

Evolutionary Dynamics in Finite Populations Mix Rapidly

Ioannis Panageas*

Piyush Srivastava[†]

Nisheeth K. Vishnoi[‡]

Abstract

In this paper we prove that the mixing time of a broad class of evolutionary dynamics in finite, unstructured populations is roughly logarithmic in the size of the state space. An important special case of such a stochastic process is the Wright-Fisher model from evolutionary biology (with selection and mutation) on a population of size N over m genotypes. Our main result implies that the mixing time of this process is $O(\log N)$ for all mutation rates and fitness landscapes, and solves the main open problem from [4]. In particular, it significantly extends the main result in [18] who proved this for $m = 2$. Biologically, such models have been used to study the evolution of viral populations with applications to drug design strategies countering them. Here the time it takes for the population to reach a steady state is important both for the estimation of the steady-state structure of the population as well in the modeling of the treatment strength and duration. Our result, that such populations exhibit rapid mixing, makes both of these approaches sound.

Technically, we make a novel connection between Markov chains arising in evolutionary dynamics and dynamical systems on the probability simplex. This allows us to use the local and global stability properties of the fixed points of such dynamical systems to construct a contractive coupling in a fairly general setting. We expect that our mixing time result would be useful beyond the evolutionary biology setting, and the techniques used here would find applications in bounding the mixing times of Markov chains which have a natural underlying dynamical system.

1 Introduction

Evolutionary dynamical systems are central to the sciences due to their versatility in modeling a wide variety of biological, social and cultural phenomena, see [11].

Such dynamics are often used to capture the deterministic, *infinite population* setting, and are typically the first step in our understanding of seemingly complex processes. However, real populations are finite and often lend themselves to substantial stochastic effects (such as random drift) and it is often important to understand these effects as the population size varies. Hence, stochastic or *finite population* versions of evolutionary dynamical systems are appealed to in order to study such phenomena. While there are many ways to translate a deterministic dynamical system into a stochastic one, one thing remains common: the mathematical analysis becomes much harder as differential equations are easier to analyze and understand than stochastic processes.

Consider the example of the error-prone evolution of an unstructured, asexual haploid population in evolutionary biology; this is also the main motivation for this work. Each individual in the population could be one of m types. An individual of type i has a *fitness* (that translates to the ability to reproduce) which is specified by a positive integer a_i , and captured as a whole by a diagonal $m \times m$ matrix A whose (i, i) th entry is a_i . The reproduction is error-prone and this is captured by an $m \times m$ stochastic matrix Q whose (i, j) th entry captures the probability that the j th type will mutate to the i th type during reproduction.¹ The population is assumed to be infinite and its evolution deterministic. The population is assumed to be *unstructured*, meaning that only the type of each member of the population matters and, thus, it is sufficient to track the fraction of each type. One can then track the fraction of each type at step t of the evolution by a vector $\mathbf{x}^{(t)} \in \Delta_m$ (the probability simplex of dimension m) whose evolution is then governed by the difference equation $\mathbf{x}^{(t+1)} = \frac{QA\mathbf{x}^{(t)}}{\|QA\mathbf{x}^{(t)}\|_1}$. Of interest is the steady state² or the limiting distribution of this process and how it changes as one changes the evolutionary parameters Q and A . This model was proposed in the pioneering work of Eigen and co-authors [6, 7]. Importantly, this particular dynamical system has found use in modeling rapidly evolving viral

*Ioannis Panageas, Georgia Institute of Technology. Email: ioannis@gatech.edu. Supported by NSF grant CCF-1415496, CCF-1415498.

[†]Piyush Srivastava, California Institute of Technology. Email: piyushs@caltech.edu. Supported by NSF grant CCF-1319745.

[‡]Nisheeth K. Vishnoi, École Polytechnique Fédérale de Lausanne (EPFL). Email: nisheeth.vishnoi@epfl.ch.

¹We follow the convention that a matrix is stochastic if its *columns* sum up to 1.

²Note that there is a unique steady state, called the *quasispecies*, when $QA > 0$.

populations (such as HIV), which in turn has guided drug and vaccine design strategies. As a result, these dynamics are well-studied; see [4, 17] for an in depth discussion.

However, viral populations are generally not infinite and show stochastic effects: e.g., the effective population size of HIV-1 in infected individuals is approximately $10^3 - 10^6$ [2, 8], which is believed to be responsible for the strongly stochastic nature of its evolution. Several researchers have studied stochastic versions of Eigen’s deterministic evolution equations [1, 3, 10, 12–16, 19]. One such stochastic version, motivated by the Wright-Fisher model in population genetics, was studied by Dixit et al. [4]. Here, the population is again assumed to be unstructured and fixed to a size N . Thus, after normalization, the composition of the population is captured by a random point in Δ_m ; say $\mathbf{X}^{(t)}$ at time t . How does one generate $\mathbf{X}^{(t+1)}$ in this model when the parameters are still described by the matrices Q and A as in the infinite population case? To do this, in the replication (R) stage, one first replaces an individual of type i in the current population by a_i individuals of type i : the total number of individuals of type i in the intermediate population is therefore $a_i N X_i^{(t)}$. In the selection (S) stage, the population is culled back to size N by sampling N individuals from this intermediate population. In analogy with the Wright-Fisher model, we assume in this paper that the N individuals are sampled with replacement.³ Finally, since the evolution is error prone, in the mutation (M) stage, one then mutates each individual in this intermediate population independently and stochastically according to the matrix Q . The vector $\mathbf{X}^{(t+1)}$ then is the normalized frequency vector of the resulting population.

While this stochastic (RSM) model captures the effect of population size, it is useful only if the mixing time, or the time it takes for the population to reach (close to) its steady state, is much smaller than the size of the state space. For example, in simulations, samples from close to steady state are needed and this is only computationally feasible when the mixing time is small. Moreover, the efficacy of drug design strategies depends on the time it takes the population to evolve to steady state – the mixing time therefore models the minimum required duration of treatment. However, the number of states in the RSM process is roughly N^m (when m is small compared to N), and a mixing time that grows too fast as a function of the size of the state space can therefore be prohibitively large. For example, even for a

small constant $m = 40$ and a population of size 10,000, the number of states can grow to more than 2^{300} !

The importance of obtaining rigorous bounds for mixing time of RSM model was first pointed out in [4]. Rigorous mixing time results are far and few; they have either ignored mutation, assumed that the model is *neutral* (i.e., types have the same fitness), or moved to the diffusion limit which requires both mutation and selection pressure to be *weak*. Recently, the mixing time was shown to be roughly $\log N$ for the $m = 2$ case (for all Q and A) when all other parameters of the model are constants when compared to N , see [18]. As discussed in the technical overview, there are significant hurdles to extend this result to $m > 2$ and this problem has remained open since it was raised in [4].

1.1 Our contribution. In this paper we prove that the mixing time of the RSM model on an unstructured population of size N , parametrized by $m \times m$ matrices Q and A is $O(\log N)$ when all other parameters of the model are constants when compared to N .

Interestingly this result turns out to be a corollary of a rapid mixing result for a more general class of evolutionary dynamics in finite populations which should have applicability beyond evolutionary biology. Before we describe the result we present the setup: consider an infinite population whose evolution is described by a function $f : \Delta_m \mapsto \Delta_m$ where the population at time t , captured by the vector $\mathbf{x}^{(t)} \in \Delta_m$, evolves according to the rule $\mathbf{x}^{(t+1)} = f(\mathbf{x}^{(t)})$. Inspired by the Wright-Fisher model, we convert an infinite population dynamics to a finite one, say of size N , as follows: the fraction of each type in the population at time t is captured by a random vector $\mathbf{X}^{(t)} \in \Delta_m$ and the population at time $t + 1$ is obtained by sampling N times independently from the distribution $f(\mathbf{X}^{(t)})$. Our main result is that, under mild assumptions on f , the mixing time is bounded by $O(\log N)$ when all parameters involving f are assumed to be constants when compared to N .

Theorem 1.1 (Informal statement - see Theorem 3.6). *Let $f : \Delta_m \mapsto \Delta_m$ be an evolution function which satisfies certain mild conditions explained below, and consider the stochastic evolution guided by f on a population of size N . Then, the stochastic evolution converges to its stationary distribution in time $O(\log N)$.*

We first explain the conditions on f required by the theorem; see Definition 3.3 for a formal description. The key conditions are 1) that f has a “unique fixed point” τ which lies in the interior of Δ_m , 2) that the unique fixed point of f shows “contraction near the fixed point” and 3) that f shows “convergence to fixed point”. While it is clear what (1) and (3) mean, the second condition

³Culling via sampling without replacement was considered in [4], but the Wright-Fisher inspired sampling with replacement is the natural model for culling in the more general setting that we consider in this paper.

roughly means that there is a $k > 0$ such that if $J(\mathbf{x})$ denotes the Jacobian of f at \mathbf{x} , then for all \mathbf{x} which are near $\boldsymbol{\tau}$, $\|J^k(\mathbf{x})\|_1 < 1$. For technical reasons, we also need that f is twice differentiable in the interior of Δ_m . While we would have liked the contractive condition to hold for $k = 1$, unfortunately, there are simple examples where this is false. As we shall see later, the fact that we cannot take $k = 1$ is the main source of all the technical difficulties that arise in our work.

While we do not prove that these assumptions on f are necessary for the theorem to hold, there is some intuitive justification for them as we discuss below. Although the case of f having multiple fixed points is interesting in its own right, one might expect the stochastic population in this setting to exhibit slow mixing, since it can get stuck close to one of the fixed points. We therefore concentrate on the case where there is a unique fixed point and leave it as an open problem to study our theorem in the case of multiple fixed points formally. The assumption that this fixed point be located in the interior is motivated by the fact that this is the case in most applications, e.g. Eigen’s model has this property when $Q > 0$ and $A > 0$, entry-wise. The “contraction near the fixed point” condition ensures that the fixed point is indeed a steady state of the deterministic evolution in the sense of being asymptotically stable: once the evolution reaches close enough to the fixed point, it converges to the latter. Together with the “convergence to fixed point” condition (which, again, is satisfied by models such as Eigen’s model), this condition also ensures that the behavior of the deterministic system is free of exotic features, e.g. the presence of cycles, which may present a barrier to the fast mixing of the stochastic evolution. The smoothness condition is of a more technical nature, but models our expectation that any evolutionary dynamics should not be too susceptible to small changes in the population profile.

Several other remarks are in order. 1) We treat m , a bound on the derivatives of f , and the rate of contraction as constants and, hence, we do not optimize the dependence on these parameters in this paper. We leave it as an open problem to determine to what extent our results can be generalized when m is a growing function of N . It would be quite interesting to get a bound on the mixing time which depends polynomially on m . In this paper, our emphasis instead is on circumventing the obstacles that arose in previous attempts to prove a mixing time of close to $O(\log N)$ for general models: earlier works in this direction either put stringent conditions on the parameters of the model (e.g., Dixit et al. had to place very strong conditions on the matrices Q and A for their mixing time result to hold), or

were valid only when the number of genotypes was very small (e.g. Vishnoi [18] required the condition $m = 2$). 2) It is not obvious that the desired mixing time result for the RSM model is a corollary of the main theorem and we give a proof in Section 3.2. This result should be of independent interest in population genetics. 3) A natural next step is to study the evolution of *structured* populations. Roughly, this setting extends the RSM model by introducing an additional input parameter, a graph on N vertices. The graph provides structure to the population by locating each individual at a vertex, and the main difference from the RSM model is that at time $t + 1$, an individual determines its new vertex by sampling with replacement from among its neighbors in the graph at time t ; see [9] for more details. Here, it is no longer sufficient to just keep track of the fraction of each type. The RSM model can be seen as a special case when the underlying graph is the complete graph on N vertices, so that the locations of the individuals in the population are of no consequence. Our results do not seem to apply directly to this setting and it is a challenging open problem to prove bounds in the general graph setting.

Finally, we give a quick comparison of our techniques with that of [18] who proved a similar result for $m = 2$. (See the technical overview below for a more detailed discussion.) While [18] also used the underlying deterministic process, it was only to bring two copies of the Markov chain close enough (about a distance $\sqrt{1/N}$). Subsequently, his argument involved the construction of an ad-hoc coupling which contracts when $m = 2$. However, there appear to be serious hurdles when one tries to generalize this coupling for $m > 2$. We bypass these obstacles by again resorting to the properties of f listed above, however, in novel ways. We now move on to illustrating our key techniques.

2 Technical overview

We analyze the mixing time of our stochastic process by studying the time required for evolutions started at two arbitrary starting states $\mathbf{X}^{(0)}$ and $\mathbf{Y}^{(0)}$ to collide. More precisely, let \mathcal{C} be any Markovian coupling of two stochastic evolutions \mathbf{X} and \mathbf{Y} , both guided by a smooth contractive evolution f , which are started at $\mathbf{X}^{(0)}$ and $\mathbf{Y}^{(0)}$. Let T be the first (random) time such that $\mathbf{X}^{(T)} = \mathbf{Y}^{(T)}$. It is well known that if it can be shown that $\mathbb{P}[T > t] \leq 1/4$ for every pair of starting states $\mathbf{X}^{(0)}$ and $\mathbf{Y}^{(0)}$ then $t_{\text{mix}}(1/4) \leq t$. We show that such a bound on $\mathbb{P}[T > t]$ holds if we couple the chains using the optimal coupling of two multinomial distributions (see Section 3 for a definition of this coupling).

Our starting point is the observation that the optimal coupling and the definition of the evolutions

implies that for any time t ,

$$(2.1) \quad \mathbb{E} \left[\left\| \mathbf{X}^{(t+1)} - \mathbf{Y}^{(t+1)} \right\|_1 \mid \mathbf{X}^{(t)}, \mathbf{Y}^{(t)} \right] = \left\| f(\mathbf{X}^{(t)}) - f(\mathbf{Y}^{(t)}) \right\|_1.$$

Now, if f were *globally* contractive, so that the right hand side of eq. (2.1) was always bounded above by $\rho' \left\| \mathbf{X}^{(t)} - \mathbf{Y}^{(t)} \right\|_1$ for some constant $\rho' < 1$, then we would get that the expected distance between the two copies of the chains contracts at a constant rate. Since the minimum possible positive ℓ_1 distance between two copies of the chain is $1/N$, this would have implied an $O(\log N)$ mixing time using standard arguments. However, such a global assumption on f , which is equivalent to requiring that the Jacobian J of f satisfies $\|J(\mathbf{x})\|_1 < 1$ for all $\mathbf{x} \in \Delta_m$, is far too strong. In particular, it is not satisfied by standard systems such as Eigen’s dynamics discussed above.

Nevertheless, these dynamics do satisfy a more local version of the above condition. That is, they have a unique fixed point $\boldsymbol{\tau}$ to which they converge quickly, and *in the vicinity* of this fixed point, *some* form of contraction holds. These conditions motivate the “unique fixed point”, “contraction near the fixed point”, and the “convergence to fixed point” conditions in our definition of a smooth contractive evolution (Definition 3.3). However, crucially, the “contraction near the fixed point” condition, inspired from the definition of “asymptotically stable” fixed points in dynamical systems, is weaker than the stepwise contraction condition described in the last paragraph, even in the vicinity of the fixed point. As we shall see shortly, this weakening is essential for generalizing the earlier results of [18] to the $m > 2$ case, but comes at the cost of making the analysis more challenging.

However, we first describe how the “convergence to fixed point” condition is used to argue that the chains come close to the fixed point in $O(\log N)$ time. This step of our argument is the only one technically quite similar to the development in [18]; our later arguments need to diverge widely from that paper. Although this step is essentially an iterated application of appropriate concentration results along with the fact that the “convergence to fixed point” condition implies that the deterministic evolution f comes close to the fixed point $\boldsymbol{\tau}$ at an exponential rate, complications arise because f can amplify the effect of the random perturbations that arise at each step. In particular, if $L > 1$ is the maximum of $\|J(\mathbf{x})\|_1$ over Δ_m , then after ℓ steps, a random perturbation can become amplified by a factor of L^ℓ . As such, if ℓ is taken to be too large, these accumulated errors can swamp the progress

made due to the fast convergence of the deterministic evolution to the fixed point. These considerations imply that the argument can only be used for $\ell = \ell_0 \log N$ steps for some small constant ℓ_0 , and hence we are only able to get the chains within $\Theta(N^{-\gamma})$ distance of the fixed point, where $\gamma < 1/3$ is a small constant. In particular, the argument cannot be carried out all the way down to distance $O(1/N)$, which, if possible, would have been sufficient to show that the coupling time is small with high probability. Nevertheless, it does allow us to argue that both copies of the chain enter an $O(N^{-\gamma})$ neighborhood of the fixed point in $O(\log N)$ steps.

At this point, [18] showed that in the $m = 2$ case, one could take advantage of the contractive behavior near the fixed point to construct a coupling obeying eq. (2.1) in which the right hand side was indeed contractive: in essence, this amounted to a proof that $\|J\|_1 < 1$ was indeed satisfied in the small $O(N^{-\gamma})$ neighborhood reached at the end of the last step. This allowed [18] to complete the proof using standard arguments, after some technicalities about ensuring that the chains remained for a sufficiently long time in the neighborhood of the fixed point had been taken care of.

The situation however changes completely in the $m > 2$ case. It is no longer possible to argue in general that $\|J(\mathbf{x})\|_1 < 1$ when \mathbf{x} is in the vicinity of the fixed point, even when there is fast convergence to the fixed point. Instead, we have to work with a weaker condition (the “contraction to the fixed point” condition alluded to earlier) which only implies that there is a positive integer k , possibly larger than 1, such that in some vicinity of the fixed point, $\|J^k\|_1 < 1$. In the setting used by [18], k could be taken to be 1, and hence it could be argued via eq. (2.1) that the distance between the two coupled copies of the chains contracts in each step. This argument however does not go through when only a k th power of J is guaranteed to be contractive while J itself could have $1 \rightarrow 1$ norm larger than 1. This inability to argue stepwise contraction is the major technical obstacle in our work when compared to the work of [18], and the source of all the new difficulties that arise in this more general setting.

As a first step toward getting around the difficulty of not having stepwise contraction, we prove Theorem 4.1, which shows that the eventual contraction after k steps can be used to ensure that the distance between two evolutions $\mathbf{x}^{(t)}$ and $\mathbf{y}^{(t)}$ close to the fixed point contracts by a factor $\rho^k < 1$ over an *epoch* of k steps (where k is as described in the last paragraph), even when the evolutions undergo arbitrary perturbations $\mathbf{u}^{(t)}$ and $\mathbf{v}^{(t)}$ at each step, *provided* that the difference $\mathbf{u}^{(t)} - \mathbf{v}^{(t)}$ between the two perturbations is small compared to the difference $\mathbf{x}^{(t-1)} - \mathbf{y}^{(t-1)}$ between the evolutions at the

previous step. The last condition actually asks for a relative notion of smallness, i.e., it requires that

$$(2.2) \quad \left\| \boldsymbol{\xi}^{(t)} \right\|_1 := \left\| \mathbf{u}^{(t)} - \mathbf{v}^{(t)} \right\|_1 \leq \delta \left\| \mathbf{x}^{(t-1)} - \mathbf{y}^{(t-1)} \right\|_1,$$

where δ is a constant specified in the theorem. Note that the theorem is a statement about deterministic evolutions against possibly adversarial perturbations, and does *not* require $\mathbf{u}^{(t)}$ and $\mathbf{v}^{(t)}$ to be stochastic, but only that they follow the required conditions on the difference of the norm (in addition to the implied condition that the evolution $\mathbf{x}^{(t)}$ and $\mathbf{y}^{(t)}$ remain close to the fixed point during the epoch).

Thus, in order to use Theorem 4.1 for showing that the distance $\left\| \mathbf{X}^{(t)} - \mathbf{Y}^{(t)} \right\|_1$ between the two coupled chains contracts after every k iterations of eq. (2.1), we need to argue that the required condition on the perturbations in eq. (2.2) holds with high probability over a given epoch during the coupled stochastic evolution of $\mathbf{X}^{(t)}$ and $\mathbf{Y}^{(t)}$. (In fact, we also need to argue that the two chains individually remain close to the fixed point, but this is easier to handle and we ignore this technicality in this proof overview; the details appear in the formal proofs in Section 5.)

However, at this point, a complication arises from the fact that Theorem 4.1 requires the difference $\boldsymbol{\xi}^{(t)}$ between the perturbations at time t to be bounded *relative* to the difference $\left\| \mathbf{X}^{(t-1)} - \mathbf{Y}^{(t-1)} \right\|_1$ at time $t-1$. In other words, the upper bounds required on the $\boldsymbol{\xi}^{(t)}$ become more stringent as the two chains come closer to each other. This fact creates a trade-off between the probability with which the condition in eq. (2.2) can be enforced in an epoch, and the required *lower* bound on the distance between the chains required during the epoch so as to ensure that probability (this trade-off is technically based on Lemma 3.5). To take a couple of concrete examples, when $\left\| \mathbf{X}^{(t)} - \mathbf{Y}^{(t)} \right\|_1$ is $\Omega(\log N/N)$ in an epoch, we can ensure that eq. (2.2) remains valid with probability at least $1 - N^{-\Theta(1)}$ (see the discussion following Theorem 4.4), so that with high probability $\Omega(\log N)$ consecutive epochs admit a contraction allowing the distance between the chains to come down from $\Theta(N^{-\gamma})$ at the end of the first step to $\Theta(\log N/N)$ at the end of this set of epochs.

Ideally, we would have liked to continue this argument till the distance between the chains is $\Theta(1/N)$ and (due to the properties of the optimal coupling) they have a constant probability of colliding in a single step. However, due to the trade-off referred to earlier, when we know only that $\left\| \mathbf{X}^{(t)} - \mathbf{Y}^{(t)} \right\|_1$ is $\Omega(1/N)$ during the epoch, we can only guarantee the condition of eq. (2.2) with probability $\Theta(1)$ (see the discussion following the proof of Theorem 4.4). Thus, we

cannot claim directly that once the distance between the chains is $O(\log N/N)$, the next $\Omega(\log \log N)$ epochs will exhibit contraction in distance leading the chain to come as close as $O(1/N)$ with a high enough positive probability. To get around this difficulty, we consider $O(\log \log N)$ epochs with successively weaker guaranteed upper bounds on $\left\| \mathbf{X}^{(t)} - \mathbf{Y}^{(t)} \right\|_1$. Although the weaker lower bounds on the distances lead in turn to weaker concentration results when Theorem 4.4 is applied, we show that this trade-off is such that we can choose these progressively decreasing guarantees so that after this set of epochs, the distance between the chains is $O(1/N)$ with probability that it is small but at least a *constant*. Since the previous steps, i.e., those involving making both chains come within distance $O(N^{-\gamma})$ of the fixed point (for some small constant $\gamma < 1$), and then making sure that the distance between them drops to $O(\log N/N)$, take time $O(\log N)$ with probability $1 - o(1)$, we can conclude that under the optimal coupling, the collision or coupling time T satisfies

$$(2.3) \quad \mathbb{P}[T > O(\log N)] \leq 1 - q,$$

for some small enough constant q , irrespective of the starting states $\mathbf{X}^{(0)}$ and $\mathbf{Y}^{(0)}$ (note that here we are also using the fact that once the chains are within distance $O(1/N)$, the optimal coupling has a constant probability of causing a collision in a single step). The lack of dependence on the starting states allows us to iterate eq. (2.3) for $\Theta(1)$ consecutive “blocks” of time $O(\log N)$ each to get

$$\mathbb{P}[T > O(\log N)] \leq \frac{1}{4},$$

which gives us the claimed mixing time.

3 Preliminaries and formal statement of results

In preparation for formally stating our main result, we now discuss some preliminary notation, definitions and technical tools that will be used later on. We then formally state our main theorem in Sections 3.1 and 3.2.

Notation. We denote the probability simplex on a set of size m as Δ_m . Vectors in \mathbb{R}^m , and probability distributions in Δ_m are both denoted in boldface, and x_j denotes the j th co-ordinate of a given vector \mathbf{x} . Time indices are denoted by superscripts. Thus, a time indexed scalar s at time t is denoted as $s^{(t)}$, while a time indexed vector \mathbf{x} at time t is denoted as $\mathbf{x}^{(t)}$. The letters \mathbf{X} and \mathbf{Y} (with time superscripts and co-ordinate subscripts, as appropriate) will be used to denote random vectors. Scalar random vectors and matrices are denoted by other capital letters. Boldface $\mathbf{1}$ denotes a vector all whose entries are 1.

Operators and norms. For any square matrix M , we denote by $\|M\|_1$ its $1 \rightarrow 1$ norm defined as

$\max_{\|v\|_1=1} \|Mv\|_1$, by $\|M\|_2$ its *operator norm* defined as $\max_{\|v\|_2=1} \|Mv\|_2$, and by $\text{sp}(M)$ its *spectral radius* defined as the maximum of the absolute values of its eigenvalues. 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 3.1 (Gelfand’s formula, specialized to the $1 \rightarrow 1$ norm). *For any square matrix M , we have*

$$\text{sp}(M) = \lim_{l \rightarrow \infty} \|M^l\|_1^{1/l}.$$

Derivatives. Let $g : \mathbb{R}^m \rightarrow \mathbb{R}^m$ be any differentiable function, whose co-ordinates are denoted as $g_i : \mathbb{R}^m \rightarrow \mathbb{R}$. The *Jacobian* $J(\mathbf{x})$ of g at a point $\mathbf{x} \in \mathbb{R}^m$ is the $m \times m$ matrix whose (i, j) entry is $\frac{\partial g_i(\mathbf{x})}{\partial x_j}$. The *Hessian* of a twice differentiable function $h : \mathbb{R}^m \rightarrow \mathbb{R}$ at a given point \mathbf{x} is the $m \times m$ symmetric matrix whose (i, j) entry is $\frac{\partial^2 h(\mathbf{x})}{\partial x_i \partial x_j}$. We use the following special case of Taylor’s theorem.

Theorem 3.2 (Taylor’s theorem, truncated). *Let $g : \mathbb{R}^m \rightarrow \mathbb{R}^m$ be a twice differentiable function, and let $J(\mathbf{z})$ denote the Jacobian of g at \mathbf{z} . Let $\mathbf{x}, \mathbf{y} \in \mathbb{R}^m$ be two points, and suppose there exists a positive constant B such that at every point on the line segment joining \mathbf{x} to \mathbf{y} , the Hessians of each of the m co-ordinates g_i of g have operator norm at most $2B$. Then, there exists a $\mathbf{v} \in \mathbb{R}^m$ such that*

$$g(\mathbf{x}) = g(\mathbf{y}) + J(\mathbf{y})(\mathbf{x} - \mathbf{y}) + \mathbf{v},$$

and $|v_i| \leq B \|\mathbf{x} - \mathbf{y}\|_2^2$ for each $i \in [m]$.

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 co-ordinate is equal to \mathbf{p} and that on the second co-ordinate is equal to \mathbf{q} . A simple, if trivial, example of a coupling is the joint distribution obtained by sampling the two co-ordinates independently, one from \mathbf{p} and the other from \mathbf{q} .

Couplings allow a very useful dual characterization of the total variation distance, as stated in the following well known lemma.

Lemma 3.3 (Coupling lemma). *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} .

Moreover, the coupling in the lemma can be explicitly described. We use this coupling extensively in our arguments, and hence we record some of its useful properties here.

Definition 3.1 (Optimal coupling). Let $\mathbf{p}, \mathbf{q} \in \Delta_m$ be two probability distributions on m objects. For each $i \in [m]$, let $s_i := \min(p_i, q_i)$, and $s := \sum_{i=1}^m s_i$. Sample U, V, W independently at random as follows:

$$\begin{aligned} \mathbb{P}[U = i] &= \frac{s_i}{s}, \quad \mathbb{P}[V = i] = \frac{p_i - s_i}{1 - s}, \\ \text{and } \mathbb{P}[W = i] &= \frac{q_i - s_i}{1 - s}, \text{ for all } i \in [m]. \end{aligned}$$

We then sample (independent of U, V, W) a Bernoulli random variable H with mean s . The sample (A, B) given by the coupling is (U, U) if $H = 1$ and (V, W) otherwise.

It is easy to verify that $A \sim \mathbf{p}$, $B \sim \mathbf{q}$ and $\mathbb{P}[A = B] = s = 1 - \|\mathbf{p} - \mathbf{q}\|_{\text{TV}}$. Another easily verified but important property is that for any $i \in [m]$

$$\mathbb{P}[A = i, B \neq i] = \begin{cases} 0 & \text{if } p_i < q_i, \\ p_i - q_i & \text{if } p_i \geq q_i. \end{cases}$$

Definition 3.2 (Mixing time). Let \mathcal{M} be an ergodic Markov chain on a finite state space Ω with stationary distribution $\boldsymbol{\pi}$. Then, the *mixing time* $t_{\text{mix}}(\varepsilon)$ is defined as the smallest time such that for any starting state $S^{(0)}$, the distribution of the state $S^{(t)}$ at time t is within total variation distance ε of $\boldsymbol{\pi}$. The term mixing time is also used for $t_{\text{mix}}(\varepsilon)$ for a *fixed* values of $\varepsilon < 1/2$. For concreteness, in the rest of this paper, we use t_{mix} to refer to $t_{\text{mix}}(1/e)$ (though any other constant smaller than $1/2$ could be chosen as well in place of $1/e$ without changing any of the claims).

A standard technique for obtaining upper bounds on mixing times is to use the Coupling Lemma above. Suppose $S_1^{(t)}$ and $S_2^{(t)}$ are two evolutions of an ergodic chain M such that their evolutions are coupled according to some coupling \mathcal{C} . Let T be the stopping time such that $S_1^{(T)} = S_2^{(T)}$. Then, if it can be shown that $\mathbb{P}[T > t] \leq 1/e$ for every pair of starting states $(S_1^{(0)}, S_2^{(0)})$, then it follows that $t_{\text{mix}} := t_{\text{mix}}(1/e) \leq t$.

Concentration. We now discuss some concentration results that are used extensively in our later arguments. We begin with some standard Chernoff-Hoeffding type bounds.

Theorem 3.4 (Chernoff-Hoeffding bounds [5]). *Let Z_1, Z_2, \dots, Z_N be i.i.d Bernoulli random variables with mean μ . We then have*

1. When $0 < \delta \leq 1$,

$$\mathbb{P} \left[\left| \frac{1}{N} \sum_{i=1}^N Z_i - \mu \right| > \mu\delta \right] \leq 2 \exp(-N\mu\delta^2/3).$$

2. When $\delta \geq 1$,

$$\mathbb{P} \left[\left| \frac{1}{N} \sum_{i=1}^N Z_i - \mu \right| > \mu\delta \right] \leq \exp(-N\mu\delta/3).$$

An important tool in our later development is the following lemma, which bounds additional “discrepancies” that can arise when one samples from two distribution \mathbf{p} and \mathbf{q} using an optimal coupling. The important feature for us is the fact that the additional discrepancy (denoted as \mathbf{e} in the lemma) is bounded as a *fraction* of the “initial discrepancy” $\|\mathbf{p} - \mathbf{q}\|_1$. However, such relative bounds on the discrepancy are less likely to hold when the initial discrepancy itself is very small, and hence, there is a trade-off between the lower bound that needs to be imposed on the initial discrepancy $\|\mathbf{p} - \mathbf{q}\|_1$, and the desired probability with which the claimed relative bound on the additional discrepancy \mathbf{e} is to hold. The lemma makes this delicate trade-off precise.

Lemma 3.5. *Let \mathbf{p} and \mathbf{q} be probability distributions on a universe of size m , so that $\mathbf{p}, \mathbf{q} \in \Delta_m$. Consider an optimal coupling of the two distributions, and let \mathbf{x} and \mathbf{y} be random frequency vectors with m coordinates (normalized to sum to 1) obtained by taking N independent samples from the coupled distributions, so that $\mathbb{E}[\mathbf{x}] = \mathbf{p}$ and $\mathbb{E}[\mathbf{y}] = \mathbf{q}$. Define the random error vector \mathbf{e} as*

$$\mathbf{e} := (\mathbf{x} - \mathbf{y}) - (\mathbf{p} - \mathbf{q}).$$

Suppose $c > 1$ and t (possibly dependent upon N) are such that $\|\mathbf{p} - \mathbf{q}\|_1 \geq \frac{ctm}{N}$. We then have

$$\|\mathbf{e}\|_1 \leq \left(\frac{2}{\sqrt{c}} \right) \|\mathbf{p} - \mathbf{q}\|_1$$

with probability at least $1 - 2m \exp(-t/3)$.

Proof. The properties of the optimal coupling of the distributions \mathbf{p} and \mathbf{q} imply that since the N coupled samples are taken independently,

1. $|x_i - y_i| = \frac{1}{N} \sum_{j=1}^N R_j$, where R_j are i.i.d. Bernoulli random variables with mean $|p_i - q_i|$, and
2. $x_i - y_i$ has the same sign as $p_i - q_i$.

The second fact implies that $|e_i| = |x_i - y_i| - |(p_i - q_i)|$. By applying the concentration bounds from Theorem 3.4 to the first fact, we then get (for any arbitrary $i \in [m]$)

$$\begin{aligned} \mathbb{P} \left[|e_i| > \sqrt{\frac{t}{N|p_i - q_i|}} \cdot |p_i - q_i| \right] &\leq 2 \exp(-t/3), \\ &\text{if } \frac{t}{N|p_i - q_i|} \leq 1, \text{ and} \\ \mathbb{P} \left[|e_i| > \frac{t}{N|p_i - q_i|} \cdot |p_i - q_i| \right] &\leq \exp(-t/3), \\ &\text{if } \frac{t}{N|p_i - q_i|} > 1. \end{aligned}$$

One of the two bounds applies to every $i \in [m]$ (except those i for which $|p_i - q_i| = 0$, but in those cases, we have $|e_i| = 0$, so the bounds below will apply nonetheless). Thus, taking a union bound over all the indices, we see that with probability at least $1 - 2m \exp(-t/3)$, we have

$$\begin{aligned} \|\mathbf{e}\|_1 = \sum_{i=1}^m |e_i| &\leq \sqrt{\frac{t}{N}} \sum_{i=1}^m \sqrt{|p_i - q_i|} + \frac{tm}{N} \\ (3.4) \end{aligned}$$

$$\begin{aligned} &\leq \sqrt{\frac{tm}{N}} \sqrt{\|\mathbf{p} - \mathbf{q}\|_1} + \frac{tm}{M} \\ (3.5) \end{aligned}$$

Here, eq. (3.4) uses the Cauchy-Schwarz inequality to bound the first term while eq. (3.5) uses the hypothesis in the lemma that $\frac{ctm}{N} \leq \|\mathbf{p} - \mathbf{q}\|_1$. The claim of the lemma follows since $c > 1$. \square

3.1 Main theorem. We are now ready to state our main theorem. We begin by formally defining the conditions on the evolution function required by the theorem.

Definition 3.3 (Smooth contractive evolution). A function $f : \Delta_m \rightarrow \Delta_m$ is said to be a (L, B, ρ) smooth contractive evolution if it has the following properties:

Smoothness f is twice differentiable in the interior of Δ_m . Further, the Jacobian J of f satisfies $\|J(x)\|_1 \leq L$ for every \mathbf{x} in the interior of Δ_m , and the operator norms of the Hessians of its coordinates are uniformly bounded above by $2B$ at all points in the interior of Δ_m .

Unique fixed point f has a unique fixed point $\boldsymbol{\tau}$ in Δ_m which lies in the interior of Δ_m .

Contraction near the fixed point At the fixed point $\boldsymbol{\tau}$, the Jacobian $J(\boldsymbol{\tau})$ of f satisfies

$$\text{sp}(J(\boldsymbol{\tau})) < \rho < 1.$$

Convergence to fixed point For every $\epsilon > 0$, there exists an ℓ such that for any $\mathbf{x} \in \Delta_m$,

$$\|f^\ell(\mathbf{x}) - \boldsymbol{\tau}\|_1 < \epsilon.$$

Remark 3.1. Note that the last condition implies that $\|f^t(\mathbf{x}) - \boldsymbol{\tau}\|_1 = O(\rho^t)$ in the light of the previous condition and the smoothness condition (see Lemma A.1). Also, it is easy to see that the last two conditions imply the uniqueness of the fixed point, i.e., the second condition. However, the last condition on global convergence does not by itself imply the third condition on contraction near the fixed point. Consider, e.g., $g : [-1, 1] \rightarrow [-1, 1]$ defined as $g(x) = x - x^3$. The unique fixed point of g in its domain is 0, and we have $g'(0) = 1$, so that the third condition is not satisfied. On the other hand, the last condition is satisfied, since for $x \in [-1, 1]$ satisfying $|x| \geq \epsilon$, we have $|g(x)| \leq |x|(1 - \epsilon^2)$. In order to construct a function $f : [0, 1] \rightarrow [0, 1]$ with the same properties, we note that the range of g is $[-x_0, x_0]$ where $x_0 = 2/(3\sqrt{3})$, and consider $f : [0, 1] \rightarrow [0, 1]$ defined as $f(x) = x_0 + g(x - x_0)$. Then, the unique fixed point of f in $[0, 1]$ is x_0 , $f'(x_0) = g'(0) = 1$, the range of f is contained in $[0, 2x_0] \subseteq [0, 1]$, and f satisfies the fourth condition in the definition but does not satisfy the third condition.

Given an f which is a smooth contractive evolution, and a population parameter N , we can define a 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 frequency vector over $[m]$. Then

$$\mathbf{x}^{(t+1)} := \frac{1}{N} \mathbf{z}^{(t)}.$$

Theorem 3.6 (Main Theorem). *Let f be a (L, B, ρ) smooth contractive evolution, and consider the stochastic evolution guided by f on a population of size N . Then, there exists c_0 and N_0 depending upon L, B and ρ and f such that for any $N > N_0$, the mixing time $t_{\text{mix}}(1/e)$ of the stochastic evolution is at most $c_0((\log N))$.*

3.2 The RSM model as a special case. We now show that the Eigen or RSM model discussed in the introduction is a special case of the abstract model defined in the last section, and hence satisfies the mixing time bound in Theorem 3.6. Our first step is to show that the RSM model can be seen as a stochastic evolution guided by the function f defined by $f(\mathbf{p}) = \frac{(QAp)_t}{\|QAp\|_1}$, where Q and A are matrices with positive entries, with Q stochastic (i.e., columns summing up to 1), as described

in the introduction. We will then show that this f is a smooth contractive evolution, which implies that Theorem 3.6 applies to the RSM process.

We begin by recalling the definition of the RSM process. Given a starting population of size N on m types represented by a $1/N$ -integral probability vector $\mathbf{p} = (p_1, p_2, \dots, p_m)$, the RSM process produces the population at the next step by independently sampling N times from the following process:

1. Sample a type T from the probability distribution $\frac{A\mathbf{p}}{\|A\mathbf{p}\|_1}$.
2. Mutate T to the result type S with probability Q_{ST} .

We now show that sampling from this process is exactly the same as sampling from the multinomial distribution $f(\mathbf{p}) = \frac{QAp}{\|QAp\|_1}$. To do this, we only need to establish the following claim:

Claim 3.7. *For any type $t \in [m]$, $\mathbb{P}[S = t] = \frac{\sum_j Q_{tj} A_{jj} p_j}{\sum_j A_{jj} p_j} = \frac{(QAp)_t}{\|QAp\|_1}$.*

Proof. We first note that $\|QAp\|_1 = \sum_{ij} Q_{ij} A_{jj} p_j = (\sum_i Q_{ij}) \cdot (\sum_j A_{jj} p_j) = \sum_j A_{jj} p_j = \|A\mathbf{p}\|_1$, where in the last equality we used the fact that the columns of Q sum up to 1. Now, we have

$$\begin{aligned} \mathbb{P}[S = t] &:= \sum_i^m Q_{ti} \cdot \frac{(A\mathbf{p})_i}{\|A\mathbf{p}\|_1} = \frac{\sum_{i=1}^m Q_{ti} A_{ii} p_i}{\sum_j A_{jj} p_j} \\ &= \frac{(QAp)_t}{\|QAp\|_1}. \quad \square \end{aligned}$$

From Claim 3.7, we see that producing N independent samples from the process described above (which corresponds exactly to the RSM model) produces the same distribution as producing N independent samples from the distribution $\frac{(QAp)}{\|QAp\|_1}$. Thus, the RSM process is a stochastic evolution guided by $f(\mathbf{x}) := \frac{QAx}{\|QAx\|_1}$. We now proceed to verify that this f is a smooth contractive evolution. We first note that the ‘‘smoothness’’ condition is directly implied by the definition of f . For the ‘‘uniqueness of fixed point’’ condition, we observe that every fixed point of $\frac{QAx}{\|QAx\|_1}$ in the simplex Δ_m must be an eigenvector of QA . Since QA is a matrix with positive entries, the Perron-Frobenius theorem implies that it has a unique positive eigenvector \mathbf{v} (for which we can assume without loss of generality that $\|\mathbf{v}\|_1 = 1$) with a positive eigenvalue λ_1 . Therefore $f(\mathbf{x})$ has a unique fixed point $\boldsymbol{\tau} = \mathbf{v}$ in the simplex Δ_m which is in its interior. The Perron-Frobenius theorem also implies that for every $\mathbf{x} \in \Delta_m$, $\lim_{t \rightarrow \infty} (QA)^t \mathbf{x} / \lambda_1^t \rightarrow \mathbf{v}$. In fact, this convergence can be made uniform over Δ_m

(meaning that given an $\epsilon > 0$ we can choose t_0 such that for all $t > t_0$, $\|(QA)^t \mathbf{x} / \lambda_1^t - \mathbf{v}\|_1 < \epsilon$ for all $\mathbf{x} \in \Delta_m$) since each point $\mathbf{x} \in \Delta_m$ is a convex combination of the extreme points of Δ_m and the left hand side is a linear function of \mathbf{x} . From this uniform convergence, it then follows easily that $\lim_{t \rightarrow \infty} f^t(\mathbf{x}) = \mathbf{v}$, and that the convergence in this limit is also uniform. The “convergence to fixed point” condition follows directly from this observation.

Finally, we need to establish that the spectral radius of the Jacobian $J := J(\mathbf{v})$ of f at its fixed point is less than 1. A simple computation shows that the Jacobian at \mathbf{v} is $J = \frac{1}{\lambda_1}(I - V)QA$ where V is the matrix each of whose columns is the vector \mathbf{v} . Since QA has positive entries, we know from the Perron-Frobenius theorem that λ_1 as defined above is real, positive, and strictly larger in magnitude than any other eigenvalue of QA . Let $\lambda_2, \lambda_3, \dots, \lambda_m$ be the other, possibly complex, eigenvalues arranged in decreasing order of magnitude (so that $\lambda_1 > |\lambda_2|$). We now establish the following claim from which it immediately follows that $\text{sp}(J) = \frac{|\lambda_2|}{\lambda_1} < 1$ as required.

Claim 3.8. *The eigenvalues of $M := (I - V)QA$ are $\lambda_2, \lambda_3, \dots, \lambda_m, 0$.*

Proof. Let D be the Jordan canonical form of QA , so that $D = U^{-1}QA U$ for some invertible matrix U . Note that D is an upper triangular matrix with $\lambda_1, \lambda_2, \dots, \lambda_m$ on the diagonal. Further, the Perron-Frobenius theorem applied to QA implies that λ_1 is an eigenvalue of both algebraic and geometric multiplicity 1, so that we can assume that the topmost Jordan block in D is of size 1 and is equal to λ_1 . Further, we can assume the corresponding first column of U is equal to the corresponding positive eigenvector \mathbf{v} satisfying $\|\mathbf{v}\|_1 = 1$. It therefore follows that $U^{-1}V = U^{-1}\mathbf{v}\mathbf{1}^T$ is the matrix $\mathbf{e}_1\mathbf{1}^T$, where \mathbf{e}_1 is the first standard basis vector.

Now, since U is invertible, M has the same eigenvalues as $U^{-1}MU = (U^{-1} - U^{-1}V)QA U = (I - \mathbf{e}_1\mathbf{1}^T U)D$, where in the last line we use $UD = QA U$. Now, note that all rows except the first of the matrix $\mathbf{e}_1\mathbf{1}^T U$ are zero, and its (1, 1) entry is 1 since the first column of U is \mathbf{v} , which in turn is chosen so that $\mathbf{1}^T \mathbf{v} = 1$. Thus, we get that $(I - \mathbf{e}_1\mathbf{1}^T U)D$ is an upper triangular matrix with the same diagonal entries as D except that its (1, 1) entry is 0. Since the (1, 1) entry of D was λ_1 while its other diagonal entries were $\lambda_2, \lambda_3, \dots, \lambda_m$, it follows that the eigenvalues of $(I - \mathbf{e}_1\mathbf{1}^T U)D$ (and hence those of M) are $\lambda_2, \lambda_3, \dots, \lambda_m, 0$, as claimed. \square

We thus see that the RSM process satisfies the condition of being guided by a smooth contractive evolution and hence has the mixing time implied by Theorem 3.6.

4 Perturbed evolution near the fixed point

As discussed in Section 2, the crux of the proof of our main theorem is analyzing how the distance between two copies of a stochastic evolution guided by a smooth contractive evolution evolves in the presence of small perturbations at every step. In this section, we present our main tool, Theorem 4.1, to study this phenomenon. We then describe how the theorem, which itself is presented in a completely deterministic setting, applies to stochastic evolutions.

Fix any (L, B, ρ) -smooth contractive evolution f on Δ_m , with fixed point τ . As we noted in Section 2, since the Jacobian of f does not necessarily have operator (or $1 \rightarrow 1$) norm less than 1, we cannot argue that the effect of perturbations shrinks in every step. Instead, we need to argue that the condition on the spectral radius of the Jacobian of f at its fixed point implies that there is eventual contraction of distance between the two evolutions, even though this distance might increase in any given step. Indeed, the fact that the spectral radius $\text{sp}(J)$ of the Jacobian at the fixed point τ of f is less than $\rho < 1$ implies that a suitable iterate of f has a Jacobian with operator (and $1 \rightarrow 1$) norm less than 1 at τ . This is because Gelfand’s formula (Theorem 3.1) implies that for all large enough positive integers k' ,

$$\|J^{k'}(\tau)\|_1 < \rho^{k'}.$$

We now use the above condition to argue that after k steps in the vicinity of the fixed point, there is indeed a contraction of the distance between two evolutions guided by f , even in the presence of adversarial perturbations, as long as those perturbations are small. The precise statement is given below; the vectors $\xi^{(i)}$ in the theorem model the small perturbations.

Theorem 4.1 (Perturbed evolution). *Let f be a (L, B, ρ) -smooth contractive evolution, and let τ be its fixed point. For all positive integers $k > k_0$ (where k_0 is a constant that depends upon f) there exist $\epsilon, \delta \in (0, 1]$ depending upon f and k for which the following is true. Let $(\mathbf{x}^{(i)})_{i=0}^k, (\mathbf{y}^{(i)})_{i=0}^k$, and $(\xi^{(i)})_{i=1}^k$ be sequences of vectors with $\mathbf{x}^{(i)}, \mathbf{y}^{(i)} \in \Delta_m$ and $\xi^{(i)}$ orthogonal to $\mathbf{1}$, which satisfy the following conditions:*

1. (**Definition**). *For $1 \leq i \leq k$, there exist vectors $\mathbf{u}^{(i)}$ and $\mathbf{v}^{(i)}$ such that*

$$\begin{aligned} \mathbf{x}^{(i)} &= f(\mathbf{x}^{(i-1)}) + \mathbf{u}^{(i)}, & \mathbf{y}^{(i)} &= f(\mathbf{y}^{(i-1)}) + \mathbf{v}^{(i)}, \\ & & \text{and } \xi^{(i)} &= \mathbf{u}^{(i)} - \mathbf{v}^{(i)}. \end{aligned}$$

2. (**Closeness to fixed point**). *For $0 \leq i \leq k$, $\|\mathbf{x}^{(i)} - \tau\|_1 \leq \epsilon$ and $\|\mathbf{y}^{(i)} - \tau\|_1 \leq \epsilon$.*

3. (**Small perturbations**). For $1 \leq i \leq k$, $\|\xi^{(i)}\|_1 \leq \delta \|\mathbf{x}^{(i-1)} - \mathbf{y}^{(i-1)}\|_1$.

Then, we have

$$\|\mathbf{x}^{(k)} - \mathbf{y}^{(k)}\|_1 \leq \rho^k \|\mathbf{x}^{(0)} - \mathbf{y}^{(0)}\|_1.$$

In the theorem, the vectors $\mathbf{x}^{(i)}$ and $\mathbf{y}^{(i)}$ model the two chains, while the vectors $\mathbf{u}^{(i)}$ and $\mathbf{v}^{(i)}$ model the individual perturbations from the evolution dictated by f . The theorem says that if the perturbations $\xi^{(i)}$ to the distance are not too large, then the distance between the two chains indeed contracts after every k steps.

Proof. As observed above, we can use Gelfand's formula to conclude that there exists a positive integer k_0 (depending upon f) such that we have $\|J(\tau)^k\|_1 < \rho^k$ for all $k > k_0$. This k_0 will be the sought k_0 in the theorem, and we fix some appropriate $k > k_0$ for the rest of the paper.

Since f is twice differentiable, J is continuous on Δ_m . This implies that the function on Δ_m^k defined by $\mathbf{z}_1, \mathbf{z}_2, \dots, \mathbf{z}_k \mapsto \prod_{i=1}^k J(\mathbf{z}_i)$ is also continuous. Hence, there exist $\epsilon_1, \epsilon_2 > 0$ smaller than 1 such that if $\|\mathbf{z}_i - \tau\| \leq \epsilon_1$ for $1 \leq i \leq k$ then

$$(4.6) \quad \left\| \prod_{i=1}^k J(\mathbf{z}_i) \right\|_1 \leq \rho^k - \epsilon_2.$$

Further, since Δ_m is compact and f is continuously differentiable, $\|J\|_1$ is bounded above on Δ_m by some positive constant L , which we assume without loss of generality to be greater than 1. Similarly, since f has bounded second derivatives, it follows from the multivariate Taylor's theorem that there exists a positive constant B (which we can again assume to be greater than 1) such that for any $\mathbf{x}, \mathbf{y} \in \Delta_m$, we can find a vector ν such that $\|\nu\|_1 \leq Bm \|\mathbf{x} - \mathbf{y}\|_2^2 \leq Bm \|\mathbf{x} - \mathbf{y}\|_1^2$ such that

$$(4.7) \quad f(\mathbf{x}) = f(\mathbf{y}) + J(\mathbf{y})(\mathbf{x} - \mathbf{y}) + \nu.$$

We can now choose

$$\epsilon = \min \left\{ \epsilon_1, \frac{\epsilon_2}{4Bmk(L+1)^{k-1}} \right\} \leq 1, \text{ and}$$

$$\delta = 2Bm\epsilon \leq 1.$$

With this setup, we are now ready to proceed with the proof. Our starting point is the use of a first order Taylor expansion to control the error $\mathbf{x}^{(i)} - \mathbf{y}^{(i)}$ in terms of $\mathbf{x}^{(i-1)} - \mathbf{y}^{(i-1)}$. Indeed, eq. (4.7) when applied to this situation (along with the hypotheses of the theorem)

yields for any $1 \leq i \leq k$ that

$$(4.8) \quad \begin{aligned} \mathbf{x}^{(i)} - \mathbf{y}^{(i)} &= f(\mathbf{x}^{(i-1)}) - f(\mathbf{y}^{(i-1)}) + \xi^{(i)} \\ &= J(\mathbf{y}^{(i-1)})(\mathbf{x}^{(i-1)} - \mathbf{y}^{(i-1)}) + (\nu^{(i)} + \xi^{(i)}), \end{aligned}$$

where $\nu^{(i)}$ satisfies $\|\nu^{(i)}\|_1 \leq Bm \|\mathbf{x}^{(i-1)} - \mathbf{y}^{(i-1)}\|_1^2$. Before proceeding, we first take note of a simple consequence of eq. (4.8). Taking the ℓ_1 norm of both sides, and using the conditions on the norms of $\nu^{(i)}$ and $\xi^{(i)}$, we have

$$\begin{aligned} \|\mathbf{x}^{(i)} - \mathbf{y}^{(i)}\|_1 &\leq \|\mathbf{x}^{(i-1)} - \mathbf{y}^{(i-1)}\|_1 \\ &\quad \cdot \left(L + \delta + Bm \|\mathbf{x}^{(i-1)} - \mathbf{y}^{(i-1)}\|_1 \right). \end{aligned}$$

Since both $\mathbf{x}^{(i-1)}$ and $\mathbf{y}^{(i-1)}$ are within distance ϵ of τ by the hypothesis of the theorem, the above calculation and the definition of ϵ and δ imply that

$$\begin{aligned} \|\mathbf{x}^{(i)} - \mathbf{y}^{(i)}\|_1 &\leq (L + 4Bm\epsilon) \|\mathbf{x}^{(i-1)} - \mathbf{y}^{(i-1)}\|_1 \\ &\leq (L + 1) \|\mathbf{x}^{(i-1)} - \mathbf{y}^{(i-1)}\|_1, \end{aligned}$$

where in the last inequality we use $4Bm\epsilon \leq \epsilon_2 \leq 1$. This, in turn, implies via induction that for every $1 \leq i \leq k$,

$$(4.9) \quad \|\mathbf{x}^{(i)} - \mathbf{y}^{(i)}\|_1 \leq (L + 1)^i \|\mathbf{x}^{(0)} - \mathbf{y}^{(0)}\|_1.$$

We now return to the proof. By iterating eq. (4.8), we can control $\mathbf{x}^{(k)} - \mathbf{y}^{(k)}$ in terms of a product of k Jacobians, as follows:

$$\begin{aligned} \mathbf{x}^{(k)} - \mathbf{y}^{(k)} &= \left(\prod_{i=1}^{i=k} J(\mathbf{y}^{(k-i)}) \right) (\mathbf{x}^{(0)} - \mathbf{y}^{(0)}) \\ &\quad + \sum_{i=1}^k \left(\prod_{j=0}^{j=k-i-1} J(\mathbf{y}^{(k-j)}) \right) (\xi^{(i)} + \nu^{(i)}). \end{aligned}$$

Since $\|\mathbf{y}^{(i)} - \tau\|_1 \leq \epsilon$ by the hypothesis of the theorem, we get from eq. (4.6) that the leftmost term in the above sum has ℓ_1 norm less than $(\rho^k - \epsilon_2) \|\mathbf{x}^{(0)} - \mathbf{y}^{(0)}\|_1$.

We now proceed to estimate the terms in the summation. Our first step to use the conditions on the norms of $\nu^{(i)}$ and $\xi^{(i)}$ and the fact that $\|J\|_1 \leq L$ uniformly to obtain the upper bound

$$\begin{aligned} \sum_{i=1}^k L^{k-i} \left(\|\mathbf{x}^{(i-1)} - \mathbf{y}^{(i-1)}\|_1 \right. \\ \left. \cdot (Bm \|\mathbf{x}^{(i-1)} - \mathbf{y}^{(i-1)}\|_1 + \delta) \right). \end{aligned}$$

Now, recalling that $\mathbf{x}^{(i)}$ and $\mathbf{y}^{(i)}$ are both within an ϵ ℓ_1 -neighborhood of $\boldsymbol{\tau}$ so that $\|\mathbf{x}^{(i)} - \mathbf{y}^{(i)}\|_1 \leq 2\epsilon$, we can estimate the above upper bound as follows:

$$\begin{aligned} & \sum_{i=1}^k L^{k-i} \left(\|\mathbf{x}^{(i-1)} - \mathbf{y}^{(i-1)}\|_1 \right. \\ & \quad \cdot (Bm \|\mathbf{x}^{(i-1)} - \mathbf{y}^{(i-1)}\|_1 + \delta) \Big) \\ & \leq (L+1)^{k-1} \|\mathbf{x}^{(0)} - \mathbf{y}^{(0)}\|_1 \\ & \quad \cdot \sum_{i=1}^k (Bm \|\mathbf{x}^{(i-1)} - \mathbf{y}^{(i-1)}\|_1 + \delta) \\ & \leq k(L+1)^{k-1} (2Bm\epsilon + \delta) \|\mathbf{x}^{(0)} - \mathbf{y}^{(0)}\|_1 \\ & \leq \epsilon_2 \|\mathbf{x}^{(0)} - \mathbf{y}^{(0)}\|_1, \end{aligned}$$

where the first inequality is an application of eq. (4.9), and the last uses the definitions of ϵ and δ . Combining with the upper bound of $(\rho^k - \epsilon_2) \|\mathbf{x}^{(0)} - \mathbf{y}^{(0)}\|_1$ obtained above for the first term, this yields the result. \square

Remark 4.1. Note that k in the theorem can be chosen as large as we want. However, for simplicity, we fix some $k > k_0$ in the rest of the discussion, and revisit the freedom of choice of k only toward the end of the proof of the main theorem (Theorem 3.6) on page 14.

4.1 Evolution under random perturbations. We now explore some consequences of the above theorem for stochastic evolutions. Our main goal in this subsection is to highlight the subtleties that arise in ensuring that the third ‘‘small perturbations’’ condition during a random evolution, and strategies that can be used to avoid them. However, we first begin by showing the second condition, that of ‘‘closeness to the fixed point’’ is actually quite simple to maintain. It will be convenient to define for this purpose the notion of an *epoch*, which is simply the set of $(k+1)$ initial and final states of k consecutive steps of a stochastic evolution.

Definition 4.1 (Epoch). Let f be a smooth contractive evolution and let k be as in the statement of Theorem 4.1 when applied to f . An *epoch* is a set of $k+1$ consecutive states in a stochastic evolution guided by f . By a slight abuse of terminology we also use the same term to refer to a set of $k+1$ consecutive states in a *pair* of stochastic evolutions guided by f that have been coupled using the optimal coupling.

Suppose we want to apply Theorem 4.1 to a pair of stochastic evolutions guided by f . Recall the parameter ϵ in the statement of Theorem 4.1. Ideally, we would like to show that if both the states in the pair at the

beginning of an epoch are within some distance $\epsilon' < \epsilon$ of the fixed point $\boldsymbol{\tau}$, then (1) all the consequent steps in the epoch are within distance ϵ of the fixed point (so that the closeness condition in the theorem is satisfied), and more importantly (2) that the states at the last step of the epoch are *again* within the same distance ϵ' of the fixed point, so that we have the ability to apply the theorem to the *next* epoch. Of course, we also need to ensure that the condition on the perturbations being true also holds during the epoch, but as stated above, this is somewhat more tricky to maintain than the closeness condition, so we defer its discussion to later in the section. Here, we state the following lemma which shows that the closeness condition can indeed be maintained at the end of the epoch.

Lemma 4.2 (Remaining close to the fixed point).

Let $w < w' < 1/3$ be fixed constants. Consider a stochastic evolution $\mathbf{X}^{(0)}, \mathbf{X}^{(1)}, \dots$ on a population of size N guided by a (L, B, ρ) -smooth contractive evolution $f : \Delta_m \rightarrow \Delta_m$ with fixed point $\boldsymbol{\tau}$. Suppose $\alpha > 1$ is such that $\|\mathbf{X}^{(0)} - \boldsymbol{\tau}\|_1 \leq \frac{\alpha}{N^w}$. If N is chosen large enough (as a function of L, α, m, k, w, w' and ρ), then with probability at least $1 - 2mk \exp(-N^{w'}/2)$ we have

- $\|\mathbf{X}^{(i)} - \boldsymbol{\tau}\|_1 \leq \frac{(\alpha+m)(L+1)^i}{N^w}$, for $1 \leq i \leq k-1$.
- $\|\mathbf{X}^{(k)} - \boldsymbol{\tau}\|_1 \leq \frac{\alpha}{N^w}$.

To prove the lemma, we need the following simple concentration result the proof of which is deferred to Appendix B.

Lemma 4.3. Let $\mathbf{X}^{(0)}, \mathbf{X}^{(1)}, \dots$ be a stochastic evolution on a population of size N which is guided by a (L, B, ρ) -smooth contractive evolution $f : \Delta_m \rightarrow \Delta_m$ with fixed point $\boldsymbol{\tau}$. For any $t > 0$ and $\gamma \leq 1/3$, it holds with probability at least $1 - 2mt \exp(-N^\gamma/2)$ that

$$\|\mathbf{X}^{(i)} - f^i(\mathbf{X}^{(0)})\|_1 \leq \frac{(L+1)^i m}{N^\gamma} \text{ for all } 1 \leq i \leq t.$$

Proof of Lemma 4.2. Lemma 4.3 implies that with probability at least $1 - 2mk \exp(-N^{w'}/2)$ we have

$$(4.10) \quad \|\mathbf{X}^{(i)} - f^i(\mathbf{X}^{(0)})\|_1 \leq \frac{(L+1)^i m}{N^w}, \text{ for } 1 \leq i \leq k.$$

On the other hand, the fact that $\max \|\mathbf{J}(\mathbf{x})\|_1 \leq L$ implies that

$$\begin{aligned} \|f^{i+1}(\mathbf{X}^{(0)}) - \boldsymbol{\tau}\|_1 &= \|f^{i+1}(\mathbf{X}^{(0)}) - f(\boldsymbol{\tau})\|_1 \\ &\leq L \|f^i(\mathbf{X}^{(0)}) - \boldsymbol{\tau}\|_1, \end{aligned}$$

so that

$$(4.11) \quad \begin{aligned} \left\| f^i(\mathbf{X}^{(0)}) - \tau \right\|_1 &\leq L^i \left\| \mathbf{X}^{(0)} - \tau \right\|_1 \\ &\leq \frac{\alpha L^i}{N^w}, \text{ for } 1 \leq i \leq k. \end{aligned}$$

Combining eqs. (4.10) and (4.11), we already get the first item in the lemma. However, for $i = k$, we can do much better than the above estimate (and indeed, this is the most important part of the lemma). Recall the parameter ϵ in Theorem 4.1. If we choose N large enough so that

$$(4.12) \quad \frac{(\alpha + m)(L + 1)^k}{N^w} \leq \epsilon,$$

then the above argument shows that with probability at least $1 - 2mk \exp(-N^{w'})$, the sequences $\mathbf{y}^{(i)} = f^i(\mathbf{X}^{(0)})$ and $\mathbf{z}^{(i)} = \tau$ (for $0 \leq i \leq k$) satisfy the hypotheses of Theorem 4.1: the perturbations in this case are simply 0. Hence, we get that

$$(4.13) \quad \left\| f^k(\mathbf{X}^{(0)}) - \tau \right\|_1 \leq \rho^k \left\| \mathbf{X}^{(0)} - \tau \right\|_1 \leq \frac{\alpha \rho^k}{N^w}.$$

Using eq. (4.13) with the $i = k$ case of eq. (4.10), we then have

$$\left\| \mathbf{X}^{(k)} - \tau \right\|_1 \leq \frac{(L + 1)^k m}{N^{w'}} + \frac{\rho^k \alpha}{N^w}.$$

Thus, (since $\rho < 1$) we only need to choose N so that

$$(4.14) \quad N^{w'-w} \geq \frac{(L + 1)^k m}{\alpha(1 - \rho^k)}$$

in order to get the second item in the lemma. Since $w > 0$ and $w' > w$, it follows that all large enough N will satisfy the conditions in both eqs. (4.12) and (4.14), and this completes the proof. \square

4.2 Controlling the size of random perturbations. We now address the “small perturbations” condition of Theorem 4.1. For a given smooth contractive evolution f , let α, w, w' be any constants satisfying the hypotheses of Lemma 4.2 (the precise values of these constants will be specified in the next section). For some N as large as required by the lemma, consider a pair $\mathbf{X}^{(t)}, \mathbf{Y}^{(t)}$ of stochastic evolutions guided by f on a population of size N , which are coupled according to the optimal coupling. Now, let us call an epoch *decent* if the first states $\mathbf{X}^{(0)}$ and $\mathbf{Y}^{(0)}$ in the epoch satisfy $\left\| \mathbf{X}^{(0)} - \tau \right\|_1, \left\| \mathbf{Y}^{(0)} - \tau \right\|_1 \leq \alpha N^{-w}$. The lemma (because of the choice of N made in eq. (4.12)) shows that if an epoch is decent, then except with probability that is sub-exponentially small in N ,

1. the current epoch satisfies the “closeness to fixed point” condition in Theorem 4.1, and
2. the *next* epoch is decent as well.

Thus, the lemma implies that if a certain epoch is decent, then with all but sub-exponential (in N) probability, a polynomial (in N) number of subsequent epochs are also decent, and hence satisfy the “closeness to fixed point” condition of Theorem 4.1. Hypothetically, if these epochs also satisfied the “small perturbation” condition, then we would be done, since in such a situation, the distance between the two chains will drop to less than $1/N$ within $O(\log N)$ time, implying that they would collide. This would in turn imply a $O(\log N)$ mixing time.

However, as alluded to above, ensuring the “small perturbations” condition turns out to be more subtle. In particular, the fact that the perturbations $\boldsymbol{\xi}^{(i)}$ need to be *multiplicatively* smaller than the actual differences $\left\| \mathbf{x}^{(i)} - \mathbf{y}^{(i)} \right\|_1$ pose a problem in achieving adequate concentration, and we cannot hope to prove that the “small perturbations” condition holds with very high probability over an epoch when the starting difference $\left\| \mathbf{X}^{(0)} - \mathbf{Y}^{(0)} \right\|_1$ is very small. As such, we need to break the arguments into two stages based on the starting differences at the start of the epochs lying in the two stages.

To make this more precise (and to state a result which provides examples of the above phenomenon and will also be a building block in the coupling proof), we define the notion of an epoch being *good* with a *goal* g . As before let $\mathbf{X}^{(t)}$ and $\mathbf{Y}^{(t)}$ be two stochastic evolutions guided by f which are coupled according to the optimal coupling, and let $\boldsymbol{\xi}^{(i)}$ be the perturbations as defined in Theorem 4.1. Then, we say that a decent epoch (which we can assume, without loss of generality, to start at $t = 0$) is *good* with goal g if one of following two conditions holds. Either (1) there is a j , $0 \leq j \leq k - 1$ such that $\left\| f(\mathbf{X}^{(j)}) - f(\mathbf{Y}^{(j)}) \right\|_1 \leq g$, or otherwise, (2) it holds that the next epoch is also decent, and, further

$$\left\| \boldsymbol{\xi}^{(i)} \right\|_1 \leq \delta \left\| \mathbf{X}^{(i)} - \mathbf{Y}^{(i)} \right\|_1 \text{ for } 0 \leq i \leq k,$$

where δ again is as defined in Theorem 4.1. Note that if an epoch is good with goal g , then either the expected difference between the two chains drops below g sometime during the epoch, or else, all conditions of Theorem 4.1 are satisfied during the epoch, and the distance between the chains drops by a factor of ρ^k . Further, in terms of this notion, the preceding discussion can be summarized as “the probability of an epoch being good depends upon the goal g , and can be small if g is too small”. To make this concrete, we prove the following theorem which quantifies this trade-off between the size

of the goal g and the probability with which an epoch is good with that goal.

Theorem 4.4 (Goodness with a given goal). *Let the chains $\mathbf{X}^{(t)}$, $\mathbf{Y}^{(t)}$, and the quantities $N, m, w, w', k, L, \epsilon$ and δ be as defined above, and let $\beta < (\log N)^2$. If N is large enough, then a decent epoch is good with goal $g := \frac{4L^2 m \beta}{\delta^2 N}$ with probability at least $1 - 2mk \left(\exp(-\beta/3) + \exp(-N^{w'}/2) \right)$.*

Proof. Let $\mathbf{X}^{(0)}$ and $\mathbf{Y}^{(0)}$ denote the first states in the epoch. Since the current epoch is assumed to be decent, Lemma 4.2 implies that with probability at least $1 - 2mk \exp(-N^{w'}/2)$, the ‘‘closeness to fixed point’’ condition of Theorem 4.1 holds throughout the epoch, and the next epoch is also decent. If there is a $j \leq k - 1$ such that $\|f(\mathbf{X}^{(j)}) - f(\mathbf{Y}^{(j)})\|_1 \leq g$, then the current epoch is already good with goal g . So let us assume that

$$\begin{aligned} \left\| f(\mathbf{X}^{(i-1)}) - f(\mathbf{Y}^{(i-1)}) \right\|_1 &\geq g \\ &= \frac{4L}{\delta^2} \cdot \frac{\beta m}{N} \text{ for } 1 \leq i \leq k. \end{aligned}$$

However, in this case, we can apply the concentration result in Lemma 3.5 with $c = 4L^2/\delta^2$ and $t = \beta$ to get that with probability at least $1 - 2mk \exp(-\beta/3)$,

$$\begin{aligned} \left\| \boldsymbol{\xi}^{(i)} \right\|_1 &\leq \frac{\delta}{L} \left\| f(\mathbf{X}^{(i-1)}) - f(\mathbf{Y}^{(i-1)}) \right\|_1 \\ &\leq \delta \left\| \mathbf{X}^{(i-1)} - \mathbf{Y}^{(i-1)} \right\|_1 \text{ for } 1 \leq i \leq k. \end{aligned}$$

Hence, both conditions (‘‘closeness to fixed point’’ and ‘‘small perturbations’’) for being good with goal g hold with the claimed probability. \square

Note that we need to take β to a large constant, at least $\Omega(\log(mk))$, even to make the result non-trivial. In particular, if we take $\beta = 3 \log(4mk)$, then if N is large enough, the probability of success is at least $1/e$. However, with a slightly larger goal g , it is possible to reduce the probability of an epoch not being good to $o_N(1)$: if we choose $\beta = \log N$, then a decent epoch is good with the corresponding goal with probability at least $1 - N^{-1/4}$, for N large enough. In the next section, we use both these settings of parameters in the above theorem to complete the proof of the mixing time result. As described in Section 2, the two settings above will be used in different stages of the evolution of two coupled chains in order to argue that the time to collision of the chains is indeed small.

5 Proof of the main theorem: Analyzing the coupling time

Our goal is now to show that if we couple two stochastic evolutions guided by the same smooth contractive

evolution f using the optimal coupling, then irrespective of their starting positions, they reach the same state in a small number of steps, with reasonably high probability. More precisely, our proof would be structured as follows. Fix any starting states $\mathbf{X}^{(0)}$ and $\mathbf{Y}^{(0)}$ of the two chains, and couple their evolutions according to the optimal coupling. Let T be the first time such that $\mathbf{X}^{(T)} = \mathbf{Y}^{(T)}$. Suppose that we establish that $\mathbb{P}[T < t] \geq q$, where t and p do not depend upon the starting states $(\mathbf{X}^{(0)}, \mathbf{Y}^{(0)})$. Then, we can dovetail this argument for ℓ ‘‘windows’’ of time t each to see that $\mathbb{P}[T > \ell \cdot t] \leq (1 - q)^\ell$: this is possible because the probability bounds for T did not depend upon the starting positions $(\mathbf{X}^{(0)}, \mathbf{Y}^{(0)})$ and hence can be applied again to the starting positions $(\mathbf{X}^{(t)}, \mathbf{Y}^{(t)})$ if $\mathbf{X}^{(t)} \neq \mathbf{Y}^{(t)}$. By choosing ℓ large enough so that $(1 - q)^\ell$ is at most $1/e$ (or any other constant less than $1/2$), we obtain a mixing time of ℓt . We therefore proceed to obtain an upper bound on $\mathbb{P}[T < t]$ for some $t = \Theta(\log N)$.

As discussed earlier, we need to split the evolution of the chains into several stages in order to complete the argument outlined above. We now describe these four different stages. Recall that f is assumed to be a (L, B, ρ) -smooth contractive evolution. Without loss of generality we assume that $L > 1$. The parameter r appearing below is a function of these parameters and k and is defined in Lemma A.1. Further, as we noted after the proof of Theorem 4.1, k can be chosen to be as large as desired. We now exercise this choice by choosing k to be large enough so that

$$(5.15) \quad \rho^k \leq e^{-1}.$$

The other parameters below are chosen to ease the application of the framework developed in the previous section.

1. **Approaching the fixed point.** We define T_{start} to be the first time such that

$$\left\| \mathbf{X}^{(T_{\text{start}}+i)} - \boldsymbol{\tau} \right\|_1, \left\| \mathbf{Y}^{(T_{\text{start}}+i)} - \boldsymbol{\tau} \right\|_1 \leq \frac{\alpha}{N^w} \text{ for } 0 \leq i \leq k - 1,$$

where $\alpha := m + r$ and $w = \min\left(\frac{1}{6}, \frac{\log(1/\rho)}{6 \log(L+1)}\right)$. We show below that

$$(5.16) \quad \mathbb{P}[T_{\text{start}} > t_{\text{start}} \log N] \leq 4mkt_o \log N \exp\left(-N^{1/3}\right),$$

where $t_{\text{start}} := \frac{1}{6 \log(L+1)}$. The probability itself is upper bounded by $\exp(-N^{1/4})$ for N large enough.

2. **Coming within distance** $\Theta\left(\frac{\log N}{N}\right)$. Let $\beta_0 := (8/\rho^k) \log(17mk)$ and $h = \frac{4L^2m}{\delta^2}$. Then, we define T_0 to be the smallest number of steps needed after T_{start} such that either

$$\left\| \mathbf{X}^{(T_{\text{start}}+T_0)} - \mathbf{Y}^{(T_{\text{start}}+T_0)} \right\|_1 \leq \frac{h\beta_0 \log N}{N}$$

or

$$\left\| f(\mathbf{X}^{(T_{\text{start}}+T_0)}) - f(\mathbf{Y}^{(T_{\text{start}}+T_0)}) \right\|_1 \leq \frac{h\beta_0 \log N}{N(1+\delta)}.$$

We prove below that when N is large enough

$$(5.17) \quad \mathbb{P}[T_0 > kt_0 \log N] \leq \frac{1}{N^{\beta_0/7}},$$

where $t_0 := \frac{1}{k \log(1/\rho)}$.

3. **Coming within distance** $\Theta(1/N)$. Let β_0 and h be as defined in the last item. We now define a sequence of $\ell_1 := \left\lceil \frac{\log \log N}{k \log(1/\rho)} \right\rceil$ random variables $T_1, T_2, \dots, T_{\ell_1}$. We begin by defining the stopping time $S_0 := T_{\text{start}} + T_0$. For $i \geq 1$, T_i is defined to be the smallest number of steps after S_{i-1} such that the corresponding stopping time $S_i := S_{i-1} + T_i$ satisfies either

$$\left\| \mathbf{X}^{(S_i)} - \mathbf{Y}^{(S_i)} \right\|_1 \leq \frac{h\rho^{ik}\beta_0 \log N}{N}$$

or

$$\left\| f(\mathbf{X}^{(S_i)}) - f(\mathbf{Y}^{(S_i)}) \right\|_1 \leq \frac{h\rho^{ik}\beta_0 \log N}{N(1+\delta)}.$$

Note that T_i is defined to be 0 if setting $S_i = S_{i-1}$ already satisfies the above conditions. Define $\beta_i = \rho^{ik}\beta_0$. We prove below that when N is large enough

$$(5.18) \quad \mathbb{P}[T_i > k+1] \leq 4mk \exp(-(\beta_i \log N)/8),$$

for $1 \leq i \leq \ell_1$.

4. **Collision.** Let β_0 and h be as defined in the last two items. Note that after time S_{ℓ_1} , we have

$$\begin{aligned} & \left\| f(\mathbf{X}^{(S_{\ell_1})}) - f(\mathbf{Y}^{(S_{\ell_1})}) \right\|_1 \\ & \leq \frac{Lh\beta_{\ell_1} \log N}{N} = \frac{Lh\beta_0 \rho^{k\ell_1} \log N}{N} \leq \frac{Lh\beta_0}{N}. \end{aligned}$$

Then, from the properties of the optimal coupling we have that $\mathbf{X}^{(S_{\ell_1}+1)} = \mathbf{Y}^{(S_{\ell_1}+1)}$ with probability at least $\left(1 - \frac{Lh\beta_0}{2N}\right)^N$ which is at least $\exp(-L\beta_0 h)$ when N is so large that $N > hL\beta_0$.

Assuming eqs. (5.16) to (5.18), we can complete the proof of Theorem 3.6 as follows.

Proof of Theorem 3.6. Let $\mathbf{X}^{(0)}, \mathbf{Y}^{(0)}$ be the arbitrary starting states of two stochastic evolutions guided by f , whose evolution is coupled using the optimal coupling. Let T be the minimum time t satisfying $\mathbf{X}^{(t)} = \mathbf{Y}^{(t)}$. By the Markovian property and the probability bounds in items 1 to 4 above, we have (for large enough N)

$$\begin{aligned} & \mathbb{P}[T \leq t_{\text{start}} \log N + kt_0 \log N + (k+1)\ell_1] \\ & \geq \mathbb{P}[T_{\text{start}} \leq t_{\text{start}} \log N] \cdot \mathbb{P}[T_0 \leq kt_0 \log N] \\ & \quad \cdot \left(\prod_{i=1}^{\ell_1} \mathbb{P}[T_i \leq k+1] \right) \cdot e^{-L\beta_0 h} \\ & \geq e^{-L\beta_0 h} \left(1 - \exp(-N^{1/4}) \right) (1 - N^{-\beta/7}) \\ & \quad \cdot \prod_{i=1}^{\ell_1} (1 - 4mk \exp(-(\rho^{ik}\beta_0 \log N)/8)) \\ & \geq \exp(-L\beta_0 h - 1) \\ & \quad \cdot \left(1 - 4mk \sum_{i=1}^{\ell_1} \exp(-(\rho^{ik}\beta_0 \log N)/8) \right) \end{aligned}$$

where the last inequality is true for large enough N . Applying Lemma C.1 to the above sum (with the parameters x and α in the lemma defined as $x = \exp(-(\beta_0 \log N)/8)$ and $\alpha = \rho^k \leq 1/e$ by the assumption in eq. (5.15)), we can put an upper bound on it as follows:

$$\begin{aligned} \sum_{i=1}^{\ell_1} \exp(-(\rho^{ik}\beta_0 \log N)/8) & \leq \frac{1}{\exp(\rho^k\beta_0/8) - 1} \\ & = \frac{1}{17mk - 1} \leq \frac{1}{16mk}, \end{aligned}$$

where the first inequality follows from the lemma and the fact that $\frac{\log \log N}{k \log(1/\rho)} \leq \ell_1 \leq 1 + \frac{\log \log N}{k \log(1/\rho)}$, and the last inequality uses the definition of β_0 and $m, k \geq 1$. Thus, for large enough N , we have

$$\mathbb{P}[T \leq c \log N] \geq q,$$

where $c := 2(t_{\text{start}} + kt_0)$ and $q := (3/4) \exp(-L\beta_0 h - 1)$. Since this estimate does not depend upon the starting states, we can bootstrap the estimate after every $c \log N$ steps to get

$$\mathbb{P}[T > c\ell \log N] < (1-q)^\ell \leq e^{-q\ell},$$

which shows that

$$\frac{c}{q} \cdot \log N$$

is the mixing time of the chain for total variation distance $1/e$, when N is large enough. \square

We now proceed to prove the claimed equations, starting with eq. (5.16). Let $t := t_{\text{start}} \log N$ for convenience of notation. From Lemma A.1 we have

$$\left\| f^t(\mathbf{X}^{(0)}) - \tau \right\|_1 \leq r \rho^t.$$

On the other hand, applying Lemma 4.3 to the chain $\mathbf{X}^{(0)}, \mathbf{X}^{(1)}, \dots$ with $\gamma = 1/3$, we have

$$\left\| f^t(\mathbf{X}^{(0)}) - \mathbf{X}^{(t)} \right\|_1 \leq \frac{(L+1)^t m}{N^{1/3}}$$

with probability at least $1 - 2mt \exp(-N^{1/3})$. From the triangle inequality and the definition of t , we then see that with probability at least $1 - 2mt \exp(-N^{1/3})$, we have

$$\begin{aligned} \left\| \mathbf{X}^{(t)} - \tau \right\|_1 &\leq \frac{m}{N^{1/6}} + \frac{r}{N^{t_{\text{start}} \log(1/\rho)}} \\ &\leq \frac{m+r}{N^w} = \frac{\alpha}{N^w}, \end{aligned}$$

where α, w are as defined in item 1 above. Now, if we instead looked at the chain starting at some $i < k$, the same result would hold for $\mathbf{X}^{(t+i)}$. Further, the same analysis applies also to $\mathbf{Y}^{(t+i)}$. Taking an union bound over these $2k$ events, we get the required result.

Before proceeding with the proof of the other two equations, we record an important consequences of eq. (5.16). Let w, α be as defined above, and let $w' > w$ be such that $w' < 1/3$. Recall that an epoch starting at time 0 is decent if both $\mathbf{X}^{(t)}$ and $\mathbf{Y}^{(t)}$ are within distance α/N^w of τ .

Observation 5.1. *For large enough N , it holds with probability at least $1 - \exp(-N^{w'/4})$ that for $1 \leq i \leq kN$, $\mathbf{X}^{(T_{\text{start}}+i)}, \mathbf{Y}^{(T_{\text{start}}+i)}$ are within ℓ_1 distance α/N^w of τ .*

Proof. We know from item 1 that the epochs starting at times $T_{\text{start}} + i$ for $0 \leq i < k$ are all decent. For large enough N , Lemma 4.2 followed by a union bound implies that the N consecutive epochs starting at $T+j+k\ell$ where $\ell \leq N$ and $0 \leq j \leq N$ are also all decent with probability at least $1 - 2mk^2 N \exp(-N^{w'/2})$, which upper bounds the claimed probability for large enough N . \square

We denote by E the event that the epochs starting at $T_{\text{start}} + i$ for $1 \leq i \leq kN$ are all decent. The above observation says that $P(E) \geq 1 - \exp(-N^{w'/4})$ for N large enough.

We now consider T_0 . Let $g_0 := \frac{h\beta_0 \log N}{N(1+\delta)}$, where h is as defined in items 2 and 3 above. From Theorem 4.4 followed by an union bound, we see that the first $t_1 \log N$ consecutive epochs starting at $T_{\text{start}}, T_{\text{start}} + k, T_{\text{start}} + 2k, \dots$ are good with goal g_0

(they are already known to be decent with probability at least $P(E)$ from the above observation) with probability at least

$$1 - 2mkt_1 \left(N^{-\beta_0/(3(1+\delta))} + \exp(-N^{w'}/4) \right) \log N - P(-E),$$

which is larger than $1 - N^{-\beta_0/7}$ for N large enough (since $\delta < 1$). Now, if we have $\left\| f(\mathbf{X}^{(i)}) - f(\mathbf{Y}^{(i)}) \right\|_1 \leq g$ for some time i during these $t_1 \log N$ good epochs then $T_0 \leq kt_1 \log N$ follows immediately. Otherwise, the goodness condition implies that the hypotheses of Theorem 4.1 are satisfied across all these epochs, and we get

$$\begin{aligned} &\left\| \mathbf{X}^{(T_{\text{start}}+kt_1 \log N)} - \mathbf{Y}^{(T_{\text{start}}+kt_1 \log N)} \right\|_1 \\ &\leq \rho^{kt_1 \log N} \left\| \mathbf{X}^{(T_{\text{start}})} - \mathbf{Y}^{(T_{\text{start}})} \right\|_1 \\ &\leq \rho^{kt_1 \log N} \frac{\alpha}{N^w} \leq g_0 \leq \frac{h\beta_0 \log N}{N}, \end{aligned}$$

where the second last inequality is true for large enough N .

Finally, we analyze T_i for $i \geq 1$. For this, we need to consider cases according to the state of the chain at time S_{i-1} . However, we first observe that plugging our choice of h into Theorem 4.4 shows that any decent epoch is good with goal $g_i := \frac{h\beta_i \log N}{N(1+\delta)}$ with probability at least

$$1 - 2mk \left(\exp(-(\beta_i \log N)/(3(1+\delta))) + \exp(-N^{w'}/4) \right),$$

which is at least $1 - 2mk \exp(-(\beta_i \log N)/7)$ for N large enough (since $\delta < 1$). Further, since we can assume via the above observation that all the epochs we consider are decent with probability at least $P(E)$, it follows that the epoch starting at S_{i-1} (and also the one starting at $S_{i-1} + 1$) is good with goal g_i with probability at least

$$\begin{aligned} p &:= 1 - 2mk \exp(-(\beta_i \log N)/7) - P(-E) \\ &\geq 1 - 2mk \exp(-(\beta_i \log N)/8), \end{aligned}$$

where the last inequality holds whenever $\beta_i \leq \log N$ and N is large enough (we will use at most one of these two epochs in each of the exhaustive cases we consider below). Note that if at any time $S_{i-1} + j$ (where $j \leq k+1$) during one of these two good epochs it happens that $\left\| f(\mathbf{X}^{(S_{i-1}+j)}) - f(\mathbf{Y}^{(S_{i-1}+j)}) \right\|_1 \leq g_i$, then we immediately get $T_i \leq k+1$ as required. We can therefore assume that this does not happen, so that the hypotheses of Theorem 4.1 are satisfied across these epochs.

Now, the first case to consider is $\left\| \mathbf{X}^{(S_{i-1})} - \mathbf{Y}^{(S_{i-1})} \right\|_1 \leq \frac{h\beta_{i-1} \log N}{N}$. Since we are

assuming that Theorem 4.1 is satisfied across the epoch starting at S_{i-1} , we get

$$(5.19) \quad \begin{aligned} & \left\| \mathbf{X}^{(S_{i-1}+k)} - \mathbf{Y}^{(S_{i-1}+k)} \right\|_1 \\ & \leq \rho^k \left\| \mathbf{X}^{(S_{i-1})} - \mathbf{Y}^{(S_{i-1})} \right\|_1 \\ & \leq \rho^k \frac{h\beta_{i-1} \log N}{N} = \frac{h\beta_i \log N}{N}. \end{aligned}$$

Thus, in this case, we have $T_i \leq k$ with probability at least p as defined in the last paragraph.

Even simpler is the case $\left\| f(\mathbf{X}^{(S_{i-1})}) - f(\mathbf{Y}^{(S_{i-1})}) \right\|_1 \leq \frac{h\beta_i \log N}{N}$ in which case T_i is zero by definition. Thus the only remaining case left to consider is

$$\begin{aligned} \frac{h\beta_i \log N}{N} & < \left\| f(\mathbf{X}^{(S_{i-1})}) - f(\mathbf{Y}^{(S_{i-1})}) \right\|_1 \\ & \leq \frac{h\beta_{i-1} \log N}{N(1+\delta)}. \end{aligned}$$

Since $h = \frac{4L^2 m}{\delta^2}$, the first inequality allows us to use Lemma 3.5 with the parameters c and t in that lemma set to $c = 4/\delta^2$ and $t = \beta_i L^2 \log N$, and we obtain

$$\begin{aligned} & \left\| \mathbf{X}^{(S_{i-1}+1)} - \mathbf{Y}^{(S_{i-1}+1)} \right\|_1 \\ & \leq (1+\delta) \left\| f(\mathbf{X}^{(S_{i-1})}) - f(\mathbf{Y}^{(S_{i-1})}) \right\|_1 \\ & \leq \frac{h\beta_{i-1} \log N}{N}, \end{aligned}$$

with probability at least $1 - 2m \exp(-(\beta_i L^2 \log N)/3)$. Using the same analysis as the first case from this point onward (the only difference being that we need to use the epoch starting at $S_{i-1} + 1$ instead of the epoch starting at S_{i-1} used in that case), we get that

$$\begin{aligned} \mathbb{P}[T_i \leq 1+k] & \geq p - 2m \exp(-(\beta_i L^2 \log N)/3) \\ & \geq 1 - 4mk \exp(-(\beta_i \log N)/8). \end{aligned}$$

since $L, k > 1$. Together with eq. (5.19), this completes the proof of eq. (5.18).

References

- [1] ALVES, D., AND FONTANARI, J. F. Error threshold in finite populations. *Phys. Rev. E* 57, 6 (June 1998), 7008 – 7013.
- [2] BALAGAM, R., SINGH, V., SAGI, A. R., AND DIXIT, N. M. Taking multiple infections of cells and recombination into account leads to small within-host effective-population-size estimates of HIV-1. *PLoS ONE* 6, 1 (01 2011), e14531.

- [3] CAMPOS, P. R. A., AND FONTANARI, J. F. Finite-size scaling of the quasispecies model. *Phys. Rev. E* 58 (1998), 2664–2667.
- [4] DIXIT, N., SRIVASTAVA, P., AND VISHNOI, N. K. A finite population model of molecular evolution: Theory and computation. *J. Comput. Biol.* 19, 10 (2012), 1176–1202.
- [5] DUBHASHI, D. P., AND PANCONESI, A. *Concentration of Measure for the Analysis of Randomized Algorithms*. Cambridge University Press, 2009.
- [6] EIGEN, M. Selforganization of matter and the evolution of biological macromolecules. *Die Naturwissenschaften* 58 (1971), 456–523.
- [7] EIGEN, M., AND SCHUSTER, P. The hypercycle, a principle of natural self-organization. Part A: Emergence of the hypercycle. *Die Naturwissenschaften* 64 (1977), 541–565.
- [8] KOUYOS, R. D., ALTHAUS, C. L., AND BONHOEFFER, S. Stochastic or deterministic: What is the effective population size of HIV-1? *Trends Microbiol.* 14, 12 (2006), 507 – 511.
- [9] LIEBERMAN, E., HAUERT, C., AND NOWAK, M. A. Evolutionary dynamics on graphs. *Nature* 433, 7023 (Jan. 2005), 312–316.
- [10] MUSSO, F. A stochastic version of Eigen’s model. *Bull. Math. Biol.* 73 (2011), 151 – 180.
- [11] NOWAK, M. *Evolutionary Dynamics*. Harvard University Press, 2006.
- [12] NOWAK, M., AND SCHUSTER, P. Error thresholds of replication in finite populations-mutation frequencies and the onset of muller’s ratchet. *J. Theor. Biol.* 137 (1989), 375–395.
- [13] RAY, T. S., PAYNE, K. A., AND MOSELEY, L. L. Role of finite populations in determining evolutionary dynamics. *Phys. Rev. E* 77, 2 (Feb 2008), 021909.
- [14] SAAKIAN, D. B., ROZANOVA, O., AND AKMETZHANOV, A. Dynamics of the Eigen and the Crow-Kimura models for molecular evolution. *Phys. Rev. E* 78, 4 (Oct 2008), 041908.
- [15] TRIPATHI, K., BALAGAM, R., VISHNOI, N. K., AND DIXIT, N. M. Stochastic simulations suggest that HIV-1 survives close to its error threshold. *PLoS Comput. Biol.* 8, 9 (2012), e1002684.
- [16] VAN NIMWEGEN, E., CRUTCHFIELD, J. P., AND MITCHELL, M. Statistical dynamics of the Royal Road Genetic Algorithm. *Theor. Comput. Sci.* 229 (1999), 41 – 102.
- [17] VISHNOI, N. K. Evolution without sex, drugs and Boolean functions. <http://theory.epfl.ch/vishnoi/> (2013).
- [18] VISHNOI, N. K. The speed of evolution. In *Proc. 26th Annual ACM-SIAM Symp. Discret. Algorithms (SODA)* (2015), pp. 1590–1601.
- [19] WILKE, C. Quasispecies theory in the context of population genetics. *BMC Evol. Biol.* 5, 1 (2005), 44.

A Proofs omitted from Section 3.1

Lemma A.1 (Exponential convergence). *Let f be a smooth contractive evolution, and let τ and ρ be as in the conditions described in Section 3.1. Then, there exist a positive r such that for every $\mathbf{z} \in \Delta_m$, and every positive integer t ,*

$$\|f^t(\mathbf{z}) - \tau\|_1 \leq r\rho^t.$$

Proof. Let ϵ and k be as defined in Theorem 4.1. From the ‘‘convergence to the fixed point’’ condition, we know that there exists an ℓ such that for all $\mathbf{z} \in \Delta_m$,

$$(A.1) \quad \|f^\ell(\mathbf{z}) - \tau\|_1 \leq \frac{\epsilon}{L^k}.$$

Note that this implies that $f^{\ell+i}(\mathbf{z})$ is within distance ϵ of τ for $i = 0, 1, \dots, k$, so that Theorem 4.1 can be applied to the sequence of vectors $f^\ell(\mathbf{z}), f^{\ell+1}(\mathbf{z}), \dots, f^{\ell+k}(\mathbf{z})$ and $\tau, f(\tau) = \tau, \dots, f^k(\tau) = \tau$ (the perturbations are simply 0). Thus, we get

$$\|f^{\ell+k}(\mathbf{z}) - \tau\|_1 \leq \rho^k \|f^\ell(\mathbf{z}) - \tau\|_1 \leq \frac{\rho^k \epsilon}{L^k}.$$

Since $\rho < 1$, we see that the epoch starting at $\ell + k$ also satisfies eq. (A.1) and hence we can iterate this process. Using also the fact that the $1 \rightarrow 1$ norm of the Jacobian of f is at most L (which we can assume without loss of generality to be at least 1), we therefore get for every $\mathbf{z} \in \Delta_m$, and every $i \geq 0$ and $0 \leq j < k$

$$\begin{aligned} \|f^{\ell+ik+j}(\mathbf{z}) - \tau\|_1 &\leq \rho^{ki+j} \frac{L^j}{\rho^j} \|f^\ell(\mathbf{z}) - \tau\|_1 \\ &\leq \rho^{ki+j+\ell} \frac{L^{j+\ell}}{\rho^{j+\ell}} \|\mathbf{z} - \tau\|_1 \\ &\leq \rho^{ki+j+\ell} \frac{L^{k+\ell}}{\rho^{k+\ell}} \|\mathbf{z} - \tau\|_1 \end{aligned}$$

where in the last line we use the facts that $L > 1, \rho < 1$ and $j < k$. Noting that any $t \geq \ell$ is of the form $\ell + ki + j$ for some i and j as above, we have shown that for every $t \geq \ell$ and every $\mathbf{z} \in \Delta_m$

$$(A.2) \quad \|f^t(\mathbf{z}) - \tau\|_1 \leq \left(\frac{L}{\rho}\right)^{k+\ell} \rho^t \|\mathbf{z} - \tau\|_1.$$

Similarly, for $t < \ell$, we have, for any $\mathbf{z} \in \Delta_m$

$$(A.3) \quad \begin{aligned} \|f^t(\mathbf{z}) - \tau\|_1 &\leq L^t \|\mathbf{z} - \tau\|_1 \\ &\leq \left(\frac{L}{\rho}\right)^t \rho^t \|\mathbf{z} - \tau\|_1 \end{aligned}$$

$$(A.4) \quad \leq \left(\frac{L}{\rho}\right)^\ell \rho^t \|\mathbf{z} - \tau\|_1,$$

where in the last line we have again used $L > 1, \rho < 1$ and $t < \ell$. From eqs. (A.2) and (A.4), we get the claimed result with $r := \left(\frac{L}{\rho}\right)^{k+\ell}$. \square

B Proofs omitted from Section 4

Proof of Lemma 4.3. Fix a co-ordinate $j \in [m]$. Since $\mathbf{X}^{(i)}$ is the normalized frequency vector obtained by taking N independent samples from the distribution $f(\mathbf{X}^{(i-1)})$, Hoeffding’s inequality yields that

$$\begin{aligned} \mathbb{P} \left[\left| X_j^{(i)} - f(\mathbf{X}^{(i-1)})_j \right| > N^{-\gamma} \right] &\leq 2 \exp(-N^{1-2\gamma}) \\ &\leq 2 \exp(-N^\gamma), \end{aligned}$$

where the last inequality holds because $\gamma \leq 1/3$. Taking a union bound over all $j \in [m]$, we therefore have that for any fixed $i \leq t$,

$$(B.5) \quad \mathbb{P} \left[\left\| \mathbf{X}^{(i)} - f(\mathbf{X}^{(i-1)}) \right\|_1 > \frac{m}{N^\gamma} \right] \leq 2m \exp(-N^\gamma).$$

For ease of notation let us define the quantities $s^{(i)} := \left\| \mathbf{X}^{(i)} - f^i(\mathbf{X}^{(0)}) \right\|_1$ for $0 \leq i \leq t$. Our goal then is to show that it holds with high probability that $s^{(i)} \leq \frac{(L+1)^i m}{N^\gamma}$ for all i such that $0 \leq i \leq t$.

Now, by taking a union bound over all values of i in eq. (B.5), we see that the following holds for all i with probability at least $1 - 2mt \exp(-N^\gamma)$:

$$(B.6) \quad \begin{aligned} s^{(i)} &= \left\| \mathbf{X}^{(i)} - f^i(\mathbf{X}^{(0)}) \right\|_1 \\ &\leq \left\| \mathbf{X}^{(i)} - f(\mathbf{X}^{(i-1)}) \right\|_1 \\ &\quad + \left\| f(\mathbf{X}^{(i-1)}) - f^i(\mathbf{X}^{(0)}) \right\|_1 \\ &\leq \frac{m}{N^\gamma} + Ls^{(i-1)}, \end{aligned}$$

where the first term is estimated using the probabilistic guarantee from eq. (B.5) and the second using the upper bound on the $1 \rightarrow 1$ norm of the Jacobian of f . However, eq. (B.6) implies that $s^{(i)} \leq \frac{m(L+1)^i}{N^\gamma}$ for all $0 \leq i \leq t$, which is what we wanted to prove. To see the former claim, we proceed by induction. Since $s_0 = 0$, the claim is trivially true in the base case. Assuming the claim is true for $s^{(i)}$, we then apply eq. (B.6) to get

$$\begin{aligned} s^{(i+1)} &\leq \frac{m}{N^\gamma} + Ls^{(i)} \leq \frac{m}{N^\gamma} (1 + L(L+1)^i) \\ &\leq \frac{m}{N^\gamma} \cdot (L+1)^{i+1}. \quad \square \end{aligned}$$

C Sums with exponentially decreasing exponents

The following technical lemma is used in the proof of Theorem 3.6.

Lemma C.1. *Let x, α be positive real numbers less than 1 such that $\alpha < \frac{1}{e}$. Let ℓ be a positive integer, and define*

$y := x^{\alpha^{\ell}}$. Then

$$\sum_{i=0}^{\ell} x^{\alpha^i} \leq \frac{y}{1-y}.$$

Proof. Note that since both x and α are positive and less than 1, so is y . We now have

$$\begin{aligned} \sum_{i=0}^{\ell} x^{\alpha^i} &= \sum_{i=0}^{\ell} x^{\alpha^{\ell-i}} = y \sum_{i=0}^{\ell} y^{\alpha^{-i}-1} \\ &\leq y \sum_{i=0}^{\ell} y^{i \log(1/\alpha)}, \\ &\quad \text{since } 0 < y \leq 1 \text{ and } \alpha^{-i} \geq 1 + i \log(1/\alpha), \\ &\leq y \sum_{i=0}^{\ell} y^i, \text{ since } 0 < y \leq 1 \text{ and } \alpha < 1/e, \\ &\leq \frac{y}{1-y}. \quad \square \end{aligned}$$