



Reducing metastable continuous-space Markov chains to Markov chains on a finite set

Nils Berglund

► To cite this version:

Nils Berglund. Reducing metastable continuous-space Markov chains to Markov chains on a finite set. 2023. hal-04042150

HAL Id: hal-04042150

<https://hal.science/hal-04042150>

Preprint submitted on 12 May 2023

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.



Distributed under a Creative Commons Attribution 4.0 International License

Reducing metastable continuous-space Markov chains to Markov chains on a finite set

Nils Berglund

Abstract

We consider continuous-space, discrete-time Markov chains on \mathbb{R}^d , that admit a finite number N of metastable states. Our main motivation for investigating these processes is to analyse random Poincaré maps, which describe random perturbations of ordinary differential equations admitting several periodic orbits. We show that under a few general assumptions, which hold in many examples of interest, the kernels of these Markov chains admit N eigenvalues exponentially close to 1, which are separated from the remainder of the spectrum by a spectral gap that can be quantified. Our main result states that these Markov chains can be approximated, uniformly in time, by a finite Markov chain with N states. The transition probabilities of the finite chain are exponentially close to first-passage probabilities at neighbourhoods of metastable states, when starting in suitable quasistationary distributions.

Date. March 22, 2023. Version with corrected typos, May 12, 2023.

2020 Mathematical Subject Classification. 60J05, 37H05 (primary), 60J35, 34F05 (secondary).

Keywords and phrases. Random Poincaré map, metastability, quasistationary distributions, trace process, large deviations.

1 Introduction and informal statement of results

In this work, we are concerned with continuous-space, discrete-time Markov chains, depending on a small parameter $\sigma \geq 0$, which reduce to a deterministic map when $\sigma = 0$. Our main motivation for considering these processes is related to the notion of random Poincaré maps. Consider a stochastic differential equation (SDE) in \mathbb{R}^d , which is a weak-noise perturbation of an ordinary differential equation (ODE), admitting a finite number of asymptotically stable periodic orbits. In the deterministic limiting case, it is useful to introduce a surface of section Σ , transverse to the flow, and to study the sequence of returns of an orbit to Σ . This allows in particular to study stability and bifurcations of periodic orbits in a systematic way.

A similar notion can be introduced in the stochastic case, taking some care in defining what one means by returns to Σ : one has to require that sample paths make some excursion away from Σ between returns, in order to avoid accumulation of intersection points. To our knowledge, this notion appeared first in the works [WK90] by Weiss and Knobloch, and [HM09] by Hitczenko and Medvedev. In [BL12], random Poincaré maps were used to study the distribution of small oscillations in the stochastic FitzHugh–Nagumo equation. In [HM13], they allowed to characterise the effect of noise on elliptic bursting. Random Poincaré maps also proved useful in the analysis of mixed-mode oscillations in systems such as the stochastic Koper model, featuring a folded-node singularity [BGK15], and in determining the distribution of transition points through an unstable periodic orbit [BG14].

The work [BB17] initiated a more systematic study of random Poincaré maps, from the point of view of spectral theory. Its main result is that under a metastable hierarchy assumption

on the N stable periodic orbits, the kernels describing random Poincaré maps have exactly N eigenvalues that are exponentially close to 1. In addition, the remaining part of the spectrum is separated from those N leading eigenvalues by a spectral gap, scaling like the logarithm of the noise intensity. Asymptotic expressions for the leading eigenvalues and eigenfunctions in terms of committor functions were also obtained in [BB17].

The present work concerns a more general class of continuous-space Markov chains, which contain random Poincaré maps as a particular case, but are not limited to them. For instance, they also include randomly perturbed deterministic maps, of the form

$$X_{n+1} = \Pi(X_n) + \sigma \xi_{n+1} ,$$

where Π is a deterministic map defined on a subset of \mathbb{R}^d , and the ξ_n are independent, identically distributed random variables. Deterministic iterated maps are common in applications such as population dynamics and epidemiology, and it is natural to study their perturbation by weak noise.

Main results. We now give an informal statement of the main assumptions and results of this work. A precise formal statement is given in Sections 2 and 3 below. We consider Markov chains on $\mathcal{X}_0 \subset \mathbb{R}^d$, with kernel K_σ , where σ measures the noise intensity. In the deterministic case $\sigma = 0$, we assume that

$$K_0(x, A) = \mathbb{1}_{\{\Pi(x) \in A\}} ,$$

for a deterministic map $\Pi : \mathcal{X}_0 \rightarrow \mathcal{X}_0$. Our main assumptions are the following.

1. **Deterministic limit:** The deterministic map Π leaves a compact set $\mathcal{X} \subset \mathcal{X}_0$ invariant. It has N asymptotically stable fixed points in \mathcal{X} , and all its other limit sets are unstable fixed points. The aim of this assumption is to ensure that the asymptotic dynamics spends most of the time near a finite set of fixed points.
2. **Large-deviation principle:** For positive σ , the kernel K_σ has a smooth density, and it obeys a large-deviation principle with good rate function I . The rate function I will be used to define a notion of quasipotential, that describes the exponential asymptotics of transition times between metastable sets.
3. **Recurrence:** For $\sigma > 0$, the Markov chain is positive Harris recurrent. This means in particular that it will reach any open set in a time having finite expectation. In particular, when starting anywhere in the compact set \mathcal{X} , the expected return time to \mathcal{X} is bounded by a finite quantity $E_{\mathcal{X}}(\sigma)$. This assumption is needed for the existence of a spectral gap, between the N leading eigenvalues, and the remainder of the spectrum of K_σ .
4. **Positivity:** The process satisfies a uniform positivity condition in the neighbourhood of the stable fixed points of Π . This is a more technical property, defined in Section 2.6 below, which essentially amounts to a lower bound of Doeblin type on transition densities. This assumption guarantees that the process conditioned on remaining near a stable fixed point relaxes to a so-called quasistationary distribution.

While the first two assumptions are quite natural, it may seem more difficult to ensure the last two assumptions. However, we will show in Section 4 that they are actually satisfied under quite weak conditions for the processes we are interested in, namely random Poincaré maps and randomly perturbed iterated maps. In particular, we will show that $E_{\mathcal{X}}(\sigma)$ is at most of

order $\log(\sigma^{-1})$ in these cases, while the large-deviation principle implies that it is always at least sub-exponential in σ , in the sense that $e^{-\eta/\sigma^2} E_{\mathcal{X}}(\sigma) \rightarrow 0$ as $\sigma \rightarrow 0$ for any $\eta > 0$.

Our first main result reads as follows.

Proposition 1.1 (Proposition 3.1). *For sufficiently small positive σ , the kernel K_{σ} has exactly N eigenvalues which are exponentially close to 1. All remaining eigenvalues have a modulus smaller than $\varrho = e^{-c/E_{\mathcal{X}}(\sigma)}$ for some constant $c > 0$.*

Note that [BB17, Thm. 3.2] provides sharper bounds on the N leading eigenvalues, under a more restrictive metastable hierarchy assumption (essentially, all transitions between fixed points should happen on different exponential timescales). In that case, those eigenvalues can also be shown to be real. Here, however, we do not make such an assumption.

The existence of a spectral gap already shows that after a time of order $1/E_{\mathcal{X}}(\sigma)$, the process will be close to a finite-dimensional subspace of the space of measures on \mathcal{X}_0 . The difficulty is that it is not straightforward to connect this finite-dimensional space to quantities that have a probabilistic interpretation. Our second main result provides such a connection. To state it, we introduce the *trace process* $(X_{\tau^+,n})_{n \geq 0}$ of the Markov chain. Here τ^+,n denotes the time of n th return of the chain to a suitably defined neighbourhood \mathcal{M} of the set of stable fixed points, given by a union of neighbourhoods B_i of these points.

Theorem 1.2 (Theorem 3.3). *Let $m(\sigma)$ be a function satisfying*

$$\lim_{\sigma \rightarrow 0} \sigma^2 \log m(\sigma) = \theta$$

for a sufficiently small parameter $\theta > 0$. Then the probability of $X_{\tau^+,nm(\sigma)}$ belonging to B_j , when starting in B_i , is well-approximated by the probability of a Markov chain $(Y_n)_{n \geq 0}$ with values in $\{1, \dots, N\}$ being in state j . The transition probabilities of this Markov chain are given, up to exponentially small multiplicative errors, by the probability of the trace process first hitting B_j at time $m(\sigma)$, when starting in a quasistationary distribution on B_i .

More precisely, we will show that there exists a linear map \mathcal{L} between measures on \mathcal{M} and measures on $\{1, \dots, N\}$, such that $\mathbb{P}Y_n^{-1} = \mathcal{L}(\mathbb{P}X_{\tau^+,nm(\sigma)}^{-1})$ for all $n \in \mathbb{N}_0$.

We refer to the statement of Theorem 3.3 below for a precise formulation of what we mean by being well-approximated. Essentially, the difference between the distributions of the two processes is bounded uniformly in time by an exponentially small quantity. The result is thus mostly useful on long timescales, when the process has had an opportunity to explore several metastable states. Then our result states that whenever the finite Markov chain Y_n is in state j with a probability that is not exponentially small, the process $X_{\tau^+,nm(\sigma)}$ will belong to B_j with a probability that is exponentially close to it.

Related results. The problem of approximating Markov processes by Markov chains on a finite set has been investigated for a long time, in particular in the case of SDEs. The idea is already present in the monograph [FW98] by Freidlin and Wentzell, where Markov chain approximations are used for instance to investigate the exit problem, and to approximate invariant measures. There is however no quantitative statement on how well the Markov chain approximates the original process directly.

The works [BEGK04, BGK05] by Bovier, Eckhoff, Gaynard and Klein investigate reversible diffusion processes, governed by gradient SDEs of the form

$$dx_t = -\nabla V(x_t) dt + \sigma dW_t, \quad (1.1)$$

where V is a confining multiwell potential. These articles use a potential-theoretic approach, that was originally limited to the reversible case, but was extended by Landim, Mariani and Seo to more general diffusions [LMS19]. One result in [BGK05] is that under a metastable hierarchy assumption, the expectations of transition times between certain well-chosen metastable sets are close to expectations of similar transitions in a finite Markov chain.

Because of their importance in simulation algorithms in molecular dynamics, in particular in kinetic Monte Carlo algorithms [Vot07], these results prompted a series of works aiming at obtaining precise descriptions of the exit location of solutions of SDEs from metastable sets, see in particular the works by Di Gesù, Lelièvre, Le Peutrec and Nectoux [DGLLPN19, DGLLPN20, LLPN22]. These authors also emphasized the importance of quasistationary distributions (QSDs) in metastable states for the reduction problem [DGLLPN16]. See for instance the work [CV23] by Champagnat and Villemonais for a recent review on QSDs. In parallel, results on the spectrum of reversible diffusions of the form (1.1) have been extended to non-reversible diffusions by Le Peutrec and Michel [LPM20], using methods from semiclassical analysis.

In a different direction, many works have investigated the metastable behaviour of continuous-time Markov chains on countable sets, arising either in statistical physics, or as spatial discretisation of SDEs. In [BL10], Beltrán and Landim introduced in particular the idea of a trace process to obtain a reduced description of the dynamics, while in [BL15] they introduced a martingale method to study the convergence of sequences of such processes with increasingly large state spaces. In [LLM18], Landim, Loulakis and Mouragui obtained convergence of finite-dimensional distributions of the so-called order parameter to those of a finite Markov chain. See [Lan19] for an overview of these results, and [LS18] for related results on sequences of discretisations of SDEs.

Finally, a recent approach based on solutions of Poisson equations managed to show convergence of time-rescaled solutions to metastable SDEs to finite Markov chains, in the limit of the noise intensity going to zero. See the work [RS18] by Rezakhanlou and Seo for the reversible case, and the work [LS22] by Lee and Seo for non-reversible cases with known invariant measure, of Gibbs type. An overview is found in [Seo20].

The main difference between the present work and those mentioned above, apart from the fact that it concerns continuous-time Markov chains instead of SDEs or Markov chains on countable spaces, is that it splits the approximation question into two separate problems. The first one, which is the main focus of this work, is to show that there exists a finite Markov chain that provides a good approximation to the metastable process. The second one is to obtain sharp asymptotics on transition probabilities of the finite Markov chain. This question is addressed here only in the sense of logarithmic equivalence, which naturally follows from the large-deviation principles. Sharper asymptotics will hopefully be determined in the future. The present results show, however, that it is sufficient to obtain such sharper asymptotics when starting in suitable QSDs.

We finally remark that there are many works analysing the dynamics of singularly perturbed Markov chains, such as $(Y_n)_{n \geq 0}$. See for instance [Sch68, HH92, AL99, YZ05, BLR16, FK17].

Structure of the paper. Section 2 contains the detailed set-up of the Markov processes we are interested in, states the four main assumptions, and introduces useful objects such as the trace process and quasistationary distributions. Section 3 contains the precise statements of the two main results mentioned above. In Section 4, we show that most of the main assumptions do hold quite generally in the case of the two main applications we have in mind, namely randomly perturbed iterated maps and random Poincaré maps. Sections 5 and 6 contain the

proofs of the two main results. Finally, the appendix contains the proofs of some auxiliary results used in Sections 2 and 4.

Acknowledgments. The author would like to thank Manon Baudel for useful discussions that inspired the early stages of this work. This work was supported by the ANR project PERISTOCH, ANR-19-CE40-0023.

2 Set-up and assumptions

Let $\mathcal{X}_0 \subset \mathbb{R}^d$ be an open, connected domain, and denote its Borel σ -algebra $\mathcal{B}(\mathcal{X}_0)$ by \mathcal{S}_0 . Our object of study are families $\{K_\sigma\}_{0 \leq \sigma < \sigma_0}$ of Markov kernels on $(\mathcal{X}_0, \mathcal{S}_0)$, such that K_0 is a singular kernel, called the deterministic limit, while for $\sigma > 0$ the kernel K_σ is positive Harris recurrent and admits a continuous density.

We denote by $(X_n^\sigma)_{n \geq 0} = (X_n)_{n \geq 0}$ the Markov chain with kernel K_σ , starting from some specified initial distribution μ , and write $\mathbb{P}^\mu\{\cdot\}$ and $\mathbb{E}^\mu[\cdot]$ for the associated law and expectations. If $\mu = \delta_x$, we simply write $\mathbb{P}^x\{\cdot\}$ and $\mathbb{E}^x[\cdot]$. Given $A \in \mathcal{S}_0$ we will sometimes use the notation

$$\mathbb{E}^A[\cdot] = \sup_{x \in A} \mathbb{E}^x[\cdot].$$

For any set $A \in \mathcal{S}_0$, we denote by

$$\tau_A(x) = \inf\{n \geq 0 : X_n \in A\} \quad \text{and} \quad \tau_A^+(x) = \inf\{n \geq 1 : X_n \in A\}$$

the hitting time of A and return time to A of $(X_n)_{n \geq 0}$ starting in x (with the convention that $\inf \emptyset = \infty$). Note that $\tau_A^+(x) = \tau_A(x)$ whenever $x \notin A$, while $0 = \tau_A(x) < \tau_A^+(x)$ when $x \in A$. We will drop the argument x whenever it is clear from the context.

The kernel K_σ induces two Markov semigroups in the standard way: for any bounded measurable test function $\varphi \in L^\infty$, we have

$$(K_\sigma \varphi)(x) = \int_{\mathcal{X}_0} K_\sigma(x, dy) \varphi(y) = \mathbb{E}^x[\varphi(X_1)],$$

while for any (signed) measure $\mu \in L^1$ we have

$$(\mu K_\sigma)(dy) = \int_{\mathcal{X}_0} \mu(dx) K_\sigma(x, dy) = \mathbb{P}^\mu\{X_1 \in dy\}.$$

For $n \in \mathbb{N}$, we denote by K_σ^n the n -fold kernel, defined by $K_\sigma^1 = K_\sigma$ and

$$K_\sigma^{n+1}(x, A) = \int_{\mathcal{X}_0} K_\sigma^n(x, dz) K_\sigma(z, A) = \mathbb{P}^x\{X_{n+1} \in A\}.$$

We are going to need a number of more precise assumptions, which are detailed in the next subsections. These concern the deterministic limit kernel K_0 , a large-deviation principle for $\sigma \rightarrow 0$, as well as positive Harris recurrence and local uniform positivity assumptions guaranteeing convergence to a unique invariant distribution.

2.1 Singular deterministic limit

Let $\Pi : \mathcal{X}_0 \rightarrow \mathcal{X}_0$ be a map of class \mathcal{C}^2 . Note that we do not assume that Π is invertible. We would like K_0 to describe the deterministic dynamical system $X_{n+1} = \Pi(X_n)$, which amounts to setting

$$K_0(x, A) = \mathbb{1}_{\{\Pi(x) \in A\}} \quad \forall A \in \mathcal{S}_0.$$

The forward semigroup then takes the form of the composition operator

$$(K_0 \varphi)(x) = \int_{\mathcal{X}_0} \mathbb{1}_{\{\Pi(x) \in dy\}} \varphi(y) = (\varphi \circ \Pi)(x),$$

sometimes called *Koopman operator* in the physics literature. The backward semigroup is given by the *pushforward* operator

$$(\mu K_0)(A) = \int_{\mathcal{X}_0} \mu(dx) \mathbb{1}_{\{\Pi(x) \in A\}} = \mu(\Pi^{-1}(A)),$$

which is known as the *transfer operator* or *Ruelle–Perron–Frobenius operator* in dynamical systems theory.

For $A \in \mathcal{S}_0$, we write $\Pi^0(A) = A$, and define inductively, for any $n \geq 1$, $\Pi^n(A) = \Pi \circ \Pi^{n-1}(A)$ and $\Pi^{-n}(A) = \{x \in \mathcal{X}_0 : \Pi^n(x) \in A\}$ (note that the last set may be empty). The ω -*limit set* $\omega(x)$ of $x \in \mathcal{X}_0$ is the set of accumulation points of the *forward orbit* $(\Pi^n(x))_{n \geq 0}$ as $n \rightarrow \infty$. The α -*limit set* $\alpha(x)$ of x is defined as the set of accumulation points of the *backward orbit* $(\Pi^{-n}(x))_{n \geq 0}$.

A *fixed point* x^* of Π (that is, a point $x^* \in \mathcal{X}_0$ satisfying $\Pi(x^*) = x^*$) is called *linearly asymptotically stable* if the Jacobian matrix $\partial_x \Pi(x^*)$ has a spectral radius strictly smaller than 1, and *linearly unstable* if it has a spectral radius strictly larger than 1.

Assumption DET (Deterministic limit). There exists a bounded, open connected set $\mathcal{X} \subset \mathcal{X}_0$ such that $\Pi(\mathcal{X}) \subset \mathcal{X}$. The map Π admits finitely many limit sets in \mathcal{X} , which are either linearly asymptotically stable fixed points, denoted x_1^*, \dots, x_N^* , or linearly unstable fixed points. ♣

For each $j = 1, \dots, N$, we let B_j be a closed set, containing x_j^* in its interior, and such that $\Pi(B_j) \subset B_j$. We will assume that the diameter of all B_j is bounded by a constant $\delta > 0$, which we are going to take small, but which is independent of σ . We denote by

$$\mathcal{M} = \bigcup_{j=1}^N B_j$$

the *metastable set* of the process.

Remark 2.1. The case of Π admitting finitely many periodic points x_i^* of bounded minimal period m_i as ω -limit sets (i.e., $\Pi^{m_i}(x_i^*) = x_i^*$ and $\Pi^n(x_i^*) \neq x_i^*$ for $1 \leq n \leq m_i - 1$) can be covered by considering the iterated kernel K_σ^m instead of K_σ , where m is the least common multiple of the periods m_i of all periodic points. ◇

2.2 Large-deviation principle

Our second assumption concerns the behaviour of the kernel K_σ for small positive σ .

Assumption LDP (Large-deviation principle). The kernel K_σ satisfies a large-deviation principle (LDP) with good rate function I . That is, there exists a lower semi-continuous function $I : \mathcal{X}_0 \times \mathcal{X}_0 \rightarrow \mathbb{R}_+$, with compact level sets, such that

$$\liminf_{\sigma \rightarrow 0} \sigma^2 \log K_\sigma(x, O) \geq - \inf_{y \in O} I(x, y) \quad (2.1)$$

holds for any open set $O \in \mathcal{S}_0$ and any $x \in \mathcal{X}_0$, and

$$\limsup_{\sigma \rightarrow 0} \sigma^2 \log K_\sigma(x, C) \leq - \inf_{y \in C} I(x, y) \quad (2.2)$$

holds for any closed set $C \in \mathcal{S}_0$ and any $x \in \mathcal{X}_0$. Furthermore, $I(x, y) = 0$ if and only if $y = \Pi(x)$, and I is continuous at (x^*, x^*) whenever $\Pi(x^*) = x^*$. \clubsuit

With any sequence (x_0, x_1, \dots, x_n) of points in \mathcal{X}_0 , we associate the rate function

$$I(x_0, x_1, \dots, x_n) = \sum_{j=1}^n I(x_{j-1}, x_j).$$

Then the probability of the Markov chain visiting small neighbourhoods of x_0, \dots, x_n in this particular order is logarithmically equivalent to $e^{-I(x_0, \dots, x_n)/\sigma^2}$. The *quasipotential* between two points x and y is then defined as

$$V(x, y) = \inf_{n \geq 1} \inf_{x_1, \dots, x_{n-1} \in \mathcal{X}_0} I(x, x_1, \dots, x_{n-1}, y). \quad (2.3)$$

It represents the cost of going from x to y in arbitrary time.

For $1 \leq i \neq j \leq N$, we denote by

$$H(i, j) = V(x_i^*, x_j^*) \quad (2.4)$$

the quasipotential between the stable fixed points x_i^* and x_j^* , and we define

$$H_0 = \min_{i \neq j} H(i, j). \quad (2.5)$$

An important role will be played by the so-called *committor functions* $\mathbb{P}^x\{\tau_{B_j}^+ < \tau_{B_i}^+\}$ between different balls B_i and B_j . The following standard consequence of the LDP shows that for starting points $x \in B_i$, $\mathbb{P}^x\{\tau_{B_j}^+ < \tau_{B_i}^+\}$ behaves like $e^{-H(i, j)/\sigma^2}$. We give its proof in Appendix A.

Proposition 2.2 (Large-deviation estimates on committor functions). *For any $\eta > 0$, there exists $\delta_0 > 0$ such that, if the diameter of the sets B_i satisfies $\delta < \delta_0$, then*

$$\begin{aligned} \liminf_{\sigma \rightarrow 0} \sigma^2 \log \mathbb{P}^x\{\tau_{B_j}^+ < \tau_{B_i}^+\} &\geq -H(i, j) - \eta, \\ \limsup_{\sigma \rightarrow 0} \sigma^2 \log \mathbb{P}^x\{\tau_{B_j}^+ < \tau_{B_i}^+\} &\leq -H(i, j) + \eta \end{aligned}$$

holds for all $x \in B_i$.

2.3 Positive Harris recurrence

Properties of irreducibility, recurrence and positive recurrence can be defined in terms of hitting and return times, as discussed in [MT92]. Given a σ -finite reference measure μ such that $\mu(\mathcal{X}_0) > 0$, the process $(X_n)_{n \geq 0}$ is μ -irreducible if $\mathbb{P}^x\{\tau_A^+ < \infty\} > 0$ whenever $\mu(A) > 0$. It is *Harris recurrent* if $\mathbb{P}^x\{\tau_A^+ < \infty\} = 1$ whenever $\mu(A) > 0$, which is equivalent to the process visiting A infinitely often. In this case, it is known [Num84] that the process admits an essentially unique invariant measure π_0 , and for any $A \in \mathcal{S}_0$ with $\pi_0(A) > 0$ and any measurable $f \geq 0$, one has

$$\pi_0(f) := \int_{\mathcal{X}_0} f(x) \pi_0(dx) = \int_A \pi_0(dx) \mathbb{E}^x \left[\sum_{n=1}^{\tau_A^+} f(X_n) \right]. \quad (2.6)$$

If π_0 can be normalized to a probability measure, the process is called *positive Harris recurrent*. Setting $f = 1$ in (2.6) shows that this is the case if $\mathbb{E}^A[\tau_A^+] < \infty$ for some A with $0 < \mu(A) < \infty$. An important role will be played by the quantity

$$E_{\mathcal{X}}(\sigma) := \mathbb{E}^{\mathcal{X}}[\tau_{\mathcal{X}}^+] = \sup_{x \in \mathcal{X}} \mathbb{E}^x[\tau_{\mathcal{X}}^+] . \quad (2.7)$$

Here we will make the simplifying assumption, which is motivated by the applications we have in mind, that for $\sigma > 0$, $K_{\sigma}(x, \cdot)$ is absolutely continuous with respect to Lebesgue measure. This will allow us to take Lebesgue measure as reference measure. By further assuming that the density k_{σ} of K_{σ} is continuous and strictly positive in $\mathcal{X} \times \mathcal{X}$, we guarantee that $K_{\sigma}(x, A) > 0$ whenever $A \subset \mathcal{X}$ has positive Lebesgue measure, which amounts to an ellipticity condition. In addition, these A are *petite* sets in the sense of [MT92, Sect. 3].

Assumption REC (Density and positive Harris recurrence). Whenever $\sigma > 0$, K_{σ} admits a density k_{σ} with respect to Lebesgue measure, that is,

$$K_{\sigma}(x, A) = \int_A k_{\sigma}(x, y) \, dy$$

for any $x \in \mathcal{X}_0$ and any $A \in \mathcal{S}_0$. The density k_{σ}^n of K_{σ}^n is continuous and strictly positive in \mathcal{X} for all $n \in \mathbb{N}$. Furthermore, K_{σ} is positive Harris recurrent (with respect to Lebesgue measure), and there exists $\sigma_0 > 0$ such that $E_{\mathcal{X}}(\sigma) < \infty$ for all $\sigma \in (0, \sigma_0]$. \clubsuit

It follows from [MT92, Thm. 4.6] that a sufficient condition for positive Harris recurrence is that there exist a *Lyapunov function* $U : \mathcal{X}_0 \rightarrow \mathbb{R}_+$, going to infinity as $x \rightarrow \infty$, and constants $\varepsilon > 0$, $a \geq 0$ satisfying the discrete drift condition

$$(K_{\sigma}U)(x) \leq U(x) - \varepsilon + a\mathbb{1}_{\{x \in \mathcal{X}\}} . \quad (2.8)$$

In addition, [MT92, Thm. 4.3] shows that (2.8) implies the bound

$$\mathbb{E}^x[\tau_{\mathcal{X}}^+] \leq \frac{1}{\varepsilon} U(x) \quad \forall x \in \mathcal{X}_0 ,$$

so that $E_{\mathcal{X}}(\sigma)$ is indeed finite. Note that if U is a Lyapunov function for K_0 , then it is a good candidate for being a Lyapunov function for small positive σ .

Remark 2.3. An alternative to assuming the existence of a Lyapunov function is to work with the process conditioned on staying in the set \mathcal{X} forever, via Doob's h -transform (see for instance [BB17, App. B]). This has a negligible effect on spectral-theoretic results if we assume that there exists a constant $\theta > 0$ such that

$$\min_{1 \leq i \leq N} V(x_i^{\star}, y) \geq \max_{1 \leq i \neq j \leq N} H(i, j) + \theta$$

holds for all $y \in \mathcal{X}_0 \setminus \mathcal{X}$. \diamond

The following result shows that the LDP also provides a rough estimate, of order e^{η/σ^2} with arbitrarily small $\eta > 0$, for the mean hitting time of the metastable set \mathcal{M} when starting in \mathcal{X} . Its proof is postponed to Appendix A.

Proposition 2.4 (Mean hitting time of \mathcal{M}). *For any $\eta > 0$, there exist $\sigma_0, \delta_0 > 0$ such that one has $E_{\mathcal{X}}(\sigma) \leq e^{\eta/\sigma^2}$, provided $0 < \sigma < \sigma_0$ and the diameter of the B_i is bounded by δ_0 .*

We will however see that in many practical situations, it is possible to show that $E_{\mathcal{X}}(\sigma)$ is much smaller, typically of order $\log(\sigma^{-1})$, which yields better spectral gap estimates.

2.4 Trace process

A very important process is going to be the *trace process* on a recurrent set $A \in \mathcal{S}_0$ (i.e., such that $\mathbb{P}^x\{\tau_A^+ < \infty\} = 1$ for all $x \in A$).

Definition 2.5 (Trace process). *Let A be a positive recurrent set. The trace process on A is defined as the Markov chain monitored only when staying in A . Its transition kernel is given by*

$${}_AK_\sigma(x, B) = \mathbb{P}^x\{X_{\tau_A^+} \in B\}$$

for any $B \in \mathcal{S}_0$. We denote this process by ${}_A(X_n)_{n \geq 0}$.

Note that owing to the strong Markov property, ${}_AK_\sigma$ is a markovian kernel, meaning that ${}_AK_\sigma(x, A) = 1$ for all $x \in A$. Since A is recurrent, it is also a stochastic kernel on A . It can be rewritten in the form

$${}_AK_\sigma(x, B) = \sum_{n \geq 1} \mathbb{P}^x\{\tau_A^+ = n, X_n \in B\},$$

and thus for $\sigma > 0$ it admits the density

$${}_Ak_\sigma(x, y) = \sum_{n \geq 1} \mathbb{P}^x\{\tau_A^+ = n\} k_\sigma^n(x, y) \mathbb{1}_{\{x \in A, y \in A\}}. \quad (2.9)$$

Assume from now on that A is positive recurrent (i.e., $\mathbb{E}^x[\tau_A^+] < \infty$ for all $x \in A$). Applying (2.6) to $f = \mathbb{1}_B$ for $B \subset A$, we obtain

$$\pi_0(B) = \int_A \pi_0(dx) \mathbb{E}^x \left[\sum_{n=1}^{\tau_A^+} \mathbb{1}_{\{X_n \in B\}} \right] = \int_A \pi_0(dx) \mathbb{P}^x\{X_{\tau_A^+} \in B\},$$

showing that the restriction of π_0 to A is invariant under the trace process. It follows that the measure ${}_A\pi_0$ defined by

$${}_A\pi_0(B) = \frac{\pi_0(B)}{\pi_0(A)} \quad \forall B \in \mathcal{S}_0: B \subset A$$

is an invariant probability measure of the trace process on A .

Remark 2.6 (Transitivity of the trace). One easily checks that the trace enjoys the following transitivity property: if $B \subset A$, then ${}_B({}_A(X_n))_{n \geq 0} = {}_B(X_n)_{n \geq 0}$. \diamond

It will be more convenient to work with kernels defined on a bounded set. This can be achieved by considering, instead of the original kernel K_σ , the kernel ${}_XK_\sigma$ of the trace process on the bounded set \mathcal{X} , which contains essentially the same dynamic information owing to Assumption REC.

To lighten the notation, we will from now on simply write K instead of ${}_XK_\sigma$, the parameter $\sigma > 0$ being always fixed at a sufficiently small value. The density of K , denoted by k , is continuous and strictly positive in \mathcal{X} by Assumption REC. The Borel σ -algebra of \mathcal{X} will be denoted $\mathcal{B}(\mathcal{X}) = \mathcal{S}$, and for any $A \in \mathcal{S}$ we write A^c instead of $\mathcal{X} \setminus A$.

Since \mathcal{X} is bounded and k is continuous, K is a compact operator (that is, it maps every closed set in \mathcal{S} to a relatively compact set, i.e., a set with compact closure). The Riesz–Schauder theorem [RS80, Thm. VI.15] ensures that K has discrete spectrum, with all eigenvalues except possibly 0 having finite multiplicity. The eigenvalues are roots of the Fredholm determinant, introduced in [Fre03]. Jentzsch's extension of the Perron–Frobenius theorem [Jen12] states that the eigenvalue of largest module is real and positive, and that the associated eigenfunctions can be taken real and positive as well.

We will denote by $(\lambda_i)_{i \in \mathbb{N}_0}$ the eigenvalues of K , ordered by decreasing modulus, and by π_i and ϕ_i the left and right eigenfunctions, that is

$$(\pi_i K)(x) = \lambda_i \pi_i(x) \quad \text{and} \quad (K \phi_i)(x) = \lambda_i \phi_i(x)$$

for all $i \in \mathbb{N}_0$. We normalise the eigenfunctions in such a way that

$$\pi_i(\phi_j) := \int_{\mathcal{X}} \pi_i(x) \phi_j(x) dx = \delta_{ij},$$

which implies that the kernels with density $\phi_i(x)\pi_i(y)$ are projectors on invariant subspaces of K . In case the set of eigenfunctions is complete and all nonzero eigenvalues have equal algebraic and geometric multiplicity, we have the spectral decomposition

$$k^n(x, y) = \sum_{i \geq 0} \lambda_i^n \phi_i(x) \pi_i(y) \quad \forall n \in \mathbb{N}.$$

If some geometric multiplicities are smaller than the corresponding algebraic multiplicities, this decomposition will contain nontrivial Jordan blocks.

Since K is stochastic ($K(x, \mathcal{X}) = 1$ for all $x \in \mathcal{X}$), we have in particular $\lambda_0 = 1$, while π_0 is the density of the invariant distribution of the process, and ϕ_0 is identically equal to 1. In what follows, we will usually identify signed measures and their density.

2.5 Killed process and QSDs

Given $A \in \mathcal{S}$, we denote by K_A the kernel of the process $(X_n^A)_{n \geq 0}$ killed upon leaving A . Its density has the expression

$$k_A(x, y) = k(x, y) \mathbb{1}_{\{x \in A, y \in A\}}.$$

If A^c has positive Lebesgue measure, this is a substochastic process, which can be turned into a stochastic process on $A \cup \{\partial\}$, where ∂ denotes a cemetery state. The killing time of the process is given for all $x \in A$ by $\tau_{\partial}(x) = \tau_{A^c}(x) = \tau_{A^c}^+(x)$.

Fredholm theory also applies to K_A , and we denote its eigen-elements by λ_i^A , π_i^A and ϕ_i^A . A major difference in the substochastic case is that the *principal eigenvalue* λ_0^A is strictly smaller than 1. The left eigenfunction π_0^A is a *quasi-ergodic distribution* (QED) of the process, meaning that it satisfies

$$\mathbb{P}^{\pi_0^A} \{X_n^A \in B \mid \tau_{A^c} > n\} = \pi_0^A(B) \quad \forall B \in \mathcal{S}, \forall n \in \mathbb{N}.$$

It can also be checked that the killing time, when starting in the QED, is geometrically distributed with success probability $(1 - \lambda_0^A)$, that is,

$$\mathbb{P}^{\pi_0^A} \{\tau_{A^c} = n\} = (\lambda_0^A)^{n-1} (1 - \lambda_0^A) \quad \forall n \in \mathbb{N} \quad \text{and} \quad \mathbb{E}^{\pi_0^A} [\tau_{A^c}] = \frac{1}{1 - \lambda_0^A}.$$

If the spectral-gap condition $|\lambda_1^A| < \lambda_0^A$ is satisfied, then one also has

$$\lim_{n \rightarrow \infty} \mathbb{P}^x \{X_n^A \in B \mid \tau_{A^c} > n\} = \pi_0^A(B)$$

for all $x \in A$ and all $B \in \mathcal{S}$, meaning that π_0^A is also a *quasistationary distribution* (QSD). We refer to [CMSM13, BG16, CV16, DGLLPN16, CV23] for proofs and further details on QSDs.

2.6 Uniform positivity

The last assumption we need is a form of ergodicity condition, which is a particular case of the uniform positivity condition used in [Bir57], and a variant of Doeblin's condition for Markov chains suitable for substochastic processes (see also [HM11] for related results).

Definition 2.7 (Uniform positivity condition). *We say that a (sub)stochastic Markov kernel K_A on A with density k_A satisfies a uniform positivity condition with parameters $n \in \mathbb{N}$ and $L > 1$ if*

$$\sup_{x \in A} k_A^n(x, y) \leq L \inf_{x \in A} k_A^n(x, y) \quad (2.10)$$

holds for all $y \in A$.

Remark 2.8. A more general uniform positivity condition one encounters in the literature is that $s(x)\nu(B) \leq K^n(x, B) \leq Ls(x)\nu(B)$ for a positive function s and a positive measure ν . The form we use here corresponds to a constant s , which is sufficient for our purposes since we are going to apply it to sets A on which $K^n(x, \cdot)$ is bounded below. \diamond

We will only need uniform positivity to hold for certain trace processes killed upon hitting some metastable sets. More precisely, given $1 \leq i \leq N$, let $\mathcal{M}K_{\sigma, B_i}$ be the kernel of the trace process on \mathcal{M} , killed when it hits $\mathcal{M} \setminus B_i$ (which is equivalent to the trace process leaving B_i).

Assumption POS (Uniform positivity). There exist a constant $L \in (1, 2)$, independent of σ , and an integer $n_0(\sigma)$, such that for each $1 \leq i \leq N$, the kernel $\mathcal{M}K_{\sigma, B_i}$ satisfies a uniform positivity condition on B_i with parameters $n_0(\sigma)$ and L . Furthermore, for any $\eta > 0$, there exists $\sigma_0(\eta) > 0$ such that

$$n_0(\sigma) \leq e^{\eta/\sigma^2}$$

holds for all $\sigma \in (0, \sigma_0]$. ♣

At first glance, it might seem difficult to prove that such a condition holds. In practice, however, we will often be in the following situation. We have a bad upper bound on the oscillation of $x \mapsto \mathcal{M}k_{\sigma, B_i}(x, y)$ valid on the whole domain (typically, this bound has order e^{C/σ^2} for some $C > 0$), but we also have a much smaller bound, uniform in σ , when x is only allowed to vary on a small ball, typically of radius σ^2 . The two bounds can then be combined into a much better one by using a coupling argument, see Proposition B.1 in Appendix B.1.

3 Main results

We assume throughout this section, without further mention, that the kernel $K = \mathcal{X}K_\sigma$ satisfies Assumptions DET, LDP, REC and POS. Our first main result concerns the spectrum of K . We give its proof in Section 5.

Proposition 3.1 (Spectral gap estimate). *For any $\eta > 0$, there exist $\sigma_0 > 0$ and $\delta_0 > 0$ such that, if $\sigma \in (0, \sigma_0]$ and the diameter of the sets B_i is bounded by δ_0 , then the kernel K has exactly N eigenvalues outside the disc $\{\lambda \in \mathbb{C} : |\lambda| \leq \varrho\}$, where*

$$\varrho = \exp\left\{-\frac{\log 2}{4E_{\mathcal{X}}(\sigma)}\right\}.$$

Furthermore, these N eigenvalues all belong to the disc of radius $e^{-[H_0 - \eta]/\sigma^2}$, centred in 1.

Remark 3.2 (Sharper estimates on the N first eigenvalues). In [BB17], we obtained sharper estimates on the N largest eigenvalues, in terms of committor functions between the B_i , under a more restrictive condition on the $H(i, j)$. The condition requires that the B_i can be ordered in such a way that

$$\min_{j < i} H(i, j) \leq \min_{k < i} \min_{j \leq i, j \neq k} H(k, j) - \theta \quad \forall i \in \{1, \dots, N\}$$

holds for some $\theta > 0$. Since we do not make this assumption here, it is necessary to give a new proof of Proposition 3.1 in the current situation. The proof uses however the same tools as in [BB17]. \diamond

One consequence of Proposition 3.1 is that the variables of the sequence $(X_n)_{n \geq 0}$ will be at distance decreasing like ϱ^n from a sequence $((X_{\text{tr}})_n)_{n \geq 0}$, generated by the truncated kernel K_{tr} , obtained by projecting K on the space associated with its N largest eigenvalues. This truncated kernel is given, in the basis of eigenfunctions of K , by a matrix of size N . However, it is not immediately clear how this approximate sequence relates to the sequence of visited B_i . The following approximation result clarifies that point.

Theorem 3.3 (Approximation by a finite Markov chain). *There exist constants $C, \theta_0 > 0$ such that the following holds for all $\theta \in (0, \theta_0]$. Let $i \in \{1, \dots, N\}$, and let $m = m(\sigma)$ satisfy*

$$\lim_{\sigma \rightarrow 0} \sigma^2 \log(m(\sigma)) = \theta.$$

Then for any $\eta > 0$, there exist $\sigma_0 > 0$ and $\delta_0 > 0$, such that if $\sigma \in (0, \sigma_0]$ and the diameter of the sets B_i is bounded by δ_0 , then for any $x \in B_i$, one has

$$|\mathbb{P}^x \{X_{\tau_{\mathcal{M}}^{+,nm}} \in B_j\} - \mathbb{P}^i \{Y_n = j\}| \leq C(e^{-[\hat{H}_{\min} - \eta]/\sigma^2} + \varrho^{nm}) \quad (3.1)$$

for all $n \in \mathbb{N}$ and all $j \in \{1, \dots, N\}$. Here \hat{H}_{\min} is a constant satisfying $\hat{H}_{\min} \geq H_0 - (N-1)\theta$, and $(Y_n)_{n \geq 0}$ is the Markov chain with transition matrix P , whose matrix elements satisfy

$$P_{ij} = \mathbb{P}^{\pi_0^{B_i}} \{X_{\tau_{\mathcal{M}}^{+,m}} \in B_j\} [1 + \mathcal{O}(e^{-[\theta - \eta]/\sigma^2})], \quad (3.2)$$

where $\pi_0^{B_i}$ is the QSD of the trace process $\mathcal{M}(X_n)$ killed when leaving B_i .

We give the proof in Section 6, which also contains more precise information on the constants θ_0 and \hat{H}_{\min} . Theorem 6.14 also provides the relation

$$\mathbb{P}^i \{Y_n = j\} = \mathbb{E}^{\mu_i} [\psi_j(X_{\tau_{\mathcal{M}}^{+,nm}})],$$

where the μ_i and ψ_j are suitable measures and test functions, and $\tau_{\mathcal{M}}^{+,n}$ is the n th return time to the metastable set \mathcal{M} . This shows in which way the original and reduced process are coupled. In fact, there exists a linear map \mathcal{L} from the space of measures on \mathcal{X} to those on $\{1, \dots, N\}$ such that

$$\mathbb{P}^i Y_n^{-1} = \mathcal{L}(\mathbb{P}^{\mu_i} X_{\tau_{\mathcal{M}}^{+,nm}}^{-1}) \quad \forall n \in \mathbb{N}_0,$$

given by $\mathcal{L}(\mu_j) = \delta_j$. The map \mathcal{L} is of course highly non-injective, since it maps an infinite-dimensional space to a space of dimension N . Its kernel is a complement of the space of measures spanned by μ_1, \dots, μ_N , given by the space of measures μ such that $\mathbb{E}^\mu[\psi_j] = 0$ for $j = 1, \dots, N$.

Relation (3.1) shows in which sense the sequence of visited balls B_i is close to the Markov chain $(Y_n)_{n \geq 0}$. Note that the error term ρ^{nm} converges to 0 as n increases. It actually becomes negligible as soon as

$$n \geq \hat{H}_{\min} \frac{E_{\mathcal{X}}(\sigma)}{\sigma^2 m(\sigma)},$$

which already happens for $n \geq 1$ if one applies Proposition 2.4 with η small enough.

The important part of the error term in (3.1) is thus given by $C e^{-[\hat{H}_{\min} - \eta]/\sigma^2}$. The point is that this error is *uniform* in time n . Thus at any given time n , we know that the trace process is likely to be in a ball B_i whenever the probability $\mathbb{P}^i\{Y_n = j\}$ is not exponentially small. This information becomes useful on time scales that are long compared to the typical time of transitions between metastable sets.

The process $(X_{t_{\mathcal{M}}^{+,nm}})_{n \geq 0}$ can thus be approximated, up to an exponentially small error that is uniform in time, by a Markov chain with transition probabilities P_{ij} . Note that the error in the expression (3.2) for these probabilities is multiplicative. Our analysis does not provide more explicit expressions for these transition probabilities than the large-deviation estimate in Proposition 6.1, but it shows that it is sufficient to know the probabilities of hitting the different balls B_j when starting in the QSD on each B_i . One may hope that future development of the theory will provide sharper estimates.

4 Applications

In this section, we show that most of the main assumptions are automatically satisfied for the two main applications we have in mind, namely randomly perturbed iterated maps, and random Poincaré maps.

4.1 Iterated maps with additive noise

Let $\mathcal{X}_0 = \mathbb{R}^d$ and consider the Markov chain given by

$$X_{n+1} = \Pi(X_n) + \sigma \xi_{n+1},$$

where $\Pi : \mathbb{R}^d \rightarrow \mathbb{R}^d$ satisfies Assumption DET, and the ξ_n are i.i.d. random variables taking values in \mathbb{R}^d . A typical example would be that the ξ_n are centred, normal random variables with positive definite covariance matrix Σ (that is, we assume $c_- \|\zeta\|^2 \leq \langle \zeta, \Sigma \zeta \rangle c_+ \leq \|\zeta\|^2$ for all $\zeta \in \mathbb{R}^d$, where $c_+ \geq c_- > 0$).

The transition kernel of the chain $(X_n)_{n \geq 0}$ is given by

$$K_\sigma(x, A) = \mathbb{P}\{\Pi(x) + \sigma \xi_1 \in A\} = \mathbb{P}\{\sigma \xi_1 \in A - \Pi(x)\} \quad \forall A \in \mathcal{S}_0.$$

We now examine the four assumptions one by one.

Assumption DET. The existence of a set \mathcal{X} invariant under the map Π is a classical growth condition that holds true for many discrete-time dynamical systems. Let us assume for simplicity that \mathcal{X} can be taken as a ball $\mathcal{B}(R_0) = \{x \in \mathbb{R}^d : \|x\| < R_0\}$. For later use, we shall make the somewhat stronger assumption that Π maps any ball $\mathcal{B}(R)$ of radius $R \geq R_0$ into a smaller ball, namely there exists $\varepsilon_0 > 0$ such that

$$\|\Pi(x)\|^2 \leq \|x\|^2 - \varepsilon_0 \quad \forall x : \|x\| \geq R_0. \quad (4.1)$$

Checking the conditions on limit sets is in general no easy task, as it requires a good understanding of fixed points and periodic orbits, their basins of attraction, and there stable and unstable manifolds. However they are known to hold for a number of systems. See for instance [CFLM06] for a non-trivial dynamical systems arising from genetic regulatory networks.

Assumption LDP. Assume the random variable ξ_1 satisfies a large-deviation principle with good rate function I_0 . Then it is immediate to see that $K_\sigma(x, \cdot)$ satisfies (2.1) and (2.2) with the rate function

$$I(x, y) = I_0(y - \Pi(x)).$$

We see that I vanishes only if $y = \Pi(x)$ provided $I(x) > 0$ for $x \neq 0$, and is continuous at fixed points whenever I_0 is continuous at 0.

In particular, if ξ_1 has a centred normal distribution with covariance matrix Σ , then

$$I(x, y) = \frac{1}{2} \langle y - \Pi(x), \Sigma^{-1}(y - \Pi(x)) \rangle \quad (4.2)$$

satisfies all required properties.

Assumption REC. Assume ξ_1 has a continuous density p . Then we see that K_σ admits the density

$$k_\sigma(x, y) = \frac{1}{\sigma^d} p\left(\frac{y - \Pi(x)}{\sigma}\right),$$

as required. Furthermore, taking $U(x) = \|x\|^2$ as Lyapunov function, we obtain

$$\begin{aligned} (K_\sigma U)(x) &= \mathbb{E}^x [\|\Pi(x) + \sigma \xi_1\|^2] \\ &= \|\Pi(x)\|^2 + 2\sigma \langle \Pi(x), \mathbb{E}[\xi] \rangle + \sigma^2 \mathbb{E}[\|\xi_1\|^2]. \end{aligned}$$

Thus if we assume that ξ_1 has zero mean and its components have bounded variance, it follows from (4.1) that the discrete drift condition (2.8) is satisfied provided $\sigma^2 < \varepsilon_0 [\mathbb{E}[\|\xi_1\|^2]]^{-1}$.

These properties clearly hold in the case of Gaussian ξ_i , for which the density is

$$k_\sigma(x, y) = \frac{1}{\mathcal{N}} e^{-I(x, y)/\sigma^2}, \quad \mathcal{N} = (2\pi\sigma^2)^{d/2} (\det \Sigma)^{1/2} \quad (4.3)$$

with I given by (4.2).

Assumption POS. In the case where ξ_1 follows a normal law, the following result based on the coupling argument in Proposition B.1 shows that the positivity condition holds for sufficiently small diameter of the B_i .

Proposition 4.1 (Positivity for Gaussian noise). *Assume ξ_1 follows a centred, normal law with positive definite covariance matrix Σ . Then there exist $\delta_0, \sigma_0 > 0$ such that, if the B_i have a diameter bounded by δ_0 and $0 < \sigma < \sigma_0$, then Assumption POS is satisfied for $n_0(\sigma)$ of order $\log(\sigma^{-1})$.*

PROOF: See Appendix B.1. □

Recall that Proposition 2.4 shows that $E_{\mathcal{X}}(\sigma)$ is bounded by any exponential e^{η/σ^2} if σ and the B_i are small enough. In fact, we can do much better, and show that this expectation has order $\log(\sigma^{-1})$ if we assume that the deterministic system does not admit any heteroclinic cycles. A *heteroclinic orbit* from an unstable fixed point z_1^* to an unstable fixed point z_2^* is an

orbit whose α -limit set is equal to z_1^* and whose ω -limit set is equal to z_2^* . A *heteroclinic cycle* between unstable fixed points z_1^*, \dots, z_n^* is a set of heteroclinic orbits connecting z_1^* to z_2^* , z_2^* to z_3^* , \dots , z_{n-1}^* to z_n^* and z_n^* to z_1^* .

Proposition 4.2 (Expected hitting time of \mathcal{M}). *Assume ξ_1 follows a centred, normal law with positive definite covariance matrix Σ , and the deterministic dynamical system generated by Π_0 has no heteroclinic cycles. Then there exist constants $c_0, \sigma_0, \delta_0 > 0$ such that*

$$E_{\mathcal{X}}(\sigma) \leq c_0 \log(\sigma^{-1})$$

holds for $0 < \sigma < \sigma_0$ and $0 < \delta < \delta_0$.

PROOF: See Appendix B.2. □

The reason we exclude heteroclinic cycles is that the system may spend times longer than $\log(\sigma^{-1})$ in their neighbourhood. Note that SDEs with heteroclinic cycles have been investigated, for instance, in [Bak11]. Results from that work may be transposed to the present situation, to analyse that point in more detail.

Based on what is known in the continuous-time case [BB17], similar results are expected to hold for more general systems with state-dependent noise, of the form

$$X_{n+1} = \Pi(X_n) + \sigma g(X_n) \xi_{n+1},$$

provided g satisfies an ellipticity condition (that is, $g(x)g(x)^\dagger$ should be positive definite). If g fails to be elliptic at certain points, a more careful analysis becomes necessary.

4.2 Random Poincaré maps

Consider a stochastic differential equation on $\mathcal{D}_0 \subset \mathbb{R}^{d+1}$ of the form

$$dz_t = f(z_t) dt + \sigma g(z_t) dW_t, \tag{4.4}$$

where $f : \mathcal{D}_0 \rightarrow \mathbb{R}^{d+1}$ is a vector field of class \mathcal{C}^2 , $g : \mathcal{D}_0 \rightarrow \mathbb{R}^{(d+1) \times k}$ is of class \mathcal{C}^1 and $(W_t)_{t \geq 0}$ is a k -dimensional standard Wiener process. Assume further that the deterministic ordinary differential equation

$$\dot{z} = f(z) \tag{4.5}$$

admits $N \geq 2$ linearly asymptotically stable periodic orbits $\Gamma_1, \dots, \Gamma_N$, and that there exists a smooth d -dimensional manifold Σ that all Γ_i intersect transversally (cf. [BB17, Sect. 2.2]).

The *random Poincaré map* associated with this system describes the sequence (X_0, X_1, \dots) of successive intersections of a sample path $(z_t)_{t \geq 0}$ of the SDE (4.4) with Σ . To obtain a well-defined process, these intersections should be separated by excursions away from Σ , which can be achieved by requiring the sample path to visit another section Σ' , disjoint from Σ , between two consecutive X_i (see [BB17, Sect. 2.3]). The strong Markov property implies that the sequence $(X_n)_{n \geq 0}$ forms a Markov chain which, under suitable assumptions on f and g , is of the form studied here.

Assumption DET. This assumption is fulfilled if the deterministic system (4.5) admits a positively invariant, bounded open connected set $\mathcal{D} \subset \mathcal{D}_0$, intersecting Σ , and the limit sets of (4.5) are given by the Γ_i and finitely many linearly unstable stationary points or unstable orbits. We can then take $\mathcal{X} = \mathcal{D} \cap \Sigma$, and Π maps a point in $x \in \Sigma$ to the point where the positive orbit of x first returns to Σ . Furthermore, $x_i^* = \Gamma_i \cap \Sigma$, and the intersections of the unstable periodic orbits with Σ are the unstable fixed points of Π .

Assumption LDP. Assume the diffusion coefficient g satisfies an ellipticity condition, that is, there exist constants $c_+ \geq c_- > 0$ such that the diffusion matrix $D(z) = g(z)g(z)^\dagger$ satisfies

$$c_- \|\xi\|^2 \leq \langle \xi, D(z)\xi \rangle \leq c_+ \|\xi\|^2 \quad (4.6)$$

for all $z \in \mathcal{D}$ and $\xi \in \mathbb{R}^{d+1}$. Then Wentzell–Freidlin theory [FW98] provides the existence of a sample-path LDP with rate function

$$\mathcal{I}_{[0,T]}(\gamma) = \begin{cases} \frac{1}{2} \int_0^T (\dot{\gamma}_s - f(\gamma_s))^\dagger D(\gamma_s)^{-1} (\dot{\gamma}_s - f(\gamma_s)) \, ds & \text{if } \gamma \in H^1, \\ +\infty & \text{otherwise.} \end{cases}$$

Note that if g fails to satisfy the ellipticity condition (4.6), an LDP may still hold, but its rate function is given by a variational principle (obtained by applying the contraction principle to Schilder's theorem for scaled Brownian motion).

This continuous-time LDP induces, by the contraction principle, a discrete-time LDP with rate function

$$I(x, y) = \inf_{T > 0} \inf_{\gamma: x \rightarrow y} \mathcal{I}_{[0,T]}(\gamma),$$

where the second infimum runs over paths connecting points x and y in Σ in time T , and making an excursion via Σ' .

Assumption REC. The ellipticity condition (4.6) ensures that the kernel K_σ of the Markov chain $(X_n)_{n \geq 0}$ admits a continuous density k_σ as shown in [BAKS84]. In fact, a weaker hypo-ellipticity condition is sufficient. As for Harris recurrence, it follows from a continuous-time analogue of the discrete drift condition (2.8) [MT93]. Namely, there should exist a function $V: \mathcal{D}_0 \rightarrow \mathbb{R}_+$ of class \mathcal{C}^2 , diverging as $\|x\| \rightarrow \infty$, and constants $c > 0$ and $d \geq 0$ such that

$$(\mathcal{L}V)(z) \leq -c + d \mathbb{1}_{\{z \in \mathcal{D}\}} \quad \forall z \in \mathcal{D}_0,$$

where \mathcal{L} is the infinitesimal generator of the diffusion (4.4).

Assumption POS. The uniform positivity condition (2.10) of the trace process can again be proved to hold by applying the coupling argument of Proposition B.1. Instead of using Harnack inequalities for the density of a Gaussian random variable, one can use Harnack inequalities satisfied by harmonic functions, see [BB17, Sect. 5.1] and [BG14, Sect. 5.3]. The parameter $n_0(\sigma)$ in the uniform positivity condition has again order $\log(\sigma^{-1})$.

We also have an analogue of Proposition 4.2 on the expected hitting time of the metastable set \mathcal{M} .

Proposition 4.3 (Expected hitting time of \mathcal{M}). *Assume the deterministic dynamical system (4.5) has no heteroclinic cycles. Then there exist constants $c_0, \sigma_0, \delta_0 > 0$ such that*

$$E_{\mathcal{X}}(\sigma) \leq c_0 \log(\sigma^{-1})$$

holds for $0 < \sigma < \sigma_0$ and $0 < \delta < \delta_0$.

PROOF: See [BB17, Cor. 8.13]. This work excluded the existence of heteroclinic orbits between unstable periodic orbits, but the same arguments as in the proof of Proposition 4.2 show that the absence of heteroclinic cycles is sufficient. \square

5 Proof of Proposition 3.1

In this section, we give the proof of the spectral gap result stated in Proposition 3.1, by adapting the proof of [BB17, Thm 3.2] to the weaker assumptions of the present work. We start with a simple but useful a priori estimate.

Lemma 5.1. *There exist constants $\theta_0, \sigma_0 > 0$ such that*

$$\sup_{x \in \mathcal{M}} \mathbb{P}^x \{X_1 \in \mathcal{M}^c\} \leq e^{-\theta_0/\sigma^2} \quad (5.1)$$

holds for all $\sigma \leq \sigma_0$.

PROOF: Pick an $x \in B_i \subset \mathcal{M}$. Since B_i is assumed to be positively invariant under the map Π and x_i^* is asymptotically stable, $\Pi(x)$ belongs to B_i , and its distance to ∂B_i is bounded below. The claim thus follows from the large-deviation principle. \square

Fix $x \in \mathcal{X}$ and $m_0 \in \mathbb{N}$. By Markov's inequality and the definition (2.7) of $E_{\mathcal{X}}(\sigma)$, we have

$$\mathbb{P}^x \{\tau_{\mathcal{M}}^+ > m_0\} \leq \frac{1}{m_0} \mathbb{E}^x[\tau_{\mathcal{M}}^+] \leq \frac{1}{m_0} E_{\mathcal{X}}(\sigma).$$

This shows that

$$\begin{aligned} \mathbb{P}^x \{X_{m_0} \notin \mathcal{M}\} &\leq \mathbb{P}^x \{\tau_{\mathcal{M}}^+ > m_0\} + \mathbb{P}^x \{X_{m_0} \notin \mathcal{M}, \tau_{\mathcal{M}}^+ \leq m_0\} \\ &\leq \frac{1}{m_0} E_{\mathcal{X}}(\sigma) + \sup_{m_1 \leq m_0} \sup_{y \in \mathcal{M}} \mathbb{P}^y \{X_{m_1} \in \mathcal{M}^c\}. \end{aligned}$$

The second term is exponentially small by (5.1), and is thus bounded by $\frac{1}{4}$ for σ small enough. For sufficiently small σ , we thus have

$$\mathbb{P}^x \{X_{m_0} \notin \mathcal{M}\} \leq \frac{1}{2}$$

provided $m_0 \geq 4E_{\mathcal{X}}(\sigma)$.

5.1 Feynman–Kac representation formulas

We will now rely on Feynman–Kac representation formulas for eigenfunctions of the kernel K , as used in [BB17, Sect. 4], to show that there are only N eigenvalues outside a given disc in the complex plane. Writing $\tilde{X}_n = X_{nm_0}$ for the time-diluted Markov chain and $\tilde{\tau}_{\mathcal{M}}^+$ for the corresponding first-hitting time of \mathcal{M} , we obtain from [BB17, Lem. 4.1] that the Laplace transform $\mathbb{E}^x[e^{u\tilde{\tau}_{\mathcal{M}}^+}]$ exists whenever $|e^{-u}| \geq \frac{1}{2}$. By [BB17, Cor. 4.3], we know that (e^{-u}, ϕ) is an eigenpair of K^{m_0} for $|e^{-u}| > \frac{1}{2}$ if, and only if, one has

$$((K^u)^{m_0} \phi)(x) = e^{-u} \phi(x) \quad \forall x \in \mathcal{M},$$

where K^u is the kernel defined by

$$K^u(x, A) = \mathbb{E}^x \left[e^{u(\tau_{\mathcal{M}}^+ - 1)} \mathbb{1}_{\{X_{\tau_{\mathcal{M}}^+} \in A\}} \right] \quad \forall A \in \mathcal{B}(\mathcal{M}).$$

Note that K^0 is equal to the kernel $_{\mathcal{M}}K$ of the trace process on \mathcal{M} . It follows that (e^{-u}, ϕ) is an eigenpair of K for $|e^{-u}| > (\frac{1}{2})^{1/m_0}$ if, and only if, one has

$$(K^u \phi)(x) = e^{-u} \phi(x) \quad \forall x \in \mathcal{M}. \quad (5.2)$$

Remark 5.2. It may seem unusual that the variable u appears both in the kernel K^u and the eigenvalue e^{-u} . This is not a problem, however, since the eigenvalue problem (5.2) can be considered as the system

$$K^u \phi = \lambda \phi, \quad \lambda = e^{-u}$$

for two unknowns λ and e^{-u} . ◇

The idea is now to compare the kernel K^u to the simpler kernel K^\star , defined for $A \subset \mathcal{M}$ by

$$K^\star(x, A) = \sum_{i=1}^N \mathbb{1}_{\{x \in B_i\}} \mathbb{P}^{\hat{\pi}_0^{B_i}} \{X_{\tau_{\mathcal{M}}^+} \in A\}.$$

Here we recall that $\hat{\pi}_0^{B_i}$ denotes the quasistationary distribution of the trace process on \mathcal{M} killed when leaving B_i . Since

$$(K^\star \phi)(x) = \sum_{i=1}^N \mathbb{1}_{\{x \in B_i\}} \mathbb{E}^{\hat{\pi}_0^{B_i}} [\phi(X_{\tau_{\mathcal{M}}^+})],$$

the kernel K^\star has finite rank. Indeed, its image is the N -dimensional space of functions $\phi : \mathcal{M} \rightarrow \mathbb{R}$ that are constant on each B_i . Therefore, K^\star has at most N nonzero eigenvalues. These eigenvalues are exactly those of the N by N stochastic matrix P^\star with elements

$$P_{ij}^\star = \mathbb{P}^{\hat{\pi}_0^{B_i}} \{X_{\tau_{\mathcal{M}}^+} \in B_j\}. \quad (5.3)$$

Note that Proposition 2.2 implies that for any $\eta > 0$, there exists a $\sigma_0(\eta) > 0$ such that one has

$$e^{-(H_0+\eta)/\sigma^2} \leq \mathbb{P}^x \{\tau_{\mathcal{M} \setminus B_i} \leq n\} \leq n e^{-(H_0-\eta)/\sigma^2} \quad (5.4)$$

for any $n \in \mathbb{N}$, any $x \in B_i$ and all $\sigma < \sigma_0(\eta)$, where H_0 has been introduced in (2.5). This shows in particular that the matrix elements (5.3) satisfy

$$P_{ij} \leq e^{-(H_0-\eta)/\sigma^2} \quad \text{for } i \neq j. \quad (5.5)$$

5.2 Norm estimates on kernels

In order to compare kernels, we will need a norm on the space of (signed) kernels on \mathcal{M} . If Q is such a kernel with density q , we write

$$\|Q\| = \sup_{x \in \mathcal{M}} \int_{\mathcal{M}} |q(x, y)| dy = \sup_{x \in \mathcal{M}} |Q(x, \mathcal{M})|.$$

One easily checks that this is a subordinate norm, given by

$$\|Q\| = \sup_{\varphi \in L^\infty: \|\varphi\|_\infty=1} \|Q\varphi\|_\infty = \sup_{\mu \in L^1: \|\mu\|_1=1} \|\mu Q\|_1.$$

In particular, (5.5) implies that the kernel $R = K^\star - \text{id}$ satisfies

$$\|R\| \leq 2(N-1) e^{-(H_0-\eta)/\sigma^2}. \quad (5.6)$$

This allows us to bound the resolvent of K^\star . Indeed, for $z \in \mathbb{C} \setminus \{1\}$, we have

$$(z \text{id} - K^\star)^{-1} = \frac{1}{z-1} \left(\text{id} - \frac{1}{z-1} R \right)^{-1} = \frac{1}{z-1} \sum_{n \geq 0} \frac{1}{(z-1)^n} R^n.$$

The Neumann series converges whenever $|z - 1| > \|R\|$, in which case we have

$$\|(z \text{id} - K^\star)^{-1}\| \leq \frac{1}{|z - 1| - \|R\|}. \quad (5.7)$$

It follows that all eigenvalues of K^\star are contained in the closed disc of radius $\|R\|$ centred in 1.

Our aim is now to compare K^u and K^\star in two steps. Firstly, [BB17, Prop. 6.1], slightly adapted to allow for complex u , shows that for any $m \in \mathbb{N}$,

$$\|(K^u)^m - (K^0)^m\| \leq \left(1 + \frac{|1 - e^{-u}| \mathbb{E}^{\mathcal{M}}[\tau_{\mathcal{M}}^+ - 1]}{1 - |1 - e^{-u}| \mathbb{E}^{\mathcal{M}^c}[\tau_{\mathcal{M}}^+]}\right)^m - 1, \quad (5.8)$$

which holds provided the denominator is strictly positive. This is indeed the case provided $|1 - e^{-u}| < E_{\mathcal{X}}(\sigma)^{-1}$. Secondly, [BB17, Prop. 6.7] shows that for any $m \in \mathbb{N}$,

$$\|(K^0)^m - (K^\star)^m\| \leq \sup_{1 \leq i \leq N} R_i,$$

where

$$\begin{aligned} R_i &= \|\phi_0^{B_i} - 1\| + 2|\lambda_1^{B_i}|^m + 2 \frac{1 - |\lambda_1^{B_i}|^m}{1 - |\lambda_1^{B_i}|} \mathbb{P}^{B_i}\{\tau_{\mathcal{M} \setminus B_i}^+ < \tau_{B_i}^+\} \\ &\quad + m(m-1) \mathbb{P}^{B_i}\{\tau_{\mathcal{M} \setminus B_i}^+ < \tau_{B_i}^+\} \mathbb{P}^{\mathcal{M} \setminus B_i}\{\tau_{B_i}^+ < \tau_{\mathcal{M} \setminus B_i}^+\}. \end{aligned} \quad (5.9)$$

Here the $\phi_k^{B_i}$ and $\lambda_k^{B_i}$ denote eigen-elements of the trace process on \mathcal{M} killed upon leaving B_i . These can be estimated thanks to the uniform positivity condition POS. First note that integrating (2.10) against $\phi_0^A(y)$, yields the very rough bound

$$\sup_{x \in A} \phi_0^A \leq L \inf_{x \in A} \phi_0^A(x).$$

With the normalisation $\pi_0^A(\phi_0^A) = 1$, this yields $L^{-1} \leq \phi_0^A(x) \leq L$ for all $x \in A$. A much sharper bound is then provided by the following estimates.

Proposition 5.3 (Spectral gap and oscillation of ϕ_0^A). *Let K_A be the kernel of the process killed upon leaving A . Assume its density k_A satisfies the uniform positivity condition (2.10) with parameters $n_0(\sigma) \in \mathbb{N}$ and $L \in (1, 2)$. Then the spectral gap satisfies*

$$\left(\frac{|\lambda_1^A|}{\lambda_0^A}\right)^{n_0(\sigma)} \leq L - \frac{\inf_{x \in A} \mathbb{P}^x\{\tau_{A^c} > n_0(\sigma)\}}{(\lambda_0^A)^{n_0(\sigma)}}. \quad (5.10)$$

Furthermore, the oscillation of the principal eigenfunction satisfies

$$\|\phi_0^A - 1\| := \sup_{x \in A} |\phi_0^A(x) - 1| \leq L^3 \left| 1 - \frac{\inf_{x \in A} \mathbb{P}^x\{\tau_{A^c} > n_0(\sigma)\}}{(\lambda_0^A)^{n_0(\sigma)}} \right|. \quad (5.11)$$

PROOF: The spectral gap estimate (5.10) is proved in [BB17, Prop. 5.1]. The estimate (5.11) is proved in [BB17, Prop. 5.5], where the constant M in that result can be taken equal to L thanks to the *a priori* bound $L^{-1} \leq \phi_0^A(x) \leq L$. \square

We are going to apply these bounds to the kernel $\mathcal{M}K$ with $A = B_i$. Note that (5.4) already allows us to bound several terms in (5.9) by exponentially small quantities. Furthermore, the definition of $\hat{\pi}_0^{B_i}$ implies that for any $x \in B_i$, one has

$$\begin{aligned} 1 - \frac{\mathbb{P}^x\{\tau_{\mathcal{M} \setminus B_i} > n_0(\sigma)\}}{(\hat{\lambda}_0^{B_i})^{n_0(\sigma)}} &= 1 - \frac{\mathbb{P}^x\{\tau_{\mathcal{M} \setminus B_i} > n_0(\sigma)\}}{\mathbb{P}^{\hat{\pi}_0^{B_i}}\{\tau_{\mathcal{M} \setminus B_i} > n_0(\sigma)\}} \\ &= \frac{\mathbb{P}^x\{\tau_{\mathcal{M} \setminus B_i} \leq n_0(\sigma)\} - \mathbb{P}^{\hat{\pi}_0^{B_i}}\{\tau_{\mathcal{M} \setminus B_i} \leq n_0(\sigma)\}}{1 - \mathbb{P}^{\hat{\pi}_0^{B_i}}\{\tau_{\mathcal{M} \setminus B_i} \leq n_0(\sigma)\}}. \end{aligned}$$

Combining (5.11) and (5.4) we get

$$\|\hat{\phi}_0^{B_i} - 1\| \leq \frac{L^3 n_0(\sigma) e^{-(H_0 - \eta)/\sigma^2}}{1 - e^{-(H_0 + \eta)/\sigma^2}}.$$

A similar argument, based on the bound (5.10), shows that

$$|\hat{\lambda}_1^{B_i}| \leq \left(\frac{L - 1 + n_0(\sigma) e^{-(H_0 - \eta)/\sigma^2}}{1 - e^{-(H_0 + \eta)/\sigma^2}} \right)^{1/n_0(\sigma)}, \quad (5.12)$$

Plugging the last two estimates into (5.9) yields

$$\begin{aligned} \|(K^0)^m - (K^*)^m\| &\leq \frac{L^3 n_0(\sigma) e^{-(H_0 - \eta)/\sigma^2}}{1 - e^{-(H_0 + \eta)/\sigma^2}} + \left(\frac{L - 1 + n_0(\sigma) e^{-(H_0 - \eta)/\sigma^2}}{1 - e^{-(H_0 + \eta)/\sigma^2}} \right)^{m/n_0(\sigma)} \\ &\quad + 2m e^{-(H_0 - \eta)/\sigma^2} + m^2 e^{-2(H_0 - \eta)/\sigma^2}, \end{aligned} \quad (5.13)$$

where we have bounded the fraction in (5.9) above by m .

5.3 Resolvent estimate

We can now apply the following classical resolvent estimate, see for instance the argument presented in [BB17, Sect. 7.1], which is based on [GG01, Cor. 8.2] and [GGK03, Prop. 4.2].

Lemma 5.4. *Let K_1 and K_2 be compact linear operators. Let Γ be a contour in the complex plane, encircling k eigenvalues of K_1 . Let*

$$\begin{aligned} \gamma &= \min\{\|(z \text{id} - K_1)^{-1}\|^{-1} : z \in \Gamma\}, \\ C &= \frac{1}{\pi} \int_{\Gamma} \|(z \text{id} - K_1)^{-1}\|^2 dz. \end{aligned}$$

If $\|K_2 - K_1\| < \min\{\frac{1}{2}\gamma, C^{-1}\}$, then K_2 has exactly k eigenvalues inside the contour Γ .

We now apply this lemma to $K_1 = (K^*)^m$, and $K_2 = (K^u)^m$. The same argument as the one yielding (5.7) shows that

$$\|(z \text{id} - (K^*)^m)^{-1}\| \leq \frac{1}{|z - 1| + 1 - (1 + \|R\|)^m}.$$

It follows that all N nonzero eigenvalues of $(K^*)^m$ are contained in a disc of radius $(1 + \|R\|)^m - 1$, centred in 1. Given $r \geq 2[(1 + \|R\|)^m - 1]$, Lemma 5.4 applied to the contour Γ of radius r , centred in 1 shows that $(K^u)^m$ has exactly N eigenvalues inside Γ , provided

$$\|(K^u)^m - (K^*)^m\| \leq r. \quad (5.14)$$

We now make some convenient choices for various parameters. First of all, we assume that

$$\eta \leq \frac{1}{3} \min\{H_0, \theta_0\},$$

and take σ small enough to guarantee that $n_0(\sigma) \leq e^{\eta/\sigma^2}$ and $E_{\mathcal{X}}(\sigma) \leq e^{\eta/\sigma^2}$, which is possible by Assumption POS and Proposition 2.4. Since $\delta = \frac{1}{2}(L-1) > 0$, we may define

$$m = \left\lceil \frac{1}{\sigma^2} \max\left\{m_0, \frac{H_0 n_0(\sigma)}{\log(\delta^{-1})}\right\} \right\rceil. \quad (5.15)$$

Note that this implies $\delta^{m/n_0(\sigma)} \leq e^{-H_0/\sigma^2}$. Then it follows from (5.13) that

$$\|(K^0)^m - (K^\star)^m\| \leq \left(1 + C_1 \frac{n_0(\sigma)}{\sigma^2}\right) e^{-(H_0 - \eta)/\sigma^2}$$

for some numerical constant C_1 . Next we note that for any $x \in \mathcal{M}$, one has

$$\mathbb{E}^x[\tau_{\mathcal{M}}^+ - 1] \leq \mathbb{P}^x\{X_1 \in \mathcal{M}^c\} E_{\mathcal{X}}(\sigma) \leq e^{-\theta_0/\sigma^2} E_{\mathcal{X}}(\sigma).$$

If we further impose the condition

$$r \leq \frac{1}{2E_{\mathcal{X}}(\sigma)}, \quad (5.16)$$

then (5.8) yields

$$\|(K^u)^m - (K^0)^m\| \leq (1 + 2r \mathbb{E}^{\mathcal{M}}[\tau_{\mathcal{M}}^+ - 1])^m - 1 \leq (1 + 2r e^{-\theta_0/\sigma^2} E_{\mathcal{X}}(\sigma))^m - 1$$

If $2mr e^{-\theta_0/\sigma^2} E_{\mathcal{X}}(\sigma)$ is bounded, this quantity has order $mr e^{-\theta_0/\sigma^2} E_{\mathcal{X}}(\sigma)$, so that

$$\|(K^u)^m - (K^0)^m\| \leq C_2 \frac{n_0(\sigma) E_{\mathcal{X}}(\sigma)}{\sigma^2} e^{-\theta_0/\sigma^2} r \leq \frac{r}{2}$$

thanks to our bounds on $n_0(\sigma)$ and $E_{\mathcal{X}}(\sigma)$. We may thus set $r = e^{-(H_0 - 2\eta)/\sigma^2}$, which satisfies both (5.14) and (5.16). One furthermore checks that the bound (5.6) implies that $r \geq 2[(1 + \|R\|)^m - 1]$. We can thus conclude that K^u and K^\star have the same number of eigenvalues in the disc $\{|z - 1| < r\}$.

By a similar argument, K^u has no eigenvalues in any contour that does not contain 0, and stays at distance at least r from 1. It follows that K^u has exactly N nonzero eigenvalues (counting multiplicity).

Recall finally that K and K^u have the same eigenvalues outside a disc of radius $(\frac{1}{2})^{1/m_0}$. It follows that K^m and $(K^u)^m$ have the same number of eigenvalues outside a disc of radius $(\frac{1}{2})^{m/m_0}$. The choice (5.15) of m implies that this disc does not intersect the disc of radius r centred in 1, which concludes the proof. \square

6 Proof of Theorem 3.3

6.1 Large-deviation estimates for the trace process on \mathcal{M}

The trace process on \mathcal{M} is given by the sequence $(X_{\tau_{\mathcal{M}}^{+,n}})_{n \in \mathbb{N}}$, where

$$\tau_{\mathcal{M}}^{+,1} = \tau_{\mathcal{M}}^+, \quad \tau_{\mathcal{M}}^{+,n+1} = \inf\{m > \tau_{\mathcal{M}}^{+,n} : X_m \in \mathcal{M}\}.$$

Owing to Proposition 2.2, for any $\eta > 0$ there exist $\sigma_0(\eta) > 0$ and $\delta_0(\eta) > 0$ such that

$$e^{-[H(i,j)+\eta]/\sigma^2} \leq \mathbb{P}^x \{X_{\tau_{\mathcal{M}}^+} \in B_j\} \leq e^{-[H(i,j)-\eta]/\sigma^2} \quad (6.1)$$

holds for all $x \in B_i$, $j \neq i$ and all $\sigma < \sigma_0(\eta)$, provided the diameter of the balls B_ℓ is bounded by $\delta_0(\eta)$.

We will use several properties of the quasipotential H defined in (2.4). First note that H satisfies the triangle inequality

$$H(i, \ell) + H(\ell, j) \geq H(i, j) \quad \forall i, j, \ell \in \{1, \dots, N\},$$

where we have extended H by setting $H(i, i) = 0$ for all $i \in \{1, \dots, N\}$.

We call *path* a tuple $\gamma = (\gamma_0, \gamma_1, \dots, \gamma_{p-1}, \gamma_p) \in \{1, \dots, N\}^{p+1}$ whose consecutive elements are different. Its *length* is defined to be $|\gamma| := p$, and its *cost* is

$$V(\gamma) = H(\gamma_0, \gamma_1) + H(\gamma_1, \gamma_2) + \dots + H(\gamma_{p-1}, \gamma_p).$$

We write $\gamma : i \rightarrow j$ if $\gamma_0 = i$ and $\gamma_p = j$. We say that γ is an *optimal path* from i to j , and write $\gamma : i \rightarrow j$, if

$$V(\gamma) = H(i, j).$$

and for any $\varepsilon > 0$, there exists a sequence of points in the definition (2.3) of the quasipotential that visits all B_{γ_k} with γ_k an element of γ , and whose cost is smaller than $H(i, j) + \varepsilon$. Note that there may be more than one optimal path $\gamma : i \rightarrow j$.

We will also use the notation

$$\hat{H}_0 = \min_{\gamma : i \rightarrow j, V(\gamma) > H(i, j)} [V(\gamma) - H(i, j)] \quad (6.2)$$

for the minimal difference between the costs of a non-optimal path and an optimal path from i to j . The minimum is reached, even though the set of paths $\gamma : i \rightarrow j$ is infinite, because

$$V(\gamma) \geq |\gamma| H_0, \quad (6.3)$$

and an optimal path $\gamma : i \rightarrow j$ can have length $N - 1$ at most.

Proposition 6.1. *For any $\eta > 0$, there exist $\sigma_0(\eta), \delta_0(\eta) > 0$ and a constant C_N depending only on N such that*

$$\begin{aligned} \mathbb{P}^x \{X_{\tau_{\mathcal{M}}^+} \in B_j\} &\leq \sum_{\gamma : i \rightarrow j} \binom{n}{|\gamma|} e^{-[H(i,j)-|\gamma|\eta]/\sigma^2} + C_N e^{-[H(i,j)+\hat{H}_0-N\eta]/\sigma^2} \\ \mathbb{P}^x \{X_{\tau_{\mathcal{M}}^+} \in B_j\} &\geq \sum_{\gamma : i \rightarrow j} \binom{n}{|\gamma|} e^{-[H(i,j)+|\gamma|\eta]/\sigma^2} [1 - e^{-(H_0-\eta)/\sigma^2}]^{n-|\gamma|} \end{aligned}$$

holds for any $i \neq j$, $x \in B_i$ and $n \geq 1$, provided $\sigma < \sigma_0(\eta)$ and the diameter of the B_k is bounded by $\delta_0(\eta)$.

PROOF: To any trajectory $(X_{\tau_{\mathcal{M}}^+}^{+,k})_{0 \leq k \leq n}$ from x to B_j , we associate a path $\gamma = (i, \ell_1, \dots, \ell_p = j) : i \rightarrow j$ and an increasing sequence $0 = k_0 < k_1 < k_2 < k_p \leq n$ of jump times, such that

$$X_{\tau_{\mathcal{M}}^+}^{+,k} \in B_{\gamma_\ell} \quad \text{for } k_\ell \leq k < k_{\ell+1}.$$

The path γ simply indicates the sequence of visited balls. Then we have

$$\mathbb{P}^x \{X_{\tau_{\mathcal{M}}^{+,n}} \in B_j\} = \sum_{\substack{\gamma: i \rightarrow j \\ |\gamma| \leq n}} \sum_{0 < k_1 < \dots < k_p} Q_{k_1, \dots, k_p}(x), \quad (6.4)$$

where we have set $p = |\gamma|$, and

$$Q_{k_1, \dots, k_p}(x) = \mathbb{P}^x \{X_{\tau_{\mathcal{M}}^{+,k}} \in B_{\gamma_\ell}, k_\ell \leq k < k_{\ell+1}, 0 \leq \ell \leq p-1\}.$$

We want to use the fact that the sum (6.4) is dominated by optimal paths $\gamma : i \rightarrow j$. Let γ be such an optimal path, of length p . Then (6.1) yields the upper bound

$$Q_{k_1, \dots, k_p}(x) \leq \prod_{\ell=0}^{p-1} \sup_{y \in B_{\gamma_\ell}} \mathbb{P}^y \{X_{\tau_{\mathcal{M}}^{+,k}} \in B_{\gamma_{\ell+1}}\} \leq \prod_{\ell=0}^{p-1} e^{-[H(\gamma_\ell, \gamma_{\ell+1}) - \eta]/\sigma^2} \leq e^{-[H(i, j) - p\eta]/\sigma^2}. \quad (6.5)$$

As a lower bound, we have

$$\begin{aligned} Q_{k_1, \dots, k_p}(x) &\geq \prod_{\ell=0}^{p-1} \left(\inf_{y \in B_{\gamma_\ell}} \mathbb{P}^y \{X_{\tau_{\mathcal{M}}^{+,k}} \in B_{\gamma_{\ell+1}}\} \right) \prod_{\ell=0}^p \left(\inf_{y \in B_{\gamma_\ell}} \mathbb{P}^y \{X_{\tau_{\mathcal{M}}^{+,k}} \notin B_{\gamma_\ell}\} \right)^{k_{\ell+1} - k_\ell - 1} \\ &\geq [1 - e^{-(H_0 - \eta)/\sigma^2}]^{n-p} e^{-[H(i, j) + p\eta]/\sigma^2}. \end{aligned} \quad (6.6)$$

Since both bounds are independent of the sequence of jump times (k_1, \dots, k_p) , summing over all these sequences simply multiplies the bounds by their number. This number is exactly the number of compositions of $n+1$ into $p+1$ parts, which is known to be equal to the binomial coefficient $\binom{n}{p}$.

It remains to bound the contribution of non-optimal paths. Here we distinguish the cases $|\gamma| \leq N$, and $|\gamma| > N$. In the first case, we use (6.2), while in the second case, we use (6.3) and bound the resulting sum by a geometric series. The constant C_N bounds the number of paths of length N , and can be taken of order N^N . \square

6.2 The finite rank kernel K^\star

In what follows, it will be convenient to use the physicists' bra-ket notation, in which a signed measure μ , viewed as a row vector, is denoted $\langle \mu|$, while a test function f , viewed as a column vector, is denoted $|f\rangle$. Recall that the kernel K^\star is defined by

$$K^\star(x, dy) = \sum_{i=1}^N \mathbb{1}_{\{x \in B_i\}} \mathbb{P}^{\hat{\pi}_0^{B_i}} \{X_{\tau_{\mathcal{M}}^{+,k}} \in dy\}. \quad (6.7)$$

Denote by $\mathcal{E}_\infty^\star \subset L^\infty(\mathcal{M})$ its right image. This is an N -dimensional vector space, admitting the explicit basis

$$\mathcal{E}_\infty^\star = \text{span}(|\mathbb{1}_{B_1}\rangle, \dots, |\mathbb{1}_{B_N}\rangle).$$

In other words, \mathcal{E}_∞^\star is the vector space of bounded measurable functions which are constant on each B_i . In particular, we have

$$(K^\star |\mathbb{1}_{B_j}\rangle)(x) = \sum_{i=1}^N \mathbb{1}_{\{x \in B_i\}} \mathbb{P}^{\hat{\pi}_0^{B_i}} \{X_{\tau_{\mathcal{M}}^{+,k}} \in B_j\} = \sum_{i=1}^N \mathbb{1}_{\{x \in B_i\}} \langle \hat{\pi}_0^{B_i} | K^0 | \mathbb{1}_{B_j}\rangle$$

where $K^0 = \mathcal{M}K$ denotes the trace of K on \mathcal{M} . This can be rewritten as

$$K^\star |\mathbb{1}_{B_j}\rangle = \sum_{i=1}^N |\mathbb{1}_{B_i}\rangle \langle \hat{\pi}_0^{B_i} | K^0 | \mathbb{1}_{B_j}\rangle = \Pi^\star K^0 |\mathbb{1}_{B_j}\rangle, \quad (6.8)$$

where

$$\Pi^\star = \sum_{i=1}^N |\mathbb{1}_{B_i}\rangle \langle \pi_0^{B_i}| \quad (6.9)$$

is the projector on \mathcal{E}_∞^\star ; indeed, $(\Pi^\star)^2 = \Pi^\star$, owing to the orthonormality relation

$$\langle \pi_0^{B_i} | \mathbb{1}_{B_j} \rangle = \int_{B_j} \pi_0^{B_i}(\mathrm{d}x) = \delta_{ij}.$$

Relation (6.8) shows that $K^\star = \Pi^\star K^0$ holds on \mathcal{E}_∞^\star .

Since K^\star involves the QSDs $\pi_0^{B_i}$, it is natural to introduce the dual space

$$\mathcal{E}_1^\star = \text{span}(\langle \pi_0^{B_1} |, \dots, \langle \pi_0^{B_N} |) \subset L^1(\mathcal{M}).$$

Note that the kernel K^\star , as defined in (6.7), does not necessarily leave \mathcal{E}_1^\star invariant. This is because even if X_0 is distributed according to the QSD $\langle \pi_0^{B_i} |$, conditionally on $X_1 \in B_j$ with $j \neq i$, X_1 need not be distributed according to the QSD $\langle \pi_0^{B_j} |$. However, the kernel

$$\hat{K}^\star = K^\star \Pi^\star = \sum_{j=1}^N K^\star |\mathbb{1}_{B_j}\rangle \langle \pi_0^{B_j}| = \sum_{i,j=1}^N |\mathbb{1}_{B_i}\rangle \langle \pi_0^{B_i}| K^0 |\mathbb{1}_{B_j}\rangle \langle \pi_0^{B_j}|$$

does leave \mathcal{E}_1^\star invariant. The probabilistic interpretation of \hat{K}^\star is that it acts as K^\star , but in addition it projects the distribution of X_1 on the QSD $\langle \pi_0^{B_j} |$ whenever $X_1 \in B_j$. Note that $\Pi^\star K^\star = K^\star$, so that we have $(\hat{K}^\star)^n = (K^\star)^n \Pi^\star$ for all $n \in \mathbb{N}$. Since $\hat{K}^\star |\mathbb{1}_{B_j}\rangle = K^\star |\mathbb{1}_{B_j}\rangle$, (6.8) implies

$$\langle \pi_0^{B_i} | \hat{K}^\star | \mathbb{1}_{B_j} \rangle = \langle \pi_0^{B_i} | K^\star | \mathbb{1}_{B_j} \rangle = \langle \pi_0^{B_i} | K^0 | \mathbb{1}_{B_j} \rangle = \mathbb{P}^{\pi_0^{B_i}} \{X_{\tau_{\mathcal{M}}^+} \in B_j\},$$

showing that \hat{K}^\star , K^\star and K^0 coincide when viewed as kernels acting on the invariant spaces \mathcal{E}_∞^\star and \mathcal{E}_1^\star .

In what follows, we will be interested in processes in which time has been sped up by a factor $m = m(\sigma)$. These will involve the kernel

$$K_m^\star = \Pi^\star (K^0)^m \Pi^\star,$$

which satisfies

$$\langle \pi_0^{B_i} | K_m^\star | \mathbb{1}_{B_j} \rangle = \langle \pi_0^{B_i} | (K^0)^m | \mathbb{1}_{B_j} \rangle = \mathbb{P}^{\pi_0^{B_i}} \{X_{\tau_{\mathcal{M}}^+, m} \in B_j\}. \quad (6.10)$$

The following large-deviation estimate is an immediate consequence of Proposition 6.1.

Lemma 6.2. *Assume $m = m(\sigma)$ satisfies*

$$\lim_{\sigma \rightarrow 0} \sigma^2 \log m(\sigma) = \theta \quad (6.11)$$

for some $\theta \in (0, H_0)$. Let p be the length of the longest optimal path $\gamma : i \rightarrow j$, and let

$$H_\theta(i, j) = H(i, j) - p\theta.$$

Then for any $\eta > 0$, there exist $\sigma_0(\eta), \delta_0(\eta) > 0$ such that

$$e^{-(H_\theta(i, j) + \eta)/\sigma^2} \leq \mathbb{P}^x \{X_{\tau_{\mathcal{M}}^+, m} \in B_j\} \leq e^{-(H_\theta(i, j) - \eta)/\sigma^2} \quad (6.12)$$

holds for any $i \neq j$ and $x \in B_i$, provided $\sigma < \sigma_0(\eta)$ and the diameter of the B_k is bounded by $\delta_0(\eta)$. Furthermore, if $(N - 2)\theta \leq \hat{H}_0$ then H_θ satisfies the triangle inequality

$$H_\theta(i, \ell) + H_\theta(\ell, j) \geq H_\theta(i, j) \quad \forall i, j, \ell \in \{1, \dots, N\}. \quad (6.13)$$

PROOF: To show (6.12), it suffices to recall that all optimal paths have a length bounded by N . Therefore, the binomial coefficient $\binom{m}{p}$ is logarithmically equivalent to m^p , so that the sum over optimal paths yields a prefactor equivalent to $e^{p\theta/\sigma^2}$.

To prove (6.13), we distinguish between two cases. If $H(i, \ell) + H(\ell, j) = H(i, j)$, then ℓ lies on an optimal path $\gamma : i \rightarrow j$. This path is thus the concatenation of optimal paths $\gamma_1 : i \rightarrow \ell$ and $\gamma_2 : \ell \rightarrow j$, so that

$$H_\theta(i, j) = H(i, j) - (|\gamma_1| + |\gamma_2|)\theta = H_\theta(i, \ell) + H_\theta(\ell, j).$$

The other possibility is that $H(i, \ell) + H(\ell, j) \geq H(i, j) + \hat{H}_0$. For optimal paths $\gamma : i \rightarrow j$, $\gamma_1 : i \rightarrow \ell$ and $\gamma_2 : \ell \rightarrow j$, one gets

$$H_\theta(i, \ell) + H_\theta(\ell, j) - H_\theta(i, j) \geq \hat{H}_0 - (|\gamma_1| + |\gamma_2| - |\gamma|)\theta.$$

Since $|\gamma_1| + |\gamma_2| \leq N - 1$ and $|\gamma| \geq 1$, the result follows. \square

6.3 The truncated kernel K_{tr}^0

Denote by λ_k^0 , $|\phi_k^0\rangle$ and $\langle\pi_k^0|$ the orthonormalised eigen-elements of K^0 , and introduce the truncated kernel K_{tr}^0 associated with the N largest eigenvalues.

We denote the right and left invariant subspaces of K_{tr}^0 by

$$\begin{aligned}\mathcal{E}_\infty^0 &= \text{span}(|\phi_0^0\rangle, \dots, |\phi_{N-1}^0\rangle), \\ \mathcal{E}_1^0 &= \text{span}(\langle\pi_0^0|, \dots, \langle\pi_{N-1}^0|).\end{aligned}$$

Our aim is now to construct another basis of the subspaces \mathcal{E}_∞^0 and \mathcal{E}_1^0 , which is close to the basis formed by the QSDs $\langle\pi_0^{B_i}|$ and the indicators $|\mathbb{1}_{B_j}\rangle$. A natural idea is to set, for some $m \in \mathbb{Z}^*$,

$$\langle\mu_i| = \langle\pi_0^{B_i}|(K_{\text{tr}}^0)^m, \quad |\psi_j\rangle = (K_{\text{tr}}^0)^{-m}|\mathbb{1}_{B_j}\rangle,$$

where $(K_{\text{tr}}^0)^{-1}$ is the generalised inverse of K_{tr}^0 , and $(K_{\text{tr}}^0)^{-m} = ((K_{\text{tr}}^0)^{-1})^m$. Indeed, we then have $\langle\mu_i| \in \mathcal{E}_1^0$ and $|\psi_j\rangle \in \mathcal{E}_\infty^0$ by construction. Unfortunately, the basis is not orthonormal, because $(K_{\text{tr}}^0)^{-1}K_{\text{tr}}^0 = K_{\text{tr}}^0(K_{\text{tr}}^0)^{-1} = \Pi^0$, where

$$\Pi^0 = \sum_{k=0}^{N-1} |\phi_k^0\rangle\langle\pi_k^0| \tag{6.14}$$

is the projector on the invariant subspaces of K_{tr}^0 . Therefore, in general we will have

$$\langle\mu_i|\psi_j\rangle = \langle\pi_0^{B_i}|\Pi^0|\mathbb{1}_{B_j}\rangle \neq \delta_{ij},$$

A solution to this problem is to modify the definition of $\langle\pi_0^{B_i}|$ and $|\psi_j\rangle$ as follows.

Lemma 6.3. *Let $\Pi_\perp^0 = \text{id} - \Pi^0$. If $\langle\pi_0^{B_i}|\Pi_\perp^0\Pi^0 \neq 0$ for $j = 1, \dots, N$, then the basis defined by*

$$\langle\mu_i| = \langle\pi_0^{B_i}|[\text{id} - \Pi_\perp^0\Pi^0]^{-1}\Pi^0, \quad |\psi_j\rangle = \Pi^0|\mathbb{1}_{B_j}\rangle$$

satisfies $\langle\mu_i|\psi_j\rangle = \delta_{ij}$ for all $i, j \in \{1, \dots, N\}$. Furthermore,

$$\sum_{i=1}^N |\psi_i\rangle\langle\mu_i| = \Pi^0. \tag{6.15}$$

PROOF: Since $\langle \dot{\pi}_0^{B_i} | \Pi^\star = \langle \dot{\pi}_0^{B_i} |$, we have

$$\langle \dot{\pi}_0^{B_i} | [\text{id} - \Pi_\perp^0 \Pi^\star] = \langle \dot{\pi}_0^{B_i} | [\text{id} - \Pi^\star + \Pi^0 \Pi^\star] = \langle \dot{\pi}_0^{B_i} | \Pi^0 \Pi^\star \neq 0$$

for $j = 1, \dots, N$, so that $\langle \mu_j |$ is indeed well-defined. Furthermore,

$$\begin{aligned} \langle \mu_i | \psi_j \rangle &= \langle \dot{\pi}_0^{B_i} | [\text{id} - \Pi_\perp^0 \Pi^\star]^{-1} \Pi^0 | \mathbb{1}_{B_j} \rangle \\ &= \langle \dot{\pi}_0^{B_i} | [\text{id} - \Pi_\perp^0 \Pi^\star]^{-1} [\text{id} - \Pi_\perp^0] | \mathbb{1}_{B_j} \rangle \\ &= \langle \dot{\pi}_0^{B_i} | [\text{id} - \Pi_\perp^0 \Pi^\star]^{-1} [\text{id} - \Pi_\perp^0 \Pi^\star] | \mathbb{1}_{B_j} \rangle \\ &= \langle \dot{\pi}_0^{B_i} | \mathbb{1}_{B_j} \rangle \\ &= \delta_{ij}, \end{aligned}$$

where we have used the fact that $\Pi^\star | \mathbb{1}_{B_j} \rangle = | \mathbb{1}_{B_j} \rangle$ to obtain the third line. As a consequence, the left-hand side of (6.15) is a projector of rank N . Since it is of the form $\Pi^0 M \Pi^0$ for a linear operator M , its left and right images are given by \mathcal{E}_1^0 and \mathcal{E}_∞^0 , which implies (6.14). \square

6.4 Comparison of transition probabilities

Our aim is now to show that for an appropriate $m = m(\sigma) \in \mathbb{N}$, one has

$$\langle \mu_i | (K_{\text{tr}}^0)^m | \psi_j \rangle \sim \langle \dot{\pi}_0^{B_i} | (K^0)^m | \mathbb{1}_{B_j} \rangle = \mathbb{P}^{\dot{\pi}_0^{B_i}} \{ X_{\tau_{\mathcal{M}}^+, m} \in B_j \}.$$

The Neumann series representation of $[\text{id} - \Pi_\perp^0 \Pi^\star]^{-1}$ and the definition (6.9) of Π^\star yield

$$\begin{aligned} \langle \mu_i | (K_{\text{tr}}^0)^m | \psi_j \rangle &= \langle \dot{\pi}_0^{B_i} | [\text{id} - \Pi_\perp^0 \Pi^\star]^{-1} (K_{\text{tr}}^0)^m | \mathbb{1}_{B_j} \rangle \\ &= \sum_{n \geq 0} \langle \dot{\pi}_0^{B_i} | (\Pi_\perp^0 \Pi^\star)^n (K_{\text{tr}}^0)^m | \mathbb{1}_{B_j} \rangle \\ &= \langle \dot{\pi}_0^{B_i} | (K_{\text{tr}}^0)^m | \mathbb{1}_{B_j} \rangle + \sum_{\ell=1}^N \sum_{n \geq 1} \langle \dot{\pi}_0^{B_i} | (\Pi_\perp^0 \Pi^\star)^n | \mathbb{1}_{B_\ell} \rangle \langle \dot{\pi}_0^{B_\ell} | (K_{\text{tr}}^0)^m | \mathbb{1}_{B_j} \rangle. \end{aligned} \quad (6.16)$$

For an appropriate m , the first term on the right-hand side is indeed close to $\langle \dot{\pi}_0^{B_i} | (K^0)^m | \mathbb{1}_{B_j} \rangle$. We thus need to bound the remaining terms. We introduce the notation

$$\varepsilon_{ij} = \varepsilon_{ij}^{(1)} = \langle \dot{\pi}_0^{B_i} | \Pi_\perp^0 \Pi^\star | \mathbb{1}_{B_j} \rangle = \langle \dot{\pi}_0^{B_i} | \Pi^\star - \Pi^0 | \mathbb{1}_{B_j} \rangle = \delta_{ij} - \langle \dot{\pi}_0^{B_i} | \Pi^0 | \mathbb{1}_{B_j} \rangle. \quad (6.17)$$

For every $n \geq 2$, (6.9) allows us to write

$$\varepsilon_{ij}^{(n)} = \langle \dot{\pi}_0^{B_i} | (\Pi_\perp^0 \Pi^\star)^n | \mathbb{1}_{B_j} \rangle = \sum_{\ell_1, \dots, \ell_{n-1}=1}^N \varepsilon_{i\ell_1} \varepsilon_{\ell_1 \ell_2} \dots \varepsilon_{\ell_{n-1} j}.$$

With these notations, (6.16) becomes

$$\langle \mu_i | (K_{\text{tr}}^0)^m | \psi_j \rangle = \langle \dot{\pi}_0^{B_i} | (K_{\text{tr}}^0)^m | \mathbb{1}_{B_j} \rangle + \sum_{\ell=1}^N \langle \dot{\pi}_0^{B_\ell} | (K_{\text{tr}}^0)^m | \mathbb{1}_{B_j} \rangle \sum_{n \geq 1} \varepsilon_{i\ell}^{(n)}. \quad (6.18)$$

The kernel $K_m^\star = \Pi^\star (K^0)^m \Pi^\star$ has the same image as \hat{K}^\star and Π^\star . Thus the Riesz projector formalism shows that

$$\Pi^\star - \Pi^0 = \frac{1}{2\pi i} \int_{\Gamma} [(z \text{id} - K_m^\star)^{-1} - (z \text{id} - (K_{\text{tr}}^0)^m)^{-1}] dz, \quad (6.19)$$

provided Γ is a contour in the complex plane encircling all (nonzero) eigenvalues of K^\star and K_{tr}^0 . One option would be to use a resolvent identity and a bound on $\|K^\star - K_{\text{tr}}^0\|$, but this would yield an estimate which is uniform in i, j , which is not sharp enough for our purpose.

To obtain a sharper bound, we note that the Cayley–Hamilton theorem implies that if K is an operator of finite rank N , then K^N is a linear combination of $\text{id}, K, K^2, \dots, K^{N-1}$. This implies that the resolvent of K can also be expressed in terms of a finite number of powers of K , as shows the following result.

Lemma 6.4. *Let K be an operator of finite rank N , and let $\lambda_1, \dots, \lambda_N$ be its nonzero eigenvalues. Then the resolvent of K can be written in the form*

$$(z \text{id} - K)^{-1} = \frac{1}{c_K(z)} \sum_{n=0}^{N-1} \alpha_n(z) K^n, \quad (6.20)$$

where $c_K(z) = \det(z \text{id} - K) = \prod_{k=1}^N (z - \lambda_k)$ is the characteristic polynomial of K , and $\alpha_n(z)$ is a polynomial of degree $N - 1 - n$ in z . More precisely, one has

$$c_K(z) = \sum_{n=0}^N c_n z^n \quad \Rightarrow \quad \alpha_n(z) = \sum_{i=0}^{N-1-n} c_{i+n+1} z^i.$$

PROOF: Multiply (6.20) by $(z \text{id} - K)$ and use the relations $z \alpha_n(z) - \alpha_{n-1}(z) = -c_n$, $\alpha_{N-1}(z) = c_N$, $z \alpha_0(z) = c_K(z) - c_0$ and $c_K(K) = 0$. Also see for instance [Hou98], which gives an iterative construction implying the above expression for the $\alpha_n(z)$. \square

The key estimates that will allow us to control ε_{ij} are the following two propositions. They will allow us to control the error made when projecting the law of the process on its QSD every m steps, and thus to compare matrix elements involving $(K^0)^m$ and K_m^\star . This approach is somewhat related in spirit to the one used in [MOS89]. To lighten notations, we write

$$\varrho_i = \frac{|\lambda_1^{B_i}|}{\lambda_0^{B_i}}$$

for the spectral gap of the trace process killed upon leaving B_i , where $\lambda_1^{B_i}$ is the next-to-leading eigenvalue of this process.

Proposition 6.5. *For any $\eta > 0$, there exist $\sigma_0(\eta), \delta_0(\eta) > 0$ such that*

$$\mathbb{P}^x \{X_{\tau_{\mathcal{M}}^{+,m}} \in B_j\} = \mathbb{P}^{\tilde{\pi}_0^{B_i}} \{X_{\tau_{\mathcal{M}}^{+,m}} \in B_j\} [1 + r_{\eta,m}(\sigma)] \quad (6.21)$$

holds for any $m \in \mathbb{N}$, any $i, j \in \{1, \dots, N\}$ and any $x \in B_i$, provided $\sigma < \sigma_0(\eta)$ and the diameter of the B_ℓ is bounded by $\delta_0(\eta)$. There exists a constant C , independent of σ, m and η , such that the remainder in (6.21) satisfies

$$|r_{\eta,m}(\sigma)| \leq C \left[\varrho_i^{m_1} + \frac{\binom{m}{p} - \binom{m-m_1}{p}}{\binom{m-m_1}{p}} \frac{e^{2p\eta/\sigma^2}}{[1 - e^{-(H_0-\eta)/\sigma^2}]^{m-p}} + \delta_{ij} m e^{-(H_0-\eta)/\sigma^2} \right] \quad (6.22)$$

for any $m_1 < m$, where p is the length of the longest optimal path $\gamma: i \rightarrow j$.

PROOF: In the case $i = j$, the large-deviation principle shows that

$$\mathbb{P}^x \{X_{\tau_{\mathcal{M}}^{+,m}} \in B_i\} = 1 - \mathcal{O}(m e^{-(H_0-\eta)/\sigma^2}),$$

so that the result follows at once by integrating this relation against $\hat{\pi}_0^{B_i}(x)$.

In the case $i \neq j$, we consider first the case where the optimal path from i to j has length 1. Then the decomposition (6.4) of the transition probability involves only m terms $Q_k(x)$, and the bounds (6.5) and (6.6) reduce to

$$\left[1 - e^{-(H_0 - \eta)/\sigma^2}\right]^{m-1} e^{-(H(i,j) + \eta)/\sigma^2} \leq Q_k(x) \leq e^{-(H(i,j) - \eta)/\sigma^2} \quad (6.23)$$

uniformly in k , which provides an upper bound on the ratio between the largest and smallest $Q_k(x)$. We now split the $Q_k(x)$ into “bad” and “good” terms, the good ones being those with $k \geq m_1$ chosen sufficiently large that the process has time to relax to the QSD $\hat{\pi}_0^{B_i}$ before making the transition to B_j . With this splitting, we have

$$\begin{aligned} \sum_{k=m_1+1}^m Q_k(x) &\leq \sum_{k=1}^m Q_k(x) \leq \sum_{k=m_1+1}^m Q_k(x) + \sum_{k=1}^{m_1} Q_k(x) \\ &\leq \sum_{k=m_1+1}^m Q_k(x) \left[1 + \frac{m_1}{m - m_1} \frac{e^{2\eta/\sigma^2}}{\left[1 - e^{-(H_0 - \eta)/\sigma^2}\right]^{m-1}} \right]. \end{aligned}$$

In order to derive a sharper estimate for the good Q_k , we rewrite them in the form

$$Q_k(x) = \int_{B_j} \int_{B_i} (\hat{k}^{B_i})^{k-1}(x, y) k^0(y, z) (\hat{K}^{B_j})^{n-k}(z, B_j) dy dz.$$

Using the facts that

$$\begin{aligned} (\hat{k}^{B_i})^{k-1}(x, y) &= (\hat{\lambda}_0^{B_i})^{k-1} \hat{\pi}_0^{B_i}(y) + \mathcal{O}(|\hat{\lambda}_1^{B_i}|^{k-1}), \\ \int_{B_i} \hat{\pi}_0^{B_i}(x_1) (\hat{k}^{B_i})^{k-1}(x_1, y) dx_1 &= (\hat{\lambda}_0^{B_i})^{k-1} \hat{\pi}_0^{B_i}(y), \end{aligned}$$

we obtain

$$Q_k(x) = \int_{B_i} \hat{\pi}_0^{B_i}(x_1) Q_k(x_1) dx_1 \left[1 + \mathcal{O}(\hat{\rho}_i^{k-1}) \right].$$

Collecting terms and bounding the contribution of non-optimal paths as in the proof of Proposition 6.1, we get

$$\frac{\mathbb{P}^x\{X_{\tau_{\mathcal{M}}^+, m} \in B_j\}}{\mathbb{P}^{\hat{\pi}_0^{B_i}}\{X_{\tau_{\mathcal{M}}^+, m} \in B_j\}} = 1 + \mathcal{O}(\hat{\rho}_i^{k-1}) + \mathcal{O}\left(\frac{m_1}{m - m_1} \frac{e^{2\eta/\sigma^2}}{\left[1 - e^{-(H_0 - \eta)/\sigma^2}\right]^{m-1}}\right)$$

as claimed. To extend the proof to optimal paths of length larger than 1, we proceed in an analogous way, where the bad terms are those for which $k_1 \leq m_1$. The ratio of binomial coefficients in (6.22) is the ratio between bad terms and good terms. \square

Proposition 6.6. *Fix $i \neq j$ and a bounded, measurable, real-valued test function F supported in B_j . There exists a constant C such that for any $x \in B_i$ and any $m_1 < m$, one has*

$$\frac{\mathbb{E}^x[F(X_{\tau_{\mathcal{M}}^+, m})]}{\mathbb{P}^x\{X_{\tau_{\mathcal{M}}^+, m} \in B_j\}} = [1 - p_{\eta, m}(\sigma)] \mathbb{E}^{\hat{\pi}_0^{B_j}}[F] [1 + q_{\eta, m}(\sigma)] + p_{\eta, m}(\sigma) \mathbb{E}^x[\mathbb{E}^{X_{\tau_{\mathcal{M}}^+, m}}[F]], \quad (6.24)$$

where the error terms satisfy

$$|p_{\eta, m}(\sigma)| \leq C \frac{\binom{m}{p} - \binom{m-m_1}{p}}{\binom{m-m_1}{p}} \frac{e^{2p\eta/\sigma^2}}{\left[1 - e^{-(H_0 - \eta)/\sigma^2}\right]^{m-p}}, \quad |q_{\eta, m}(\sigma)| \leq C \left[\hat{\rho}_j^{m_1} + m e^{-(H_0 - \eta)/\sigma^2} \right]. \quad (6.25)$$

Here p is again the length of the longest optimal path $\gamma : i \rightarrow j$. If $i = j$, then the bound (6.24) holds with $|p_{\eta, m}(\sigma)| \leq m e^{-(H_0 - \eta)/\sigma^2}$.

PROOF: For $i \neq j$, let us consider again the case $p = 1$ first. We introduce the stopping time

$$\tau = \inf\{n > 0: X_{\tau_{\mathcal{M}}^{+,n}} \in B_j\} - 1.$$

Since the optimal path from i to j visits j only once, at the very end of the path, τ will be with overwhelming probability equal to the last time the process visits $\mathcal{M} \setminus \{j\}$. Then we write

$$\mathbb{E}^x[F(X_{\tau_{\mathcal{M}}^{+,m}})] = \sum_{k=0}^{m-m_1} E_k(x) + \sum_{k=m-m_1+1}^{m-1} E_k(x), \quad (6.26)$$

where

$$E_k(x) = \mathbb{E}^x[\mathbb{1}_{\{\tau=k\}} F(X_{\tau_{\mathcal{M}}^{+,k}})].$$

Note that

$$\hat{Q}_k(x) = \mathbb{P}^x\{\tau = k\} = \mathbb{P}^x\{X_{\tau_{\mathcal{M}}^{+,k}} \notin B_j, X_{\tau_{\mathcal{M}}^{+,k+1}} \in B_j, \dots, X_{\tau_{\mathcal{M}}^{+,m}} \in B_j\}$$

has similar properties as $Q_k(x)$ in the previous proof. In particular, it again satisfies (6.23) owing to the large-deviation principle. In the same spirit as in the previous proof, we consider the $E_k(x)$ in the first sum in (6.26) as good terms, and those in the second sum as bad terms.

To estimate the good terms, we write for $k \leq m - m_1$

$$E_k(x) = \int_{B_j} \int_{\mathcal{M} \setminus B_j} (\mathring{k}^{B_i})^{k-1}(x, y) k^0(y, z) (\mathring{K}^{B_j})^{m-k}(z, F) dy dz,$$

where

$$\begin{aligned} (\mathring{K}^{B_j})^{m-k}(z, F) &= \mathbb{E}^z[F(X_{\tau_{\mathcal{M}}^{+,m-k}})] \\ &= (\mathring{\lambda}_0^{B_j})^{m-k} \mathbb{E}^{\mathring{\pi}_0^{B_j}}[F][1 + \mathcal{O}(\mathring{\varrho}_j^{m-k})]. \end{aligned}$$

Since $\mathring{\lambda}_0^{B_j} = 1 - \mathcal{O}(e^{-(H_0 - \eta)/\sigma^2})$, it follows that

$$E_k(x) = \mathbb{P}^x\{\tau = k - 1\} \mathbb{E}^{\mathring{\pi}_0^{B_j}}[F][1 + \mathcal{O}(\mathring{\varrho}_j^{m-k}) + \mathcal{O}(m e^{-(H_0 - \eta)/\sigma^2})],$$

so that the sum of good terms satisfies

$$\sum_{k=0}^{m-m_1} E_k(x) = \mathbb{P}^x\{\tau \leq m - m_1\} \mathbb{E}^{\mathring{\pi}_0^{B_j}}[F][1 + \mathcal{O}(\mathring{\varrho}_j^{m_1}) + \mathcal{O}(m e^{-(H_0 - \eta)/\sigma^2})].$$

This implies the result, with

$$p_{\eta, m}(\sigma) = \mathbb{P}^x\{\tau > m - m_1 \mid X_{\tau_{\mathcal{M}}^{+,m}} \in B_j\} = \frac{\sum_{k=m-m_1+1}^{m-1} \hat{Q}_k(x)}{\sum_{k=0}^{m-1} \hat{Q}_k(x)},$$

which satisfies indeed the bound (6.25) with $p = 1$, thanks to (6.23). The case of general p then follows in a similar way, by counting the number of good and bad terms.

In the case $i = j$, the result follows by distinguishing the cases where $X_{\tau_{\mathcal{M}}^{+,k}} \in B_i$ for all $k \leq m$, and the unlikely complementary event. \square

Corollary 6.7. *For any $\eta > 0$, there exist $\sigma_0(\eta), \delta_0(\eta) > 0$ such that*

$$\left| \frac{\langle \mathring{\pi}_0^{B_i} | (K^0)^{nm} | \mathbb{1}_{B_j} \rangle}{\langle \mathring{\pi}_0^{B_i} | (K_m^*)^n | \mathbb{1}_{B_j} \rangle} - 1 \right| \leq R_{\eta, m, n}(\sigma) := \left[1 + q_{\eta, m}(\sigma) + p_{\eta, m}(\sigma) r_{\eta, nm}(\sigma) \right]^{n-1} - 1$$

holds for all $\sigma < \sigma_0$ and all i, j , provided the diameter of the B_k is bounded by $\delta_0(\eta)$.

PROOF: For $n = 1$, the result follows from (6.10) with $R_{\eta,m,1}(\sigma) = 0$. To prove the result for $n \geq 2$, we first observe that

$$\begin{aligned} \langle \hat{\pi}_0^{B_i} | (K^0)^{nm} | \mathbb{1}_{B_j} \rangle &= \sum_{\ell=1}^N \mathbb{E}^{\hat{\pi}_0^{B_i}} \left[\mathbb{1}_{\{X_{\tau_{\mathcal{M}}^+,m} \in B_\ell\}} \mathbb{P}^{X_{\tau_{\mathcal{M}}^+,m}} \{X_{\tau_{\mathcal{M}}^+, (n-1)m} \in B_j\} \right] \\ &= \sum_{\ell=1}^N \mathbb{E}^{\hat{\pi}_0^{B_i}} [F_\ell(X_{\tau_{\mathcal{M}}^+,m})], \end{aligned}$$

where

$$F_\ell(x) = \mathbb{1}_{\{x \in B_\ell\}} \mathbb{P}^x \{X_{\tau_{\mathcal{M}}^+, (n-1)m} \in B_j\}.$$

Proposition 6.6 shows that

$$\begin{aligned} \mathbb{E}^{\hat{\pi}_0^{B_i}} [F_\ell(X_{\tau_{\mathcal{M}}^+,m})] &= \mathbb{P}^{\hat{\pi}_0^{B_i}} \{X_{\tau_{\mathcal{M}}^+,m} \in B_\ell\} \\ &\quad \times \left[(1 - p_{\eta,m}) \mathbb{E}^{\hat{\pi}_0^{B_\ell}} [F_\ell] (1 + q_{\eta,m}) + p_{\eta,m} \mathbb{E}^{\hat{\pi}_0^{B_i}} \left[\mathbb{E}^{X_{\tau_{\mathcal{M}}^+,m}} [F_\ell] \right] \right]. \end{aligned}$$

Now we note that

$$\mathbb{E}^{\hat{\pi}_0^{B_\ell}} [F_\ell] = \mathbb{P}^{\hat{\pi}_0^{B_\ell}} \{X_{\tau_{\mathcal{M}}^+, (n-1)m} \in B_\ell\},$$

while Proposition 6.5 implies that for any $x \in B_\ell$, one has

$$\mathbb{E}^x [F_\ell] = \mathbb{P}^{\hat{\pi}_0^{B_\ell}} \{X_{\tau_{\mathcal{M}}^+, (n-1)m} \in B_\ell\} [1 + r_{\eta, (n-1)m}].$$

It follows that

$$\begin{aligned} \mathbb{E}^{\hat{\pi}_0^{B_i}} [F_\ell(X_{\tau_{\mathcal{M}}^+,m})] &= \mathbb{P}^{\hat{\pi}_0^{B_i}} \{X_{\tau_{\mathcal{M}}^+,m} \in B_\ell\} \mathbb{P}^{\hat{\pi}_0^{B_\ell}} \{X_{\tau_{\mathcal{M}}^+, (n-1)m} \in B_\ell\} [1 + R_n] \\ &= \langle \hat{\pi}_0^{B_i} | (K^0)^m | \mathbb{1}_{B_\ell} \rangle \langle \hat{\pi}_0^{B_\ell} | (K^0)^{(n-1)m} | \mathbb{1}_{B_j} \rangle [1 + R_n], \end{aligned} \quad (6.27)$$

where the remainder

$$R_n = (1 - p_{\eta,m})(1 + q_{\eta,m}) + p_{\eta,m}(1 + r_{\eta, (n-1)m}) - 1$$

satisfies

$$0 \leq R_n \leq q_{\eta,m} + p_{\eta,m} r_{\eta, (n-1)m}.$$

Summing (6.27) over ℓ shows that

$$\langle \hat{\pi}_0^{B_i} | (K^0)^{nm} | \mathbb{1}_{B_j} \rangle = \langle \hat{\pi}_0^{B_i} | (K^0)^m \Pi^* (K^0)^{(n-1)m} | \mathbb{1}_{B_j} \rangle [1 + R_n],$$

and the result follows by induction on n . □

Corollary 6.8. *Assume m satisfies (6.11). Then there exists a constant C such that for any $\eta > 0$*

$$|\varepsilon_{ij}| \leq C \left[\delta_{ij} e^{-(H_0 - \eta)/\sigma^2} + N e^{-[H_\theta(i,j) - \eta]/\sigma^2} [R_{\eta,m,N}(\sigma) + \varrho^m + e^{-(H_0 - \eta)/\sigma^2}] \right] \quad (6.28)$$

holds for $1 \leq i, j \leq N$, provided σ and the B_k are sufficiently small as a function of η .

PROOF: Let Γ be a contour encircling the (nonzero) eigenvalues of $(K_{\text{tr}}^0)^m$, and staying at a distance of order 1 from all eigenvalues of $(K^0)^m$. Note that this is possible for σ small enough by Proposition 3.1. Recall that $K_{\text{tr}}^0 = \Pi^0 K^0$, where Π^0 is the Riesz projector associated with Γ . Using (6.17), (6.19), Lemma 6.4 and Corollary 6.7, we obtain

$$\begin{aligned} \varepsilon_{ij} &= \sum_{n=0}^{N-1} \frac{1}{2\pi i} \int_{\Gamma} \langle \dot{\pi}_0^{B_i} | \left[\frac{\alpha_n^*(z)}{c^*(z)} (K_m^*)^n - \frac{\alpha_n^0(z)}{c^0(z)} (K_{\text{tr}}^0)^{nm} \right] | \mathbb{1}_{B_j} \rangle dz \\ &= \sum_{n=1}^{N-1} \frac{1}{2\pi i} \int_{\Gamma} \left[\frac{\alpha_n^*(z)}{c^*(z)} [1 + \mathcal{O}(R_{\eta,m,n}(\sigma))] - \frac{\alpha_n^0(z)}{c^0(z)} [1 + \mathcal{O}(\varrho^{nm})] \right] dz \langle \dot{\pi}_0^{B_i} | (K^0)^{nm} | \mathbb{1}_{B_j} \rangle \\ &\quad + \frac{1}{2\pi i} \int_{\Gamma} \left[\frac{\alpha_0^*(z)}{c^*(z)} - \frac{\alpha_0^0(z)}{c^0(z)} \right] dz \delta_{ij}, \end{aligned}$$

where $c^*(z)$ and the $\alpha_n^*(z)$ are the coefficients of the decomposition (6.20) of K_m^* , and $c^0(z)$ and the $\alpha_n^0(z)$ are those of the decomposition of $(K_{\text{tr}}^0)^m$. The contour Γ has been chosen such that the characteristic polynomials $c^*(z)$ and $c^0(z)$ are bounded away from 0. Proposition 3.1 shows that the eigenvalues of K^* and K^0 are at distance $\mathcal{O}(e^{-(H_0-\eta)/\sigma^2})$ from each other. This shows that $\alpha_n^*(z)/c^*(z) = (\alpha_n^0(z)/c^0(z)) [1 + \mathcal{O}(e^{-(H_0-\eta)/\sigma^2})]$ on the contour Γ . Hence the result follows from Lemma 6.2. \square

We can now choose a value of m_1 yielding the smallest possible error terms. Since N is a finite constant, we no longer indicate the dependence of the error terms on N . The bound (5.12) implies that

$$|\dot{\lambda}_1^{B_i}| \leq \delta^{1/n_0(\sigma)} = \exp \left\{ -\frac{\log(\delta^{-1})}{n_0(\sigma)} \right\}$$

for a constant $\delta < 1$ related to L . The choice

$$m_1 = \frac{H_0 n_0(\sigma)}{\log(\delta^{-1}) \sigma^2}$$

then yields

$$\dot{\varrho}_i^{m_1} \leq \frac{|\dot{\lambda}_1^{B_i}|^{m_1}}{(\dot{\lambda}_0^{B_i})^{m_1}} \leq 2 e^{-H_0/\sigma^2}. \quad (6.29)$$

Since m is assumed to satisfy (6.11), we have $2m_1 \leq m \leq e^{(H_0-\eta)/\sigma^2}$ for η small enough. Further note that for these m_1 and p or order 1,

$$\frac{\binom{m}{p} - \binom{m-m_1}{p}}{\binom{m-m_1}{p}} = \mathcal{O} \left(p \frac{m_1}{m} \right).$$

Substituting in (6.22) yields

$$|r_{\eta,m}(\sigma)| \leq C_1 \left[e^{-H_0/\sigma^2} + \frac{1}{m} \frac{n_0(\sigma)}{\sigma^2} e^{2\eta/\sigma^2} \right] \leq 2C_1 e^{-(\theta-3\eta)/\sigma^2}.$$

The error term $q_{\eta,m}(\sigma)$ also satisfies (6.29), while $p_{\eta,m}(\sigma)$ is at most of order $r_{\eta,m}(\sigma)$. Therefore,

$$|R_{\eta,m,N}(\sigma)| \leq C_2 e^{-2(\theta-3\eta)/\sigma^2}.$$

Furthermore, Proposition 3.1 implies that for these m , ϱ^m is negligible with respect to $r_{\eta,m}(\sigma)$, provided $2\eta < \theta$. Writing

$$\hat{H}_{\theta}(i, j) = H_0 \delta_{ij} + H_{\theta}(i, j)(1 - \delta_{ij}),$$

we can rewrite the bound (6.28) as

$$|\varepsilon_{ij}| \leq e^{-(\hat{H}_\theta(i,j)+\theta-4\eta)/\sigma^2}.$$

Note that the fact that the error $R_{\eta,m,N}$ involves the product $p_{\eta,m}r_{\eta,m}$ instead of only one of these terms has improved the accuracy of the approximation.

Remark 6.9. Corollary 6.8 shows in particular that the assumption that $\langle \pi_0^{B_i} | \Pi^0 \Pi^\star \neq 0$, made in Lemma 6.3, is satisfied for small enough σ . Indeed, writing Π^0 as a contour integral as in the proof of the Corollary shows that $\langle \pi_0^{B_i} | \Pi^0 | \mathbb{1}_{B_j} \rangle = \delta_{ij} + \mathcal{O}(e^{-[H_0-\eta]/\sigma^2})$. Therefore, $\langle \pi_0^{B_i} | \Pi^0 \Pi^\star = \sum_{j=1}^N \langle \pi_0^{B_i} | \Pi^0 | \mathbb{1}_{B_j} \rangle \langle \pi_0^{B_j} |$ is exponentially close to $\langle \pi_0^{B_i} |$. \diamond

Corollary 6.10. *For any $\eta \in (0, \theta)$, there exist $\sigma_0, \delta_0 > 0$ such that for all $1 \leq i \neq j \leq N$,*

$$\langle \mu_i | (K_{\text{tr}}^0)^m | \psi_j \rangle = \langle \pi_0^{B_i} | (K^0)^m | \mathbb{1}_{B_j} \rangle [1 + \mathcal{O}(e^{-(\theta-\eta)/\sigma^2})],$$

provided $\sigma < \sigma_0$ and all B_k have a diameter smaller than δ_0 .

PROOF: Proceeding by induction on n and using the triangle inequality (6.13), one easily obtains the bounds

$$|\varepsilon_{ij}^{(n)}| \leq e^{-[\hat{H}_\theta(i,j)+n(\theta-4\eta)]/\sigma^2} \quad \forall n \geq 1, \quad \sum_{n \geq 1} |\varepsilon_{ij}^{(n)}| \leq 2e^{-[\hat{H}_\theta(i,j)+\theta-4\eta]/\sigma^2}. \quad (6.30)$$

It follows that the remainder in (6.18) satisfies

$$\left| \sum_{\ell=1}^N \langle \pi_0^{B_i} | (K_{\text{tr}}^0)^m | \mathbb{1}_{B_\ell} \rangle \sum_{n \geq 1} \varepsilon_{\ell j}^{(n)} \right| \leq 2Ne^{-[\hat{H}_\theta(i,j)+\theta-5\eta]/\sigma^2},$$

provided σ and the B_k are sufficiently small, depending on η . On the other hand, we have

$$\langle \pi_0^{B_i} | (K_{\text{tr}}^0)^m | \mathbb{1}_{B_j} \rangle = \langle \pi_0^{B_i} | (K^0)^m | \mathbb{1}_{B_j} \rangle + \mathcal{O}(\varrho^m) \geq e^{-[\hat{H}_\theta(i,j)+\eta]/\sigma^2} + \mathcal{O}(\varrho^m).$$

Proposition 3.1 shows that for m as above, the error term $\mathcal{O}(\varrho^m)$ is indeed negligible, yielding the claimed exponentially small multiplicative error, after redefining η . \square

The following result shows in which sense the new basis vectors $\langle \mu_i |$ and $|\psi_j\rangle$ are close to $\langle \pi_0^{B_i} |$ and $|\mathbb{1}_{B_j}\rangle$.

Proposition 6.11. *The basis vectors satisfy*

$$\langle \mu_i | \mathbb{1}_{B_j} \rangle = \delta_{ij} \quad \text{and} \quad \langle \pi_0^{B_i} | \psi_j \rangle = \delta_{ij} - \varepsilon_{ij} \quad (6.31)$$

for all $1 \leq i, j \leq N$. Furthermore, for any $\eta > 0$, one has

$$\|\psi_j - \mathbb{1}_{B_j}\|_\infty = \sup_{x \in \mathcal{M}} |\psi_j(x) - \mathbb{1}_{B_j}(x)| \leq e^{-[\hat{H}_j - \eta]/\sigma^2} \quad (6.32)$$

provided σ and the diameters of the B_i are small enough, where

$$\hat{H}_j = \min_{i \neq j} \left[H(i, j) - \max_{\gamma: i \rightarrow j} |\gamma| \right] \geq H_0 - (N-1)\theta. \quad (6.33)$$

PROOF: The relations (6.31) follow immediately from the definitions, since

$$\begin{aligned}\langle \mu_i | \mathbb{1}_{B_j} \rangle &= \langle \dot{\pi}_0^{B_i} | [\text{id} - \Pi_\perp^0 \Pi^\star]^{-1} \Pi^0 | \mathbb{1}_{B_j} \rangle = \langle \mu_i | \psi_j \rangle = \delta_{ij}, \\ \langle \dot{\pi}_0^{B_i} | \psi_j \rangle &= \langle \dot{\pi}_0^{B_i} | \Pi^0 | \mathbb{1}_{B_j} \rangle = \delta_{ij} - \varepsilon_{ij}.\end{aligned}$$

In order to prove (6.32), we proceed as in the proof of Corollary 6.8, writing for $x \in B_i$

$$\begin{aligned}\psi_j(x) - \delta_{ij} &= \langle \delta_x | \Pi^0 - \Pi^\star | \mathbb{1}_{B_j} \rangle \\ &= \sum_{n=0}^{N-1} \frac{1}{2\pi i} \int_{\Gamma} \langle \delta_x | \left[\frac{\alpha_n^0(z)}{c^0(z)} (K_{\text{tr}}^0)^{nm} - \frac{\alpha_n^\star(z)}{c^\star(z)} (K_m^\star)^n \right] | \mathbb{1}_{B_j} \rangle dz.\end{aligned}$$

Propositions 3.1 and 6.5 imply that for $0 \leq n \leq N-1$, one has

$$\begin{aligned}\langle \delta_x | (K_{\text{tr}}^0)^{nm} | \mathbb{1}_{B_j} \rangle &= \mathbb{P}^x \{X_{\tau_{\mathcal{M}}^{+,nm}} \in B_j\} + \mathcal{O}(\varrho^{nm}) \\ &= \mathbb{P}^{\dot{\pi}_0^{B_i}} \{X_{\tau_{\mathcal{M}}^{+,nm}} \in B_j\} [1 + r_{\eta,m}(\sigma)] + \mathcal{O}(\varrho^{nm}),\end{aligned}$$

while the definition of K_m^\star implies

$$\langle \delta_x | (K_m^\star)^n | \mathbb{1}_{B_j} \rangle = \langle \dot{\pi}_0^{B_i} | (K_m^\star)^n | \mathbb{1}_{B_j} \rangle = \mathbb{P}^{\dot{\pi}_0^{B_i}} \{X_{\tau_{\mathcal{M}}^{+,nm}} \in B_j\} [1 + r_{\eta,m}(\sigma)].$$

Substituting, we find

$$|\psi_j(x) - \delta_{ij}| = \mathcal{O}(e^{-[\hat{H}_\theta(i,j) - \eta]/\sigma^2}).$$

The expression (6.33) of \hat{H}_j follows from the definition of $\hat{H}_\theta(i, j)$, and the fact that $\hat{H}(i, i) = H_0$. \square

Remark 6.12. Getting an L^1 -estimate on the difference $\langle \mu_i | - \langle \dot{\pi}_0^{B_i} |$ would require a sharper, pointwise estimate on densities, than in Corollary 6.7. Indeed, we have

$$\begin{aligned}\langle \mu_i | - \langle \dot{\pi}_0^{B_i} | &= \langle \dot{\pi}_0^{B_i} | \Pi^0 - \langle \dot{\pi}_0^{B_i} | + \sum_{n \geq 1} \langle \dot{\pi}_0^{B_i} | (\Pi_\perp^0 \Pi^\star)^n \Pi^0 \\ &= \langle \dot{\pi}_0^{B_i} | [\Pi^0 - \Pi^\star] + \sum_{j=1}^N \sum_{n \geq 1} \varepsilon_{ij}^{(n)} \langle \dot{\pi}_0^{B_j} | \Pi^0.\end{aligned}$$

Using (6.30), one can bound the L^1 -norm of the double sum by an exponentially small term. However, estimating the L^1 -norm of $\langle \dot{\pi}_0^{B_i} | [\Pi^0 - \Pi^\star]$ would require a pointwise estimate of

$$\langle \dot{\pi}_0^{B_i} | [(K^0)^{nm} - (K_m^\star)^n]$$

instead of an integral estimate as in Corollary 6.7. \diamond

6.5 Proof of the main approximation result

Let m be as in the previous section. Define a matrix P of dimension $N \times N$ with elements

$$P_{ij} = \langle \mu_i | (K_{\text{tr}}^0)^m | \psi_j \rangle. \quad (6.34)$$

Corollary 6.10 shows that

$$P_{ij} = \mathbb{P}^{\dot{\pi}_0^{B_i}} \{X_{\tau_{\mathcal{M}}^{+,m}} \in B_j\} [1 + \mathcal{O}(e^{-(\theta-\eta)/\sigma^2})].$$

Lemma 6.13. *P is a stochastic matrix for sufficiently small σ .*

PROOF: First note that $\sum_{j=1}^N |\mathbb{1}_{B_j}\rangle = |\phi_0^0\rangle$, since both are identically equal to 1. It follows that

$$\sum_{j=1}^N |\psi_j\rangle = \Pi^0 \sum_{j=1}^N |\mathbb{1}_{B_j}\rangle = \Pi^0 |\phi_0^0\rangle = |\phi_0^0\rangle,$$

and thus

$$\sum_{j=1}^N P_{ij} = \langle \mu_i | (K_{\text{tr}}^0)^m | \phi_0^0 \rangle = \langle \mu_i | \phi_0^0 \rangle = \sum_{j=1}^N \langle \mu_i | \psi_j \rangle = 1.$$

Furthermore, the P_{ij} are clearly positive if σ is small enough. \square

Let $(Y_n)_{n \geq 0}$ be the Markov chain with transition matrix P . Then Theorem 3.3 follows directly from Theorem 6.14 below. Here expectations and probabilities with respect to a signed measure are interpreted as differences of these quantities with respect to the positive and negative parts of that measure.

Theorem 6.14. *If X_n starts with the (signed) distribution μ_i , then*

$$\mathbb{E}^{\mu_i} \left[\psi_j \left(X_{\tau_{\mathcal{M}}^{+,nm}} \right) \right] = \mathbb{P}^i \{ Y_n = j \} \quad (6.35)$$

holds for all $n \geq 0$ and all $j \in \{1, \dots, N\}$. As a consequence,

$$\mathbb{P}^{\mu_i} \{ X_{\tau_{\mathcal{M}}^{+,nm}} \in B_j \} = \mathbb{P}^i \{ Y_n = j \} [1 + \mathcal{O}(e^{-[\hat{H}_j - \eta]/\sigma^2})] + \mathbb{P}^i \{ Y_n \neq j \} \mathcal{O}(e^{-[\hat{H}_j - \eta]/\sigma^2}) \quad (6.36)$$

for any $\eta > 0$, provided σ and the B_i are small enough. Furthermore, for all $x \in B_i$, one has

$$\mathbb{P}^x \{ X_{\tau_{\mathcal{M}}^{+,nm}} \in B_j \} = \mathbb{P}^i \{ Y_n = j \} + \mathcal{O}(e^{-[\hat{H}_{\min} - \eta]/\sigma^2}) + \mathcal{O}(\rho^{nm}), \quad (6.37)$$

where $\hat{H}_{\min} = \min_{\ell} \hat{H}_{\ell} \geq H_0 - (N-1)\theta$.

PROOF: The first claim (6.35) follows from (6.34) by taking the n th power of P , and using the completeness relation (6.15). The second claim (6.36) is a consequence of the decomposition

$$\mathbb{P}^i \{ Y_n = j \} = \int_{B_j} \mathbb{P}^{\mu_i} \{ X_{\tau_{\mathcal{M}}^{+,nm}} \in dy \} \psi_j(y) + \sum_{\ell \neq j} \int_{B_{\ell}} \mathbb{P}^{\mu_i} \{ X_{\tau_{\mathcal{M}}^{+,nm}} \in dy \} \psi_j(y).$$

Indeed, writing P_{ij}^n for the left-hand side and $Q_{ij}^n = \mathbb{P}^{\mu_i} \{ X_{\tau_{\mathcal{M}}^{+,nm}} \in B_j \}$, Proposition 6.11 yields

$$P_{ij}^n = \sum_{\ell=1}^N Q_{i\ell}^n [\delta_{\ell j} + r_{\ell j}],$$

where $r_{\ell j} = \mathcal{O}(e^{-[\hat{H}_j - \eta]/\sigma^2})$ for all ℓ . This is equivalent to the matrix equation $P^n = Q^n [\text{id} + R]$, which can be inverted using the Neumann series for $[\text{id} + R]^{-1}$. The resulting expression of Q^n in terms of P^n and R is equivalent to (6.36).

To obtain (6.37), we write

$$\begin{aligned} \mathbb{P}^x \{ X_{\tau_{\mathcal{M}}^{+,nm}} \in B_j \} &= \langle \delta_x | (K^0)^{nm} | \mathbb{1}_{B_j} \rangle \\ &= \langle \delta_x | (K_{\text{tr}}^0)^{nm} | \mathbb{1}_{B_j} \rangle + \langle \delta_x | (K_{\perp}^0)^{nm} | \mathbb{1}_{B_j} \rangle, \end{aligned}$$

where $K_{\perp}^0 = K^0 - K_{\text{tr}}^0$. The second term on the right-hand side decreases like the nm th power of the spectral gap ρ . As for the first term, it can be written

$$\begin{aligned} \langle \delta_x | (K_{\text{tr}}^0)^{nm} | \mathbb{1}_{B_j} \rangle &= \langle \delta_x | \Pi^0 (K_{\text{tr}}^0)^{nm} | \mathbb{1}_{B_j} \rangle \\ &= \sum_{\ell=1}^N \langle \delta_x | \psi_{\ell} \rangle \langle \mu_{\ell} | (K_{\text{tr}}^0)^{nm} | \mathbb{1}_{B_j} \rangle \\ &= \sum_{\ell=1}^N \psi_{\ell}(x) \mathbb{P}^{\mu_{\ell}} \{ X_{\tau_{\mathcal{M}}^+, nm} \in B_j \}. \end{aligned}$$

If $x \in B_i$, the term $\ell = i$ can be estimated by (6.36), while the other terms are exponentially small by Proposition 6.11. \square

A Other proofs for Section 2

A.1 Proof of Proposition 2.2

Since I is continuous at (x_i^*, x_i^*) and (x_j^*, x_j^*) and $I(x_i^*, x_j^*) = I(x_j^*, x_i^*) = 0$, we can find $\delta > 0$ such that $I(y_1, y_2) \leq \eta/6$ for all $y_1, y_2 \in B_i$, and similarly for points $z_1, z_2 \in B_j$. This implies that

$$H(i, j) - \frac{\eta}{2} \leq V(y, z) \leq H(i, j) + \frac{\eta}{2}$$

holds for all $y \in B_i$ and all $z \in B_j$. Consider now the increasing sequence of events

$$\Gamma_n = \{ \tau_{B_j}^+(x) < \tau_{B_i}^+(x), \tau_{B_j}^+(x) \leq n \} = \bigcup_{m=1}^n \left([(B_i \cup B_j)^c]^{m-1} \times B_j \times \mathcal{X}_0^{n-m} \right).$$

Then the LDP for paths (x, x_1, \dots, x_n) yields

$$-\inf_{\tilde{\Gamma}_n} I(x, \cdot) \leq \liminf_{\sigma \rightarrow 0} \sigma^2 \log \mathbb{P}^x(\Gamma_n) \leq \limsup_{\sigma \rightarrow 0} \sigma^2 \log \mathbb{P}^x(\Gamma_n) \leq -\inf_{\tilde{\Gamma}_n} I(x, \cdot).$$

Since $\mathbb{P}^x(\Gamma_n)$ is increasing in n , to prove the lower bound, it suffices to find $n \geq 2$, points $x_1, \dots, x_{n-1} \in (B_i \cup B_j)^c$ and $z \in B_j$ such that $I(x, x_1, \dots, x_{n-1}, z) \leq H(i, j) + \eta$. To this end, let $y \in B_i$ and $z \in B_j$ be the points minimizing V . Since $V(y, z) \leq H(i, j) + \eta/2$, there exist n, x_1, \dots, x_{n-1} such that $I(y, x_1, \dots, x_{n-1}, z) \leq H(i, j) + 3\eta/4$. We can assume that $x_1, \dots, x_{n-1} \notin B_i \cup B_j$ since otherwise there would exist a cheaper way to connect these sets. Replacing y by x increases I by at most $\eta/6$, yielding the required path.

To prove the upper bound, we have to show that for any n , and any path $(x_1, \dots, x_n) \in \tilde{\Gamma}_n$, $I(x, x_1, \dots, x_n) \geq H(i, j) - \eta$. This follows from the fact that $V(x, y) \geq H(i, j) - \eta$ for all $y \in B_j$, since $V(x, y)$ involves the infimum over a larger set. \square

A.2 Proof of Proposition 2.4

In the spirit of [FW98, Chapt. 6, Thm. 5.1], we first construct a path of finite length n_0 from x to \mathcal{M} , whose rate function I is bounded by $\eta/2$. In the case where the ω -limit set $\omega(x)$ is one of the stable fixed points x_i^* , there exists $n_0 \in \mathbb{N}$ such that $\Pi^{n_0}(x) \in B_i$. Setting $x_n = \Pi^n(x)$, we have $I(x, x_1, \dots, x_n) = 0$. If $\omega(x)$ is an unstable fixed point y^* , we can find $n_1 \in \mathbb{N}$ such that $\|\Pi^{n_1}(x) - y^*\| \leq \delta$. Since the stable manifolds of all unstable fixed points have codimension at least 1, they cannot contain any open subset of \mathcal{X} . Thus there exists a point y_1 at distance at

most δ from y^\star such that $\omega(y_1)$ is a stable fixed point x_i^\star . Setting $x_n = \Pi^n(x)$ and $y_n = \Pi^{n-1}(y_1)$, we obtain the existence of $n_2 \in \mathbb{N}$ such that $y_{n_2} \in B_i$ and

$$I(x, x_1, \dots, x_{n_1}, y_1, \dots, y_{n_2}) = I(x_{n_1}, y_1).$$

The continuity of I at (y^\star, y^\star) implies that we can assume $I(x_{n_1}, y_1) \leq \eta/4$ by making δ small enough. The large-deviation lower bound implies that if $n_0 = n_1 + n_2$, then

$$\liminf_{\sigma \rightarrow 0} \sigma^2 \log \mathbb{P}^x \{\tau_{\mathcal{M}}^+ \leq n_0\} \geq -\frac{\eta}{2},$$

so that there exists $\sigma_0 > 0$ such that $\mathbb{P}^x \{\tau_{\mathcal{M}}^+ \leq n_0\} \geq e^{-\eta/\sigma^2}$ holds for all $x \in \mathcal{X}$ and all $\sigma < \sigma_0$. To extend this result to an estimate on expected return times, we use the fact that for any sets $A, B, C \in \mathcal{S}_0$ such that $B \cap C = \emptyset$, one has

$$\mathbb{E}^A[\tau_B^+] \leq \mathbb{E}^A[\tau_{B \cup C}^+] + \mathbb{P}^A\{\tau_C^+ < \tau_B^+\} \mathbb{E}^C[\tau_B^+],$$

where we write $\mathbb{P}^A\{\cdot\} = \sup_{x \in A} \mathbb{P}^x\{\cdot\}$. For a proof, see for instance [BB17, Lem. 8.9] (the proof only requires $B \cap C = \emptyset$). Taking $A = \mathcal{X}$, $B = \mathcal{M}$ and $C = \mathcal{X} \setminus \mathcal{M}$, bounding $\mathbb{E}^C[\tau_B^+]$ by $\mathbb{E}^A[\tau_B^+]$, and rearranging, we obtain

$$\mathbb{E}^{\mathcal{X}}[\tau_{\mathcal{M}}^+] \leq \frac{\mathbb{E}^{\mathcal{X}}[\tau_{\mathcal{X}}^+]}{\mathbb{P}^{\mathcal{X}}\{\tau_{\mathcal{M}}^+ < \tau_{\mathcal{X} \setminus \mathcal{M}}^+\}}.$$

The same relation holds when the τ^+ are replaced by the return times $\hat{\tau}^+$ of the diluted process $(X_{nn_0})_{n \in \mathbb{N}}$. This yields

$$\mathbb{E}^{\mathcal{X}}[\tau_{\mathcal{M}}^+] \leq n_0 \mathbb{E}^{\mathcal{X}}[\hat{\tau}_{\mathcal{M}}^+] \leq \frac{n_0 \mathbb{E}^{\mathcal{X}}[\hat{\tau}_{\mathcal{X}}^+]}{\mathbb{P}^{\mathcal{X}}\{\hat{\tau}_{\mathcal{M}}^+ < \hat{\tau}_{\mathcal{X} \setminus \mathcal{M}}^+\}}. \quad (\text{A.1})$$

Observe that for any $x \in \mathcal{X}$,

$$\begin{aligned} \mathbb{P}^x\{\hat{\tau}_{\mathcal{M}}^+ < \hat{\tau}_{\mathcal{X} \setminus \mathcal{M}}^+\} &\geq \mathbb{P}^x\{1 = \hat{\tau}_{\mathcal{M}}^+ < \hat{\tau}_{\mathcal{X} \setminus \mathcal{M}}^+\} \\ &= \mathbb{P}^x\{X_{n_0} \in \mathcal{M}\} \\ &= \mathbb{P}^x\{\tau_{\mathcal{M}}^+ \leq n_0\} - \mathbb{P}^x\{\tau_{\mathcal{M}}^+ \leq n_0, X_{n_0} \notin \mathcal{M}\} \\ &\geq \mathbb{P}^x\{\tau_{\mathcal{M}}^+ \leq n_0\} \left[1 - \sup_{k \leq n_0} \mathbb{P}^{\mathcal{M}}\{X_k \notin \mathcal{M}\}\right]. \end{aligned}$$

The supremum is exponentially small by Lemma 5.1. Since n_0 is independent of σ and $\mathbb{E}^{\mathcal{X}}[\tau_{\mathcal{X}}^+]$ is uniformly bounded, the result follows. \square

B Proofs for Section 4

B.1 Proof of Proposition 4.1

An important tool in the proof is the following coupling argument.

Proposition B.1 (Coupling argument). *Let K_A be a submarkovian kernel on a set A , and denote its killing time by τ_{A^c} . Assume that there exist constants $r, \eta > 0$ such that the density k_A of K_A satisfies the Harnack inequality*

$$\sup_{x \in A: \|x - x_0\| \leq r} k_A(x, y) \leq (1 + \eta) \inf_{x \in A: \|x - x_0\| \leq r} k_A(x, y)$$

for all $x_0, y \in A$. Let $(\hat{X}_n^x)_{n \geq 0}$ be the process with kernel K_A conditioned on staying in A , defined by

$$\mathbb{P}\{\hat{X}_n^x \in B\} = \frac{K_A^n(x, B)}{K_A^n(x, A)}$$

for any Borel set $B \subset A$. Assume that for any $x_1 \neq x_2 \in A$, there exists a coupling between $(\hat{X}_n^{x_1})_{n \geq 0}$ and $(\hat{X}_n^{x_2})_{n \geq 0}$ such that the stopping time

$$N(x_1, x_2) = \inf\{n \geq 1 : \|\hat{X}_n^{x_2} - \hat{X}_n^{x_1}\| \leq r\}$$

is almost surely finite, and define

$$\rho_n = \sup_{x_1 \neq x_2 \in A} \mathbb{P}\{N(x_1, x_2) > n\}.$$

Then k_A satisfies for every $n \in \mathbb{N}$ a uniform positivity condition with parameters n and

$$L = \frac{(1 + \eta)^2 + \rho_{n-1} \sup_{y \in A} \left(\frac{\sup_{x \in A} k_A(x, y)}{\inf_{x \in A} k_A(x, y)} \right)}{\inf_{x \in A} \mathbb{P}^x\{\tau_{A^c} > n\}}. \quad (\text{B.1})$$

PROOF: See [BG14, Prop. 5.9] and [BB17, Prop. 5.4]. \square

Fix $1 \leq i \leq k \leq N$. To apply the above coupling argument, we need to estimate the various terms appearing in (B.1). We first claim that for any $\eta > 0$, there exist $C, r > 0$ such that the two relations

$$\sup_{x \in B_i} \mathcal{M} k_{\sigma, B_i}(x, y) \leq e^{C/\sigma^2} \inf_{x \in B_i} \mathcal{M} k_{\sigma, B_i}(x, y) \quad (\text{B.2})$$

$$\sup_{x \in B_i : \|x - x_0\| \leq r\sigma^2} \mathcal{M} k_{\sigma, B_i}(x, y) \leq (1 + \eta) \inf_{x \in B_i : \|x - x_0\| \leq r\sigma^2} \mathcal{M} k_{\sigma, B_i}(x, y) \quad (\text{B.3})$$

hold for all $x_0, y \in B_i$. First note that the Gaussian density (4.3) of the original kernel satisfies

$$\frac{k_{\sigma}(x_1, y)}{k_{\sigma}(x_2, y)} = \exp\left\{\frac{I(x_2, y) - I(x_1, y)}{\sigma^2}\right\},$$

where $I(x_2, y) - I(x_1, y) = \langle \Sigma^{-1}y, \Pi(x_1) - \Pi(x_2) \rangle + \frac{1}{2} \langle \Pi(x_2), \Sigma^{-1}\Pi(x_2) \rangle - \frac{1}{2} \langle \Pi(x_1), \Sigma^{-1}\Pi(x_1) \rangle$. This quantity is bounded by a constant C for all $x_1, x_2 \in B_i$ and y in a bounded set, and has order $r\sigma^2$ if in addition $\|x_1 - x_2\| \leq r\sigma^2$. Hence k_{σ} satisfies (B.2) and (B.3) if $r = r(\eta)$ is small enough.

In order to extend this to $\mathcal{M} k_{\sigma}$, we use the fact that for all $n \in \mathbb{N}$ and $x_1, y \in B_i$, we have

$$\mathbb{P}^{x_1}\{\tau_{B_i}^+ = n\} k_{\sigma}^n(x_1, y) = \int_{B_i^c} k_{\sigma}(x_1, z) \mathbb{P}^z\{\tau_{B_i}^+ = n - 1\} k_{\sigma}^{n-1}(z, y) dz.$$

The Laplace method shows that the integral is dominated by z of order 1 at most. We can thus bound $k_{\sigma}(x_1, z)$ above by $e^{C/\sigma^2} k_{\sigma}(x_2, z)$ for any $x_2 \in B_i$, and by $(1 + \eta) k_{\sigma}(x_2, z)$ if $\|x_1 - x_2\| \leq r\sigma^2$. Together with the expression (2.9) for the density of the trace process, this shows that $\mathcal{M} k_{\sigma}$ also satisfies (B.2) and (B.3).

Regarding the effect of the killing, denote by $\mathcal{M} \tau_{B_i^c}^+$ the killing time of the trace process and observe that Proposition 2.2 yields

$$\mathbb{P}^{B_i}\{\mathcal{M} \tau_{B_i^c}^+ = 1\} = \mathbb{P}^{B_i}\{\tau_{\mathcal{M} \setminus B_i}^+ < \tau_{B_i}^+\} \leq e^{-H/\sigma^2}$$

for an $H > 0$. By the Markov property, we get $\mathbb{P}^{B_i} \{ \tau_{B_i^c}^+ \leq n \} \leq n e^{-H/\sigma^2}$ for any $n \in \mathbb{N}$. Since $k_\sigma(x, y)$ is bounded below by $e^{-c\delta_0^2/\sigma^2}$ for $x, y \in B_i$, this shows that the killing has a negligible effect for δ_0 small enough. It also shows that the denominator in (B.1) is close to 1.

It thus remains to show that ρ_{n-1} in (B.1) can be made exponentially small for an n of order $\log(\sigma^{-1})$. Let $(X_n^{x_1})_{n \geq 0}$ and $(X_n^{x_2})_{n \geq 0}$ denote the original processes driven by the same noise $(\xi_n)_{n \geq 1}$, and starting respectively from x_1 and x_2 . With this coupling, writing $Y_n = X_n^{x_2} - X_n^{x_1}$, we see that

$$Y_{n+1} = \Pi(X_n^{x_1} + Y_n) - \Pi(X_n^{x_1}) = A_n Y_n + b_n(Y_n),$$

where $A_n = \partial_x \Pi(X_n^{x_1})$ and $\|b_n(y)\| \leq M\|y\|^2$ for bounded y and some constant $M > 0$. Since $\partial_x \Pi(x_i^*)$ has spectral radius strictly smaller than 1, there exists $0 < \varrho_1 < 1$ such that A_n has spectral radius bounded by ϱ_1 for any n such that $X_n^{x_1} \in B_i$. Thus there exists a norm $\|\cdot\|'$, equivalent to the Euclidean norm, such that $\|A_n y\|' \leq \varrho_2 \|y\|'$ for these n , where $\varrho_2 < 1$. Taking δ_0 small enough that $M\|x_2 - x_1\|' < 1 - \varrho_2$, we conclude that $\|Y_1\|' \leq \varrho \|x_2 - x_1\|'$.

Since $\Pi(B_i) \subset B_i$ (where the inclusion is strict for δ_0 small enough), there exists $\kappa > 0$ such that $\mathbb{P}^{B_i} \{X_1 \notin B_i\} \leq e^{-\kappa/\sigma^2}$, and thus $\mathbb{P}^{B_i} \{\exists \ell \leq n: X_\ell \notin B_i\} \leq n e^{-\kappa/\sigma^2}$ for any $n \in \mathbb{N}$. Hence the coupled trace processes conditioned on staying in B_i satisfy

$$\mathbb{P} \{ \|\hat{X}_n^{x_2} - \hat{X}_n^{x_1}\|' > \varrho^n \|x_2 - x_1\|' \} \leq \frac{2n e^{-\kappa/\sigma^2}}{1 - 2n e^{-\kappa/\sigma^2}} \leq 3n e^{-\kappa/\sigma^2} \quad (\text{B.4})$$

for σ small enough. Let $N(x_1, x_2) = \inf\{n \geq 1: \|\hat{X}_n^{x_2} - \hat{X}_n^{x_1}\|' \leq r(\eta)\sigma^2\}$, and let $n_1(\sigma)$ be such that $\text{diam}(B_i)\varrho^{n_1(\sigma)} \leq r(\eta)\sigma^2$. Note that $n_1(\sigma)$ has order $\log(\sigma^{-1})$, and that $\mathbb{P}\{N(x_1, x_2) > n_1(\sigma)\}$ is bounded above in (B.4). Applying the Markov property at times which are multiples of $n_1(\sigma)$, we obtain

$$\rho_{\ell n_1(\sigma)} = \mathbb{P}\{N(x_1, x_2) > \ell n_1(\sigma)\} \leq (3n_1(\sigma) e^{-\kappa/\sigma^2})^\ell.$$

Choosing ℓ such that $\ell\kappa > C$, the result follows with $n_0(\sigma) = \ell n_1(\sigma)$, taking η small enough. \square

B.2 Proof of Proposition 4.2

The proof is essentially an adaptation to the discrete-time setting of results in [BB17, Sect. 8.2 and 8.3], which rely in part on methods from [BG06, Chapt. 5]. Since several proofs simplify in the present setting, we believe it is worth giving details here. In view of the upper bound (A.1) for $\mathbb{E}^{\mathcal{X}}[\tau_{\mathcal{M}}^+]$ and the fact that $\mathbb{E}^{\mathcal{X}}[\tau_{\mathcal{X}}^+]$ is bounded by Assumption REC, it is sufficient to show that $\mathbb{P}^{\mathcal{X}}\{\tau_{\mathcal{M}}^+ \leq n_0\}$ is bounded away from 1 for some n_0 of order $\log(\sigma^{-1})$.

Given $x \in \mathcal{X}$, we first give an estimate for the probability of the sample path $(X_n)_{n \geq 0}$ starting in x deviating from the deterministic orbit $(X_n^{\text{det}})_{n \geq 0}$ defined by $X_n^{\text{det}} = \Pi^n(x)$.

Lemma B.2. *There exist constants $\kappa > 0$ and $h_0 > 0$ such that for any $n \in \mathbb{N}$,*

$$\mathbb{P}^x \left\{ \max_{1 \leq i \leq n} \|X_i - X_i^{\text{det}}\| > h \right\} \leq n \exp \left\{ -\frac{\kappa h^2}{\sigma^2 G_n^{(2)}} \right\}$$

holds whenever $0 < h < h_0/G_n^{(1)}$, where $G_n^{(\ell)} = \sum_{i=0}^{n-\ell} D_n^{\ell i}$ with $D_n = \max_{1 \leq i \leq n} \|\partial_x \Pi(X_i^{\text{det}})\|$.

PROOF: The difference $Y_n = X_n - X_n^{\text{det}}$ satisfies $Y_0 = 0$ and

$$\begin{aligned} Y_{n+1} &= \Pi(X_n^{\text{det}} + Y_n) - \Pi(X_n^{\text{det}}) + \sigma \xi_{n+1} \\ &= A_n Y_n + b(Y_n) + \sigma \xi_{n+1}, \end{aligned}$$

where we have set $A_n = \partial_x \Pi(X_n^{\det})$, and $\|b(y)\| \leq M\|y\|^2$ for bounded y and some $M > 0$. Consider first the linearized equation

$$Y_{n+1}^0 = A_n Y_n^0 + \sigma \xi_{n+1}, \quad Y_0^0 = 0.$$

Its solution can be written

$$Y_n^0 = \sigma \sum_{i=1}^n B_{ni} \xi_i,$$

where $B_{ni} = A_{n-1} \dots A_i$ if $i < n$ and $B_{nn} = \text{id}$ is the identity matrix. Thus Y_n^0 is a centred Gaussian random variable with covariance matrix $\sigma^2 \Sigma_n$, where

$$\Sigma_n = \sum_{i=1}^n B_{ni} \Sigma B_{ni}^\dagger.$$

We have $\|B_{ni}\| \leq D_n^{n-i}$ and $\|\Sigma_n\| \leq \|\Sigma\| G_n^{(2)}$. This implies that $\mathbb{P}\{\|Y_n^0\| > h\} \leq e^{-\kappa_0 h^2 / (G_n^{(2)} \sigma^2)}$ for some $\kappa_0 > 0$, and thus

$$\mathbb{P}\left\{\max_{1 \leq i \leq n} \|Y_i^0\| > h\right\} = \mathbb{P}\left\{\bigcup_{i=1}^n \{\|Y_i^0\| > h\}\right\} \leq n \exp\left\{-\frac{\kappa_0 h^2}{G_n^{(2)} \sigma^2}\right\}. \quad (\text{B.5})$$

To extend this estimate to Y_n , we write $Y_n = Y_n^0 + R_n$ and note that we have

$$R_{n+1} = A_n R_n + b(Y_n) \quad \Rightarrow \quad R_n = \sum_{i=2}^n B_{ni} b(Y_{i-1}).$$

Setting $\tau = \inf\{n \geq 1 : \|Y_n\| > h\}$, we have $\|b(Y_{n \wedge \tau})\| \leq Mh^2$ and $\|R_{n \wedge \tau}\| \leq G_n^{(1)} Mh^2$. For any decomposition $h = H_0 + H_1$ with $H_0, H_1 > 0$, we have

$$\mathbb{P}\{\tau < n\} \leq \mathbb{P}\left\{\max_{1 \leq i \leq n} \|Y_i^0\| > H_0\right\} + \mathbb{P}\left\{\max_{1 \leq i \leq n \wedge \tau} \|R_i\| > H_1\right\}.$$

The first term on the right-hand side can be estimated with (B.5), while the second one vanishes if $H_1 \geq G_n^{(1)} Mh^2$. The result thus follows by setting $H_0 = h(1 - G_n^{(1)} Mh)$. \square

Corollary B.3. *Let $x \in \mathcal{X}$ be such that $\omega(x)$ is a stable fixed point x_i^* . Then there exist $n(x) < \infty$ and $\kappa(x) > 0$ such that $\mathbb{P}^x\{\tau_{\mathcal{M}}^+ \geq n(x)\} \leq n(x) e^{-\kappa(x)/\sigma^2}$.*

PROOF: By definition of ω -limit sets, there exists $n(x)$ such that $\|\Pi^{n(x)}(x) - x_i^*\| < \delta/2$. It is thus sufficient to apply Lemma B.2 with $h = \delta/2$. \square

This bound deteriorates when x approaches an unstable fixed point of Π (or the stable manifold of such a fixed point), because $n(x)$ diverges and $\kappa(x)$ tends to 0. We thus have to treat these cases separately. In doing so, we will repeatedly use the following elementary estimate.

Lemma B.4. *Let A, B be two disjoint sets in \mathcal{S}_0 . Then for any $n_1, n_2 \in \mathbb{N}$,*

$$\mathbb{P}^{A \cup B}\{\tau_{(A \cup B)^c}^+ \geq n_1 + n_2\} \leq \mathbb{P}^A\{\tau_{A^c}^+ \geq n_1\} + \mathbb{P}^B\{\tau_{B^c}^+ \geq n_2\} + \mathbb{P}^B\{\tau_A^+ < \tau_{(A \cup B)^c}^+\}.$$

PROOF: When starting in B , consider separately the cases $\tau_{(A \cup B)^c}^+ = \tau_{A^c}^+$ and $\tau_{(A \cup B)^c}^+ = \tau_{B^c}^+$. When starting in A , distinguish the cases $\tau_{A^c}^+ \geq n_1$ and $\tau_{A^c}^+ < n_1$, and use the bound for starting points in B . \square

To estimate exit probabilities from the neighbourhood of an unstable equilibrium point, we proceed in two steps, considering first the exit from a small neighbourhood of size $\sigma^{3/4}$, and then the exit from a larger neighbourhood of size δ .

Lemma B.5. *Let \mathcal{S} be a neighbourhood of diameter $\sigma^{3/4}$ of an unstable fixed point z_j^* . Then there exist constants $c_1, C_1 > 0$ such that*

$$\mathbb{P}^{\mathcal{S}}\{\tau_{\mathcal{S}^c}^+ > c_1 \log(\sigma^{-1})\} \leq C_1 \log(\sigma^{-1}) \sigma^{1/2}.$$

PROOF: We may assume that $z_j^* = 0$. Let λ_+ be the module of the largest eigenvalue of $\partial_x \Pi(z^*)$, and let m be its multiplicity. There exists a linear change of variables $X_n \mapsto (Y_n, Z_n)$ such that

$$\begin{aligned} Y_{n+1} &= A_+ Y_n + b_+(Y_n, Z_n) + \sigma \xi_{n+1}^+, & Y_0 &= y, \\ Z_{n+1} &= A_- Z_n + b_-(Y_n, Z_n) + \sigma \xi_{n+1}^-, & Z_0 &= z, \end{aligned} \quad (\text{B.6})$$

where A_+ is a square matrix of size m , all of whose eigenvalues are equal to λ_+ , all eigenvalues of A_- are strictly smaller in module than λ_+ , $\|b_{\pm}(y, z)\| \leq M(\|y\|^2 + \|z\|^2)$ for bounded y and z , and the ξ_n^{\pm} are nondegenerate Gaussian random variables. Let Y_n^0 obey the linearized dynamics $Y_{n+1}^0 = A_+ Y_n^0 + \sigma \xi_{n+1}^+$. Similarly to the Lemma B.2, we have $Y_n = Y_n^0 + R_n$, where

$$Y_n^0 = A_+^n y + \sigma \sum_{i=1}^n A_+^{n-i} \xi_i^+, \quad R_n = \sum_{i=2}^n A_+^{n-i} b_+(Y_{i-1}, Z_{i-1}).$$

A similar decomposition $Z_n = Z_n^0 + Q_n$ holds for the second component. Let Σ_+ denote the covariance matrix of the ξ_i^+ . The law of Y_n^0 is Gaussian with covariance matrix

$$\text{Cov}(Y_n^0) = \sigma^2 \sum_{i=1}^n A_+^{n-i} \Sigma_+ (A_+^{\dagger})^{n-i}$$

We have $\det \text{Cov}(Y_n^0) \geq c(\sigma \lambda_+^n)^{2m}$ for some $c > 0$, so that there exists $C > 0$ such that

$$\mathbb{P}\{\|Y_n^0\| < h\} \leq C \left(\frac{h}{\sigma \lambda_+^n} \right)^m$$

for all $n \in \mathbb{N}$ and $h > 0$. Setting $X_i^0 = (Y_i^0, Z_i^0)$ we have for any $h_1 > 0$

$$\begin{aligned} \mathbb{P}\left\{\max_{1 \leq i \leq n} \|X_i\| < h\right\} &\leq \mathbb{P}\left\{\max_{1 \leq i \leq n} \|X_i^0\| < h + h_1\right\} + \mathbb{P}\left\{\max_{1 \leq i \leq n} \|(R_i, Q_i)\| \geq h_1, \max_{1 \leq i \leq n} \|X_i\| < h\right\}, \\ &\leq \mathbb{P}\left\{\|Y_n^0\| < h + h_1\right\} + \mathbb{P}\left\{\max_{1 \leq i \leq n} \|(R_i, Q_i)\| \geq h_1, \max_{1 \leq i \leq n} \|X_i\| < h\right\}. \end{aligned}$$

The second term on the right-hand side vanishes if we set $h_1 = C_1 M h^2 \lambda_+^n$ for a sufficiently large constant C_1 . Choosing $h = \sigma^{3/4}$ and n such that $\lambda_+^n \geq \sigma^{-3/4}$, we obtain the result. \square

Lemma B.6. *Let \mathcal{U} be a neighbourhood of diameter δ of an unstable fixed point z_j^* . Then there exist constants $c_2, C_2 > 0$ such that*

$$\mathbb{P}^{\mathcal{U}}\{\tau_{\mathcal{U}^c}^+ > c_2 \log(\sigma^{-1})\} \leq C_2 \log(\sigma^{-1}) \sigma^{1/2}.$$

PROOF: We may use a similar coordinate system as in (B.6), except that now Y contains all unstable directions, while Z contains the marginally stable and stable ones. The center-stable manifold theorem allows us to assume that $b_+(0, z) = 0$ and $\|b_+(y, z)\| \leq M(\|y\|^2 + \|y\|\|z\|)$ in \mathcal{U} for some $M > 0$. The Lyapunov function $U_n = \|Y_n\|^2$ satisfies

$$U_{n+1} = \|A_+ Y_n\|^2 + [2\langle b_+, A_+ Y_n \rangle + \|b_+\|^2] + 2\sigma\langle A_+ Y_n + b_+, \xi_{n+1}^+ \rangle + \sigma^2 \|\xi_{n+1}^+\|^2.$$

All eigenvalues of A_+ have a module strictly larger than 1, showing that $\|A_+ Y_n\|^2 \geq \lambda_+ U_n$ for some $\lambda_+ > 1$. The term in square brackets has order $U_n^{3/2} + U_n \|Z_n\|$. Thus for small enough δ , there exists $\tilde{\lambda}_+ > 1$ such that

$$U_{n+1} \geq \tilde{\lambda}_+ U_n + \sigma g(X_n) \eta_{n+1} + \sigma^2 \|\xi_{n+1}^+\|^2,$$

where $\|g(x)\| \leq M_1 U_n^{1/2}$ for some $M_1 > 0$, and η_{n+1} is a centred Gaussian random variable of bounded variance. Let $\mathcal{K} = \{(y, z) \in \mathcal{U} : \sigma^{3/4} \leq \|y\| \leq \delta\}$. For $n \leq \tau_{\mathcal{K}^c}^+$, we obtain that $V_n = U_n^{1/2}$ satisfies

$$V_{n+1} \geq \tilde{\lambda}_+^{1/2} V_n + \sigma \tilde{g}(X_n) \eta_{n+1},$$

where \tilde{g} is bounded in \mathcal{K} . It follows that for $n \leq \tau_{\mathcal{K}^c}^+$,

$$V_n \geq \tilde{\lambda}_+^{n/2} [V_0 + \sigma \zeta_n],$$

where $V_0 \geq \sigma^{3/4}$ and $\zeta_n = \sum_{i=1}^n \tilde{\lambda}_+^{(n-i)/2} \tilde{g}(X_{i-1}) \eta_i$ has bounded variance. Chebyshev's inequality shows that

$$\mathbb{P} \left\{ \min_{1 \leq i \leq n \wedge \tau_{\mathcal{K}^c}^+} \frac{\zeta_i}{\tilde{\lambda}_+^{i/2} - 1} < -\sigma^{3/4} \right\} \leq n C \sigma^{1/2}$$

for some $C > 0$. Taking n of order $\log(\sigma^{-1})$ such that $\tilde{\lambda}_+^{i/2} > \delta/\sigma^{3/4}$, this yields the existence of constants $c_1, C_1 > 0$ such that

$$\begin{aligned} \mathbb{P}^{\mathcal{K}} \{ \tau_{\mathcal{K}^c}^+ > c_1 \log(\sigma^{-1}) \} &\leq C_1 \log(\sigma^{-1}) \sigma^{1/2}, \\ \mathbb{P}^{\mathcal{K}} \{ \tau_{\mathcal{K}^c}^+ < \tau_{\mathcal{U}^c}^+ \} &\leq C_1 \log(\sigma^{-1}) \sigma^{1/2}. \end{aligned} \quad (\text{B.7})$$

Applying Lemma B.4 with $A = \mathcal{S}$ and $B = \mathcal{K}$ yields the claimed result with $\mathcal{K} \cup \mathcal{S}$ instead of \mathcal{U} . The result can be extended to \mathcal{U} by showing that Lemma B.5 also applies to the larger set $\mathcal{U} \setminus \mathcal{K} = \{(y, z) \in \mathcal{U} : \|y\| \leq \sigma^{3/4}\}$, using the better bounds on b_+ due to the centre-stable manifold theorem, and analysing a slightly more general recursion for Y_n with time-dependent linear part. \square

To finish the proof, we have to deal with the possible existence of heteroclinic orbits. Denote the unstable fixed points by z_1^*, \dots, z_M^* . Let \mathcal{U}_i be a ball of diameter δ centred in z_i^* , with δ small enough for Lemma B.6 to apply. We denote the union of all \mathcal{U}_i by \mathcal{U} . Define

$$\begin{aligned} \tau_A^{\det}(x) &= \inf \{ n \geq 1 : \Pi^n(x) \in A \}, \\ \mathcal{A}_i &= \{ x \in \mathcal{X} \setminus \mathcal{U} : X_{\tau_{\mathcal{U}}^{\det}} \in \mathcal{U}_i \}. \end{aligned}$$

The set \mathcal{A}_i contains part of the stable manifold of z_i^* . Note that \mathcal{A}_i contains no fixed points of Π , showing, by Lemma B.2, that $\mathbb{P}^{\mathcal{A}_i} \{ \tau_{\mathcal{A}_i^c}^+ \geq n \}$ is exponentially small for some bounded n . Furthermore, the proof of Lemma B.6, in particular (B.7), shows that

$$\mathbb{P}^{\mathcal{U}_i} \{ \tau_{\mathcal{A}_i}^+ < \tau_{(\mathcal{A}_i \cup \mathcal{U}_i)^c}^+ \} \leq C_1 \log(\sigma^{-1}) \sigma^{1/2}.$$

This is because deterministic orbits starting on the boundary of \mathcal{K} on which $Y = \delta$ cannot enter \mathcal{A}_i (recall that there are no heteroclinic cycles). Thus Lemma B.4 shows that there exist constants $c_3, C_3 > 0$ such that

$$\mathbb{P}^{\mathcal{U}_i \cup \mathcal{A}_i} \{ \tau_{(\mathcal{U}_i \cup \mathcal{A}_i)^c}^+ > c_3 \log(\sigma^{-1}) \} \leq C_3 \log(\sigma^{-1}) \sigma^{1/2}.$$

When leaving \mathcal{U}_i , it may happen that a trajectory enters the domain of attraction \mathcal{A}_j of another unstable fixed point, due to the existence of a heteroclinic orbit from z_i^* to z_j^* . In this case we write $i < j$. Extending this relation by transitivity yields a strict partial order relation, owing to the fact that there are no heteroclinic cycles. Lemma B.2 implies that whenever $i < j$ or i and j are not related, the probability, when starting from $\mathcal{U}_j \cup \mathcal{A}_j$, to hit $\mathcal{U}_i \cup \mathcal{A}_i$ before $(\mathcal{U}_j \cup \mathcal{A}_j)^c$ is exponentially small. Repeated application of Lemma B.4 shows that, if $\hat{\mathcal{U}} = \bigcup_{i=1}^M (\mathcal{U}_i \cup \mathcal{A}_i)$, then

$$\mathbb{P}^{\hat{\mathcal{U}}} \{ \tau_{\hat{\mathcal{U}}^c}^+ > c_4 \log(\sigma^{-1}) \} \leq C_4 \log(\sigma^{-1}) \sigma^{1/2}$$

holds for constants $c_4, C_4 > 0$. In $\mathcal{X} \setminus \hat{\mathcal{U}}$, we can apply Corollary B.3 with a uniformly bounded $n(x)$, which finishes the proof, applying one last time Lemma B.4. \square

References

- [AL99] Konstantin E. Avrachenkov and Jean B. Lasserre, *The fundamental matrix of singularly perturbed Markov chains*, Adv. in Appl. Probab. **31** (1999), no. 3, 679–697. MR 1742689 (2000m:60087)
- [Bak11] Yuri Bakhtin, *Noisy heteroclinic networks*, Probab. Theory Relat. Fields **150** (2011), no. 1-2, 1–42 (English).
- [BAKS84] Gérard Ben Arous, Shigeo Kusuoka, and Daniel W. Stroock, *The Poisson kernel for certain degenerate elliptic operators*, J. Funct. Anal. **56** (1984), no. 2, 171–209. MR 738578
- [BB17] Manon Baudel and Nils Berglund, *Spectral Theory for Random Poincaré Maps*, SIAM J. Math. Anal. **49** (2017), no. 6, 4319–4375. MR 3719020
- [BEGK04] Anton Bovier, Michael Eckhoff, Véronique Gayrard, and Markus Klein, *Metastability in reversible diffusion processes. I. Sharp asymptotics for capacities and exit times*, J. Eur. Math. Soc. (JEMS) **6** (2004), no. 4, 399–424. MR 2094397
- [BG06] Nils Berglund and Barbara Gentz, *Noise-induced phenomena in slow-fast dynamical systems. A sample-paths approach*, Probability and its Applications (New York), Springer-Verlag London, Ltd., London, 2006. MR 2197663
- [BG14] ———, *On the noise-induced passage through an unstable periodic orbit II: General case*, SIAM J. Math. Anal. **46** (2014), no. 1, 310–352. MR 3151386
- [BG16] Alessandra Bianchi and Alexandre Gaudillière, *Metastable states, quasi-stationary distributions and soft measures*, Stochastic Process. Appl. **126** (2016), no. 6, 1622–1680. MR 3483732
- [BGK05] Anton Bovier, Véronique Gayrard, and Markus Klein, *Metastability in reversible diffusion processes. II. Precise asymptotics for small eigenvalues*, J. Eur. Math. Soc. (JEMS) **7** (2005), no. 1, 69–99. MR 2120991
- [BGK15] Nils Berglund, Barbara Gentz, and Christian Kuehn, *From random Poincaré maps to stochastic mixed-mode-oscillation patterns*, J. Dynam. Differential Equations **27** (2015), no. 1, 83–136. MR 3317393
- [Bir57] Garrett Birkhoff, *Extensions of Jentzsch's theorem*, Trans. Amer. Math. Soc. **85** (1957), 219–227.

- [BL10] J. Beltrán and C. Landim, *Tunneling and metastability of continuous time Markov chains*, J. Stat. Phys. **140** (2010), no. 6, 1065–1114 (English).
- [BL12] Nils Berglund and Damien Landon, *Mixed-mode oscillations and interspike interval statistics in the stochastic FitzHugh–Nagumo model*, Nonlinearity **25** (2012), no. 8, 2303–2335.
- [BL15] J. Beltrán and C. Landim, *A martingale approach to metastability*, Probab. Theory Relat. Fields **161** (2015), no. 1-2, 267–307 (English).
- [BLR16] Volker Betz and Stéphane Le Roux, *Multi-scale metastable dynamics and the asymptotic stationary distribution of perturbed markov chains*, Stochastic Processes and their Applications (2016).
- [CFLM06] R. Coutinho, B. Fernandez, R. Lima, and A. Meyroneinc, *Discrete time piecewise affine models of genetic regulatory networks*, J. Math. Biol. **52** (2006), no. 4, 524–570. MR 2235517
- [CMSM13] Pierre Collet, Servet Martínez, and Jaime San Martín, *Quasi-stationary distributions*, Probability and its Applications (New York), Springer, Heidelberg, 2013. MR 2986807
- [CV16] Nicolas Champagnat and Denis Villemonais, *Exponential convergence to quasi-stationary distribution and Q-process*, Probab. Theory Related Fields **164** (2016), no. 1-2, 243–283. MR 3449390
- [CV23] Nicolas Champagnat and Denis Villemonais, *General criteria for the study of quasi-stationarity*, Electronic Journal of Probability **28** (2023), no. none, 1 – 84.
- [DGLLPN16] Giacomo Di Gesù, Tony Lelièvre, Dorian Le Peutrec, and Boris Nectoux, *Jump Markov models and transition state theory: the quasi-stationary distribution approach*, Faraday Discussion **195** (2016), 469–495.
- [DGLLPN19] ———, *Sharp asymptotics of the first exit point density*, Ann. PDE **5** (2019), no. 1, 174 (English), Id/No 5.
- [DGLLPN20] ———, *The exit from a metastable state: concentration of the exit point distribution on the low energy saddle points. I*, J. Math. Pures Appl. (9) **138** (2020), 242–306 (English).
- [FK17] M. Freidlin and L. Koralov, *Metastable distributions of Markov chains with rare transitions*, J. Stat. Phys. **167** (2017), no. 6, 1355–1375 (English).
- [Fre03] Ivar Fredholm, *Sur une classe d'équations fonctionnelles*, Acta Math. **27** (1903), no. 1, 365–390. MR 1554993
- [FW98] M. I. Freidlin and A. D. Wentzell, *Random perturbations of dynamical systems*, second ed., Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences], vol. 260, Springer-Verlag, New York, 1998, Translated from the 1979 Russian original by Joseph Szücs. MR 1652127
- [GG01] Israel Gohberg and Seymour Goldberg, *Basic operator theory*, Birkhäuser Boston, Inc., Boston, MA, 2001, Reprint of the 1981 original. MR 1843182
- [GGK03] Israel Gohberg, Seymour Goldberg, and Marinus A. Kaashoek, *Basic classes of linear operators*, Birkhäuser Verlag, Basel, 2003. MR 2015498
- [HH92] Refael Hassin and Moshe Haviv, *Mean passage times and nearly uncoupled Markov chains*, SIAM J. Discrete Math. **5** (1992), no. 3, 386–397. MR 1172747 (93m:60146)
- [HM09] Pawel Hitczenko and Georgi S. Medvedev, *Bursting oscillations induced by small noise*, SIAM J. Appl. Math. **69** (2009), no. 5, 1359–1392. MR 2487064
- [HM11] Martin Hairer and Jonathan C. Mattingly, *Yet another look at Harris' ergodic theorem for Markov chains*, Seminar on Stochastic Analysis, Random Fields and Applications VI, Progr. Probab., vol. 63, Birkhäuser/Springer Basel AG, Basel, 2011, pp. 109–117. MR 2857021
- [HM13] Pawel Hitczenko and Georgi S. Medvedev, *The Poincaré map of randomly perturbed periodic motion*, J. Nonlinear Sci. **23** (2013), no. 5, 835–861. MR 3101836

- [Hou98] Shui-Hung Hou, *A simple proof of the Leverrier-Faddeev characteristic polynomial algorithm*, SIAM Rev. **40** (1998), no. 3, 706–709. MR 1642776
- [Jen12] Robert Jentzsch, *Über Integralgleichungen mit positivem Kern*, J. f. d. reine und angew. Math. **141** (1912), 235–244.
- [Lan19] Claudio Landim, *Metastable Markov chains*, Probab. Surv. **16** (2019), 143–227 (English).
- [LLM18] Claudio Landim, Michail Loulakis, and Mustapha Mourragui, *Metastable Markov chains: from the convergence of the trace to the convergence of the finite-dimensional distributions*, Electron. J. Probab. **23** (2018), 34 (English), Id/No 95.
- [LLPN22] Tony Lelièvre, Dorian Le Peutrec, and Boris Nectoux, *The exit from a metastable state: concentration of the exit point distribution on the low energy saddle points. II*, Stoch. Partial Differ. Equ., Anal. Comput. **10** (2022), no. 1, 317–357 (English).
- [LMS19] Claudio Landim, Mauro Mariani, and Insuk Seo, *Dirichlet's and Thomson's principles for non-selfadjoint elliptic operators with application to non-reversible metastable diffusion processes*, Arch. Ration. Mech. Anal. **231** (2019), no. 2, 887–938. MR 3900816
- [LPM20] Dorian Le Peutrec and Laurent Michel, *Sharp spectral asymptotics for nonreversible metastable diffusion processes*, Probab. Math. Phys. **1** (2020), no. 1, 3–53 (English).
- [LS18] Claudio Landim and Insuk Seo, *Metastability of nonreversible random walks in a potential field and the Eyring-Kramers transition rate formula*, Commun. Pure Appl. Math. **71** (2018), no. 2, 203–266 (English).
- [LS22] Jungkyoung Lee and Insuk Seo, *Non-reversible metastable diffusions with Gibbs invariant measure. II: Markov chain convergence*, J. Stat. Phys. **189** (2022), no. 2, 34 (English), Id/No 25.
- [MOS89] Fabio Martinelli, Enzo Olivieri, and Elisabetta Scoppola, *Small random perturbations of finite- and infinite-dimensional dynamical systems: unpredictability of exit times*, J. Statist. Phys. **55** (1989), no. 3-4, 477–504. MR 1003525 (91f:60105)
- [MT92] Sean P. Meyn and R. L. Tweedie, *Stability of Markovian processes. I. Criteria for discrete-time chains*, Adv. in Appl. Probab. **24** (1992), no. 3, 542–574. MR 1174380
- [MT93] ———, *Stability of Markovian processes. III. Foster-Lyapunov criteria for continuous-time processes*, Adv. in Appl. Probab. **25** (1993), no. 3, 518–548. MR 1234295
- [Num84] Esa Nummelin, *General irreducible Markov chains and nonnegative operators*, Cambridge Tracts in Mathematics, vol. 83, Cambridge University Press, Cambridge, 1984.
- [RS80] Michael Reed and Barry Simon, *Methods of modern mathematical physics.*, second ed., vol. I. Functional analysis, Academic Press, Inc. [Harcourt Brace Jovanovich, Publishers], New York, 1980. MR 751959
- [RS18] Fraydoun Rezakhanlou and Insuk Seo, *Scaling limit of small random perturbation of dynamical systems*, Preprint arXiv/1812.02069, 2018.
- [Sch68] Paul J. Schweitzer, *Perturbation theory and finite Markov chains*, J. Appl. Probability **5** (1968), 401–413. MR 0234527 (38 #2844)
- [Seo20] Insuk Seo, *Analysis of metastable behavior via solutions of Poisson equations*, RIMS Kôkyûroku Bessatsu **B79** (2020), 1–17 (English).
- [Vot07] A.F. Voter, *Radiation effects in solids*, ch. Introduction to the kinetic Monte Carlo method, pp. 1–23, Dordrecht: Springer, 2007.
- [WK90] Jeffrey B. Weiss and Edgar Knobloch, *A stochastic return map for stochastic differential equations*, J. Stat. Phys. **58** (1990), no. 5-6, 863–883 (English).
- [YZ05] G.G. Yin and Q. Zhang, *Discrete-time Markov chains: Two-time-scale methods and applications*, Springer, 2005.

Contents

1	Introduction and informal statement of results	1
2	Set-up and assumptions	5
2.1	Singular deterministic limit	6
2.2	Large-deviation principle	6
2.3	Positive Harris recurrence	7
2.4	Trace process	9
2.5	Killed process and QSDs	10
2.6	Uniform positivity	11
3	Main results	11
4	Applications	13
4.1	Iterated maps with additive noise	13
4.2	Random Poincaré maps	15
5	Proof of Proposition 3.1	17
5.1	Feynman–Kac representation formulas	17
5.2	Norm estimates on kernels	18
5.3	Resolvent estimate	20
6	Proof of Theorem 3.3	21
6.1	Large-deviation estimates for the trace process on \mathcal{M}	21
6.2	The finite rank kernel K^\star	23
6.3	The truncated kernel K_{tr}^0	25
6.4	Comparison of transition probabilities	26
6.5	Proof of the main approximation result	33
A	Other proofs for Section 2	35
A.1	Proof of Proposition 2.2	35
A.2	Proof of Proposition 2.4	35
B	Proofs for Section 4	36
B.1	Proof of Proposition 4.1	36
B.2	Proof of Proposition 4.2	38

Institut Denis Poisson (IDP)
 Université d'Orléans, Université de Tours, CNRS – UMR 7013
 Bâtiment de Mathématiques, B.P. 6759
 45067 Orléans Cedex 2, France
E-mail address: nils.berglund@univ-orleans.fr