} with $\text{sp}(J[\mathbf{z}]) < \rho < 1$, i.e.,

$$\|f^{t+k}(\mathbf{x}) - \mathbf{z}\|_1 \approx \|J^k[\mathbf{z}]\|_1 \|f^t(\mathbf{x}) - \mathbf{z}\|_1 \leq \rho^k \|f^t(\mathbf{x}) - \mathbf{z}\|_1.$$

To do that, we use Gelfand's formula, Taylor's theorem and continuity of $J[\mathbf{x}]$ where \mathbf{x} lies in a neighborhood of the fixed point \mathbf{z} . Hence, due to the above contraction of the ℓ_1 norm and the concentration of Chernoff bounds, it takes a long time for the chain $\mathbf{X}^{(t)}$ to get out of the region of attraction of the fixed point \mathbf{z} . Technically, the error that aggregates due to the randomness of the stochastic evolution guided by f does not become large due to the convergence of the series $\sum_{i=0}^{\infty} \rho^i$. Hence, we focus on the error probability, namely the probability the stochastic evolution guided by f deviates a lot from the dynamical system with rule f if both have same starting point after one step. Since this probability is exponentially small, i.e., it holds that

$$\|f(\mathbf{X}^{(0)}) - \mathbf{X}^{(1)}\|_1 > \varepsilon m$$

with probability at most $2me^{-2\varepsilon^2 N}$, an exponential number of steps is required for the above to be violated. Finally, as we have shown that it takes exponential time to get out of the region of attraction of a stable fixed point \mathbf{z} we do the following easy (common) trick. Since the function has at least two fixed points, we start the Markov chain very close to the fixed point that its neighborhood has mass at most $1/2$ in the stationary distribution (this can happen since we have at least 2 fixed points that are well separated). Then, after exponential number of steps, it will follow that the total variation distance between the distribution of the chain and the stationary will be at least $1/4$.

4.3 Overview of Theorem 9

Below we give the necessary ingredients of the proof of Theorem 9. Our previous results, along with some analysis on the fixed points of g (function of Linguistic Dynamics) suffice to show the phase transition result. To prove Theorem 9, initially we show that the model (finite population) is essentially a stochastic evolution (see Definition 3) guided by g as defined in Section 2.1 and proceed as follows: We prove that in the interval $0 < \tau < \tau_c$, the function g has multiple fixed points whose Jacobian have spectral radius less than 1. Therefore due to Theorem 7 discussed above, the mixing time will be exponential in N . For $\tau = \tau_c$ a bifurcation takes place which results in function g of linguistic dynamics having only one fixed point inside simplex (specifically, the uniform point $(1/m, \dots, 1/m)$). In dynamical systems, a local bifurcation occurs when a parameter (in particular the mutation parameter τ) change causes two (or more) fixed points to collide or the stability of an equilibrium (or fixed point) to change. To prove fast mixing in the case $\tau_c < \tau \leq 1/m$, we make use of the result in [20] (see Theorem 5). One of the assumptions is that the dynamical system with g as update rule needs to converge to the unique fixed point for all initial points in simplex. To prove convergence to the unique fixed point, we define a Lyapunov function P such that

$$P(g(\mathbf{x})) > P(\mathbf{x}) \text{ unless } \mathbf{x} \text{ is a fixed point.} \quad (1)$$

As a consequence, the (infinite population) linguistic dynamics converge to the unique fixed point $(1/m, \dots, 1/m)$. To show Equation (1), we use an inequality that dates back in 1967 (see Theorem 10, [1]), which intuitively states the discrete analogue of proving that for a gradient system $\frac{d\mathbf{x}}{dt} = \nabla V(\mathbf{x})$ it is true that $\frac{dV}{dt} \geq 0$.

References

- 1 L. Baum and J. Eagon. An inequality with applications to statistical estimation for probabilistic functions of Markov processes and to a model for ecology. *Bull. Amer. Math. Soc.*, 73:360–363, 1967.
- 2 Michel Benaïm. Dynamics of stochastic approximation algorithms. In *Seminaire de probabilites XXXIII*, pages 1–68. Springer, 1999.
- 3 Erick Chastain, Adi Livnat, Christos Papadimitriou, and Umesh Vazirani. Algorithms, games, and evolution. *Proceedings of the National Academy of Sciences*, 2014.
- 4 Noam A. Chomsky. Rules and Representations. *Behavioral and Brain Sciences*, 3(127):1–61, 1980.
- 5 Narendra Dixit, Piyush Srivastava, and Nisheeth K. Vishnoi. A finite population model of molecular evolution: Theory and computation. *Journal of Computational Biology*, 19(10):1176–1202, 2012.
- 6 Devdatt P. Dubhashi and Alessandro Panconesi. *Concentration of Measure for the Analysis of Randomized Algorithms*. Cambridge University Press, 2009.
- 7 Richard Durrett. *Probability models for DNA sequence evolution*. Springer, 2008.
- 8 Stewart N Ethier and Thomas G Kurtz. *Markov processes: characterization and convergence*, volume 282. John Wiley & Sons, 2009.
- 9 Warren J. Ewens. *Mathematical Population Genetics I. Theoretical Introduction*. Springer, 2004.
- 10 W.T. Fitch. *The Evolution of Language*. Approaches to the Evolution of Language. Cambridge University Press, 2010.
- 11 J. Hofbauer and K. Sigmund. *Evolutionary Games and Population Dynamics*. Cambridge University Press, 1998.
- 12 Natalia L. Komarova and Martin A. Nowak. Language dynamics in finite populations. *Journal of Theoretical Biology*, 221(3):445–457, 2003.
- 13 Erwin Kreyszig. *Introductory Functional Analysis with Applications*. Wiley, 1978.
- 14 David A. Levin, Yuval Peres, and Elizabeth L. Wilmer. *Markov chains and mixing times*. American Mathematical Society, 2008.
- 15 Yun Long, Asaf Nachmias, and Yuval Peres. Mixing time power laws at criticality. In *Proceedings of the 48th Annual IEEE Symposium on Foundations of Computer Science, FOCS'07*, pages 205–214, Washington, DC, USA, 2007. IEEE Computer Society.
- 16 J. Maynard-Smith and E. Szathmáry. *The Major Transitions in Evolution*. New York: Oxford University Press, 1997.
- 17 Ruta Mehta, Ioannis Panageas, and Georgios Piliouras. Natural selection as an inhibitor of genetic diversity: Multiplicative weights updates algorithm and a conjecture of haploid genetics. In *Innovations in Theoretical Computer Science*, 2015.
- 18 M.A. Nowak. *Evolutionary Dynamics*. Harvard University Press, 2006.
- 19 Martin A. Nowak, Natalia L. Komarova, and Partha Niyogi. Evolution of universal grammar. *Science*, 2001.
- 20 Ioannis Panageas, Piyush Srivastava, and Nisheeth K. Vishnoi. Evolutionary dynamics in finite populations mix rapidly. *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms*, 2016.
- 21 Christos H. Papadimitriou and Nisheeth K. Vishnoi. On the computational complexity of limit cycles in dynamical systems. In *Proceedings of the 2016 Innovations in Theoretical Computer Science, Cambridge, MA, USA, January 14-16, 2016*, page 403, 2016.
- 22 Robin Pemantle. When are touchpoints limits for generalized Pólya urns? *Proceedings of the American Mathematical Society*, pages 235–243, 1991.

63:14 Mixing Time of Markov Chains, Dynamical Systems and Evolution

- 23 Georgios Piliouras, Carlos Nieto-Granda, Henrik I. Christensen, and Jeff S. Shamma. Persistent patterns: Multi-agent learning beyond equilibrium and utility. In *AAMAS*, pages 181–188, 2014.
- 24 Kushal Tripathi, Rajesh Balagam, Nisheeth K. Vishnoi, and Narendra M. Dixit. Stochastic simulations suggest that HIV-1 survives close to its error threshold. *PLoS Comput Biol*, 8(9):e1002684, 09 2012.
- 25 Nisheeth K. Vishnoi. The speed of evolution. In *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1590–1601, 2015.
- 26 Nicholas C Wormald. Differential equations for random processes and random graphs. *The annals of applied probability*, pages 1217–1235, 1995.