



HAL
open science

CUTOFF AT THE " ENTROPIC TIME " FOR SPARSE MARKOV CHAINS

Charles Bordenave, Pietro Caputo, Justin Salez

► **To cite this version:**

Charles Bordenave, Pietro Caputo, Justin Salez. CUTOFF AT THE " ENTROPIC TIME " FOR SPARSE MARKOV CHAINS. Probability Theory and Related Fields, 2019, 173 (1-2), pp.261-292. 10.1007/s00440-018-0834-0 . hal-01391939

HAL Id: hal-01391939

<https://hal.science/hal-01391939>

Submitted on 4 Nov 2016

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

CUTOFF AT THE “ENTROPIC TIME” FOR SPARSE MARKOV CHAINS

CHARLES BORDENAVE, PIETRO CAPUTO, JUSTIN SALEZ

ABSTRACT. We study convergence to equilibrium for a large class of Markov chains in random environment. The chains are sparse in the sense that in every row of the transition matrix P the mass is essentially concentrated on few entries. Moreover, the random environment is such that rows of P are independent and such that the entries are exchangeable within each row. This includes various models of random walks on sparse random directed graphs. The models are generally non reversible and the equilibrium distribution is itself unknown. In this general setting we establish the cutoff phenomenon for the total variation distance to equilibrium, with mixing time given by the logarithm of the number of states times the inverse of the average row entropy of P . As an application, we consider the case where the rows of P are i.i.d. random vectors in the domain of attraction of a Poisson-Dirichlet law with index $\alpha \in (0, 1)$. Our main results are based on a detailed analysis of the weight of the trajectory followed by the walker. This approach offers an interpretation of cutoff as an instance of the concentration of measure phenomenon.

1. INTRODUCTION

Given a $n \times n$ stochastic matrix P with unique invariant law π , and an initial state $i \in [n]$, one may consider the total-variation distance to equilibrium after t iterations, $\|P^t(i, \cdot) - \pi\|_{\text{TV}}$. The time at which this decreasing function of t falls below a given $\varepsilon \in (0, 1)$ is known as the *mixing time*:

$$t_{\text{MIX}}^{(i)}(\varepsilon) := \inf \{t \geq 0: \|P^t(i, \cdot) - \pi\|_{\text{TV}} \leq \varepsilon\}.$$

Estimating this quantity is often a difficult task. The purpose of this paper is to relate it to the following simple information-theoretical statistics, which we call the *entropic time*:

$$t_{\text{ENT}} := \frac{\log n}{\text{H}} \quad \text{where} \quad \text{H} := -\frac{1}{n} \sum_{i,j=1}^n P(i, j) \log P(i, j). \quad (1)$$

H is the average row entropy of the matrix P . Our finding is that, in a certain sense, “most” sparse stochastic matrices have mixing time roughly given by t_{ENT} , regardless of the choice of the precision $\varepsilon \in (0, 1)$ and the initial state $i \in [n]$.

To give a precise meaning to the previous assertion, we define the following model of *Random Stochastic Matrix*. For each $i \in [n]$, let $p_{i1} \geq \dots \geq p_{in} \geq 0$ be given ranked numbers such that $\sum_{j=1}^n p_{ij} = 1$, and define the $n \times n$ random stochastic matrix P by

$$P(i, j) := p_i \sigma_i^{-1}(j), \quad (1 \leq i, j \leq n), \quad (2)$$

where $\sigma = (\sigma_i)_{1 \leq i \leq n}$ is a collection of n independent, uniform random permutations of $[n]$, which we refer to as the *environment*. We sometimes write P_σ instead of P to emphasize the dependence on the environment. Note that the average row entropy $\text{H} = -\frac{1}{n} \sum_{i,j=1}^n p_{ij} \log p_{ij}$ of this random matrix is deterministic. To study large-size asymptotics, we let the input parameters $(p_{ij})_{1 \leq i, j \leq n}$

implicitly depend on n and consider the limit as $n \rightarrow \infty$. Our focus is on the sparse and non-degenerate regime defined below. It might help the reader to think of all these parameters as taking values in $\{0\} \cup [\varepsilon, 1 - \varepsilon]$ for some fixed $\varepsilon \in (0, 1)$, so that the number of non-zero entries in each row is bounded independently of n . However, we will only impose the following weaker conditions:

1. **Sparsity** (in every row, the mass is essentially concentrated on a few entries):

$$H = \mathcal{O}(1) \quad \text{and} \quad \max_{i \in [n]} \sum_{j=1}^n p_{ij} (\log p_{ij})^2 = o(\log n). \quad (3)$$

2. **Non-degeneracy** (in most rows, the mass is not concentrated on a single entry):

$$\limsup_{n \rightarrow \infty} \left\{ \frac{1}{n} \sum_{i=1}^n \mathbf{1}_{p_{i1} > 1 - \varepsilon} \right\} \xrightarrow{\varepsilon \rightarrow 0^+} 0. \quad (4)$$

These conditions imply in particular that $t_{\text{ENT}} = \Theta(\log n)$ as $n \rightarrow \infty$. Our main result states that around the entropic time t_{ENT} , the distance to equilibrium undergoes the following sharp transition, henceforth referred to as a *uniform cutoff* (to emphasize the insensitivity to the initial state). A remark on notation: below we say that an event that depends on n holds *with high probability* if the probability of this event converges to 1 as $n \rightarrow \infty$; we use $\xrightarrow{\mathbf{P}}$ to indicate convergence *in probability*.

Theorem 1 (Uniform cutoff at the entropic time). *Under the above assumptions, the Markov chain defined by P has, with high probability, a unique stationary distribution π . Moreover, for $t = \lambda t_{\text{ENT}} + o(t_{\text{ENT}})$ with λ fixed as $n \rightarrow \infty$,*

$$\lambda < 1 \quad \implies \quad \min_{i \in [n]} \|P^t(i, \cdot) - \pi\|_{\text{TV}} \xrightarrow[n \rightarrow \infty]{\mathbf{P}} 1 \quad (5)$$

$$\lambda > 1 \quad \implies \quad \max_{i \in [n]} \|P^t(i, \cdot) - \pi\|_{\text{TV}} \xrightarrow[n \rightarrow \infty]{\mathbf{P}} 0. \quad (6)$$

Equivalently, for any fixed $\varepsilon \in (0, 1)$, we have

$$\max_{i \in [n]} \left| \frac{t_{\text{MIX}}^{(i)}(\varepsilon)}{t_{\text{ENT}}} - 1 \right| \xrightarrow[n \rightarrow \infty]{\mathbf{P}} 0.$$

Remark 1 (Invariant measure). *The stationary distribution π appearing in the above theorem is itself a non-trivial random object. To overcome this difficulty, we will in fact prove the statements (5)-(6) with π replaced by the explicit approximation*

$$\hat{\pi}(j) := \frac{1}{n} \sum_{i \in [n]} P^{\lfloor \frac{t_{\text{ENT}}}{10} \rfloor}(i, j). \quad (7)$$

The uniformity over the initial state then automatically ensures that the true invariant measure π is unique with high probability and satisfies

$$\|\pi - \hat{\pi}\|_{\text{TV}} \xrightarrow[n \rightarrow \infty]{\mathbf{P}} 0. \quad (8)$$

Thus, once (5)-(6) have been obtained for $\hat{\pi}$, the same conclusion extends to π .

Let us illustrate our result with a special case.

Example 1 (Random walk on random digraphs). *When $p_{i1} = \dots = p_{id_i} = \frac{1}{d_i}$ and $p_{i,d_i+1} = \dots = p_{in} = 0$ for some integers $d_1, \dots, d_n \geq 1$, the random matrix P may be interpreted as the transition matrix of the random walk on a uniform random directed graph with n vertices*

and out-degrees d_1, \dots, d_n (loops are allowed). The average row entropy is then simply the average logarithmic degree $H = \frac{1}{n} \sum_{i=1}^n \log d_i$. Assumption (3) requires that $H = \mathcal{O}(1)$ and that the maximum out-degree Δ satisfies $\Delta = e^{o(\sqrt{\log n})}$, while Assumption (4) simply asks for the proportion of degree-one vertices to vanish. Notice that, because of the possibility of vertices with zero in-degree, the random matrix P may, with uniformly positive probability, fail to be irreducible. However, under the above conditions, Theorem 1 ensures that with high probability there is a unique stationary distribution and the walk exhibits uniform cutoff at time $(\log n)/H$.

Interesting applications of Theorem 1 can be obtained by taking the input parameters (p_{ij}) also random, provided the main assumptions (3)-(4) are satisfied with high probability. The following theorem is concerned with the case where the rows $\{(p_{i1}, \dots, p_{in}), i = 1, \dots, n\}$ are i.i.d. random vectors in the domain of attraction of a Poisson-Dirichlet law.

Theorem 2. *Let $\omega = (\omega_{ij})_{1 \leq i, j < \infty}$ be i.i.d. positive random variables whose tail distribution function $G(t) = \mathbb{P}(\omega_{ij} > t)$ is regularly varying at infinity with index $\alpha \in (0, 1)$, i.e., for each $\lambda > 0$,*

$$\frac{G(\lambda t)}{G(t)} \xrightarrow[t \rightarrow \infty]{} \lambda^{-\alpha}. \quad (9)$$

Then as $n \rightarrow \infty$, the n -state Markov chain with transition matrix

$$P(i, j) := \frac{\omega_{ij}}{\omega_{i1} + \dots + \omega_{in}}, \quad (1 \leq i, j \leq n)$$

has with high probability a unique stationary distribution π , and exhibits uniform cutoff at time $\frac{\log n}{h(\alpha)}$ in the sense of (5)-(6), where $h(\alpha)$ is defined in terms of the digamma function $\psi = \frac{\Gamma'}{\Gamma}$ by

$$h(\alpha) := \psi(1) - \psi(1 - \alpha) = \int_0^\infty \frac{e^{\alpha t} - 1}{e^t - 1} dt. \quad (10)$$

Let us briefly sketch the main ideas behind the proof of our results. The essence of the sharp transition described in Theorem 1 lies in a quenched *concentration of measure* phenomenon in the trajectory space that can be roughly described as follows; we refer to Section 2 for more details. Let $i = X_0, X_1, X_2, \dots$ denote the trajectory of the random walk with transition matrix P and starting point $i \in [n]$ and let Q_i denote the associated quenched law, that is the law of the trajectory for a fixed realization of the environment σ . Define the trajectory weight

$$\rho(t) := P(X_0, X_1) \cdots P(X_{t-1}, X_t).$$

In other words, $\rho(t)$ is the probability of the followed trajectory up to time t . Theorem 4 below establishes that for $t = \Theta(\log n)$, with high probability with respect to the environment, it is very likely, uniformly in the starting point i , that $\log \rho(t) \sim -Ht$. More precisely, we prove that for any $\varepsilon > 0$,

$$\max_{i \in [n]} Q_i \left(\rho(t) \notin \left[e^{-(1+\varepsilon)Ht}, e^{-(1-\varepsilon)Ht} \right] \right) \xrightarrow[n \rightarrow \infty]{\mathbf{P}} 0. \quad (11)$$

In particular, at $t = t_{\text{ENT}}$ one has $\log \rho(t) \sim -\log n$. As we will see in Section 3, the lower bound (5) is a rather direct consequence of the concentration result (11). Indeed, we will check that if the invariant probability measure has its atoms $\pi(j), j \in [n]$, roughly of order $\mathcal{O}(1/n)$ then we cannot have reached equilibrium by time t if with high probability $\rho(t) \gg 1/n$. The proof of the upper bound (6) requires a more detailed investigation of the structure of the set of trajectories that the random walker is likely to follow. As explained in Section 4, this allows us to obtain a

sharp comparison between the transition probability $P^t(i, j)$ and the approximate equilibrium $\hat{\pi}(j)$ defined in (7).




FIGURE 1. Distance to equilibrium along time for the $n \times n$ random matrix in Theorem 2, with $\text{Pareto}(\alpha)$ entry distribution, i.e. $\mathbb{P}(w_{ij} > t) = (t \vee 1)^{-\alpha}$. Here, $n = 10^4$ and $\alpha = 0,3$ (red), $\alpha = 0,5$ (blue) and $\alpha = 0,7$ (green). Note that the function h increases continuously from $h(0) = 0$ to $h(\infty) = \infty$: the more “spread-out” the transition probabilities are, the faster the chain mixes.

1.1. Related work. Theorem 1 describes a sharp transition in the approach to equilibrium, visible on Figure 1: the total variation distance drops from the maximal value 1 to the minimal value 0 on a time scale that is asymptotically negligible with respect to the mixing time. This is an instance of the so-called cutoff phenomenon, a remarkable property shared by several models of finite Markov chains. Since its original discovery by Diaconis, Shashahani, and Aldous in the context of card shuffling around 30 years ago [12, 1, 2], the problem of characterizing the Markov chains exhibiting cutoff has attracted much attention. We refer to [10, 3, 16] for an introduction. While the phenomenon is now rather well understood in various specific settings, see e.g. [11, 13] for the case of birth and death chains, a general characterization is still unknown and its nature remains somewhat elusive (but see [4] for an interesting interpretation in the reversible case).

Recently, some attention has shifted from “specific” to “generic” instances: instead of being fixed, the sequence of transition matrices itself is drawn at random from a certain distribution, and the cutoff phenomenon is shown to occur for almost every realization. An important example is that of random walks on random graphs: in their influential paper [18], Lubetzky and Sly established cutoff in random d -regular graphs, for both the simple random walk and the non-backtracking random walk. We refer also to [5] and [6] for important breakthroughs regarding graphs with given degree sequences and the giant component of an Erdős-Renyi random graph.

The above mentioned references are all concerned with the reversible case of undirected graphs, where the associated simple random walk and non-backtracking random walk have explicitly known stationary distributions. In our recent work [8], we investigated the non-reversible case of random walk on sparse directed graphs with given bounded degree sequences. Despite the lack of direct information on the stationary distribution, we obtained a detailed description of the cutoff behavior in such cases. The present paper considerably extends these results by establishing cutoff for a large class of non-reversible sparse stochastic matrices, not necessarily arising as the transition matrix of the random walk on a graph. The proof of our main results here follows a strategy that is closely related to the one we introduced in [8]. However, due to the general assumptions on the transition probabilities, the same combinatorial arguments do not always apply and a finer analysis of the trajectory weights is required.

The eigenvalues and singular values of the random stochastic matrix appearing in Theorem 2 were analyzed very recently in [7]: under a slightly stronger assumption than (9), the associated empirical distributions are shown to converge to some deterministic limits, characterized by a certain recursive distributional equation. The numerical simulations given therein seem to indicate that the spectral gap should also converge to a non-zero limit, and the authors formulate an explicit conjecture (see [7, Remark 1.3]). However, the results in [7] do not allow one to infer something quantitative about the distance to equilibrium. Indeed, the relation between spectrum and mixing for non-reversible chains is rather loose, and one would certainly need more precise information on the structure of the eigenvectors – as done in, e.g., [17]. The proof of Theorem 2 relies entirely on the the general result of Theorem 1 and makes no use of spectral theory. As detailed in Lemma 16 below, the expression (10) for $h(\alpha)$ coincides with the expected value of $-\log \xi$ where ξ has law $\text{Beta}(1 - \alpha, \alpha)$. That should be expected in light of the fact that a size-biased pick from the Poisson Dirichlet law is Beta-distributed [19].

2. QUENCHED LAW OF LARGE NUMBERS FOR PATH WEIGHTS

The main result of this section can be interpreted as a quenched law of large numbers for the logarithm of the total weight of the path followed by the random walk; see Theorem 4 below.

2.1. Uniform unlikeliness. Consider a collection $\sigma = (\sigma_i)_{1 \leq i \leq n}$ of n independent random permutations, referred to as the environment, and a $[n]$ -valued process $X = (X_t)_{t \geq 0}$ whose conditional law, given the environment, is that of a Markov chain with transition matrix (2) and initial law uniform on $[n]$. Our main object of interest will be the sequence of *weights* $W = (W_t)_{t \geq 1}$ seen along the trajectory, and the associated total weight up to time t :

$$W_t := P(X_{t-1}, X_t), \quad \rho(t) := \prod_{s=1}^t W_s. \quad (12)$$

Write Q for the conditional law of the pair (X, W) given the environment. Note that it is a random probability measure on the *trajectory space* $\mathcal{E} = [n]^{\{0,1,\dots\}} \times [0, 1]^{\{1,2,\dots\}}$ equipped with the natural product σ -algebra of events. A generic point of \mathcal{E} will be denoted (x, w) , where $x = (x_0, x_1, x_2, \dots)$ and $w = (w_1, w_2, \dots)$. For example, the trajectorial event “a transition with

weight less than $n^{-\gamma}$ occurs within the first t steps" will be denoted

$$\mathcal{A} = \{(x, w) \in \mathcal{E} : \min(w_1, \dots, w_t) < n^{-\gamma}\}, \quad (13)$$

$$\text{and } Q(\mathcal{A}) = \frac{1}{n} \sum_{i_0 \in [n]} \cdots \sum_{i_t \in [n]} \prod_{s=1}^t P(i_{s-1}, i_s) \left(1 - \prod_{u=1}^t \mathbf{1}_{\{P(i_{u-1}, i_u) \geq n^{-\gamma}\}} \right).$$

We let also $Q_i(\cdot) := Q(\cdot | X_0 = i)$ be the law starting at $i \in [n]$. Recall that all objects are implicitly indexed by the size-parameter n , and asymptotic statements are understood in the $n \rightarrow \infty$ limit. We call a trajectorial event \mathcal{A} *uniformly unlikely* if

$$\max_{i \in [n]} Q_i(\mathcal{A}) \xrightarrow[n \rightarrow \infty]{\mathbf{P}} 0. \quad (14)$$

Lemma 3. *For $t = \mathcal{O}(\log n)$ and $\gamma = \Theta(1)$, the event \mathcal{A} from (13) is uniformly unlikely.*

Proof. A union bound implies the deterministic estimate

$$\max_{i \in [n]} Q_i(\mathcal{A}) \leq t \max_{i \in [n]} \sum_{j=1}^n p_{ij} \mathbf{1}_{(p_{ij} < n^{-\gamma})}.$$

Since $u \mapsto (\log u)^2$ is decreasing on $(0, 1)$,

$$\max_{i \in [n]} Q_i(\mathcal{A}) \leq \frac{t}{(\gamma \log n)^2} \max_{i \in [n]} \sum_{j=1}^n p_{ij} (\log p_{ij})^2.$$

The conclusion follows from the assumption (3). \square

Our main task in the rest of this section will be to establish:

Theorem 4 (Trajectories of length t have weight roughly e^{-Ht}). *For $t = \Theta(\log n)$ and fixed $\varepsilon > 0$, the event $\{\rho(t) \notin [e^{-(1+\varepsilon)Ht}, e^{-(1-\varepsilon)Ht}]\}$ is uniformly unlikely.*

Let us observe here, for future reference, that if $\theta: \mathcal{E} \rightarrow \mathcal{E}$ is the operator that shifts $x = (x_0, x_1, \dots)$ and $w = (w_1, w_2, \dots)$ to $x' = (x_1, x_2, \dots)$ and $w' = (w_2, w_3, \dots)$ respectively, then, for any $i \in [n]$, $t \in \mathbb{N}$ and any event $\mathcal{A} \subset \mathcal{E}$

$$Q_i(\theta^{-t}\mathcal{A}) = \sum_{j \in [n]} P^t(i, j) Q_j(\mathcal{A}) \leq \max_{j \in [n]} Q_j(\mathcal{A}), \quad (15)$$

where $\theta^{-t}\mathcal{A} = \{(x, w) \in \mathcal{E} : \theta^t(x, w) \in \mathcal{A}\}$. Thus, uniform unlikeliness propagates through time.

2.2. Sequential generation. By averaging the quenched probability $Q(\cdot)$ with respect to the environment, one obtains the so-called *annealed* probability, which we denote by \mathbb{P} . In symbols, letting \mathbb{E} denote the associated expectation, for any event \mathcal{A} in the trajectory space:

$$\mathbb{E}[Q(\mathcal{A})] = \frac{1}{n} \sum_{i=1}^n \mathbb{E}[Q_i(\mathcal{A})] = \mathbb{P}((X, W) \in \mathcal{A}).$$

Markov's inequality offers a way to reduce the problem of estimating the worst-case quenched probability $\max_{i \in [n]} Q_i(\mathcal{A})$ of a trajectorial event $\mathcal{A} \subset \mathcal{E}$ to that of controlling the corresponding annealed quantity, at the cost of an extra factor of n : for any $\delta > 0$,

$$\mathbb{P}\left(\max_{i \in [n]} Q_i(\mathcal{A}) > \delta\right) \leq \frac{1}{\delta} \mathbb{E}\left[\sum_{i=1}^n Q_i(\mathcal{A})\right] = \frac{n}{\delta} \mathbb{P}((X, W) \in \mathcal{A}). \quad (16)$$

The analysis of the right-hand side may often be simplified by the observation that the pair (X, W) can be constructed sequentially, together with the underlying environment σ , as follows:

initially, $\text{Dom}(\sigma_i) = \text{Ran}(\sigma_i) = \emptyset$ for all $i \in [n]$, and X_0 is drawn uniformly from $[n]$; then for each $t \geq 1$,

#1. Set $i = X_{t-1}$ and draw an index $j \in [n]$ at random with probability p_{ij} .

#2. If $j \notin \text{Dom}(\sigma_i)$, then extend σ_i by setting $\sigma_i(j) = k$, where k is uniform in $[n] \setminus \text{Ran}(\sigma_i)$.

#3. In either case, $\sigma_i(j)$ is now well defined: set $X_t = \sigma_i(j)$ and $W_t = p_{ij}$.

Let us illustrate the strength of this sequential construction on an important trajectorial feature. A path $(x_0, \dots, x_t) \in [n]^{t+1}$ naturally induces a directed graph with vertex set $V = \{x_0, \dots, x_t\} \subset [n]$ and edge set $E = \{(x_0, x_1), \dots, (x_{t-1}, x_t)\} \subset [n] \times [n]$. As a rule, below we neglect possible multiplicities in the edge set E , that is every repeated edge from the path appears only once in E . We define the *tree-excess* of the path (x_0, \dots, x_t) as

$$\text{TX}(x_0, \dots, x_t) = 1 + |E| - |V|.$$

Here $|V|$ and $|E|$ denote the cardinalities of V and E . In particular, $\text{TX}(x_0, \dots, x_t) = 0$ if and only if (x_0, \dots, x_t) is a *simple path* in the usual graph-theoretical sense, while $\text{TX}(x_0, \dots, x_t) = 1$ if and only if the edge set of (x_0, \dots, x_t) has a single cycle (the path may turn around it more than once).

Lemma 5 (Tree-excess). *For $t = o(n^{1/4})$, $\{\text{TX}(X_0, \dots, X_t) \geq 2\}$ is uniformly unlikely.*

Proof. In the sequential generation process, we have $\text{TX}(X_0, \dots, X_t) = \xi_1 + \dots + \xi_t$, where $\xi_t \in \{0, 1\}$ indicates whether or not, during the t^{th} iteration, the random index k appearing at line #2 is actually drawn and satisfies $k \in \{X_0, \dots, X_{t-1}\}$. The conditional chance of this, given the past, is at most

$$\frac{|\{X_0, \dots, X_{t-1}\}| - |\text{Ran}(\sigma_i)|}{n - |\text{Ran}(\sigma_i)|} \leq \frac{t}{n}.$$

Thus, $\text{TX}(X_0, \dots, X_t)$ is stochastically dominated by a Binomial $(t, \frac{t}{n})$. In particular, for $r \in \mathbb{N}$,

$$\mathbb{P}(\text{TX}(X_0, \dots, X_t) \geq r) \leq \binom{t}{r} \left\{ \frac{t}{n} \right\}^r \leq \frac{1}{r!} \left\{ \frac{t^2}{n} \right\}^r. \quad (17)$$

Now, let $n \rightarrow \infty$: since $t = o(n^{1/4})$, the right-hand side is $o(\frac{1}{n})$ already for $r = 2$, and (16) concludes. \square

2.3. Approximation by i.i.d. samples. Consider the modified process (X^*, W^*) obtained by resetting $\text{Dom}(\sigma_i) = \text{Ran}(\sigma_i) = \emptyset$ before every execution of line #2, thereby suppressing any time dependency: the environment is locally regenerated *afresh* at each step. In particular, the pairs $(X_{t-1}^*, W_t^*)_{t \geq 1}$ are i.i.d. with law

$$\mathbb{P}(X_0^* = i, W_1^* \geq t) = \sum_{j=1}^n \frac{p_{ij}}{n} \mathbf{1}_{p_{ij} \geq t} \quad (1 \leq i \leq n, t \geq 0). \quad (18)$$

By construction, the modified process and the original one can be coupled in such a way that they coincide until the time

$$T := \inf\{t \geq 0: \text{TX}(X_0, \dots, X_t) = 1\}, \quad (19)$$

that is the first time a state gets visited for the second time. Thus, on the event $\{T \geq t\}$,

$$(X_0^*, \dots, X_t^*) = (X_0, \dots, X_t) \quad \text{and} \quad (W_1^*, \dots, W_t^*) = (W_1, \dots, W_t). \quad (20)$$

We exploit this observation to establish a preliminary step towards Theorem 4. Notice that the estimate below becomes trivial if the parameters p_{ij} are such that $p_{ij} \leq 1 - \varepsilon$ for some fixed $\varepsilon > 0$. In the general case it relies on the non-degeneracy assumption (4).

Lemma 6. *If $t = \Theta(\log n)$ and $\delta = o(1)$, then $\{\rho(t) > e^{-\delta t}\}$ is uniformly unlikely.*

Proof. Call (x_0, \dots, x_t) a *cycle* if (x_0, \dots, x_{t-1}) is simple and $x_t = x_0$. We will show:

- (i) for t and δ as above, $\mathcal{B} := \{\rho(t) > e^{-\delta t}, \text{TX}(X_0, \dots, X_t) = 0\}$ is uniformly unlikely;
- (ii) for $\delta = o(1)$, $\mathcal{C}_\delta := \{\exists s \geq 1: (X_0, \dots, X_s) \text{ is a cycle}, \rho(s) > e^{-\delta s}\}$ is uniformly unlikely.

Let us first show that this is sufficient to conclude the proof. Indeed, the event

$$\mathcal{A} := \{\rho(t) > e^{-\delta t}\} = \{(x, w) \in \mathcal{E} : w_1 \cdots w_t > e^{-\delta t}\},$$

can be partitioned according to the size of $\text{TX}(x_0, \dots, x_t)$. Therefore

$$\mathcal{A} \subset \mathcal{B} \cup \{\text{TX}(x_0, \dots, x_t) = 1, w_1 \cdots w_t > e^{-\delta t}\} \cup \{\text{TX}(x_0, \dots, x_t) \geq 2\}.$$

The event $\{\text{TX}(x_0, \dots, x_t) \geq 2\}$ is uniformly unlikely, thanks to Lemma 5. The event \mathcal{B} is uniformly unlikely, by (i) above. The event $\{\text{TX}(x_0, \dots, x_t) = 1, w_1 \cdots w_t > e^{-\delta t}\}$ on the other hand is contained in the union of the following three events:

- $\{\text{TX}(x_0, \dots, x_{\lfloor t/3 \rfloor}) = 0, w_1 \cdots w_{\lfloor t/3 \rfloor} > e^{-\delta t}\}$
- $\{\text{TX}(x_{\lceil 2t/3 \rceil}, \dots, x_t) = 0, w_{\lceil 2t/3 \rceil} \cdots w_t > e^{-\delta t}\}$
- $\{\text{TX}(x_0, \dots, x_t) = \text{TX}(x_0, \dots, x_{\lfloor t/3 \rfloor}) = \text{TX}(x_{\lceil 2t/3 \rceil}, \dots, x_t) = 1, w_1 \cdots w_t > e^{-\delta t}\}$

The first two cases are uniformly unlikely by (i) and by the observation (15). To handle the third case, observe that if $\text{TX}(x_0, \dots, x_t) = 1$, then the path (x_0, \dots, x_t) can be rewritten as $(x_0, \dots, x_a, \dots, x_{a+r\ell}, \dots, x_t)$, where (x_0, \dots, x_a) and $(x_{a+r\ell}, \dots, x_t)$ are simple paths, while the path $(x_a, \dots, x_{a+r\ell})$ consists of r complete turns around a cycle of length ℓ . Here $a \geq 0$, $r, \ell \geq 1$ and $a + r\ell \leq t$. If $\text{TX}(x_0, \dots, x_{\lfloor t/3 \rfloor}) = \text{TX}(x_{\lceil 2t/3 \rceil}, \dots, x_t) = 1$, then the two simple paths must have lengths less than $t/3$ and therefore $r\ell > t/3$. If $\rho = w_{a+1} \cdots w_{a+\ell}$ is the weight associated to one turn around the cycle, then $w_1 \cdots w_t > e^{-\delta t}$ implies $\rho^r > e^{-\delta t}$ and therefore $\rho > e^{-3\delta\ell}$. It follows that the shifted trajectory $\theta^{\lfloor t/3 \rfloor}(x, w)$ must belong to $\mathcal{C}_{3\delta}$. Using (15) and (ii) above, this is uniformly unlikely.

It remains to prove (i) and (ii). By (20), we have

$$\{(X, W) \in \mathcal{B}\} \subset \{W_1^* \cdots W_t^* > e^{-\delta t}\} \subset \left\{ \sum_{s=1}^t \mathbf{1}_{(W_s^* < e^{-2\delta})} < \frac{t}{2} \right\}.$$

Now $\sum_{s=1}^t \mathbf{1}_{(W_s^* < e^{-2\delta})}$ is Binomial(t, q) with $q = \mathbb{P}(W_1^* < e^{-2\delta})$. Thus, Bennett's inequality yields

$$\mathbb{P}((X, W) \in \mathcal{B}) \leq e^{-t\phi(q)}, \tag{21}$$

for some universal function $\phi: [0, 1] \rightarrow \mathbb{R}_+$ that diverges at 1^- (more precisely, [9, Theorem 2.9] gives $\phi(q) = \sigma^2 h((q - 1/2)/\sigma^2)$ with $\sigma^2 = q(1 - q)$, $h(x) = (x + 1) \log(x + 1) - x$ for $x \geq 0$ and 0 otherwise). From (18),

$$1 - q = \mathbb{P}(W_1^* \geq e^{-2\delta}) = \frac{1}{n} \sum_{i,j=1}^n p_{ij} \mathbf{1}_{p_{ij} \geq e^{-2\delta}}.$$

Now, let $n \rightarrow \infty$. Since $\delta \rightarrow 0$, the assumption (4) ensures that $q \rightarrow 1$, so that (21) implies $\mathbb{P}((X, W) \in \mathcal{B}) = o(\frac{1}{n})$. From the first moment argument (16) one obtains part (i).

To prove part (ii), observe that the coupling (20) implies

$$\{(X, W) \in \mathcal{C}_\delta\} \subset \bigcup_{s \geq 1} \{W_1^* \cdots W_s^* > e^{-\delta s}, X_s^* = X_0^*\}.$$

Since X_s^* is independent of the other variables and uniform, the argument for (21) shows that

$$\mathbb{P}((X, W) \in \mathcal{C}_\delta) \leq \frac{1}{n} \sum_{s \geq 1} e^{-s\phi(q)} = \frac{1}{n(e^{\phi(q)} - 1)}. \quad (22)$$

Letting $n \rightarrow \infty$, the conclusion follows as above. \square

2.4. Proof of Theorem 4. The event $\mathcal{A} = \{\rho(t) \notin [e^{-(1+\varepsilon)Ht}, e^{-(1-\varepsilon)Ht}]\}$ can be written as

$$\mathcal{A} = \left\{ \left| 1 - \frac{1}{Ht} \sum_{s=1}^t \log \frac{1}{W_s} \right| > \varepsilon \right\}.$$

We are going to prove the uniform unlikeliness of \mathcal{A} for any fixed $\varepsilon > 0$ and $t = \Theta(\log n)$. First note that, by (17), the random time T defined in (19) satisfies

$$\mathbb{P}(T \leq t) \leq \frac{t^2}{n} = o(1).$$

Combining this with (20), we see that

$$\mathbb{P}((X, W) \in \mathcal{A}) = \mathbb{P} \left(\left| 1 - \frac{1}{Ht} \sum_{s=1}^t \log \frac{1}{W_s^*} \right| > \varepsilon \right) + o(1), \quad (23)$$

where (W_1^*, \dots, W_t^*) are i.i.d. with law determined by (18). Now, the variable $Y := \frac{1}{H} \log \frac{1}{W_1^*}$ has mean 1 by definition of H . From (18) and the assumption (3), one has $\mathbb{E}[Y^2] = o(\log n)$. In particular, the variance of Y satisfies $\text{Var}(Y) = o(\log n)$. Therefore,

$$\mathbb{E} \left[\left(1 - \frac{1}{Ht} \sum_{s=1}^t \log \frac{1}{W_s^*} \right)^2 \right] = \frac{1}{t} \text{Var}(Y) \xrightarrow{n \rightarrow \infty} 0. \quad (24)$$

By Chebychev’s inequality, (24) and (23) already show that $\mathbb{P}((X, W) \in \mathcal{A}) \rightarrow 0$. However, this is not enough to guarantee the uniform unlikeliness of \mathcal{A} , due to the extra factor n appearing on the RHS of (16). To overcome this difficulty, we will use a more elaborate approach, based on the following higher-order version of (16). For any event \mathcal{B} in the trajectory space, for any $\delta > 0$ and $k \in \mathbb{N}$,

$$\mathbb{P} \left(\max_{i \in [n]} Q_i(\mathcal{B}) > \delta \right) \leq \frac{1}{\delta^k} \mathbb{E} \left[\sum_{i=1}^n (Q_i(\mathcal{B}))^k \right] = \frac{n}{\delta^k} \mathbb{P} \left(\bigcap_{\ell=1}^k \{(X^\ell, W^\ell) \in \mathcal{B}\} \right), \quad (25)$$

where the processes $(X^1, W^1), \dots, (X^k, W^k)$ are formed by generating a random environment σ and a uniform state $\mathcal{I} \in [n]$, and conditionally on that, by running k independent P_σ -Markov chains in the same environment σ , with the same starting node \mathcal{I} . We will fix $\delta > 0$ and prove that for suitable choices of the event \mathcal{B} , the right-hand side of (25) is $o(1)$ for

$$k := \left\lfloor \frac{\log n}{2 \log(1/\delta)} \right\rfloor. \quad (26)$$

First observe that the variables $(X_s^1, W_s^1)_{0 \leq s \leq t}, \dots, (X_s^k, W_s^k)_{0 \leq s \leq t}$ can again be constructed sequentially, together with σ : pick \mathcal{I} uniformly in $[n]$, set $X_0^1 = \mathcal{I}$, and construct $(X_s^1, W_s^1)_{1 \leq s \leq t}$ by repeating t times the instructions #1, #2 and #3 of subsection 2.2. Then set $X_0^2 = \mathcal{I}$, construct $(X_s^2, W_s^2)_{1 \leq s \leq t}$ similarly (without re-initializing the environment), and so on. Note that kt iterations are performed in total. The union of the graphs induced by the first j paths

$(X_1^\ell, \dots, X_t^\ell)$, $\ell = 1, \dots, j$, forms a certain graph $G_j = (V_j, E_j)$, and the argument used for Lemma 5 shows that $\text{TX}(G_j) := 1 + |E_j| - |V_j|$ satisfies

$$\mathbb{P}(\text{TX}(G_j) \geq 2) \leq \frac{(kt)^4}{2n^2} = o\left(\frac{\delta^k}{n}\right),$$

where we use the fact that by (26) one has $\delta^k = \Theta(n^{-1/2})$. In view of (25), this reduces our task to showing that $\mathbb{P}(B_k) = o\left(\frac{\delta^k}{n}\right)$, where, for any $j = 1, \dots, k$, we define the event

$$B_j := \{\text{TX}(G_j) \leq 1\} \cap \{(X^1, W^1) \in \mathcal{B}\} \cap \dots \cap \{(X^j, W^j) \in \mathcal{B}\}. \quad (27)$$

Note that $B_k \subset B_{k-1} \dots \subset B_1$. We will actually show that $\mathbb{P}(B_j|B_{j-1}) = o(1)$ uniformly in $2 \leq j \leq k$ and that $\mathbb{P}(B_1) = o(1)$. This will be enough to conclude, since for $k = \Theta(\log n)$ one has

$$\mathbb{P}(B_k) = \mathbb{P}(B_1) \prod_{j=2}^k \mathbb{P}(B_j|B_{j-1}) = o\left(\frac{\delta^k}{n}\right).$$

To prove Theorem 4 we now apply the above strategy with two choices of the event \mathcal{B} .

Uniform unlikeliness of $\{\rho(t) < e^{-(1+\varepsilon)Ht}\}$. Define the event

$$\mathcal{B} := \left\{W_1 \dots W_t < e^{-(1+\varepsilon)Ht}\right\} \cap \left\{\min(W_1, \dots, W_t) > n^{-\gamma}\right\}, \quad \gamma := \frac{\varepsilon t}{4t_{\text{ENT}}}.$$

We use the method described above, i.e., we prove that $\mathbb{P}(B_1) = o(1)$ and

$$\mathbb{P}(B_\ell|B_{\ell-1}) = o(1), \quad (28)$$

uniformly in $2 \leq \ell \leq k$, with k given by (26) and B_k defined as in (27). Notice that once (28) has been proved, the previous observations together with Lemma 3 imply that the event $\{\rho(t) < e^{-(1+\varepsilon)Ht}\}$ is uniformly unlikely, thus establishing one half of Theorem 4.

To prove (28), first observe that $\mathbb{P}(B_1)$ is bounded from above by (23), so that $\mathbb{P}(B_1) = o(1)$ follows from (24). Next, fix $2 \leq \ell \leq k$, assume that the first $\ell - 1$ walks have already been sequentially generated and that $B_{\ell-1}$ holds, and let us evaluate the conditional probability that $(X^\ell, W^\ell) \in \mathcal{A}$. We distinguish between two scenarios, depending on the random times

$$\tau := \inf \left\{s \geq 1: (X_{s-1}^\ell, X_s^\ell) \notin E_{\ell-1}\right\} \quad \text{and} \quad \tau' := \inf \left\{s \geq 0: W_1^\ell \dots W_s^\ell \leq n^{-\gamma}\right\}.$$

Since $n^{-\gamma} = e^{-\varepsilon Ht/4}$, we may clearly restrict to the case $t \geq \tau'$, otherwise the event $\rho(t) < e^{-(1+\varepsilon)Ht}$ is trivially false.

Case I: $\tau' < \tau$ and $t \geq \tau'$. Let F denote the event $\{\tau' < \tau\} \cap \{t \geq \tau'\}$. We show that $\mathbb{P}(F|B_{\ell-1}) = o(1)$. For any $1 \leq s \leq t$, let \mathcal{G}_s denote the set of directed paths in the graph $G_{\ell-1}$, with length s and starting node \mathcal{I} . The condition $\text{TX}(G_{\ell-1}) \leq 1$ ensures that $G_{\ell-1}$ is a directed tree with at most one extra edge. Thus, for every vertex $v \in V_{\ell-1}$ there are at most 2 directed paths of length s from the given vertex \mathcal{I} to v . It follows that $|\mathcal{G}_s| \leq 2|V_{\ell-1}| \leq 2kt$. If F holds, and $\tau' = s$, then $(X_0^\ell, \dots, X_s^\ell)$ is one of the paths in \mathcal{G}_s with weight at most $n^{-\gamma}$. By definition, each such path has conditional probability at most $n^{-\gamma}$ to be actually followed by the ℓ th walk. Summing over the possible values of τ' , we find that the conditional probability of F is less than $2kt^2 n^{-\gamma} = o(1)$.

Case II: $\tau \leq \tau' \leq t$. Let F' denote the event $\{\tau \leq \tau' \leq t\}$. We show that $\mathbb{P}(B_\ell \cap F'|B_{\ell-1}) = o(1)$. On the event F' one has $W_1^\ell \dots W_{\tau-1}^\ell > n^{-\gamma}$. Since \mathcal{B} includes the condition $\min(W_1, \dots, W_t) > n^{-\gamma}$, and therefore $W_\tau > n^{-\gamma}$, for (X^ℓ, W^ℓ) to fall in \mathcal{B} we must have

$$W_{\tau+1}^\ell \dots W_t^\ell < n^{2\gamma} e^{-H(1+\varepsilon)t} = e^{-H(1+\frac{\varepsilon}{2})t}. \quad (29)$$

Now, the condition $j \notin \text{Dom}(\sigma_i)$ in line #2 of the sequential generation process is actually satisfied when the ℓ th walk exits $G_{\ell-1}$, so X_τ is constructed by sampling $\sigma_i(j)$ uniformly in $[n] \setminus \text{Ran}(\sigma_i)$. Since $\sum_i |\text{Ran}(\sigma_i)| \leq kt$, this random choice and the subsequent ones can be coupled with i.i.d. samples from the uniform law on $[n]$ at a total-variation cost less than $\frac{kt^2}{n} = o(1)$. This induces a coupling between $W_{\tau+1}^\ell \cdots W_t^\ell$ and a product of (less than t) i.i.d. variables with law (18), and it follows from (23)-(24) that (29) occurs with probability $o(1)$.

Uniform unlikeliness of $\{\rho(t) > e^{-(1-\varepsilon)Ht}\}$. Let us define the event

$$\mathcal{B} := \left\{ W_1 \cdots W_t > e^{-(1-\varepsilon)Ht} \right\} \cap \left\{ W_1 \cdots W_s \leq (\log n)^{-4} \right\}, \quad s := \left\lfloor \frac{\varepsilon t}{2 - \varepsilon} \right\rfloor.$$

We use the same method as above, with this new definition of \mathcal{B} . Notice that if we prove that \mathcal{B} is uniformly unlikely, then it follows from Lemma 6 that $\{\rho(t) > e^{-(1-\varepsilon)Ht}\}$ is also uniformly unlikely, thus completing the proof Theorem 4.

We need to prove (28) with the current definition of the sets B_j ; see (27). First observe that $\mathbb{P}(B_1) = o(1)$ follows again as in (23)-(24). Next, fix $2 \leq \ell \leq k$, assume that the first $\ell - 1$ walks have already been sequentially generated and that $B_{\ell-1}$ holds, and let us evaluate the conditional probability that $(X^\ell, W^\ell) \in \mathcal{B}$. As before, we let τ be the first exit from $G_{\ell-1}$. We distinguish two cases.

Case I: $\tau > s$. We proceed as in case I above. If $B_\ell \cap \{\tau > s\}$ holds, then (X_0, \dots, X_s) must be one of the paths in the set \mathcal{G}_s , with weight at most $(\log n)^{-4}$. As before, there are less than $2kt$ possible paths, each having conditional probability at most $(\log n)^{-4}$ to be actually followed. Therefore, $\mathbb{P}(B_\ell \cap \{\tau > s\} | B_{\ell-1}) \leq 2kt(\log n)^{-4} = o(1)$.

Case II: $\tau \leq s$. On this event, reasoning as in case II above, one sees that $(W_{s+1}^\ell, \dots, W_t^\ell)$ can be coupled with $(t - s)$ i.i.d. variables with law (18) with an error $o(1)$ in total variation, and (23)-(24) then implies that their product will be below $e^{-(1-\frac{\varepsilon}{2})H(t-s)}$ with probability $1 - o(1)$. But $e^{-(1-\frac{\varepsilon}{2})H(t-s)} \leq e^{-(1-\varepsilon)Ht}$ by our choice of s .

3. PROOF OF THE LOWER BOUND IN THEOREM 1

In this section we prove the simpler half of Theorem 1, namely the lower bound (5). We shall actually prove (5) with π replaced by the probability $\hat{\pi}$ given in (7); see Remark 1.

Fix the environment σ , an arbitrary probability measure ν on $[n]$, $t \in \mathbb{N}$, $\theta \in (0, 1)$ and $i, j \in [n]$. Since $P^t(i, j) = Q_i(X_t = j)$, we have

$$P^t(i, j) \geq Q_i(X_t = j, \rho(t) \leq \theta). \quad (30)$$

If equality holds in this inequality, then clearly

$$\nu(j) - Q_i(X_t = j, \rho(t) \leq \theta) \leq [\nu(j) - P^t(i, j)]_+,$$

where $[x]_+ := \max(x, 0)$. On the other-hand, if the inequality (30) is strict, then there must exist a path of length t from i to j with weight $> \theta$, implying that $P^t(i, j) > \theta$ and hence that

$$\nu(j) - Q_i(X_t = j, \rho(t) \leq \theta) \leq \nu(j) \mathbf{1}_{P^t(i, j) > \theta}.$$

In either case, we have

$$\nu(j) - Q_i(X_t = j, \rho(t) \leq \theta) \leq [\nu(j) - P^t(i, j)]_+ + \nu(j) \mathbf{1}_{P^t(i, j) > \theta}.$$

Summing over all $j \in [n]$, the left hand side above yields the probability $Q_i(\rho(t) > \theta)$, while the first term in the right hand side gives the total variation norm $\|\nu - P^t(i, \cdot)\|_{\text{TV}}$. On the other hand, the Cauchy–Schwarz and Markov inequalities imply

$$\left(\sum_{j \in [n]} \nu(j) \mathbf{1}_{P^t(i, j) > \theta} \right)^2 \leq \sum_{j \in [n]} \nu(j)^2 \sum_{\ell \in [n]} \mathbf{1}_{P^t(i, \ell) > \theta} \leq \frac{1}{\theta} \sum_{j \in [n]} \nu(j)^2.$$

Summarizing,

$$Q_i(\rho(t) > \theta) \leq \|\nu - P^t(i, \cdot)\|_{\text{TV}} + \sqrt{\frac{1}{\theta} \sum_{j \in [n]} \nu(j)^2}. \quad (31)$$

We now specialize to $\theta = \frac{\log^3 n}{n}$ and $\nu = \hat{\pi}$ as in (7). If $t = (\lambda + o(1))t_{\text{ENT}}$ with $0 < \lambda < 1$ fixed, then for some $\varepsilon > 0$ one has $e^{-(1+\varepsilon)Ht} > \theta$ for all n large enough. Therefore, from Theorem 4, we have

$$\min_{i \in [n]} Q_i(\rho(t) > \theta) \xrightarrow[n \rightarrow \infty]{\mathbf{P}} 1.$$

To conclude the proof, it remains to verify that the square-root term in (31) converges to zero in probability. Below, we prove the stronger estimate

$$\mathbb{E} \left[\sum_{j \in [n]} \hat{\pi}(j)^2 \right] = o(\theta). \quad (32)$$

Fix $h := \lfloor \frac{t_{\text{ENT}}}{10} \rfloor$. The left-hand-side of (32) may be rewritten as $\mathbb{P}(X_h = Y_h)$, where conditionally on the environment σ , X and Y denote two independent P_σ –Markov chains, each starting from the uniform distribution on $[n]$. To evaluate this annealed probability, we generate the chains sequentially, together with the environment, as follows: we pick X_0 uniformly in $[n]$, and construct (X_1, \dots, X_h) by repeating t times the instructions #1, #2 and #3 of subsection 2.2. We then pick Y_0 uniformly at random in $[n]$, and construct (Y_1, \dots, Y_h) similarly, without re-initializing the environment. Now, observe that $\{X_h = Y_h\} \subset \{S \leq h\}$, where

$$S = \inf \{s \geq 0: Y_s \in \{X_0, \dots, X_h, Y_0, \dots, Y_{s-1}\}\}.$$

By uniformity of the random choices made at each execution of the instruction #2, we have for $0 \leq s \leq h$,

$$\mathbb{P}(S = s) \leq \frac{|\{X_0, \dots, X_h, Y_0, \dots, Y_{s-1}\}|}{n} \leq \frac{2h + 1}{n}.$$

By a union bound, we see that $\mathbb{P}(S \leq h) \leq \frac{2(h+1)^2}{n}$, which is $o(\theta)$ thanks to our choice of θ .

4. PROOF OF THE UPPER BOUND IN THEOREM 1

The overall strategy of the proof is similar to that introduced in [8]. Before entering the details of the proof, let us give a brief overview of the main steps involved.

Fix the environment and, for every $i, j \in [n]$, define a suitable set of *nice* paths $\mathcal{N}_t(i, j)$ that go from i to j in t steps, where $t = (\lambda + o(1))t_{\text{ENT}}$, with $\lambda > 1$. Call $P_0^t(i, j)$ the probability that the walk started at i arrives in j after t steps by following one of the paths in $\mathcal{N}_t(i, j)$. Clearly, $P_0^t(i, j) \leq P^t(i, j)$, and therefore, for any probability ν on $[n]$, any $\delta > 0$, one has

$$\|\nu - P^t(i, \cdot)\|_{\text{TV}} = \sum_{j \in [n]} [\nu(j) - P^t(i, j)]_+ \leq \sum_{j \in [n]} \left[\nu(j)(1 + \delta) + \frac{\delta}{n} - P_0^t(i, j) \right]_+. \quad (33)$$

Suppose now that, for some $\delta > 0$, and some ν , for all $i, j \in [n]$, one has

$$P_0^t(i, j) \leq (1 + \delta)\nu(j) + \frac{\delta}{n}. \quad (34)$$

In this case we can compute the sum in (33) to obtain, for all $i \in [n]$,

$$\|\nu - P^t(i, \cdot)\|_{\text{TV}} \leq q(i) + 2\delta, \quad (35)$$

where $q(i)$ is the probability that a walk of length t started at i does not follow one of the nice paths in $\mathcal{N}_t(i) = \cup_j \mathcal{N}_t(i, j)$, i.e.

$$q(i) = \sum_{j \in [n]} (P^t(i, j) - P_0^t(i, j)).$$

We want to prove that $\|\nu - P^t(i, \cdot)\|_{\text{TV}} \xrightarrow{\mathbf{P}} 0$ when $\nu = \hat{\pi}$; see Remark 1. Thus, roughly speaking, the key to the proof is to define the set of nice paths $\mathcal{N}_t(i, j)$ in such a way that:

- (1) $q(i)$ vanishes in probability, and
- (2) for any $\delta > 0$ the bound (34) holds with high probability if we choose $\nu = \hat{\pi}$.

The definition of nice paths will be given in Subsection 4.2. Below, we start with some preliminary facts. Throughout this section we will use the following notation.

Notation. We fix $0 < \varepsilon < 1/20$, and

$$t := (1 + \varepsilon + o(1)) t_{\text{ENT}}, \quad (36)$$

Moreover, we set

$$h := \left\lfloor \frac{t_{\text{ENT}}}{10} \right\rfloor, \quad \underline{H} = H(1 - \frac{\varepsilon}{2}) \quad \text{and} \quad \bar{H} = H(1 + \varepsilon). \quad (37)$$

For any path $\mathbf{p} := (x_0, \dots, x_s) \in [n]^{s+1}$, $s \in \mathbb{N}$, the weight of \mathbf{p} is defined by

$$w(\mathbf{p}) = P(x_0, x_1) \cdots P(x_{s-1}, x_s). \quad (38)$$

4.1. The forward graph $\mathcal{G}_x(s)$ and the spanning tree $\mathcal{T}_x(s)$. For integer $s \geq 1$ and $x \in [n]$ we call $\mathcal{G}_x(s)$ the weighted directed graph spanned by the set of directed paths \mathbf{p} with at most s edges, starting at x , and with weight $w(\mathbf{p}) \geq e^{-\bar{H}s}$. We can construct $\mathcal{G}_x(s)$, together with a spanning tree $\mathcal{T}_x(s)$, as follows. We start at $\mathcal{G}^0 = \mathcal{T}^0 = x$ and define a process $(\mathcal{G}^0, \mathcal{T}^0), (\mathcal{G}^1, \mathcal{T}^1), \dots$, which stops at some random time κ , and we define $\mathcal{G}_x(s) = \mathcal{G}^\kappa$ and $\mathcal{T}_x(s) = \mathcal{T}^\kappa$. As in Subsection 2.2, we will add oriented edges one by one, using sequential generation. Initially, $\text{Dom}(\sigma_y) = \text{Ran}(\sigma_y) = \emptyset$ for all $y \in [n]$. When $j \notin \text{Dom}(\sigma_y)$, we interpret (y, j) as a free half-edge exiting y to be matched with a free half-edge z to be chosen uniformly among the vertices $z \in [n] \setminus \text{Ran}(\sigma_y)$. If we are at $(\mathcal{G}^\ell, \mathcal{T}^\ell)$, to obtain $(\mathcal{G}^{\ell+1}, \mathcal{T}^{\ell+1})$ the iterative step is as follows:

- 1) Consider all nodes y of \mathcal{G}^ℓ together with their available half-edges (y, j) , $j \notin \text{Dom}(\sigma_y)$. The weight of such half-edge is defined as

$$\hat{w}(y, j) := w(\mathbf{p}) p_{y,j},$$

where \mathbf{p} is the unique path in \mathcal{T}^ℓ from x to y . Pick (y, j) with maximal weight $\hat{w}(y, j)$, among all available half-edges such that: (i) y is at graph distance at most $s - 1$ from x , and (ii) the weight satisfies $\hat{w}(y, j) \geq e^{-\bar{H}s}$. If this set is empty, then the process stops and we set $\kappa = \ell$.

- 2) Extend σ_y by setting $\sigma_y(j) = z$, where z is uniform in $[n] \setminus \text{Ran}(\sigma_y)$.
- 3) Add the weighted directed edge (y, z) , with weight $p_{y,z}$, to the graph \mathcal{G}^ℓ ; add it also to \mathcal{T}^ℓ if z was not already a vertex of \mathcal{G}^ℓ . This defines $\mathcal{T}^{\ell+1}$ and $\mathcal{G}^{\ell+1}$.

Notice that $\mathcal{T}_x(s)$ is a spanning tree of $\mathcal{G}_x(s)$, and that $\mathcal{G}_x(s)$ is indeed the union of all directed paths p with at most s edges, starting at x , and such that $w(p) \geq e^{-\bar{H}s}$. We start our analysis of $\mathcal{G}_x(s)$ and $\mathcal{T}_x(s)$ with a deterministic lemma.

Lemma 7. *Fix $x \in [n]$ and $s \in \mathbb{N}$, and consider the generation process defined above. The weight \hat{w}_ℓ of the half-edge picked at the ℓ -th iteration of step 1 satisfies*

$$\hat{w}_\ell \leq \frac{s}{\ell}.$$

In particular, the random time κ satisfies

$$\kappa \leq s e^{\bar{H}s}.$$

Proof. Consider the following new tree, say $\tilde{\mathcal{T}}^\ell$, obtained as \mathcal{T}^ℓ in the above process except that at step 3 if z has already been seen, we create anyway a new fictitious leaf node. Then both \mathcal{G}^ℓ and $\tilde{\mathcal{T}}^\ell$ have exactly ℓ edges. Let F denote the set of all leaf nodes $\tilde{\mathcal{T}}^\ell$. Thus F consists of all leaf nodes of \mathcal{T}^ℓ plus all the fictitious leaf nodes introduced above. By construction:

$$\sum_{p: x \rightarrow F} w(p) \leq 1, \quad (39)$$

where the sum runs over all directed paths in $\tilde{\mathcal{T}}^\ell$ from the root x to a leaf node in F . Note also that the chosen weights at step 1 for $\ell = 1, 2, \dots$ are non-increasing: $\hat{w}_{\ell-1} \geq \hat{w}_\ell$. Hence, any p from the sum in (39) satisfies $w(p) \geq \hat{w}_\ell$. Since there is a unique path p for each leaf node in F , it follows from (39) that $|F|\hat{w}_\ell \leq 1$. Each path p has length at most s , and their union spans $\tilde{\mathcal{T}}^\ell$. Since there are a total of ℓ edges one must have $\ell \leq s|F|$. Therefore $\ell \leq s/\hat{w}_\ell$ as desired. For the second statement, we use that for $\ell = \kappa$, $\hat{w}_\ell \geq e^{-\bar{H}s}$. \square

Let as usual $\text{TX}(\mathcal{G}_x(s)) := 1 + |E| - |V|$ denote the tree excess of the directed graph $\mathcal{G}_x(s)$, where E is the set of edges and V is the set of vertices of $\mathcal{G}_x(s)$. Note that $|E| = \kappa$, that $\text{TX}(\mathcal{G}_x(s)) = 0$ iff $\mathcal{G}_x(s) = \mathcal{T}_x(s)$, and that $\text{TX}(\mathcal{G}_x(s)) \leq 1$ iff $\mathcal{G}_x(s)$ is a directed tree except for at most one extra edge. Remark also that if $s \leq (1-\varepsilon)t_{\text{ENT}}$, then the number of vertices in $\mathcal{G}_x(s)$ satisfies $|V| = o(n)$. Indeed, there are at most $\kappa + 1$ vertices, and by Lemma 7, $\kappa \leq t_{\text{ENT}} e^{\bar{H}(1-\varepsilon)t_{\text{ENT}}} = \mathcal{O}(n^{1-\varepsilon^2} \log n)$.

Lemma 8. *Denote by S_0 the set of all $x \in [n]$ such that $\text{TX}(\mathcal{G}_x(2h)) \leq 1$, where h is defined in (37). Then with high probability $S_0 = [n]$, that is $\mathbb{P}(S_0 = [n]) = 1 - o(1)$.*

Proof. We can use the same argument as in the proof of Lemma 5. Consider the stage $(\mathcal{G}^\ell, \mathcal{T}^\ell) \mapsto (\mathcal{G}^{\ell+1}, \mathcal{T}^{\ell+1})$ of the above sequential generation process. The conditional chance, given the past stages, that the vertex z in step 3 is already a vertex of \mathcal{G}^ℓ is at most $(\ell + 1)/n$. Hence, if $m = \lceil se^{\bar{H}s} \rceil$, from Lemma 7, the tree excess of $\mathcal{G}_x(s)$ is stochastically upper bounded by $\text{Binomial}(m, (m + 1)/n)$. As in (17), the probability that the tree excess is larger than 1 is bounded by

$$\frac{1}{2} \left(\frac{m(m+1)}{n} \right)^2.$$

For $s = 2h$, the latter is $o(1/n)$ since $m^4 = o(n)$ which follows from $4\bar{H}2h < (84/100) \log n$ (since $\varepsilon < 1/20$). \square

4.2. Nice trajectories. We will first show that for most starting states $x \in [n]$, it is likely that the walker spends its first $(1 - \varepsilon)t_{\text{ENT}}$ steps in $\mathcal{T}_x(s)$ (Lemma 11) and does not come back to it for a long time (Lemma 12). We start by identifying these good starting points x .

Lemma 9 (Good states). *Let S_\star be the set of all $x \in [n]$ such that $\text{TX}(\mathcal{G}_x(h)) = 0$. For any $s = \Theta(\log n)$, the event $\{X_s \notin S_\star\}$ is uniformly unlikely.*

Proof. From the Markov property, it is sufficient to prove the claim for $s \leq h$ and $s = \Theta(\log n)$. By Lemma 8, we may further assume that $S_0 = [n]$. Consider the trajectory (X_0, \dots, X_s) started at $X_0 = x$. The event that $X_s \notin S_\star$ is contained in the union of the events $A = \{\rho(s) \notin [e^{-\bar{H}s}, e^{-\underline{H}s}]\}$ and $A^c \cap B$ where $B = \{(X_0, \dots, X_s) \in \mathcal{P}\}$ and \mathcal{P} is the set of paths starting from x of length s in $\mathcal{G}_x(2h)$, whose end point is not in S_\star . Since $\mathcal{G}_x(2h)$ is a tree except for at most one directed edge, and $s \leq h$, then \mathcal{P} has cardinality at most 2. Hence, $Q_x(A^c \cap B) \leq 2e^{-\underline{H}s} = o(1)$. Finally, Theorem 4 asserts that A is uniformly unlikely. \square

The next corollary implies that it is enough to check that the upper bound (6) holds uniformly over S_\star rather than over all of $[n]$.

Corollary 10. *For all integers $u \geq s = \Theta(\log n)$, for any probability ν on $[n]$:*

$$\max_{x \in [n]} \|P^u(x, \cdot) - \nu\|_{\text{TV}} \leq \max_{x \in S_\star} \|P^{u-s}(x, \cdot) - \nu\|_{\text{TV}} + o_{\mathbf{P}}(1),$$

where $o_{\mathbf{P}}(1)$ denotes a random variable that converges to zero in probability, as $n \rightarrow \infty$.

Proof. Notice that

$$\|P^u(x, \cdot) - \nu\|_{\text{TV}} \leq Q_x(X_s \notin S_\star) + \max_{y \in S_\star} \|P^{u-s}(y, \cdot) - \nu\|_{\text{TV}}.$$

Taking maximum over $x \in [n]$ and using Lemma 9 concludes the proof. \square

Theorem 4 implies that the trajectory started at x is likely to remain in $\mathcal{G}_x(s)$ for a long time. We now prove that it is also likely that the trajectory stays in $\mathcal{T}_x(s)$ if $x \in S_\star$ and s is not too large.

Lemma 11. *If $\varepsilon t_{\text{ENT}} \leq s \leq (1 - \varepsilon)t_{\text{ENT}}$, then*

$$\max_{x \in S_\star} Q_x((X_0, \dots, X_s) \notin \mathcal{T}_x(s)) \xrightarrow[n \rightarrow \infty]{\mathbf{P}} 0. \quad (40)$$

Proof. By construction, there are only two ways that the trajectory exits $\mathcal{T}_x(s)$: either (i) the weight of the trajectory $\rho(s)$ is below $e^{-\bar{H}s}$ or (ii) (X_0, \dots, X_s) has used an edge in $\mathcal{G}_x(s) \setminus \mathcal{T}_x(s)$, that is, there exists $1 \leq u \leq s$ such that $(x_{u-1}, x_u) \in \mathcal{G}_x(s) \setminus \mathcal{T}_x(s)$. The event depicted in (i) is uniformly unlikely by Theorem 4. We should thus treat the event (ii).

We may follow the argument of [8, Proposition 12]. Fix $x \in [n]$, and consider the sequential generation process $(\mathcal{G}^0, \mathcal{T}^0), (\mathcal{G}^1, \mathcal{T}^1), \dots$ defined above. Define a new process $(M_\ell)_{\ell \geq 0}$ by $M_0 = 0$ and

$$M_{\ell+1} = M_\ell + \mathbf{1}(\ell < \kappa) \mathbf{1}(z_\ell \in \mathcal{G}^\ell) \widehat{w}_\ell,$$

where $\widehat{w}_\ell = \widehat{w}(y_\ell, j_\ell)$ is the weight of the half-edge (y_ℓ, j_ℓ) picked in step 1 and $z_\ell = \sigma_{y_\ell}(j_\ell)$ is the vertex picked in step 2. In words: M_ℓ is the cumulative weight of all half-edges that are matched in $\mathcal{G}^\ell \setminus \mathcal{T}^\ell$. In particular the probability of the scenario described in point (ii) above is bounded above by M_κ . Thus, to conclude the proof of Lemma 11, it is sufficient to prove that for any fixed $\delta > 0$, $M_\kappa \geq \delta$ is unlikely, uniformly over $x \in S_\star$. By construction, $M_h = 0$ for

$x \in S_\star$, hence it is sufficient to prove that $M_\kappa - M_h \geq \delta$ is uniformly unlikely. Note that we may further assume that

$$\ell \geq h \implies \widehat{w}_\ell \leq \frac{\delta}{2}, \quad (41)$$

since the complementary event entails the existence of a path of length $h = \Theta(\log n)$ and weight at least $\frac{\delta}{2} = \Omega(1)$ starting at x , which is uniformly unlikely by Lemma 6. In other words, we may safely replace \widehat{w}_ℓ with $\widehat{w}_\ell \wedge \frac{\delta}{2}$ in the definition of M , for all $\ell \geq h$. For this modified definition of M , this ensures that

$$0 \leq M_{\ell+1} - M_\ell \leq \frac{\delta}{2}.$$

for all $\ell \geq h$. We then claim that for h as in (37) and any fixed $\delta > 0$, uniformly in $x \in [n]$,

$$\mathbb{P}(M_\kappa \geq M_h + \delta) = o\left(\frac{1}{n}\right). \quad (42)$$

To prove (42), let \mathcal{F}_ℓ be the natural filtration associated to the process $(\mathcal{G}^0, \mathcal{T}^0), (\mathcal{G}^1, \mathcal{T}^1), \dots$. If $|\mathcal{G}^\ell|$ is the number of nodes in \mathcal{G}^ℓ , then

$$\begin{aligned} \mathbb{E}\left[M_{\ell+1} - M_\ell \mid \mathcal{F}_\ell\right] &= \mathbf{1}(\ell < \kappa) \frac{\widehat{w}_\ell |\mathcal{G}^\ell|}{n - |\text{Ran}(\sigma_{y_\ell})|}, \\ \mathbb{E}\left[(M_{\ell+1} - M_\ell)^2 \mid \mathcal{F}_\ell\right] &= \mathbf{1}(\ell < \kappa) \frac{\widehat{w}_\ell^2 |\mathcal{G}^\ell|}{n - |\text{Ran}(\sigma_{y_\ell})|}. \end{aligned}$$

We now use that $|\text{Ran}(\sigma_{y_\ell})| \leq \ell$, $|\mathcal{G}^\ell| \leq \ell + 1$. Moreover, by Lemma 7, $\widehat{w}_\ell \leq s/\ell$, and $\kappa \leq se^{\overline{H}s} \leq t_{\text{ENT}} n^{1-\varepsilon^2}$. Therefore,

$$\begin{aligned} \sum_{\ell \geq 0} \mathbb{E}\left[M_{\ell+1} - M_\ell \mid \mathcal{F}_\ell\right] &= \mathcal{O}\left((\log n)^2 n^{-\varepsilon^2}\right) =: a. \\ \sum_{\ell \geq 0} \mathbb{E}\left[(M_{\ell+1} - M_\ell)^2 \mid \mathcal{F}_\ell\right] &= \mathcal{O}\left((\log n)^3 n^{-1}\right) =: b. \end{aligned}$$

The martingale version of Bennett's inequality [15, Theorem 1.6] gives

$$\mathbb{P}(M_\kappa - M_h \geq a + \delta) \leq (2eb)^2.$$

Since $a = o(1)$ and $b = n^{-1+o(1)}$, this concludes the proof of (42). \square

Lemma 12. *Suppose $u, s = \Theta(\log n)$ are such that $s \leq u \wedge (1 - \varepsilon)t_{\text{ENT}}$. Then*

$$\max_{x \in [n]} Q_x(\{(X_0, \dots, X_s) \in \mathcal{T}_x(s)\} \cap \{(X_{s+1}, \dots, X_u) \cap \mathcal{T}_x(s) \neq \emptyset\}) \xrightarrow[n \rightarrow \infty]{\mathbf{P}} 0.$$

Proof. We use a version of the method explained in (25), with $k = \mathcal{O}(\log n)$ as in (26) and

$$\mathcal{B} = \{(X_0, \dots, X_s) \in \mathcal{T}_x(s)\} \cap \{(X_{s+1}, \dots, X_u) \cap \mathcal{T}_x(s) \neq \emptyset\} \cap \{\rho(s) \leq e^{-\overline{H}s}\},$$

Thanks to Theorem 4, the intersection with $\{\rho(s) \leq e^{-\overline{H}s}\}$ is not restrictive. Consider k trajectories $(X^1, W^1), \dots, (X^k, W^k)$ all started from $X_0^\ell = \mathcal{I}$ for any $1 \leq \ell \leq k$ where \mathcal{I} is picked uniformly in $[n]$. For $1 \leq \ell \leq k$, consider the sequence of non-increasing events,

$$B_\ell := \{(X^1, W^1) \in \mathcal{B}\} \cap \dots \cap \{(X^\ell, W^\ell) \in \mathcal{B}\}.$$

As explained below (27), it is sufficient to prove that $\mathbb{P}(B_\ell | B_{\ell-1}) = o(1)$, uniformly in $1 \leq \ell \leq k$. We will show the stronger uniform bounds: $\mathbb{P}_{\mathcal{F}}(B_1) = o(1)$ and, uniformly in $2 \leq \ell \leq k$,

$$\mathbb{P}_{\mathcal{F}}(B_\ell | B_{\ell-1}) = o(1), \quad (43)$$

where $\mathbb{P}_{\mathcal{F}}(\cdot) = \mathbb{P}(\cdot | \mathcal{F})$ and \mathcal{F} is the σ -algebra generated by the random variables $\mathcal{I}, \mathcal{G}_{\mathcal{I}}(s)$, and $\mathcal{T}_{\mathcal{I}}(s)$. If B_ℓ holds, then two disjoint cases may occur:

- (i) either $(X_0^\ell, \dots, X_s^\ell)$ is one of the trajectories (X_0^i, \dots, X_s^i) , $1 \leq i \leq \ell - 1$, in $\mathcal{T}_{\mathcal{I}}(s)$,
- (ii) or, $(X_0^\ell, \dots, X_s^\ell)$ is a new trajectory in $\mathcal{T}_{\mathcal{I}}(s)$ and $(X_{s+1}^\ell, \dots, X_t^\ell) \cap \mathcal{T}_{\mathcal{I}}(s)$ is not empty.

In the case $\ell = 1$ of course only the second scenario occurs. If (i) holds, then on the event $B_{\ell-1}$, $(X_0^\ell, \dots, X_s^\ell)$ is one of the at most $\ell - 1$ distinct trajectories in $\mathcal{T}_{\mathcal{I}}(s)$ each of weight at most $e^{-\mathbb{H}s}$. Hence, the probability of this case is upper bounded by $ke^{-\mathbb{H}s} = o(1)$. If (ii) holds, then the node X_s^ℓ has never been visited before and we may couple $(X_{s+1}^\ell, \dots, X_u^\ell)$ with $u - s$ i.i.d. samples from the uniform law on $[n]$ at a total-variation cost less than $\frac{ku^2}{n} = o(1)$; see the proof of Theorem 4. If this coupling occurs, then the chance of intersecting $\mathcal{T}_{\mathcal{I}}(s)$ is at most $(u - s)|\mathcal{T}_{\mathcal{I}}(s)|/n$. The latter is $o(1)$ since $|\mathcal{T}_{\mathcal{I}}(s)| \leq se^{\mathbb{H}s} \leq sn^{1-\varepsilon^2}$ by Lemma 7. This concludes the proof of (43). \square

We turn to the definition of nice trajectories. Let ε, h , and t be fixed as in (36)-(37). Set also

$$s := t - h.$$

Since $0 < \varepsilon < 1/20$, for n large enough,

$$s \leq (1 - \varepsilon)t_{\text{ENT}}.$$

For a given $x \in [n]$ and $y \notin \mathcal{G}_x(s)$, call $\mathcal{G}_y^x(h)$ the graph spanned by trajectories in $\mathcal{G}_y(h)$ which do not intersect nodes in $\mathcal{G}_x(s)$. We denote by S_x^* the set of $y \notin \mathcal{G}_x(s)$ such that $\text{TX}(\mathcal{G}_y^x(h)) = 0$. The set $\mathcal{N}_t(x)$ of *nice paths* is defined as the subset of all paths $\mathbf{p} = (x_0, x_1, \dots, x_t) \in [n]^{t+1}$, such that:

- 1) $w(\mathbf{p}) \leq n^{-1-\varepsilon/4}$;
- 2) $x_0 = x$ and $(x_0, \dots, x_s) \in \mathcal{T}_x(s)$;
- 3) $P(x_s, x_{s+1}) \geq n^{-\varepsilon/8}$.
- 4) $x_{s+1} \in S_{x_{s+1}}^*$ and $(x_{s+1}, \dots, x_t) \in \mathcal{G}_{x_{s+1}}^x(h)$.

Combining Lemma 3, Lemma 9, Lemma 11, Lemma 12 and (15), we have proved:

Proposition 13. *For ε, h, s, t as above, we have*

$$\max_{x \in S_x^*} Q_x((X_0, \dots, X_t) \notin \mathcal{N}_t(x)) \xrightarrow[n \rightarrow \infty]{\mathbf{P}} 0. \quad (44)$$

4.3. Upper bound. Define

$$P_0^t(x, y) = \sum_{\mathbf{p} \in \mathcal{N}_t(x, y)} w(\mathbf{p}), \quad (45)$$

where $\mathcal{N}_t(x, y) \subset \mathcal{N}_t(x)$ is the subset of nice paths such that $x_t = y$.

Proposition 14. *Let $0 < \varepsilon, t$ be as in (36), and $\hat{\pi}$ as in (7). For any $\delta > 0$, with high probability*

$$P_0^t(x, y) \leq (1 + \delta)\hat{\pi}(y) + \frac{\delta}{n} \quad \forall x, y \in [n]. \quad (46)$$

Notice that if Proposition 14 is available, then the argument in (34)-(35) allows us to estimate, with high probability,

$$\|\hat{\pi} - P^t(x, \cdot)\|_{\text{TV}} \leq q(x) + 2\delta, \quad (47)$$

where $q(x) = Q_x((X_0, \dots, X_t) \notin \mathcal{N}_t(x))$. From Proposition 13, uniformly in $x \in S_x^*$, one has $q(x) \xrightarrow{\mathbf{P}} 0$. This proves that (6) holds uniformly in $x \in S_x^*$. Using Corollary 10, this concludes the proof of the upper bound in Theorem 1.

Proof of Proposition 14. Consider the set $\mathcal{V}_x(s)$ of all nodes at distance s from x in the tree $\mathcal{T}_x(s)$. Any such node must be a leaf by construction. We define the set $\mathcal{L}_x(s)$ as the collection of pairs (u, k) , where $u \in \mathcal{V}_x(s)$ and $k \in [n]$. An element of $\mathcal{L}_x(s)$ is regarded as an half-edge (u, k) , with weight $\widehat{w}(u, k)$. Given $v \in S_\star^x$, by definition there is at most one path of length h from v to y in $\mathcal{G}_v^x(h)$. If such path exists, we call it $p_\star(v; y)$. Then, any $p \in \mathcal{N}_t(x, y)$ must be of the form $(x, \dots, u) \circ (u, v) \circ p_\star(v; y)$, where (x, \dots, u) is the unique path connecting x to u in $\mathcal{T}_x(s)$, for some $u \in \mathcal{V}_x(s)$ and some $v \in S_\star^x$. Here \circ denotes the natural concatenation of paths. Therefore,

$$P_0^t(x, y) = \sum_{(u, k) \in \mathcal{L}_x(s)} \widehat{w}(u, k) \sum_{v \in S_\star^x} w(p_\star(v; y)) \mathbf{1}_{\widehat{w}(u, k)w(p_\star(v; y)) \leq n^{-1-\varepsilon/4}} \mathbf{1}_{p_{u, k} \geq n^{-\varepsilon/8}} \mathbf{1}_{\sigma_u(k) = v}. \quad (48)$$

Let \mathcal{F} denote the σ -algebra generated by all the random permutations $\{\sigma_z, z \notin \mathcal{V}_x(s)\}$. A crucial observation is that the quantities $\widehat{w}(u, k)$, $w(p_\star(v; y))$, and the sets $\mathcal{L}_x(s)$, S_\star^x are all \mathcal{F} -measurable. Notice also that by construction one has

$$\frac{1}{n} \sum_{v \in S_\star^x} w(p_\star(v; y)) \leq \widehat{\pi}(y), \quad (49)$$

and

$$\sum_{(u, k) \in \mathcal{L}_x(s)} \widehat{w}(u, k) \leq 1. \quad (50)$$

Moreover, conditioned on \mathcal{F} the remaining permutations σ_u , $u \in \mathcal{V}_x(s)$, are independent and satisfy $\sigma_u(k) = y$ with probability $1/n$ for all k, y . It follows from (49)-(50) that

$$\mathbb{E}_{\mathcal{F}} P_0^t(x, y) \leq \widehat{\pi}(y), \quad (51)$$

where $\mathbb{E}_{\mathcal{F}}$ is the conditional expectation associated to \mathcal{F} . Notice also that we may write (48) as

$$P_0^t(x, y) = \sum_{u \in \mathcal{V}_x(s)} f(u, \sigma_u),$$

where

$$f(u, \sigma_u) := \sum_{k=1}^n \widehat{w}(u, k) w(p_\star(\sigma_u(k); y)) \mathbf{1}_{\widehat{w}(u, k)w(p_\star(\sigma_u(k); y)) \leq n^{-1-\varepsilon/4}} \mathbf{1}_{p_{u, k} \geq n^{-\varepsilon/8}} \mathbf{1}_{\sigma_u(k) \in S_\star^x}.$$

Since there are at most $n^{\varepsilon/8}$ indices k such that $p_{u, k} \geq n^{-\varepsilon/8}$, we have

$$0 \leq f(u, \sigma_u) \leq M = n^{\varepsilon/8} n^{-1-\varepsilon/4} = n^{-1-\varepsilon/8}.$$

Thus using Bernstein's inequality (see e.g. [9, Corollary 2.11]), for $a > 0$

$$\mathbb{P}_{\mathcal{F}}(P_0^t(x, y) - \mathbb{E}_{\mathcal{F}} P_0^t(x, y) \geq a) \leq \exp\left(-\frac{a^2}{2M(\mathbb{E}_{\mathcal{F}} P_0^t(x, y) + a)}\right).$$

Applying the above to $a = \delta \widehat{\pi}(y) + \frac{\delta}{n}$ and using (51), writing $r = n\widehat{\pi}(y)$ one finds

$$\mathbb{P}_{\mathcal{F}}\left(P_0^t(x, y) \geq (1 + \delta)\widehat{\pi}(y) + \frac{\delta}{n}\right) \leq \exp\left(-\frac{\delta^2 n^{\varepsilon/8} (r + 1)^2}{2(r(1 + \delta) + \delta)}\right).$$

Optimizing over $r \geq 0$ one has that for some constant $c(\delta) > 0$:

$$\mathbb{P}_{\mathcal{F}}\left(P_0^t(x, y) \geq (1 + \delta)\widehat{\pi}(y) + \frac{\delta}{n}\right) \leq \exp\left(-c(\delta)n^{\varepsilon/8}\right).$$

This ends the proof of Proposition 14. \square

5. PROOF OF THEOREM 2

Let $\omega = (\omega_{ij})_{1 \leq i, j < \infty}$ be i.i.d. positive random variables whose tail distribution function $G(t) = \mathbb{P}(\omega_{ij} > t)$ satisfies (9) for some $\alpha \in (0, 1)$, and consider the random transition matrix

$$P_n(i, j) := \frac{\omega_{ij}}{\omega_{i1} + \cdots + \omega_{in}} \quad (1 \leq i, j \leq n). \quad (52)$$

Permuting entries within a row clearly leaves the distribution of P_n unchanged. Therefore, P_n is of the form (2), but with the parameters $(p_{ij})_{1 \leq i, j \leq n}$ now being random. In order to apply our Theorem 1 and obtain Theorem 2, we only have to establish that almost-surely,

$$\frac{1}{n} \sum_{i, j=1}^n P_n(i, j) \log P_n(i, j) \xrightarrow{n \rightarrow \infty} h(\alpha); \quad (53)$$

$$\max_{i \in [n]} \sum_{j=1}^n P_n(i, j) (\log P_n(i, j))^2 = o(\log n); \quad (54)$$

$$\limsup_{n \rightarrow \infty} \left\{ \frac{1}{n} \sum_{i, j=1}^n \mathbf{1}_{P_n(i, j) > 1-\varepsilon} \right\} \xrightarrow{\varepsilon \rightarrow 0^+} 0. \quad (55)$$

The proof will rely on the following estimates on the random probability vector $(P_n(1, 1), \dots, P_n(1, n))$.

Lemma 15 (Uniform sparsity). *For each $\beta \in (\alpha, 1)$, there exists $\lambda > 0$ such that*

$$\sup_{n \geq 1} \mathbb{E} \left[\exp \left\{ \lambda \sum_{j=1}^n (P_n(1, j))^\beta \right\} \right] < \infty. \quad (56)$$

Lemma 16 (Beta asymptotics). *Let ξ_n be distributed as a size-biased pick from the random sequence $(P_n(1, 1), \dots, P_n(1, n))$, i.e., for any measurable $g: [0, 1] \rightarrow [0, \infty]$,*

$$\mathbb{E}[g(\xi_n)] = \mathbb{E} \left[\sum_{j=1}^n P_n(1, j) g(P_n(1, j)) \right] = n \mathbb{E}[P_n(1, 1) g(P_n(1, 1))].$$

Then $\xi_n \xrightarrow[n \rightarrow \infty]{d} \xi$, where ξ has the Beta($1 - \alpha, \alpha$)–density:

$$f_\alpha(u) = \frac{(1-u)^{\alpha-1} u^{-\alpha}}{\Gamma(\alpha)\Gamma(1-\alpha)}, \quad (0 < u < 1).$$

Before we establish those Lemmas, let us quickly see how they imply the three almost-sure conditions stated above. For any $0 < \varepsilon, \beta < 1$, we have

$$\sum_{j=1}^n P_n(i, j) (\log P_n(i, j))^2 \leq (\log \varepsilon)^2 + \sup_{p \in [0, \varepsilon]} \left\{ p^{1-\beta} (\log p)^2 \right\} \sum_{j=1}^n (P_n(i, j))^\beta, \quad (57)$$

where we have simply split the summands according to whether $P_n(i, j) \leq \varepsilon$ or not. Note that the supremum on the right-hand side can be made arbitrarily small by choosing ε small enough. Claim (54) follows, since for $\beta > \alpha$, Lemma 15 ensures that almost-surely as $n \rightarrow \infty$,

$$\max_{i \in [n]} \left\{ \sum_{j=1}^n (P_n(i, j))^\beta \right\} = \mathcal{O}(\log n). \quad (58)$$

We now turn to (53). The row entropies $\left\{ -\sum_{j=1}^n P_n(i, j) \log P_n(i, j) \right\}_{1 \leq i \leq n}$ are independent, $[0, \log n]$ –valued random variables with mean $-\mathbb{E}[\log \xi_n]$, where ξ_n is as in Lemma 16. Therefore,

Azuma-Hoeffding's inequality ensures that almost-surely as $n \rightarrow \infty$,

$$\frac{1}{n} \sum_{i,j=1}^n P_n(i,j) \log P_n(i,j) = \mathbb{E}[\log \xi_n] + o(1).$$

In view of (57), Lemma 15 is more than enough to ensure the uniform integrability of $(\log \xi_n)_{n \geq 1}$. Together with the weak convergence $\xi_n \rightarrow \xi$ stated in Lemma 16, this implies

$$\mathbb{E}[\log \xi_n] \xrightarrow{n \rightarrow \infty} \mathbb{E}[\log \xi].$$

It is classical that the expected logarithm of a Beta($1 - \alpha, \alpha$) is $\psi(\alpha) - \psi(1) = -h(\alpha)$, and (53) follows.

The proof of (55) is similar: for each $\varepsilon < \frac{1}{2}$, the random variables $\left\{ \sum_{j=1}^n \mathbf{1}_{P_n(i,j) \geq 1-\varepsilon} \right\}_{1 \leq i \leq n}$ are independent, $[0, 1]$ -valued and with mean $\mathbb{E}[\xi_n^{-1} \mathbf{1}_{\xi_n \geq 1-\varepsilon}]$. Therefore, Azuma-Hoeffding's inequality ensures that almost-surely as $n \rightarrow \infty$,

$$\begin{aligned} \frac{1}{n} \sum_{i,j=1}^n \mathbf{1}_{P_n(i,j) \geq 1-\varepsilon} &= \mathbb{E}[\xi_n^{-1} \mathbf{1}_{\xi_n \geq 1-\varepsilon}] + o(1) \\ &= \mathbb{E}[\xi^{-1} \mathbf{1}_{\xi \geq 1-\varepsilon}] + o(1), \end{aligned}$$

where the second line follows from Lemma 16 and the fact that the Beta distribution is atom-free. It remains to note that $\mathbb{E}[\xi^{-1} \mathbf{1}_{\xi \geq 1-\varepsilon}] \rightarrow 0$ as $\varepsilon \rightarrow 0$, since $\mathbb{P}(\xi \in (0, 1)) = 1$. We now turn to the proof of Lemmas 15 and 16.

5.1. Proof of Lemma 15. Our starting point is a classical result on regularly varying functions (see, e.g., [14, Theorem VIII.9.1]), which asserts that as $t \rightarrow \infty$,

$$\mathbb{E} \left[(\omega_{11} \wedge t)^\beta \right] \sim \frac{\beta}{\beta - \alpha} t^\beta \mathbb{P}(\omega_{11} > t).$$

In particular, there exists a constant $c_\beta < \infty$ such that for all $t > 0$,

$$\mathbb{E} \left[\left(\frac{\omega_{11}}{t} \wedge 1 \right)^\beta \right] \leq c_\beta \mathbb{P}(\omega_{11} > t).$$

Since $(\omega_{11}, \dots, \omega_{1n})$ are i.i.d., we immediately obtain that for any $J \subset [n]$,

$$\mathbb{E} \left[\prod_{j \in J} \left(\frac{\omega_{1j}}{t} \wedge 1 \right)^\beta \right] \leq c_\beta^{|J|} \mathbb{P} \left(\min_{j \in J} \omega_{1j} > t \right).$$

This formula holds for any $t > 0$, and we may choose $t = \max_{j \in [n] \setminus J} \omega_{1j}$, since the latter is independent of $(\omega_{1j})_{j \in J}$. With this choice of t , we clearly have $P_n(1, j) \leq \frac{\omega_{1j}}{t} \wedge 1$ and therefore

$$\mathbb{E} \left[\prod_{j \in J} (P_n(1, j))^\beta \right] \leq c_\beta^{|J|} \mathbb{P} \left(\min_{j \in J} \omega_{1j} > \max_{j \in [n] \setminus J} \omega_{1j} \right).$$

Write A_J for the event on the RHS. Clearly, the $\{A_J : J \subset [n], |J| = k\}$ are pairwise disjoint. Thus,

$$\sum_{J \subset [n], |J|=k} \mathbb{E} \left[\prod_{j \in J} (P_n(1, j))^\beta \right] \leq c_\beta^k. \quad (59)$$

This is enough to conclude. Indeed, using $e^{\lambda x} \leq 1 + (e^\lambda - 1)x$ for $x \in [0, 1]$, we have

$$\begin{aligned} \mathbb{E} \left[\exp \left\{ \lambda \sum_{j=1}^n (P_n(1, j))^\beta \right\} \right] &\leq \mathbb{E} \left[\prod_{j=1}^n \left(1 + (e^\lambda - 1) (P_n(1, j))^\beta \right) \right] \\ &\leq \sum_{k=0}^n (e^\lambda - 1)^k \sum_{J \subset [n], |J|=k} \mathbb{E} \left[\prod_{j \in J} (P_n(1, j))^\beta \right] \\ &\leq \sum_{k=0}^n \left(c_\beta (e^\lambda - 1) \right)^k, \end{aligned}$$

which is bounded uniformly in n as long as $\lambda < \log \left(1 + \frac{1}{c_\beta} \right)$.

5.2. Proof of Lemma 16. Since the $(\xi_n)_{n \geq 1}$ are $[0, 1]$ -valued, it is enough to prove $\mathbb{E}[\xi_n^p] \rightarrow \mathbb{E}[\xi^p]$ for each $p \geq 0$. We first rewrite both sides as follows:

$$\mathbb{E}[\xi_n^p] = \int_0^1 n \mathbb{P} \left(\{P_n(1, 1)\}^{p+1} > u \right) du = \int_0^1 n \mathbb{P} (P_n(1, 1) > u) (p+1) u^p du, \quad (60)$$

$$\mathbb{E}[\xi^p] = \int_0^1 \frac{f_\alpha(u)}{u} u^{p+1} du = \int_0^1 \kappa \left(\frac{1-u}{u} \right)^\alpha (p+1) u^p du, \quad (61)$$

where $\kappa^{-1} = \Gamma(1+\alpha)\Gamma(1-\alpha)$, and where we have used the change of variables $u \mapsto u^{p+1}$ for (60) and an integration by parts for (61). Comparing these two lines, our goal reduces to proving that

$$\forall u \in (0, 1), \quad n \mathbb{P} (P_n(1, 1) > u) \xrightarrow{n \rightarrow \infty} \kappa \left(\frac{1-u}{u} \right)^\alpha. \quad (62)$$

Indeed, the convergence of (60) to (61) then follows by dominated convergence since for $\beta \in (\alpha, 1)$,

$$n \mathbb{P} (P_n(1, 1) > u) = \mathbb{E} \left[\sum_{j=1}^n \mathbf{1}_{P_n(1, j) > u} \right] \leq u^{-\beta} \mathbb{E} \left[\sum_{j=1}^n \{P_n(1, j)\}^\beta \right] \leq c_\beta u^{-\beta}, \quad (63)$$

by (59). We may now fix $0 < u < 1$ and focus on (62). Our regular variation assumption on G yields

$$R(s) := \frac{G\left(\frac{us}{1-u}\right)}{G(s)} \xrightarrow{s \rightarrow \infty} \left(\frac{1-u}{u} \right)^\alpha. \quad (64)$$

In particular, $s \mapsto R(s)$ is bounded on $(0, \infty)$. Now, since ω_{11} is independent of $S_n := \omega_{12} + \dots + \omega_{1n}$,

$$\mathbb{P} (P_n(1, 1) > u) = \mathbb{P} \left(\omega_{11} > \frac{u S_n}{1-u} \right) = \mathbb{E} \left[G \left(\frac{u S_n}{1-u} \right) \right] = \mathbb{E} [G(S_n) R(S_n)].$$

Observing that the right-hand side simplifies to $\mathbb{E}[G(S_n)]$ when $u = \frac{1}{2}$, we deduce that

$$\mathbb{P} (P_n(1, 1) > u) - \left(\frac{1-u}{u} \right)^\alpha \mathbb{P} \left(P_n(1, 1) > \frac{1}{2} \right) = \mathbb{E} \left[G(S_n) \left\{ R(S_n) - \left(\frac{1-u}{u} \right)^\alpha \right\} \right].$$

Since S_n increases almost-surely to $+\infty$ as $n \rightarrow \infty$ and since R is bounded, (64) implies that

$$\mathbb{E} \left[\left\{ R(S_n) - \left(\frac{1-u}{u} \right)^\alpha \right\}^2 \right] \xrightarrow{n \rightarrow \infty} 0,$$

by dominated convergence. On the other hand, since G is decreasing, we have

$$\mathbb{E} \left[\{G(S_n)\}^2 \right] \leq \mathbb{E} \left[\left\{ G \left(\max_{2 \leq j \leq n} \omega_{1j} \right) \right\}^2 \right] = \mathbb{P} \left[\min(\omega_{11}, \omega_{1(n+1)}) > \max_{2 \leq j \leq n} \omega_{1j} \right] \leq \binom{n+1}{2}^{-1},$$

by symmetry. Invoking the Cauchy-Schwarz inequality, we conclude that

$$n\mathbb{P}(P_n(1, 1) > u) - \left(\frac{1-u}{u} \right)^\alpha n\mathbb{P} \left(P_n(1, 1) > \frac{1}{2} \right) \xrightarrow{n \rightarrow \infty} 0.$$

This is not quite (62), as it is not yet clear that $n\mathbb{P}(P_n(1, 1) > \frac{1}{2}) \rightarrow \kappa$. However, one may still insert this into (60) and invoke the domination (63) to obtain that

$$\mathbb{E}[\xi_n^p] - \frac{n}{\kappa} \mathbb{P} \left(P_n(1, 1) > \frac{1}{2} \right) \mathbb{E}[\xi^p] \xrightarrow{n \rightarrow \infty} 0.$$

But now the special case $p = 0$ shows that $n\mathbb{P}(P_n(1, 1) > \frac{1}{2}) \rightarrow \kappa$, which completes the proof.

REFERENCES

- [1] D. Aldous. Random walks on finite groups and rapidly mixing Markov chains. In *Seminar on probability, XVII*, volume 986 of *Lecture Notes in Math.*, pages 243–297. Springer, Berlin, 1983.
- [2] D. Aldous and P. Diaconis. Shuffling cards and stopping times. *American Mathematical Monthly*, pages 333–348, 1986.
- [3] D. Aldous and J. Fill. Reversible Markov chains and random walks on graphs, 2002.
- [4] R. Basu, J. Hermon, and Y. Peres. Characterization of cutoff for reversible Markov chains. *ArXiv e-prints*, Sept. 2014.
- [5] A. Ben-Hamou and J. Salez. Cutoff for non-backtracking random walks on sparse random graphs. *ArXiv e-prints*, 2015.
- [6] N. Berestycki, E. Lubetzky, Y. Peres, and A. Sly. Random walks on the random graph. *ArXiv e-prints*, 2015.
- [7] C. Bordenave, P. Caputo, D. Chafaï, and D. Piras. Spectrum of large random Markov chains: heavy-tailed weights on the oriented complete graph. *ArXiv e-prints*, Oct. 2016.
- [8] C. Bordenave, P. Caputo, and J. Salez. Random walk on sparse random digraphs. *arXiv preprint arXiv:1508.06600*, 2015.
- [9] S. Boucheron, G. Lugosi, and P. Massart. *Concentration inequalities*. Oxford University Press, Oxford, 2013. A nonasymptotic theory of independence, With a foreword by Michel Ledoux.
- [10] P. Diaconis. The cutoff phenomenon in finite Markov chains. *Proc. Nat. Acad. Sci. U.S.A.*, 93(4):1659–1664, 1996.
- [11] P. Diaconis and L. Saloff-Coste. Separation cut-offs for birth and death chains. *The Annals of Applied Probability*, 16(4):2098–2122, 2006.
- [12] P. Diaconis and M. Shahshahani. Generating a random permutation with random transpositions. *Probability Theory and Related Fields*, 57(2):159–179, 1981.
- [13] J. Ding, E. Lubetzky, and Y. Peres. Total variation cutoff in birth-and-death chains. *Probability theory and related fields*, 146(1-2):61–85, 2010.
- [14] W. Feller. *An introduction to probability theory and its applications. Vol. II*. Second edition. John Wiley & Sons, Inc., New York-London-Sydney, 1971.
- [15] D. A. Freedman. On tail probabilities for martingales. *The Annals of Probability*, pages 100–118, 1975.
- [16] D. A. Levin, Y. Peres, and E. L. Wilmer. *Markov chains and mixing times*. American Mathematical Soc., 2009.
- [17] E. Lubetzky and Y. Peres. Cutoff on all Ramanujan graphs. *ArXiv e-prints*, July 2015.
- [18] E. Lubetzky and A. Sly. Cutoff phenomena for random walks on random regular graphs. *Duke Mathematical Journal*, 153(3):475–510, 2010.
- [19] J. Pitman and M. Yor. The two-parameter poisson-dirichlet distribution derived from a stable subordinator. *The Annals of Probability*, pages 855–900, 1997.