



HAL
open science

Universality of cutoff for exclusion with reservoirs

Justin Salez

► **To cite this version:**

Justin Salez. Universality of cutoff for exclusion with reservoirs. *Annals of Probability*, 2023, 51 (2), 10.1214/22-AOP1600 . hal-03902891

HAL Id: hal-03902891

<https://hal.science/hal-03902891v1>

Submitted on 6 Jan 2025

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.

Universality of cutoff for exclusion with reservoirs

Justin Salez*

January 13, 2022

Abstract

We consider the reversible exclusion process with reservoirs on arbitrary networks. We characterize the spectral gap, mixing time, and mixing window of the process, in terms of certain simple statistics of the underlying network. Among other consequences, we establish a non-conservative analogue of Aldous’s spectral gap conjecture, and we show that cutoff occurs if and only if the product condition is satisfied. We illustrate this by providing explicit cutoffs on discrete lattices of arbitrary dimensions and boundary conditions, which substantially generalize recent one-dimensional results. We also obtain cutoff phenomena in relative entropy, Hilbert norm, separation distance and supremum norm. Our proof exploits negative dependence in a novel, simple way to reduce the understanding of the whole process to that of single-site marginals. We believe that this approach will find other applications.

Contents

1	Introduction	2
1.1	Setup	2
1.2	Results	4
1.3	Examples	8
2	Perturbations of product measures	9
2.1	Framework	10
2.2	Upper-bound	10
2.3	Lower-bound	12
3	Application to exclusion with reservoirs	13
3.1	Graphical construction	13
3.2	Analysis of the perturbed region	15
3.3	Putting things together	17

*CEREMADE, CNRS, UMR 7534, Université Paris-Dauphine, PSL University, 75016 Paris, France

1 Introduction

The exclusion process [30, 20] is a classical model of interacting random walks in which indistinguishable particles attempt to evolve independently on a graph, except that their jumps are canceled if the destination is already occupied. Here we consider the non-conservative variant of the model, where particles may additionally be created and annihilated at certain vertices, modeling contact with an external reservoir. We refer the reader to the papers [4, 18] and the more recent works [11, 12] for motivations and background on this process.

1.1 Setup

Networks. Let us first specify the geometry of the model and introduce some terminology. Throughout the paper, we consider a *network* $G = (V, c, \kappa)$ consisting of

- a finite set V whose elements are called *vertices*;
- a symmetric array $c: V \times V \rightarrow \mathbb{R}_+$, whose entries are called *conductances*;
- a function $\kappa: V \rightarrow \mathbb{R}_+$, whose entries are called *external rates*.

The support of $c(\cdot, \cdot)$ constitutes the set of *edges* of the network, along which particles can move. The support of $\kappa(\cdot)$ represents the *boundary* of the network, where particles can be created or annihilated due to contact with an external *reservoir*. To ensure irreducibility, we will always assume that the network G is connected, and that its boundary is not empty. A simple way to produce networks consists in “cutting them out” from some fixed, locally finite graph $\mathbb{G} = (\mathbb{V}, \mathbb{E})$. More precisely, we choose a finite connected subset $V \subset \mathbb{V}$ and, for $i, j \in V$, we set

$$c(i, j) := \mathbf{1}_{\{i, j\} \in \mathbb{E}} \quad \text{and} \quad \kappa(i) := \sum_{k \in \mathbb{V} \setminus V} \mathbf{1}_{\{i, k\} \in \mathbb{E}}.$$

We call $G = (V, c, \kappa)$ the *network induced by V in \mathbb{G}* . An important example to have in mind, and to which we shall come back later, is the network induced by an hypercube $V = [n]^d$ in the d -dimensional square lattice $\mathbb{G} = \mathbb{Z}^d$.

State space and generator. As usual, the state of the system is represented by a binary vector $x = (x_i)_{i \in V}$, with $x_i = 1$ if the vertex i is occupied and $x_i = 0$ if it is empty. We let $x^{i \leftrightarrow j}$, $x^{i, 1}$ and $x^{i, 0}$ denote the vectors obtained from x by respectively swapping the i -th and j -th coordinates (exchange), resetting the i -th coordinate to 1 (creation), and resetting it to 0 (annihilation). With this notation, the *exclusion process with reservoir density $\rho \in (0, 1)$ on the network G* is the

continuous-time Markov chain with state space $\mathcal{X} = \{0, 1\}^V$ and whose infinitesimal generator \mathcal{L} acts on observables $f: \mathcal{X} \rightarrow \mathbb{R}$ as follows:

$$(\mathcal{L}f)(x) := \frac{1}{2} \sum_{i,j \in V} c(i,j) [f(x^{i \leftrightarrow j}) - f(x)] + \sum_{i \in V} \kappa(i) [\rho f(x^{i,1}) + (1 - \rho)f(x^{i,0}) - f(x)].$$

In words, every pair of vertices $\{i, j\}$ exchange contents at rate $c(i, j)$, and every vertex i resamples its content afresh at rate $\kappa(i)$ according to \mathcal{B}_ρ , the Bernoulli distribution with mean ρ .

Convergence to equilibrium. The generator \mathcal{L} is clearly irreducible and reversible w.r.t. the product measure $\pi = \mathcal{B}_\rho^{\otimes V}$. As a consequence, the associated semi-group $(\mathcal{P}_t)_{t \geq 0}$ *mixes*:

$$\forall x, y \in \mathcal{X}, \quad \mathcal{P}_t(x, y) \xrightarrow[t \rightarrow \infty]{} \pi(y).$$

The time-scale on which this convergence occurs is measured by the so-called *mixing times*

$$t_{\text{MIX}}(\varepsilon) := \min \left\{ t \geq 0: \max_{x \in \mathcal{X}} \|\mathcal{P}_t(x, \cdot) - \pi\|_{\text{TV}} \leq \varepsilon \right\}, \quad \varepsilon \in (0, 1), \quad (1)$$

where $\|\mu - \pi\|_{\text{TV}} = \max_{A \subseteq \mathcal{X}} |\mu(A) - \pi(A)|$ denotes the total-variation distance between two probability measures μ, π on \mathcal{X} . The following stronger measures of discrepancy will also be considered:

- Separation distance: $\text{sep}(\mu, \pi) := \max_{x \in \mathcal{X}} \left(1 - \frac{\mu(x)}{\pi(x)} \right)$;
- Relative entropy: $d_{\text{KL}}(\mu || \pi) := \sum_{x \in \mathcal{X}} \mu(x) \log \frac{\mu(x)}{\pi(x)}$;
- Hilbert norm: $\|\frac{\mu}{\pi} - 1\|_{L^2_\pi}^2 := \sum_{x \in \mathcal{X}} \pi(x) \left(\frac{\mu(x)}{\pi(x)} - 1 \right)^2$;
- Supremum norm: $\|\frac{\mu}{\pi} - 1\|_\infty := \max_{x \in \mathcal{X}} \left| \frac{\mu(x)}{\pi(x)} - 1 \right|$.

We refer the reader to the books [25, 19]) for details on those natural quantities. Understanding how the fundamental parameter $t_{\text{MIX}}(\varepsilon)$ depends on the size and geometry of the underlying network constitutes a natural and important question, to which the present paper is devoted.

State of the art. While the convergence to equilibrium of the exclusion process *without* reservoirs has received a considerable attention (see, e.g., [31, 6, 26, 27, 16, 15, 17, 2, 13]), only very little has been said about the non-conservative version of the model. In fact, there seems to be only two examples of networks on which the mixing time of the exclusion process with reservoirs is known. The first is the segment of length n with a reservoir at one end ($V = [n]$, $c(i, j) = \mathbf{1}_{|i-j|=1}$, $\kappa(i) = \mathbf{1}_{i=n}$), for which Gantert, Nestoridi and Schmid [11] recently established the estimate

$$t_{\text{MIX}}(\varepsilon) = \frac{2n^2 \log n}{\pi^2} + o(n^2 \log n), \quad (2)$$

as $n \rightarrow \infty$, for any fixed $\varepsilon, \rho \in (0, 1)$. The proof crucially relies on a coupling used by Lacoïn for the conservative version of the problem [16], which seems specific to the one-dimensional setup. The second example is the variant of the above network where reservoirs are present at both ends of the segment ($\kappa(i) = \mathbf{1}_{i=1} + \mathbf{1}_{i=n}$). Using a new and promising application of Yau’s celebrated relative entropy method [32], Gonçalves, Jara, Marinho and Menezes [12] proved that

$$t_{\text{MIX}}(\varepsilon) = \frac{n^2 \log n}{2\pi^2} + cn^2 + o(n^2), \quad (3)$$

as $n \rightarrow \infty$, where $c = c(\varepsilon, \rho) \in \mathbb{R}$ is an explicit constant. In both cases, the remarkable fact that the leading order term does not depend on $\varepsilon \in (0, 1)$ reflects a sharp transition to equilibrium known as a *cutoff*; see [8] for an introduction. We emphasize that the works [11, 12] are not limited to the above estimates: [11] initiates the more delicate study of the non-reversible setup, where the reservoir densities differ at the two end-points of the segment, while [12] provides a very detailed picture of the system started from any smooth initial condition. To the best of our knowledge, however, the current understanding of the mixing properties of the exclusion with reservoirs is limited to the two aforementioned one-dimensional networks.

Our contribution. In the present paper, we consider the exclusion process on an *arbitrary* network G . In this level of generality, we determine the spectral gap, mixing time, and mixing window of the process, in terms of certain simple spectral statistics of G . As a by-product, we completely characterize the occurrence of the cutoff phenomenon, and obtain multi-dimensional generalizations of (2) and (3). We also establish cutoff phenomena in relative entropy, Hilbert norm, separation distance and supremum norm. Our proof exploits negative dependence in a novel, simple way to reduce the understanding of the whole process to that of single-site marginals. We believe that this approach will find other applications.

1.2 Results

The main message of our paper is that the mixing properties of the high-dimensional operator \mathcal{L} are entirely governed by those of a much lower-dimensional object, namely the $V \times V$ matrix

$$\Delta(i, j) := \begin{cases} c(i, j) & \text{if } j \neq i \\ -\kappa(i) - \sum_{k \in V} c(i, k) & \text{if } j = i. \end{cases}$$

We call Δ the *Laplace matrix* of the network $G = (V, c, \kappa)$. It describes the evolution of a single random walker moving on V according to the conductances $c(\cdot, \cdot)$ and killed at the space-varying rate $\kappa(\cdot)$. Writing τ for the time at which the walker is killed, we consider the key statistics

$$\mathfrak{z}_i(t) := \mathbb{P}_i(\tau > t),$$

where the notation \mathbb{P}_i indicates that the walk starts at i . Note that \mathfrak{z} solves the differential equation

$$\frac{d\mathfrak{z}}{dt} = \Delta\mathfrak{z}, \quad \mathfrak{z}(0) = \mathbf{1}, \quad (4)$$

i.e. $\mathfrak{z}(t) = e^{t\Delta}\mathbf{1}$. Our main result asserts that the distance to equilibrium of the exclusion process with reservoirs at any time t is uniformly controlled, in a two-sided way, by the simple quantity $\|\mathfrak{z}(t)\|^2 := \sum_{i \in V} \mathfrak{z}_i^2(t)$. Below and throughout the paper, we use the following short-hands:

$$\rho_\star := \min(\rho, 1 - \rho), \quad x_\star := \begin{cases} \mathbf{1} & \text{if } \rho \leq \frac{1}{2} \\ \mathbf{0} & \text{else.} \end{cases}$$

Theorem 1 (Two-sided estimate). *From the extremal initial state x_\star , we have the lower-bound*

$$\|\mathcal{P}_t(x_\star, \cdot) - \pi\|_{\text{TV}} \geq \frac{\|\mathfrak{z}(t)\|^2}{4 + \|\mathfrak{z}(t)\|^2},$$

at any time $t \geq 0$. Conversely, we have the uniform Hilbert-norm upper-bound

$$\max_{x \in \mathcal{X}} \left\| \frac{\mathcal{P}_t(x, \cdot)}{\pi} - 1 \right\|_{L_\pi^2} \leq \sqrt{\exp\left(\frac{\|\mathfrak{z}(t)\|^2}{\rho_\star}\right) - 1}.$$

The first estimate asserts that $\|\mathfrak{z}(t)\|$ needs to be small in order for the exclusion process to be well-mixed, at least from the extremal initial state x_\star . The second (and much more surprising) bound asserts that this condition actually also suffices to guarantee mixing from *any* initial condition, even when the distance to equilibrium is measured in the strong L_π^2 -norm. Thus, the mixing time of the process is essentially the time at which $\|\mathfrak{z}(t)\|$ becomes small. In order to study this quantity, observe that (4) implies the spectral expression

$$\|\mathfrak{z}(t)\|^2 = \sum_{k=1}^{|V|} e^{-2\lambda_k t} \langle \psi_k, \mathbf{1} \rangle^2, \quad (5)$$

where $0 < \lambda_1 \leq \dots \leq \lambda_{|V|}$ denote the eigenvalues of the symmetric positive-definite matrix $-\Delta$, and $\psi_1, \dots, \psi_{|V|}$ a corresponding orthonormal basis of eigenvectors. It follows that the mixing properties of our interacting particle system are entirely dictated by the spectral statistics of Δ and, most particularly, by the Perron eigen-pair (λ_1, ψ_1) , henceforth denoted simply (λ, ψ) . This has a number of important consequences, which we now enumerate (see Section 3.3 for details).

Spectral gap. The most fundamental parameter of a reversible Markov generator \mathcal{L} is arguably its *spectral gap* or *Poincaré constant*, defined as the second smallest eigenvalue of $-\mathcal{L}$. This constant provides quantitative controls on a variety of properties of the process, including concentration of measure (via Poincaré's inequality), isoperimetry (via Cheeger's inequality) and mixing (via contraction in the L_π^2 -norm), see the books [25, 19] for details. Our main estimate implies the following pleasant surprise, which does not seem to have been noted before.

Corollary 2 (Spectral gap). *The spectral gap of the exclusion process with reservoir density $\rho \in (0, 1)$ on G coincides with the smallest eigenvalue λ of the matrix $-\Delta$.*

Recall that \mathcal{L} is a $2^{|V|}$ -dimensional operator describing the joint evolution of many interacting particles, while Δ is a $|V|$ -dimensional matrix describing the motion of a single particle. The drastic dimensionality reduction stated in Corollary 2 can be seen as a non-conservative analogue of the celebrated Aldous’ *spectral gap conjecture*, now proved by Caputo, Liggett and Richthammer [6]. See Hermon and Salez [14] for a similar conclusion in the case of the Zero-Range process. Corollary 2 can be used to explicitly compute the spectral gap in various examples, see Section 1.3.

Window and cutoff. A second notable consequence of Theorem 1 is the following universal estimate on the *width* of the mixing window, i.e., the time-scale during which the system moves from being barely mixed to being completely mixed.

Corollary 3 (Mixing window). *There is a constant $c = c(\varepsilon, \rho)$, not depending on G , such that*

$$t_{\text{MIX}}(\varepsilon) - t_{\text{MIX}}(1 - \varepsilon) \leq \frac{c}{\lambda}.$$

This result implies the following characterization of cutoff for exclusion processes with reservoirs. Here and throughout the paper, when considering a sequence of networks $(G_n)_{n \geq 1}$ instead of a fixed network G , we naturally index all relevant quantities by n (e.g., $|V_n|$, $t_{\text{MIX}}^{(n)}$, λ_n , ψ_n).

Corollary 4 (Characterization of cutoff). *Consider the exclusion process with fixed reservoir density $\rho \in (0, 1)$ on a sequence of networks $(G_n)_{n \geq 1}$. Then the cutoff phenomenon*

$$\forall \varepsilon \in (0, 1), \quad \frac{t_{\text{MIX}}^{(n)}(1 - \varepsilon)}{t_{\text{MIX}}^{(n)}(\varepsilon)} \xrightarrow{n \rightarrow \infty} 1, \quad (6)$$

occurs if and only if the sequence satisfies the so-called “product condition”:

$$\lambda_n \times t_{\text{MIX}}^{(n)}(1/4) \xrightarrow{n \rightarrow \infty} +\infty. \quad (7)$$

Moreover, this remains valid if ρ varies with n , as long as it stays bounded away from 0 and 1.

The value $1/4$ appearing in (7) is arbitrary, and can be replaced by any other precision $\varepsilon \in (0, 1)$. The interest of this criterion is that it only involves orders of magnitude: unlike the definition (6), it can be checked without having to determine the precise pre-factor in front of mixing times. The product condition is well known to be necessary for cutoff, along any sequence of reversible chains [19, Proposition 18.4]. In the 2004 AIM workshop on mixing times, Peres [28] conjectured that it is also sufficient. Unfortunately, counter-examples have been constructed; see [7, Section 6] or [19,

Example 18.7]. However, sufficiency has been established for all birth-and-death chains [9] and, more generally, all random walks on trees [3]. Corollary 4 adds the exclusion process with reservoirs to this short list. For general chains, finding an effective sufficient criterion for cutoff remains the most fundamental problem in the area of mixing times. See [29] for a first step in this direction.

Mixing time. Corollary 4 can be used to predict the occurrence of a cutoff, but does not say *where* it occurs. Fortunately, Theorem 1 also implies sharp mixing-time estimates. We start with the following universal upper-bound.

Corollary 5 (Upper-bound). *There is a constant $c = c(\varepsilon, \rho)$, not depending on G , such that*

$$t_{\text{MIX}}(\varepsilon) \leq \frac{\log |V| + c}{2\lambda}. \quad (8)$$

Moreover, this remains valid if the distance to equilibrium is measured in the stronger L_π^2 sense.

This result essentially asserts that the mixing time of the whole system is at most that of $|V|$ independent random walkers. Proving a similar result for the conservative version of the model constitutes a long-standing open problem, see Oliveira [27], Alon and Kozma [2], or Hermon and Pymar [13] for partial progress. The generic bound (8) happens to be sharp in many cases, as we will now see. Recall that ψ denotes the eigenvector of $-\Delta$ corresponding to the smallest eigenvalue λ . By the Perron-Frobenius theorem, ψ is the only eigenvector (up to scalar multiplication) whose coordinates all have the same sign. The normalized vector $\bar{\psi} := \psi / \langle \psi, \mathbf{1} \rangle$ is known as the *quasi-stationary distribution* of the network G : it is the large-time limit of the distribution of a random walker on G conditioned on not having been killed yet (see, e.g., [1, Section 3.6.5]).

Corollary 6 (Lower-bound). *There is a constant $c = c(\varepsilon)$, not depending on G, ρ , such that*

$$t_{\text{MIX}}(\varepsilon) \geq \frac{\log (\langle \psi, \mathbf{1} \rangle^2) - c}{2\gamma}. \quad (9)$$

Note that the quantity $\langle \psi, \mathbf{1} \rangle^2 = \frac{1}{\|\psi\|^2}$ measures how balanced the unit vector ψ is: it ranges from 1 (when ψ concentrates on a single entry) to $|V|$ (when all entries of ψ are equal). In particular, the bounds (8) and (9) match as soon as ψ is sufficiently *delocalized*, in the following precise sense.

Corollary 7 (Delocalization implies cutoff). *Consider the exclusion process with fixed reservoir density $\rho \in (0, 1)$ on any sequence of networks $(G_n)_{n \geq 1}$ satisfying the delocalization condition*

$$\langle \psi_n, \mathbf{1} \rangle^2 \geq |V_n|^{1-o(1)}.$$

Then, cutoff occurs at the following time:

$$t_{\text{MIX}}^{(n)}(\varepsilon) \underset{n \rightarrow \infty}{\sim} \frac{\log |V_n|}{2\lambda_n}. \quad (10)$$

Here again, this remains valid if ρ varies with n , as long as it is bounded away from 0 and 1.

Other metrics. Finally, let us briefly discuss what happens when the total-variation distance appearing in (1) is replaced by one of the stronger divergences mentioned above. First recall that, for any measure $\mu \in \mathcal{X}$, we have

$$2\|\mu - \pi\|_{\text{TV}}^2 \leq d_{\text{KL}}(\mu|\pi) \leq \left\| \frac{\mu}{\pi} - 1 \right\|_{L^2_\pi}^2,$$

where the first inequality is Pinsker's inequality and the second is simply $\log u \leq u - 1$. As a consequence, the total-variation cutoff announced in Corollary 7 also occurs in relative entropy and in the Hilbert norm. However, it turns out that when the distance to equilibrium is measured in separation distance or in the supremum norm, one has to wait twice longer. More precisely, our methods imply the following analogue of Theorem 1, in which $\|\mathfrak{z}(t)\|^2$ is essentially replaced with $\|\mathfrak{z}(t/2)\|^2$. Note that the separation distance is, always a lower bound on the supremum distance.

Corollary 8 (Separation and supremum norm). *In separation distance, we have the lower-bound*

$$\text{sep}(\mathcal{P}_t(x_\star, \cdot), \pi) \geq \frac{\|\mathfrak{z}(\frac{t}{2})\|^2}{1 + \|\mathfrak{z}(\frac{t}{2})\|^2},$$

for any $t \geq 0$. On the other hand, the supremum norm satisfies the upper-bound

$$\max_{x, y \in \mathcal{X}} \left| \frac{\mathcal{P}_t(x, y)}{\pi(y)} - 1 \right| \leq \exp\left(\frac{\|\mathfrak{z}(\frac{t}{2})\|^2}{\rho_\star}\right) - 1.$$

In particular, the total-variation cutoff announced in Corollary 7 also holds in separation distance and in the uniform norm, but without the 2 in the denominator of (10).

1.3 Examples

To illustrate the above results, we now specialize them to two important examples. The first one is the network induced by an arbitrary box $V = [n_1] \times \cdots \times [n_d]$ in the d -dimensional lattice $\mathbb{G} = \mathbb{Z}^d$. We call this network the box of dimensions $n_1 \times \cdots \times n_d$ with *open boundaries*. Note that the special case $d = 1$ is precisely the model studied in [12], for which (3) was established. The second example is the natural *semi-open* variant where the ambient lattice $\mathbb{G} = \mathbb{Z}^d$ is replaced by the d -dimensional semi-lattice $\mathbb{G} = \mathbb{N}^d$, where $\mathbb{N} = \{1, 2, \dots\}$. The special case $d = 1$ is then precisely the one for which (2) was established in [11]. Our results imply the following high-dimensional generalization of those two results.

Corollary 9 (Spectral gap, mixing time and cutoff on boxes). *For the exclusion process with reservoir density $\rho \in (0, 1)$ on a box of dimension $n_1 \times \cdots \times n_d$ with open boundaries, we have*

$$\lambda = 2 \sum_{k=1}^d \left[1 - \cos\left(\frac{\pi}{n_k + 1}\right) \right], \quad \psi(i_1, \dots, i_d) = \prod_{k=1}^d \sqrt{\frac{2}{n_k + 1}} \sin\left(\frac{\pi i_k}{n_k + 1}\right),$$

while for the semi-open variant, we have

$$\lambda = 2 \sum_{k=1}^d \left[1 - \cos \left(\frac{\pi}{2n_k + 1} \right) \right], \quad \psi(i_1, \dots, i_d) = \prod_{k=1}^d \sqrt{\frac{4}{2n_k + 1}} \sin \left(\frac{\pi i_k}{2n_k + 1} \right).$$

Consequently, in both cases, the total-variation mixing time satisfies

$$\frac{\log |V| - cd}{2\lambda} \leq t_{\text{MIX}}(\varepsilon) \leq \frac{\log |V| + c}{2\lambda},$$

where $c = c(\varepsilon, \rho)$ does not depend on d, n_1, \dots, n_d . In particular, there is cutoff at time (10) along any sequence of boxes $(G_n)_{n \geq 1}$ with diverging average logarithmic side-length, i.e.

$$\frac{\log |V_n|}{d_n} \xrightarrow{n \rightarrow \infty} +\infty.$$

For example, if G_n denotes the box of dimension $n \times \dots \times n$ with open boundaries, then

$$t_{\text{MIX}}^{(n)}(\varepsilon) = \frac{n^2 \log n}{2\pi^2} + \mathcal{O}(n^2),$$

where the implicit constant depends only on $\varepsilon, \rho \in (0, 1)$. For the semi-open version, we obtain

$$t_{\text{MIX}}^{(n)}(\varepsilon) = \frac{2n^2 \log n}{\pi^2} + \mathcal{O}(n^2).$$

We emphasize that those estimates are valid even if the ambient dimension $d = d_n$ varies with n . The reservoir density ρ is allowed to vary as well, as long as it is bounded away from 0 and 1. Moreover, a cutoff also occurs at the above times when the distance to equilibrium is measured in relative entropy or Hilbert norm, and twice later when measured in separation or uniform norm.

One could also consider the hybrid case where the box has open boundaries in certain directions and semi-open boundaries in others. The above expressions for λ, ψ adapt in the obvious way. For example, on a $n \times n$ grid with boundaries on three of the four sides, one finds

$$t_{\text{MIX}}(\varepsilon) = \frac{4n^2 \log n}{5\pi^2} + \mathcal{O}(n^2).$$

Finally, we emphasize that boxes are just a particular example chosen for comparison with the existing literature [12, 11]: our general results will yield explicit asymptotics on any network for which we can compute or estimate the spectral statistics λ and ψ . This is particularly easy in the case of Cartesian products, thanks to an obvious tensorization property of Δ (see (24) below).

2 Perturbations of product measures

Without loss of generality, we henceforth assume that the vertex set V is $[n] = \{1, \dots, n\}$. Theorem 1 happens to be a special case of a new and general two-sided estimate, which is valid for all *negatively dependent perturbations* of product measures. We establish this general result in the present section, and will specialize it to the exclusion process with reservoirs in Section 3.

2.1 Framework

To what extent can a random binary vector be statistically distinguished from a perturbed version where a few coordinates have been modified? To formalize this question, consider a random binary vector $X^\star = (X_1^\star, \dots, X_n^\star)$ distributed according to the reference measure $\pi = \mathcal{B}_\rho^{\otimes n}$, and let (Y, Z) be an arbitrary pair of random binary vectors, independent of X^\star . We think of (Y, Z) as *noise variables*, which we use to perturb X^\star as follows: Z indicates which coordinates of X^\star are to be modified, and Y specifies the values to be used for replacement. Specifically, for $i \in [n]$, we set

$$X_i := (1 - Z_i)X_i^\star + Z_iY_i. \quad (11)$$

At least intuitively, the law μ of the perturbed vector $X = (X_1, \dots, X_n)$ should be close to the reference law π as long as the support of Z (the perturbed region) is sufficiently *small* and *delocalized*, regardless of the noise Y . Our goal here is to quantify this statement by a precise estimate on the distance $\|\mu - \pi\|_{\text{TV}}$, as a function of the mean noise vector

$$\mathfrak{z} := (\mathbb{E}[Z_1], \dots, \mathbb{E}[Z_n]).$$

We will see that, for a broad class of perturbations, the correct answer is given in a two-sided way by the Euclidean norm $\|\mathfrak{z}\| = \sqrt{\mathfrak{z}_1^2 + \dots + \mathfrak{z}_n^2}$.

2.2 Upper-bound

A naive way to quantify the impact of the perturbation consists in using the Wasserstein bound

$$\|\mu - \pi\|_{\text{TV}} \leq \mathbb{E} \left[\sum_{i=1}^n |X_i - X_i^\star| \right] \leq \sum_{i=1}^n \mathbb{E}[Z_i] = \langle \mathfrak{z}, \mathbf{1} \rangle. \quad (12)$$

Albeit simple and general, this estimate is too pessimistic, because it only focuses on the *total size* of the perturbed region Z , and not on its *localized/delocalized* nature in space: intuitively, zeroing a fixed, deterministic entry should be much easier to detect (in total-variation distance) than zeroing a uniformly chosen entry. Yet, (12) does not distinguish at all between those two situations.

To get a feeling of how much better the answer could be for delocalized perturbations, let us investigate the elementary but instructive case where the coordinates of Z are independent, while Y is deterministically equal to the extremal vector x_\star . Under this simplifying assumption, the resulting law π is clearly a product measure, so an easy and classical computation yields

$$\|\mu - \pi\|_{\text{TV}} \leq \frac{1}{2} \left\| \frac{\mu}{\pi} - 1 \right\|_{L_\pi^2} = \frac{1}{2} \sqrt{\prod_{i=1}^n \left(1 + \frac{1 - \rho_\star}{\rho_\star} \mathfrak{z}_i^2 \right)} - 1 \quad (13)$$

In the small-perturbation regime where $\mathfrak{z} \rightarrow \mathbf{0}$, the right-hand side decays like $\|\mathfrak{z}\|$, which constitutes a considerable improvement over the linear dependency predicted by (12).

Of course, the explicit computation of the L_π^2 -norm in (13) crucially relies on independence, and getting sharp estimates beyond the product case constitutes a notoriously challenging task. For certain perturbations of spin systems, a sophisticated and powerful approach called *information percolation* was developed by Lubetzky and Sly in an impressive series of papers [21, 22, 24, 23]. Here our main – and much more elementary – finding is that the trivial bound (13) still holds if one assumes that the coordinates of Z are *negatively dependent* (ND) in the following sense:

$$\forall S \subseteq [n], \quad \mathbb{E} \left[\prod_{i \in S} Z_i \right] \leq \prod_{i \in S} \mathbb{E}[Z_i]. \quad (14)$$

More precisely, we have the following universal upper-bound.

Lemma 10 (Upper-bound on the impact of ND perturbations). *If Z is ND, then*

$$\left\| \frac{\mu}{\pi} - 1 \right\|_{L_\pi^2} \leq \sqrt{\prod_{i=1}^n \left(1 + \frac{1 - \rho_\star}{\rho_\star} \mathfrak{z}_i^2 \right) - 1}.$$

Proof. Since X^\star is independent of (Y, Z) and has law $\pi = \mathcal{B}_\rho^{\otimes n}$, we have

$$\mu(y) = \mathbb{E} \left[\left(\prod_{i \notin Z} \mathcal{B}_\rho(y_i) \right) \left(\prod_{i \in Z} \mathbf{1}_{Y_i=y_i} \right) \right].$$

Denoting by (Y', Z') an independent copy of (Y, Z) , we deduce that

$$\mu^2(y) = \mathbb{E} \left[\left(\prod_{i \notin Z} \mathcal{B}_\rho(y_i) \right) \left(\prod_{i \notin Z'} \mathcal{B}_\rho(y_i) \right) \left(\prod_{i \in Z} \mathbf{1}_{Y_i=y_i} \right) \left(\prod_{i \in Z'} \mathbf{1}_{Y'_i=y_i} \right) \right].$$

Dividing through by $\pi(y) = \mathcal{B}_\rho(y_1) \cdots \mathcal{B}_\rho(y_n)$ and simplifying, we obtain

$$\begin{aligned} \frac{\mu^2(y)}{\pi(y)} &= \mathbb{E} \left[\left(\prod_{i \in Z \cap Z'} \frac{1}{\mathcal{B}_\rho(y_i)} \right) \left(\prod_{i \notin Z \cup Z'} \mathcal{B}_\rho(y_i) \right) \left(\prod_{i \in Z} \mathbf{1}_{Y_i=y_i} \right) \left(\prod_{i \in Z'} \mathbf{1}_{Y'_i=y_i} \right) \right] \\ &\leq \mathbb{E} \left[\frac{1}{\rho_\star^{|Z \cap Z'|}} \left(\prod_{i \notin Z \cup Z'} \mathcal{B}_\rho(y_i) \right) \left(\prod_{i \in Z} \mathbf{1}_{Y_i=y_i} \right) \left(\prod_{i \in Z' \setminus Z} \mathbf{1}_{Y'_i=y_i} \right) \right], \end{aligned}$$

because $\min \{\mathcal{B}_\rho(0), \mathcal{B}_\rho(1)\} = \rho_\star$. Summing over all $y \in \{0, 1\}^n$, we arrive at

$$\left\| \frac{\mu}{\pi} \right\|_{L_\pi^2}^2 \leq \mathbb{E} \left[\frac{1}{\rho_\star^{|Z \cap Z'|}} \right].$$

Finally, recalling that Z, Z' are i.i.d. and satisfy (14), we can write

$$\begin{aligned}
\mathbb{E} \left[\frac{1}{\rho_\star^{|Z \cap Z'|}} \right] &= \mathbb{E} \left[\prod_{i=1}^n \left(1 + \frac{1 - \rho_\star}{\rho_\star} Z_i Z'_i \right) \right] \\
&= \sum_{S \subseteq [n]} \left(\frac{1 - \rho_\star}{\rho_\star} \right)^{|S|} \mathbb{E} \left[\prod_{i \in S} Z_i Z'_i \right] \\
&\leq \sum_{S \subseteq [n]} \left(\frac{1 - \rho_\star}{\rho_\star} \right)^{|S|} \left(\prod_{i \in S} \mathbb{E}[Z_i] \right)^2 \\
&= \prod_{i=1}^n \left(1 + \frac{1 - \rho_\star}{\rho_\star} \mathfrak{z}_i^2 \right).
\end{aligned}$$

Since $\|\frac{\mu}{\pi} - 1\|_{L_\pi^2}^2 = \|\frac{\mu}{\pi}\|_{L_\pi^2}^2 - 1$, the claim is proved. \square

Remark 1 (Sharpness). *As already explained (or deduced from a careful examination of the above proof), the inequality in Lemma 10 is an equality when the coordinates of Z are independent, while Y is deterministically set to x_\star . Thus, Lemma 10 is sharp, and states that the L^2 impact of negatively-dependent perturbations is maximized in the product case.*

We now complement the above upper-bound with a matching lower-bound.

2.3 Lower-bound

Lemma 10 states that μ is uniformly close to π (even in the strong L_π^2 sense) whenever $\|\mathfrak{z}\|$ is small. Conversely, we now show that μ is uniformly *far* from π whenever $\|\mathfrak{z}\|$ is large, at least in the extreme case $Y = x_\star$. We here only need the special case $|S| = 2$ of the ND condition (14), namely

$$\forall i \neq j, \quad \text{Cov}(Z_i, Z_j) \leq 0. \quad (15)$$

Lemma 11 (Matching lower bound). *Assume that Z satisfies (15), and that $Y = x_\star$. Then,*

$$\|\mu - \pi\|_{\text{TV}} \geq \frac{\|\mathfrak{z}\|^2}{4 + \|\mathfrak{z}\|^2}.$$

Proof. A simple way to bound total variation from below consists in applying the general inequality

$$\|\mu - \pi\|_{\text{TV}} \geq \frac{(\mathbb{E}[f(X)] - \mathbb{E}[f(X^\star)])^2}{(\mathbb{E}[f(X)] - \mathbb{E}[f(X^\star)])^2 + 2\text{Var}[f(X)] + 2\text{Var}[f(X^\star)]}, \quad (16)$$

to an appropriate *distinguishing statistics* $f: \{0, 1\}^n \rightarrow \mathbb{R}$, see [19, Proposition 7.8]. Here we choose

$$f(x) := \sum_{i=1}^n \mathfrak{z}_i (x_i - \rho),$$

where we recall that $\mathfrak{z}_i := \mathbb{E}[Z_i]$. Since the coordinates of X^* are i.i.d. with law \mathcal{B}_ρ , we have

$$\begin{aligned}\mathbb{E}[f(X^*)] &= 0; \\ \text{Var}[f(X^*)] &= \rho(1-\rho)\|\mathfrak{z}\|^2 \leq \frac{\|\mathfrak{z}\|^2}{4}.\end{aligned}$$

On the other hand, we know that $X_i = (1 - Z_i)X_i^* + Z_i\mathbf{1}_{\rho < 1/2}$ with X^*, Z independent. Thus, $\mathbb{E}[X_i] = \rho + \mathfrak{z}_i(\mathbf{1}_{\rho < 1/2} - \rho)$, $\text{Cov}(X_i, X_j) = (1 - \rho_*)^2 \text{Cov}(Z_i, Z_j) \leq 0$ for $i \neq j$, and

$$\begin{aligned}\mathbb{E}^2[f(X)] &= \|\mathfrak{z}\|^4(\mathbf{1}_{\rho < 1/2} - \rho)^2 \geq \frac{\|\mathfrak{z}\|^4}{4}; \\ \text{Var}[f(X)] &\leq \sum_{i=1}^n \mathfrak{z}_i^2 \text{Var}(X_i) \leq \frac{\|\mathfrak{z}\|^2}{4}.\end{aligned}$$

Inserting those estimates into (16) readily yields the claimed bound. \square

3 Application to exclusion with reservoirs

We now show that our general perturbation theory for negatively-dependent measures applies to the special case of the exclusion process with reservoirs. More precisely,

1. We show in Section 3.1 that at any given time $t \geq 0$ and from any given initial state $x \in \mathcal{X}$, the distribution $\mu = \mathcal{P}_t(x, \cdot)$ of the system is a perturbation of the equilibrium measure π in the sense of Definition (11), for a certain noise vector (Y, Z) that we explicitate.
2. We show in Section 3.2 that the perturbed region Z is negatively-dependent in the sense of Assumption (14), and that its marginals satisfy (4), thereby implying Theorem 1.
3. Finally, in Section 3.3, we detail the arguments that lead from Theorem 1 to Corollaries 2-9.

3.1 Graphical construction

Let $\mathbf{X} = (X(t))_{t \geq 0}$ be an exclusion process with reservoir density ρ on G , starting from an arbitrary state $x \in \mathcal{X}$. The *graphical construction* provides a standardized representation of \mathbf{X} in the form

$$\mathbf{X} := \Psi \left(x, (\Xi_i)_{1 \leq i \leq n}, (\Xi_{ij})_{1 \leq i < j \leq n}, (\xi_k)_{k \geq 1} \right), \quad (17)$$

where the variables $\Xi_1, \dots, \Xi_n, \Xi_{11}, \dots, \Xi_{nn}, \xi_1, \xi_2, \dots$ are independent and as follows:

- Ξ_{ij} is a Poisson point process of rate $c(i, j)$ specifying the exchange times between i and j .
- Ξ_i is a Poisson point process of rate $\kappa(i)$ specifying the resampling times at site i .

- ξ_k is a \mathcal{B}_ρ -variable specifying the new value to be assigned when the k -th resampling occurs.

From this data, the trajectory $\mathbf{X} = (X(t))_{t \geq 0}$ is deterministically obtained as the unique right-continuous, $\{0, 1\}^n$ -valued function which equals x at time 0, is constant outside the locally finite set $\Xi := \Xi_1 \cup \dots \cup \Xi_n \cup \Xi_{11} \cup \dots \cup \Xi_{nn}$, and jumps at $t \in \Xi$ as follows:

- (i) if $t \in \Xi_{ij}$, then $X(t) = (X(t^-))^{i \leftrightarrow j}$.
- (ii) if $t \in \Xi_i$ and t is the k -th smallest point in $\Xi_1 \cup \dots \cup \Xi_n$, then $X(t) = (X(t^-))^{i, \xi_k}$.

This rigorously specifies the measurable map Ψ appearing in (17). We may now couple \mathbf{X} with a stationary process \mathbf{X}^* by setting

$$\mathbf{X}^* := \Psi \left(\zeta, (\Xi_i)_{1 \leq i \leq n}, (\Xi_{i,j})_{1 \leq i < j \leq n}, (\xi_k)_{k \geq 1} \right), \quad (18)$$

where ζ denote a π -distributed random variable independent of $\Xi_1, \dots, \Xi_n, \Xi_{11}, \dots, \Xi_{nn}, \xi_1, \xi_2, \dots$. In order to compare \mathbf{X} and \mathbf{X}^* , we introduce two auxiliary processes:

$$\begin{aligned} \mathbf{Y} &:= \Psi(x, (\Xi_i)_{1 \leq i \leq n}, (\Xi_{i,j})_{1 \leq i < j \leq n}, \mathbf{0}); \\ \mathbf{Z} &:= \Psi(\mathbf{1}, (\Xi_i)_{1 \leq i \leq n}, (\Xi_{i,j})_{1 \leq i < j \leq n}, \mathbf{0}). \end{aligned}$$

Note that \mathbf{Y}, \mathbf{Z} are exclusion processes with reservoir density $\rho = 0$ starting from x and $\mathbf{1}$, respectively. The next lemma shows that at any time $t \geq 0$, the random vector $X(t)$ is the perturbation of $X^*(t)$ induced by the noise $(Y(t), Z(t))$, in the sense of Definition (11).

Lemma 12 (Perturbative structure of the exclusion process). *At any time $t \geq 0$, the random vector $X^*(t)$ is π -distributed and independent of (\mathbf{Y}, \mathbf{Z}) . Moreover, we have*

$$\forall i \in [n], \quad X_i(t) = (1 - Z_i(t)) X_i^*(t) + Z_i(t) Y_i(t). \quad (19)$$

Finally, in this formula, we can replace $Y_i(t)$ by 1 when $x = \mathbf{1}$, and by 0 when $x = \mathbf{0}$.

Proof. The law $\pi = \mathcal{B}_\rho^{\otimes n}$ is trivially preserved under swapping two coordinates or replacing a coordinate with a fresh \mathcal{B}_ρ -distributed variable. Thus, the conditional law of $X^*(t)$ given the point processes $\Xi_1, \dots, \Xi_n, \Xi_{11}, \dots, \Xi_{nn}$ is π . Since \mathbf{Y}, \mathbf{Z} are measurable functions of $\Xi_1, \dots, \Xi_n, \Xi_{11}, \dots, \Xi_{nn}$, the first claim follows. For the second, it suffices to note that the identity (19) holds at time $t = 0$ (both sides being equal to x_i) and is preserved at each discontinuity time $t \in \Xi$ (in case (i) above, the transposition $i \leftrightarrow j$ is simultaneously applied to the four vectors $X^*(t), X(t), Y(t), Z(t)$, and in case (ii) we have $Z_i(t) = 0$ and $X_i(t) = X_i^*(t) = \xi^k$). Finally, observe that the process \mathbf{Y} is equal to \mathbf{Z} in the case $x = \mathbf{1}$, and to $\mathbf{0}$ in the case $x = \mathbf{0}$. \square

Remark 2 (Strong stationary time). *The above lemma readily implies that the random variable $T := \inf \{t \geq 0: \mathbf{Z}(t) = \mathbf{0}\}$ is a strong stationary time, i.e.*

$$\forall (t, y) \in \mathbb{R}_+ \times \mathcal{X}, \quad \mathbb{P}(X(t) = y, T \leq t) = \pi(y)\mathbb{P}(T \leq t).$$

In particular, this classically implies the separation-distance bound

$$\forall t \in \mathbb{R}_+, \quad \text{sep}(\mathcal{P}_t(x, \cdot), \pi) \leq \mathbb{P}(T > t),$$

with equality in the extreme cases $x = \mathbf{0}$ and $x = \mathbf{1}$; see [19, Lemma 6.12 and Proposition 6.14].

3.2 Analysis of the perturbed region

In order to apply the results of Section 2, we must verify that at any time $t \geq 0$, the perturbed region $Z(t)$ meets our negative dependence (ND) requirement (14). By construction, the process $\mathbf{Z} = (Z(t))_{t \geq 0}$ is an exclusion process with reservoir density $\rho = 0$, starting from $Z(0) = \mathbf{1}$. Thus, the claim is a special case of the following general result which, for the conservative variant of the model, was established by Liggett [20, Proposition VIII.1.7] (see Borcea, Brändén and Liggett [5, Proposition 5.1] for a considerable refinement).

Lemma 13 (Negative dependence for exclusion with reservoirs). *Let $\mathbf{X} = (X(t))_{t \geq 0}$ denote an exclusion process with reservoir density $\rho \in [0, 1]$ on an arbitrary network G , and suppose that the initial random vector $X(0)$ is ND. Then, so is $X(t)$ for all $t \geq 0$.*

Proof. Recall that the law $\mu(t)$ of $X(t)$ is given by $\mu(t) = \mu(0)e^{t\mathcal{L}}$, where \mathcal{L} is the generator defined in the introduction. Thus, we want to show that \mathcal{L} is ND-preserving in the following sense:

$$\mu \text{ is ND} \implies \left(\forall t \geq 0, \mu e^{t\mathcal{L}} \text{ is ND} \right).$$

Here is a simple observation that will substantially reduce our task: if $\mathcal{L}_1, \mathcal{L}_2$ are two ND-preserving generators on $\{0, 1\}^n$, then so is any superposition of the form $\mathcal{L} = \lambda_1\mathcal{L}_1 + \lambda_2\mathcal{L}_2$ with $\lambda_1, \lambda_2 \geq 0$. Indeed, Trotter product formula [10, p. 33] asserts that for all $t \geq 0$,

$$e^{t(\lambda_1\mathcal{L}_1 + \lambda_2\mathcal{L}_2)} = \lim_{k \rightarrow \infty} \left(e^{\frac{\lambda_1 t}{k}\mathcal{L}_1} e^{\frac{\lambda_2 t}{k}\mathcal{L}_2} \right)^k,$$

and the claim follows because the ND property is preserved under weak convergence. Consequently, we only need to separately prove Lemma 13 in the following three elementary cases:

- (i) $\mathcal{L}f(x) = f(x^{i,1}) - f(x)$ (creation at i)
- (ii) $\mathcal{L}f(x) = f(x^{i,0}) - f(x)$ (annihilation at i)

(iii) $\mathcal{L}f(x) = f(x^{i \leftrightarrow j}) - f(x)$ (exchange between i and j)

To do so, we suppose that $X(0)$ is ND. In case (i)-(ii), we have the representation

$$X(t) = \begin{cases} X(0)^{i,b} & \text{if } \mathcal{N}(t) \geq 1 \\ X(0) & \text{else,} \end{cases}$$

where $(\mathcal{N}(t))_{t \geq 0}$ denotes a unit-rate Poisson process independent of $X(0)$, and with $b = 1$ in the creation case and $b = 0$ in the annihilation case (ii). Thus, the desired inequality

$$\mathbb{E} \left[\prod_{k \in S} X_k(t) \right] \leq \prod_{k \in S} \mathbb{E} [X_k(t)] \quad (20)$$

is trivially satisfied if the set S does not contain i . On the other hand, if S contains i , then

$$\begin{aligned} \mathbb{E} \left[\prod_{k \in S} X_k(t) \right] &= b(1 - e^{-t}) \mathbb{E} \left[\prod_{k \in S \setminus \{i\}} X_k(0) \right] + e^{-t} \mathbb{E} \left[\prod_{k \in S} X_k(0) \right] \\ &\leq b(1 - e^{-t}) \prod_{k \in S \setminus \{i\}} \mathbb{E} [X_k(0)] + e^{-t} \prod_{k \in S} \mathbb{E} [X_k(0)] \\ &= \prod_{k \in S} \mathbb{E} [X_k(t)], \end{aligned}$$

as desired. In case (iii), we have

$$X(t) = \begin{cases} X(0)^{i \leftrightarrow j} & \text{if } \mathcal{N}(t) \text{ is odd} \\ X(0) & \text{else} \end{cases}$$

In particular, (20) trivially holds if S contains neither i nor j . On the other hand, if S contains i but not j (or vice versa), then writing $\theta_t = \mathbb{P}(\mathcal{N}(t) \text{ is even})$, we have

$$\begin{aligned} \mathbb{E} \left[\prod_{k \in S} X_k(t) \right] &= \theta_t \mathbb{E} \left[X_j(0) \prod_{k \in S \setminus \{i\}} X_k(0) \right] + (1 - \theta_t) \mathbb{E} \left[\prod_{k \in S} X_k(0) \right] \\ &\leq \theta_t \mathbb{E} [X_j(0)] \prod_{k \in S \setminus \{i\}} \mathbb{E} [X_k(0)] + (1 - \theta_t) \prod_{k \in S} \mathbb{E} [X_k(0)] \\ &= \prod_{k \in S} \mathbb{E} [X_k(t)], \end{aligned}$$

as desired. Finally, if S contains both i and j , then

$$\mathbb{E} \left[\prod_{k \in S} X_k(t) \right] = \mathbb{E} \left[\prod_{k \in S} X_k(0) \right] \leq \prod_{k \in S} \mathbb{E} [X_k(0)] = \mathbb{E} [X_i(0)] \mathbb{E} [X_j(0)] \prod_{k \in S \setminus \{i, j\}} \mathbb{E} [X_k(0)],$$

so, the desired inequality (20) boils down to $\mathbb{E}[X_i(0)]\mathbb{E}[X_j(0)] \leq \mathbb{E}[X_i(t)]\mathbb{E}[X_j(t)]$. But since

$$\begin{aligned}\mathbb{E}[X_i(t)] &= (1 - \theta_t)\mathbb{E}[X_i(0)] + \theta_t\mathbb{E}[X_j(0)] \\ \mathbb{E}[X_j(t)] &= (1 - \theta_t)\mathbb{E}[X_j(0)] + \theta_t\mathbb{E}[X_i(0)],\end{aligned}$$

the claim further simplifies to $\theta_t(1 - \theta_t)(\mathbb{E}[X_i(0)] - \mathbb{E}[X_j(0)])^2 \geq 0$, which clearly holds. \square

To establish Theorem 1, it now only remains to prove that the marginals $\mathfrak{z}_i(t) := \mathbb{E}[Z_i(t)]$ satisfy the differential equation (4). This is exactly the special case $\rho = 0, X(0) = \mathbf{1}$ of the following result.

Lemma 14 (Single-site marginals). *Let $\mathbf{X} = (X(t))_{t \geq 0}$ denote an exclusion process with reservoir density $\rho \in [0, 1]$ on an arbitrary network G . Then, the mean function $\mathfrak{z}: t \mapsto \mathbb{E}[X(t)]$ solves*

$$\frac{d\mathfrak{z}}{dt} = \Delta(\mathfrak{z} - \rho).$$

Proof. Dynkin's formula asserts that for any observable $f: \mathcal{X} \rightarrow \mathbb{R}$,

$$\frac{d}{dt}\mathbb{E}[f(X(t))] = \mathbb{E}[(\mathcal{L}f)(X(t))].$$

Now, for the i -th projection $f(x) = x_i$, we readily compute

$$\forall x \in \mathcal{X}, \quad (\mathcal{L}f)(x) = -\kappa(i)(x_i - \rho) + \sum_{j=1}^n c(i, j)(x_j - x_i). \quad (21)$$

Since the right-hand side is precisely the i -th coordinate of the vector $\Delta(x - \rho)$, we are done. \square

3.3 Putting things together

We now have all the ingredients needed to prove the results announced in the introduction.

Proof of Theorem 1. Consider the perturbation $(X^*(t), X(t), Y(t), Z(t))$ defined in Section 3.1 and analyzed in Section 3.2. For this perturbation, Lemma 10 reads

$$\begin{aligned}\left\| \frac{\mathcal{P}_t(x, \cdot)}{\pi} - 1 \right\|_{L^2_\pi} &\leq \sqrt{\prod_{i=1}^n \left(1 + \frac{1 - \rho_\star}{\rho_\star} \mathfrak{z}_i^2 \right) - 1} \\ &\leq \sqrt{\exp\left(\frac{1 - \rho_\star}{\rho_\star} \|\mathfrak{z}\|^2\right) - 1} \\ &\leq \sqrt{\exp\left(\frac{\|\mathfrak{z}\|^2}{\rho_\star}\right) - 1}.\end{aligned}$$

Since $x \in \mathcal{X}$ is arbitrary, the upper-bound is proved. The lower bound is precisely Lemma 11. \square

Proof of Corollary 2. The spectral expansion (5) implies

$$\forall t \geq 0, \quad \|\mathfrak{z}(t)\|^2 \leq |V|e^{-2\lambda t}. \quad (22)$$

Inserting this into the upper-bound provided by Theorem 1 and letting $t \rightarrow \infty$, we see that

$$\max_{x \in \mathcal{X}} \left\| \frac{\mathcal{P}_t(x, \cdot)}{\pi} - 1 \right\|_{L^2_\pi} \leq e^{-\lambda t + o(t)}.$$

Since the spectral gap $\gamma(\mathcal{L})$ of a reversible generator \mathcal{L} coincides with the asymptotic exponential decay rate of the distance to equilibrium, we deduce that $\gamma(\mathcal{L}) \geq \lambda$. Conversely, consider the function $f: \mathcal{X} \rightarrow \mathbb{R}$ defined by $f(x) := \langle \psi, x - \rho \rangle$ for all $x \in \mathcal{X}$. Recalling (21), we have

$$\begin{aligned} (\mathcal{L}f)(x) &= \langle \psi, \Delta(x - \rho) \rangle \\ &= \langle \Delta\psi, x - \rho \rangle \\ &= -\lambda \langle \psi, x - \rho \rangle \\ &= -\lambda f(x), \end{aligned}$$

where the second line uses the symmetry of Δ . Thus, $-\lambda$ is an eigenvalue of \mathcal{L} , hence $\gamma(\mathcal{L}) \leq \lambda$. \square

Proof of Corollary 3. Fix $\varepsilon \in (0, \frac{1}{2})$ and set $t := t_{\text{MIX}}(1 - \varepsilon)$. By the lower-bound in Theorem 1,

$$\frac{\|\mathfrak{z}(t)\|^2}{4 + \|\mathfrak{z}(t)\|^2} \leq 1 - \varepsilon.$$

This implies that $\|\mathfrak{z}(t)\|^2 \leq 4/\varepsilon$, so the upper-bound in Theorem 1 yields

$$\max_{x \in \mathcal{X}} \left\| \frac{\mathcal{P}_t(x, \cdot)}{\pi} - 1 \right\|_{L^2_\pi} \leq \exp\left(\frac{2}{\varepsilon\rho_\star}\right).$$

Since the spectral gap λ of a reversible generator \mathcal{L} coincides with the exponential contraction rate of the L^2_π -distance to equilibrium, we deduce that for all $s \geq 0$,

$$\max_{x \in \mathcal{X}} \left\| \frac{\mathcal{P}_{t+s}(x, \cdot)}{\pi} - 1 \right\|_{L^2_\pi} \leq \exp\left(\frac{2}{\varepsilon\rho_\star} - \lambda s\right).$$

Choosing $s = \frac{3}{\lambda\varepsilon\rho_\star}$ makes the right-hand side less than ε . Recalling the Cauchy-Schwarz inequality

$$\|\mu - \pi\|_{\text{TV}} \leq \frac{1}{2} \left\| \frac{\mu}{\pi} - 1 \right\|_{L^2_\pi}, \quad (23)$$

valid for any probability measure μ on \mathcal{X} , we conclude that $t_{\text{MIX}}(\varepsilon) \leq t + s$. This is exactly the claim, with $c := 3/(\varepsilon\rho_\star)$ \square

Proof of Corollary 4. It is classical that the product condition is necessary for cutoff, see [19, Proposition 18.4]. Conversely, if the product condition holds, then by Corollary 3 we have

$$t_{\text{MIX}}^{(n)}(1/4) \gg t_{\text{MIX}}^{(n)}(\varepsilon) - t_{\text{MIX}}^{(n)}(1 - \varepsilon),$$

for any fixed $\varepsilon \in (0, \frac{1}{4})$, which precisely mean that there is cutoff. \square

Proof of Corollary 5. Fix $c \geq 0$ and set $t := \frac{\log|V|+c}{2\lambda}$. In view of (22), we have $\|\mathfrak{z}(t)\|^2 \leq e^{-c}$. Consequently, the upper-bound in Theorem 1 implies

$$\max_{x \in \mathcal{X}} \left\| \frac{\mathcal{P}_t(x, \cdot)}{\pi} - 1 \right\|_{L_\pi^2} \leq \sqrt{\exp\left(\frac{e^{-c}}{\rho_\star}\right) - 1}.$$

Choosing $c = c(\varepsilon, \rho)$ such that the right-hand side equals ε concludes the proof. \square

Proof of Corollary 6. Fix $\varepsilon \in (0, 1)$ and set $t = t_{\text{MIX}}(\varepsilon)$. By the lower bound in Theorem 1, we have

$$\frac{\|\mathfrak{z}(t)\|^2}{4 + \|\mathfrak{z}(t)\|^2} \leq \varepsilon,$$

which implies $\|\mathfrak{z}(t)\|^2 \leq \frac{4}{1-\varepsilon}$. On the other hand, we have $\|\mathfrak{z}(t)\|^2 \geq \langle \psi, \mathbf{1} \rangle^2 e^{-2\lambda t}$, by (5). Combining these two inequalities, we deduce that

$$t \geq \frac{1}{2\lambda} \log\left(\frac{\langle \psi, \mathbf{1} \rangle^2 (1 - \varepsilon)}{4}\right).$$

This is precisely the claim, with $c := \log \frac{4}{1-\varepsilon}$. \square

Proof of Corollary 7. The claim readily follows from Corollaries 5 and 6. \square

Proof of Corollary 8. An easy and classical consequence of reversibility (see [25, p. 120]) is that

$$\left| \frac{\mathcal{P}_t(x, y)}{\pi(y)} - 1 \right| \leq \left\| \frac{\mathcal{P}_{t/2}(x, \cdot)}{\pi(\cdot)} - 1 \right\|_{L_\pi^2} \left\| \frac{\mathcal{P}_{t/2}(y, \cdot)}{\pi(\cdot)} - 1 \right\|_{L_\pi^2},$$

for all $t \geq 0$ and $x, y \in \mathcal{X}$. Thus, the claimed upper-bound follows from Theorem 1. We now turn to the lower-bound from $x = x_\star$. Using the short-hand $S := \langle Z(t), \mathbf{1} \rangle$, we have by Remark 2,

$$\begin{aligned} \text{sep}(\mathcal{P}_t(x_\star, \cdot), \pi) &= \mathbb{P}(S > 0) \\ &\geq \frac{\mathbb{E}^2[S]}{\text{Var}(S) + \mathbb{E}^2[S]} \\ &\geq \frac{\mathbb{E}[S]}{1 + \mathbb{E}[S]}, \end{aligned}$$

where the second line uses Cantelli's one-sided improvement of Chebychev's inequality, and the third the fact that $\text{Var}(S) \leq \mathbb{E}[S]$ when S is a sum of negatively correlated Bernoulli random variables. To conclude, it only remains to note that

$$\mathbb{E}[S] = \langle \mathfrak{z}(t), \mathbf{1} \rangle = \|\mathfrak{z}(t/2)\|^2,$$

because $\mathfrak{z}(t) = e^{t\Delta}\mathbf{1}$ and Δ is symmetric. □

Proof of Corollary 9. By the Perron-Frobenius theorem, ψ is characterized as the only eigenvector of Δ all of whose coordinates have the same sign. Thus, it is enough to check that the formula for ψ proposed in the claim defines an eigenvector. This is well known (and immediate to check) in the case $d = 1$. The general case follows by observing that the Laplace matrix $\Delta = \Delta^{(n_1, \dots, n_d)}$ of a box of dimensions $n_1 \times \dots \times n_d$ (either with open, or with semi-open boundaries) tensorizes as follows:

$$\Delta^{(n_1, \dots, n_d)} = \Delta^{(n_1)} \oplus \dots \oplus \Delta^{(n_d)}, \tag{24}$$

where \oplus denotes the Kronecker sum of matrices. The rest of the claim is a straightforward application of our general results listed above. □

Acknowledgment. This work was partially supported by Institut Universitaire de France.

References

- [1] David Aldous and James Allen Fill. Reversible markov chains and random walks on graphs, 2002. Unfinished monograph, recompiled 2014, available at <http://www.stat.berkeley.edu/~aldous/RWG/book.html>.
- [2] Gil Alon and Gady Kozma. Comparing with octopi. *Ann. Inst. Henri Poincaré Probab. Stat.*, 56(4):2672–2685, 2020.
- [3] Riddhipratim Basu, Jonathan Hermon, and Yuval Peres. Characterization of cutoff for reversible Markov chains. *Ann. Probab.*, 45(3):1448–1487, 2017.
- [4] L. Bertini, A. De Sole, D. Gabrielli, G. Jona-Lasinio, and C. Landim. Large deviations for the boundary driven symmetric simple exclusion process. *Math. Phys. Anal. Geom.*, 6(3):231–267, 2003.
- [5] Julius Borcea, Petter Brändén, and Thomas M. Liggett. Negative dependence and the geometry of polynomials. *J. Amer. Math. Soc.*, 22(2):521–567, 2009.

- [6] Pietro Caputo, Thomas M. Liggett, and Thomas Richthammer. Proof of Aldous' spectral gap conjecture. *J. Amer. Math. Soc.*, 23(3):831–851, 2010.
- [7] Guan-Yu Chen and Laurent Saloff-Coste. The cutoff phenomenon for ergodic Markov processes. *Electronic Journal of Probability*, 13(3):26–78, 2008.
- [8] Persi Diaconis. The cutoff phenomenon in finite Markov chains. *Proc. Nat. Acad. Sci. U.S.A.*, 93(4):1659–1664, 1996.
- [9] Jian Ding, Eyal Lubetzky, and Yuval Peres. Total variation cutoff in birth-and-death chains. *Probability theory and related fields*, 146(1-2):61–85, 2010.
- [10] Stewart N. Ethier and Thomas G. Kurtz. *Markov processes*. Wiley Series in Probability and Mathematical Statistics: Probability and Mathematical Statistics. John Wiley & Sons, Inc., New York, 1986. Characterization and convergence.
- [11] Nina Gantert, Evita Nestoridi, and Dominik Schmid. Mixing times for the simple exclusion process with open boundaries, 2021.
- [12] Patrícia Gonçalves, Milton Jara, Rodrigo Marinho, and Otávio Menezes. Sharp convergence to equilibrium for the ssep with reservoirs, 2021.
- [13] Jonathan Hermon and Richard Pymar. The exclusion process mixes (almost) faster than independent particles. *Ann. Probab.*, 48(6):3077–3123, 2020.
- [14] Jonathan Hermon and Justin Salez. A version of Aldous' spectral-gap conjecture for the zero range process. *Ann. Appl. Probab.*, 29(4):2217–2229, 2019.
- [15] Hubert Lacoin. The cutoff profile for the simple exclusion process on the circle. *Ann. Probab.*, 44(5):3399–3430, 2016.
- [16] Hubert Lacoin. Mixing time and cutoff for the adjacent transposition shuffle and the simple exclusion. *Ann. Probab.*, 44(2):1426–1487, 2016.
- [17] Hubert Lacoin. The simple exclusion process on the circle has a diffusive cutoff window. *Ann. Inst. Henri Poincaré Probab. Stat.*, 53(3):1402–1437, 2017.
- [18] C. Landim, A. Milanés, and S. Olla. Stationary and nonequilibrium fluctuations in boundary driven exclusion processes. *Markov Process. Related Fields*, 14(2):165–184, 2008.

- [19] David A. Levin and Yuval Peres. *Markov chains and mixing times*. American Mathematical Society, Providence, RI, 2017. Second edition of [MR2466937], With contributions by Elizabeth L. Wilmer, With a chapter on “Coupling from the past” by James G. Propp and David B. Wilson.
- [20] Thomas M. Liggett. *Interacting particle systems*. Classics in Mathematics. Springer-Verlag, Berlin, 2005. Reprint of the 1985 original.
- [21] Eyal Lubetzky and Allan Sly. Cutoff for the Ising model on the lattice. *Invent. Math.*, 191(3):719–755, 2013.
- [22] Eyal Lubetzky and Allan Sly. Cutoff for general spin systems with arbitrary boundary conditions. *Comm. Pure Appl. Math.*, 67(6):982–1027, 2014.
- [23] Eyal Lubetzky and Allan Sly. An exposition to information percolation for the Ising model. *Ann. Fac. Sci. Toulouse Math. (6)*, 24(4):745–761, 2015.
- [24] Eyal Lubetzky and Allan Sly. Information percolation and cutoff for the stochastic Ising model. *J. Amer. Math. Soc.*, 29(3):729–774, 2016.
- [25] Ravi Montenegro and Prasad Tetali. Mathematical aspects of mixing times in Markov chains. *Found. Trends Theor. Comput. Sci.*, 1(3):x+121, 2006.
- [26] Ben Morris. The mixing time for simple exclusion. *Ann. Appl. Probab.*, 16(2):615–635, 2006.
- [27] Roberto Imbuzeiro Oliveira. Mixing of the symmetric exclusion processes in terms of the corresponding single-particle random walk. *Ann. Probab.*, 41(2):871–913, 2013.
- [28] Y Peres. American institute of mathematics (AIM) research workshop “sharp thresholds for mixing times” (Palo Alto, December 2004). *Summary available at <http://www.aimath.org/WWN/mixingtimes>*.
- [29] Justin Salez. Cutoff for non-negatively curved markov chains, 2021.
- [30] Frank Spitzer. Interaction of Markov processes. *Advances in Math.*, 5:246–290 (1970), 1970.
- [31] David Bruce Wilson. Mixing times of Lozenge tiling and card shuffling Markov chains. *Ann. Appl. Probab.*, 14(1):274–325, 2004.
- [32] Horng-Tzer Yau. Relative entropy and hydrodynamics of Ginzburg-Landau models. *Lett. Math. Phys.*, 22(1):63–80, 1991.