The Long-Run Behavior of Regularized Learning in Games under Uncertainty

Waïss Azizian[1,2,3,*]
waiss.azizian@morganstanley.com

Pierre-Louis Cauvin[2,3,*]
pierre-louis.cauvin@univ-grenoble-alpes.fr

Franck Iutzeler[4]
franck.iutzeler@math.univ-toulouse.fr

Jérôme Malick[3]
jerome.malick@cnrs.fr

Panayotis Mertikopoulos[2,5]
panayotis.mertikopoulos@imag.fr

[1] Machine Learning Research, Morgan Stanley
[2] Univ. Grenoble Alpes, CNRS, Inria, Grenoble INP, LIG, Grenoble, France
[3] Univ. Grenoble Alpes, CNRS, Inria, Grenoble INP, LJK, Grenoble, France
[4] Univ. Toulouse, INSA Toulouse, CNRS, IMT, Toulouse, France
[5] Archimedes, Athena Research Center

Abstract

In this paper, we investigate the long-run behavior of discounted regularized learning with stochastic gradient feedback in general, non-concave games. Specifically, we focus on a family of implicitly regularized exponential / multiplicative weight update schemes, and we seek to determine which actions—or, more generally, which recurrent patterns of play—are more likely to arise in the long run. We approach this question through the lens of large deviations theory and randomly perturbed dynamical systems, and we obtain a precise characterization of the distribution of the process: in the long run, it follows a Boltzmann-Gibbs law with temperature equal to the method's learning rate, and energy levels determined by the game and the statistics of the noise. Concretely, we show that the distribution of play concentrates exponentially around the dynamics' *internally chain transitive* (ICT) sets—i.e., irreducible invariant sets containing no smaller attractors—and the probability of visiting such a set depends exponentially on its energy. As a result, unstable ICT sets are exponentially less likely to be visited than stable ones, and the sequence of play is exponentially concentrated around the problem's "ground state", where energy is minimized. In this manner, learning acts as a selection mechanism: with exponentially high probability, the ground state is the *only* outcome observed in the long run, even in the presence of multiple equilibria and other attractors.

1 Introduction

What is the long-run behavior of learning in games?
Which patterns of play are more likely to be observed, and by how much?

These questions have been a mainstay of non-cooperative game theory since the early days of the field, with Nash [88, p. 21] famously noting that "it is unnecessary to assume that [the players of

*Equal contribution.

40th Conference on Neural Information Processing Systems (NeurIPS 2026).

the game] have full knowledge of its structure or the ability and inclination to go through a complex reasoning process to calculate an equilibrium. But the participants are still supposed to adapt by accumulating empirical information on the relative advantages of the various pure strategies at their disposal". In this way, the hope was that learning would lead to justification of equilibrium play through repeated interactions: self-interested, myopic agents would eventually learn to behave "as if" they were fully rational in the first place.

Albeit plausible, this expectation is ultimately misleading. A first (dynamical) obstruction was discovered by Hart & Mas-Colell [48, 49], who showed that uncoupled dynamics cannot converge to Nash equilibrium in all games. Then, a few years later, the seminal work of Daskalakis et al. [28] showed that the problem of computing a Nash equilibrium is PPAD-complete—so a learning-based justification of Nash equilibrium is likewise beyond reach. As a result, to recover a partially positive answer, the focus shifted to more tractable solution concepts (like coarse correlated equilibria), more relaxed modes of convergence (e.g., in terms of time-averages), and/or learning in specific classes of games (such as potential, zero-sum, or monotone games).

Our paper takes a complementary viewpoint. Instead of targeting a specific notion of equilibrium (Nash, correlated, or other), we take the players' learning dynamics as the object of interest, and we aim to characterize their long-run behavior—whether this is an equilibrium, cycle, or other recurrent pattern of play. Specifically, we seek to provide concrete answers to the following questions:

Do the players' dynamics settle around certain outcomes? If so, which ones?

These questions are largely motivated by the growing use of game-theoretic models in machine learning and data science—from multi-agent reinforcement learning to generative adversarial networks (GANs) and online ad auctions. In this context, the system's training dynamics are fairly simple to describe, but equilibrium convergence guarantees are typically absent. This is precisely our objective: to characterize the limit behavior of learning in general, non-concave games, and to identify which recurrent patterns of play are more likely to arise in the long run.

Our contributions. More concretely, our goal is to quantify the long-run distribution of play under regularized learning in the most comprehensive manner possible. For this reason, we consider a broad class of continuous games with smooth, non-concave payoff functions defined over (products of) simplices—a class containing, among others, all finite games, Markov / stochastic games, box-constrained GANs, etc.—and we focus on a family of discounted exponential / multiplicative weights (EW) update schemes run with stochastic gradient feedback and a constant learning rate.

Our driving observation is that, when run with a constant learning rate, the induced trajectory of play does not converge—it *concentrates*. In particular, players spend long intervals of time near certain "metastable components" interleaved with fast transitions between them. These transitions usually follow the mean, continuous-time flow of the dynamics, *but not always:* with positive probability, the process may go *against the flow*, returning infinitely often to previously visited regions.

In light of this, we first focus on identifying these "metastable components". To that end, we adapt and extend the continuous-time theory of large deviations and randomly perturbed dynamical systems of Freidlin & Wentzell [36] and Kifer [63, 64] to establish a large deviation principle (LDP) for the players' learning process in *discrete* time. Using this LDP, we are able to estimate the "rare event" probability of going against the mean dynamics for a period of time that is sufficient to reach another metastable component. In turn, this allows us to show that the metastable components of the process are *internally chain transitive* (ICT) under the associated mean dynamics, i.e., up to small jumps, the mean dynamics may wander around freely within such a set, without collapsing to a smaller subset.

Armed with this characterization, our main results can be summarized as follows (cf. Theorems 1–5):

1. **Concentration on ICT sets.** In the long run, the induced sequence of play spends an exponentially small amount of time away from the ICT sets of the associated mean dynamics.
2. **Boltzmann–Gibbs law.** The sequence of play follows a Boltzmann–Gibbs law with temperature η: each metastable component $\mathcal{K}$ is observed with probability proportional to $e^{-E(\mathcal{K})/\eta}$, where $E(\mathcal{K})$ is an "energy level" determined by the underlying game and the noise entering the process.
3. **Avoidance of unstable points and sets.** Every unstable ICT set is "dominated" by a stable one, which is visited exponentially more often.

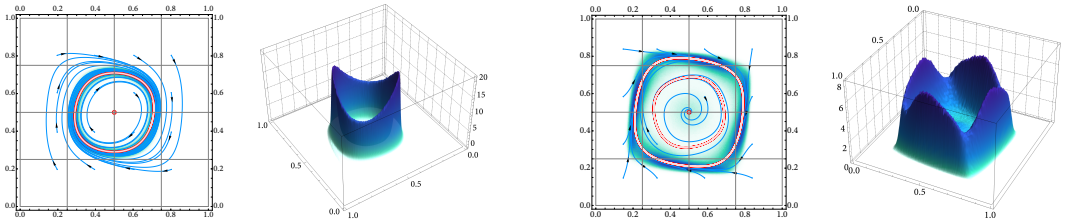


Figure 1: Learning in two "almost bilinear" zero-sum games (cf. Section 4 for details). In both cases, the heatmap and the density landscape depict the observed long-run distribution of play after 10^3 iterations averaged over 10^4 sample runs with learning rate $\eta = 0.1$; the mean flow of the system is drawn in light blue lines, and its ICT sets are highlighted in red. Even though there are multiple ICT sets (left: one stable, one unstable; right: two stable, one unstable), the problem's ground state arises as the single, exponentially more probable outcome.

4. **Selection of the ground state.** The trajectory of play eventually settles down with exponentially high probability near the system's *ground state*, that is, the set of points that minimize its energy.

Finally, we also derive an explicit expression for the long-run distribution of play when the noise is Gaussian, and we use this expression to design a simple modification of the EW algorithm which concentrates exponentially around potential maximizers in potential games, and social welfare maximizers in common interest ones. [For a numerical illustration of our results, see Figs. 1 and 2.]

The game-theoretic significance of the above is (at least) twofold. First and foremost, our results provide a quantitative characterization of the dynamics' probable outcomes. These metastable components include the game's equilibria, but also recurrent patterns of play, like periodic orbits (as in Matching Pennies), limit cycles (as in Jordan's Matching Pennies), etc. In general, there is no game-theoretic way to assign a value to these sets, nor a way to guess which of these sets may be more or less likely to be observed in the long run. The derived Boltzmann–Gibbs law *assigns* such a value—their *energy*—thus providing a precise quantitative answer to this question.

More than that, our results show that learning ultimately acts as a selection mechanism. Out of a mélange of equilibria, limit cycles, and other ICT sets, the long-run distribution of play assigns all but an exponentially small fraction of probability to the ground state of the process; moreover, the sharpness of this concentration is controlled by the players' learning rate, becoming sharper for small η. Generically—that is, in the absence of payoff ties and other degeneracies—the ground state is a single metastable component, which may thus be regarded as the "fate" of the dynamics. We find this interpretation particularly appealing, as it identifies the outcome that arises in the long run with exponentially high probability, even in the presence of multiple equilibria or other attractors.

Related work. The literature on learning in games is vast, so we briefly review here some of the most closely related threads, and we refer the reader to Appendix B for a more in-depth look.

The first important distinction that needs to be made is between learning with *perfect* (deterministic) vs. *imperfect* (stochastic) gradient feedback. The deterministic literature is itself extremely broad, and contains both global and local convergence results, depending on whether the game admits an amenable potential-like structure or not; for an entry point to this prodigious literature, see [27, 30, 32, 44, 81, 84, 87] and citing works. However, the phenomena that we examine are inherently stochastic in nature, so they admit no relevant counterpart in a deterministic context.

In the stochastic setting, there are two further distinctions that need to be made. First, there is a stark conceptual difference between running an algorithm with constant vs. *vanishing* learning rate. In the latter case, the theory of stochastic approximation essentially precludes recurrence and concentration phenomena, so play typically converges to some (random) ICT set [13, 16, 67]; in this regard, the vanishing and constant learning rate regimes are mutually orthogonal, in both results and techniques. For a representative—but, by necessity, far from exhaustive—sample of works on regularized learning with a vanishing learning rate, see [14, 15, 26, 39, 54, 56, 72–74, 83, 86] and references therein.

In the stochastic, constant learning rate regime—which is often preferred in practice owing to its simplicity and superior empirical performance—a large corpus of literature has focused on games with a potential-like structure, providing results for reaching an ε-approximate Nash equilibrium in a time-averaged or "best-iterate" sense, cf. [21, 71, 94, 108] and cited works. In the absence of such a structure, most results focus on the convergence properties of the empirical frequency of play relative to some correlated notion of equilibrium, cf. [4, 18, 19, 31, 34, 35, 45, 47] as well as

the textbook of Cesa-Bianchi & Lugosi [24]. However, these works do not shed any light on the *actual* trajectory of play—especially in non-concave games—so they are likewise orthogonal to our analysis and results. To the best of our knowledge, the only results on the long-run distribution of play come from Huang & Zhang [57], Loizou et al. [76] and Lotidis et al. [78] for monotone games, and from Vlatakis-Gkaragkounis et al. [104] for weakly monotone variational inequalities; however, these results are essentially convex in nature, so they cannot be used to understand the metastability landscape of general non-concave games—and, in particular, the transitions from one metastable component to another.

The works which are closest to our own are the recent papers by Azizian et al. [9, 10] who likewise leverage the Freidlin–Wentzell–Kifer (FWK) theory of large deviations to quantify the long-run distribution of stochastic gradient descent (SGD) in non-convex minimization problems. Our approach shares this common ancestor with [9, 10], but with several major differences: (*a*) first, there is no potential-like structure in general games, implying in particular that the metastable components of the process are not zeros of the driving flow; (*b*) there is no immediate mechanism relating metastability to chain transitivity (as in the case of critical points of minimization problems); and (*c*) unstable components of non-stationary points are significantly more complicated to treat than saddle points. These challenges are both technical and conceptual, so we discuss them in detail in Appendix A.

Finally, we should stress that our results concern the *discounted*—or *regularized*—regime of learning with exponential weights, that is, when past observations are discounted over more recent ones. In the context of finite games, this scheme goes back (at least) to Leslie & Collins [73] while, more recently, in stochastic games, it is known as "magnetic mirror descent" [98]; in all cases, it is a regularized version of entropic mirror descent. The difference between the discounted and undiscounted settings is significant, not only mathematically, but also conceptually, as it may lead to different predictions of the dynamics' long-run behavior—chief among these differences being that the latter may converge to equilibria or other invariant sets on the boundary of the game's strategy space, whereas the former cannot. For completeness, we provide a detailed account of these differences in Appendix B.

2 Problem setup and blanket assumptions

2.1. Continuous games. Throughout the sequel, we will focus on games with a finite number of players $i \in \mathcal{N} = \{1, \dots, N\}$, each selecting an action x_i from some compact convex subset $\mathcal{X}_i$ of a d_i-dimensional Euclidean space $\mathcal{V}_i \cong \mathbb{R}^{d_i}$. Aggregating over all players, we will write $\mathcal{X} = \prod_i \mathcal{X}_i$ for the space of all *action profiles* $x = (x_1, \dots, x_N)$ and $d = \sum_i d_i$ for the dimension of the *ambient space* $\mathcal{V} = \prod_i \mathcal{V}_i$. To simplify notation, we will also use the shorthand $x = (x_i; x_{-i})$ to highlight the action of player $i \in \mathcal{N}$ against the action profile $x_{-i} = (x_j)_{j \neq i}$ of all other players $j \in \mathcal{N}_{-i} := \mathcal{N} \setminus \{i\}$.

Given a profile $x \in \mathcal{X}$, the reward of each player $i \in \mathcal{N}$ is determined by an associated *payoff function* $u_i \colon \mathcal{X} \to \mathbb{R}$. In terms of regularity, we will assume that each u_i is C^2-smooth, and we will write

$$v_i(x) = \nabla_{x_i} u_i(x_i; x_{-i}) \quad \text{and} \quad v(x) = (v_1(x), \dots, v_N(x)) \tag{1}$$

for the *individual gradient field* of player $i \in \mathcal{N}$ and the ensemble thereof, respectively.[1]

A *continuous game in normal form* is defined as a tuple $\mathcal{G} \equiv \mathcal{G}(\mathcal{N}, \mathcal{X}, u)$ with players, actions, and payoff functions as above. In this general context, we say that (*a*) $\mathcal{G}$ is a *min-max* game if $\mathcal{N} = \{1, 2\}$ and $u_1 = -u_2$; and (*b*) $\mathcal{G}$ is a *potential game* if it admits a *potential function* $f \colon \mathcal{X} \to \mathbb{R}$ such that

$$u_i(x_i; x_{-i}) - u_i(x_i'; x_{-i}) = f(x_i; x_{-i}) - f(x_i'; x_{-i}) \quad \text{for all } x, x' \in \mathcal{X} \text{ and all } i \in \mathcal{N}. \tag{2}$$

Finally, we say that x^* is a *Nash equilibrium* (NE) of $\mathcal{G}$ if it is *unilaterally stable* in the sense that

$$u_i(x^*) \geq u_i(x_i; x_{-i}^*) \quad \text{for all } x_i \in \mathcal{X}_i \text{ and all } i \in \mathcal{N}. \tag{NE}$$

To strike a sensible balance between concreteness and generality, we will specialize in the rest of our paper to continuous games played over products of simplices, i.e., games where each player's action space $\mathcal{X}_i$ admits a simplicial decomposition of the form

$$\mathcal{X}_i = \prod_{k=1}^{m_i} \Delta(\mathcal{A}_{ik}) \text{ for a collection of finite sets } \mathcal{A}_{ik},\ k = 1, \dots, m_i, \text{ with } \sum_k |\mathcal{A}_{ik}| = d_i. \tag{3}$$

We detail below two of the most prominent examples of such games, *finite* and *stochastic games*:

[1] To enable the use of ordinary derivatives, we tacitly assume here that the players' payoff functions are defined on a neighborhood of $\mathcal{X}$. This is done for convenience only; none of our results depend on this device.

Example 1 (Finite games)**.** A finite game $\Gamma \equiv \Gamma(\mathcal{N}, \mathcal{A}, u)$ consists of the same primitives as a continuous game, the only difference being that the players' action sets $\mathcal{A}_i$, $i = 1, \dots, N$, are now finite, and "actions" $\alpha_i \in \mathcal{A}_i$ are traditionally called "pure strategies". In the *mixed extension* $\bar{\Gamma}$ of Γ, every player draws a pure strategy based on some *mixed strategy* $x_i \in \Delta(\mathcal{A}_i)$ and their corresponding *mixed* payoff is given by the multilinear expectation $u_i(x) = \sum_{\alpha \in \mathcal{A}} x_\alpha u_i(\alpha)$ where $x_\alpha = \prod_i x_{i\alpha_i}$ is the joint probability of playing the pure strategy profile $\alpha = (\alpha_1, \dots, \alpha_N)$. In this way, the mixed extension of a finite game can be seen as a continuous game with action spaces $\mathcal{X}_i \equiv \Delta(\mathcal{A}_i)$. ✦

Example 2 (Stochastic / Markov games)**.** Going back to Shapley [97], a *stochastic* (or *Markov*) *game* is a standard model for multi-agent reinforcement learning, and consists of the following primitives: *(i)* a finite set of *states* $s \in \mathcal{S} = \{1, \dots, S\}$; *(ii)* for each state $s \in \mathcal{S}$, a (finite) *state game* $\Gamma^s = (N, \mathcal{A}, u^s)$; and *(iii)* a *transition matrix* $P(s' \mid s, \alpha)$ which indicates the probability of transitioning from state s to state s' when the players play the pure action profile $\alpha = (\alpha_1, \dots, \alpha_N) \in \mathcal{A}$. Then, during play, if the game is at state s, each player $i \in \mathcal{N}$ selects an action $\alpha_i \in \mathcal{A}_i$ based on a *policy* $\pi_i \colon \mathcal{S} \to \Delta(\mathcal{A}_i)$, that is, a recipe prescribing a state-dependent mixed strategy $\pi_i(s) \in \Delta(\mathcal{A}_i)$. Subsequently, each player gets a reward of $u_i^s(\alpha_1, \dots, \alpha_N)$, the game transitions to a new state s' based on the transition matrix P, and the process repeats.

In the so-called "episodic setting", there is a finite probability q that the game terminates at any given round, thus giving rise to an "episode" of (random) duration T. In this case, the mean cumulative reward of each player at the end of an episode is given by the *value function*

$$V_i(\pi) \coloneqq \mathbb{E}\left[\sum_{n=0}^{T} u_i^{s_n}(\alpha_n) \,\middle|\, s_0 \sim \rho \right] \tag{4}$$

where $\pi = (\pi_1, \dots, \pi_N)$ is the player's policy profile, and ρ is the initial state distribution. In this way, a stochastic game can be seen as a continuous game $\mathcal{G} \equiv \mathcal{G}(\mathcal{N}, \Pi, V)$ played over the players' *policy spaces* $\Pi_i = \Delta(\mathcal{A}_i)^{\mathcal{S}}$ with payoff functions $V_i \colon \Pi \to \mathbb{R}$. ✦

For completeness, we detail some more examples of the class of games under study in Appendix C.1.

2.2. The learning process. We now proceed to describe our core scheme for learning in continuous games. At a high level, the process unfolds according to the following sequence of events:

1. At each instance $n = 0, 1, \dots$, every player $i \in \mathcal{N}$ selects an action $x_{i,n} \in \mathcal{X}_i$.
2. Players receive their rewards and feedback, they update their actions, and the process repeats.

The two main elements of the players' learning process are thus *a*) the feedback that players get; and *b*) how they update their actions. We discuss these elements in turn, beginning with the former.

The feedback process. In terms of information, we will assume that, at each stage $n = 0, 1, \dots$ of the process, every player $i \in \mathcal{N}$ receives as feedback a stochastic payoff gradient of the form

$$\hat{v}_{i,n} = \mathsf{V}_i(x_n; \omega_n) \quad \text{or, aggregating over all players,} \quad \hat{v}_n = \mathsf{V}(x_n; \omega_n) \tag{5a}$$

where, in obvious notation, we set $\hat{v}_n = (\hat{v}_{i,n})_{i \in \mathcal{N}}$ and $\mathsf{V}(x_n; \omega_n) = (\mathsf{V}_i(x_n; \omega_n))_{i \in \mathcal{N}}$. In the above, *(i)* the sequence ω_n, $n = 0, 1, \dots$, is an i.i.d. sequence of random seeds drawn from some compact subset Ω of $\mathbb{R}^p$, and *(ii)* each V_i, $i \in \mathcal{N}$, is a *stochastic gradient oracle* (SGO) which, when queried at $x \in \mathcal{X}$ with seed $\omega \in \Omega$, returns a random vector of the form

$$\mathsf{V}_i(x; \omega) = v_i(x) + \mathsf{Z}_i(x; \omega) \tag{SGO}$$

where the generic "error term" $\mathsf{Z}_i(x; \omega)$ captures all sources of noise and randomness in the process. Our standing assumptions for the aggregate error $\mathsf{Z} = (\mathsf{Z}_1, \dots, \mathsf{Z}_N)$ will then be as follows:

Assumption 1. The error term $\mathsf{Z} \colon \mathcal{X} \times \Omega \to \mathbb{R}^d$ of (SGO) is C^2-smooth and *proper*, i.e.,

$$\mathbb{E}[\mathsf{Z}(x; \omega)] = 0 \;\text{ and }\; \operatorname{cov}(\mathsf{Z}(x; \omega)) \succ 0 \;\text{ for all } x \in \mathcal{X}. \tag{6}$$

This basic assumption has two main consequences: *(i)* that the players' feedback consists of *unbiased* gradient estimates (i.e., with no inherent systematic bias or offset); and *(ii)* that the noise does not vanish identically at any point $x \in \mathcal{X}$. This last requirement is tied to the fact that we are working in a bona fide stochastic setting, clearly separated from the deterministic, perfect feedback regime. In this regard, the oracle model above is sufficiently flexible to account for most established frameworks for learning under randomness and uncertainty, including e.g., minibatch sampling (where ω represents the sampled minibatch) and/or cases where gradients are obtained by physical measurements in an underlying stochastic system (e.g., as in data networks).

Regularized exponential weights. In terms of learning, we will focus throughout on a family of *regularized exponential / multiplicative weights* (REW) algorithms of the general form

$$y_{n+1} = y_n + \eta(\hat{v}_n - \lambda y_n) \qquad x_{n+1} = \operatorname{softmax}(y_{n+1}) \tag{REW}$$

where $\eta > 0$ is the algorithm's step-size, $\lambda > 0$ is a regularization parameter, and the softmax operator collects the components of the *multinomial logit choice* (MLC) map $\Lambda\colon \mathbb{R}^d \to \mathcal{X}$ defined here as

$$[\operatorname{softmax}(y)]_{i\alpha_i} \coloneqq \Lambda_{i\alpha_i}(y) \equiv \frac{\exp(y_{i\alpha_i})}{\sum_{\beta_i \in \mathcal{A}_{ik}} \exp(y_{i\beta_i})} \tag{MLC}$$

for all $\alpha_i \in \mathcal{A}_{ik}$, $k = 1, \dots, m_i$, $i \in \mathcal{N}$. When $\lambda = 0$, (REW) reduces to the landmark exponential / multiplicative weights algorithm introduced by Vovk [105] and Littlestone & Warmuth [75], and subsequently pioneered by Auer et al. [6, 7], Freund & Schapire [37] and Beck & Teboulle [12] in the context of multi-armed bandits, zero-sum games, and convex optimization, respectively; for a representative sample of more recent works, see [5, 29, 50, 66, 81, 94]. The regularized variant that we consider here for $\lambda > 0$ goes back to Leslie & Collins [73], and it has likewise been studied extensively in the literature, cf. [1, 26, 55, 70, 92, 98, 101] and references therein.

To gain some intuition on the method, note that the gradient update step in (REW) can be unrolled as

$$y_{n+1} = (1 - \eta\lambda)^{n+1} y_0 + \eta \sum_{k=0}^{n} (1 - \eta\lambda)^{n-k} \hat{v}_k \,. \tag{7}$$

In other words, (REW) maintains an exponentially discounted aggregate of gradient steps in the problem's dual space $\mathcal{Y} \coloneqq \mathbb{R}^d$, and then performs an "exponentiated projection" step to generate an action $x_n \in \mathcal{X}$. In this regard, for $\lambda > 0$, the method exhibits a "recency bias" whereby past information is discounted in favor of more recent gradient estimates. To maintain this natural discount mechanism, we will assume in the sequel that (REW) is run with a positive—but small—regularization weight λ; we discuss the role of λ in detail in Section 3.1 and Appendix B.

3 Mean dynamics and large deviations

With all these preliminaries in place, we move now to characterize the long-run behavior of (REW) and, in particular, its concentration properties. To that end, our chief figure of merit will be the *mean occupation measure* of x_n, defined here as

$$\pi_n(\mathcal{B}) = \mathbb{E}\left[\frac{1}{n} \sum_{k=0}^{n-1} \mathbb{1}\{x_k \in \mathcal{B}\}\right] \quad \text{for every Borel } \mathcal{B} \subseteq \mathcal{X}. \tag{8}$$

In words, $\pi_n(\mathcal{B})$ simply measures the expected fraction of time that x_n has spent in $\mathcal{B}$ up to time n, so the long-run statistics of (REW) can be understood by means of the limiting distribution $\lim_{n\to\infty} \pi_n$. Accordingly, setting aside for the moment all questions of existence and uniqueness of this limit, our goal will be to determine the subsets of $\mathcal{X}$ where π_n concentrates in the long run, and the degree of this concentration. Because the formal statement of our results cannot be separated from the analysis used to derive them, we begin with a high-level roadmap of our approach:

- Our first step is to identify the *mean dynamics* of the process, that is, a continuous-time dynamical system describing the mean evolution of play under (REW).
- Next, we derive a *large deviation principle* (LDP) for (REW), quantifying its "most likely" trajectories as well as the probability of going against the flow for a long enough period of time.
- In turn, this allows us to establish a crisp dynamical characterization of the *metastable components* of (REW), that is, those subsets of $\mathcal{X}$ where x_n spends all but a negligible fraction of its time: any metastable component of (REW) is *internally chain transitive* under the mean dynamics, i.e., it is a union of orbits of the mean dynamics connected by arbitrarily small jumps.
- We then leverage this characterization to construct the so-called *metastability chain* of (REW), a finite-state Markov chain whose transitions capture the probability of moving from one metastable component to another, and which defines the *energy* of each metastable component.

Using these energy levels, we establish a Boltzmann–Gibbs law for the long-run behavior of π_n through a combination of Markov chain and exponential concentration arguments. In the rest of this section and the next, we provide a condensed account of the technical scaffolding required to make these statements precise; for the details, we refer the reader to Appendices C–I.

3.1. The mean dynamics. Our starting point is the observation that (REW) can be recast as $(y_{n+1} - y_n)/\eta = v(\Lambda(y_n)) - \lambda y_n + \mathsf{Z}(x_n;\omega_n)$, i.e., as a noisy Euler discretization of the *mean dynamics*

$$\dot{y} = V(y) - \lambda y \tag{MD}$$

where, recalling that $x_n = \Lambda(y_n)$, we write $V(y) \equiv v(\Lambda(y))$ for the payoff field of the game expressed in terms of the dual, gradient-tracking variables $y \in \mathcal{Y}$. The terminology "mean dynamics" alludes here to the fact that, for small η, it is plausible to expect that the noise $\mathsf{Z}(x_n;\omega_n)$ in (REW) "washes out" in the long run, so the mean asymptotic behavior of y_n under (REW) may be captured by that of (MD). Hence, by a straightforward differentiation (which we detail in Lemma C.1), we readily get that the corresponding mean dynamics for x_n are given by the *regularized replicator dynamics* [26]

$$\dot{x}_{i\alpha_i} = x_{i\alpha_i}\Big[v_{i\alpha_i}(x) - \textstyle\sum_{\beta_i\in\mathcal{A}_{ik}} x_{i\beta_i} v_{i\beta_i}(x)\Big] - \lambda x_{i\alpha_i}\Big[\log x_{i\alpha_i} - \textstyle\sum_{\beta_i\in\mathcal{A}_{ik}} x_{i\beta_i}\log x_{i\beta_i}\Big] \tag{RRD}$$

for all $\alpha_i \in \mathcal{A}_{ik}$, $k = 1,\dots,m_i$, and all $i \in \mathcal{N}$.

The dynamics (RRD) were originally derived by Coucheney et al. [26] in the context of learning in finite games, and are the natural candidate for the "mean" asymptotic behavior of x_n; as such, it is important to pause for a moment in order to consider their game-theoretic properties.

The most important one is that, contrary to the undiscounted regime $\lambda = 0$—in which case (RRD) reduces to the standard replicator dynamics of Taylor & Jonker [99]—*Nash equilibria are not stationary* under (RRD) for $\lambda > 0$. Instead, as was shown by Coucheney et al. [26], the interior rest points of (RRD) are the game's *logit equilibria* (LE), that is, the Nash equilibria of the entropically regularized game $\mathcal{G}^\lambda \equiv (\mathcal{N},\mathcal{X},u^\lambda)$ with payoff functions

$$u_i^\lambda(x) = u_i(x) - \lambda h_i(x_i) \quad \text{where} \quad h_i(x_i) = \textstyle\sum_{\alpha_i\in\mathcal{A}_i} x_{i\alpha_i}\log x_{i\alpha_i}\,. \tag{$\mathcal{G}^\lambda$}$$

The notion of a logit equilibrium dates back to van Damme [102] and McKelvey & Palfrey [80], who introduced it as a "bounded rationality" approximant to the game's Nash equilibria; in particular, as was shown by Govindan et al. [46], logit equilibria are $\mathcal{O}(\lambda)$-Nash equilibria.

Remark 1. In light of the above, the regularization weight λ plays a double role: it represents at the same time the players' *tendency to forget* (as the discount factor for past information) and their *tendency to explore* (as the mixing factor in $\mathcal{G}^\lambda$). Thus, even though our results apply to the entire spectrum $\lambda \in (0,\infty)$, we will mostly be interested in the "small λ" regime, where the players forget and explore "just enough" so that Nash and logit equilibria practically coincide. We stress here however that the limiting case $\lambda = 0$ is fundamentally different; for a detailed discussion, see Appendix B. ✦

3.2. Large deviations. With now proceed to quantify the link between (REW) and (RRD), that is, to determine how probable it is for x_n to remain close to the flow of (RRD), and how (im)probable it is to go against it. To estimate the probability of such rare—but significant—events, we provide below a *large deviation principle* (LDP) for (REW) in the spirit of Freidlin & Wentzell [36].

An important obstacle to this is that the updates $x_n \leftarrow x_{n+1}$ are not additive, so it is difficult to work directly with the process x_n. One could instead work with the dual process y_n—whose updates *are* additive—but the logit map $\Lambda\colon \mathcal{Y} \to \mathcal{X}$ is not invertible, and this creates difficulties for the theory: for example, y_n could run off to infinity even as its primal image $x_n = \Lambda(y_n)$ remains unchanged.

For this reason, we begin by reformulating (REW) in a space of lower dimension which *can* be mapped to $\mathcal{X}$ in a one-to-one way. To that end, we pick an "exclusion coordinate" $\hat{\alpha}_{ik}$ for each $k = 1,\dots,m_i$, $i = 1,\dots,N$, we let $\mathcal{A}_{ik}^* := \mathcal{A}_{ik}\setminus\{\hat{\alpha}_{ik}\}$, and we introduce the *state variables*

$$z_{i\alpha_i} = y_{i\alpha_i} - y_{i\hat{\alpha}_{ik}} \quad \text{for all } \alpha_i \in \mathcal{A}_{ik}^* \text{ and all } k = 1,\dots,m_i,\ i \in \mathcal{N}. \tag{9}$$

We then define the associated *state space* $\mathcal{Z} = \prod_i \mathbb{R}^{d_i - m_i}$ (where state variables live), and the *reduced logit map* $Q\colon \mathcal{Z} \to \mathcal{X}$ given in coordinates by the expression

$$Q_{i\alpha_i}(z) = \Big(1 + \textstyle\sum_{\beta_i\in\mathcal{A}_{ik}^*}\exp(z_{i\beta_i})\Big)^{-1} \cdot \begin{cases} 1 & \text{if } \alpha_i = \hat{\alpha}_{ik}, \\ \exp(z_{i\alpha_i}) & \text{if } \alpha_i \in \mathcal{A}_{ik}^*. \end{cases} \tag{10}$$

With all this in hand, we show in Appendix C that (*a*) Q is a diffeomorphism from $\mathcal{Z}$ to $\operatorname{ri}\mathcal{X}$; and (*b*) as a result, (REW) can be rewritten equivalently in Robbins–Monro form as

$$z_{n+1} = z_n + \eta(F(z_n) + U_n - \lambda z_n) \tag{RM}$$

where the drift and noise terms—$F(z_n)$ and U_n respectively—are defined in components as

$$F_{i\alpha_i}(z) = v_{i\alpha_i}(Q(z)) - v_{i\hat{\alpha}_{ik}}(Q(z)) \quad \text{and} \quad U_{i\alpha_i,n} = \mathsf{Z}_{i\alpha_i}(Q(z_n);\omega_n) - \mathsf{Z}_{i\hat{\alpha}_{ik}}(Q(z_n);\omega_n)\,. \tag{11}$$

Armed with this tight reformulation of (REW), and proceeding in the spirit of the continuous-time theory of Freidlin & Wentzell [36], we consider the *Hamiltonian* and *Lagrangian* functions

$$\mathcal{H}(z,p) := \log \mathbb{E}_\omega\left[e^{\langle p,\mathsf{F}(z;\omega)\rangle}\right] \quad \text{and} \quad \mathcal{L}(z,w) := \mathcal{H}^*(z,w) = \max_{p\in\mathcal{Z}}\{\langle w,p\rangle - \mathcal{H}(z,p)\} \tag{12}$$

where $\mathsf{F}(z;\omega) = F(z) - \lambda z + \mathsf{U}(z;\omega)$ is the corresponding oracle signal in (RM) and $\mathsf{U}_{i\alpha_i}(z;\omega) = \mathsf{Z}_{i\alpha_i}(Q(z);\omega) - \mathsf{Z}_{i\hat{\alpha}_{ik}}(Q(z);\omega)$ is the induced (random) error for coordinate $\alpha_i \in \mathcal{A}_i^*$. Then, given a continuous curve $\gamma(t) \in \mathcal{Z}$, $t \in [0,T]$, we define the problem's (normalized) *action functional* as

$$\mathcal{S}_T[\gamma] = \int_0^T \mathcal{L}(\gamma(t),\dot{\gamma}(t))\,dt \tag{13}$$

with the convention that $\mathcal{S}_T[\gamma] = \infty$ if γ is not absolutely continuous. The key property of this definition—and the linchpin of FWK theory—is that the action $\mathcal{S}_T[\gamma]$ provides a measure of the "likelihood" of observing a trajectory of (RM) near γ, with lower action values indicating exponentially higher likelihood. At a high level, this intuition ultimately translates to the *large deviation principle*

$$\mathbb{P}(\sup_{t\in[0,T]}\|z(t)-\gamma(t)\| < \delta) \approx \exp(-\mathcal{S}_T[\gamma]/\eta) \tag{LDP}$$

where $\delta > 0$ is a fixed precision level, and $z(t) = z_n + (t/\eta - n)(z_{n+1} - z_n)$ for $t \in [\eta n, \eta(n+1)]$ is the affine interpolation of z_n (for a formal statement and proof, see Appendices D.5 and I).

This estimate provides a concrete link between z_n and its mean dynamics, so it plays a pivotal role in our analysis. In particular, a first key consequence is that, if a trajectory of play x_n, $n = 0, 1, \dots$, has a high action value (via the transformed process $z_n = Q^{-1}(x_n)$), then it becomes exponentially less probable to be observed relative to lower-action ones. In view of this "least action principle", the cost of moving from one point to another under (REW) can be quantified by means of the *quasi-potential*

$$P(x,x') = \inf\{\mathcal{S}_T[\gamma] : T \in \mathbb{N},\, \gamma \in \Gamma_T(\mathcal{Z}) \;\text{ and }\; Q(\gamma(0)) = x,\, Q(\gamma(T)) = x'\} \tag{14}$$

where $\Gamma_T(\mathcal{Z})$ is the space of continuous curves $\gamma\colon [0,T] \to \mathcal{Z}$. This notion—extending and adapting [9, 36]—captures the transition cost incurred by the "most probable" path from x to x'. As such, it plays a key part in identifying the "trapping regions" of $\mathcal{X}$; we make this precise in the next section.

4 Metastability analysis and main results

4.1. Metastable components. In view of the above, it is reasonable to expect that two points $x, x' \in \mathcal{X}$ are equally likely to be observed under (REW) if the transitions between them are costless. This naturally induces the equivalence relation

$$x \sim x' \iff P(x,x') = P(x',x) = 0\,, \tag{15}$$

which essentially means that x_n may move freely between x and x'. This motivates the following:

Definition 1. An equivalence class $\mathcal{C}$ of $\sim$ is called *basic*—or a *metastable component*—if it consists of (*a*) at least two distinct points of $\mathcal{X}$; or (*b*) a single stationary point of (RRD).

Building on the above, a natural way of measuring the energy cost of moving from one metastable component to another is to look at the least costly transition between elements of said sets, viz.

$$P(\mathcal{C},\mathcal{C}') =: \inf\{P(x,x') : x \in \mathcal{C}, x' \in \mathcal{C}'\}\,. \tag{16}$$

In this way, (LDP) suggests that the probability of moving from $\mathcal{C}$ to $\mathcal{C}'$ scales approximately as $\exp(-P(\mathcal{C},\mathcal{C}')/\eta)$, i.e., high-cost transitions are exponentially less likely than low-cost ones. We make this statement precise in Appendix F.3 under the following technical assumption:

Assumption 2. There are finitely many metastable components, $\mathcal{C}_i$, $i = 1, \dots, K$, and the induced pairwise transition costs are all finite, i.e., $P_{ij} := P(\mathcal{C}_i,\mathcal{C}_j) < \infty$ for all $i, j = 1, \dots, K$.

This requirement is imposed solely for technical convenience, and to streamline our presentation—in particular, to avoid complicated statements and arguments involving non-communicating classes. Then, modulo this standing assumption, we have the following exponential concentration result.

Theorem 1. *Let $\mathcal{M}^\delta$ be a δ-neighborhood of the metastable region $\mathcal{M} := \bigcup_i \mathcal{C}_i$ of the system. Then there exists some $c \equiv c_\delta > 0$ such that, for small enough $\eta > 0$ and large enough n, we have*

$$\pi_n(\mathcal{M}^\delta) \geq 1 - e^{-c/\eta}. \tag{17}$$

Theorem 1 shows that, in a very precise sense, the system's metastable components $\mathcal{C}_i$ are the only ones that matter in the long run: they act as "traps", and x_n spends all but an exponentially small fraction of the time in their vicinity. In this regard, it is essential to examine if there is a precise link between the system's metastable components and the mean dynamics (RRD).

To do so, we will need some basic definitions from dynamical systems, beginning with the *flow* $\Phi\colon \mathbb{R}_+ \times \mathcal{X} \to \mathcal{X}$ of (RRD), i.e., the map $(t, x) \mapsto \Phi_t(x)$ such that $\Phi_t(x)$, $t \geq 0$, is the solution orbit to (RRD) starting at x. Then, given a nonempty closed subset $\mathcal{S}$ of $\mathcal{X}$, we say that:

1. $\mathcal{S}$ is *invariant* if $\Phi_t(\mathcal{S}) = \mathcal{S}$ for all $t \geq 0$.
2. $\mathcal{S}$ is (*Lyapunov*) *stable* if, for every neighborhood $\mathcal{U}$ of $\mathcal{S}$, there exists a neighborhood $\mathcal{U}'$ of $\mathcal{S}$ such that $\Phi_t(\mathcal{U}') \subseteq \mathcal{U}$ for all $t \geq 0$; otherwise, if $\mathcal{S}$ is not stable, $\mathcal{S}$ is called *unstable*.
3. $\mathcal{S}$ is an *attractor* if it is invariant and there exists a compact neighborhood $\mathcal{K}$ of $\mathcal{S}$ such that $\operatorname{dist}(\Phi_t(x), \mathcal{S}) \to 0$ as $t \to \infty$ uniformly in $x \in \mathcal{K}$.
4. $\mathcal{S}$ is *internally chain transitive* (ICT) if it is invariant and $\Phi|_{\mathcal{S}}$ admits no proper attractors in $\mathcal{S}$.

Equivalently, ICT sets may be viewed as "minimal connected periodic orbits up to arbitrarily small jumps" [13, Prop. 5.3], so they comprise the irreducible building blocks of the dynamics. Proposition 1 below collects some properties of the ICT sets of (RRD) which are standard in the literature:

Proposition 1. *For all $\lambda > 0$, we have:*

- *Every ICT set of* (RRD) *is contained in the* (*relative*) *interior of $\mathcal{X}$.*
- *The logit equilibria of the game are ICT under* (RRD).
- *In finite min-max games,* (RRD) *admits a unique ICT set—the game's unique logit equilibrium.*
- *If the game admits a potential f, the ICT sets of* (RRD) *consist of critical points of the regularized potential $f^\lambda := f - \lambda h$.*

In our context, the importance of ICT sets is captured by the following characterization:

Theorem 2. *Every metastable set $\mathcal{C}_i$, $i = 1, \dots, K$, of* (REW) *is an ICT set of* (RRD).

Thus, combining Theorems 1 and 2, we obtain a crisp answer to our original question of where x_n settles down: *with exponentially high probability, the trajectory of play concentrates around the ICT sets of the system's mean dynamics, irrespective of the statistics of the noise or other considerations.*

These sets can be equilibria, periodic orbits, limit cycles, or any other repeating pattern of play—but they are always invariant under (RRD), and they always have the internal recurrence property that any point thereof can be reached by any other point. The question that remains is the degree of concentration per ICT set. We address this below.

4.2. The metastability chain and the long-run distribution of play. To quantify which metastable components are more likely to arise in the long run, recall that (LDP) indicates that the probability of moving from $\mathcal{C}_i$ to $\mathcal{C}_j$ scales as $e^{-P_{ij}/\eta}$. Moreover, since x_n spends a negligible amount of time away from the system's metastable components (by Theorem 1), it stands to reason that the asymptotic distribution of x_n can be approximated by a finite-state Markov chain with state space $\mathcal{K} = \{1, \dots, K\}$ and transition probabilities $p_{ij} = e^{-P_{ij}/\eta}$, $i, j \in \mathcal{K}$. We call this the *metastability chain* of the system, and we represent it as a complete weighted directed graph $\mathcal{G}$ on $\mathcal{K}$ with edge weights $(p_{ij})_{i,j\in\mathcal{K}}$.

Making this analogy precise is a highly intricate affair, which we detail in Appendix G. On the other hand, analyzing the metastability chain itself is more straightforward and relies on established Markov chain techniques. Specifically, by the Markov chain tree theorem [52, 65, 69], the stationary distribution ν_i of the metastability chain can be expressed as $\nu_i \propto \sum_{T_i \in \mathcal{T}_i} w(T_i)$, where $\mathcal{T}_i$ is the set of spanning trees of $\mathcal{G}$ that point to i, and $w(T_i) = \prod_{(j,k)\in T_i} p_{jk} = \exp\big(-\sum_{(j,k)\in T_i} P_{jk}/\eta\big)$. When η is small, the dominant contribution in ν_i is that of the minimum-weight spanning tree, so we define the *energy* of $\mathcal{K}_i$ as

$$E_i = \min_{T_i \in \mathcal{T}_i} \sum_{(j,k)\in T_i} P_{jk}\,. \tag{18}$$

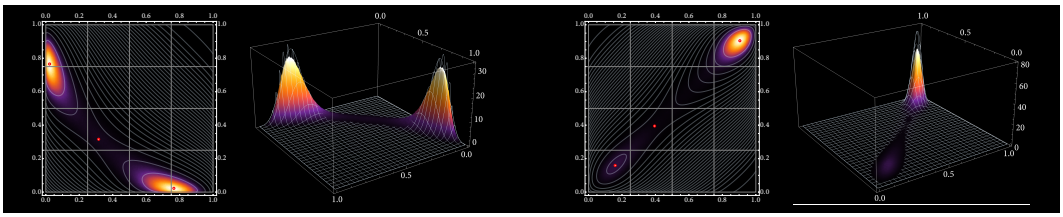


Figure 2: Learning with shaped Gaussian noise as per Section 4.3 in a 2×2 congestion game (left) and a 2×2 coordination game (right; payoff-dominant Nash at top right). The heatmaps and density landscapes represent the observed distribution of play; the superimposed contours and wireframes indicate our theoretical prediction.

The terminology "energy" is justified by our main result below, which shows that π_n follows an approximate Boltzmann–Gibbs law with temperature η and energy levels E_i [68, 90].

Theorem 3. *Fix a tolerance level $\varepsilon > 0$, let $\mathcal{C}_i, \mathcal{C}_j$ be two metastable components of the system, and let $\mathcal{C}_i^\delta, \mathcal{C}_j^\delta$ be δ-neighborhoods thereof. Then, for small enough $\delta, \eta > 0$ and large enough n, we have:*

$$\left|\eta \log \pi_n(\mathcal{C}_i^\delta) + E_i - E_0\right| \leq \varepsilon \quad \text{and} \quad \left|\eta \log\big(\pi_n(\mathcal{C}_i^\delta)/\pi_n(\mathcal{C}_j^\delta)\big) + E_i - E_j\right| \leq \varepsilon\,, \tag{BG}$$

or, more concisely, $\pi_n(\mathcal{C}_i^\delta) \propto \exp(-(E_i + \mathcal{O}(\varepsilon))/\eta)$.

Theorem 3 provides a sharp characterization of the distribution of probability mass over the metastable sets of the process, but it does not provide an off-the-shelf selection criterion. We provide two such criteria below, beginning with the separation between stable and unstable sets.

Theorem 4. *Let $\mathcal{C}_i$ be a metastable component that is not Lyapunov stable. Then there exists a metastable component $\mathcal{C}_j$ that is Lyapunov stable and a constant $c \equiv c_{ij} > 0$ such that*

$$\pi_n(\mathcal{C}_i^\delta)\big/\pi_n(\mathcal{C}_j^\delta) \leq \exp(-c/\eta) \tag{19}$$

for all sufficiently small $\delta, \eta > 0$ and large enough n.

In words, Theorem 4 means that, with exponentially high probability, unstable sets are actively selected against under (REW). Such behavior is well-documented for *strict* saddle-points of unconstrained gradient systems [8, 40, 60, 85, 91] but, to the best of our knowledge, there is no comparable result for *constrained*, *non-gradient* systems, and *without any strictness requirement* to boot.

Our last result sharpens this selection mechanism considerably, by showing that, in the long run, only the system's *ground state* $\mathcal{C}_0 = \bigcup_{i \in \arg\min_j E_j} \mathcal{C}_i$ is observed with high probability.

Theorem 5. *Let $\mathcal{C}_0^\delta$ be a δ-neighborhood of the ground state $\mathcal{C}_0$. Then there exists a constant $c \equiv c_\delta > 0$ such that, for all sufficiently small $\delta, \eta > 0$ and large enough n, we have $\pi_n(\mathcal{C}_0^\delta) \geq 1 - e^{-c/\eta}$.*

In words, Theorem 5 shows that learning ultimately concentrates at a *single* energy level, which, generically consists of a single metastable component—what one could call the *outcome of the game*. This outcome is selected over a highly diverse miscellany of equilibria, limit cycles and other ICT sets, a behavior which comes in stark contrast with the case of a vanishing step-size $\eta_n \to 0$ and/or no regularization ($\lambda = 0$). In the latter case, play may converge to an initialization-dependent ICT set [13–15, 54, 83, 86], so there is no selection at hand. We find this particularly appealing as it shows that regularization induces in this way an "implicit bias" in game-theoretic learning.

4.3. An application: Potential games with Gaussian noise. We conclude our discussion with a concrete illustration of the energy landscape of a potential game with gradient feedback perturbed by zero-mean Gaussian noise with state-dependent covariance $\Sigma(x)$. At an informal level—and deferring the details to Appendix C.5 it is possible to show in this case that (*a*) the metastable components of the process are *precisely* the critical components of $f_\lambda := f - \lambda h$; and (*b*) the system's energy levels are

$$E_i = -f_\lambda(\mathcal{C}_i) + \mathcal{O}\big(\|\Sigma^{-1} - \operatorname{Jac} Q\|_\infty\big)\,. \tag{20}$$

In addition to its predictive value, this expression is also important from the viewpoint of algorithm design: it shows that if players *inject* Gaussian noise with shaped covariance $\Sigma(x) \propto (\operatorname{diag}(x) - x \otimes x)^{-1}$, the energy levels of the process are $-f_\lambda(\mathcal{C}_i)$. For small enough λ, this means that the ground state of the system is precisely the maximum of the potential, so (REW) will select in this case the game's potential maximizer—and, likewise, the social welfare maximizer in common interest games. Extending this steering scheme to more general games would be a significant advance in the field.

Acknowledgments and Disclosure of Funding

This research was supported in part by the French National Research Agency (ANR) in the framework of the PEPR IA FOUNDRY project (ANR-23-PEIA-0003). PM is also a member of the Archimedes Research Unit/Athena RC, and was partially supported by project MIS 5154714 of the National Recovery and Resilience Plan Greece 2.0 funded by the European Union under the NextGenerationEU Program.

References

[1] Abe, K., Ariu, K., Sakamoto, M., and Iwasaki, A. Adaptively perturbed mirror descent for learning in games. In *ICML '24: Proceedings of the 41st International Conference on Machine Learning*, 2024.

[2] Ahunbay, M. Ş. First-order (coarse) correlated equilibria in non-concave games. In *STOC '26: Proceedings of the 58th Annual ACM Symposium on Theory of Computing*, 2026.

[3] Ahunbay, M. Ş. and Bichler, M. Semicoarse correlated equilibria and LP-based guarantees for gradient dynamics in normal-form games. In *EC '25: Proceedings of the 26th ACM Conference on Economics and Computation*, 2025.

[4] Anagnostides, I., Daskalakis, C., Farina, G., Fishelson, M., Golowich, N., and Sandholm, T. Near-optimal no-regret learning for correlated equilibria in multi-player general-sum games. In *STOC '22: Proceedings of the 54th Annual ACM Symposium on Theory of Computing*, 2022.

[5] Arora, S., Hazan, E., and Kale, S. The multiplicative weights update method: A meta-algorithm and applications. *Theory of Computing*, 8(1):121–164, 2012.

[6] Auer, P., Cesa-Bianchi, N., Freund, Y., and Schapire, R. E. Gambling in a rigged casino: The adversarial multi-armed bandit problem. In *Proceedings of the 36th Annual Symposium on Foundations of Computer Science*, 1995.

[7] Auer, P., Cesa-Bianchi, N., and Fischer, P. Finite-time analysis of the multiarmed bandit problem. *Machine Learning*, 47:235–256, 2002.

[8] Azizian, W., Iutzeler, F., Malick, J., and Mertikopoulos, P. The rate of convergence of Bregman proximal methods: Local geometry vs. regularity vs. sharpness. *SIAM Journal on Optimization*, 34(3):2440–2471, September 2024.

[9] Azizian, W., Iutzeler, F., Malick, J., and Mertikopoulos, P. What is the long-run distribution of stochastic gradient descent? A large deviations analysis. In *ICML '24: Proceedings of the 41st International Conference on Machine Learning*, 2024.

[10] Azizian, W., Iutzeler, F., Malick, J., and Mertikopoulos, P. The global convergence time of stochastic gradient descent in non-convex landscapes: Sharp bounds via large deviations. In *ICML '25: Proceedings of the 42nd International Conference on Machine Learning*, 2025.

[11] Bailey, J. P., Nagarajan, S. G., and Piliouras, G. Stochastic multiplicative weights updates in zero-sum games. https://arxiv.org/abs/2110.02134, 2021.

[12] Beck, A. and Teboulle, M. Mirror descent and nonlinear projected subgradient methods for convex optimization. *Operations Research Letters*, 31(3):167–175, 2003.

[13] Benaïm, M. Dynamics of stochastic approximation algorithms. In Azéma, J., Émery, M., Ledoux, M., and Yor, M. (eds.), *Séminaire de Probabilités XXXIII*, volume 1709 of *Lecture Notes in Mathematics*, pp. 1–68. Springer Berlin Heidelberg, 1999.

[14] Benaïm, M., Hofbauer, J., and Sorin, S. Stochastic approximations and differential inclusions. *SIAM Journal on Control and Optimization*, 44(1):328–348, 2005.

[15] Benaïm, M., Hofbauer, J., and Sorin, S. Stochastic approximations and differential inclusions, part II: Applications. *Mathematics of Operations Research*, 31(4):673–695, 2006.

[16] Benveniste, A., Métivier, M., and Priouret, P. *Adaptive Algorithms and Stochastic Approximations*. Springer, 1990.

[17] Blackwell, D. An analog of the minimax theorem for vector payoffs. *Pacific Journal of Mathematics*, 6:1–8, 1956.

[18] Blum, A. and Mansour, Y. Learning, regret minimization, and equilibria. In Nisan, N., Roughgarden, T., Tardos, É., and Vazirani, V. V. (eds.), *Algorithmic Game Theory*, chapter 4. Cambridge University Press, 2007.

[19] Blum, A. and Mansour, Y. From external to internal regret. *Journal of Machine Learning Research*, 8:1307–1324, December 2007.

[20] Blume, L. E. The statistical mechanics of strategic interaction. *Games and Economic Behavior*, 5:387–424, 1993.

[21] Brown, N., Lerer, A., Gross, S., and Sandholm, T. Deep counterfactual regret minimization. In *ICML '19: Proceedings of the 36th International Conference on Machine Learning*, 2019.

[22] Cai, Y., Daskalakis, C., Luo, H., Wei, C.-Y., and Zheng, W. On tractable Φ-equilibria in non-concave games. In *NeurIPS '24: Proceedings of the 38th International Conference on Neural Information Processing Systems*, 2024.

[23] Cai, Y., Daskalakis, C., Luo, H., Wei, C.-Y., and Zheng, W. Proximal regret and proximal correlated equilibria: A new tractable solution concept for online learning and games. In *STOC '26: Proceedings of the 58th Annual ACM Symposium on Theory of Computing*, 2026.

[24] Cesa-Bianchi, N. and Lugosi, G. *Prediction, Learning, and Games*. Cambridge University Press, 2006.

[25] Coste, M. An introduction to o-minimal geometry, 1999.

[26] Coucheney, P., Gaujal, B., and Mertikopoulos, P. Penalty-regulated dynamics and robust learning procedures in games. *Mathematics of Operations Research*, 40(3):611–633, August 2015.

[27] Daskalakis, C. and Panageas, I. The limit points of (optimistic) gradient descent in min-max optimization. In *NeurIPS '18: Proceedings of the 32nd International Conference of Neural Information Processing Systems*, 2018.

[28] Daskalakis, C., Goldberg, P. W., and Papadimitriou, C. H. The complexity of computing a Nash equilibrium. *SIAM Journal on Computing*, 39(1):195–259, 2009.

[29] Daskalakis, C., Deckelbaum, A., and Kim, A. Near-optimal no-regret algorithms for zero-sum games. *Games and Economic Behavior*, 92:327–348, 2015.

[30] Daskalakis, C., Ilyas, A., Syrgkanis, V., and Zeng, H. Training GANs with optimism. In *ICLR '18: Proceedings of the 2018 International Conference on Learning Representations*, 2018.

[31] Daskalakis, C., Fishelson, M., and Golowich, N. Near-optimal no-regret learning in general games. In *NeurIPS '21: Proceedings of the 35th International Conference on Neural Information Processing Systems*, 2021.

[32] Diakonikolas, J., Daskalakis, C., and Jordan, M. I. Efficient methods for structured nonconvex-nonconcave min-max optimization. In *AISTATS '21: Proceedings of the 24th International Conference on Artificial Intelligence and Statistics*, 2021.

[33] Douc, R., Moulines, E., Priouret, P., and Soulier, P. *Markov chains*, volume 4. Springer, 2018.

[34] Farina, G., Anagnostides, I., Luo, H., Lee, C.-W., Kroer, C., and Sandholm, T. Near-optimal no-regret learning dynamics for general convex games. In *NeurIPS '22: Proceedings of the 36th International Conference on Neural Information Processing Systems*, 2022.

[35] Foster, D. P. and Vohra, R. V. Calibrated learning and correlated equilibrium. *Games and Economic Behavior*, 21(1): 40–55, October 1997.

[36] Freidlin, M. I. and Wentzell, A. D. *Random Perturbations of Dynamical Systems*. Springer, 1 edition, 1998.

[37] Freund, Y. and Schapire, R. E. Adaptive game playing using multiplicative weights. *Games and Economic Behavior*, 29:79–103, 1999.

[38] Fudenberg, D. and Harris, C. Evolutionary dynamics with aggregate shocks. *Journal of Economic Theory*, 57(2): 420–441, August 1992.

[39] Fudenberg, D. and Kreps, D. M. Learning mixed equilibria. *Games and Economic Behavior*, 5(320-367), 1993.

[40] Ge, R., Huang, F., Jin, C., and Yuan, Y. Escaping from saddle points – Online stochastic gradient for tensor decomposition. In *COLT '15: Proceedings of the 28th Annual Conference on Learning Theory*, 2015.

[41] Giannou, A., Vlatakis-Gkaragkounis, E. V., and Mertikopoulos, P. Survival of the strictest: Stable and unstable equilibria under regularized learning with partial information. In *COLT '21: Proceedings of the 34th Annual Conference on Learning Theory*, 2021.

[42] Giannou, A., Vlatakis-Gkaragkounis, E. V., and Mertikopoulos, P. The convergence rate of regularized learning in games: From bandits and uncertainty to optimism and beyond. In *NeurIPS '21: Proceedings of the 35th International Conference on Neural Information Processing Systems*, 2021.

[43] Giannou, A., Lotidis, K., Mertikopoulos, P., and Vlatakis-Gkaragkounis, E. V. On the convergence of policy gradient methods to Nash equilibria in general stochastic games. In *NeurIPS '22: Proceedings of the 36th International Conference on Neural Information Processing Systems*, 2022.

[44] Gidel, G., Berard, H., Vignoud, G., Vincent, P., and Lacoste-Julien, S. A variational inequality perspective on generative adversarial networks. In *ICLR '19: Proceedings of the 2019 International Conference on Learning Representations*, 2019.

[45] Gordon, G. J., Greenwald, A., and Marks, C. No-regret learning in convex games. In *ICML '08: Proceedings of the 25th International Conference on Machine Learning*, 2008.

[46] Govindan, S., Reny, P. J., and Robson, A. J. A short proof of Harsanyi's purification theorem. *Games and Economic Behavior*, 45:369–374, 2003.

[47] Hart, S. and Mas-Colell, A. A simple adaptive procedure leading to correlated equilibrium. *Econometrica*, 68(5): 1127–1150, September 2000.

[48] Hart, S. and Mas-Colell, A. Uncoupled dynamics do not lead to Nash equilibrium. *American Economic Review*, 93(5): 1830–1836, 2003.

[49] Hart, S. and Mas-Colell, A. Stochastic uncoupled dynamics and Nash equilibrium. *Games and Economic Behavior*, 57: 286–303, 2006.

[50] Héliou, A., Cohen, J., and Mertikopoulos, P. Learning with bandit feedback in potential games. In *NIPS '17: Proceedings of the 31st International Conference on Neural Information Processing Systems*, 2017.

[51] Hernández-Lerma, O. and Lasserre, J. B. *Markov chains and invariant probabilities*, volume 211. Springer Science & Business Media, 2003.

[52] Hill, T. L. Studies in irreversible thermodynamics iv. diagrammatic representation of steady state fluxes for unimolecular systems. *Journal of theoretical biology*, 10(3):442–459, 1966.

[53] Hofbauer, J. and Imhof, L. A. Time averages, recurrence and transience in the stochastic replicator dynamics. *The Annals of Applied Probability*, 19(4):1347–1368, 2009.

[54] Hofbauer, J. and Sandholm, W. H. On the global convergence of stochastic fictitious play. *Econometrica*, 70(6): 2265–2294, November 2002.

[55] Hopkins, E. and Posch, M. Attainability of boundary points under reinforcement learning. *Games and Economic Behavior*, 53(1):110–125, March 2005.

[56] Hsieh, Y.-P., Mertikopoulos, P., and Cevher, V. The limits of min-max optimization algorithms: Convergence to spurious non-critical sets. In *ICML '21: Proceedings of the 38th International Conference on Machine Learning*, 2021.

[57] Huang, K. and Zhang, S. New first-order algorithms for stochastic variational inequalities. *SIAM Journal on Optimization*, 32(4):2745–2772, 2022.

[58] Imhof, L. A. The long-run behavior of the stochastic replicator dynamics. *The Annals of Applied Probability*, 15(1B): 1019–1045, 2005.

[59] Ioffe, A. D., Tikhomirov, V. M., and Makowski, K. *Theory of Extremal Problems*. North Holland, Amsterdam, NL, 1979.

[60] Jin, C., Ge, R., Netrapalli, P., Kakade, S. M., and Jordan, M. I. How to escape saddle points efficiently. In *ICML '17: Proceedings of the 34th International Conference on Machine Learning*, 2017. URL http://proceedings.mlr.press/v70/jin17a.html.

[61] Kallenberg, O. *Foundations of Modern Probability*, volume 99 of *Probability Theory and Stochastic Modelling*. Springer, 2021.

[62] Kelly, F. P., Maulloo, A. K., and Tan, D. K. H. Rate control for communication networks: shadow prices, proportional fairness and stability. *Journal of the Operational Research Society*, 49(3):237–252, March 1998.

[63] Kifer, Y. *Random Perturbations of Dynamical Systems*. Birkhäuser, Boston, MA, 1988.

[64] Kifer, Y. A discrete-time version of the Wentzell-Freidlin theory. *Annals of Probability*, 18(4):1676–1692, October 1990.

[65] Kirchhoff, G. Ueber die auflösung der gleichungen, auf welche man bei der untersuchung der linearen vertheilung galvanischer ströme geführt wird. *Annalen der Physik*, 148(12):497–508, 1847.

[66] Kleinberg, R., Piliouras, G., and Tardos, É. Multiplicative updates outperform generic no-regret learning in congestion games. In *STOC '09: Proceedings of the 41st annual ACM symposium on the Theory of Computing*, 2009.

[67] Kushner, H. J. and Yin, G. G. *Stochastic approximation algorithms and applications*. Springer-Verlag, New York, NY, 1997.

[68] Landau, L. D. and Lifshitz, E. M. Statistical physics. In *Course of Theoretical Physics*, volume 5. Pergamon Press, Oxford, 1976.

[69] Leighton, F. and Rivest, R. Estimating a probability using finite memory. *IEEE Transactions on Information Theory*, 32(6):733–742, 1986.

[70] Leonardos, S., Piliouras, G., and Spendlove, K. Exploration-exploitation in multi-agent dompetition: Convergence with bounded rationality. In *NeurIPS '21: Proceedings of the 35th International Conference on Neural Information Processing Systems*, 2021.

[71] Leonardos, S., Overman, W., Panageas, I., and Piliouras, G. Global convergence of multi-agent policy gradient in Markov potential games. In *ICLR '22: Proceedings of the 2022 International Conference on Learning Representations*, 2022.

[72] Leslie, D. S. and Collins, E. J. Convergent multiple-timescales reinforcement learning algorithms in normal form games. *The Annals of Applied Probability*, 13(4):1231–1251, November 2003.

[73] Leslie, D. S. and Collins, E. J. Individual Q-learning in normal form games. *SIAM Journal on Control and Optimization*, 44(2):495–514, 2005.

[74] Leslie, D. S. and Collins, E. J. Generalised weakened fictitious play. *Games and Economic Behavior*, 56(2):285–298, August 2006.

[75] Littlestone, N. and Warmuth, M. K. The weighted majority algorithm. *Information and Computation*, 108(2):212–261, 1994.

[76] Loizou, N., Berard, H., Gidel, G., Mitliagkas, I., and Lacoste-Julien, S. Stochastic gradient descent-ascent and consensus optimization for smooth games: Convergence analysis under expected co-coercivity. In *NeurIPS '21: Proceedings of the 35th International Conference on Neural Information Processing Systems*, 2021.

[77] Lotidis, K., Mertikopoulos, P., Bambos, N., and Blanchet, J. H. Invariance and concentration properties of gradient-based learning in games. In *CDC '25: Proceedings of the 64th IEEE Annual Conference on Decision and Control*, 2025.

[78] Lotidis, K., Mertikopoulos, P., Bambos, N., and Blanchet, J. H. Multi-agent learning under uncertainty: Recurrence vs. concentration. In *NeurIPS '25: Proceedings of the 39th International Conference on Neural Information Processing Systems*, 2025.

[79] Lotidis, K., Mertikopoulos, P., Bambos, N., and Blanchet, J. H. Robust equilibria in continuous games: From strategic to dynamic robustness robust equilibria in continuous games: From strategic to dynamic robustness robust equilibria in continuous games: From strategic to dynamic robustness. In *NeurIPS '25: Proceedings of the 39th International Conference on Neural Information Processing Systems*, 2025.

[80] McKelvey, R. D. and Palfrey, T. R. Quantal response equilibria for normal form games. *Games and Economic Behavior*, 10(6):6–38, 1995.

[81] Mertikopoulos, P. Regret, equilibrium, and learning in games: A guided tour. In Sorin, S. and von Stengel, B. (eds.), *Equilibria in Games: Existence, Selection, and Dynamics*. Cambridge University Press, 2026.

[82] Mertikopoulos, P. and Moustakas, A. L. The emergence of rational behavior in the presence of stochastic perturbations. *The Annals of Applied Probability*, 20(4):1359–1388, July 2010.

[83] Mertikopoulos, P. and Zhou, Z. Learning in games with continuous action sets and unknown payoff functions. *Mathematical Programming*, 173(1-2):465–507, January 2019.

[84] Mertikopoulos, P., Lecouat, B., Zenati, H., Foo, C.-S., Chandrasekhar, V., and Piliouras, G. Optimistic mirror descent in saddle-point problems: Going the extra (gradient) mile. In *ICLR '19: Proceedings of the 2019 International Conference on Learning Representations*, 2019.

[85] Mertikopoulos, P., Hallak, N., Kavis, A., and Cevher, V. On the almost sure convergence of stochastic gradient descent in non-convex problems. In *NeurIPS '20: Proceedings of the 34th International Conference on Neural Information Processing Systems*, 2020.

[86] Mertikopoulos, P., Hsieh, Y.-P., and Cevher, V. A unified stochastic approximation framework for learning in games. *Mathematical Programming*, 203:559–609, January 2024.

[87] Mokhtari, A., Ozdaglar, A., and Pattathil, S. A unified analysis of extra-gradient and optimistic gradient methods for saddle point problems: proximal point approach. In *AISTATS '20: Proceedings of the 23rd International Conference on Artificial Intelligence and Statistics*, 2020.

[88] Nash, J. F. *Non-cooperative games*. PhD thesis, Princeton University, 1950.

[89] Nisan, N., Roughgarden, T., Tardos, É., and Vazirani, V. V. (eds.). *Algorithmic Game Theory*. Cambridge University Press, 2007.

[90] Nishimori, H. *Statistical Physics of Spin Glasses and Information Processing*. Oxford University Press, 2001.

[91] Pemantle, R. Nonconvergence to unstable points in urn models and stochastic aproximations. *Annals of Probability*, 18 (2):698–712, April 1990.

[92] Perolat, J., Munos, R., Lespiau, J.-B., Omidshafiei, S., Rowland, M., Ortega, P., Burch, N., Anthony, T., Balduzzi, D., Vylder, B. D., Piliouras, G., Lanctot, M., and Tuyls, K. From poincaré recurrence to convergence in imperfect information games: Finding equilibrium via regularization. In *ICML '21: Proceedings of the 38th International Conference on Machine Learning*, 2021.

[93] Pradelski, B. S. R. and Young, H. P. Learning efficient nash equilibria in distributed systems. *Games and Economic Behavior*, 75(2):882–897, 2012.

[94] Rakhlin, A. and Sridharan, K. Optimization, learning, and games with predictable sequences. In *NIPS '13: Proceedings of the 27th International Conference on Neural Information Processing Systems*, 2013.

[95] Rigollet, P. and Hütter, J.-C. High-dimensional statistics. *arXiv preprint arXiv:2310.19244*, 2023.

[96] Sandholm, W. H. and Staudigl, M. Large deviations and stochastic stability in the small noise double limit. *Theoretical Economics*, 11(1):279–355, September 2016.

[97] Shapley, L. S. Stochastic games. *Proceedings of the National Academy of Sciences of the USA*, 39:1095–1100, 1953.

[98] Sokota, S., D'Orazio, R., Kolter, Z., Loizou, N., Lanctot, M., Mitliagkas, I., Brown, N., and Kroer, C. A unified approach to reinforcement learning, quantal response equilibria, and two-player zero-sum games. In *ICLR '23: Proceedings of the 2023 International Conference on Learning Representations*, 2023.

[99] Taylor, P. D. and Jonker, L. B. Evolutionary stable strategies and game dynamics. *Mathematical Biosciences*, 40(1-2): 145–156, 1978.

[100] Tullock, G. Efficient rent seeking. In Tollison, J. M. B. R. D. and Tullock, G. (eds.), *Toward a theory of the rent-seeking society*. Texas A&M University Press, 1980.

[101] Tuyls, K., 't Hoen, P. J., and Vanschoenwinkel, B. An evolutionary dynamical analysis of multi-agent learning in iterated games. *Autonomous Agents and Multi-Agent Systems*, 12(1):115–153, 2006.

[102] van Damme, E. *Stability and perfection of Nash equilibria*. Springer-Verlag, Berlin, 1987.

[103] van den Dries, L. and Miller, C. Geometric categories and o-minimal structures. *Duke Mathematical Journal*, 84(2), August 1996.

[104] Vlatakis-Gkaragkounis, E. V., Giannou, A., Chen, Y., and Xie, Q. Stochastic methods in variational inequalities: Ergodicity, bias and refinements. In *AISTATS '24: Proceedings of the 27th International Conference on Artificial Intelligence and Statistics*, 2024.

[105] Vovk, V. G. Aggregating strategies. In *COLT '90: Proceedings of the 3rd Workshop on Computational Learning Theory*, pp. 371–383, 1990.

[106] Young, H. P. *Strategic Learning and its Limits*. Oxford University Press, 2004.

[107] Young, H. P. Learning by trial and error. *Games and Economic Behavior*, 65(2):626–643, March 2009.

[108] Zinkevich, M., Johanson, M., Bowling, M., and Piccione, C. Regret minimization in games with incomplete information. In *NIPS '07: Proceedings of the 21st International Conference on Neural Information Processing Systems*, 2007.

Appendix contents

A Appendix overview **15**

B Further related work **17**

C Game-theoretic elements and calculations **19**

D Setup and preliminary results **29**

E Equivalence classes **33**

F Transitions **55**

G Long-run distribution **69**

H Potential games with Gaussian noise **85**

I Proofs of the large deviation principles **96**

J Numerical illustrations **102**

A Appendix overview

A.1. Overview. This appendix is structured as follows. Appendix B presents additional related work. The remaining appendices are devoted to the proofs of our results.

Most of the proofs—Appendices D–I—are carried out in a general abstract setting that strips away the specific structure of games played over products of simplices, allowing the core arguments to be developed in full generality. Appendix C explains how the setting and results from the main text fit into our general framework: specifically, how (REW) can be reformulated as a Robbins–Monro iteration and how the simplicial structure of the game can be abstracted *via* a diffeomorphic transform, thereby allowing the main results of Sections 3 and 4 to be recovered as special cases of the general theorems established in the subsequent appendices.

The remaining appendices operate entirely within this general framework. Appendix D introduces the setup, blanket assumptions, and basic properties of the Hamiltonian and Lagrangian, and states the large deviation principle (LDP) for the process, whose proofs are deferred to Appendix I. Appendix E introduces the equivalence classes, which are the central objects of the analysis: we provide multiple equivalent characterizations, establish their key structural properties (compactness, invariance, closure), and show that they are included in the internally chain transitive (ICT) sets of the mean dynamics. Appendix F studies the transitions of the metastability chain, establishing exponential bounds on both transition times and transition probabilities between equivalence classes. Appendix G contains the central results on the long-run distribution of the process, including the Boltzmann–Gibbs law for the invariant measure and the mean occupation measure, the avoidance of unstable components, and the concentration around the ground state. Appendix H specializes these results to the case of potential games with Gaussian noise, deriving an explicit expression for the energy levels and identifying conditions under which the ground state coincides with the maximizer of the potential.

A.2. Technical challenges. In order to establish our results in the general, non-potential setting of multi-player games, we need to overcome a series of major challenges, both structural and methodological. We outline the main ones below.

Metastable components. The metastable components of the process—the equivalence classes of Appendices E.2 and E.3—are intricate objects: they may contain limit cycles, quasi-periodic orbits, or other recurrent patterns of the mean dynamics, and they need not reduce to fixed points of the vector field. Already defining these objects precisely and establishing their basic properties (alternative characterization and relationship to ICT sets) requires substantial work, and constitutes the bulk of Appendix E.

Regularity near metastable components. A key step in our analysis is obtaining fine regularity properties of the quasi-potential near the metastable components. Since the vector field need not

vanish there, this turns out to be a considerable challenge. In particular, the continuity of the one-step rate function near equivalence classes—which plays a central role in the transition probability estimates—requires a technical derivation developed in Appendix E.

Stability analysis and long-run distribution. Stable components play a crucial role both in the description of the asymptotic distribution and in the proofs used to derive it. In our setting, since stability is a purely dynamical notion—Lyapunov stability of a compact invariant set under the mean flow—obtaining the quantitative estimates required is challenging. In particular, the equivalence between dynamical instability and the Freidlin–Wentzell notion of non-attractors requires a careful new argument, which is the main technical contribution of Appendix G.1. The remaining derivations in Appendix G, though standard, are included in full to ensure the correctness of the arguments.

Lyapunov analysis. Without an underlying potential function, no natural coercive Lyapunov function is available. To guarantee that the process does not spend too much time at infinity, we develop Appendix F.1 which is able to accommodate dissipative but non-gradient vector fields.

Potential games with Gaussian noise. The results of Appendix H require tracking the Jacobian of the diffeomorphism Q through the noise covariance, due to the composition of the potential f with Q. This leads to a nontrivial covariance geometry.

Large deviations and transition probabilities. The proofs in Appendix I and the transition probability estimates in Appendix F.3 follow a similar overall structure to standard results, but we include complete proofs for the sake of self-containedness.

All of the above challenges were absent in Azizian et al. [9], the closest work to our own, which focused on SGD in unbounded Euclidean spaces and thus operated in a significantly more structured setting. Taken together, they necessitate a wide range of new technical developments, which are detailed in the appendices that follow.

B Further related work

In this appendix, we discuss some works which, while relevant to the overall scope of our work, did not otherwise fit the main narrative. We organize these works by thread below, while noting that several of these works touch on several threads, so this classification is, by necessity, imperfect.

B.1. Learning and stochastic stability. Our work relies substantially on estimating the invariant measures of Markov chains, an idea that has surfaced several times in game theory and game-theoretic learning. The first instance that we are aware of is the work of Blume [20] on strategic interactions in a lattice, which introduced the process of log-linear in potential games (a variant of Metropolis sampling) and established convergence to a Gibbs measure on the values of the game's potential. This is related to the "shaping" technique that we outline in Section 4.3 but, otherwise, the two schemes are completely different—as are the results.

Stochastic stability ideas were also employed in the "mood-based" learning schemes popularized by Young [106, 107] for learning pure Nash equilibria by trial-and-error—and, later, by Pradelski & Young [93], for concentrating play in a neighborhood of the game's welfare-maximizing Nash equilibrium. More recently, Sandholm & Staudigl [96] considered a model of stochastic evolution under general noisy best-response protocols, and derived the system's behavior in the small noise / large population double limit. This last work is notable in that it also uses the theory of Freidlin & Wentzell [36] to derive the stochastic stability properties of the limit distribution but, otherwise, the setting and results are completely different.

B.2. Undiscounted regularized learning. In the literature, the term "regularized learning" is often used to refer to a variant of (REW) run with a different regularizer—the so-called "follow-the-regularized-leader" method—and *no* discounting. The difference between the discounted and undiscounted settings is significant, not only mathematically, but also conceptually.

Perhaps the most important differences can be seen in the context of finite games. Here, undiscounted regularized learning can—and in fact does—converge globally in potential games [50], and locally to strict Nash equilibria [41, 42] in general games—in fact, even in general stochastic games [43] and/or continuous games [79]. In the absence of an implicit regularization / discounting mechanism the two settings are orthogonal, and could lead to misunderstandings about what the long-run behavior of learning actually looks like.

In the undiscounted regime, the work which is perhaps closest to our own is the preprint of Bailey et al. [11] which examines the long-run distribution of undiscounted exponential / multiplicative weights in finite two-player zero-sum games with pure payoff vector feedback—that is, at each stage, each player is made aware of the payoff they would have obtained by playing any of their pure strategies against the pure strategy profile chosen by their opponents. Bailey et al. [11] established a transience property of the resulting dynamics, meaning in particular that any concentration occurs at the boundary. By contrast, in the context of continuous games, Lotidis et al. [77, 78] established recurrence for certain undiscounted variants of regularized learning, but they only considered the strongly monotone case for this result, so there are no phenomena of metastability complicating the analysis.

Some versions of the above questions have also been studied in a diffusive context in continuous time, but there is no meaningful comparison of results here; for the interested reader, a good entry point to this literature would be the original works of Fudenberg & Harris [38], Hofbauer & Imhof [53], Imhof [58] and [82].

B.3. Coarse notions of equilibrium. A related line of recent work with a measure-theoretic flavor concerns equilibrium relaxations for non-concave games, motivated in part by the fact that standard no-regret dynamics may fail to converge to meaningful stationary solution concepts in this setting. The first such contribution is the work of Cai et al. [22], who introduced the notion of a "tractable Φ-equilibrium" for non-concave games and showed how these can be approached—in the sense of Blackwell [17]—by suitable gradient-based dynamics. Independently and shortly thereafter, Ahunbay [2] proposed a first-order variant of coarse correlated equilibrium for continuous games, based on local first-order deviations rather than arbitrary global deviations. Subsequent refinements appeared in two closely related directions: Cai et al. [23] developed the framework of proximal regret and proximal correlated equilibria, while Ahunbay & Bichler [3] introduced the notion of semicoarse

correlated equilibrium as a refinement of coarse correlated equilibrium (CCE) specialized to finite games.

Viewed abstractly, these works examine the empirical frequency of play—an atomic measure on the game's joint action space—and they show that, under different assumptions, all the limit points of said distribution lie in a certain subset of the space of measures that admits a variational characterization similar to that of coarse correlated equilibria in finite games. In this regard, the "coarse" nature of these equilibrium relaxations alludes to the fact that it is not possible, under said analyses, to determine the actual limit points of the empirical frequency of play—only the long-run satisfaction of a variational characterization thereof.

C Game-theoretic elements and calculations

C.1 Examples of games 19
C.2 Reformulation of REW dynamics 19
C.3 Main Results 23
C.4 Potential games in reduced coordinates 24
C.5 Potential games with truncated Gaussian noise 27

We recall from the main text that we are considering continuous games played over products of simplices. That is, each player $i \in \mathcal{N} = \{1, \dots, N\}$ selects an action x_i from $\mathcal{X}_i = \prod_{k=1}^{m_i} \Delta(\mathcal{A}_{ik})$ for a collection of finite sets $\mathcal{A}_{ik}$, $1 \leq k \leq m_i$, with $\sum_k |\mathcal{A}_{ik}| = d_i$. We then write $\mathcal{X} = \prod_i \mathcal{X}_i$ for the space of all action profiles, and $d = \sum_i d_i$ for the dimension of its ambient space. Given a profile $x \in \mathcal{X}$, the reward of each player $i \in \mathcal{N}$ is then determined by a C^2-smooth payoff function $u_i \colon \mathcal{X} \to \mathbb{R}$, which generates individual gradient fields of the form $v_i(x) = \nabla_{x_i} u_i(x_i; x_{-i})$ and $v(x) = (v_1(x), \dots, v_N(x))$.

C.1. Examples of games. In this section, we outline a range of examples that fall in the general framework of Section 2 in order to showcase the applications of the class of games under study.

Example C.1 (Cournot oligopolies)**.** Consider N firms, each producing a quantity $x_i \in [0, C_i]$ of some good, up to the firm's capacity $C_i > 0$, $i = 1, \dots, N$. The good is then priced as a function of the total supply provided to the market following the model $P(x) = a - b \sum_{i=1}^{N} x_i$, where $a, b > 0$ are market-related constants. The net utility of the i-th firm is then given by the total profit $u_i(x) = x_i P(x) - c_i x_i$, where $c_i > 0$ denotes the firm's unit production cost. The resulting game $\mathcal{G} = \mathcal{G}(\mathcal{N}, \mathcal{X}, u)$ is known as a Cournot oligopoly, and it plays a fundamental role in economic theory and market design. ✦

Example C.2 (Fisher markets and Kelly auctions)**.** In a Fisher market [89], there is a set $\mathcal{N} = \{1, \dots, N\}$ of *buyers*—or *players*—that seek to share a set of m perfectly divisible *goods* or *resources* (e.g., ad space, CPU/GPU runtime, bandwidth, etc.). The allocation mechanism for these goods follows a proportionally fair price-setting rule that is sometimes referred to as a *Kelly auction* [62] or a *Tullock contest* [100]: each player $i \in \mathcal{N}$ bids x_{ik} per unit of the k-th good up to the player's individual budget, assumed for simplicity to be equal to 1 for all players, that is, $\sum_{k=1}^{m} x_{ik} \leq 1$ for all $i \in \mathcal{N}$. The price p_k of the k-th good is then set to be equal to the players' bids plus a *reserve bid* $b_k \geq 0$, that is, $p_k = b_k + \sum_{i \in \mathcal{N}} x_{ik}$; then, each player gets a prorated fraction $w_{ik} = x_{ik}/p_k$ of each good. Thus, factoring in the marginal gain g_{ik} that each player $i \in \mathcal{N}$ gets per unit of the k-th resource, the player's utility is given by

$$u_i(x) = \sum_{k=1}^{m} [g_{ik} w_{ik} - x_{ik}] = \sum_{k=1}^{m} \left[\frac{g_{ik} x_{ik}}{b_k + \sum_{j \in \mathcal{N}} x_{jk}} - x_{ik} \right] \tag{C.1}$$

Since each player's space of bids is the "corner-of-cube" simplex $\{x_i \in \mathbb{R}^m_+ : \sum_{k=1}^{m} x_{ik} \leq 1\}$, the resulting game is easily seen to be an example of the class of games under study. ✦

C.2. Reformulation of REW dynamics. In this section, we focus on *regularized exponential/multiplicative weights* (REW) algorithms of the general form

$$y_{n+1} = y_n + \eta(\hat{v}_n - \lambda y_n) \qquad x_{n+1} = \Lambda(y_{n+1}) \tag{EW$_\lambda$}$$

where $\eta > 0$ is the algorithm's step-size, $\lambda > 0$ is a regularization parameter, and $\Lambda \colon \mathbb{R}^d \to \mathcal{X}$ denotes the *logit choice* (LC) map, defined component-wise as

$$\Lambda_{i\alpha_i}(y) = \frac{\exp(y_{i\alpha_i})}{\sum_{\beta_i \in \mathcal{A}_{ik}} \exp(y_{i\beta_i})} \qquad \text{for all } \alpha_i \in \mathcal{A}_{ik},\ k = 1, \dots, m_i,\ i \in \mathcal{N}. \tag{C.2}$$

As in the main text, we define continuous-time dynamics capturing the mean behavior of y_n under (REW) as

$$\dot{y} = V(y) - \lambda y \tag{MD}$$

where $V(y) \equiv v(\Lambda(y))$.

We begin with showing that the corresponding mean dynamics for x_n are given by the regularized replicator dynamics.

Lemma C.1. *If y_t is a solution to (MD), then $x_t = \Lambda(y_t)$ satisfies the regularized replicator dynamics*

$$\dot{x}_{i\alpha_i} = x_{i\alpha_i}\Big[v_{i\alpha_i}(x) - \textstyle\sum_{\beta_i\in\mathcal{A}_{ik}} x_{i\beta_i} v_{i\beta_i}(x)\Big] - \lambda x_{i\alpha_i}\Big[\log x_{i\alpha_i} - \textstyle\sum_{\beta_i\in\mathcal{A}_{ik}} x_{i\beta_i}\log x_{i\beta_i}\Big] \tag{RD$_\lambda$}$$

for all $\alpha_i \in \mathcal{A}_{ik}$, $k = 1,\dots,m_i$, and all $i \in \mathcal{N}$.

Proof. Fix $i \in \mathcal{N}$ and $1 \leq k \leq m_i$. Then, for every $\alpha_i, \beta_i \in \mathcal{A}_{ik}$, we compute

$$\frac{\partial\,\Lambda_{i\alpha_i}(y)}{\partial y_{i\alpha_i}} = \Lambda_{i\alpha_i}(y)\big(1 - \Lambda_{i\alpha_i}(y)\big) \quad\text{and}\quad \frac{\partial\,\Lambda_{i\alpha_i}(y)}{\partial y_{i\beta_i}} = -\,\Lambda_{i\alpha_i}(y)\,\Lambda_{i\beta_i}(y). \tag{C.3}$$

In particular, since $\Lambda_{i\alpha_i}(y)$ depends only on $y_{i\beta_i}$, $\beta_i \in \mathcal{A}_{ik}$, the chain rule applied to $x_t = \Lambda(y_t)$ yields

$$\begin{aligned}
\dot{x}_{i\alpha_i} &= \sum_{\beta_i\in\mathcal{A}_{ik}} \frac{\partial\,\Lambda_{i\alpha_i}(y)}{\partial y_{i\beta_i}}\dot{y}_{i\beta_i} && \text{(C.4)}\\
&= \Lambda_{i\alpha_i}(y)\left[\dot{y}_{i\alpha_i} - \sum_{\beta_i\in\mathcal{A}_{ik}} \Lambda_{i\beta_i}(y)\dot{y}_{i\beta_i}\right] && \text{(C.5)}\\
&= x_{i\alpha_i}\left[v_{i\alpha_i}(x) - \sum_{\beta_i\in\mathcal{A}_{ik}} x_{i\beta_i} v_{i\beta_i}(x)\right] - \lambda x_{i\alpha_i}\left[y_{i\alpha_i} - \sum_{\beta_i\in\mathcal{A}_{ik}} x_{i\beta_i} y_{i\beta_i}\right], && \text{(C.6)}
\end{aligned}$$

and it only remains to note that

$$\begin{aligned}
\log x_{i\alpha_i} - \sum_{\beta_i\in\mathcal{A}_{ik}} x_{i\alpha_i}\log x_{i\alpha_i} &= y_{i\alpha_i} - \log\sum_{\gamma_i\in\mathcal{A}_{ik}} e^{y_{i\gamma_i}} - \sum_{\beta_i\in\mathcal{A}_{ik}} x_{i\beta_i}\Big(y_{i\beta_i} - \log\sum_{\gamma_i\in\mathcal{A}_{ik}} e^{y_{i\gamma_i}}\Big) && \text{(C.7)}\\
&= y_{i\alpha_i} - \sum_{\beta_i\in\mathcal{A}_{ik}} x_{i\beta_i} y_{i\beta_i}, && \text{(C.8)}
\end{aligned}$$

where we have used $\sum_{\beta_i} x_{i\beta_i} = 1$. ∎

As mentioned in the main text, neither the primal trajectories x_n nor the dual process y_n are well-suited to study the long-run behaviors of (REW) since the logit map Λ is not invertible. This is due to the fact that $\mathcal{X}$ effectively lives in a space of lower dimension than $\mathcal{Y}$. For this reason, we begin by reformulating (REW) in this space of lower dimension which *can* be mapped to $\mathcal{X}$ in a one-to-one way.

For every $i \in \mathcal{N}$, $1 \leq k \leq m_i$, we fix an arbitrary baseline coordinate $\hat{\alpha}_{ik} \in \Delta(\mathcal{A}_{ik})$, define $\mathcal{A}_{ik}^* \coloneqq \mathcal{A}_{ik}\backslash\{\hat{\alpha}_{ik}\}$ and introduce the state variables

$$z_{i\alpha_i} = y_{i\alpha_i} - y_{i\hat{\alpha}_{ik}} \quad \text{for all } \alpha_i \in \mathcal{A}_{ik}^* \text{ and every } y \in \mathbb{R}^{|\mathcal{A}_{ik}|}. \tag{C.9}$$

Letting $d_{ik} = |\mathcal{A}_{ik}|$ and $m = \sum_i m_i$, we then define the associated state spaces

$$\mathcal{Z}_{ik} = \mathbb{R}^{d_{ik}-1}, \quad \mathcal{Z}_i = \prod_{k=1}^{m_i}\mathcal{Z}_{ik} = \mathbb{R}^{d_i - m_i}, \quad \mathcal{Z} = \prod_{i\in\mathcal{N}}\mathcal{Z}_i = \mathbb{R}^{d-m}, \tag{C.10}$$

and the *reduced logit map* $Q\colon \mathcal{Z} \to \mathcal{X}$ given in coordinates by the expression

$$Q_{i\alpha_i}(z) = \begin{cases} \dfrac{\exp(z_{i\alpha_i})}{1+\sum_{\beta_i\in\mathcal{A}_{ik}^*}\exp(z_{i\beta_i})} & \text{if } \alpha_i \in \mathcal{A}_{ik}^*,\\[2ex] \dfrac{1}{1+\sum_{\beta_i\in\mathcal{A}_{ik}^*}\exp(z_{i\beta_i})} & \text{if } \alpha_i = \hat{\alpha}_{ik}. \end{cases} \tag{C.11}$$

Unlike the standard Λ map, Q now defines a diffeomorphism.

Lemma C.2. *The reduced logit map Q is a diffeomorphism from $\mathcal{Z}$ to the relative interior* $\operatorname{ri}\mathcal{X}$ *of $\mathcal{X}$.*

Proof. First, note that for any fixed $i \in \mathcal{N}$, $1 \leq k \leq m_i$, $\alpha_i \in \mathcal{A}_{ik}$, $Q_{i\alpha_i}$ depends only on $z_{i\beta_i}$, $\beta_i \in \mathcal{A}_{ik}$. As a consequence, it suffices to prove that Q is a diffeomorphism from $\mathcal{Z}_{ik}$ to the relative interior of the simplex $\mathcal{X}_{ik} \coloneqq \Delta(\mathcal{A}_{ik})$. In particular, we can drop the indices i and k for notational convenience, and can therefore assume that $\mathcal{X}$ is a $(d-1)$-dimensional simplex while $\mathcal{Z} = \mathbb{R}^{d-1}$.

Moreover, from standard topological results, the relative interior $\operatorname{ri}\mathcal{X}$ is diffeomorphic to the interior $\operatorname{int}\tilde{\mathcal{X}}$ of the set $\tilde{\mathcal{X}} := \{x \in \mathbb{R}^{d-1} : x_\alpha \geq 0, \sum_{\beta\in\mathcal{A}^*} x_\beta \leq 1\}$, obtained by projecting $\mathcal{X}$ onto $d-1$ of its coordinates. As such, it only remains to show that the projection of Q onto $d-1$ of its coordinates is a diffeomorphism; that is, we want to establish that the map $\tilde{Q}\colon \mathcal{Z} \to \operatorname{int}\tilde{\mathcal{X}}$ defined as

$$\tilde{Q}_\alpha(z) = \frac{\exp(z_\alpha)}{1 + \sum_{\beta\in\mathcal{A}^*}\exp(z_\beta)} \quad \text{for all } \alpha \in \mathcal{A}^* \tag{C.12}$$

is diffeomorphic. To begin with, note that $\tilde{Q}$ admits a continuous inverse

$$\tilde{Q}^{-1}(\tilde{x}) = \log(\tilde{x}_\alpha) - \log\big(1 - \sum_{\beta\in\mathcal{A}^*}\tilde{x}_\beta\big), \quad \tilde{x} \in \operatorname{int}\tilde{\mathcal{X}}, \tag{C.13}$$

so $\tilde{Q}$ is a bijection from $\mathcal{Z}$ to $\operatorname{int}\tilde{\mathcal{X}}$. Then, we derive

$$\frac{\partial\tilde{Q}_\alpha(z)}{\partial z_\alpha} = \tilde{Q}_\alpha(z)\big(1 - \tilde{Q}_\alpha(z)\big) \quad \text{and} \quad \frac{\partial\tilde{Q}_\alpha(z)}{\partial z_\beta} = -\tilde{Q}_\alpha(z)\tilde{Q}_\beta(z), \tag{C.14}$$

which therefore leads to

$$\operatorname{Jac}\tilde{Q}(z) = \operatorname{diag}(\tilde{Q}(y)) - \tilde{Q}(y)\tilde{Q}(y)^\top. \tag{C.15}$$

The matrix formula $\det(A + uv^\top) = (1 + v^\top A^{-1}u)\det(A)$ can then be used to obtain

$$\begin{aligned}
\det(\operatorname{Jac}\tilde{Q}(z)) &= \big(1 - \tilde{Q}(z)^\top \operatorname{diag}(1/\tilde{Q}(z))\tilde{Q}(z)\big)\det\big(\operatorname{diag}(\tilde{Q}(z))\big) && \text{(C.16)}\\
&= \big(1 - \tilde{Q}(z)^\top\mathbf{1}\big)\prod_{\alpha\in\mathcal{A}^*}\tilde{Q}_\alpha(z) && \text{(C.17)}\\
&= \big(1 - \sum_{\alpha\in\mathcal{A}^*}\tilde{Q}_\alpha(z)\big)\prod_{\alpha\in\mathcal{A}^*}\tilde{Q}_\alpha(z), && \text{(C.18)}
\end{aligned}$$

which is strictly positive since $\tilde{Q}$ maps $\mathcal{Z}$ to $\operatorname{int}\tilde{\mathcal{X}}$. It follows that $\tilde{Q}\colon \mathcal{Z} \to \operatorname{int}\tilde{\mathcal{X}}$ defines a diffeomorphism, hence concluding the proof. ∎

Moreover, the trajectories of play generated by (REW) can now be mapped in a one-to-one correspondence to trajectories in $\mathcal{Z}$.

Lemma C.3. *The primal trajectories x_n of (REW) satisfy $x_n = Q(z_n)$ where*

$$z_{n+1} = z_n + \eta(F^*(z_n) + U_n - \lambda z_n), \tag{RM}$$

with drift F and noise U_n defined as

$$\begin{aligned}
F^*_{i\alpha_i}(z) &= v_{i\alpha_i}(Q(z)) - v_{i\hat{\alpha}_{ik}}(Q(z)) && \text{(C.19)}\\
U_{i\alpha_i,n} &= U_{i\alpha_i}(z_n,\omega_n) = \mathsf{Z}_{i\alpha_i}(Q(z_n);\omega_n) - \mathsf{Z}_{i\hat{\alpha}_{ik}}(Q(z_n);\omega_n)
\end{aligned}$$

for every $\alpha_i \in \mathcal{A}_{i,k}$, $1 \leq k \leq m_i$ and $i \in \mathcal{N}$.

Equation (RM) is the main dynamics that is studied in the following appendices. More precisely, it corresponds to Eq. (SD) in Appendix D.1.

Proof. Let x_n and y_n denote the primal and dual trajectories of (REW), respectively. We then define an auxiliary discrete process z_n such that

$$z_{i\alpha_i,n} = y_{i\alpha_i,n} - y_{i\hat{\alpha}_{ik},n} \quad \text{for all } \alpha_i \in \mathcal{A}^*_{ik},\ 1 \leq k \leq m_i, i \in \mathcal{N}. \tag{C.20}$$

Fix $i \in \mathcal{N}$ and $1 \leq k \leq m_i$. Note that by definition of Λ and Q, we have

$$\begin{aligned}
Q_{i\alpha_i}(z_n) = \frac{\exp(z_{i\alpha_i,n})}{1 + \sum_{\beta_i\in\mathcal{A}^*_{ik}}\exp(z_{i\beta_i,n})} &= \frac{\exp(y_{i\alpha_i,n})\exp(-y_{i\hat{\alpha}_{ik},n})}{1 + \exp(-y_{i\hat{\alpha}_{ik},n})\sum_{\beta_i\in\mathcal{A}^*_{ik}}\exp(y_{i\beta_i,n})} && \text{(C.21)}\\
&= \frac{\exp(y_{i\alpha_i,n})}{\sum_{\beta\in\mathcal{A}_{ik}}\exp(y_{i\beta_i,n})} && \text{(C.22)}\\
&= \Lambda_{i\alpha_i}(y_n) && \text{(C.23)}
\end{aligned}$$

for any $\alpha_i \in \mathcal{A}^*_{ik}$. Similarly, for $\alpha_i = \hat{\alpha}_{ik}$,

$$Q_{i\hat{\alpha}_{ik}}(z_n) = \frac{1}{1 + \sum_{\beta\in\mathcal{A}^*_{ik}}\exp(z_{i\beta_i,n})} = \frac{\exp(y_{i\hat{\alpha}_{ik},n})}{\sum_{\beta\in\mathcal{A}_{ik}}\exp(y_{i\beta_i,n})} = \Lambda_{i\hat{\alpha}_{ik}}(y_n). \tag{C.24}$$

In particular, this means that $x_n = \Lambda(y_n) = Q(z_n)$. As a result, we can also derive

$$\begin{aligned}
&z_{i\alpha_i,n+1} \\
&= y_{i\alpha_i,n+1} - y_{i\hat{\alpha}_{ik},n+1} \\
&= y_{i\alpha_i,n} - y_{i\hat{\alpha}_{ik},n} + \eta\big[v_{i\alpha_i}(x_n) - v_{i\hat{\alpha}_{ik}}(x_n) + \mathsf{Z}_{i\alpha_i}(x_n;\omega_n) - \mathsf{Z}_{i\hat{\alpha}_{ik}}(x_n;\omega_n - \lambda\big(y_{i\alpha_i,n} - y_{i\hat{\alpha}_{ik},n}\big)\big] \\
&= z_{i\alpha_i,n} + \eta\big[F^*_{i\alpha_i}(z_n) + U_{i\alpha_i,n} - \lambda z_{i\alpha_i,n}\big], \qquad \text{(C.25)}
\end{aligned}$$

which concludes the proof. ∎

We collect below some regularity properties of the drift and noise terms.

Lemma C.4. *Let $F(z) = F^*(z) - \lambda z$.*

(*a*) *The vector field F is C^1-differentiable and Lipschitz-continuous.*

(*b*) *Dissipativity: There exists a constant $M > 0$ such that $\langle z, F(z)\rangle \le -\frac{\lambda}{2}\|z\|^2 + \frac{2M^2}{\lambda}$ for all $z \in \mathcal{Z}$.*

(*c*) *Bounded noise: The noise function U is C^2-differentiable and $\sup_{z\in\mathcal{Z},\omega\in\Omega}\|U(z,\omega)\| < \infty$.*

(*d*) *Properness: $\mathbb{E}[U(z,\omega)] = 0$ and $\operatorname{cov}(U(z,\omega)) \succ 0$ for all $z \in \mathcal{Z}$.*

Proof. (a) Follows immediately since v and Q are both C^1-differentiable and Lipschitz-continuous. (b) Since $Q(z)$ belongs to the compact set $\mathcal{X}$, $F^*(y)$ is uniformly bounded by some constant $M > 0$. We therefore get $\langle z, F(z)\rangle \le -\lambda\|z\|^2 + M\|z\|$. If $\|z\| \ge 2M/\lambda$ then we have $\langle z, F(z)\rangle \le -\lambda\|z\|^2/2$, while $\langle z, F(z)\rangle \le -\lambda\|z\|^2 + 2M^2/\lambda$ when $\|z\| \le 2M/\lambda$. Combining both cases finally yields the result. (c) Follows from the fact that $Q(z)$ belongs to $\mathcal{X}$ compact and that Q and Z are both C^2-smooth. (d) This is implied by the properness of Z, see Assumption 1. ∎

Similarly to what we have done for the discrete trajectories y_n, the mean behavior of z_n can then be described through the continuous-time dynamics

$$\dot{z} = F^*(z) - \lambda z = F(z). \qquad \text{(MRM)}$$

Since these dynamics are mapped to those of x_t through the diffeomorphism Q, they will be used extensively in the sequel in order to obtain quantitative properties on x_n.

In particular, we note that the vector field F generates a unique *flow* $\Phi^{\mathcal{Z}}$ on $\mathcal{Z}$ given by $\Phi^{\mathcal{Z}}_t(z_0) \coloneqq z_t$. Similarly, the trajectories of (RRD) define a flow $\Phi^{\mathcal{X}}$ on $\operatorname{ri}\mathcal{X}$ by $\Phi^{\mathcal{X}}_t(x_0) \coloneqq x_t$. For a rigorous definition of flows and their properties, we refer to Appendix E.1. We can then define the notion of *internally chain transitive* sets following Benaïm [13]:

Definition C.1 (Internal chain transitivity)**.** Let $\mathcal{V}$ be a vector space and Φ a flow on $\mathcal{V}$. A compact set $\mathcal{K} \subseteq \mathcal{V}$ invariant under Φ is said to be internally chain transitive (ICT) for Φ if, for every $x, y \in \mathcal{K}$ and all $\varepsilon > 0$, $T > 0$, there exists a finite sequence $x_0, \dots, x_m \in \mathcal{K}$ such that $x_0 = x$, $x_m = y$ and $t_1, \dots, t_m \ge T$ such that

$$\max_{1\le i\le m} \|\Phi_{t_i}(x_{i-1}) - x_i\| < \varepsilon. \qquad \text{(C.26)}$$

Equivalently, a compact set $\mathcal{K} \subseteq \mathcal{V}$ is internally chain transitive if and only if it is invariant and such that $\Phi|_{\mathcal{K}}$ admits no proper attractors in $\mathcal{K}$ [13, Proposition 5.3].

We conclude this section by establishing that the ICT sets with respect to $\Phi^{\mathcal{X}}$ are in a one-to-one correspondence with those of $\Phi^{\mathcal{Z}}$.

Theorem C.1. *A set $\mathcal{A} \subseteq \operatorname{ri}\mathcal{X}$ is internally chain transitive for $\Phi^{\mathcal{X}}$ if and only if $Q^{-1}(\mathcal{A})$ is internally chain transitive for $\Phi^{\mathcal{Z}}$.*

Proof. The properties established in Lemma C.4 imply that the discrete dynamics (RM) satisfy the abstract framework of Appendix D. In particular, one obtains from Proposition E.3 that the ICTs of $\Phi^{\mathcal{X}}$ and $\Phi^{\mathcal{Z}}$ are mapped through the diffeomorphism Q. ∎

C.3. Main Results. In this section, we now restate the main results of the main text and indicate where their proofs can be found in the appendix.

Theorem 1. *Let $\mathcal{M}^\delta$ be a δ-neighborhood of the metastable region $\mathcal{M} := \bigcup_i \mathcal{C}_i$ of the system. Then there exists some $c \equiv c_\delta > 0$ such that, for small enough $\eta > 0$ and large enough n, we have*

$$\pi_n(\mathcal{M}^\delta) \geq 1 - e^{-c/\eta}. \tag{17}$$

Theorem 2. *Every metastable set $\mathcal{C}_i$, $i = 1, \dots, K$, of (REW) is an ICT set of (RRD).*

Theorem 3. *Fix a tolerance level $\varepsilon > 0$, let $\mathcal{C}_i, \mathcal{C}_j$ be two metastable components of the system, and let $\mathcal{C}_i^\delta, \mathcal{C}_j^\delta$ be δ-neighborhoods thereof. Then, for small enough $\delta, \eta > 0$ and large enough n, we have:*

$$\left|\eta \log \pi_n(\mathcal{C}_i^\delta) + E_i - E_0\right| \leq \varepsilon \quad \textit{and} \quad \left|\eta \log\big(\pi_n(\mathcal{C}_i^\delta)/\pi_n(\mathcal{C}_j^\delta)\big) + E_i - E_j\right| \leq \varepsilon\,, \tag{BG}$$

or, more concisely, $\pi_n(\mathcal{C}_i^\delta) \propto \exp(-(E_i + \mathcal{O}(\varepsilon))/\eta)$.

Theorem 4. *Let $\mathcal{C}_i$ be a metastable component that is not Lyapunov stable. Then there exists a metastable component $\mathcal{C}_j$ that is Lyapunov stable and a constant $c \equiv c_{ij} > 0$ such that*

$$\pi_n(\mathcal{C}_i^\delta)/\pi_n(\mathcal{C}_j^\delta) \leq \exp(-c/\eta) \tag{19}$$

for all sufficiently small $\delta, \eta > 0$ and large enough n.

Theorem 5. *Let $\mathcal{C}_0^\delta$ be a δ-neighborhood of the ground state $\mathcal{C}_0$. Then there exists a constant $c \equiv c_\delta > 0$ such that, for all sufficiently small $\delta, \eta > 0$ and large enough n, we have $\pi_n(\mathcal{C}_0^\delta) \geq 1 - e^{-c/\eta}$.*

Their proofs are obtained by combining the properties of the dynamics with the general results on stochastic approximation processes developed in the following appendices. More precisely, (RM) corresponds to (SD) in Appendix D.1, and the properties of the drift and noise terms established in Lemma C.4 ensures that (RM) satisfies the assumptions of Appendix D (with Assumption 5 being implied by Assumption 2 by definition). In particular, the *metastable components* introduced in the main text (Definition 1) coincide with the equivalence classes of Appendix E.3 (see also Remark E.1) mapped through Q.

Then, we invoke the results of Section G.3.3 with the diffeomorphism Q (Lemma C.2). Equation (G.128) from Section G.3.3 is now exactly the process $(x_n)_n$ from (RRD) and this yields the proofs of the main results. Also note that the transition costs (16) exactly match the ones defined in the following sections (see Eq. (F.74) in Appendix F.3).

More precisely, each of the main theorems is obtained as follows:

Proof of Theorem 1. We apply Theorem G.11 to Eq. (RM) along with the diffeomorphism Q. The result of Theorem G.11 thus directly applies to the process x_n and its occupation measures Eq. (8). ∎

Proof of Theorem 2. Since the equivalence classes of (RM) are internally chain transitive by Theorem E.2, it follows from Theorem C.1 that the metastable components are internally chain transitive themselves. ∎

Proof of Theorem 3. We apply Theorem G.9 to Eq. (RM) along with the diffeomorphism Q. The result of Theorem G.9 thus directly applies to the process x_n and its occupation measures Eq. (8). ∎

Proof of Theorem 4. We apply Theorem G.10 to Eq. (RM) along with the diffeomorphism Q. The result of Theorem G.10 thus directly applies to the process x_n and its occupation measures Eq. (8). Note that $\widetilde{\mathcal{K}}_i$ unstable under (RRD) is equivalent to $\widetilde{\mathcal{K}}_i$ being unstable under Eq. (G.128) since (RRD) exactly corresponds to the lifted dynamics of Eq. (G.128). ∎

Proof of Theorem 5. We apply Theorem G.12 to Eq. (RM) along with the diffeomorphism Q. The result of Theorem G.12 thus directly applies to the process x_n and its occupation measures Eq. (8). ∎

C.4. Potential games in reduced coordinates. We now record how the potential-game structure is expressed in the reduced coordinates introduced in Appendix C.

In this section, we thus assume that $\mathcal{G}$ is a potential game with potential f (cf. Eq. (2)).

Lemma C.5. *For every $x \in \operatorname{ri}\mathcal{X}$, every $i \in \mathcal{N}$, every $1 \le k \le m_i$, and every vector $w \in \mathbb{R}^{|\mathcal{A}_{ik}|}$ such that*

$$\sum_{\alpha_i \in \mathcal{A}_{ik}} w_{\alpha_i} = 0, \tag{C.27}$$

we have

$$\sum_{\alpha_i \in \mathcal{A}_{ik}} v_{i\alpha_i}(x) w_{\alpha_i} = \sum_{\alpha_i \in \mathcal{A}_{ik}} \partial_{x_{i\alpha_i}} f(x) w_{\alpha_i} . \tag{C.28}$$

Equivalently, for every $x \in \operatorname{ri}\mathcal{X}$, $i \in \mathcal{N}$, and $1 \le k \le m_i$, there exists a scalar $c_{ik}(x) \in \mathbb{R}$ such that

$$v_{i\alpha_i}(x) - \partial_{x_{i\alpha_i}} f(x) = c_{ik}(x) \qquad \text{for all } \alpha_i \in \mathcal{A}_{ik}. \tag{C.29}$$

In particular, for every $\alpha_i \in \mathcal{A}_{ik}^*$,

$$v_{i\alpha_i}(x) - v_{i\hat{\alpha}_{ik}}(x) = \partial_{x_{i\alpha_i}} f(x) - \partial_{x_{i\hat{\alpha}_{ik}}} f(x) . \tag{C.30}$$

Proof. Fix $i \in \mathcal{N}$, $1 \le k \le m_i$, and $x \in \operatorname{ri}\mathcal{X}$. We consider a unilateral tangent direction supported on the block $\mathcal{A}_{ik}$, so let $w \in \mathbb{R}^{|\mathcal{A}_{ik}|}$ satisfy $\sum_{\alpha_i} w_{\alpha_i} = 0$, and extend it by zero to the other coordinates of player i. Since $x \in \operatorname{ri}\mathcal{X}$, the unilateral perturbation $(x_i + \epsilon w; x_{-i})$ belongs to $\mathcal{X}$ for all ϵ in a neighborhood of zero. Since the payoff functions are smooth, the potential identity may be differentiated along feasible unilateral directions. Applying the potential-game identity (2) to x and $(x_i + tw; x_{-i})$ gives

$$u_i(x_i + \epsilon w; x_{-i}) - u_i(x) = f(x_i + \epsilon w; x_{-i}) - f(x). \tag{C.31}$$

Dividing by ϵ and letting $\epsilon \to 0$ yields (C.28).

Now, (C.28) means that the vector

$$v_{i\alpha_i}(x) - \partial_{x_{i\alpha_i}} f(x) \,, \alpha_i \in \mathcal{A}_{ik} \tag{C.32}$$

is orthogonal to the set of vectors

$$\left\{ w \in \mathbb{R}^{|\mathcal{A}_{ik}|} : \sum_{\alpha_i \in \mathcal{A}_{ik}} w_{\alpha_i} = 0 \right\}. \tag{C.33}$$

Since the orthogonal complement of the set (C.33) is the span of the all-ones vector on this block, we conclude that there exists a scalar $c_{ik}(x)$ such that (C.29) holds.

Finally, subtracting the coordinate $\hat{\alpha}_{ik}$ from the coordinate $\alpha_i \in \mathcal{A}_{ik}^*$ gives (C.30), which concludes the proof. ∎

Lemma C.6. *Let $y \in \mathbb{R}^d$ and set $x = \Lambda(y)$. Then, for every $i \in \mathcal{N}$, every $1 \le k \le m_i$, and every $\alpha_i \in \mathcal{A}_{ik}^*$,*

$$\begin{aligned} V_{i\alpha_i}(y) - V_{i\hat{\alpha}_{ik}}(y) &= v_{i\alpha_i}(x) - v_{i\hat{\alpha}_{ik}}(x) \\ &= \partial_{x_{i\alpha_i}} f(x) - \partial_{x_{i\hat{\alpha}_{ik}}} f(x) . \end{aligned} \tag{C.34}$$

Moreover, let

$$\mathcal{X}' := \prod_{i \in \mathcal{N}} \prod_{k=1}^{m_i} \left\{ x_{ik} \in \mathbb{R}^{|\mathcal{A}_{ik}^*|} : x_{i\alpha_i} > 0 \text{ for all } \alpha_i \in \mathcal{A}_{ik}^*, \quad \sum_{\alpha_i \in \mathcal{A}_{ik}^*} x_{i\alpha_i} < 1 \right\} \tag{C.35}$$

be the product of projected simplex interiors, and define $\iota \colon \mathcal{X}' \to \operatorname{ri}\mathcal{X}$ by

$$\iota_{i\alpha_i}(x) = \begin{cases} x_{i\alpha_i}, & \text{if } \alpha_i \in \mathcal{A}_{ik}^*, \\ 1 - \displaystyle\sum_{\beta_i \in \mathcal{A}_{ik}^*} x_{i\beta_i}, & \text{if } \alpha_i = \hat{\alpha}_{ik}. \end{cases} \tag{C.36}$$

Define $\tilde{f} \colon \mathcal{X}' \to \mathbb{R}$ by

$$\tilde{f}(x) = f(\iota(x)). \tag{C.37}$$

Then, for every $x \in \mathcal{X}'$, every $i \in \mathcal{N}$, every $1 \le k \le m_i$, and every $\alpha_i \in \mathcal{A}_{ik}^*$,

$$\partial_{x_{i\alpha_i}} \tilde{f}(x) = \partial_{x_{i\alpha_i}} f(\iota(x)) - \partial_{x_{i\hat{\alpha}_{ik}}} f(\iota(x)). \tag{C.38}$$

Proof. The first identity follows directly from the definition $V(y) = v(\Lambda(y))$. The second one is exactly the identity (C.30) of Lemma C.5 applied at $x = \Lambda(y)$.

It remains to verify the reduced-coordinate formulation. Fix $x \in \mathcal{X}'$, $i \in \mathcal{N}$, $1 \le k \le m_i$, and $\alpha_i \in \mathcal{A}^*_{ik}$. By the definition (C.37) of $\tilde{f}$ and the chain rule, we have

$$\partial_{x_{i\alpha_i}} \tilde{f}(x) = \sum_{j \in \mathcal{N}} \sum_{l=1}^{m_j} \sum_{\beta_j \in \mathcal{A}_{jl}} \partial_{x_{j\beta_j}} f(\iota(x)) \partial_{x_{i\alpha_i}} \iota_{j\beta_j}(x). \tag{C.39}$$

Moreover, by the definition (C.36) of ι,

$$\partial_{x_{i\alpha_i}} \iota_{j\beta_j}(x) = \begin{cases} 1, & \text{if } j = i \text{ and } \beta_j = \alpha_i, \\ -1, & \text{if } j = i,\ l = k, \text{ and } \beta_j = \hat{\alpha}_{ik}, \\ 0, & \text{otherwise.} \end{cases} \tag{C.40}$$

Substituting this expression in (C.39) gives

$$\partial_{x_{i\alpha_i}} \tilde{f}(x) = \partial_{x_{i\alpha_i}} f(\iota(x)) - \partial_{x_{i\hat{\alpha}_{ik}}} f(\iota(x)), \tag{C.41}$$

which is (C.38) and concludes the proof. ∎

Lemma C.7. *With the notation of Lemma C.6, , for every $z \in \mathcal{Z}$, for every $i \in \mathcal{N}$, every $1 \le k \le m_i$, and every $\alpha_i \in \mathcal{A}^*_{ik}$,*

$$F^*_{i\alpha_i}(z) = \partial_{x_{i\alpha_i}} \tilde{f}(\tilde{Q}(z)) \,. \tag{C.42}$$

Proof. Fix $z \in \mathcal{Z}$, $i \in \mathcal{N}$, $1 \le k \le m_i$, and $\alpha_i \in \mathcal{A}^*_{ik}$. By definition of the reduced drift,

$$F^*_{i\alpha_i}(z) = v_{i\alpha_i}(Q(z)) - v_{i\hat{\alpha}_{ik}}(Q(z)). \tag{C.43}$$

Combining this expression with Eqs. (C.30) and (C.38) and using that $\iota \circ \tilde{Q} = Q$ gives (C.42). ∎

Lemma C.8. *With the notations of Lemma C.6, let $\tilde{Q} \colon \mathcal{Z} \to \mathcal{X}'$ denote the projection of Q on the non-exclusion coordinates, and define $\Phi_{\text{logit}} \colon \mathcal{Z} \to \mathbb{R}$ by*

$$\Phi_{\text{logit}}(z) := \sum_{i \in \mathcal{N}} \sum_{k=1}^{m_i} \log\left(1 + \sum_{\alpha_i \in \mathcal{A}^*_{ik}} \exp(z_{i\alpha_i})\right). \tag{C.44}$$

Then Φ_{logit} is C^∞ and strictly convex, and

$$\nabla \Phi_{\text{logit}}(z) = \tilde{Q}(z) \qquad \text{for all } z \in \mathcal{Z}. \tag{C.45}$$

Moreover, the restriction of its convex conjugate Φ^*_{logit} to $\mathcal{X}'$ is given by

$$\Phi^*_{\text{logit}}(x) = \sum_{i \in \mathcal{N}} \sum_{k=1}^{m_i} \Bigg(\sum_{\alpha_i \in \mathcal{A}^*_{ik}} x_{i\alpha_i} \log x_{i\alpha_i} + \left(1 - \sum_{\alpha_i \in \mathcal{A}^*_{ik}} x_{i\alpha_i}\right) \log\left(1 - \sum_{\alpha_i \in \mathcal{A}^*_{ik}} x_{i\alpha_i}\right)\Bigg). \tag{C.46}$$

In particular, Φ^*_{logit} is C^1 on $\mathcal{X}'$ and

$$\nabla \Phi^*_{\text{logit}}(x) = \tilde{Q}^{-1}(x) \qquad \text{for all } x \in \mathcal{X}'. \tag{C.47}$$

The proof is standard, but we give the details to make the role of the exclusion action explicit. In particular, the computation is just the usual duality between the log-sum-exp function and the negative entropy, written in non-exclusion coordinates.

Proof. The statement is block-separable, so it is enough to fix $i \in \mathcal{N}$ and $1 \le k \le m_i$. In this proof, we drop the corresponding indices for notational convenience. For every $\alpha \in \mathcal{A}^*$, we have

$$\partial_{z_\alpha} \Phi_{\mathrm{logit}}(z) = \frac{\exp(z_\alpha)}{1 + \sum_{\beta \in \mathcal{A}^*} \exp(z_\beta)} = \tilde{Q}_\alpha(z), \tag{C.48}$$

which gives (C.45). Moreover, writing $z = \tilde{Q}(z)$, we get

$$\nabla^2 \Phi_{\mathrm{logit}}(z) = \mathrm{diag}(z) - zz^\top. \tag{C.49}$$

Thus, for every $u \in \mathbb{R}^{|\mathcal{A}^*|}$,

$$u^\top \nabla^2 \Phi_{\mathrm{logit}}(z) u = \sum_{\alpha \in \mathcal{A}^*} z_\alpha u_\alpha^2 - \left(\sum_{\alpha \in \mathcal{A}^*} z_\alpha u_\alpha \right)^2. \tag{C.50}$$

Since $z_\alpha > 0$ for all $\alpha \in \mathcal{A}^*$ and $\sum_{\alpha \in \mathcal{A}^*} z_\alpha < 1$, the Cauchy–Schwarz inequality gives

$$\left(\sum_{\alpha \in \mathcal{A}^*} z_\alpha u_\alpha \right)^2 \le \left(\sum_{\alpha \in \mathcal{A}^*} z_\alpha \right) \left(\sum_{\alpha \in \mathcal{A}^*} z_\alpha u_\alpha^2 \right) < \sum_{\alpha \in \mathcal{A}^*} z_\alpha u_\alpha^2 \tag{C.51}$$

whenever $u \neq 0$. Therefore $\nabla^2 \Phi_{\mathrm{logit}}(z)$ is positive definite for every $z \in \mathcal{Z}$, which shows that Φ_{logit} is strictly convex.

It remains to identify the conjugate. Fix $x \in \mathcal{X}'$ and set

$$x_{\hat{\alpha}} = 1 - \sum_{\alpha \in \mathcal{A}^*} x_\alpha. \tag{C.52}$$

Since $\tilde{Q}$ is a diffeomorphism from $\mathcal{Z}$ to $\mathcal{X}'$, there exists a unique $z \in \mathcal{Z}$ such that $\tilde{Q}(z) = x$, namely

$$z_\alpha = \log x_\alpha - \log x_{\hat{\alpha}} \qquad \text{for all } \alpha \in \mathcal{A}^*. \tag{C.53}$$

At this point, the map

$$z' \mapsto \langle z', x \rangle - \Phi_{\mathrm{logit}}(z') \tag{C.54}$$

has vanishing gradient. Since Φ_{logit} is strictly convex, this point is the unique maximizer in the definition of $\Phi^*_{\mathrm{logit}}(x)$. Moreover,

$$\Phi_{\mathrm{logit}}(z) = \log\left(1 + \sum_{\alpha \in \mathcal{A}^*} \frac{x_\alpha}{x_{\hat{\alpha}}}\right) = -\log x_{\hat{\alpha}}. \tag{C.55}$$

Consequently,

$$\begin{aligned} \Phi^*_{\mathrm{logit}}(x) &= \langle z, x \rangle - \Phi_{\mathrm{logit}}(z) \\ &= \sum_{\alpha \in \mathcal{A}^*} x_\alpha \log x_\alpha + x_{\hat{\alpha}} \log x_{\hat{\alpha}}. \end{aligned} \tag{C.56}$$

Summing the latter identity over all blocks gives (C.46).

Finally, differentiating this expression gives, for every $\alpha \in \mathcal{A}^*$,

$$\partial_{x_\alpha} \Phi^*_{\mathrm{logit}}(x) = \log x_\alpha - \log x_{\hat{\alpha}} = \tilde{Q}_\alpha^{-1}(x). \tag{C.57}$$

This also shows that Φ^*_{logit} is C^1 on $\mathcal{X}'$, which concludes the proof. ∎

Lemma C.9. *The full drift*

$$F(z) = F^*(z) - \lambda z\,, \tag{C.58}$$

also satisfies, for every $i \in \mathcal{N}$, every $1 \le k \le m_i$, and every $\alpha_i \in \mathcal{A}^*_{ik}$,

$$F_{i\alpha_i}(z) = \partial_{x_{i\alpha_i}} \tilde{f}_\lambda(\tilde{Q}(z)) \tag{C.59}$$

where, for every $x \in \mathcal{X}'$,

$$\tilde{f}_\lambda(x) \coloneqq \tilde{f}(x) - \lambda \Phi^*_{\mathrm{logit}}(x)\,. \tag{C.60}$$

Note Eq. (C.60) is reminiscent of Eq. ($\mathcal{G}^\lambda$).

Proof. This follows from Lemma C.7 and Lemma C.8 using that

$$\nabla \Phi^*_{\mathrm{logit}}(\tilde{Q}(z)) = \tilde{Q}^{-1}(\tilde{Q}(z)) = z. \tag{C.61}$$

∎

C.5. Potential games with truncated Gaussian noise. We now precisely state the result mentioned in Section 4.3. We assume that noise Eq. (11) in Eq. (RM) is Gaussian and obtain explicit expressions for the energies of the metastable components. More precisely, as it will be made formal in Appendix D, for technical reasons, we require the noise to be bounded so we consider Gaussian noise truncated to a ball of large enough radius. The most general statement is given in Theorem C.2, and the more specific statement from the main text is given in Corollary C.1.

Theorem C.2. *Assume that, for every $z \in \mathcal{Z}$,*

$$U(z,\omega) \sim N(0,\Sigma(z)) \quad \text{conditioned on} \quad U(z,\omega) \in \mathbb{B}(0,R(z)), \tag{C.62}$$

where $\Sigma(z)$ is positive definite, and where Σ and R are continuous. Let $\sigma\colon \mathbb{R} \to (0,+\infty)$ be continuous, and define

$$g(z) \coloneqq 2\alpha\big(\tilde{f}_\lambda(\tilde{Q}(z))\big), \qquad \alpha' = \sigma^{-2}. \tag{C.63}$$

There exists a compact $\mathcal{K} \subseteq \mathcal{Z}$ depending explicitly on Σ, R, and $\tilde{f}_\lambda$ such that the following holds.

Denote by g_i the value taken by g on $\mathcal{K}_i$ and define

$$\delta \coloneqq \sup_{z\in\mathcal{K}} \frac{e^{-\frac{R(z)^2}{16\|\Sigma(z)\|}}(\operatorname{tr}\Sigma(z) + \|\Sigma(z)\|)2^{d-m+3}}{\lambda_{\min}(\Sigma(z))}, \tag{C.64}$$

$$\zeta \coloneqq \sup_{z\in\mathcal{K}} \Big\| I - \Sigma^{1/2}(z)\sigma^{-2}\big(\tilde{f}_\lambda(\tilde{Q}(z))\big)\operatorname{Jac}\tilde{Q}(z)\Sigma^{1/2}(z)\Big\|. \tag{C.65}$$

Assume that $0 < \delta \le \frac{1}{2}$, that $0 \le \zeta < 1$, and that, for every $z \in \mathcal{K}$,

$$\sqrt{\|\Sigma(z)\|\|\Sigma(z)^{-1}\|}\|F(z)\| \le \frac{R(z)}{8}. \tag{C.66}$$

Finally, set

$$Q_{\mathcal{K}} \coloneqq \max_{i,j\in\mathcal{I}} B_{i,j}, \qquad M_{\mathcal{K}} \coloneqq \max_{i,j\in\mathcal{I}} |g_i - g_j|, \tag{C.67}$$

and

$$\Delta_{\mathcal{K}} \coloneqq (|\mathcal{I}|-1)\frac{2(\delta+\zeta)Q_{\mathcal{K}} + (\delta+\zeta-\delta\zeta)M_{\mathcal{K}}}{(1-\delta)(1-\zeta)}. \tag{C.68}$$

Then, for every $i \in \mathcal{I}$,

$$E_i = c_{\mathcal{K}} - g_i + r_i, \qquad 0 \le r_i \le \Delta_{\mathcal{K}}, \tag{C.69}$$

where $c_{\mathcal{K}}$ is independent of i.

Proof. By Lemma C.9, for every $z \in \mathcal{Z}$,

$$F(z) = \nabla\tilde{f}_\lambda(\tilde{Q}(z)). \tag{C.70}$$

Moreover, by Lemma C.8, the map $\tilde{Q}$ is the gradient of the strictly convex function Φ_{logit}. Hence, in the notation of Appendix H.4, the reduced vector field is of the form

$$F = \nabla\tilde{f}_\lambda \circ \tilde{Q} = \nabla\tilde{f}_\lambda \circ \nabla\Phi_{\text{logit}}. \tag{C.71}$$

We may therefore apply Theorem H.1 with $x \leftarrow z$, $f \leftarrow \tilde{f}_\lambda$, and $Q \leftarrow \tilde{Q}$. The assumption on U is exactly the truncated Gaussian noise assumption required there, and the correspondence developed in the previous section identifies the equivalence classes, quasipotentials, and invariant potentials in the reduced game coordinates. This gives the result. ∎

We instantiate this general statement with the setting from the main text Section 4.3. To follow the notation of Section 4.3, we define $f_\lambda \coloneqq f - \lambda h$.

Corollary C.1. *Assume that, for every $z \in \mathcal{Z}$,*

$$U(z,\omega) \sim N(0,\Sigma(Q(z))) \quad \text{conditioned on} \quad U(z,\omega) \in \mathbb{B}(0,R(z)), \tag{C.72}$$

where Σ is positive definite, and where Σ and R are continuous.

There exists a compact $\mathcal{K} \subseteq \mathcal{Z}$ depending explicitly on Σ, R, and $\tilde{f}_\lambda$ such that the following holds.

Denote by g_i the value taken by g on $\mathcal{K}_i$ and define

$$\delta \coloneqq \sup_{\mathcal{K}} \frac{e^{-\frac{R^2}{16\|\Sigma\circ Q\|}}\,(\operatorname{tr}\Sigma\circ Q + \|\Sigma\circ Q\|)2^{d-m+3}}{\lambda_{\min}(\Sigma\circ Q)}, \tag{C.73}$$

$$\zeta \coloneqq \sup_{Q(\mathcal{K})} \left\|\Sigma^{-1} - \operatorname{Jac}\tilde{Q}\right\|. \tag{C.74}$$

Assume that $0 < \delta \le \frac{1}{2}$, that $0 \le \zeta < \sup_{\mathcal{K}}\|\Sigma\|^{-1}$, and that, for every $z \in \mathcal{K}$,

$$\sqrt{\|\Sigma(Q(z))\|\|\Sigma(Q(z))^{-1}\|}\|\nabla f_\lambda(Q(z))\| \le \frac{R(z)}{8}. \tag{C.75}$$

Finally, set

$$Q_{\mathcal{K}} \coloneqq \max_{i,j} P(\mathcal{C}_i, \mathcal{C}_j)\,, \qquad M_{\mathcal{K}} \coloneqq \max_{i,j}|f_\lambda(\mathcal{C}_i) - f_\lambda(\mathcal{C}_j)|, \tag{C.76}$$

and

$$\Delta_{\mathcal{K}} \coloneqq K\frac{2(\delta+\zeta)Q_{\mathcal{K}} + (\delta+\zeta-\delta\zeta)M_{\mathcal{K}}}{(1-\delta)(1-\zeta)}. \tag{C.77}$$

Then, for every $i \in \mathcal{I}$,

$$E_i = -f_\lambda(\mathcal{C}_i) + c_{\mathcal{K}} + r_i, \qquad 0 \le r_i \le \Delta_{\mathcal{K}}, \tag{C.78}$$

where $c_{\mathcal{K}}$ is independent of i.

Note that in Section 4.3 in the main text, the truncation error δ is omitted for simplicity. Also, in the main text, the covariance mismatch error ζ was informally written in terms of the covariance matrix of Z and Jac Q rather than in terms of Σ^{-1} and Jac $\tilde{Q}$ for clarity.

Proof. We instantiate Theorem C.2 with $\sigma \equiv 1$ so that $\alpha(u) = u$ and $g = \tilde{f}_\lambda$. We have that

$$\tilde{f}_\lambda(\tilde{Q}(z)) = \tilde{f}(\tilde{Q}(z)) - \lambda\Phi^*_{\text{logit}}(\tilde{Q}(z)) \tag{C.79}$$

$$= f(Q(z)) - \lambda h(Q(z)) \tag{C.80}$$

$$= f_\lambda(Q(z)) \tag{C.81}$$

and therefore the value of $\tilde{f}_\lambda \circ \tilde{Q}$ on $\mathcal{K}_i$ is exactly $f_\lambda(\mathcal{C}_i)$ since $\mathcal{C}_i = Q(\mathcal{K}_i)$. ∎

Towards the general setup. In this section, we have shown how the game theoretic setting of the main text can be abstracted to the general setup of Appendix D and how the main results are thus obtained by applying the results of Section G.3.3 and Appendix H.4 to the process $(x_n)_n$ defined by (RRD). (RM) indeed corresponds to (SD) in Appendix D.1, and the properties of the drift and noise terms established in Lemma C.4 ensure that (RM) satisfies the assumptions of Appendix D. We now proceed to present our general setup and to prove our main results in this setting.

D Setup and preliminary results

D.1 Setup 29
D.2 Blanket assumptions and related definitions 29
D.3 Basic properties of the process 30
D.4 Basic properties of Hamiltonians and Lagrangians 31
D.5 Large deviation principle 31

In this section, we present the abstract setup mentioned in Appendices C.2 and C.5. This general setup allows us to abstract away from the technicalities of the game theory setting and to focus on the core of the problem, which is the analysis of the stochastic dynamics and their large-deviations properties.

D.1. Setup. Throughout the sequel, we will write $\langle \cdot, \cdot \rangle$ for the standard inner product on $\mathcal{V} \equiv \mathbb{R}^d$ and $\|\cdot\|$ for the induced (Euclidean) norm. To lighten notation, we will identify $\mathcal{V}$ with its dual $\mathcal{V}^*$, and we will not formally distinguish between primal and dual vectors (though the distinction should be clear from the context). We will also write $\mathbb{B}(x, r)$ (resp. $\bar{\mathbb{B}}(x, r)$) for the open (resp. closed) ball of radius r centered at $x \in \mathcal{V}$, and we will respectively denote the open and closed δ-neighborhoods of $S \subseteq \mathcal{V}$ as

$$\mathcal{U}_\delta(S) := \{x \in \mathcal{V} : d(S, x) < \delta\} \tag{D.1a}$$

$$S_\delta := \mathrm{cl}(\mathcal{U}_\delta(S)) = \{x \in \mathcal{V} : d(S, x) \leq \delta\} \tag{D.1b}$$

The dynamics that we are interested in is

$$\begin{cases} x_0 \in \mathcal{V} \\ x_{n+1} = x_n + \eta F(x_n) + \eta U_n\,, \qquad \text{where } U_n = U(x_n, \omega_n) \end{cases} \tag{SD}$$

where $\eta > 0$ is the step-size and $(\omega_n)_{n \geq 0}$ are i.i.d. random variables taking values in some measurable space Ω.

D.2. Blanket assumptions and related definitions. *In the following, we will always assume that the upcoming Assumptions 3 and 4 hold and we will make frequent use of the equations and definitions introduced here.*

Assumption 3 (Vector field and noise regularity)**.**

(*a*) The vector field $F : \mathcal{V} \to \mathcal{V}$ is C^1-differentiable and satisfies the Lipschitz condition

$$\|F(x) - F(x')\| \leq L\|x - x'\| \quad \text{for all } x, x' \in \mathcal{V} \tag{D.2}$$

(*b*) The noise function $U : \mathcal{V} \times \Omega \to \mathcal{V}$ is C^2-differentiable with Ω compact subspace of a vector space and satisfies the growth condition

$$\sup_{x \in \mathcal{V}, \omega \in \Omega} \frac{\|U(x, \omega)\|}{1 + \|x\|} < +\infty\,. \tag{D.3}$$

(*c*) The sequence of noise samples $(\omega_n)_{n \geq 0}$ is i.i.d.

(*d*) *Properness:* $\mathbb{E}[U(x, \omega)] = 0$ and $\mathrm{cov}(U(x, \omega)) \succ 0$ for all $x \in \mathcal{V}$.

For convenience, let us define $M > 0$ such that: for all $x \in \mathcal{V}$, $\omega \in \Omega$,

$$\|F(x)\| \leq M(1 + \|x\|) \quad \text{and} \quad \|U(x, \omega)\| \leq M(1 + \|x\|)\,. \tag{D.4}$$

Remark D.1. Note that in Assumption 3, the requirement on the covariance matrix in (*d*) is in reality only required in a neighborhood of the equivalence classes.

Assumption 4 (Behavior at infinity)**.**

(*a*) *Dissipativity:* There exist constants $\kappa, C > 0$ such that

$$\langle x, F(x) \rangle \leq -\kappa\|x\|^2 + C \tag{D.5}$$

(*b*) *Sub-Gaussian noise:* The cumulant-generating function of U is bounded as

$$\log \mathbb{E}\big[e^{\langle p, U(x,\omega)\rangle}\big] \le \frac{1}{2}\overline{\sigma}^2(1+\|x\|^s)\,\|p\|^2 \quad \text{for all } p \in \mathbb{R}^d,\, x \in \mathcal{V}\,, \tag{D.6}$$

for some $s \in [0,2]$, $\overline{\sigma}^2 > 0$. If $s = 2$, assume in addition that

$$\frac{\kappa^2}{\overline{\sigma}^2} > 512 d \log 6 \tag{D.7}$$

Remark D.2. We could replace Assumption 4(*b*) with a slightly more general assumption of the form

$$\log \mathbb{E}\big[e^{\langle p, U(x,\omega)\rangle}\big] \le \frac{1}{2}\Big(\overline{\sigma}_0^2 + \overline{\sigma}_1^2\|x\|^s\Big)\,\|p\|^2 \quad \text{for all } p \in \mathbb{R}^d,\, x \in \mathcal{V}\,, \tag{D.8}$$

for some $s \in [0,2]$, $\overline{\sigma}_0^2, \overline{\sigma}_1^2 > 0$. In that case, if $s = 2$, we would need to only assume that

$$\frac{\kappa}{\overline{\sigma}_1^2} > 512 d \log 6 \tag{D.9}$$

which is a slightly weaker condition.

For convenience, we introduce the following definitions.

Definition D.1. Define $\sigma_\infty^2 : \mathbb{R}_+ \to \mathbb{R}_+^*$ as $\sigma_\infty^2(t) = \overline{\sigma}^2(1+t^{s/2})$ so that Assumption 4(*b*) can be rewritten as

$$\log \mathbb{E}\big[e^{\langle p, U(x,\omega)\rangle}\big] \le \frac{1}{2}\sigma_\infty^2(\|x\|^2)\,\|p\|^2 \quad \text{for all } p \in \mathbb{R}^d,\, x \in \mathcal{V}\,. \tag{D.10}$$

Define $\alpha_\infty : \mathbb{R}_+ \to \mathbb{R}_+^*$ as a primitive of $1/\sigma_\infty^2$, i.e., $\alpha_\infty'(t) = 1/\sigma_\infty^2(t)$ and

$$g_\infty(x) := 2\alpha_\infty(\|x\|^2) \tag{D.11}$$

Finally, note that, as mentioned in the main text (Assumption 2), we need to ensure that the number of metastable components is finite and that the transition costs in between them are well-defined and finite. This assumption is stated in a later section as Assumption 5.

D.3. Basic properties of the process.

Lemma D.1 (Growth of the iterates, I)**.** *For any $x_0 \in \mathcal{V}$, $\eta_0 > 0$, for every $\eta \in (0, \eta_0]$, for any $T \ge 1$, $0 \le n \le T\lceil \eta^{-1} \rceil$,*

$$\|x_n\| \le e^{2M(1+\eta_0)T}(1+\|x_0\|)\,. \tag{D.12}$$

Proof. For any $n \ge 0$, the triangular inequality and the growth condition of Eq. (D.4) yield

$$\begin{aligned}\|x_{n+1}\| &\le \|x_n\| + \eta\|F(x_n)\| + \eta\|U_n\| \\ &\le \|x_n\| + 2\eta M(1+\|x_n\|) \\ &= (1+2\eta M)\|x_n\| + 2\eta M.\end{aligned} \tag{D.13}$$

Solving this recursion therefore leads to

$$\|x_n\| \le (1+2\eta M)^n(1+\|x_0\|), \tag{D.14}$$

and the result then follows by noticing that, for $0 \le n \le T\lceil \eta^{-1} \rceil$,

$$(1+2\eta M)^n \le (1+2\eta M)^{T\lceil \eta^{-1}\rceil} \le e^{2MT\eta\lceil \eta^{-1}\rceil} \le e^{2MT(1+\eta)}\,. \qquad \blacksquare$$

Lemma D.2 (Growth of the iterates, II)**.** *For any $x_0 \in \mathcal{V}$, for every $\eta \in (0, (4M)^{-1}]$, for any $T \ge 1$, $0 \le n \le T\lceil \eta^{-1} \rceil$,*

$$\|x_n\| \ge e^{-(4M+1)T}\|x_0\| - 1\,. \tag{D.15}$$

Proof. For any $n \ge 0$, the triangular inequality and the growth condition of Eq. (D.4) yield

$$\begin{aligned}\|x_{n+1}\| &\ge \|x_n\| - \eta\|F(x_n)\| - \eta\|U_n\| \\ &\ge \|x_n\| - 2\eta M(1+\|x_n\|) \\ &= (1-2\eta M)\|x_n\| - 2\eta M.\end{aligned} \tag{D.16}$$

Solving this recursion therefore leads to

$$\|x_n\| \ge (1-2\eta M)^n\|x_0\| - 1, \tag{D.17}$$

and the result then follows by noticing that, for $0 \le n \le T\lceil \eta^{-1} \rceil$,

$$(1-2\eta M)^n \ge (1-2\eta M)^{T\lceil \eta^{-1}\rceil} \ge e^{-4MT\eta\lceil \eta^{-1}\rceil} \ge e^{-4MT(1+\eta)} \tag{D.18}$$

where we used that $2\eta M \le 1/2$ and $\log(1-x) \ge -2x$ on $[0, 1/2]$. $\blacksquare$

D.4. Basic properties of Hamiltonians and Lagrangians.

Definition D.2 (Hamiltonians and Lagrangians)**.** Define, for $x \in \mathcal{V}$, $w \in \mathcal{V}^*$,

$$\begin{aligned} \bar{\mathcal{H}}(x,p) &= \log \mathbb{E}[\exp(\langle p, U(x,\omega)\rangle)] \\ \mathcal{H}(x,p) &= \langle F(x), p\rangle + \bar{\mathcal{H}}(x,p) \\ \bar{\mathcal{L}}(x,w) &= \bar{\mathcal{H}}(x,\cdot)^*(w) \\ \mathcal{L}(x,w) &= \mathcal{H}(x,\cdot)^*(w) = \bar{\mathcal{L}}(x, w - F(x))\,. \end{aligned} \tag{D.19}$$

We note the following standard properties of the Hamiltonian and the Lagrangian, which are derived directly from their definition.

Lemma D.3 (Properties of $\mathcal{H}$ and $\mathcal{L}$)**.**

1. *$\mathcal{H}$ is C^2 and $\mathcal{H}(x,\cdot)$ is convex for any $x \in \mathcal{V}$.*
2. *$\mathcal{L}(x,\cdot)$ is convex for any $x \in \mathcal{V}$, $\mathcal{L}$ is lower semi-continuous (l.s.c.) on $\mathcal{V} \times \mathcal{V}^*$.*
3. *For any $x \in \mathcal{V}$, $w \in \mathcal{V}^*$, $\mathcal{H}(x,w) \le 2M(1+\|x\|)\|w\|$ and $\operatorname{dom} \mathcal{L}(x,\cdot) \subset \bar{\mathbb{B}}(0, 2M(1+\|x\|))$.*
4. *For any $x \in \mathcal{V}$, $w \in \mathcal{V}^*$, $\mathcal{L}(x,w) \ge 0$ and $\mathcal{L}(x,w) = 0 \iff w = F(x)$.*

Proof. To prove the last point, note that since $\mathcal{H}(x,\cdot)$ is convex and differentiable,

$$\mathcal{L}(x,w) = 0 \iff 0 \in \partial_w \mathcal{L}(x,w) \iff w \in \partial_w \mathcal{H}(x,0) \iff w = \nabla_w \mathcal{H}(x,0)\,, \tag{D.20}$$

but $\nabla_w \mathcal{H}(x,0) = F(x)$ because the noise has zero mean by assumption. ∎

Lemma D.4. *For any $x \in \mathcal{V}$, $w \in \mathcal{V}^*$,*

$$\mathcal{L}(x,w) \ge \frac{\|w - F(x)\|^2}{2\sigma_\infty^2(\|x\|^2)}\,. \tag{D.21}$$

Proof. For any $x \in \mathcal{V}$, $w \in \mathcal{V}^*$, we have that

$$\mathcal{H}(x,w) = \langle w, F\rangle + \bar{\mathcal{H}}(x,w) \le \langle w, F\rangle + \frac{1}{2}\sigma_\infty^2(\|x\|^2)\|w\|^2, \tag{D.22}$$

by Assumption 4(*b*), and taking the conjugate therefore yields

$$\mathcal{L}(x,w) \ge \frac{\|w - F\|^2}{2\sigma_\infty^2(\|x\|^2)}\,. \qquad ∎$$

D.5. Large deviation principle. The goal of this section is to provide large deviation principles for continuous and discrete trajectories related to our dynamics of interest (SD). From the sequences $(x_n)_{n\ge0}$ and $(\omega_n)_{n\ge0}$, we define two sequences: one discrete (in lowercase) and one continuous (in uppercase):

(*a*) The discrete "rescaled" trajectory

$$x_n^\eta := x_{n\lfloor 1/\eta\rfloor}\,. \tag{D.23a}$$

(*b*) The continuous "interpolated" trajectory, defined for any $n \ge 0$, $t \in [\eta n, \eta(n+1)]$ by

$$X_t = x_n + \left(\frac{t}{\eta} - n\right)(x_{n+1} - x_n) \tag{D.23b}$$

Note that these sequences are "accelerated" by a factor $1/\eta$ compared to the original sequences appearing in (SD), as classically done in the large deviation literature (see [36, Chap. 7.1]). In this section, we establish large deviation principles for the continuous interpolated trajectory $(X_t)_{t\ge0}$ and for the discrete rescaled sequence $(x_n^\eta)_{n\ge0}$.

We equip the space of continuous functions $[0,T] \to \mathcal{V}$, $\mathcal{C}_T := \mathcal{C}_T(\mathcal{V})$, with the distance induced by the uniform norm

$$\operatorname{dist}_{[0,T]}(\gamma,\gamma') := \sup_{t\in[0,T]} \|\gamma(t) - \gamma'_t\|\,. \tag{D.24}$$

From the Lagrangian defined in Eq. (D.19), we define on $\mathcal{C}_T$ the normalized *action functional*

$$\mathcal{S}_{[0,T]}(\gamma) = \begin{cases} \int_0^T \mathcal{L}(\gamma(t), \dot{\gamma}(t))\, dt & \text{if } \gamma \text{ is absolutely continuous,} \\ \infty & \text{otherwise,} \end{cases} \tag{D.25}$$

following Freidlin & Wentzell [36, Chap. 3.2], as a manner to quantify how "probable" a trajectory is. This will be a central quantity in the upcoming derivations.

We now make precise the link between trajectories probabilities and action functional, i.e., the continuous *large deviation principle*. We also provide a compactness result on the set of all curves with bounded action functional. These results are in the same flavor as [9, Lemma C.3, Prop. C.1] as they are at the base of the approach of Freidlin & Wentzell [36]. For clarity of exposition, proof details are deferred to Appendix I.

Corollary D.1. *Fix $T > 0$. Then:*

- *For all $s > 0$, $\mathcal{K} \subset \mathcal{V}$ compact, the set*
$$\Gamma_{[0,T]}(\mathcal{K}; s) := \{\gamma \in \mathcal{C}_T(\mathcal{V}) : \gamma(0) \in \mathcal{K}, \mathcal{S}_{[0,T]}(\gamma) \leq s\} \tag{D.26}$$
is compact and $\mathcal{S}_{[0,T]}$ is l.s.c. on $\mathcal{C}_T(\mathcal{V})$
- *For all $s, \delta, \epsilon > 0$, $\mathcal{K} \subset \mathcal{V}$ compact, there exists $\eta_0 > 0$ such that, for any $\eta \in (0, \eta_0]$, for all $x_0 \in \mathcal{K}$, $n \leq T$, $\gamma \in \Gamma_{[0,T]}(\{x_0\}; s)$, we have that*
$$\mathbb{P}_{x_0}\left(\mathrm{dist}_{[0,T]}(X, \gamma) < \delta\right) \geq \exp\left(-\frac{\mathcal{S}_{[0,T]}(\gamma) + \epsilon}{\eta}\right) \tag{D.27a}$$
and
$$\mathbb{P}_{x_0}\left(\mathrm{dist}_{[0,T]}(X, \Gamma_{[0,T]}(\{x_0\}; s)) > \delta\right) \leq \exp\left(-\frac{s - \epsilon}{\eta}\right) \tag{D.27b}$$

Proof. The first point is given by Lemma I.3 and the second one by Proposition I.1. ∎

Now, to benefit from the same type of results on discrete sequences which are our primary question, we fix some integer $T \geq 1$ and equip $\mathcal{V}^T$ (i.e., the discrete sequences of T points) with the distance

$$\mathrm{dist}_T(\xi, \zeta) = \max_{0 \leq n \leq T-1} \|\xi_n - \zeta_n\|. \tag{D.28}$$

For a discrete sequence $\xi = (\xi_0, \dots, \xi_{T-1}) \in \mathcal{V}^T$, we can now define the normalized *discrete action functional* by

$$\mathcal{A}_T(\xi) := \sum_{n=0}^{T-2} \rho(\xi_n, \xi_{n+1}), \quad \mathcal{A}_1(\xi) := 0, \tag{D.29}$$

where the cost of moving from one iteration to the next is defined for any $x, x' \in \mathcal{V}$ from the previous continuous normalized action functional (cf. Eq. (D.25)) with horizon 1 as

$$\rho(x, x') := \inf\{\mathcal{S}_{[0,1]}(\gamma) : \gamma \in \mathcal{C}_1(\mathcal{V}), \gamma(0) = x, \gamma(1) = x'\} \tag{D.30}$$

which will be another central quantity in our analysis, following the seminal work of Kifer [63].

Corollary D.2. *Fix $T \geq 1$. Then:*

- *For all $s > 0$, $\mathcal{K} \subset \mathcal{V}$ compact, the set*
$$\Gamma_T(\mathcal{K}; s) := \{\xi \in \mathcal{V}^T : \xi_0 \in \mathcal{K}, \mathcal{A}_T(\xi) \leq s\} \tag{D.31}$$
is compact and $\mathcal{A}_T$ is l.s.c. on $\mathcal{V}^T$.
- *For all $s, \delta, \epsilon > 0$, $\mathcal{K} \subset \mathcal{V}$ compact, there exists $\eta_0 > 0$ such that, for any $\eta \in (0, \eta_0]$, for all $x_0 \in \mathcal{K}$, $n \leq T$, $\xi \in \Gamma_T(\{x_0\}; s)$, we have that*
$$\mathbb{P}_{x_0}(\mathrm{dist}_T(x^\eta, \xi) < \delta) \geq \exp\left(-\frac{\mathcal{A}_T(\xi) + \epsilon}{\eta}\right) \tag{D.32a}$$
and
$$\mathbb{P}_{x_0}\left(\mathrm{dist}_T(x^\eta, \Gamma_T^{\{x_0\}}(s)) > \delta\right) \leq \exp\left(-\frac{s - \epsilon}{\eta}\right). \tag{D.32b}$$

Proof. The first point is given by Lemma I.5 and the second one by Proposition I.2. ∎

E Equivalence classes

E.1 Flow of a vector field and its properties 33
E.2 ℓ^2-chains and their equivalence classes 35
E.3 Basic equivalence classes 38
E.4 Regularity results 46
E.5 From equivalence classes to ICTs 50
E.6 Lifting through diffeomorphism 52

This appendix investigates the properties of sets of points that can be connected by null-energy paths. In contrast to the framework of Azizian et al. [9], these sets are no longer restricted to fixed points under our general vector field setup, making their analysis significantly more intricate. We begin by providing an exact characterization of null-energy paths, relating them to chains of arbitrarily small ℓ^2 cost with respect to the deterministic flow induced by the vector field. Leveraging this equivalence, we derive regularity properties for the energy displacement, and subsequently prove that the null-energy sets are internally chain transitive in the sense of Benaïm [13].

E.1. Flow of a vector field and its properties. In this section, we collect some important properties of the flow induced by F, starting with its formal definition.

Definition E.1. Define, for $x \in \mathcal{V}$, the flow Φ of F starting at x, that is,

$$\dot{\Phi}_t(x) = F(\Phi_t(x)) \quad \text{with} \quad \Phi_0(x) = x, \tag{E.1}$$

and let $\phi(x)$ be the value of this flow at time 1, i.e.,

$$\phi(x) = \Phi_1(x)\,. \tag{E.2}$$

The following standard results establish the connection between the flow and paths that minimize the normalized action functional.

Lemma E.1 (Properties of the flow)**.** *Φ is well-defined and continuous in both time and space, and, for any $T \geq 0$, $\gamma \in \mathcal{C}([0,T],\mathcal{V})$ such that $\gamma_0 = x$,*

$$\mathcal{S}_{0,T}(\gamma) = 0 \iff \gamma_t = \Phi_t(x) \quad \textit{for all } t \in [0,T]\,. \tag{E.3}$$

Proof. The well-posedness and continuity of Φ are a consequence of F being continuously differentiable and of the global Cauchy-Lipschitz (Picard–Lindelöf) theorem for ordinary differential equation (ODE). The equivalence follows from the uniqueness of the flow and Lemma D.3 since

$$\begin{aligned} \mathcal{S}_{0,T}(\gamma) = 0 &\iff \mathcal{L}(\gamma_t, \dot{\gamma}_t) = 0 \text{ almost everywhere} \\ &\iff \dot{\gamma}_t = F(\gamma_t) \text{ almost everywhere} \end{aligned} \tag{E.4}$$

and thus both γ and Φ satisfy the same ODE with the same initial condition and are therefore equal for all t by uniqueness of the solutions. ∎

Translating these properties for ϕ yields the following result.

Lemma E.2 (Properties of ϕ)**.** *ϕ is well-defined and continuous and, for any $x, x' \in \mathcal{V}$,*

$$\rho(x,x') = 0 \iff x' = \phi(x)\,. \tag{E.5}$$

Proof. The implication ($\Leftarrow$) is immediate by definition of ϕ and Lemma E.1. Now for the reverse, assume that $\rho(x,x') = 0$. Following the proof of Lemma I.5, there exists $\gamma \in \mathcal{C}([0,1],\mathcal{V})$ such that $\gamma_0 = x$, $\gamma_1 = x'$ and $\mathcal{S}_{0,1}(\gamma) = 0$. By Lemma E.1, $\gamma_t = \Phi_t(x)$ for all $t \in [0,1]$ and thus $x' = \Phi_1(x) = \phi(x)$. ∎

Since the vector field F is Lipschitz-continuous by assumption, its associated flow is as well.

Lemma E.3 (Lipschitzness of the flow)**.** *For every $T > 0$, there exists a finite constant $L(T) > 0$ such that the flow $\Phi_t(\cdot)$ is $L(T)$-Lipschitz continuous whenever $|t| \leq T$.*

Proof. This result is a direct consequence of F being Lipschitz-continuous and of Grönwall's inequality. Indeed, if we first suppose that $0 \leq t \leq T$, then the definition of the flow yields

$$\|\Phi_t(x) - \Phi_t(y)\| = \left\| \int_0^t [F(\Phi_s(x)) - F(\Phi_s(y)] ds \right\| \leq \int_0^t \|F(\Phi_s(x) - \Phi_s(y)\| ds \tag{E.6}$$

$$\leq L \int_0^t \|\Phi_s(x) - \Phi_s(y)\| ds \tag{E.7}$$

where L is the Lipschitz constant of F. By Grönwall's lemma, we therefore get

$$\|\Phi_t(x) - \Phi_t(y)\| \leq e^{Lt} \|\Phi_0(x) - \Phi_0(y)\| \leq e^{LT} \|x - y\|. \tag{E.8}$$

The same arguments can also be used with $-T \leq t \leq 0$, thus establishing that $\Phi_t(\cdot)$ is $L(T)$-Lipschitz continuous with constant $L(T) \coloneqq e^{LT}$. ∎

The following property is an additional consequence of F being Lipschitz continuous.

Lemma E.4. *For every $T \geq 0$, there exists a finite constant $C(T) \geq 0$ such that, for every $x \in \mathcal{V}$ and any path $\gamma \in \mathcal{C}([0,T], \mathcal{V})$ starting from x,*

$$\|\Phi_t(x) - \gamma_t\|^2 \leq C(T) \int_0^t \|\dot{\gamma}_s - F(\gamma_s)\|^2 ds, \quad 0 \leq t \leq T. \tag{E.9}$$

Proof. Fix $T > 0$, $x \in \mathcal{V}$, and let $\gamma \in \mathcal{C}([0,T], \mathcal{V})$ be a path starting from x. We can assume without loss of generality that γ is absolutely continuous. Indeed, if it was not the case, then the integral in the right-hand side of Eq. (E.9) would become infinite and thus the inequality would hold trivially. Since γ is absolutely continuous, we can therefore write

$$\|\Phi_t(x) - \gamma_t\| = \left\| \int_0^t (F(\Phi_s(x)) - \dot{\gamma}_s) ds \right\| \leq \int_0^t \|F(\Phi_s(x)) - \dot{\gamma}_s\| ds \tag{E.10}$$

$$\leq \int_0^t (\|\dot{\gamma}_s - F(\gamma_s)\| + \|F(\gamma_s) - F(\Phi_s(x))\|) ds \tag{E.11}$$

$$\leq \int_0^t \|\dot{\gamma}_s - F(\gamma_s)\| ds + L \int_0^t \|\Phi_s(x) - \gamma_s\| ds \tag{E.12}$$

where the last inequality comes from L-Lipschitzness of F, cf. Assumption 3(*a*). One can therefore apply the integral form of Grönwall's lemma to get

$$\|\Phi_t(x) - \gamma_t\| \leq e^{Lt} \int_0^t \|\dot{\gamma}_s - F(\gamma_s)\| ds \leq e^{LT} \int_0^t \|\dot{\gamma}_s - F(\gamma_s)\| ds \tag{E.13}$$

and Cauchy-Schwarz inequality therefore yields

$$\|\Phi_t(x) - \gamma_t\|^2 \leq e^{2LT} \left(\int_0^t \|\dot{\gamma}_s - F(\gamma_s)\| ds \right)^2 \leq T e^{2LT} \int_0^t \|\dot{\gamma}_s - F(\gamma_s)\|^2 ds, \tag{E.14}$$

which proves Eq. (E.9) with $C(T) = T e^{2LT}$. ∎

Finally, the dissipativity of the vector field guarantees the existence of a compact set that is positive invariant under the flow.

Lemma E.5. *If $\mathcal{K}_0 \subseteq \mathcal{V}$ is a compact set, then it is included in another compact $\mathcal{K} \subseteq \mathcal{V}$ such that $\Phi_t(\mathcal{K}) \subseteq \mathcal{K}$ for all $t \geq 0$.*

Proof. From Assumption 4(*a*), let $\mathcal{K}'$ be a compact set such that $\langle F(x), x \rangle \leq -\kappa \|x\|^2$ for all x outside of $\mathcal{K}'$, and take $M = \max_{x \in \mathcal{K}_0 \cup \mathcal{K}'} \|x\|$. We then consider the compact set $\mathcal{K} \coloneqq \bar{\mathbb{B}}(0, M)$, which contains both $\mathcal{K}_0$ and $\mathcal{K}'$ by construction. We proceed by contradiction to show that $\mathcal{K}$ is positively invariant: take $x \in \mathcal{K}$, $t > 0$, and assume that $\Phi_t(x) \notin \mathcal{K}$. Since $x \in \mathcal{K}$, there exists a time $0 \leq s \leq t$ such that $\Phi_s(x) \in \mathcal{K}$ and $\Phi_\tau(x) \notin \mathcal{K}$ for all $s < \tau \leq t$. As a consequence,

$$\|\Phi_t(x)\|^2 - \|\Phi_s(x)\|^2 = \int_s^t 2 \langle \Phi_\tau(x), \dot{\Phi}_\tau(x) \rangle d\tau = 2 \int_s^t \langle \Phi_\tau(x), F(\Phi_\tau(x)) \rangle d\tau \tag{E.15}$$

$$\leq -2\kappa \int_s^t \|\Phi_\tau(x)\|^2 d\tau \tag{E.16}$$

$$\leq 0, \tag{E.17}$$

where the second-to-last inequality comes from the fact that $\Phi_\tau(x) \in \mathcal{V} \setminus \mathcal{K} \subseteq \mathcal{V} \setminus \mathcal{K}'$. This leads to $\|\Phi_t(x)\| \leq \|\Phi_s(x)\| \leq M$ and thus $\Phi_t(x) \in \mathcal{K}$, which contradicts our assumptions. It follows that the compact $\mathcal{K}$ is indeed positively invariant. ∎

E.2. ℓ^2-chains and their equivalence classes. We first formalize the notion of an ℓ^2-chain. This concept is inspired by the standard pseudo-orbits studied in dynamics systems (cf. Benaïm [13]), but relies on an ℓ^2-cost rather than the typical ℓ^∞ formulation.

Definition E.2 (ℓ^2-chain)**.** We define $x \rightsquigarrow y$ if and only if for all $\varepsilon > 0, T > 0$, there exists a finite sequence of points $x_0, x_1, \ldots, x_m \in \mathcal{V}$ such that $x_0 = x$, $x_m = y$ and $t_1, \ldots, t_m \geq T$ such that

$$\sum_{i=1}^m \|\Phi_{t_i}(x_{i-1}) - x_i\|^2 < \varepsilon\,. \tag{E.18}$$

If this is the case, we then say that there is an ℓ^2-chain from x to y.

We add a technical lemma that will be useful later.

Lemma E.6. *For any $x, y \in \mathcal{V}$, $x \rightsquigarrow y$ if and only if for all $\varepsilon > 0$, $T > 0$, there exists a finite sequence of points $x_0, x_1, \ldots, x_m \in \mathcal{V}$ such that $x_m = y$ and $t_1, \ldots, t_m \geq T$ such that*

$$\|x - x_0\|^2 + \sum_{i=1}^m \|\Phi_{t_i}(x_{i-1}) - x_i\|^2 < \varepsilon\,. \tag{E.19}$$

Proof. The implication ($\Rightarrow$) is immediate by choosing $x_0 = x$. We now show the reverse implication. Fix $\varepsilon > 0, T > 0$. By Assumption 3(*a*), there exists $C = C(T) \geq 1$ such that $\Phi_T(\cdot)$ is C-Lipschitz. By assumption, there exists a finite sequence of points $x_0, x_1, \ldots, x_m \in \mathcal{V}$ such that $x_m = y$ and $t_1, \ldots, t_m \geq 2T$ such that

$$\|x - x_0\|^2 + \sum_{i=1}^m \|\Phi_{t_i}(x_{i-1}) - x_i\|^2 < \frac{\varepsilon}{C^2}\,. \tag{E.20}$$

We then consider the sequence $x, \Phi_T(x_0), x_1, \ldots, x_m$ and the times $T, t_1 - T, t_2, \ldots, t_m$. By the Lipschitzness of $\Phi_T(\cdot)$ with $C \geq 1$, we have

$$\|\Phi_T(x) - \Phi_T(x_0)\|^2 + \|\Phi_{t_1 - T}(\Phi_T(x_0)) - x_1\|^2 + \sum_{i=2}^m \|\Phi_{t_i}(x_{i-1}) - x_i\|^2 \tag{E.21}$$

$$\leq C\|x - x_0\|^2 + \|\Phi_{t_1}(x_0) - x_1\|^2 + \sum_{i=2}^m \|\Phi_{t_i}(x_{i-1}) - x_i\|^2 \tag{E.22}$$

$$\leq C\left(\|x - x_0\|^2 + \sum_{i=1}^m \|\Phi_{t_i}(x_{i-1}) - x_i\|^2\right) < \varepsilon\,, \tag{E.23}$$

which concludes the proof. ∎

In particular, ℓ^2-chains can be used to define an equivalence relation as follows:

Proposition E.1. *The relation $\leftrightsquigarrow$ defined for any $x, y \in \mathcal{V}$ by*

$$x \leftrightsquigarrow y \iff x \rightsquigarrow y \text{ and } y \rightsquigarrow x \tag{E.24}$$

is an equivalence relation on $\{x \in \mathcal{V} : x \rightsquigarrow x\}$. We denote by $[x]_{\leftrightsquigarrow}$ its equivalence classes.

Proof. Follows directly from the definition of an ℓ^2-chain. ∎

The significance of these equivalence classes will become evident throughout the remainder of the paper, as they constitute the sets around which the trajectories of the stochastic system (SD) concentrate in the long run.

We begin by proving their invariance under the flow.

Lemma E.7 (Equivalence classes are invariant)**.** *The equivalence classes $[x]_{\leftrightsquigarrow}$ are invariant under the flow; that is, $\Phi_t([x]_{\leftrightsquigarrow}) = [x]_{\leftrightsquigarrow}$ for every $t \in \mathbb{R}$.*

Proof. First note that it is sufficient to prove $\Phi_t([x]_{\leftrightsquigarrow}) \subseteq [x]_{\leftrightsquigarrow}$ for every $t \in \mathbb{R}$. Indeed, applying Φ_{-t} on both sides of the inclusion would then yield $[x]_{\leftrightsquigarrow} \subseteq \Phi_{-t}([x]_{\leftrightsquigarrow})$ for all $t \in \mathbb{R}$ by definition of the flow.

Let us pick $x \in \mathcal{V}$ such that $x \rightsquigarrow x$ and let $t \in \mathbb{R} \setminus \{0\}$ (the case $t = 0$ is trivial). We therefore need to show that $x \leftrightsquigarrow \Phi_t(x)$. Fix $\varepsilon > 0$ and $T > 0$. Since $x \rightsquigarrow x$, there exists a finite sequence of points $x_0, \dots, x_m \in \mathcal{V}$ such that $x_0 = x_m = x$ and $t_1, \dots, t_m \geq T + |t|$ such that $\sum_i \|\Phi_{t_i}(x_{i-1}) - x_i\|^2 < \varepsilon$.

- $(\Phi_t(x) \rightsquigarrow x)$: We consider the modified sequences $(x'_i)_{0 \leq i \leq m}$ and $(t'_i)_{1 \leq i \leq m}$ obtained by replacing the *first element* of $(x_i)_{0 \leq i \leq m}$ by $\Phi_t(x)$ and the *first element* of $(t_i)_{1 \leq i \leq m}$ by $t_1 - t$ respectively. Note that $t'_i \geq T$ for all $1 \leq i \leq m$ since $t_i \geq T + |t|$, and that
$$\|\Phi_{t'_1}(x'_0) - x'_1\| = \|\Phi_{t_1 - t}(\Phi_t(x)) - x_1\| = \|\Phi_{t_1}(x_0) - x_1\|, \tag{E.25}$$
which directly leads to
$$\sum_{i=1}^{m} \|\Phi_{t'_i}(x'_{i-1}) - x'_i\|^2 = \sum_{i=1}^{m} \|\Phi_{t_i}(x_{i-1}) - x_i\|^2 < \varepsilon. \tag{E.26}$$
This holds true for every $\varepsilon > 0$, $T > 0$, and thus $\Phi_t(x) \rightsquigarrow x$.
- $(x \rightsquigarrow \Phi_t(x))$:We consider the modified sequences $(x'_i)_{0 \leq i \leq m}$ and $(t'_i)_{1 \leq i \leq m}$ obtained by replacing the *last element* of $(x_i)_{0 \leq i \leq m}$ by $\Phi_t(x)$ and the *last element* of $(t_i)_{1 \leq i \leq m}$ by $t_m + t$ respectively. These sequences still satisfy $t'_i \geq T$ for all $1 \leq i \leq m$, and they further have
$$\|\Phi_{t'_m}(x'_{m-1}) - x_m\| = \|\Phi_t(\Phi_{t_m}(x_{m-1})) - \Phi_t(x)\| \leq L_t \|\Phi_{t_m}(x_{m-1}) - x\|, \tag{E.27}$$
where the last step follows from the standard fact that the flow $\Phi_t(\cdot)$ associated to a Lipschitz-continuous vector field is itself Lipschitz-continuous with some constant $L_t \geq 0$. As a consequence,
$$\sum_{i=1}^{m} \|\Phi_{t'_i}(x'_{i-1}) - x'_i\|^2 \leq \max(1, L_t^2) \sum_{i=1}^{m} \|\Phi_{t_i}(x_{i-1}) - x_i\|^2 < \max(1, L_t^2)\varepsilon, \tag{E.28}$$
and since this is true for every $\varepsilon > 0$, $T > 0$, it follows that $x \rightsquigarrow \Phi_t(x)$. ∎

Moreover, they are also automatically closed by definition.

Lemma E.8 (Equivalence classes are closed)**.** *The equivalence classes $[x]_{\leftrightsquigarrow}$ are closed.*

Proof. Fix $x \in \mathcal{V}$ and assume that there is a sequence $y_k \to y$ such that $x \leftrightsquigarrow y_k$. To show that $[x]_{\leftrightsquigarrow}$ is closed, it suffices to prove $x \rightsquigarrow y$ and $y \rightsquigarrow x$.

- $(x \rightsquigarrow y)$: Fixing $\varepsilon > 0$ and $T > 0$, the definition of $x \rightsquigarrow y_k$ implies that there exist finite sequences $x_0^k, \dots, x_{m_k}^k \in \mathcal{V}$ with $x_0^k = x$, $x_{m_k}^k = y_k$, and $t_1^k, \dots, t_{m_k}^k \geq T$ such that
$$\sum_{i=1}^{m_k} \|\Phi_{t_i^k}(x_{i-1}^k) - x_i^k\|^2 < \frac{\varepsilon}{4}. \tag{E.29}$$
Next, we pick k large enough such that $\|y_k - y\|^2 \leq \varepsilon/4$ (which is possible since $y_k \to y$ by assumption), and we construct a chain $x_0, \dots, x_m$ with $m := m_k$ by $x_i = x_i^k$ for every $i = 1, \dots, m-1$ and $x_m = y$, together with a sequence of times $t_1, \dots, t_m$ such that $t_i = t_i^k$ for all $i = 1, \dots, m$. This chain satisfies
$$\sum_{i=1}^{m} \|\Phi_{t_i}(x_{i-1}) - x_i\|^2 \tag{E.30}$$
$$= \sum_{i=1}^{m-1} \|\Phi_{t_i^k}(x_{i-1}^k) - x_i^k\|^2 + \|\Phi_{t_{m_k}^k}(x_{m_k-1}^k) - y\|^2 \tag{E.31}$$
$$\leq 2 \sum_{i=1}^{m-1} \|\Phi_{t_i^k}(x_{i-1}^k) - x_i^k\|^2 + 2\|\Phi_{t_{m_k}^k}(x_{m_k-1}^k) - x_{m_k}^k\|^2 + 2\|x_{m_k}^k - y\|^2 \tag{E.32}$$

$$= 2\sum_{i=1}^{m}\|\Phi_{t_i^k}(x_{i-1}^k) - x_i^k\|^2 + 2\|y_k - y\|^2 \tag{E.33}$$

$$< \varepsilon, \tag{E.34}$$

which implies that $x \rightsquigarrow y$.

- $(y \rightsquigarrow x)$: Fix $\varepsilon > 0, T > 0$, and let $k := k(\varepsilon, T)$ be sufficiently large such that $\|\Phi_T(y) - \Phi_T(y_k\|^2 < \varepsilon/2$. Since the equivalence class $[x]_{\leftrightsquigarrow}$ is invariant by Lemma E.7, $\Phi_T(y_k) \rightsquigarrow x$ and we can therefore construct a chain $x_0^k, \dots, x_m^k$ from $\Phi_T(y_k)$ to x with $t_1^k, \dots, t_m^k \geq T$ such that $\sum_i \|\Phi_{t_i^k}(x_{i-1}^k) - x_i^k\|^2 < \varepsilon/2$. In particular, if we consider a chain that begins at y, jumps at $\Phi_T(y_k)$ after a time $t_0 = T$, and then follows the chain $x_0^k, \dots, x_m^k$ with times $t_1^k, \dots, t_m^k \geq T$, its total cost is

$$\|\Phi_T(y) - \Phi_T(y_k)\|^2 + \sum_{i=1}^{m}\|\Phi_{t_i^k}(x_{i-1}^k) - x_i^k\|^2 < \frac{\varepsilon}{2} + \frac{\varepsilon}{2} \leq \varepsilon. \tag{E.35}$$

Since this is true for every $\varepsilon > 0$ and $T > 0$, it follows that $y \rightsquigarrow x$. ∎

The following lemma, a consequence of the dissipativity of F, will prove highly useful in the sequel.

Lemma E.9. *For every $x_0 \in \mathcal{V}$ and $\varepsilon_0 > 0$, there exist $t_0 := t_0(x_0, \varepsilon_0) > 0$ and $\mathcal{K} := \mathcal{K}(x_0, \varepsilon_0) \subseteq \mathcal{V}$ compact such that any finite chain $x_1, \dots, x_m \in \mathcal{V}$, $t_1, \dots, t_m \geq t_0$ satisfying $\max_{1\leq i\leq m}\|\Phi_{t_i}(x_{i-1}) - x_i\|^2 < \varepsilon_0$ has $x_0, \dots, x_m \in \mathcal{K}$.*

Proof. Fix $x_0 \in \mathcal{V}$ and let $\mathcal{K}_0$ be a compact set outside of which $\langle F(x), x\rangle \leq -\kappa\|x\|^2$. Then, Lemma E.5 implies that there exists a positively invariant compact set $\mathcal{K}_1$ containing both x_0 and $\mathcal{K}_0$. Moreover, one can assume without loss of generality that $\mathcal{K}_1 = \bar{\mathbb{B}}(0, M)$ for some $M > 0$.

Now, fix $\varepsilon_0 > 0$ and define $\mathcal{K} := \{x \in \mathcal{V} : d(x, \mathcal{K}_1)^2 \leq \varepsilon_0\}$. $\mathcal{K}$ is obviously a compact set, and we will show that it satisfies the conclusion of Lemma E.9. To do so, choose $t_0 = ((M + \sqrt{\varepsilon_0})^2 - M^2)/(2\kappa M^2)$ and assume that there is a finite chain $x_1, \dots, x_m \in \mathcal{V}$, $t_1, \dots, t_m \geq t_0$ such that $\max_{1\leq i\leq m}\|\Phi_{t_i}(x_{i-1}) - x_i\|^2 < \varepsilon_0$. We then proceed by induction over the index $i = 0, \dots, m$ in order to establish that $x_i \in \mathcal{K}$. First, note that it is obvious for $i = 0$ since $x_0 \in \mathcal{K}_1 \subset \mathcal{K}$ by construction. Next, assume that $x_{i-1} \in \mathcal{K}$. We first show that there is some time $0 \leq t \leq t_i$ for which $\Phi_t(x_{i-1}) \in \mathcal{K}_1$. Indeed, if we were to suppose that $\Phi_t(x_{i-1}) \notin \mathcal{K}_1$ for every $0 \leq t \leq t_i$, then we would get

$$\|\Phi_{t_i}(x_{i-1})\|^2 - \|x_{i-1}\|^2 = 2\int_0^{t_i}\langle\Phi_\tau(x_{i-1}), \dot{\Phi}_\tau(x_{i-1})\rangle d\tau \leq -2\kappa\int_0^{t_i}\|\Phi_\tau(x_{i-1})\|^2 d\tau \tag{E.36}$$

$$\leq -2\kappa M^2 t_i \tag{E.37}$$

$$\leq -2\kappa M^2 t_0 \tag{E.38}$$

so that

$$\|\Phi_{t_i}(x_{i-1})\|^2 \leq \|x_{i-1}\|^2 - 2\kappa M^2 t_0 \leq (M + \sqrt{\varepsilon_0})^2 - ((M + \sqrt{\varepsilon_0})^2 - M^2) = M^2. \tag{E.39}$$

However, this would implies that $\Phi_{t_i}(x_{i-1}) \in \mathcal{K}_1$, which is a contradiction, and thus there is some $0 \leq t \leq t_i$ such that $\Phi_t(x_{i-1}) \in \mathcal{K}_1$. By positive invariance of $\mathcal{K}_1$, it follows that $\Phi_{t_i}(x_{i-1}) \in \mathcal{K}_1$. It only remains to invoke the chain condition $\operatorname{dist}(\Phi_{t_i}(x_{i-1}), x_i)^2 < \varepsilon_0$ to obtain that $d(x_i, \mathcal{K}_1)^2 \leq \varepsilon_0$, hence showing $x_i \in \mathcal{K}$ by definition. Iterating over all indices then concludes the proof. ∎

In particular, the preceding lemma can be used to establish that the equivalence classes of $\leftrightsquigarrow$ are compact sets.

Lemma E.10 (Equivalence classes are compact)**.** *The equivalence classes $[x]_{\leftrightsquigarrow}$ are compact.*

Proof. Fixing $x \in \mathcal{V}$ and $\varepsilon_0 > 0$, let t_0 be the uniform lower bound and $\mathcal{K}$ the compact set provided by Lemma E.9. For every $y \in \mathcal{V}$ such that $x \rightsquigarrow y$, we obtain by definition the existence of a finite chain $x_0, \dots, x_m$ with $x_0 = x$, $x_m = y$ and $t_1, \dots, t_m \geq t_0$ for which

$$\|\Phi_{t_i}(x_{i-1}) - x_i\|^2 \leq \sum_{i=1}^{m}\|\Phi_{t_i}(x_{i-1}) - x_i\|^2 < \varepsilon_0. \tag{E.40}$$

This chain satisfies the condition posed in Lemma E.9, and thus $x_i \in \mathcal{K}$ for all $i = 0, \dots, m$. In particular, $x_m = y \in \mathcal{K}$. Since this compact set does not depend on the choice of y, we directly get that $[x]_{\leftrightsquigarrow} \subseteq \{y \in \mathcal{V} : x \rightsquigarrow y\} \subseteq \mathcal{K}$, and therefore the equivalence classes are relatively compact. Since they are also closed by Lemma E.8, it follows that they are in fact compact. ∎

E.3. Basic equivalence classes. We begin by formalizing the minimum-energy displacement between two points.

Definition E.3 (Kifer [63, §1.5])**.** Define, for $x, x' \in \mathcal{V}$,

$$\begin{aligned} B(x,x') &= \inf\left\{\mathcal{S}_{[0,T]}(\gamma) : \gamma \in \mathcal{C}([0,T],\mathcal{V}), \gamma_0 = x, \gamma_T = x', T \in \mathbb{N}\right\} \\ &= \inf\left\{\mathcal{A}_T(\xi) : \xi \in \mathcal{X}^T\,,\ \xi_0 = x\,,\ \xi_{T-1} = x',\ T \geq 1\right\}. \end{aligned} \tag{E.41}$$

The fact that these two expressions coincide directly comes from the definition of ρ.

This enables us to define an equivalence relation by grouping points connected by a null-energy path.

Proposition E.2. *The relation $\sim$ defined for any $x, x' \in \mathcal{V}$ by*

$$x \sim x' \iff B(x,x') = B(x',x) = 0 \tag{E.42}$$

is an equivalence relation on $\mathcal{V}$.

Proof. • Reflexivity: this follows from the definition of B by choosing $T = 0$.

- Symmetry: this follows from the definition of $\sim$.
- Transitivity: for any $x, x', x'' \in \mathcal{V}$, we have by construction of B that

$$B(x,x'') \leq B(x,x') + B(x',x'')\,. \tag{E.43}$$

In particular, if $x \sim x'$ and $x' \sim x''$, then $B(x,x'') = 0$. A symmetric argument leads to $B(x'',x) = 0$ and thus $x \sim x''$. ∎

In particular, building upon the framework of Kifer [63], we define the following equivalence classes.

Definition E.4 (Equivalence classes)**.** 1. $[x]_\sim$ are the equivalence classes of the relation $\sim$ on $\mathcal{V}$.

2. The *basic* equivalence classes are the classes $[x]_\sim$ that are either made of at least two points or made of a single fixed point of Φ.

Up to a diffeomorphism, the basic equivalence classes are exactly the *metastable components* defined in the main text (see Definition 1).

The remainder of this section is dedicated to proving that the basic equivalence classes coincide exactly with the equivalence classes $[x]_{\rightsquigarrow}$ of the relation $\rightsquigarrow$ defined in Appendix E.2.

We begin by demonstrating that if an ℓ^2-chain exists between two points, there is also a null-energy path connecting them. To achieve this, however, we first require a series of lemmas establishing equivalent characterization of ℓ^2-chains.

Through an appropriate choice of T, we may assume without loss of generality that the sum of times is arbitrarily large.

Lemma E.11 (Can assume $\sum_i t_i$ arbitrarily large)**.** *Fix $x, y \in \mathcal{V}$. Assume that for all $\varepsilon > 0$, $T > 0$, there exists a finite sequence of points $x_0, x_1, \ldots, x_m \in \mathcal{V}$ such that $x_0 = x$, $x_m = y$ and $t_1, \ldots, t_m \geq T$ such that*

$$\sum_{i=1}^{m} \|\Phi_{t_i}(x_{i-1}) - x_i\|^2 < \varepsilon. \tag{E.44}$$

Then, for all $\varepsilon > 0$, $t_0, T_0 > 0$, there exists a finite sequence of points $x_0, x_1, \ldots, x_m \in \mathcal{V}$ such that $x_0 = x$, $x_m = y$ and $t_1, \ldots, t_m \geq t_0$ such that

$$\sum_{i=1}^{m} \|\Phi_{t_i}(x_{i-1}) - x_i\|^2 < \varepsilon \quad \text{and} \quad \sum_{i=1}^{m} t_i \geq T_0. \tag{E.45}$$

Proof. It suffices to choose $T = \max(t_0, T_0)$, which yields both $t_i \geq T \geq t_0$ and $\sum_{i=1}^m t_i \geq mT \geq T_0$. ∎

Furthermore, by subdividing any large time step t_i into smaller intervals, we may always assume that the times between steps are upper-bounded.

Lemma E.12 (Can assume $t_i \le 2t_0$)**.** *Fix $x, y \in \mathcal{V}$. Assume that for all $\varepsilon > 0$, $t_0, T_0 > 0$, there exists a finite sequence of points $x_0, x_1, \dots, x_m \in \mathcal{V}$ such that $x_0 = x$, $x_m = y$ and $t_1, \dots, t_m \ge t_0$ such that*

$$\sum_{i=1}^{m} \|\Phi_{t_i}(x_{i-1}) - x_i\|^2 < \varepsilon \quad \text{and} \quad \sum_{i=1}^{m} t_i \ge T_0. \tag{E.46}$$

Then, for all $\varepsilon > 0$, $t_0, T_0 > 0$, there exists a finite sequence of points $x_0, x_1, \dots, x_m \in \mathcal{V}$ such that $x_0 = x$, $x_m = y$ and $t_0 \le t_1, \dots, t_m \le 2t_0$ such that

$$\sum_{i=1}^{m} \|\Phi_{t_i}(x_{i-1}) - x_i\|^2 < \varepsilon \quad \text{and} \quad \sum_{i=1}^{m} t_i \ge T_0. \tag{E.47}$$

Proof. Fix $x, y \in \mathcal{V}$ and $\varepsilon, t_0, T_0 > 0$, and let $x_0, \dots, x_m \in \mathcal{V}$ and $t_1, \dots, t_m \ge t_0$ be sequences satisfying Eq. (E.46). Suppose that there exists an index j and an integer $k \ge 2$ such that $kt_0 < t_j \le (k+1)t_0$. We then define a new sequence of times $(t'_i)_{1 \le i \le m+k-1}$ and states $(x'_i)_{0 \le i \le m+k-1}$ by

$$t'_i = \begin{cases} t_i & \text{if } 1 \le i \le j-1, \\ t_j/k & \text{if } j \le i \le j+k-1, \\ t_{i+1-k} & \text{if } j+k \le i \le m+k-1, \end{cases} \qquad x'_i = \begin{cases} x_i & \text{if } i \le j-1, \\ \Phi_{t'_i}(x'_{i-1}) & \text{if } j \le i \le j+k-2, \\ x_{i+1-k} & \text{if } j+k-1 \le i \le m+k-1. \end{cases} \tag{E.48}$$

For this new sequence, we have

$$\begin{aligned} \sum_{i=1}^{m+k-1} \|\Phi_{t'_i}(x'_{i-1}) - x'_i\|^2 &= \sum_{i=1}^{j-1} \|\Phi_{t_i}(x_{i-1}) - x_i\|^2 + \sum_{i=j}^{j+k-2} \|\Phi_{t'_i}(x'_{i-1}) - \Phi_{t'_i}(x'_{i-1})\|^2 \\ &\quad + \|\Phi_{t'_{j+k-1}}(x'_{j+k-2}) - x_j\|^2 + \sum_{i=j+k}^{m+k-1} \|\Phi_{t_{i+1-k}}(x_{i-k}) - x_{i+1-k}\|^2 \qquad \text{(E.49)} \\ &= \sum_{\substack{1 \le i \le m \\ i \ne j}} \|\Phi_{t_i}(x_{i-1}) - x_i\|^2 + \|\Phi_{t_j/k}(x'_{j+k-2}) - x_j\|^2. \qquad \text{(E.50)} \end{aligned}$$

Now, note that

$$\Phi_{t_j/k}(x'_{j+k-2}) = \Phi_{t_j/k}\big(\Phi_{t_j/k}(x'_{j+k-3})\big) = \underbrace{\big(\Phi_{t_j/k} \circ \cdots \circ \Phi_{t_j/k}\big)}_{k \text{ times}}(x_{j-1}) = \Phi_{t_j}(x_{j-1}) \tag{E.51}$$

by uniqueness of the flow (which implies that $\Phi_t \circ \Phi_s = \Phi_{t+s}$). In particular, we therefore get

$$\begin{aligned} \sum_{i=1}^{m+k-1} \|\Phi_{t'_i}(x'_{i-1}) - x'_i\|^2 &= \sum_{\substack{1 \le i \le m \\ i \ne j}} \|\Phi_{t_i}(x_{i-1}) - x_i\|^2 + \|\Phi_{t_j}(x_{j-1}) - x_j\|^2 \qquad \text{(E.52)} \\ &= \sum_{i=1}^{m} \|\Phi_{t_i}(x_{i-1}) - x_i\|^2 \le \varepsilon. \qquad \text{(E.53)} \end{aligned}$$

Moreover, we also have $\sum t'_i = \sum t_i \ge T_0$ and $t_0 \le t'_i \le 2t_0$ for all $j \le i \le j+k-1$ by construction. Iterating this procedure over every $t_i \ge 2t_0$, we therefore obtain a chain satisfying Eq. (E.46) and for which all times are upper-bounded by $2t_0$, hence concluding the proof. ∎

Moreover, Lemma E.9 guarantees that all elements of the chain remain confined to a sufficiently large compact set.

Lemma E.13 (Can assume chain remains in a compact)**.** *Fix $x, y \in \mathcal{V}$. Assume that for all $\varepsilon > 0, t_0, T_0 > 0$, there exists a finite sequence of points $x_0, x_1, \dots, x_m \in \mathcal{V}$ such that $x_0 = x$, $x_m = y$ and $t_0 \le t_1, \dots, t_m \le 2t_0$ such that*

$$\sum_{i=1}^{m} \|\Phi_{t_i}(x_{i-1}) - x_i\|^2 < \varepsilon \quad \text{and} \quad \sum_{i=1}^{m} t_i \ge T_0. \tag{E.54}$$

Then, for all $\varepsilon_0 > 0$, there exists $t_0 := t_0(x, \varepsilon_0) > 0$ and a compact set $\mathcal{K} := \mathcal{K}(x, \varepsilon_0)$ such that, for any $\varepsilon \in (0, \varepsilon_0]$ and $T_0 > 0$, there exists a finite sequence of points $x_0, x_1, \dots, x_m \in \mathcal{K}$ such that $x_0 = x$, $x_m = y$ and $t_0 \le t_1, \dots, t_m \le 2t_0$ such that

$$\sum_{i=1}^{m} \|\Phi_{t_i}(x_{i-1}) - x_i\|^2 < \varepsilon \quad \text{and} \quad \sum_{i=1}^{m} t_i \ge T_0. \tag{E.55}$$

Proof. This is an immediate consequence of Lemma E.9. ∎

Finally, we demonstrate that we may assume the sum of times is not only arbitrarily large, but also an integer.

Lemma E.14 (Can assume $\sum_i t_i \in \mathbb{N}$)**.** *Fix $x, y \in \mathcal{V}$. Assume that for all $\varepsilon_0 > 0$, there exists $t_0 := t_0(x, \varepsilon_0) > 0$ and a compact set $\mathcal{K} := \mathcal{K}(x, \varepsilon_0)$ such that, for any $\varepsilon \in (0, \varepsilon_0]$ and $T_0 > 0$, there exists a finite sequence of points $x_0, x_1, \dots, x_m \in \mathcal{K}$ such that $x_0 = x$, $x_m = y$ and $t_0 \le t_1, \dots, t_m \le 2t_0$ such that*

$$\sum_{i=1}^{m} \|\Phi_{t_i}(x_{i-1}) - x_i\|^2 < \varepsilon \quad \text{and} \quad \sum_{i=1}^{m} t_i \ge T_0. \tag{E.56}$$

Then, for all $\varepsilon_0 > 0$, there exists $t_0 := t_0(x, \varepsilon_0) > 0$ and a compact set $\mathcal{K} := \mathcal{K}(x, \varepsilon_0)$ such that, for any $\varepsilon \in (0, \varepsilon_0]$, there exists a finite sequence of points $x_0, x_1, \dots, x_m \in \mathcal{K}$ such that $x_0 = x$, $x_m = y$ and $t_0 \le t_1, \dots, t_m \le 2(t_0 + 1)$ such that

$$\sum_{i=1}^{m} \|\Phi_{t_i}(x_{i-1}) - x_i\|^2 < 4\varepsilon \quad \text{and} \quad \sum_{i=1}^{m} t_i \in \mathbb{N}. \tag{E.57}$$

Proof. After fixing $x, y \in \mathcal{V}$ and $\varepsilon_0 > 0$, let $t_0 := t_0(x, \varepsilon_0)$ and $\mathcal{K} := \mathcal{K}(x, \varepsilon_0)$ be, respectively, the uniform lower bound on times and the compact set provided by the assumption of Lemma E.14. We then pick $\varepsilon \in (0, \varepsilon_0]$ and $T_0 > 0$ to be determined later, and consider a finite sequence of states $x_0, \dots, x_m \in \mathcal{K}$ and times $t_0 \le t_1, \dots t_m \le 2t_0$ satisfying Eq. (E.56).

Defining $T = \sum_{i=1}^{m} t_i$, let us consider the perturbed times $t_i' = t_i + \delta$ with $\delta = \frac{\lceil T \rceil - T}{m}$. These times satisfy $\sum_{i=1}^{m} t_i' = \lceil T \rceil \in \mathbb{N}$ by construction and, since $0 \le \lceil T \rceil - T < 1$, we also have $0 \le \delta < 1/m \le 1$ which leads to $t_i' \le 2t_0 + 1 \le 2(1 + t_0)$.

Moreover, we can write

$$\begin{aligned}
\|\Phi_{t_i'}(x_{i-1}) - x_i\|^2 &= \|\Phi_{t_i'}(x_{i-1}) - \Phi_{t_i}(x_{i-1}) + \Phi_{t_i}(x_{i-1}) - x_i\|^2 && \text{(E.58)}\\
&\le 2\Big(\|\Phi_{t_i'}(x_{i-1}) - \Phi_{t_i}(x_{i-1})\|^2 + \|\Phi_{t_i}(x_{i-1}) - x_i\|^2\Big) && \text{(E.59)}\\
&\le 2\left\{\left(\int_{t_i}^{t_i'} \|\dot{\Phi}_t(x_{i-1})\|\, dt\right)^2 + \|\Phi_{t_i}(x_{i-1}) - x_i\|^2\right\} && \text{(E.60)}\\
\text{(def. of flow)} \qquad &= 2\left\{\left(\int_{t_i}^{t_i'} \|F(\Phi_t(x_{i-1}))\|\, dt\right)^2 + \|\Phi_{t_i}(x_{i-1}) - x_i\|^2\right\} && \text{(E.61)}\\
(F \text{ is } L\text{-Lipschitz}) \qquad &\le 2\left\{\left(\int_{t_i}^{t_i'} \big(\|F(0)\| + L\|\Phi_t(x_{i-1})\|\big)\, dt\right)^2 + \|\Phi_{t_i}(x_{i-1}) - x_i\|^2\right\} && \text{(E.62)}\\
&\le 2\left\{\Big(\|F(0)\| + L \max_{\substack{t_0 \le t \le 2(t_0+1)\\ z \in \mathcal{K}}} \|\Phi_t(z)\|\Big)^2 \delta^2 + \|\Phi_{t_i}(x_{i-1}) - x_i\|^2\right\} && \text{(E.63)}\\
&= 2\Big\{C(t_0, \mathcal{K})\, \delta^2 + \|\Phi_{t_i}(x_{i-1}) - x_i\|^2\Big\} && \text{(E.64)}
\end{aligned}$$

where

$$C(t_0, \mathcal{K}) = \left(\|F(0)\| + L \max_{\substack{t_0 \le t \le 2(t_0+1)\\ z \in \mathcal{K}}} \|\Phi_t(z)\|\right)^2 \tag{E.65}$$

is a finite constant by joint continuity of the flow and compactness of $[t_0, 2(t_0 + 1)] \times \mathcal{K}$. As a consequence,

$$\begin{aligned}
\sum_{i=1}^{m} \|\Phi_{t_i'}(x_{i-1}) - x_i\|^2 &\le 2\left\{C(t_0, \mathcal{K})\, \delta^2 m + \sum_{i=1}^{m} \|\Phi_{t_i}(x_{i-1}) - x_i\|^2\right\} && \text{(E.66)}\\
\text{Eq. (E.56)} \qquad &\le 2\left\{C(t_0, \mathcal{K})\, \frac{(\lceil T \rceil - T)^2}{m} + \varepsilon\right\} && \text{(E.67)}
\end{aligned}$$

$$(0 \le \lceil T \rceil - T < 1) \qquad \le 2\left(\frac{C(t_0,\mathcal{K})}{m} + \varepsilon\right). \tag{E.68}$$

Furthermore, the condition $\sum_{i=1}^m t_i \ge T_0$ yields $T_0 \le \sum_{i=1}^m t_i \le \sum_{i=1}^m 2t_0 = 2mt_0$, which leads to $m \ge \frac{T_0}{2t_0}$ and thus

$$\sum_{i=1}^m \|\Phi_{t_i'}(x_{i-1}) - x_i\|^2 \le 2\left(\frac{2C(t_0,\mathcal{K})t_0}{T_0} + \varepsilon\right). \tag{E.69}$$

In particular, if we take $T_0 = \frac{2C(t_0,\mathcal{K})t_0}{\varepsilon}$, then

$$\sum_{i=1}^m \|\Phi_{t_i'}(x_{i-1}) - x_i\|^2 \le 4\varepsilon \quad \text{and} \quad \sum_{i=1}^m t_i' \in \mathbb{N}, \tag{E.70}$$

hence establishing that the sequence of perturbed times satisfy the conclusion of Lemma E.14. ∎

Equipped with the characterization of ℓ^2-chain provided by Lemma E.14, we now prove that the energy displacement between these points is zero.

Lemma E.15. *Fix $x, y \in \mathcal{V}$. Assume that for all $\varepsilon_0 > 0$, there exists $t_0 := t_0(x,\varepsilon_0) > 0$ and a compact set $\mathcal{K} := \mathcal{K}(x,\varepsilon_0)$ such that, for any $\varepsilon \in (0,\varepsilon_0]$, there exists a finite sequence of points $x_0, x_1, \dots, x_m \in \mathcal{K}$ such that $x_0 = x$, $x_m = y$ and $t_0 \le t_1, \dots, t_m \le 2(t_0+1)$ such that*

$$\sum_{i=1}^m \|\Phi_{t_i}(x_{i-1}) - x_i\|^2 < 4\varepsilon \quad \textit{and} \quad \sum_{i=1}^m t_i \in \mathbb{N}. \tag{E.71}$$

Then $B(x,y) = 0$.

Proof. Fix $x, x' \in \mathcal{V}$ and $\varepsilon_0 > 0$. For the lower bound $t_0 > 0$ on times and the compact set $\mathcal{K}$ provided both by the condition of Lemma E.15, we then define

$$\mathcal{K}_1 := \mathcal{K}_1(t_0,\mathcal{K}) = \{\Phi_t(z) : 0 \le t \le 2(t_0+1),\ z \in \mathcal{K}\} \tag{E.72}$$

and its enlargement

$$\mathcal{K}' := \mathcal{K}'(\varepsilon_0,\mathcal{K}_1) = \{z \in \mathcal{V} : (z,\mathcal{K}_1) \le 2\sqrt{\varepsilon_0}\}. \tag{E.73}$$

By joint continuity of the flow, both $\mathcal{K}_1$ and $\mathcal{K}'$ are compact sets. In particular, applying Lemma E.20 at each $z \in \mathcal{K}'$ and extracting a sub-cover thereof yield finite constants $\delta > 0$ and $L' > 0$ such that

$$\bar{\mathcal{L}}(z,w) \le L'\|w\|^2 \quad \text{for all } z \in \mathcal{K}',\ w \in \bar{\mathbb{B}}(0,\delta). \tag{E.74}$$

By definition of the Lagrangian, this implies that

$$\mathcal{L}(z,w) = \bar{\mathcal{L}}(z, w - F(z)) \le L'\|w - F(z)\|^2 \quad \text{for all } z \in \mathcal{K}',\ w \in \bar{\mathbb{B}}(F(z),\delta). \tag{E.75}$$

Now, let us fix $\varepsilon \in (0,\varepsilon_0]$ and consider a chain of states $x_0, \dots, x_m \in \mathcal{K}$ and times $t_0 \le t_1, \dots, t_m \le 2(t_0+1)$ satisfying Eq. (E.71). For every $i = 1, \dots, m$, we then construct a path from x_{i-1} to x_i on $[0, t_i]$ by the prescription

$$\gamma_t^i = \Phi_t(x_{i-1}) + \frac{t}{t_i}\big(x_i - \Phi_{t_i}(x_{i-1})\big), \quad t \in [0,t_i]. \tag{E.76}$$

In particular, these paths satisfy

$$\|\gamma_t^i - \Phi_t(x_{i-1})\| = \frac{t}{t_i}\|x_i - \Phi_{t_i}(x_{i-1})\| \le \big(\|x_i - \Phi_{t_i}(x_{i-1})\|^2\big)^{1/2} \tag{E.77}$$

$$\le \Big(\sum_{i=1}^m \|x_i - \Phi_{t_i}(x_{i-1})\|^2\Big)^{1/2} \tag{E.78}$$

$$\le 2\sqrt{\varepsilon} \tag{E.79}$$

$$\le 2\sqrt{\varepsilon_0}. \tag{E.80}$$

Since $t \le t_i \le 2(t_0+1)$ and $x_{i-1} \in \mathcal{K}$, it follows that $\Phi_t(x_{i-1}) \in \mathcal{K}_1$ and thus $\gamma_t^i \in \mathcal{K}'$ for all $t \in [0, t_i]$ and every $i = 1, \dots, m$. Moreover, we also have

$$\|\dot\gamma_t^i - F(\gamma_t^i)\| \le \|F(\Phi_t(x_{i-1})) - F(\gamma_t^i)\| + \frac{1}{t_i}\|x_i - \Phi_{t_i}(x_{i-1})\| \tag{E.81}$$

$$\le L\|\Phi_t(x_{i-1}) - \gamma_t^i\| + \frac{1}{t_0}\|x_i - \Phi_{t_i}(x_{i-1})\| \qquad (F \text{ is } L\text{-Lipschitz}) \tag{E.82}$$

$$\le \Big(L + \frac{1}{t_0}\Big)\|x_i - \Phi_{t_i}(x_{i-1})\| \tag{E.83}$$

$$\le 2\sqrt{\varepsilon}\Big(L + \frac{1}{t_0}\Big). \tag{E.84}$$

As a consequence, if ε is chosen small enough, for instance such that

$$\sqrt{\varepsilon} \le \delta\,\frac{t_0}{2(Lt_0+1)}, \tag{E.85}$$

then $\dot\gamma_t^i \in \bar{\mathbb{B}}(F(\gamma_t^i), \delta)$ for all $t \in [0, t_i]$ and every $i = 1, \dots, m$. Eq. (E.75) can therefore be used to bound the cost of γ^i as

$$\mathcal{S}_{0,t_i}(\gamma^i) = \int_0^{t_i} \mathcal{L}(\gamma_t^i, \dot\gamma_t^i)\,dt \le L' \int_0^{t_i} \|\dot\gamma_t^i - F(\gamma_t^i)\|^2\,dt \tag{E.86}$$

$$\le L'\Big(L + \frac{1}{t_0}\Big)^2 2(t_0+1)\,\|x_i - \Phi_{t_i}(x_{i-1})\|^2, \tag{E.87}$$

where we have used that $t_i \le 2(t_0+1)$ by assumption. Now, if we define a path $\gamma \in \mathcal{C}([0, \sum t_i], \mathcal{V})$ by

$$\gamma_t = \gamma^i_{t-T_{i-1}} \quad \text{for } t \in [T_{i-1}, T_i], \qquad \text{where } T_0 = 0 \text{ and } T_i = \sum_{j\le i} t_j, \tag{E.88}$$

then its cost can be estimated as

$$\mathcal{S}_{0,T}(\gamma) = \sum_{i=1}^m \int_{T_{i-1}}^{T_i} \mathcal{L}(\gamma_t, \dot\gamma_t)\,dt = \sum_{i=1}^m \int_0^{t_i} \mathcal{L}(\gamma_t^i, \dot\gamma_t^i)\,dt \tag{E.89}$$

$$= \sum_{i=1}^m \mathcal{S}_{0,t_i}(\gamma^i) \tag{E.90}$$

$$\le L'\Big(L + \frac{1}{t_0}\Big)^2 2(t_0+1) \sum_{i=1}^m \|x_i - \Phi_{t_i}(x_{i-1})\|^2 \tag{E.91}$$

$$\le 4L'\Big(L + \frac{1}{t_0}\Big)^2 (t_0+1)\,\varepsilon. \tag{E.92}$$

Consequently, if we choose $\varepsilon = 1/k$ with k large enough, then the construction above yields a sequence of paths $\gamma_k \in \mathcal{C}([0, T_k], \mathcal{V})$ from $x_0 = x$ to $x_m = y$ with $T_k \in \mathbb{N}$ such that $\mathcal{S}_{0,T_k}(\gamma_k) \to 0$ as $k \to \infty$, which directly implies that $B(x, x') = 0$ by definition. ∎

Corollary E.1. *If $x \rightsquigarrow y$, then $B(x, y) = 0$.*

Proof. This follows automatically from Lemmas E.11 and E.15. ∎

As a corollary, we obtain that the equivalence classes of the relation $\leftrightsquigarrow$ are included in the basic equivalence classes.

Corollary E.2. *Each equivalence class $[x]_{\leftrightsquigarrow}$ is included in a basic equivalence class.*

Proof. Fix $x \in \mathcal{V}$ such that $x \rightsquigarrow x$, and let $y \in [x]_{\leftrightsquigarrow}$. Then $x \rightsquigarrow y$ and $y \rightsquigarrow x$ by definition, and thus Corollary E.1 yields $B(x, y) = B(y, x) = 0$. It follows that $[x]_{\leftrightsquigarrow} \subseteq [x]_\sim$. As such, it only remains to show that the equivalence class $[x]_\sim$ is basic in the sense of Definition E.4. If $[x]_\sim$ is made of at least two elements then it is directly basic by definition, so we assume instead that $[x]_\sim = \{x\}$. We then need to show that x is necessarily a fixed point of the flow. Since $[x]_{\leftrightsquigarrow} \subseteq [x]_\sim$, this means that we also have $[x]_{\leftrightsquigarrow} = \{x\}$, but the equivalence classes of the relation $\leftrightsquigarrow$ are invariant under the flow (cf. Lemma E.7), and thus x is indeed a fixed point, which concludes the proof. ∎

To show the reverse inclusion, we first require a technical lemma relating the action cost of a path to the change in the potential g_∞ along it.

Lemma E.16. *Fix $T > 0$. Consider $\gamma \in \mathcal{C}_T(\mathcal{V})$ such that, for all $t \in [0, T]$,*

$$\|\gamma(t)\| \geq \sqrt{\frac{2C}{\kappa}}\,. \tag{E.93}$$

Then, the following holds:

$$g_\infty(\gamma_T) - g_\infty(\gamma_0) \leq \frac{4}{\kappa}\mathcal{S}_T[\gamma] \tag{E.94}$$

Proof. Let us begin by proving an algebraic identity. Fix $x \in \mathcal{V}$ and consider the supremum:

$$\sup_{w \in \mathcal{V}^*} \langle x, w\rangle - \frac{1}{\kappa}\|w - F(x)\|^2\,, \tag{E.95}$$

which is a quadratic optimization problem with a unique solution $w = F(x) + \frac{\kappa}{2}x$. Evaluating the objective at this point then yields that

$$\sup_{w \in \mathcal{V}^*} \langle x, w\rangle - \frac{1}{\kappa}\|w - F(x)\|^2 \tag{E.96}$$

$$= \langle x, F(x)\rangle + \frac{\kappa}{4}\|x\|^2 \tag{E.97}$$

$$\leq C - \frac{\kappa}{4}\|x\|^2\,. \tag{E.98}$$

by Assumption 4(a). If in addition $\|x\| \geq \sqrt{\frac{2C}{\kappa}}$, we thus obtain that, for any $w \in \mathcal{V}^*$,

$$\langle x, w\rangle \leq \frac{1}{\kappa}\|w - F(x)\|^2\,. \tag{E.99}$$

By Lemma D.4, multiplying the above by $\sigma_\infty^2(\|x\|^2)$ then yields that

$$\frac{\langle x, w\rangle}{\sigma_\infty^2(\|x\|^2)} \leq \frac{2}{\kappa}\mathcal{L}(x, w)\,. \tag{E.100}$$

Finally, using the definition of g_∞ (Definition D.1), we have that

$$g_\infty(\gamma_T) - g_\infty(\gamma_0) = 2\int_0^T \frac{\langle \dot\gamma(t), \gamma(t)\rangle}{\sigma_\infty^2(\|\gamma(t)\|^2)}\,dt \tag{E.101}$$

$$\leq \frac{4}{\kappa}\int_0^T \mathcal{L}(\gamma(t), \dot\gamma(t))\,dt = \frac{4}{\kappa}\mathcal{S}_T[\gamma]\,. \tag{E.102}$$

■

We also need the following auxiliary result, which establishes that the paths considered in the definition of B can be restricted to a compact set.

Lemma E.17. *For every compact $\mathcal{K} \subseteq \mathcal{V}$ and all $x, y \in \mathcal{K}$, define*

$$B_{\mathcal{K}}(x, y) \coloneqq \inf\left\{\mathcal{S}_{[0,T]}(\gamma) : \gamma \in \mathcal{C}([0, T], \mathcal{K}), \gamma_0 = x, \gamma_T = y, T \in \mathbb{N}\right\}. \tag{E.103}$$

Then, for any $\mathcal{K}_1, \mathcal{K}_2 \subseteq \mathcal{V}$ compact, there exists a compact $\mathcal{K} \subseteq \mathcal{V}$ containing both $\mathcal{K}_1$ and $\mathcal{K}_2$ such that $B(\mathcal{K}_1, \mathcal{K}_2) = B_{\mathcal{K}}(\mathcal{K}_1, \mathcal{K}_2)$, where

$$B(\mathcal{K}_1, \mathcal{K}_2) \coloneqq \inf_{x \in \mathcal{K}_1, y \in \mathcal{K}_2} B(x, y) \quad \text{and} \quad B_{\mathcal{K}}(\mathcal{K}_1, \mathcal{K}_2) \coloneqq \inf_{x \in \mathcal{K}_1, y \in \mathcal{K}_2} B_{\mathcal{K}}(x, y). \tag{E.104}$$

Proof. Fix $\mathcal{K}_1, \mathcal{K}_2 \subseteq \mathcal{V}$ some compact sets. Note that $B(x, y) \leq B_{\mathcal{K}}(x, y)$ for every $x \in \mathcal{K}_1$, $y \in \mathcal{K}_2$, and all compact $\mathcal{K} \subseteq \mathcal{V}$ by definition. In particular, the direction $B(\mathcal{K}_1, \mathcal{K}_2) \leq B_{\mathcal{K}}(\mathcal{K}_1, \mathcal{K}_2)$ is immediate, and it is sufficient to assume $B(\mathcal{K}_1, \mathcal{K}_2) < \infty$ throughout the proof.

With that said let us define $M_1 \coloneqq \max_{x\in\mathcal{K}_1\cup\mathcal{K}_2}\|x\| + \sqrt{2C/\kappa}$. The compact set $\mathcal{K}' \coloneqq \bar{\mathbb{B}}(0, M_1)$ therefore contains both $\mathcal{K}_1$ and $\mathcal{K}_2$. For reasons that will become apparent shortly, we also choose a constant $M_2 > M_1$ large enough such that

$$\alpha_\infty(M_2^2) > \alpha_\infty(M_1^2) + 2\frac{B(\mathcal{K}_1,\mathcal{K}_2)+2}{\kappa} \tag{E.105}$$

and consider its associated closed ball $\mathcal{K} \coloneqq \bar{\mathbb{B}}(0, M_2)$. Such a constant indeed exists since $\alpha_\infty(t)$ is strictly increasing and converges to $+\infty$ (cf. Definition D.1).

We begin by showing that $B_{\mathcal{K}}(x,y) \le B(x,y)$ for all elements $x \in \mathcal{K}_1$, $y \in \mathcal{K}_2$ such that $B(x,y) < B(\mathcal{K}_1,\mathcal{K}_2) + 1$. To do so, we use the definition of B to obtain paths $\gamma^k \in \mathcal{C}([0,T_k],\mathcal{V})$ such that $\gamma^k_0 = x$, $\gamma^k_{T_k} = y$, $T_k \ge 1$ and $\mathcal{S}_{[0,T_k]}(\gamma^k) \le B(x,y) + 1/k$. Now, assume that there are $k \ge 2$ and $t^* \coloneqq t^*(k) \in [0,T_k]$ such that $\gamma^k_{t^*} \notin \mathcal{K}$. Then, by continuity of γ^k and since $\mathcal{K}_1 \subseteq \mathcal{K}' \subset \mathcal{K}$, there exists a time $t_0 \coloneqq t_0(k) \in [0,t^*)$ such that $\gamma^k_{t_0} \in \operatorname{bd}\mathcal{K}'$ and $\gamma^k_t \notin \mathcal{K}'$ for every $t \in (t_0,t^*]$. In particular, $\|\gamma^k_t\| \ge \sqrt{2C/\kappa}$ for all $t \in [t_0,t^+]$ by definition of $\mathcal{K}'$ and M_1, and thus we can invoke Lemma E.16 to get

$$\begin{aligned}
\mathcal{S}_{[0,T_k]}(\gamma^k) \ge \mathcal{S}_{[t_0,t^*]}(\gamma^k) &\ge \frac{\kappa}{4}\left[g_\infty(\gamma^k_{t^*}) - g_\infty(\gamma^k_{t_0})\right] && \text{(E.106)}\\
&= \frac{\kappa}{2}\left[\alpha_\infty(\|\gamma^k_{t^*}\|^2) - \alpha_\infty(\|\gamma^k_{t_0}\|^2)\right] && \text{(E.107)}\\
&\ge \frac{\kappa}{2}\left[\alpha_\infty(M_2^2) - \alpha_\infty(M_1^2)\right] && \text{(E.108)}\\
&> B(\mathcal{K}_1,\mathcal{K}_2) + 2. && \text{(E.109)}
\end{aligned}$$

But at the same time, we also have $\mathcal{S}_{[0,T_k]}(\gamma^k) \le B(x,y) + 1/k < B(\mathcal{K}_1,\mathcal{K}_2) + 2$, so this leads to a contradiction. It follows that $\gamma^k_t \in \mathcal{K}$ for every $0 \le t \le T_k$ and $k \ge 1$. Moreover, since $\inf_{k\ge1}\mathcal{S}_{[0,T_k]}(\gamma^k) \le B(x,y)$ and all paths γ^k remain in $\mathcal{K}$, this directly yields $B_{\mathcal{K}}(x,y) \le B(x,y)$ as needed.

To conclude the proof, it suffices to take a sequence $(x_\ell,y_\ell) \in \mathcal{K}_1\times\mathcal{K}_2$ such that $B(x_\ell,y_\ell) \le B(\mathcal{K}_1,\mathcal{K}_2)+1/\ell$. Indeed, this sequence then satisfies $B_{\mathcal{K}}(x_\ell,y_\ell) \le B(x_\ell,y_\ell) \le B(\mathcal{K}_1,\mathcal{K}_2)+1/\ell$ from our previous computations, and thus $B_{\mathcal{K}}(\mathcal{K}_1,\mathcal{K}_2) \le B(\mathcal{K}_1,\mathcal{K}_2)$ by taking the limit as $\ell\to\infty$. ∎

To establish the inclusion $[x]_\sim \subseteq [x]_{\leftrightsquigarrow}$, we first show any two points within a basic equivalence class can be connected by an arbitrarily long null-energy path.

Lemma E.18. *Let $[x]_\sim$ be a basic equivalence class. Then, for every $y \in [x]_\sim$, there exists a compact set $\mathcal{K} \subseteq \mathcal{V}$ such that $h_{\mathcal{K}}(x,y) = h_{\mathcal{K}}(y,x) = 0$, where*

$$h_{\mathcal{K}}(x,y) \coloneqq \liminf_{T\in\mathbb{N},T\to\infty}\{\mathcal{S}_{[0,T]}(\gamma) : \gamma\in\mathcal{C}([0,T],\mathcal{K}),\gamma_0 = x,\gamma_T = y\}. \tag{E.110}$$

Proof. Fix $[x]_\sim$ a basic equivalence class with reference state $x \in \mathcal{V}$. If $[x]_\sim$ is made of a single point, that is, if $[x]_\sim = \{x\}$, then x is necessarily a fixed point of the flow, cf. Definition E.4. In particular, $h_{\mathcal{K}}(x,x) = 0$ for any compact $\mathcal{K} \subseteq \mathcal{V}$ containing x since it suffices to take $\gamma = \Phi_\cdot(x)$ in the definition of $h_{\mathcal{K}}$.

Assume on the other hand that $[x]_\sim$ contains at least two points, and let $y \ne x$ such that $x \sim y$. Applying Lemma E.17 then yields a compact set $\mathcal{K} \subseteq \mathcal{V}$ such that $B_{\mathcal{K}}(x,y) = B(x,y) = 0$ and $B_{\mathcal{K}}(y,x) = B(y,x) = 0$. Let $\varepsilon > 0$, $T_0 > 0$ be fixed. Since $x \ne y$, this implies that there exist paths $\gamma \in \mathcal{C}([0,T_1],\mathcal{K})$ and $\varphi \in \mathcal{C}([0,T_2],\mathcal{K})$ with $T_1,T_2 \ge 1$ and such that $\gamma_0 = \varphi_{T_2} = x$, $\gamma_{T_1} = \varphi_0 = y$, $\mathcal{S}_{[0,T_1]}(\gamma) \le \varepsilon/(T_0+3)$ and $\mathcal{S}_{[0,T_2]}(\varphi) \le \varepsilon/(T_0+3)$. Note that if one follows γ then φ, this leads to a path from x to x with length T_1+T_2. As a consequence, if we let $k \in \mathbb{N}$ be the integer such that

$$(k-1)(T_1+T_2) + T_1 \le T_0 < k(T_1+T_2) + T_1, \tag{E.111}$$

then we can construct a path $\psi \in \mathcal{C}([0,T],\mathcal{K})$ from x to y with length $T = k(T_1+T_2)+T_1 \ge T_0$ by cycling k times from x to x and then by following γ toward y. The total cost of this path is

$$\mathcal{S}_{[0,T]}(\psi) \le (2k+1)\frac{\varepsilon}{T_0+3} \le \left(2\frac{T_0-T_1}{T_1+T_2}+2\right)\frac{\varepsilon}{T_0+3} \le \frac{T_0+2}{T_0+3} < \varepsilon, \tag{E.112}$$

and therefore $h_{\mathcal{K}}(x, y) = 0$ since the above holds true for all ε sufficiently small and every T_0 large enough. Moreover, if we cycle back from y to x by following φ, then this construction also yields a path from x to x with length greater than T_0 and cost smaller than ε. Hence $h_{\mathcal{K}}(x, x) = 0$. Finally, by inverting the roles of x and y we also get $h_{\mathcal{K}}(y, x) = 0$, which concludes the proof. ∎

We then establish that the existence of an arbitrarily long null-energy path is a sufficient condition for the existence of an ℓ^2-path.

Lemma E.19. *Let $\mathcal{K} \subseteq \mathcal{V}$ a compact set and $x, y \in \mathcal{K}$. If $h_{\mathcal{K}}(x, y) = 0$, then $x \rightsquigarrow y$.*

Proof. Let $x, y \in \mathcal{K}$ such that $h_{\mathcal{K}}(x, y) = 0$ and fix $\varepsilon > 0$, $T_0 > 0$. In particular, by definition of h, there exists a path $\gamma \in \mathcal{C}([0, T], \mathcal{K})$ such that $\gamma_0 = x$, $\gamma_T = y$, $T \geq T_0$ and $\mathcal{S}_{[0,T]}(\gamma) < \varepsilon/\Delta$, where

$$\Delta \coloneqq 2C(2T_0)\sigma_\infty^2(\max_{z \in \mathcal{K}} \|z\|^2) \tag{E.113}$$

and $C(2T_0)$ is the constant given by Lemma E.4. Let $m \geq 1$ be the integer such that $mT_0 \leq T < (m+1)T_0$, and let $t_1, \dots, t_m$ be the sequence defined as

$$t_i = \begin{cases} T_0 & \text{if } 1 \leq i \leq m-1, \\ T - (m-1)T_0 & \text{if } i = m. \end{cases} \tag{E.114}$$

Note that $\sum_i t_i = T$ and $T_0 \leq t_i \leq 2T_0$ for every $1 \leq i \leq m$ by construction. We then set $\tau_i = \sum_{j \leq i} t_i$ and define a sequence $x_0, \dots, x_m \in \mathcal{K} \subseteq \mathcal{V}$ by $x_0 = x$ and $x_i = \gamma_{\tau_i}$ for all $1 \leq i \leq m$. Since $\sum_i t_i = T$, this sequence also has $x_m = \gamma_T = y$.

Fix $1 \leq i \leq m$. Then an application of Lemma E.4 with $T \leftarrow 2T_0$ and $t \leftarrow t_i$ to the path $\gamma_t \leftarrow \gamma_{\tau_{i-1}+t}$ started at $\gamma_{\tau_{i-1}}$ yields

$$\begin{aligned} \|\Phi_{t_i}(x_{i-1}) - x_i\|^2 = \|\Phi_{t_i}(\gamma_{\tau_{i-1}}) - \gamma_{\tau_i}\|^2 &= \|\Phi_{t_i}(\gamma_{\tau_{i-1}}) - \gamma_{\tau_{idx-1}+t_i}\|^2 && \text{(E.115)} \\ &\leq C(2T_0)\int_0^{t_i} \|\dot\gamma_{\tau_{i-1}+s} - F(\gamma_{\tau_{i-1}+s})\|^2 ds && \text{(E.116)} \\ &= C(2T_0)\int_{\tau_{i-1}}^{\tau_i} \|\dot\gamma_s - F(\gamma_s)\|^2 ds. && \text{(E.117)} \end{aligned}$$

Summing over all indices $1 \leq i \leq m$, we therefore get

$$\sum_{i=1}^m \|\Phi_{t_i}(x_{i-1}) - x_i\|^2 \leq C(2T_0)\int_0^T \|\dot\gamma_s - F(\gamma_s)\|^2 ds, \tag{E.118}$$

where we have used that $\tau_m = \sum_i t_i = T$ by construction. Moreover, thanks to the Sub-Gaussian bound of Lemma D.4 we also have

$$\|\dot\gamma_s - F(\gamma_t)\|^2 \leq 2\sigma_\infty^2(\|\gamma_s\|^2)\mathcal{L}(\gamma_s, \dot\gamma_s). \tag{E.119}$$

Since the path γ remains in the compact $\mathcal{K}$ and since σ_∞^2 is an increasing continuous function (cf. Definition D.1), it follows that $\sigma_\infty^2(\|\gamma_s\|^2) \leq \sigma_\infty^2(\max_{z \in \mathcal{K}} \|z\|^2) < \infty$ and thus

$$\sum_{i=1}^m \|\Phi_{t_i}(x_{i-1}) - x_i\|^2 \leq 2C(2T_0)\sigma_\infty^2(\max_{z \in \mathcal{K}} \|z\|^2)\int_0^T \mathcal{L}(\gamma_s, \dot\gamma_s) ds = \Delta\, \mathcal{S}_{[0,T]}(\gamma) < \varepsilon. \tag{E.120}$$

We have therefore constructed a chain from x to y composed of states $x_0, \dots, x_m \in \mathcal{V}$ and times $t_1, \dots, t_m \geq T_0$ with total L^2-cost (strictly) smaller than ε, hence proving that $x \rightsquigarrow y$. ∎

Combining Lemmas E.18 and E.19 therefore yields the desired inclusion.

Corollary E.3. *Each basic equivalence class is included in an equivalence class $[x]_{\leftrightsquigarrow}$.*

Proof. Let $[x]_\sim$ be a basic equivalence with reference state $x \in \mathcal{V}$. Then, for every $y \in [x]_\sim$, Lemma E.18 implies that there is a compact set $\mathcal{K} \subseteq \mathcal{V}$ such that $h_{\mathcal{K}}(x, y) = h_{\mathcal{K}}(y, x) = 0$, and Lemma E.19 therefore yields both $x \rightsquigarrow y$ and $y \rightsquigarrow x$. Since this is true in particular for $y = x$, it follows that $x \leftrightsquigarrow y$ for all $y \in [x]_\sim$ and thus $[x]_\sim \subseteq [x]_{\leftrightsquigarrow}$. ∎

Consequently, we obtain that the basic equivalence classes and the equivalence classes of $\leftrightsquigarrow$ coincide.

Theorem E.1. *The equivalence classes $[x]_{\leftrightsquigarrow}$ and the basic equivalence classes coincide.*

Proof. Follows from Corollaries E.2 and E.3. ∎

Remark E.1. Since the ℓ^2-chains equivalence classes coincide with the basic equivalence classes—which correspond exactly to the metastable components introduced in the main text—we shall omit the "basic" and "ℓ^2-chains" prefixes and simply refer to them as *the equivalence classes* hereafter.

E.4. Regularity results. In this section we establish several regularity results required for our subsequent analysis.

We start by stating a basic property of the noise Lagrangian, establishing its local regularity.

Lemma E.20 (Azizian et al. [9, Lemma D.3])**.** *For any $x \in \mathcal{V}$, there exist $\delta > 0$ and finite constants $L', L'' > 0$ such that*

1. *$\bar{\mathcal{L}}$ is finite and jointly L'-Lipschitz continuous on $\mathbb{B}(x,\delta) \times \mathbb{B}(0,\delta)$;*
2. *$\bar{\mathcal{L}}(x', w) \le L''\|w\|^2$ for every $(x', w) \in \mathbb{B}(x,\delta) \times \mathbb{B}(0,\delta)$.*

Under the framework of Azizian et al. [9], the flow is generated by a gradient field $-\nabla f$; consequently, $\rho(x,x) = 0$ whenever x is a critical point of f. Because the basic equivalences classes reduce to the connected components of the critical points in this setting, ρ is uniformly continuous on neighborhoods of these classes [9, Lemma D.4]. However, under our general vector field setup, the basic equivalence classes are no longer restricted to fixed points, rendering the derivation of regularity results for ρ substantially more intricate.

The following result establishes that $\rho(x,\cdot)$ (resp. $\rho(\cdot,y)$) remains continuous, but only at $\phi(x)$ (resp. $\phi^{-1}(y)$) instead of x.

Lemma E.21. *Fix $x, y \in \mathcal{V}$. Then, for every $\varepsilon > 0$, there exists $\delta > 0$ such that*

1. *$\rho(x,z) \le \varepsilon$ for all $z \in \mathbb{B}(\phi(x),\delta)$;*
2. *$\rho(z,y) \le \varepsilon$ for all $z \in \mathbb{B}(\phi^{-1}(y),\delta)$.*

Proof. 1. Fix $x \in \mathcal{V}$ and $\varepsilon > 0$. By continuity of the flow, the set $\{\Phi_t(x) : t \in [0,1]\}$ is compact and we can then consider its enlargement

$$\mathcal{K} = \{z \in \mathcal{V} : d\big(z, \{\Phi_t(x)\}_{t\in[0,1]}\big) \le 1\}, \tag{E.121}$$

which is also compact. As a result, applying Lemma E.20 at each $z \in \mathcal{K}$ and extracting a finite cover thereof yields that the Lagrangian $\bar{\mathcal{L}}$ is jointly Lipschitz continuous on $\mathcal{K} \times \mathbb{B}(0,\delta')$ for some $\delta' > 0$.

Now, let us fix $\delta \in (0,1)$ to be determined later and let $z \in \mathbb{B}(\phi(x),\delta)$. We then define a path $\gamma \in \mathcal{C}([0,1],\mathcal{V})$ from x to z by

$$\gamma_t = \Phi_t(x) + t(z - \phi(x)). \tag{E.122}$$

First, note that

$$\|\gamma_t - \Phi_t(x)\| \le t\,\|z - \phi(x)\| \le \delta < 1, \tag{E.123}$$

so $\gamma_t \in \mathcal{K}$ for every $t \in [0,1]$. Moreover, this path also satisfies

$$\begin{aligned}
\|\dot\gamma_t - F(\gamma_t)\| = \|z - \phi(x) + F(\Phi_t(x)) - F(\gamma_t)\| &\le \|z - \phi(x)\| + \|F(\Phi_t(x)) - F(\gamma_t)\| && \text{(E.124)}\\
&\le \|z - \phi(x)\| + L\,\|\Phi_t(x) - \gamma_t\| && \text{(E.125)}\\
&\le (1 + tL)\,\|z - \phi(x)\| && \text{(E.126)}\\
&< (1 + L)\delta, && \text{(E.127)}
\end{aligned}$$

where L is the Lipschitz constant of F. As a consequence, if we let L' be the Lipschitz constant of $\bar{\mathcal{L}}$ on $\mathcal{K} \times \mathbb{B}(0,\delta')$ and if we pick $\delta = \frac{1}{L+1}\min(\delta', \varepsilon/L')$, then we obtain that

$$\mathcal{L}(\gamma_t, \dot\gamma_t) = \bar{\mathcal{L}}\big(\gamma_t, \dot\gamma_t - F(\gamma_t)\big) \le L'\|\dot\gamma_t - F(\gamma_t)\| \le L'(L+1)\delta \le \varepsilon. \tag{E.128}$$

This directly leads to the estimate

$$\rho(x,z) \le \mathcal{S}_{0,1}(\gamma) = \int_0^1 \mathcal{L}(\gamma_t, \dot{\gamma}_t)\, dt \le \varepsilon, \tag{E.129}$$

which proves our claim.

2. This can be shown using the exact same arguments by considering the compact set

$$\mathcal{K} = \{z \in \mathcal{V} : d(z, \{\Phi_t(\phi^{-1}(y))\}_{t\in[0,1]}) \le 1\}, \tag{E.130}$$

and by defining a path γ from $z \in \mathbb{B}(\phi^{-1}(y), \delta)$ to y as

$$\gamma_t = \Phi_t\big(\phi^{-1}(y)\big) + (1-t)\big(z - \phi^{-1}(y)\big) \tag{E.131}$$

instead of the one given in Eq. (E.122). ∎

Leveraging Lemma E.21 alongside the invariance and compactness of the equivalence classes, we obtain the continuity of the rate ρ at these classes.

Lemma E.22. *Let $\mathcal{K}$ be an equivalence class of the relation $\leftrightsquigarrow$. Then, for every $\varepsilon > 0$, there exists $\delta > 0$ such that*

$$\max(\rho(z,\mathcal{K}),\, \rho(\mathcal{K},z)) \le \varepsilon \quad \textit{whenever } d(z,\mathcal{K}) < \delta, \tag{E.132}$$

where

$$\rho(z,\mathcal{K}) := \inf_{y\in\mathcal{K}} \rho(z,y) \quad \textit{and} \quad \rho(\mathcal{K},z) := \inf_{x\in\mathcal{K}} \rho(x,z). \tag{E.133}$$

Proof. Let $\mathcal{K}$ be an equivalence class of the relation $\leftrightsquigarrow$ and fix $\varepsilon > 0$. By Lemma E.21 we know that:

1. For every $x \in \mathcal{K}$, there exists $\delta(x) > 0$ such that $\rho(x,z) \le \varepsilon$ whenever $z \in \mathbb{B}(\phi(x), \delta(x))$.
2. For every $y \in \mathcal{K}$, there exists $\delta(y) > 0$ such that $\rho(z,y) \le \varepsilon$ whenever $z \in \mathbb{B}\big(\phi^{-1}(y), \delta(y)\big)$.

In particular, since $\mathcal{K}$ is invariant under the flow (cf. Lemma E.7), $\phi^{-1}(\mathcal{K}) = \mathcal{K} = \phi(\mathcal{K})$ and we have

$$\mathcal{K} = \phi(\mathcal{K}) \subseteq \bigcup_{x\in\mathcal{K}} \mathbb{B}(\phi(x), \delta(x)/2) \quad \text{and} \quad \mathcal{K} = \phi^{-1}(\mathcal{K}) \subseteq \bigcup_{y\in\mathcal{K}} \mathbb{B}(\phi^{-1}(y), \delta(y)/2). \tag{E.134}$$

We can therefore write

$$\mathcal{K} \subseteq \left(\bigcup_{x\in\mathcal{K}} \mathbb{B}(\phi(x), \delta(x)/2)\right) \bigcap \left(\bigcup_{y\in\mathcal{K}} \mathbb{B}(\phi^{-1}(y), \delta(y)/2)\right) \tag{E.135}$$

$$= \bigcup_{x,y\in\mathcal{K}} \mathbb{B}(\phi(x), \delta(x)/2) \cap \mathbb{B}(\phi^{-1}(y), \delta(y)/2). \tag{E.136}$$

This means that the open sets

$$\mathcal{U}(x,y) := \mathbb{B}(\phi(x), \delta(x)/2) \cap \mathbb{B}(\phi^{-1}(y), \delta(y)/2) \quad x, y \in \mathcal{K}, \tag{E.137}$$

form an open cover of $\mathcal{K}$. But since $\mathcal{K}$ is compact thanks to Lemma E.10, it follows that there exists a finite number of elements $x_i, y_i \in \mathcal{K}$, $i \in \mathcal{I}$, such that $\mathcal{K} \subseteq \bigcup_{i\in\mathcal{I}} \mathcal{U}(x_i, y_i)$. Let $\delta = \min_{i\in\mathcal{I}} \min(\delta(x_i), \delta(y_i))$, which is a strictly positive constant since $\mathcal{I}$ is finite, and assume that $z \in \mathcal{V}$ satisfies $d(z,\mathcal{K}) < \delta/2$. Then, by definition, there exists $z' \in \mathcal{K}$ such that $\|z - z'\| < \delta/2$. Moreover, since $z' \in \mathcal{K}$, there is an index $i \in \mathcal{I}$ for which $z' \in \mathcal{U}(x_i, y_i)$. In particular,

$$\|z - \phi(x_i)\| \le \|z - z'\| + \|z' - \phi(x_i)\| < \frac{1}{2}(\delta + \delta(x_i)) \le \delta(x_i) \tag{E.138}$$

and

$$\|z - \phi^{-1}(y_i)\| \le \|z - z'\| + \|z' - \phi^{-1}(y_i)\| < \frac{1}{2}(\delta + \delta(y_i)) \le \delta(y_i), \tag{E.139}$$

which directly leads to

$$\rho(\mathcal{K}, z) \le \rho(x_i, z) \le \varepsilon \quad \text{and} \quad \rho(z,\mathcal{K}) \le \rho(z, y_i) \le \varepsilon, \tag{E.140}$$

and thus concludes the proof. ∎

Using this lemma, we now define the following neighborhoods of an equivalence class $\mathcal{K}$.

Corollary E.4. *For any $\mathcal{K}$ equivalence class of $\leftrightsquigarrow$ and $r > 0$, the set*

$$\mathcal{W}_r(\mathcal{K}) := \{x \in \mathcal{V} : \rho(x, \mathcal{K}) < r,\ \rho(\mathcal{K}, x) < r\} \tag{E.141}$$

is a neighborhood of $\mathcal{K}$.

Proof. Fix $\mathcal{K}$ an equivalence class of $\leftrightsquigarrow$. In virtue of Lemma E.22, for every $r > 0$, there exists $\delta := \delta(r) > 0$ such that $\mathcal{U}_\delta(\mathcal{K}) := \{z \in \mathcal{V} : d(z, \mathcal{K}) < \delta\} \subseteq \mathcal{W}_r(\mathcal{K})$. Since $\mathcal{U}_\delta(\mathcal{K})$ is an open set containing $\mathcal{K}$, it follows that $\mathcal{W}_r(\mathcal{K})$ is indeed a neighborhood of $\mathcal{K}$. ∎

We will also need the following two regularity results on B in the sequel.

Lemma E.23. *Let $x, y \in \mathcal{V}$.*

1. *If $\liminf_{z\to y} B(x, z) = 0$ then $\liminf_{z\to\phi^{-1}(y)} B(x, z) = 0$.*
2. *If $\liminf_{z\to x} B(z, y) = 0$ then $\liminf_{z\to\phi(x)} B(z, y) = 0$.*

Proof. 1. Fix $x, y \in \mathcal{V}$ and let (y_n) be a sequence such that $y_n \to y$ and $B(x, y_n) \to 0$. In particular, by definition of B, this means that there exist paths $\gamma^n \in \mathcal{C}([0, T_n], \mathcal{V})$ from $\gamma^n_0 = x$ to $\gamma^n_{T_n} = y_n$ such that $T_n \in \mathbb{N}$, $T_n \geq 1$ and $\mathcal{S}_{0,T_n}(\gamma^n) \leq 1/n$. The main idea behind this proof is to show that $\gamma^n_{T_n-1} \to \phi^{-1}(y)$, which would directly yields the desired result.

Since $\mathcal{S}_{0,T_n}(\gamma^n) < \infty$, it follows that $\dot\gamma^n_t \in \operatorname{dom}\mathcal{L}(\gamma^n_t, \cdot)$ for almost every $t \in [0, T^n]$, and thus Lemma D.3 leads to

$$\|\dot\gamma^n_t\| \leq 2M\big(1 + \|\gamma^n_t\|\big) \quad \text{a.e.} \tag{E.142}$$

We can then apply Grönwall's inequality to the time-reversed path $t \mapsto \gamma^n_{T_n-t}$ to get

$$\|\gamma^n_{T_n-t}\| \leq e^{2Mt}\Big(2Mt + \|\gamma^n_{T_n}\|\Big) = e^{2Mt}(2Mt + \|y_n\|). \tag{E.143}$$

In particular, picking $t = 1$ and N big enough such that $\|y_n\| \leq \|y\| + M$ for all $n \geq N$ yields

$$\|\gamma^n_{T_n-1}\| \leq e^{2M}(3Mt + \|y\|). \tag{E.144}$$

As a consequence, the set $\{\gamma^n_{T_n-1} : n \geq N\}$ is relatively compact in $\mathcal{V}$. Now, if we extract the part from γ that spans from $t = T_n - 1$ to $t = T_n$, then this results to a path $\varphi^n \in \mathcal{C}([0, 1], \mathcal{V})$ defined as $\varphi^n_t = \gamma^n_{T_n-1+t}$ and for which $\mathcal{K} := \operatorname{cl}\{\varphi^n_0 : n \geq N\}$ is a compact set. Moreover it also satisfies $\mathcal{S}_{0,1}(\varphi^n) \leq \mathcal{S}_{0,T_n}(\gamma^n) \leq 1$, and thus $\varphi^n \in \Gamma^{\mathcal{K}}_{0,1}(1)$, which is compact by Corollary D.1. Passing to a subsequence if needed, we can therefore assume that φ^n converges to some path $\varphi \in \mathcal{C}([0, 1], \mathcal{V})$ in the supremum topology. Furthermore, by lower semi-continuity of $\mathcal{S}_{0,1}$ (cf. Corollary D.1), we also get

$$\mathcal{S}_{0,1}(\varphi) \leq \liminf_{n\to\infty} \mathcal{S}_{0,1}(\varphi^n) \leq \liminf_{n\to\infty} \mathcal{S}_{0,T_n}(\gamma^n) = 0, \tag{E.145}$$

so that $\rho(\varphi_0, \varphi_1) = 0$. By the flow characterization of Lemma E.2, this implies that

$$y = \lim_{n\to\infty} y_n = \lim_{n\to\infty} \gamma^n_{T_n} = \lim_{n\to\infty} \varphi^n_1 = \varphi_1 = \phi(\varphi_0), \tag{E.146}$$

and thus $\varphi_0 = \phi^{-1}(y)$. However we also have $\varphi_0 = \lim_{n\to\infty} \gamma^n_{T_n-1}$ by construction, which finally leads to $\gamma^n_{T_n-1} \to \phi^{-1}(y)$ as needed. Consequently, if we let $z_n = \gamma^n_{T_n-1}$, then $z_n \to \phi^{-1}(y)$ and

$$B(x, z_n) \leq \mathcal{S}_{0,T_n-1}(\gamma^n) \leq \mathcal{S}_{0,T_n}(\gamma^n) \to 0, \tag{E.147}$$

hence proving our claim.

2. Proving the second claim is done in a similar fashion by extracting the part $[0, 1]$ of the path in order to show that γ^n_1 converges to $\phi(x)$. ∎

Lemma E.24. *Let $x, y \in \mathcal{V}$.*

1. *If $\liminf_{z\to y} B(x, z) = 0$ then $B(x, y) = 0$.*

2. *If* $\liminf_{z\to x} B(z,y) = 0$ *then* $B(x,y) = 0$.

Proof. 1. Fix $x, y \in \mathcal{V}$ and assume that $\liminf_{z\to y} B(x,z) = 0$. Then Lemma E.23 implies that $\liminf_{z\to\phi^{-1}(y)} B(x,z) = 0$ and thus there exist a sequence of states (z_n) and of paths $\gamma^n \in \mathcal{C}([0,T_n],\mathcal{V})$ such that $\gamma_0^n = x$, $\gamma_{T_n}^n = z_n \to \phi^{-1}(y)$, $T_n \in \mathbb{N}$ and $\mathcal{S}_{0,T_n}(\gamma^n) \to 0$. We will show that these paths can be prolonged with a portion of length 1 that goes from z_n to y with an arbitrarily small cost.

To do so, note that for any fixed $\varepsilon > 0$, the regularity of ρ established in Lemma E.21 provides a positive radius $\delta > 0$ such that $\rho(z,y) \le \varepsilon$ whenever $z \in \mathbb{B}(\phi^{-1}(y),\delta)$. In particular, if we let $N(\varepsilon)$ be large enough such that $z_n \in \mathbb{B}(\phi^{-1}(y),\delta)$ for all $n \ge N(\varepsilon)$ (which is always possible since $z_n \to \phi^{-1}(y)$ by construction), then $\rho(z_n,y) \le \varepsilon$, which means that there exists a path $\varphi^n \in \mathcal{C}([0,1],\mathcal{V})$ with $\varphi_0^n = z_n$, $\varphi_1^n = y$ and $\mathcal{S}_{0,1}(\varphi^n) = \rho(z,y) \le \varepsilon$. For any $n \ge N(\varepsilon)$, concatenating γ^n and φ^n therefore yields a new path $\psi^n \in \mathcal{C}([0,T_n+1],\mathcal{V})$ from x to y such that

$$\mathcal{S}_{0,T_n+1}(\psi^n) = \mathcal{S}_{0,T_n}(\gamma^n) + \mathcal{S}_{0,1}(\varphi^n) \le \mathcal{S}_{0,T_n}(\gamma^n) + \varepsilon \overset{n\to\infty}{\longrightarrow} \varepsilon. \tag{E.148}$$

Since $\varepsilon > 0$ was fixed arbitrarily, it follows that $\lim_{n\to\infty} \mathcal{S}_{0,T_n+1}(\psi^n) = 0$ and thus $B(x,y) = 0$.

2. The second point is obtained by invoking the same kinds of arguments. ∎

In particular, the preceding lemma allows us to prove the following result, which extends Azizian et al. [9, Lemma D.9] to accommodate our general setting, where the equivalence classes are no longer restricted to fixed points and the dynamical system is not a gradient flow.

Lemma E.25. *Let $\mathcal{K}$ be an equivalence class of $\leftrightsquigarrow$. Then, for any $\epsilon > 0$, there is some $T_0 \ge 1$ such that, for any $x, y \in \mathcal{K}$, there is $\xi \in (\mathcal{V})^T$ such that $T \le T_0$, $\xi_0 = x$, $\xi_{T-1} = y$, $\mathcal{A}_T(\xi) < \epsilon$ and* $\max_{0\le n<T} d(\xi_n,\mathcal{K}) < \epsilon$.

Proof. Fix $\epsilon > 0$. Following the proof of Lemma E.22, we know that there is a finite number of elements $x_i, y_i \in \mathcal{K}$, $i \in \mathcal{I}$, such that

$$\mathcal{K} \subseteq \bigcup_{i\in\mathcal{I}} \mathbb{B}(\phi(x_i),\delta(x_i)) \cap \mathbb{B}(\phi^{-1}(y_i),\delta(y_i)), \tag{E.149}$$

where $\delta(x_i)$ and $\delta(y_i)$ are positive constants for which

$$\rho(x_i,z) < \epsilon \text{ when } z \in \mathbb{B}(\phi(x_i),\delta(x_i)) \quad \text{and} \quad \rho(z,y_i) < \epsilon \text{ when } z \in \mathbb{B}(\phi^{-1}(y_i),\delta(y_i)).$$

We first show that the result holds for the points y_i, x_j, $i, j \in \mathcal{I}$, before explaining why it actually suffices for the general case.

Fix $i, j \in \mathcal{I}$. Since $y_i \rightsquigarrow x_j$, $B(y_i,x_j) = 0$ by Corollary E.1 and therefore there exist sequences $\xi^k \in (\mathcal{V})^{T_k}$ with $\xi_0^k = y_i, \xi_{T_k-1}^k = x_j$ such that $\mathcal{A}_{T_k}(\xi^k) \to 0$ as $k \to \infty$. For the sake of contradiction, assume that from some k, there always exists $0 \le n < T_k$ such that $d(\xi_n^k,\mathcal{K}) \ge \epsilon$. Define n_k as the smallest n such that it happens, which is necessarily greater or equal to 1. Note that, by definition, $\xi_{n_k-1}^k$ satisfies $d(\xi_{n_k-1}^k,\mathcal{K}) \le \epsilon$.

Define $\mathcal{K}' := \{x \in \mathcal{V} : d(x,\mathcal{K}) \le \varepsilon\}$ which is compact. However, for k large enough, $\rho(\xi_{n_k-1}^k,\xi_{n_k}^k) \le \mathcal{A}_{T_k}(\xi^k) \le 1$ so that $(\xi_{n_k-1}^k,\xi_{n_k}^k)$ belongs to $\Gamma_2^{\mathcal{K}'}(1)$, which is compact by Corollary D.2. By passing to a subsequence if needed, the sequence $(\xi_{n_k-1}^k,\xi_{n_k}^k)_{k\ge1}$ therefore converges to some $(x,y) \in \mathcal{V}^2$ that satisfies $d(y,\mathcal{K}) \ge \varepsilon$. Moreover, by l.s.c. of ρ (Corollary D.2 with $T = 1$), one has

$$\rho(x,y) \le \liminf_{k\to\infty} \rho(\xi_{n_k-1}^k,\xi_{n_k}^k) \le \mathcal{A}_{T_k}(\xi^k) = 0, \tag{E.150}$$

so that $\rho(x,y) = 0$.

Now, note that $(\xi_0^k,\dots,\xi_{n_k-1}^k)$ and $(\xi_{n_k}^k,\dots,\xi_{T_k-1}^k)$ are, respectively, paths from y_i to $\xi_{n_k-1}^k$ and from $\xi_{n_k}^k$ to x_j with action cost at most $\mathcal{A}_{T_k}(\xi^k)$, thus $B(y_i,\xi_{n_k-1}^k) \to 0$ and $B(\xi_{n_k}^k,x_j) \to 0$ as $k \to \infty$. But since $\xi_{n_k-1}^k$ converges to x and $\xi_{n_k}^k$ to y, it follows from Lemma E.24 that $B(y_i,x) = B(y,x_j) = 0$. Then, we also have

$$0 \le B(x_j,y) \le B(x_j,y_i) + B(y_i,x) + B(x,y) \le \rho(x,y) = 0, \tag{E.151}$$

and thus $y \sim x_j$, which means that $y \in \mathcal{K}$. This is a contradiction.

Therefore, there must exist some sequence $\xi^{i,j} \in (\mathcal{V})^{T_{i,j}}$ with $\xi^{i,j}_0 = y_i$, $\xi^{i,j}_{T_{i,j}-1} = x_j$ such that $\mathcal{A}_{T_{i,j}}(\xi^{i,j}) < \epsilon$ and $\max_{0\le n<T_{i,j}} d(\xi^{i,j}_n, \mathcal{K}) < \epsilon$.

Finally, for the general case with $x, y \in \mathcal{K}$, note that (E.149) implies that there is $i, j \in \mathcal{I}$ such that $x \in \mathbb{B}(\phi^{-1}(y_i), \delta(y_i))$ and $y \in \mathbb{B}(\phi(x_j), \delta(x_j))$. Consider $\xi \in (\mathcal{V})^{T_{i,j}+2}$ an extension of $\xi^{i,j}$ defined by

$$\xi_n = \begin{cases} x & \text{if } n = 0, \\ \xi^{i,j}_{n-1} & \text{if } 1 \le n \le T_{i,j}, \\ y & \text{if } n = T_{i,j} + 1, \end{cases} \tag{E.152}$$

that still satisfies $\max_{0\le n<T_{i,j}+2} d(\xi_n, \mathcal{K}) < \varepsilon$ since we only added points that are already in $\mathcal{K}$. Then, one has

$$\mathcal{A}_{T_{i,j}+2}(\xi) = \rho(x, y_i) + \mathcal{A}_{T_{i,j}}(\xi^{i,j}) + \rho(x_j, y) < 3\varepsilon, \tag{E.153}$$

and the length of the path ξ is upper bounded by $T_0 \coloneqq \max_{i,j\in\mathcal{I}} T_{i,j}+2$, which concludes the proof. ∎

E.5. From equivalence classes to ICTs. Let us remind the definition of an internally chain transitive set:

Definition E.5 (Internal chain transitivity, Benaïm [13])**.** A compact set $\mathcal{K} \subseteq \mathcal{V}$ invariant under Φ is said to be internally chain transitive (ICT) if, for every $x, y \in \mathcal{K}$ and all $\varepsilon > 0, T > 0$, there exists a finite sequence $x_0, \dots, x_m \in \mathcal{K}$ such that $x_0 = x$, $x_m = y$ and $t_1, \dots, t_m \ge T$ such that

$$\max_{1\le i\le m} \|\Phi_{t_i}(x_{i-1}) - x_i\| < \varepsilon. \tag{E.154}$$

There are two key distinctions between the equivalence classes of $\leftrightsquigarrow$ and internally chain transitive sets. First, chains connecting points within an internally chain transitive set must remain entirely within this set, whereas chains associated with our equivalence classes may leave the set. Second, the cost of chains in internally chain transitive sets is strictly bounded in the ℓ^∞ norm, which is generally a much weaker requirement than the ℓ^2 norm used to define the relation $\leftrightsquigarrow$.

The objective of this section is to prove that all equivalence classes of $\leftrightsquigarrow$ are, in fact, internally chain transitive.

We begin with a technical lemma that specializes Lemma E.17 to paths of unit length.

Lemma E.26. *For every compact $\mathcal{K} \subseteq \mathcal{V}$ and all $x, y \in \mathcal{K}$, define*

$$\rho_{\mathcal{K}}(x, y) \coloneqq \inf\left\{\mathcal{S}_{[0,1]}(\gamma) : \gamma \in \mathcal{C}([0, 1], \mathcal{K}), \gamma_0 = x, \gamma_1 = y\right\}. \tag{E.155}$$

Then, for any $\mathcal{K}_1, \mathcal{K}_2 \subseteq \mathcal{V}$ compact, there exists a compact $\mathcal{K} \subseteq \mathcal{V}$ containing both $\mathcal{K}_1$ and $\mathcal{K}_2$ such that $\rho(x, y) = \rho_{\mathcal{K}}(x, y)$ for all $(x, y) \in \mathcal{K}_1 \times \mathcal{K}_2$ satisfying $\rho(x, y) \le 1$.

Proof. Follows from the same arguments used in the proof of Lemma E.17. ∎

Armed with this lemma, we now demonstrate that any two points in an equivalence class $\mathcal{K}$ can be connected by chains that remain within $\mathcal{K}$ and that possess arbitrarily small ℓ^∞ cost.

Lemma E.27. *Let $\mathcal{K}$ be an equivalence class of $\leftrightsquigarrow$. Then, for every $x, y \in \mathcal{K}$ and all $\varepsilon > 0$, $T_0 > 0$, there exists a finite sequence of points $x_0, \dots, x_m \in \mathcal{K}$ such that $x_0 = x$, $x_m = y$ and $t_1, \dots, t_m \ge T_0$ such that $\max_{1\le i\le m} \|\Phi_{t_i}(x_{i-1}) - x_i\| < \varepsilon$.*

Proof. Fix $\varepsilon > 0$, $T_0 > 0$ and let $x, y \in \mathcal{K}$. For $\mathcal{K}_0 = \{z \in \mathcal{V} : d(z, \mathcal{K}) \le 1\}$, we let $\mathcal{K}' \subseteq \mathcal{V}$ be the compact set provided by Lemma E.26 with $\mathcal{K}_1 = \mathcal{K}_2 = \mathcal{K}_0$. With a fair bit of hindsight, we then define the margin

$$\Delta \coloneqq \Delta(T_0, \mathcal{K}) = 1 + L(2T) + \sqrt{2C(2T_0)\sigma^2_\infty(\max_{z\in\mathcal{K}'} \|z\|^2)} < \infty, \tag{E.156}$$

where $L(2T_0)$ is the Lipschitz constant of $\Phi_{2T_0}(\cdot)$ and $C(2T_0)$ is the constant of Lemma E.4. Moreover, without loss of generality, we suppose that $T_0 \ge 1$ is an integer and that $\varepsilon < \min(\Delta, 1)$.

Now, note that by cycling along the paths given by Lemma E.25 enough time so that the total length is sufficiently large (cf. the construction used in the proof of Lemma E.18), there exists a discrete path $\xi \in (\mathcal{V})^T$ such that $T \geq T_0 + 1$, $\xi_0 = x$, $\xi_{T-1} = y$, $\mathcal{S}_T(\xi) < (\varepsilon/\Delta)^2$ and $\max_{0\leq k<T} d(\xi_k, \mathcal{K}) < (\varepsilon/\Delta)^2$. In particular $\xi \subseteq \mathcal{K}_0$ since $\varepsilon < \Delta$, and we also have

$$\rho(\xi_k, \xi_{k+1}) \leq \sum_\ell \rho(\xi_\ell, \xi_{\ell+1}) \leq \mathcal{A}_T(\xi) < (\varepsilon/\Delta)^2 \leq 1. \tag{E.157}$$

It therefore follows from Lemma E.26 that $\rho_{\mathcal{K}'}(\xi_k, \xi_{k+1}) = \rho(\xi_k, \xi_{k+1})$ for all $0 \leq k \leq T-2$. This means that, for all $0 \leq k \leq T-2$, there exist continuous paths $\gamma^k \in \mathcal{C}([0,1], \mathcal{K}')$ such that $\gamma^k_0 = \xi_k$, $\gamma^k_1 = \xi_{k+1}$ and $\mathcal{S}_{[0,1]}(\gamma^k) = \rho(\xi_k, \xi_{k+1})$. Concatenating all these paths therefore yields $\gamma \in \mathcal{C}([0, T-1], \mathcal{K}')$ such that $\gamma_k = \xi_k$ for every $0 \leq k \leq T-1$ and with total cost

$$\mathcal{S}_{[0,T-1]}(\gamma) = \sum_{k=0}^{T-2} \mathcal{S}_{[0,1]}(\gamma^k) = \sum_{k=0}^{T-2} \rho(\xi_k, \xi_{k+1}) \leq \mathcal{A}_T(\xi) < (\varepsilon/\Delta)^2. \tag{E.158}$$

Let $m \geq 1$ be the unique integer such that $mT_0 \leq T-1 \leq (m+1)T_0$. We then define a sequence $\tilde{x}_0, \dots, \tilde{x}_m$ by $\tilde{x}_0 = x$, $\tilde{x}_m = y$ and $\tilde{x}_i = \xi_{iT_0}$ for every $1 \leq i \leq m-1$, and times $t_1, \dots, t_m$ such that $t_m = T - 1 - (m-1)T_0$ and $t_i = T_0$ for $1 \leq i < m$. In particular, we note that $T_0 \leq t_i \leq 2T_0$ by construction. If we further consider $x_0, \dots, x_m \in \mathcal{K}$ such that $x_i \in \mathrm{proj}_{\mathcal{K}}(\tilde{x}_i)$, then we get

$$\begin{aligned}
\|\Phi_{t_i}(x_{i-1}) - x_i\| &\leq \|\Phi_{t_i}(x_{i-1}) - \Phi_{t_i}(\tilde{x}_{i-1})\| + \|x_i + \tilde{x}_i\| + \|\Phi_{t_i}(\tilde{x}_{i-1}) - \tilde{x}_i\| && \text{(E.159)}\\
&\leq L(2T_0)\|x_{i-1} - \tilde{x}_{i-1}\| + \|x_i - \tilde{x}_i\| + \|\Phi_{t_i}(\tilde{x}_{i-1}) - \tilde{x}_i\| && \text{(E.160)}\\
&= L(2T_0) d(\tilde{x}_{i-1}, \mathcal{K}) + d(\tilde{x}_i, \mathcal{K}) + \|\Phi_{t_i}(\tilde{x}_{i-1}) - \tilde{x}_i\| && \text{(E.161)}\\
&\leq (L(2T_0)+1) \max_{0\leq k\leq T-1} d(\xi_k, \mathcal{K}) + \|\Phi_{t_i}(\tilde{x}_{i-1}) - \tilde{x}_i\| && \text{(E.162)}\\
&< (L(2T_0)+1)(\varepsilon/\Delta)^2 + \|\Phi_{t_i}(\tilde{x}_{i-1}) - \tilde{x}_i\|, && \text{(E.163)}
\end{aligned}$$

where we have used that $\Phi_t(\cdot)$ is $L(2T_0)$-Lipschitz for all $0 \leq t \leq 2T_0$ and that the discrete chain ξ remains at distance at most $(\varepsilon/\Delta)^2$ of $\mathcal{K}$. On the other hand, if we define $\tau_i = \sum_{j\leq i} t_i$, then applying Lemma E.4 and following the same arguments as in the proof of Lemma E.19 also yields

$$\begin{aligned}
\|\Phi_{t_i}(\tilde{x}_{i-1}) - \tilde{x}_i\|^2 = \|\Phi_{t_i}(\gamma_{\tau_{i-1}}) - \gamma_{\tau_{i-1}+t_i}\|^2 &\leq C(2T_0) \int_{\tau_{i-1}}^{\tau_i} \|\dot{\gamma}_s - F(\gamma_s)\|^2 ds && \text{(E.164)}\\
&\leq 2C(2T_0)\sigma_\infty^2 (\max_{z\in\mathcal{K}'} \|z\|^2) \mathcal{S}_{[0,T-1]}(\gamma) && \text{(E.165)}\\
&\leq 2C(2T_0)\sigma_\infty^2 (\max_{z\in\mathcal{K}'} \|z\|^2) \frac{\varepsilon^2}{\Delta^2}. && \text{(E.166)}
\end{aligned}$$

Combined with Eq. (E.163), this finally leads to

$$\begin{aligned}
\|\Phi_{t_i}(x_{i-1}) - x_i\| &< (L(2T_0)+1)(\varepsilon/\Delta)^2 + \sqrt{2C(2T_0)\sigma_\infty^2 (\max_{z\in\mathcal{K}'} \|z\|^2)}(\varepsilon/\Delta) && \text{(E.167)}\\
&\leq \frac{1}{\Delta}\left(1 + L(2T) + \sqrt{2C(2T_0)\sigma_\infty^2 (\max_{z\in\mathcal{K}'} \|z\|^2)}\right)\varepsilon && \text{(E.168)}\\
&= \varepsilon, && \text{(E.169)}
\end{aligned}$$

hence concluding the proof since $x_0, \dots, x_m$ is a chain from x to y contained in $\mathcal{K}$ by construction and $t_0, \dots, t_m \geq T_0$ for T_0 arbitrarily large. ∎

It therefore follows that the equivalence classes are internally chain transitive.

Theorem E.2 (Equivalence classes are ICTs)**.** *All equivalence classes of $\leftrightsquigarrow$ are internally chain transitive.*

Proof. Invoking Lemmas E.7, E.10 and E.27, an equivalence class $\mathcal{K}$ of $\leftrightsquigarrow$ is a compact invariant set whose members are linked by chains of arbitrarily large length that remain in $\mathcal{K}$, and thus it is internally chain-transitive by definition. ∎

E.6. Lifting through diffeomorphism. In this section, we assume that the vector field F admits the representation

$$F(z) = (\operatorname{Jac} Q)^{-1}(z)\tilde{F}(Q(z)) \tag{E.170}$$

where $Q\colon \mathcal{V} \to \operatorname{int}\mathcal{X}$ is a diffeomorphism between $\mathcal{V}$ and the interior of a compact set $\mathcal{X} \subseteq \mathcal{V}$, and $\tilde{F}\colon \operatorname{int}\mathcal{X} \to \mathcal{V}$ is a vector field on $\operatorname{int}\mathcal{X}$.

We will establish that the equivalence classes (resp. the internally chain transitive components) of F in $\mathcal{V}$ are in one-to-one correspondence with the equivalence classes (resp. the internally chain transitive components) of $\tilde{F}$ in $\operatorname{int}\mathcal{X}$. These results allow us to translate our general abstract framework to the specific (REW) algorithm in Appendix C.

We start by proving that the vector field $\tilde{F}$ induces a unique flow Ψ, which is conjugate to the flow Φ generated by F.

Lemma E.28. *The flow Ψ induced by $\tilde{F}$ is unique and satisfies*

$$\Psi_t = Q \circ \Phi_t \circ Q^{-1} \tag{E.171}$$

where Φ is the flow induced by F.

Proof. Letting $t \in \mathbb{R}$ and $x \in \operatorname{int}\mathcal{X}$, the chain rule applied to $\Psi_t := Q \circ \Phi_t \circ Q^{-1}$ yields

$$\begin{aligned}\dot{\Psi}_t(x) &= (\operatorname{Jac} Q \circ \Phi_t \circ Q^{-1})(x)\dot{\Phi}_t(Q^{-1}(x)) \\ &= (\operatorname{Jac} Q \circ \Phi_t \circ Q^{-1})(x)F((\Phi_t \circ Q^{-1})(x)) \\ &= (\operatorname{Jac} Q \circ Q^{-1} \circ \Psi_t)(x)F((Q^{-1} \circ \Psi_t)(x)) \\ &= \tilde{F}(x),\end{aligned} \tag{E.172}$$

and thus Ψ is indeed a flow of $\tilde{F}$. On other hand, if Ψ satisfies $\dot{\Psi}(x) = \tilde{F}(\Psi_t(x))$, then

$$\begin{aligned}\frac{d}{dt}[(Q^{-1} \circ \Psi_t)(x)] &= (\operatorname{Jac} Q^{-1} \circ \Psi_t)(x)\dot{\Psi}_t(x) \\ &= [(\operatorname{Jac} Q \circ Q^{-1} \circ \Psi_t)(x)]^{-1}\tilde{F}(\Psi_t(x)) \\ &= [(\operatorname{Jac} Q \circ Q^{-1} \circ \Psi_t)(x)]^{-1}(\operatorname{Jac} Q \circ Q^{-1} \circ \Psi_t)(x)F((Q^{-1} \circ \Psi_t)(x)) \\ &= F((Q^{-1} \circ \Psi_t)(x))\end{aligned} \tag{E.173}$$

and therefore $Q^{-1} \circ \Psi_t = \Phi_t \circ Q^{-1}$ by uniqueness of the flow induced by F (cf. Lemma E.1). It follows that the flow of $\tilde{F}$ is also unique. ∎

To distinguish the two spaces, we denote the ℓ^2-chain relations and equivalence classes associated with the flow Φ on $\mathcal{V}$ by $\rightsquigarrow_{\mathcal{V}}$, $\leftrightsquigarrow_{\mathcal{V}}$ and $[z]^{\mathcal{V}}_{\leftrightsquigarrow}$, respectively. Similarly, for the flow Ψ on $\operatorname{int}\mathcal{X}$, we write $\rightsquigarrow_{\mathcal{X}}$, $\leftrightsquigarrow_{\mathcal{X}}$ and $[x]^{\mathcal{X}}_{\leftrightsquigarrow}$. We first prove that the diffeomorphism Q maps the equivalence classes under Φ directly to those under Ψ.

Lemma E.29. $Q([z]^{\mathcal{V}}_{\leftrightsquigarrow}) \subseteq [Q(z)]^{\mathcal{X}}_{\leftrightsquigarrow}$.

Proof. Assume that $z \rightsquigarrow_{\mathcal{V}} z'$ for some $z, z' \in \mathcal{V}$. Then Lemma E.9 provides a compact set $\mathcal{K} \subseteq \mathcal{V}$ such that, for all $\varepsilon > 0$ small enough and all $T > 0$ large enough, there exists a finite sequence $z_0, \dots, z_m \in \mathcal{K}$ such that $z_0 = z$, $z_m = z'$ and $t_1, \dots, t_m \geq T$ such that $\sum_i \|\Phi_{t_i}(z_{i-1}) - z_i\|^2 < \varepsilon$. We now consider the lifted sequence $x_0, \dots, x_m$ defined by $x_i = Q(z_i)$. This sequence satisfies

$$\sum_{i=1}^{m} \|\Psi_{t_i}(x_{i-1}) - x_i\|^2 = \sum_{i=1}^{m} \|\Psi_{t_i}(Q(z_{i-1})) - Q(z_i)\|^2 = \sum_{i=1}^{m} \|Q(\Phi_{t_i}(z_{i-1})) - Q(z_i)\|^2, \tag{E.174}$$

where we have used the relation between Φ and Ψ derived in Lemma E.28. Moreover, all elements z_i's are included in the compact $\mathcal{K}$ by construction, and we can suppose without loss of generality that $\mathcal{K}$ is positive invariant under Φ (cf. Lemma E.5). Since Q is continuously differentiable, it follows that there exists a constant $L_Q(\mathcal{K}) > 0$ such that

$$\|Q(\Phi_{t_i}(z_{i-1})) - Q(z_i)\| \leq L_Q(\mathcal{K})\|\Phi_{t_i}(z_{i-1}) - z_i\| \tag{E.175}$$

for all $1 \le i \le m$. As a consequence,

$$\sum_{i=1}^{m} \|\Psi_{t_i}(x_{i-1}) - x_i\|^2 \le L_Q(\mathcal{K})^2 \sum_{i=1}^{m} \|\Phi_{t_i}(z_{i-1}) - z_i\|^2 < L_Q(\mathcal{K})^2 \varepsilon, \tag{E.176}$$

and it suffices to take ε sufficiently small such that $\varepsilon < 1/L_m irror(\mathcal{K})^2$ in order to obtain $Q(z) \rightsquigarrow_{\mathcal{X}} Q(z')$. Since this hold true for every $z' \in [z]^{\mathcal{V}}_{\leftrightsquigarrow}$, it follows that $Q([z]^{\mathcal{V}}_{\leftrightsquigarrow}) \subseteq [Q(z)]^{\mathcal{X}}_{\leftrightsquigarrow}$ as needed. ∎

Proving the reverse inclusion of Lemma E.29 is not immediate our results. Indeed, if $\tilde{F}$ were dissipative in the sense of Assumption 4, the argument used for Lemma E.29 could be followed ad verbatim to obtain the desired result. However, this precise dissipation condition cannot hold here, as the diffeomorphism maps the space to the interior of a compact set. Nonetheless, the following lemma demonstrates that the flow Ψ still satisfies a dissipative condition, albeit with the standard norm $\frac{1}{2}\|x\|^2$ replaced by the Lyapunov function $V(x) = \frac{1}{2}\|Q^{-1}(x)\|^2$.

Lemma E.30. *Let $V\colon \operatorname{int}\mathcal{X} \to \mathbb{R}$ be the function defined as $V(x) = \frac{1}{2}\|Q^{-1}(x)\|^2$. Then:*

1. *V is continuous and positive on $\operatorname{int}\mathcal{X}$;*
2. *$\mathcal{K}^{\mathcal{X}}_{\delta} := \{x \in \operatorname{int}\mathcal{X} : V(x) \le \delta\}$ is compact in $\operatorname{int}\mathcal{X}$ for every $\delta > 0$;*
3. *$\frac{d}{dt}V(\Psi_t(x)) \le -\kappa V(\Psi_t(x)) + C$ for every $x \in \mathcal{X}$.*

Proof. 1. Follows directly from the continuity of Q^{-1}.

2. We have $\mathcal{K}^{\mathcal{X}}_{\delta} = Q(\bar{\mathbb{B}}(0, 2\delta))$, but $\bar{\mathbb{B}}(0, 2\delta)$ is compact in $\mathcal{V}$ and $Q\colon \mathcal{V} \to \operatorname{int}\mathcal{X}$ is continuous, thus $\mathcal{K}^{\mathcal{X}}_{\delta}$ is compact.

3. For every $x \in \operatorname{int}\mathcal{X}$,

$$\begin{aligned}
\frac{d}{dt}V(\Psi_t(x)) = \frac{1}{2}\frac{d}{dt}\|Q^{-1}(\Psi_t(x))\|^2 &= \frac{1}{2}\frac{d}{dt}\|\Phi_t(Q^{-1}(x))\|^2 && \text{(E.177)}\\
&= \langle \dot{\Phi}_t(Q^{-1}(x)), \Phi_t(Q^{-1}(x))\rangle && \text{(E.178)}\\
&= \langle F(\Phi_t(Q^{-1}(x))), \Phi_t(Q^{-1}(x))\rangle && \text{(E.179)}\\
&\le -\kappa\|\Phi_t(Q^{-1}(x))\|^2 + C && \text{(E.180)}\\
&= -\kappa V(\Psi_t(x)) + C, && \text{(E.181)}
\end{aligned}$$

where the inequality follows from Assumption 4(*a*). ∎

Crucially, these properties suffice to extend Lemmas E.5 and E.9 to the flow Ψ on $\operatorname{int}\mathcal{X}$. As the arguments are identical, save for substituting the standard norm $\frac{1}{2}\|x\|^2$ with the Lyapunov function $V(x) = \frac{1}{2}\|Q^{-1}(x)\|^2$ and replacing the closed balls $\bar{\mathbb{B}}(0, M)$ with the sets $\mathcal{K}^{\mathcal{X}}_{M}$, we state the results without proof.

Lemma E.31. *If $\mathcal{K}_0 \subseteq \operatorname{int}\mathcal{X}$ is a compact set, then it is included in another compact $\mathcal{K} \subseteq \operatorname{int}\mathcal{X}$ such that $\Psi_t(\mathcal{K}) \subseteq \mathcal{K}$ for all $t \ge 0$.*

Lemma E.32. *For every $x_0 \in \operatorname{int}\mathcal{X}$ and $\varepsilon_0 > 0$, there exist $t_0 := t_0(x_0, \varepsilon_0) > 0$ and $\mathcal{K} := \mathcal{K}(x_0, \varepsilon_0) \subseteq \operatorname{int}\mathcal{X}$ compact such that any finite chain $x_1, \dots, x_m \in \operatorname{int}\mathcal{X}$, $t_1, \dots, t_m \ge t_0$ satisfying $\max_{1\le i\le m}\|\Psi_{t_i}(x_{i-1}) - x_i\|^2 < \varepsilon_0$ has $x_0, \dots, x_m \in \mathcal{K}$.*

Equipped with these two lemmas, we may now establish the reverse inclusion of Lemma E.29 by following the same line of reasoning.

Lemma E.33. $Q^{-1}([x]^{\mathcal{X}}_{\leftrightsquigarrow}) \subseteq [Q^{-1}(x)]^{\mathcal{V}}_{\leftrightsquigarrow}$.

Proof. Thanks to Lemmas E.31 and E.32, we can always suppose that the elements of an ℓ^2-chain for Ψ remain in a positive invariant compact subset, and thus we can argue as in the proof of Lemma E.29 to obtain the result. ∎

Corollary E.5. *$\mathcal{K}$ is an equivalence class of $\leftrightsquigarrow_{\mathcal{V}}$ if and only if $Q(\mathcal{K})$ is an equivalence class of $\leftrightsquigarrow_{\mathcal{X}}$.*

Proof. Follows from Lemmas E.29 and E.33. ∎

As in Definition E.5, we say that a set $\mathcal{K} \subseteq \operatorname{int} \mathcal{X}$ is internally chain transitive for the flow Ψ if it is compact and invariant under Ψ, and if any two points of $\mathcal{K}$ can be connected by arbitrarily long chains that remain within $\mathcal{K}$ and with arbitrarily small ℓ^∞ cost. The following result establishes the correspondence between internally chain transitive sets for Φ and those for Ψ.

Proposition E.3. *A set $\mathcal{K} \subseteq \mathcal{V}$ is internally chain transitive for Φ if and only if $Q(\mathcal{K})$ is internally chain transitive for Ψ.*

Proof. The compactness is immediate since Q and Q^{-1} are both continuous. For invariance under the flows, this follows directly from their conjugation relation stated in Lemma E.28. Finally, since all elements of a chain containing two points of an internally chain transitive set remain in this internally chain transitive set by definition, the same argument used in the proof of Lemma E.29 can again be invoked to obtain both inclusions. ∎

F Transitions

F.1 Lyapunov condition 55
F.2 Transition times 59
F.3 Transition probabilities 63

In this appendix, we study the transitions of the learning process between components. For this, in Appendix F.1, we ensure that the process spends little time at infinity by establishing a Lyapunov condition, which is then used in Appendix F.2 to obtain exponential bounds on the transition times between components. Finally, in Appendix F.3, we obtain exponential bounds on the transition probabilities between components, which will be used in the subsequent appendices to derive the long-run distribution of the process.

F.1. Lyapunov condition. The main result of this section is Lemma F.3, which establishes a Lyapunov condition for the process, using the dissipativity of the vector field and the sub-Gaussianity of the noise. The proof of this result relies on the algebraic identity Lemma F.2. The other results of this section follow closely Azizian et al. [9], but we include complete proofs to make this work self-contained and to verify that the arguments remain valid in our more general setting.

Definition F.1 (Stopping times for the accelerated process)**.** For any set $S \subset \mathcal{V}$, we define the hitting and exit times of S:

$$\sigma_S \coloneqq \inf\{n \geq 1 : x_n^\eta \in S\} \tag{F.1a}$$
$$\tau_S \coloneqq \inf\{n \geq 0 : x_n^\eta \notin S\}\,. \tag{F.1b}$$

We will need the following concentration lemma and its corollary.

Lemma F.1 (Part of the proof of [95, Th. 1.19])**.** *Let X be a random variable in $\mathbb{R}^d$ such that, for all $w \in \mathbb{R}^d$,*

$$\log \mathbb{E}[\exp(\langle w, X\rangle)] \leq \frac{\|w\|^2}{2}\,. \tag{F.2}$$

Then, for all $t > 0$, we have

$$\mathbb{P}\Big(\|X\|^2 \geq t\Big) \leq 6^d \exp\Big(-\frac{t}{8}\Big)\,. \tag{F.3}$$

Corollary F.1 ([9, Cor. D.2])**.** *In the context of Lemma F.1, it holds that*

$$\mathbb{E}\big[\|X\|^2\big] \leq 16\, d \log 6\,. \tag{F.4}$$

Let us begin by a simpler algebraic lemma.

Lemma F.2. *For any $x \in \mathcal{V}$, $\delta x \in \mathcal{V}$ and $\lambda > 0$, it holds that*

$$\langle x, \delta x\rangle + \lambda\|\delta x\|^2 \leq (\kappa^{-1} + 2\lambda)\|\delta x - F(x)\|^2 - \Big(\frac{\kappa}{4} - 4\lambda M^2\Big)\|x\|^2 + C + 4\lambda M^2\,. \tag{F.5}$$

Proof. On the one hand, by the Cauchy-Schwarz inequality we have,

$$\begin{aligned}\langle x, \delta x\rangle &= \langle x, \delta x - F(x)\rangle + \langle x, F(x)\rangle \\ &\leq \|x\|\|\delta x - F(x)\| + \langle x, F(x)\rangle \\ &\leq \frac{1}{\kappa}\|\delta x - F(x)\|^2 + \frac{\kappa}{4}\|x\|^2 - \frac{\kappa}{2}\|x\|^2 + C\,,\end{aligned} \tag{F.6}$$

where we used Young's inequality and the dissipativity of F in the last line (Assumption 4(a)). On the other hand, by the triangular inequality,

$$\begin{aligned}\|\delta x\|^2 &\leq 2\|\delta x - F(x)\|^2 + 2\|F(x)\|^2 \\ &\leq 2\|\delta x - F(x)\|^2 + 4M^2(1 + \|x\|^2)\,,\end{aligned} \tag{F.7}$$

where we used the growth condition on F (Eq. (D.4)). Combining Eq. (F.6) and Eq. (F.7) yields the desired result. ∎

Lemma F.3 (Lyapunov condition)**.** *Define g_∞ as in Definition D.1. Then, there exists $\mathcal{K} \subset \mathcal{V}$ compact, $\eta_0 > 0$, $c > 0$ such that, for any $\eta \le \eta_0$, $n \ge 0$, if $x_n \notin \mathcal{K}$, then, almost surely,*

$$\begin{aligned} g_\infty(x_{n+1}) - g_\infty(x_n) &\le 4\eta\kappa^{-1}\left(\frac{\|U(x_n,\omega_n)\|^2}{\sigma_\infty^2(\|x_n\|^2)} - \frac{\kappa^2}{32}\frac{\|x_n\|^2}{\sigma_\infty^2(\|x_n\|^2)}\right) \\ &\le 4\eta\kappa^{-1}\left(\frac{\|U(x_n,\omega_n)\|^2}{\sigma_\infty^2(\|x_n\|^2)} - (16d\log 6 + c)\right). \end{aligned} \tag{F.8}$$

Proof. Suppose that $\eta \in (0,1]$. Note that the last condition in Assumption 4(*b*) implies that, whether $s < 2$ or $s = 2$, we always have that

$$\liminf_{x\to\infty} \frac{\|x\|^2}{\sigma_\infty^2(\|x\|^2)} > \frac{512d\log 6}{\kappa^2}\,. \tag{F.9}$$

This guarantees that there exists $R \ge \max\left(\frac{1}{2}, \frac{1}{2}4^{\frac{s+3}{2}}M\right)$ and $c > 0$ such that, for any x such that $\|x\| \ge R$,

$$\frac{\|x\|^2}{\sigma_\infty^2(\|x\|^2)} \ge \frac{512d\log 6}{\kappa^2} + \frac{32c}{\kappa^2}\,. \tag{F.10}$$

Then, define $\mathcal{K} \coloneqq \bar{\mathbb{B}}(0, 2R+1)$.

By definition, g_∞ is twice continuously differentiable and, its Hessian satisfies, for any $x \in \mathcal{V}$,

$$\operatorname{Hess} g_\infty(x) = \operatorname{Jac}\left(\frac{4x}{\sigma_\infty^2(\|x\|^2)}\right) \tag{F.11}$$

$$= \frac{4}{\sigma_\infty^2(\|x\|^2)} I - \frac{16xx^\top((\sigma_\infty^2)'(\|x\|^2))}{(\sigma_\infty^2(\|x\|^2))^2} \tag{F.12}$$

$$\preccurlyeq \frac{4}{\sigma_\infty^2(\|x\|^2)} I\,. \tag{F.13}$$

For the sake of clarity, for any $n \ge 0$, denote by δx_n the quantity

$$\delta x_n = \frac{x_{n+1} - x_n}{\eta}\,. \tag{F.14}$$

For any $n \ge 0$, we now have, by Taylor's inequality, that

$$\begin{aligned} g_\infty(x_{n+1}) - g_\infty(x_n) &\le \eta\langle \nabla g_\infty(x_n), \delta x_n\rangle \\ &\quad + \frac{\eta^2}{2}\|\delta x_n\|^2 \sup_{t\in[0,1]} \frac{4}{\sigma_\infty^2(\|x_n + t(x_{n+1}-x_n)\|^2)}\,. \end{aligned} \tag{F.15}$$

We first focus on bounding the last term. First note that, by the blanket assumptions Eq. (D.4), $\|x_{n+1} - x_n\| \le 2\eta M(1 + \|x_n\|)$ so that, For any $t \in [0,1]$, $\eta \le (4M)^{-1}$,

$$\begin{aligned} \|x_n + t(x_{n+1} - x_n)\| &\ge \|x_n\| - \|x_{n+1} - x_n\| \\ &\ge \|x_n\| - 2\eta M(1 + \|x_n\|) \\ &\ge \frac{1}{2}(\|x_n\| - 1)\,. \end{aligned} \tag{F.16}$$

If x_n is outside of $\mathcal{K}$, we have that $\|x_n\| \ge 2R + 1$ and thus $\|x_n + t(x_{n+1} - x_n)\| \ge R$. Moreover, since $R \ge \frac{1}{2}$, x_n being outside of $\mathcal{K}$ also implies that $\frac{1}{2}\|x_n\| \ge 1$ and so $\|x_n + t(x_{n+1} - x_n)\| \ge \frac{1}{4}\|x_n\|$. By the definition of σ_∞^2 (Definition D.1), we have that, for any $t \in [0,1]$,

$$\begin{aligned} \sigma_\infty^2(\|x_n + t(x_{n+1} - x_n)\|^2) &= \overline{\sigma}^2(1 + \|x_n + t(x_{n+1} - x_n)\|^s) \\ &\ge \overline{\sigma}^2\left(1 + \left(\frac{1}{4}\right)^s \|x_n\|^s\right) \end{aligned}$$

$$\geq \overline{\sigma}^2 \left(\frac{1}{4}\right)^s (1 + \|x_n\|^s)$$
$$= \frac{1}{4^s} \sigma_\infty^2(\|x_n\|^2)\,. \tag{F.17}$$

Thus, if $x_n \notin \mathcal{K}$, Eq. (F.13) yields that

$$g_\infty(x_{n+1}) - g_\infty(x_n) \leq \eta \langle \nabla g_\infty(x_n), \delta x_n \rangle + 2\eta^2 \frac{4^s}{\sigma_\infty^2(\|x_n\|^2)} \|\delta x_n\|^2 \tag{F.18}$$
$$= \frac{2\eta}{\sigma_\infty^2(\|x_n\|^2)} \left(\langle x_n, \delta x_n \rangle + \eta 4^s \|\delta x_n\|^2 \right). \tag{F.19}$$

Now, invoking Lemma F.2 with $x \leftarrow x_n$, $\delta x \leftarrow \delta x_n$ and $\lambda \leftarrow 4^s\eta$ yields that,

$$\langle x_n, \delta x_n \rangle + \eta 4^s \|\delta x_n\|^2 \leq (\kappa^{-1} + 2 \cdot 4^s \eta) \|\delta x_n - F(x_n)\|^2 - \left(\frac{\kappa}{4} - 4^{s+1}\eta M^2\right) \|x_n\|^2 + C + 4^{s+1}\eta M^2\,. \tag{F.20}$$

Taking $\eta > 0$ small enough such that

$$2 \cdot 4^s \eta \leq \kappa^{-1} \quad \text{and} \quad 4^{s+1}\eta M^2 \leq \frac{\kappa}{8}\,, \tag{F.21}$$

we obtain,

$$\langle x_n, \delta x_n \rangle + \eta 4^s \|\delta x_n\|^2 \leq 2\kappa^{-1} \|\delta x_n - F(x_n)\|^2 - \frac{\kappa}{16} \|x_n\|^2 + \left(C + 4^{s+1}\eta M^2 - \frac{\kappa}{16}(2R+1)^2 \right) \tag{F.22}$$
$$\leq 2\kappa^{-1} \|\delta x_n - F(x_n)\|^2 - \frac{\kappa}{16} \|x_n\|^2 \tag{F.23}$$
$$= 2\kappa^{-1} \left(\|\delta x_n - F(x_n)\|^2 - \frac{\kappa^2}{32} \|x_n\|^2 \right), \tag{F.24}$$

using that $R \geq \frac{1}{2} 4^{\frac{s+3}{2}} M$ and $\eta \leq 1$.

Plugging this into Eq. (F.19) yields that, if $x_n \notin \mathcal{K}$,

$$g_\infty(x_{n+1}) - g_\infty(x_n) \leq 4\eta\kappa^{-1} \left(\frac{\|\delta x_n - F(x_n)\|^2}{\sigma_\infty^2(\|x_n\|^2)} - \frac{\kappa^2}{32} \frac{\|x_n\|^2}{\sigma_\infty^2(\|x_n\|^2)} \right) \tag{F.25}$$
$$= 4\eta\kappa^{-1} \left(\frac{\|U(x_n, \omega_n)\|^2}{\sigma_\infty^2(\|x_n\|^2)} - \frac{\kappa^2}{32} \frac{\|x_n\|^2}{\sigma_\infty^2(\|x_n\|^2)} \right), \tag{F.26}$$

where we unrolled δx_n in the last line. Using Eq. (F.10), we thus also obtain that,

$$g_\infty(x_{n+1}) - g_\infty(x_n) \leq 4\eta\kappa^{-1} \left(\frac{\|U(x_n, \omega_n)\|^2}{\sigma_\infty^2(\|x_n\|^2)} - (16d \log 6 + c) \right). \tag{F.27}$$

■

We will reuse, in later sections, the following fact that we thus state as a lemma: the sequence of iterates of (SD) is *(weak) Feller* (see e.g., [51, Def. 4.4.2]).

Lemma F.4. *The Markov chain $(x_n)_{n\geq 0}$ is (weak) Feller.*

Proof. since both F and U are continuous, for any $g\colon \mathcal{V} \to \mathbb{R}$ continuous and bounded, the function.

$$x \in \mathcal{V} \longmapsto \mathbb{E}_x[g(x_1)] = \mathbb{E}_x[g(x + \eta F(x) + \eta U(x, \omega))] \tag{F.28}$$

is still continuous and bounded. Therefore, the Markov chain $(x_n)_{n\geq 0}$ is weak-Feller. ■

Lemma F.5. *There is $\eta_0 > 0$ such that, for any $\eta \leq \eta_0$, there exists an invariant probability measure for $(x_n)_{n\geq 0}$.*

Proof. We invoke a general result on weak-Feller Markov chains that satisfy a Lyapunov condition, e.g., Hernández-Lerma & Lasserre [51, Thm. 7.2.4] or Douc et al. [33, Thm. 12.3.6].

First, Lemma F.4 ensures that $(x_n)_{n\geq 0}$ is weak-Feller.

Moreover, by Lemma F.3, there exists $\mathcal{K} \subset \mathcal{V}$ compact, $\eta_0 > 0$, $c > 0$ such that, for any $\eta \leq \eta_0$, $x_0 = x \notin \mathcal{K}$,

$$g_\infty(x_1) - g_\infty(x) \leq \eta \left(\frac{\|U(x,\omega_0)\|^2}{\sigma_\infty^2(\|x\|^2)} - (16d \log 6 + c) \right). \tag{F.29}$$

Passing to the expectation yields that, for any $x \notin \mathcal{K}$,

$$\mathbb{E}_x[g_\infty(x_1)] - g_\infty(x) \leq \eta \left(\frac{\mathbb{E}_x\left[\|U(x,\omega_0)\|^2\right]}{\sigma_\infty^2(\|x\|^2)} - (16d \log 6 + c) \right). \tag{F.30}$$

Applying Corollary F.1 with $X \leftarrow \frac{U(x,\omega_0)}{\sqrt{\sigma_\infty^2(\|x\|^2)}}$ (the conditions of application are verified from Assumption 4(*b*)) yields that

$$\mathbb{E}_x[g_\infty(x_1)] - g_\infty(x) \leq -\eta c \,. \tag{F.31}$$

Hence, for any $x \in \mathcal{V}$, it holds

$$\mathbb{E}_x[g_\infty(x_1)] - g_\infty(x) \leq -\eta c + \mathbb{1}\{x \in \mathcal{K}\} \left(\sup_{x' \in \mathcal{K}} \mathbb{E}_{x'}[g_\infty(x')] - \inf_{\mathcal{V}} g_\infty + \eta c \right), \tag{F.32}$$

with g_∞ which is not identically equal to its minimum since α_∞ is increasing.

Therefore, we can apply Hernández-Lerma & Lasserre [51, Thm. 7.2.4] with $g_\infty - \inf_{\mathcal{V}} g_\infty$, that guarantees that there exists an invariant probability measure for $(x_n)_{n\geq 0}$. ∎

Lemma F.6. *There exists a compact set $\mathcal{D} \subset \mathcal{V}$, $\eta_0 > 0$, such that for any compact set $\mathcal{D}' \subset \mathcal{V}$ such that $\mathcal{D} \subset \mathcal{D}'$, there exists $a, b > 0$, such that*

$$\forall \eta \leq \eta_0 \,,\, x \in \mathcal{D}' \setminus \mathcal{D} \,,\, n \geq 0 \,,\, \mathbb{P}_x(\sigma_{\mathcal{D}} > n) \leq \exp\left(-\frac{an}{\eta} + \frac{b}{\eta} \right). \tag{F.33}$$

Proof. This result is a consequence of Lemma F.3: there exists $\mathcal{K} \subset \mathcal{V}$ compact, $\eta_0 > 0$, $c > 0$ such that, for any $\eta \leq \eta_0$, $n \geq 0$, if $x_n \notin \mathcal{K}$, then, almost surely,

$$g_\infty(x_{n+1}) - g_\infty(x_n) \leq \eta \left(\frac{\|U(x_n,\omega_n)\|^2}{\sigma_\infty^2(\|x_n\|^2)} - (16d \log 6 + c) \right). \tag{F.34}$$

First, let us choose $\mathcal{D}$. There is some $R > 0$ such that $\mathcal{K} \subset \mathbb{B}(0,R)$. Define $\widetilde{R} := e^{4M+1}(R+1)$ and $\mathcal{D} := \bar{\mathbb{B}}(0,\widetilde{R})$. With $\eta \leq (4M)^{-1}$, if, for some $n \geq 0$, x_n is not in $\mathcal{D}$, by Lemma D.2, for any $k \leq \lceil \eta^{-1} \rceil$, x_{n+k} has norm greater or equal to R and thus x_{n+k} is not in $\mathcal{K}$ either.

Now, fix $T \geq 0$ and consider the event $\{\sigma_{\mathcal{D}} > T\}$ with $x_0 = x \in \mathcal{D}' \setminus \mathcal{D}$. This means that, for any $0 \leq n \leq T$, x_n^η is outside of $\mathcal{D}$ and so, for any $0 \leq n \leq T\lfloor \eta^{-1} \rfloor$, x_n is not in $\mathcal{K}$.

Summing Eq. (F.34) over $n = 0, \dots, T\lfloor \eta^{-1} \rfloor - 1$ yields

$$\begin{aligned} g_\infty(x_T^\eta) - g_\infty(x_0) &= g_\infty(x_{T\lfloor \eta^{-1} \rfloor}) - g_\infty(x_0) \\ &\leq \eta \sum_{n=0}^{T\lfloor \eta^{-1} \rfloor - 1} \left(\frac{\|U(x_n,\omega_n)\|^2}{\sigma_\infty^2(\|x_n\|^2)} - (16d \log 6 + c) \right). \end{aligned} \tag{F.35}$$

Define $\Delta := \sup_{\mathcal{D}' \setminus \mathcal{D}} g_\infty - \inf_{\mathcal{V} \setminus \mathcal{D}} g_\infty$ which is finite since α_∞ is increasing. By definition, we have that $g_\infty(x_T^\eta) - g_\infty(x_0) \geq -\Delta$.

Therefore, on the event $\{\sigma_{\mathcal{D}} > T\}$, we have

$$\sum_{n=0}^{T\lfloor \eta^{-1} \rfloor - 1} \frac{\|U(x_n,\omega_n)\|^2}{\sigma_\infty^2(\|x_n\|^2)} \geq T\lfloor \eta^{-1} \rfloor (16d \log 6 + c) - \frac{\Delta}{\eta} \,. \tag{F.36}$$

Now, on the whole event, consider the random variable $X \in \mathbb{R}^{T\lfloor\eta^{-1}\rfloor d}$ defined by

$$(X_{nd}+1,\dots,X_{(n+1)d}) = \frac{U(x_n,\omega_n)}{\sqrt{\sigma_\infty^2(\|x_n\|^2)}} \quad \text{for } n=0,\dots,T\lfloor\eta^{-1}\rfloor-1\,, \tag{F.37}$$

we have that $\mathbb{P}_x(\sigma_{\mathcal{D}}>T) \le \mathbb{P}_x\Big(\|X\|^2 \ge T\lfloor\eta^{-1}\rfloor(16d\log 6+c)-\frac{\Delta}{\eta}\Big)$.

Since, for any $w\in\mathcal{V}^*$,

$$\log\mathbb{E}\left[\exp\left(\left\langle w, \frac{U(x_n,\omega_n)}{\sqrt{\sigma_\infty^2(\|x_n\|^2)}}\right\rangle\right)\middle| x_0,x_1,\dots,x_n\right] \le \frac{\|w\|^2}{2}\,, \tag{F.38}$$

the random variable X satisfies the assumptions of Lemma F.1 with $d \leftarrow T\lfloor\eta^{-1}\rfloor d$.

First, suppose that $t \coloneqq T\lfloor\eta^{-1}\rfloor(16d\log 6+c)-\frac{\Delta}{\eta}$ is non-negative. Applying Lemma F.1 with this t yields that

$$\begin{aligned}\mathbb{P}_x(\sigma_{\mathcal{D}}>T) &\le \mathbb{P}_x\left(\|X\|^2 \ge T\lfloor\eta^{-1}\rfloor(16d\log 6+c)-\frac{\Delta}{\eta}\right)\\ &\le 6^{T\lfloor\eta^{-1}\rfloor d}\exp\left(-\frac{T\lfloor\eta^{-1}\rfloor(16d\log 6+c)-\frac{\Delta}{\eta}}{8}\right)\\ &= \exp\left(-\frac{T\lfloor\eta^{-1}\rfloor(8d\log 6+c)-\frac{\Delta}{\eta}}{8}\right).\end{aligned} \tag{F.39}$$

If t, defined above, is negative, then in particular $\frac{\Delta}{\eta} \ge T\lfloor\eta^{-1}\rfloor c$ so that this bounds still (trivially) holds.

Finally, in particular, for $\eta \le 1/2$, $\lfloor\eta^{-1}\rfloor \ge (2\eta)^{-1}$ so we obtain

$$\mathbb{P}_x(\sigma_{\mathcal{D}}>T) \le \exp\left(-\frac{T(8d\log 6+c)}{16\eta}+\frac{\Delta}{8\eta}\right) \tag{F.40}$$

and our proof is complete. ∎

F.2. Transition times. In this section, we leverage the Lyapunov analysis of Appendix F.1 to establish exponential bounds on the transition times of the process, both for exiting large compact sets and for exiting compact sets that do not intersect any equivalence class. The main result of this section is Lemma F.10. Most results of this section follow closely Azizian et al. [9], but we include complete proofs to make this work self-contained and to verify that the arguments remain valid in our more general setting.

We will use the following lemma, which corresponds to Kifer [63, Lem. 5.3].

Lemma F.7. *Let $\mathcal{K}\subset\mathcal{V}$ be a compact set which does not intersect any equivalence class. Then there exists $c>0$, $T\ge 1$, $\eta_0>0$, such that, for any $n>T$, $x\in\mathcal{K}$, $\eta\le\eta_0$,*

$$\mathbb{P}_x\big(\sigma_{\mathcal{V}\setminus\mathcal{K}}>n\big) = \mathbb{P}_x(\tau_{\mathcal{K}}>n) \le \exp(-c(n-T)/\eta)\,. \tag{F.41}$$

Proof. The proof is exactly the same as the proof Kifer [63, Lem. 5.3], which only uses the l.s.c. of $\mathcal{A}_T$ (Corollary D.2). ∎

The following lemma provides a convenient reformulation of the results of Lemma F.6 and Lemma F.7.

Lemma F.8.

- *There exists $\mathcal{D}\subset\mathcal{V}$ a compact set, $\eta_0>0$, such that for any $\mathcal{D}'\subset\mathcal{V}$ compact set such that $\mathcal{D}\subset\mathcal{D}'$, there exists $\lambda_0,a,b>0$ such that,*

$$\forall\eta\le\eta_0,\ \lambda\le\lambda_0,\ x\in\mathcal{D}'\setminus\mathcal{D}\,,\quad \mathbb{E}_x\left[e^{\frac{\lambda\sigma_{\mathcal{D}}}{\eta}}\right] \le e^{\frac{a\lambda}{\eta}+b}\,. \tag{F.42}$$

- *Let $\mathcal{K} \subset \mathcal{V}$ be a compact set which does not intersect any equivalence class. Then, there exists $\eta_0, \lambda_0, a, b > 0$ such that,*

$$\forall \eta \le \eta_0,\, \lambda \le \lambda_0,\, x \in \mathcal{K}\,, \quad \mathbb{E}_x\left[e^{\frac{\lambda \tau_{\mathcal{K}}}{\eta}}\right] \le e^{\frac{a\lambda}{\eta}+b}\,. \tag{F.43}$$

Proof. The proofs of both statements are very similar, so we prove only the second one for notational convenience. Fix $\mathcal{K} \subset \mathcal{V}$ a compact set which does not intersect any equivalence class. By Lemma F.7, there exists $c > 0$, $T \ge 1$ such that, for any $n > T$, $x \in \mathcal{K}$, $\eta \le \eta_0$,

$$\mathbb{P}_x(\tau_{\mathcal{K}} > n) \le \exp(-c(n-T)/\eta)\,. \tag{F.44}$$

Let us first bound $\mathbb{E}_x\left[\exp\left(\frac{\lambda(\tau_{\mathcal{K}}-T)}{\eta}\right)\right]$. We have that

$$\begin{aligned}
\mathbb{E}_x\left[\exp\left(\frac{\lambda(\tau_{\mathcal{K}}-T)}{\eta}\right)\right] &= \int_0^\infty \mathbb{P}_x\left(\exp\left(\frac{\lambda(\tau_{\mathcal{K}}-T)}{\eta}\right) > t\right) dt \\
&\le e^{\frac{\lambda}{\eta}} + \int_{e^{\frac{\lambda}{\eta}}}^\infty \mathbb{P}_x\left(\tau_{\mathcal{K}} > T + \frac{\eta \log t}{\lambda}\right) dt \\
&\le e^{\frac{\lambda}{\eta}} + \int_{e^{\frac{\lambda}{\eta}}}^\infty \mathbb{P}_x\left(\tau_{\mathcal{K}} > T + \left\lfloor \frac{\eta \log t}{\lambda} \right\rfloor\right) dt \\
&\le e^{\frac{\lambda}{\eta}} + \int_{e^{\frac{\lambda}{\eta}}}^\infty \exp\left(-\frac{c}{\eta}\left\lfloor \frac{\eta \log t}{\lambda} \right\rfloor\right) dt\,,
\end{aligned} \tag{F.45}$$

where we used Eq. (F.44) in the last inequality.

Lower bounding $\lfloor s \rfloor$ by $s - 1$, we obtain that

$$\begin{aligned}
\mathbb{E}_x\left[\exp\left(\frac{\lambda(\tau_{\mathcal{K}}-T)}{\eta}\right)\right] &\le e^{\frac{\lambda}{\eta}} + \int_{e^{\frac{\lambda}{\eta}}}^\infty \exp\left(-\frac{c}{\eta}\frac{\eta \log t}{\lambda} + \frac{c}{\eta}\right) dt \\
&= e^{\frac{\lambda}{\eta}} + \int_{e^{\frac{\lambda}{\eta}}}^\infty e^{\frac{c}{\eta}} t^{-\frac{c}{\lambda}}\, dt\,.
\end{aligned} \tag{F.46}$$

Performing the change of variable $s \leftarrow e^{-\frac{\lambda}{\eta}} t$, we obtain that

$$\mathbb{E}_x\left[\exp\left(\frac{\lambda(\tau_{\mathcal{K}}-T)}{\eta}\right)\right] \le e^{\frac{\lambda}{\eta}}\left(1 + \int_1^\infty s^{-\frac{c}{\lambda}}\, ds\right). \tag{F.47}$$

When $\lambda \le c/2$, we obtain that

$$\mathbb{E}_x\left[\exp\left(\frac{\lambda(\tau_{\mathcal{K}}-T)}{\eta}\right)\right] \le e^{\frac{\lambda}{\eta}}\left(1 + \int_1^\infty s^{-2}\, ds\right), \tag{F.48}$$

and therefore,

$$\mathbb{E}_x\left[\exp\left(\frac{\lambda \tau_{\mathcal{K}}}{\eta}\right)\right] \le e^{\frac{(T+1)\lambda}{\eta}}\left(1 + \int_1^\infty s^{-2}\, ds\right), \tag{F.49}$$

which concludes the proof. ∎

The following lemma upper-bounds the probability of exiting a compact set before visiting a neighborhood of the equivalence classes.

Lemma F.9. *Let $\mathcal{U}$ be an open set containing all equivalence classes and $\mathcal{D} \subset \mathcal{V}$ a compact set such that $\mathcal{U} \subset \mathcal{D}$. There exists $\mathcal{D}' \subset \mathcal{V}$ compact set such that $\mathcal{D} \subset \mathcal{D}'$, $\Delta > 0$, $\eta_0 > 0$ such that, for any $\eta \le \eta_0$, $x \in \mathcal{D}$,*

$$\mathbb{P}_x(\tau_{\mathcal{D}'} < \sigma_{\mathcal{U}}) \le \exp\left(-\frac{\Delta}{\eta}\right). \tag{F.50}$$

Proof. Letting $M = \max_{x \in \mathcal{D}} \|x\| + \sqrt{2C/\kappa}$, we consider the two compact sets $\mathcal{D}_0 := \bar{\mathbb{B}}(0, M)$ and $\mathcal{D}' := \bar{\mathbb{B}}(0, M+1)$, satisfying $\mathcal{D} \subset \mathcal{D}_0 \subset \mathcal{D}'$ by construction.

Since the function α_∞ is strictly increasing (cf. Definition D.1), we have

$$\Delta := \frac{\kappa}{2}\left[\alpha_\infty\left((M+1)^2\right) - \alpha_\infty\left(M^2\right)\right] > 0\,. \tag{F.51}$$

Applying Lemma F.7 to $\mathcal{K} \leftarrow \mathcal{D'}_\delta \setminus \mathcal{U}$, there exist $c > 0$, $T_0 \geq 1$, $\eta_0 > 0$ such that for any $n > T_0$, $x \in \mathcal{D}'_\delta \setminus \mathcal{U}$, $\eta \leq \eta_0$,

$$\mathbb{P}_x\left(\tau_{\mathcal{D}'_\delta \setminus \mathcal{U}} > n\right) \leq \exp(-c(n - T_0)/\eta)\,. \tag{F.52}$$

In particular, if we set $T := \left\lceil \frac{\Delta}{c} \right\rceil + T_0$, which is greater or equal than 1, and $n = T$, then we obtain that, for any $\eta \leq \eta_0$, $x \in \mathcal{D}'_\delta \setminus \mathcal{U}$,

$$\mathbb{P}_x\left(\tau_{\mathcal{D}'_\delta \setminus \mathcal{U}} \geq T\right) \leq \exp\left(-\frac{\Delta}{\eta}\right). \tag{F.53}$$

Moreover, this inequality actually holds for all $x \in \mathcal{D}'_\delta$.

We now bound $\mathbb{P}_x\left(\tau_{\mathcal{D}'_\delta} > \sigma_{\mathcal{U}}\right)$ for $x \in \mathcal{D}$ by distinguishing the cases where $\tau_{\mathcal{D}'_\delta} < T$ and $\tau_{\mathcal{D}'_\delta} \geq T$. For any $x \in \mathcal{D}$, we have that

$$\begin{aligned}
\mathbb{P}_x\left(\tau_{\mathcal{D}'_\delta} < \sigma_{\mathcal{U}}\right) &\leq \mathbb{P}_x\left(\tau_{\mathcal{D}'_\delta} < \sigma_{\mathcal{U}},\, \tau_{\mathcal{D}'_\delta} < T\right) + \mathbb{P}_x\left(\tau_{\mathcal{D}'_\delta} < \sigma_{\mathcal{U}},\, \tau_{\mathcal{D}'_\delta} \geq T\right) \\
&\leq \mathbb{P}_x\left(\tau_{\mathcal{D}'_\delta} < T\right) + \mathbb{P}_x\left(\tau_{\mathcal{D}'_\delta \setminus \mathcal{U}} \geq T\right) \\
&\leq \mathbb{P}_x\left(\tau_{\mathcal{D}'_\delta} < T\right) + \exp\left(-\frac{\Delta}{\eta}\right)
\end{aligned} \tag{F.54}$$

where we used Eq. (F.53).

We now focus on bounding the first term. For this, we first show that $\tau_{\mathcal{D}'_\delta} < T$ implies that

$$\mathrm{dist}_T\left(x^\eta, \Gamma_T^{\{x\}}(\Delta/2)\right) > \delta/2\,. \tag{F.55}$$

For the sake of contradiction, suppose that this inequality does not hold. Therefore, there must exist $\xi \in \Gamma_T^{\{x\}}(\Delta/2)$ such that $\mathrm{dist}_T(x^\eta, \xi) < \delta$. In particular, there is some $n < T$ such that $\xi_n \notin \mathcal{D}'$, so that, by Definition E.3,

$$\mathcal{A}_T(\xi) \geq \mathcal{A}_{n+1}(\xi) \geq B(\xi_0, \xi_n)\,. \tag{F.56}$$

As a consequence, $B(\xi_0, \xi_n) \leq \Delta/2$ and we can then find paths $\gamma^k \in \mathcal{C}([0, T_k], \mathcal{V})$ such that $\gamma^k_0 = \xi_0$, $\gamma^k_{T_k} = \xi_n$, and $\mathcal{S}_{[0,T_k]}(\gamma^k) \leq B(\xi_0, \xi_n) + 1/k$. Since $\gamma^k_0 \in \mathcal{D} \subset \mathcal{D}_0$, $\gamma^k_{T_k} \notin \mathcal{D}'$ and γ^k is continuous, there exists $s_0 := s_0(k) < T_k$ such that $\gamma^k_{s_0} \in \mathrm{bd}\,\mathcal{D}_0$ and $\gamma^k_s \notin \mathcal{D}_0$ for all $s \in (s_0, T_k]$. In particular, this means that $\min_{s_0 \leq s \leq T_k} \|\gamma^k_s\| \geq \sqrt{2C/\kappa}$ and thus Lemma E.16 yields

$$\begin{aligned}
\mathcal{S}_{[0,T_k]}(\gamma^k) \geq \mathcal{S}_{[s_0,T_k]}(\gamma^k) &\geq \frac{\kappa}{4}\left[g_\infty(\gamma^k_{T_k}) - g_\infty(\gamma^k_{s_0})\right] \\
&= \frac{\kappa}{2}\left[\alpha_\infty(\|\xi_n\|^2) - \alpha_\infty(\|\gamma^k_{s_0}\|^2)\right] \\
&\geq \frac{\kappa}{2}\left[\alpha_\infty\left((M+1)^2\right) - \alpha_\infty\left(M^2\right)\right] \\
&= \Delta.
\end{aligned} \tag{F.57}$$

But at the same time, we also have

$$\mathcal{S}_{[0,T_k]}(\gamma^k) \leq B(\xi_0, \xi_n) + \frac{1}{k} \leq \frac{\Delta}{2} + \frac{1}{k} < \Delta \tag{F.58}$$

for k sufficiently large, which therefore contradicts Eq. (F.57).

It then follows that

$$\begin{aligned}
\mathbb{P}\left(\tau_{\mathcal{D}'_\delta} < T\right) &\leq \mathbb{P}\left(\mathrm{dist}_T\left(x^\eta, \Gamma_T^{\{x\}}(\Delta/2)\right) > \delta/2\right) \\
&\leq \exp\left(-\frac{\Delta}{4\eta}\right)
\end{aligned} \tag{F.59}$$

where the second inequality comes from Corollary D.2 with $\delta \leftarrow \delta/2$, $s \leftarrow \Delta/2$, $\epsilon \leftarrow \Delta/4$. ∎

The following lemma is key to our analysis: it shows that trajectories of (SD) return quickly to neighborhoods of the equivalence classes.

Lemma F.10. *Let $\mathcal{U}$ be an open set containing all equivalence classes and $\mathcal{D} \subset \mathcal{V}$ a compact set such that $\mathcal{U} \subset \mathcal{D}$. Then, there is some $\eta_0, \lambda_0, a, b > 0$ such that,*

$$\forall \eta \le \eta_0,\ \lambda \le \lambda_0,\ x \in \mathcal{D}\,, \quad \mathbb{E}_x\left[e^{\frac{\lambda \sigma_{\mathcal{U}}}{\eta}}\right] \le e^{\frac{a\lambda}{\eta}+b}\,. \tag{F.60}$$

Proof. Fix $\epsilon > 0$. Without loss of generality, assume that $\mathcal{D}$ is large enough to include the compact set given by the first item of Lemma F.8 (note that the guarantee of the first item of Lemma F.8 still holds even if $\mathcal{D}$ is larger).

Apply Lemma F.9 with $\mathcal{U} \leftarrow \mathcal{U}$, $\mathcal{D} \leftarrow \mathcal{D}$ and denote by $\tilde{\mathcal{D}}$ the obtained compact and $\eta_0, \Delta > 0$ such that, for every $\eta \le \eta_0$, $x \in \mathcal{D}$,

$$\mathbb{P}_x\big(\tau_{\tilde{\mathcal{D}}} < \sigma_{\mathcal{U}}\big) \le \exp\left(-\frac{\Delta}{\eta}\right). \tag{F.61}$$

Define $r \coloneqq \sup_{x\in\tilde{\mathcal{D}}}\|x\|$ and $R = e^{8M}(1+r)$. Assuming that $\eta \le 1$, by Lemma D.1, for any $x \in \tilde{\mathcal{D}}$, the next two iterates of $(x_n^\eta)_{n\ge 0}$ satisfies $\|x_1^\eta\| \le R$ and $\|x_2^\eta\| \le R$. Define $\mathcal{D}' \coloneqq \bar{\mathbb{B}}(0,R)$.

We invoke both items of Lemma F.8 with $\mathcal{K} \leftarrow \tilde{\mathcal{D}} \setminus \mathcal{U}$ and denote by $c > 0$ a constant that satisfies the bounds of both items. In the rest of the proof, we consider η and λ smaller than the bounds given by this lemma.

Our goal is to bound, for any $T \ge 0$, the quantity,

$$s_T(\lambda,\eta) \coloneqq \sup_{x\in\mathcal{D}} \mathbb{E}_x\left[\exp\left(\frac{\lambda\sigma_{\mathcal{U}}^T}{\eta}\right)\right], \quad \text{where} \quad \sigma_{\mathcal{U}}^T \coloneqq \min(T,\sigma_{\mathcal{U}})\,. \tag{F.62}$$

Note that, by construction, $s_T(\lambda,\eta)$ is finite.

Take $x \in \mathcal{D}$. In particular, Lemma F.8 implies that $\tau_{\tilde{\mathcal{D}}\setminus\mathcal{U}} < +\infty$ almost surely for all $x \in \mathcal{D}$.

$$\begin{aligned}
\exp\left(\frac{\lambda\sigma_{\mathcal{U}}^T}{\eta}\right) &= \mathbb{1}\{x^\eta_{\tau_{\tilde{\mathcal{D}}\setminus\mathcal{U}}} \in \mathcal{U}\}\exp\left(\frac{\lambda\tau_{\tilde{\mathcal{D}}\setminus\mathcal{U}}}{\eta}\right) \\
&\qquad + \mathbb{1}\{x^\eta_{\tau_{\tilde{\mathcal{D}}\setminus\mathcal{U}}} \notin \mathcal{D}_\delta\}\exp\left(\frac{\lambda\tau_{\tilde{\mathcal{D}}\setminus\mathcal{U}}}{\eta}\right)\exp\left(\frac{\lambda\left(\sigma_{\mathcal{U}}^T - \tau_{\tilde{\mathcal{D}}\setminus\mathcal{U}}\right)}{\eta}\right) \\
&\le \mathbb{1}\{x^\eta_{\tau_{\tilde{\mathcal{D}}\setminus\mathcal{U}}} \in \mathcal{U}\}\exp\left(\frac{\lambda\tau_{\tilde{\mathcal{D}}\setminus\mathcal{U}}}{\eta}\right) \\
&\qquad + \mathbb{1}\{x^\eta_{\tau_{\tilde{\mathcal{D}}\setminus\mathcal{U}}} \notin \mathcal{D}_\delta\}\exp\left(\frac{\lambda\tau_{\tilde{\mathcal{D}}\setminus\mathcal{U}}}{\eta}\right)\exp\left(\frac{\lambda\min\left(\sigma_{\mathcal{U}} - \tau_{\tilde{\mathcal{D}}\setminus\mathcal{U}}, T\right)}{\eta}\right)
\end{aligned} \tag{F.63}$$

so that we can apply the strong Markov property to the Markov chain $(x_n^\eta)_{n\ge0}$ with stopping time $\tau_{\tilde{\mathcal{D}}\setminus\mathcal{U}}$ to obtain that

$$\begin{aligned}
\mathbb{E}_x\left[\exp\left(\frac{\lambda\sigma_{\mathcal{U}}^T}{\eta}\right)\right] \le \mathbb{E}_x\Bigg[&\mathbb{1}\{x^\eta_{\tau_{\tilde{\mathcal{D}}\setminus\mathcal{U}}} \in \mathcal{U}\}\exp\left(\frac{\lambda\tau_{\tilde{\mathcal{D}}\setminus\mathcal{U}}}{\eta}\right) \\
&+ \mathbb{1}\{x^\eta_{\tau_{\tilde{\mathcal{D}}\setminus\mathcal{U}}} \notin \mathcal{D}_\delta\}\exp\left(\frac{\lambda\tau_{\tilde{\mathcal{D}}\setminus\mathcal{U}}}{\eta}\right)\mathbb{E}_{x^\eta_{\tau_{\tilde{\mathcal{D}}\setminus\mathcal{U}}}}\left[\exp\left(\frac{\lambda\sigma_{\mathcal{U}}^T}{\eta}\right)\right]\Bigg]
\end{aligned} \tag{F.64}$$

We now bound

$$\mathbb{1}\{x^\eta_{\tau_{\tilde{\mathcal{D}}\setminus\mathcal{U}}} \notin \mathcal{D}_\delta\}\,\mathbb{E}_{x^\eta_{\tau_{\tilde{\mathcal{D}}\setminus\mathcal{U}}}}\left[\exp\left(\frac{\lambda\sigma_{\mathcal{U}}^T}{\eta}\right)\right] = \mathbb{E}_{x^\eta_{\tau_{\tilde{\mathcal{D}}\setminus\mathcal{U}}}}\left[\mathbb{1}\{x_0^\eta \notin \mathcal{D}_\delta\}\exp\left(\frac{\lambda\sigma_{\mathcal{U}}^T}{\eta}\right)\right] \tag{F.65}$$

Since x is in $\mathcal{D}$, by definition, $\tau_{\tilde{\mathcal{D}}\setminus\mathcal{U}}$ is at least equal to 1 and $x^\eta_{\tau_{\tilde{\mathcal{D}}\setminus\mathcal{U}}-1}$ is still in $\tilde{\mathcal{D}}$. By definition of $\mathcal{D}'$, $x^\eta_{\tau_{\tilde{\mathcal{D}}\setminus\mathcal{U}}}$ must still be in $\mathcal{D}'$. Therefore, the guarantee of Lemma F.8 applies and, since in particular it implies that $\sigma_{\mathcal{D}}$ is finite almost surely when the chain is started at $x^\eta_{\tau_{\tilde{\mathcal{D}}\setminus\mathcal{U}}}$, we can apply the strong Markov property to obtain that

$$\begin{aligned}\mathbb{E}_{x^\eta_{\tau_{\tilde{\mathcal{D}}\setminus\mathcal{U}}}}\left[\mathbb{1}\{x_0\notin\tilde{\mathcal{D}}\}\exp\left(\frac{\lambda\sigma^T_{\mathcal{U}}}{\eta}\right)\right] &\le \mathbb{E}_{x^\eta_{\tau_{\tilde{\mathcal{D}}\setminus\mathcal{U}}}}\left[\mathbb{1}\{x_0\notin\tilde{\mathcal{D}}\}\exp\left(\frac{\lambda\sigma_{\mathcal{D}}}{\eta}\right)\mathbb{E}_{x^\eta_{\sigma_{\mathcal{D}}}}\left[\exp\left(\frac{\lambda\sigma^T_{\mathcal{U}}}{\eta}\right)\right]\right]\\ &\le s_T(\lambda,\eta)\,\mathbb{E}_{x^\eta_{\tau_{\tilde{\mathcal{D}}\setminus\mathcal{U}}}}\left[\mathbb{1}\{x_0\notin\tilde{\mathcal{D}}\}\exp\left(\frac{\lambda\sigma_{\mathcal{D}}}{\eta}\right)\right] \qquad \text{(F.66)}\end{aligned}$$

where, for the second inequality, we used the definition of $s_T(\lambda,\eta)$ and the fact that $x_{\sigma_{\mathcal{D}}}$ is in $\mathcal{D}$.

Now, to bound the remaining expectation, note that x_0 does not belong to $\tilde{\mathcal{D}}$ and, *a fortiori*, does not belong to $\mathcal{D}$. Therefore, $\sigma_{\mathcal{D}}$ depends only on $(x^\eta_n)_{n\ge1}$ and the (weak) Markov property implies that

$$\begin{aligned}\mathbb{E}_{x^\eta_{\tau_{\tilde{\mathcal{D}}\setminus\mathcal{U}}}}\left[\mathbb{1}\{x_0\notin\tilde{\mathcal{D}}\}\exp\left(\frac{\lambda\sigma_{\mathcal{D}}}{\eta}\right)\right] &\le \mathbb{E}_{x^\eta_{\tau_{\tilde{\mathcal{D}}\setminus\mathcal{U}}}}\left[\mathbb{1}\{x_0\notin\tilde{\mathcal{D}}\}\,\mathbb{E}_{x_1}\left[\exp\left(\frac{\lambda(1+\sigma_{\mathcal{D}})}{\eta}\right)\right]\right]\\ &\le \exp\left(\frac{\lambda(1+c)}{\eta}\right)\mathbb{E}_{x^\eta_{\tau_{\tilde{\mathcal{D}}\setminus\mathcal{U}}}}\left[\mathbb{1}\{x_0\notin\tilde{\mathcal{D}}\}\right], \qquad \text{(F.67)}\end{aligned}$$

by Lemma F.8, since the $x^\eta_{\tau_{\tilde{\mathcal{D}}\setminus\mathcal{U}}-1}$ is still in $\tilde{\mathcal{D}}$ so that x_1 of the chain started at $x^\eta_{\tau_{\tilde{\mathcal{D}}\setminus\mathcal{U}}}$ is still in $\mathcal{D}'$.

Lemma F.9 then implies that,

$$\mathbb{E}_{x^\eta_{\tau_{\tilde{\mathcal{D}}\setminus\mathcal{U}}}}\left[\mathbb{1}\{x_0\notin\tilde{\mathcal{D}}\}\exp\left(\frac{\lambda\sigma_{\mathcal{D}'}}{\eta}\right)\right]\le\exp\left(\frac{\lambda(1+a)-\Delta}{\eta}+b\right) \qquad \text{(F.68)}$$

Combining these bounds with Eq. (F.64) then gives

$$\begin{aligned}\mathbb{E}_x\left[\exp\left(\frac{\lambda\sigma^T_{\mathcal{U}}}{\eta}\right)\right] &\le \left(1+s_T(\lambda,\eta)\exp\left(\frac{\lambda(1+a)-\Delta}{\eta}+b\right)\right)\mathbb{E}_x\left[\exp\left(\frac{\lambda\tau_{\tilde{\mathcal{D}}\setminus\mathcal{U}}}{\eta}\right)\right]\\ &\le \left(1+s_T(\lambda,\eta)\exp\left(\frac{\lambda(1+a)-\Delta}{\eta}+b\right)\right)\exp\left(\frac{a\lambda}{\eta}+b\right), \qquad \text{(F.69)}\end{aligned}$$

by Lemma F.8. Since this inequality is valid for any $x\in\mathcal{D}$, we have shown that

$$s_T(\lambda,\eta)\le\left(e^{\frac{a\lambda}{\eta}+b}+s_T(\lambda,\eta)e^{\frac{\lambda(1+2a)-\Delta}{\eta}+b}\right). \qquad \text{(F.70)}$$

For $\lambda\le\Delta/(2(1+2a))$ and η small enough,

$$e^{\frac{\lambda(1+2a)-\Delta}{\eta}+b}\le e^{-\frac{\Delta}{2\eta}+b}\le\frac{1}{2}, \qquad \text{(F.71)}$$

and we obtain that

$$s_T(\lambda,\eta)\le 2e^{\frac{a\lambda}{\eta}+b}. \qquad \text{(F.72)}$$

Taking $T\to+\infty$ and using Fatou's lemma yields that

$$\sup_{x\in\mathcal{D}}\mathbb{E}_x\left[\exp\left(\frac{\lambda\sigma_{\mathcal{U}}}{\eta}\right)\right]\le 2e^{\frac{a\lambda}{\eta}+b}, \qquad \text{(F.73)}$$

which concludes the proof. ∎

F.3. Transition probabilities. In this section, we derive exponential bounds on the transition probabilities between equivalence classes, which are the key ingredient in the analysis of the long-run distribution of the process. These bounds crucially rely on the regularity properties of the quasi-potential near equivalence classes established in Appendix E as well as the transition time estimates of Appendix F.2. Most results of this section follow closely Azizian et al. [9], but we include complete proofs to make this work self-contained and to verify that the arguments remain valid in our more general setting.

Starting from this section, we suppose throughout that the following assumption holds true (which corresponds to Assumption 2 in the main text):

Assumption 5.

(*a*) There are finitely many basic equivalence classes.

(*b*) $B(\mathcal{K}, \mathcal{K}') < \infty$ for every basic equivalence classes $\mathcal{K}$ and $\mathcal{K}'$.

We then write $\mathcal{K}_1, \dots, \mathcal{K}_K$ for the collection of all equivalence classes, and we define their associated *transition costs*

$$B_{i,j} = B(\mathcal{K}_i, \mathcal{K}_j) = \inf\{B(x, y) : x \in \mathcal{K}_i, y \in \mathcal{K}_j\} . \tag{F.74}$$

Moreover, we also denote as $\mathbf{K}$ the union of all $\mathcal{K}_i$, $1 \le i \le K$.

Remark F.1. The requirement that all the transition costs be finite is only to avoid degenerate cases where some transitions are impossible, and thus to keep the presentation of the results and proofs simpler. If it did not hold, the same analysis could still be carried out. We would consider the components of the graph $\mathcal{G}$ connected by edges with finite weights, proceed with the proof on each of them, and obtain the same results on each of these components. To keep the complexity of the proof reasonable, we will not consider this case here.

Moreover, note that this requirement can be satisfied in particular if, for any i, j, there exists a C^1 path γ joining $\mathcal{K}_i$ and $\mathcal{K}_j$ such that, for all t, $-F(\gamma(t))$ belongs to the interior of the closed convex hull of the support of the noise $U(\gamma(t), \omega)$:

$$-F(\gamma(t)) \in \operatorname{int} \overline{\operatorname{conv}} \operatorname{supp} U(\gamma(t), \omega) . \tag{F.75}$$

This means there is a sufficient level of noise for the probability of going from $\mathcal{K}_i$ to $\mathcal{K}_j$ with at least one path to be non-zero (though it can be vanishingly small), see Azizian et al. [9, Lem. D.12].

We adapt to our context Kifer [63, Lem. 5.4] and simplify it using ideas from Freidlin & Wentzell [36, Chap. 6].

Definition F.2 (Freidlin & Wentzell [36, Chap. 6,§2])**.** For i, j,

$$\widetilde{B}_{i,j} := \inf\left\{ \mathcal{A}_T(\xi) : T \ge 1, \xi \in \mathcal{V}^T, \xi_0 \in \mathcal{K}_i, \xi_{T-1} \in \mathcal{K}_j, \xi_n \notin \bigcup_{k \ne i,j} \mathcal{K}_k \text{ for } n = 1, \dots, T-2 \right\} . \tag{F.76}$$

We now defined an important object: the law of the (accelerated) iterated at the first time they reach some set $\mathcal{U}$ (typically a neighborhood of $\mathbf{K}$), following e.g., Douc et al. [33, Chap. 3.4].

Definition F.3 (Kifer [63, Prop. 5.3])**.** For $\mathcal{U}$ open set, with hitting time (of the accelerated sequence) $\sigma_{\mathcal{U}} := \inf\{n \ge 1 : x_n^\eta \in \mathcal{U}\}$, (as in Definition F.1), denote the law of $x^\eta_{\sigma_{\mathcal{U}}}$ started at x by $\mathbb{Q}_{\mathcal{U}}(x, \cdot)$ and the corresponding T-step transition probability by $\mathbb{Q}^T_{\mathcal{U}}(x, \cdot)$.

In words, $\mathbb{Q}_{\mathcal{U}}(x, \cdot)$ is the distributions of the x^η started at x at the first time they reach $\mathcal{U}$. Typically, $\mathbb{Q}_{\mathcal{U}}(x, \mathcal{U})$ is the probability that $\mathcal{U} \subset \mathcal{U}$ is reached first among $\mathcal{U}$. Then, $\mathbb{Q}^T_{\mathcal{U}}(x, \cdot)$ is the distribution of the x^η started at x at the T-th time they reach $\mathcal{U}$.

We first give estimates of the transition probabilities using the $\widetilde{B}_{i,j}$. We will then translate them to $B_{i,j}$.

Lemma F.11. *For any $\epsilon > 0$, for any small enough neighborhoods $\mathcal{U}_i$ of $\mathcal{K}_i$, $i = 1, \dots, K$, there is some $\eta_0 > 0$ such that for all i, j, $x \in \mathcal{U}_i$, $0 < \eta < \eta_0$,*

$$\mathbb{Q}_{\mathcal{U}}(x, \mathcal{U}_j) \le \exp\left(-\frac{\widetilde{B}_{i,j}}{\eta} + \frac{\epsilon}{\eta} \right) . \tag{F.77}$$

where we defined $\mathcal{U} := \bigcup_{i=1}^K \mathcal{U}_i$.

Proof. Assume that, without loss of generality, ϵ is small enough so that Corollary E.4 with $r \leftarrow \epsilon$ can be applied to every $\mathcal{K}_i$, $i = 1, \dots, K$. Denote by $\mathcal{W}_i$, $i = 1, \dots, K$ the corresponding neighborhoods of $\mathcal{K}_i$.

Since these $\mathcal{W}_i$'s are neighborhoods of the $\mathcal{K}_i$'s, there exists $\delta > 0$ such that $\mathcal{U}_{2\delta}(\mathcal{K}_i) \subset \mathcal{W}_i$ for all $i = 1, \dots, K$. Require then that $\mathcal{U}_i$ be contained in $\mathcal{U}_\delta(\mathcal{K}_i)$ so that $\mathcal{U}_\delta(\mathcal{U}_i) \subset \mathcal{W}_i$ for all $i = 1, \dots, K$.

Moreover, assume that $\delta > 0$ is small enough so that the neighborhoods $\mathcal{U}_\delta(\mathcal{K}_i)$, $i = 1, \dots, K$ are pairwise disjoint.

Define $0 < \delta' \leq \delta$ such that $\mathcal{U}'_\delta(\mathcal{K}_i)$ is contained in $\mathcal{U}_i$ for all $i = 1, \dots, K$.

Fix $i, j \in \mathcal{I}$ and consider $\xi \in (\mathcal{V})^T$ such that $\xi_0 \in \mathcal{U}'_\delta(\mathcal{U}_i) \subset \mathcal{W}_i$, $\xi_{T-1} \in \mathcal{U}'_\delta(\mathcal{U}_j) \subset \mathcal{W}_j$ and $\xi_n \in \mathcal{U}_{\delta'}\big(\mathcal{V} \setminus \bigcup_{k \neq i,j} \mathcal{U}_k\big)$ for all $n = 1, \dots, T-2$. By the choice of δ', ξ_n cannot be in $\bigcup_{k \neq i,j} \mathcal{K}_k$ for any $n = 1, \dots, T-2$.

By definition of $\mathcal{W}_i$ and $\mathcal{W}_j$, there are $x \in \mathcal{K}_i$, $x' \in \mathcal{K}_j$ such that $\rho(x, \xi_0) < \epsilon$, $\rho(\xi_{T-1}, x') < \epsilon$. Therefore, the path $\zeta \in (\mathcal{V})^{T+2}$ defined as $\zeta = (x, \xi_0, \xi_1, \dots, \xi_{T-1}, x')$ satisfies

$$\mathcal{A}_T(\xi) \geq \mathcal{A}_{T+2}(\zeta) - 2\epsilon\,. \tag{F.78}$$

and, by definition of $\widetilde{B}_{i,j}$, we thus obtain,

$$\mathcal{A}_T(\xi) \geq \mathcal{A}_{T+2}(\zeta) - 2\epsilon \geq \widetilde{B}_{i,j} - 2\epsilon\,. \tag{F.79}$$

Fix $x \in \mathcal{U}_i$. Let us now bound the probability

$$\mathbb{Q}_{\mathcal{U}}(x, \mathcal{U}_j) = \mathbb{P}_x\big(x^\eta_{\sigma_{\mathcal{U}}} \in \mathcal{U}_i\big)\,. \tag{F.80}$$

We have, for any $T \geq 0$,

$$\mathbb{P}_x\big(x^\eta_{\sigma_{\mathcal{U}}} \in \mathcal{U}_i\big) \leq \mathbb{P}_x\big(x^\eta_{\sigma_{\mathcal{U}}} \in \mathcal{U}_i,\, \sigma_{\mathcal{U}} < T\big) + \mathbb{P}_x(\sigma_{\mathcal{U}} \geq T)\,. \tag{F.81}$$

We first bound the second probability using Lemma F.10 applied to $\mathcal{U} \leftarrow \mathcal{U}_i$. Take T such that $\lambda_0(a - T) + \eta_0 b \leq -\widetilde{B}_{i,j}$. Then, by Markov's inequality and Lemma F.10, it holds that for all $\eta \leq \eta_0$

$$\begin{aligned}
\mathbb{P}_x(\sigma_{\mathcal{U}} \geq T) &\leq \mathbb{P}_x\left(\exp\left(\frac{\lambda_0 \sigma_{\mathcal{U}}}{\eta}\right) \geq \exp\left(\frac{\lambda_0 T}{\eta}\right)\right) \\
&\leq \exp\left(\frac{\lambda_0(a - T)}{\eta} + b\right) \\
&\leq \exp\left(\frac{-\widetilde{B}_{i,j}}{\eta}\right).
\end{aligned} \tag{F.82}$$

We now bound the term $\mathbb{P}_x\big(x^\eta_{\sigma_{\mathcal{U}}} \in \mathcal{U}_j,\, \sigma_{\mathcal{U}} < T\big)$ for this choice of T.

For this, we show that $x^\eta_{\sigma_{\mathcal{U}}} \in \mathcal{U}_j$ with $\sigma_{\mathcal{U}} < T$ implies that

$$\mathrm{dist}_T\Big(x^\eta, \Gamma_T(\{x\}; \widetilde{B}_{i,j} - 3\epsilon)\Big) > \frac{\delta'}{2}\,. \tag{F.83}$$

Indeed, on the event $x^\eta_{\sigma_{\mathcal{U}}} \in \mathcal{U}_j$ with $\sigma_{\mathcal{U}} < T$, there is some $T' \leq T$ such that $\sigma_{\mathcal{U}} = T' - 1$. If $\mathrm{dist}_T\Big(x^\eta, \Gamma_T(\{x\}; \widetilde{B}_{i,j} - 3\epsilon)\Big) > \frac{\delta'}{2}$ did not hold, this would mean that there exists $\xi \in (\mathcal{V})^{T'}$ such that $\mathrm{dist}_{T'}(x^\eta, \xi) < \delta'$, $\xi_0 = x$ and, $\mathcal{A}_{T'}(\xi) \leq \widetilde{B}_{i,j} - 3\epsilon$. In particular, ξ would also satisfy $\xi_{T'-1} \in \mathcal{U}_{\delta'}(\mathcal{U}_j)$, $\xi_n \in \mathcal{U}_{\delta'}(\mathcal{V} \setminus \mathcal{U})$ for all $n = 1, \dots, T' - 2$. This would be in direct contradiction of Eq. (F.79).

Therefore, we have that

$$\begin{aligned}
\mathbb{P}_x\big(x^\eta_{\sigma_{\mathcal{U}}} \in \mathcal{U}_j,\, \sigma_{\mathcal{U}} < T\big) &\leq \mathbb{P}_x\left(\mathrm{dist}_T\Big(x^\eta, \Gamma_T(\{x\}; \widetilde{B}_{i,j} - 3\epsilon)\Big) > \frac{\delta'}{2}\right) \\
&\leq \exp\left(-\frac{\widetilde{B}_{i,j} - 4\epsilon}{\eta}\right),
\end{aligned} \tag{F.84}$$

by Corollary D.2.

Combining this bound with Eq. (F.82) yields

$$\mathbb{P}_x(x^\eta_{\sigma_{\mathcal{U}}} \in \mathcal{U}_j) \leq \exp\left(-\frac{\widetilde{B}_{i,j} - 4\epsilon}{\eta}\right) + \exp\left(\frac{-\widetilde{B}_{i,j}}{\eta}\right), \tag{F.85}$$

which concludes the proof. ∎

Lemma F.12. *For any $\epsilon > 0$, for any neighborhoods $\mathcal{U}_i$ of $\mathcal{K}_i$, $i = 1, \dots, K$ small enough, there exists $T \geq 0$, $\eta_0 > 0$ such that for all i, j, $x \in \mathcal{U}_i$, $0 < \eta < \eta_0$,*

$$\mathbb{Q}^T_{\mathcal{U}}(x, \mathcal{U}_j) \geq \exp\left(-\frac{\widetilde{B}_{i,j}}{\eta} - \frac{\epsilon}{\eta}\right). \tag{F.86}$$

Proof. For any i, j, there exists $T_{i,j} \geq 1$, $\xi^{i,j} \in (\mathcal{V})^{T_{i,j}}$ such that $\xi^{i,j}_0 \in \mathcal{K}_i$, $\xi^{i,j}_{T_{i,j}-1} \in \mathcal{K}_j$, $\xi^{i,j}_n \notin \cup_k \mathcal{K}_k$ for all $n = 1, \dots, T_{i,j} - 2$ and $\mathcal{A}_{T_{i,j}}(\xi^{i,j}) \leq \widetilde{B}_{i,j} + \epsilon$. Define $\delta_{i,j} \coloneqq \min\left\{d\left(\xi^{i,j}_n, \cup_k \mathcal{K}_k\right) : n = 1, \dots, T_{i,j} - 2\right\}$ and $\delta \coloneqq \min_{i,j \in \mathcal{I}} \delta_{i,j}$. By construction, it holds that $\delta > 0$.

Require that $\mathcal{U}_i$ be contained in $\mathcal{W}_i \cap \mathcal{U}_{\delta/2}(\mathcal{K}_i)$ for all $i = 1, \dots, K$. Now, given such $\mathcal{U}_i$ neighborhoods of $\mathcal{K}_i$, $i = 1, \dots, K$, there exists $0 < \delta' \leq \delta/2$ such that $\mathcal{U}'_\delta(\mathcal{K}_i)$ is contained in $\mathcal{U}_i$ for all $i = 1, \dots, K$.

Apply Lemma E.25 to $\mathcal{K}_i$, $i = 1, \dots, K$ with $\epsilon \leftarrow \min(\epsilon, \delta'/2)$ and denote by T_i the bound on the length of paths obtained. Define

$$T \coloneqq \max_{i \in \mathcal{I}} T_i + 1\,. \tag{F.87}$$

Fix $i, j \in \mathcal{I}$ and $x \in \mathcal{U}_i$. Since $\mathcal{U}_i \subset \mathcal{W}_i$, there exists $y \in \mathcal{K}_i$ such that $\rho(x, y) < \epsilon$.

By Lemma E.25, there exists $n \leq T_i \leq T - 1$, $\xi \in (\mathcal{V})^n$ such that $\xi_0 = y$, $\xi_{n-1} = \xi^{i,j}_0$, $\xi_k \in \mathcal{U}_{\delta'/2}(\mathcal{K}_i)$ for all $k = 1, \dots, n-2$ and $\mathcal{A}_n(\xi) < \epsilon$.

Considering the concatenation

$$\zeta \coloneqq \left(x, \xi_0, \xi_1, \dots, \xi_{n-2}, \xi_{n-1}, \xi^{i,j}_1, \dots, \xi^{i,j}_{T_{i,j}-1}, \phi(\xi^{i,j}_{T_{i,j}-1}), \phi^2(\xi^{i,j}_{T_{i,j}-1}), \dots, \phi^{T-n-1}(\xi^{i,j}_{T_{i,j}-1})\right) \tag{F.88}$$

which is a path of length $T + T_{i,j} - 1$ made of $x \in \mathcal{U}_i$, then exactly n points in $\mathcal{U}_{\delta'/2}(\mathcal{K}_i)$ then $T_{i,j} - 2$ in $\mathcal{V} \setminus \mathcal{U}_{\delta/2}(\mathcal{U})$, $\xi^{i,j}_{T_{i,j}-1} \in \mathcal{K}_j$ followed by $T - n - 1$ points in $\mathcal{K}_j$ (with this last list omitted if $n = T - 1$). Moreover, by construction, and, by invariance of $\mathcal{K}_j$ under the flow (Lemma E.7), $\mathcal{A}_{T+T_{i,j}-1}(\zeta) \leq \widetilde{B}_{i,j} + 3\epsilon$. Therefore, if

$$\mathrm{dist}_{T+T_{i,j}-1}(x^\eta, \zeta) < \delta'/2\,, \tag{F.89}$$

with $x^\eta_0 = x$, then $x^\eta_1, \dots, x^\eta_n$ are in $\mathcal{U}_{\delta'}(\mathcal{K}_i) \subset \mathcal{U}_i$ and, since $\delta' \leq \delta/2$, $x^\eta_{n+1}, \dots, x^\eta_{n+T_{i,j}-2}$ are not in $\mathcal{U}_{\delta/4}(\mathcal{U})$, and therefore not in $\mathcal{U}$. Moreover, $x^\eta_{n+T_{i,j}-1}, \dots, x^\eta_{T+T_{i,j}-2}$ would be in $\mathcal{U}_{\delta'}(\mathcal{K}_j) \subset \mathcal{U}_j$.

Thus, all the paths x^η satisfying (F.89) with $x^\eta_0 = x$ are started at x and, after n visits to $\mathcal{U}_i$, their remaining visits to $\mathcal{U}$ up to the T-th one belong to $\mathcal{U}_j$. In particular, their T-th point that falls into $\mathcal{U}$ belongs to $\mathcal{U}_j$.

Therefore, using the definition of $\mathbb{Q}^T_{\mathcal{U}}$, we have that

$$\begin{aligned}\mathbb{Q}^T_{\mathcal{U}}(x, \mathcal{U}_j) &\geq \mathbb{P}_x\left(\mathrm{dist}_{T+T_{i,j}-1}(x^\eta, \zeta) < \delta'/2\right) \\ &\geq \exp\left(-\frac{\widetilde{B}_{i,j} + 4\epsilon}{\eta}\right),\end{aligned} \tag{F.90}$$

by Corollary D.2. ∎

Lemma F.13. *For any $i, j \in \mathcal{I}$,*

$$\begin{aligned}B_{i,j} &= \min\left\{\sum_{k=0}^{n-2} \widetilde{B}_{i_k, i_{k+1}} : i_0 = i,\, i_{n-1} = j,\, i_k \in \mathcal{I} \text{ for } k = 1, \dots, n-2, n \geq 1\right\} \\ &= \min\left\{\sum_{k=0}^{K-2} \widetilde{B}_{i_k, i_{k+1}} : i_0 = i,\, i_{K-1} = j,\, i_k \in \mathcal{I} \text{ for } k = 1, \dots, K-2\right\}.\end{aligned} \tag{F.91}$$

Proof. It suffices to show that

$$B_{i,j} = \min\left\{\sum_{k=0}^{n-2} \widetilde{B}_{i_k,i_{k+1}} : i_0 = i,\, i_{n-1} = j,\, i_k \in \mathcal{I} \text{ for } k = 1,\dots,n-2\,, n \geq 1\right\}. \tag{F.92}$$

The statement of the lemma then follows from the fact that $\widetilde{B}_{k,k} = 0$ for all $k \in \mathcal{I}$ and that shortest paths on graphs can be chosen not to visit the same node twice.

For the inequality ($\geq$), notice that any path between $\mathcal{K}_i$ and $\mathcal{K}_j$ can be decomposed into a concatenation of paths between $\mathcal{K}_i$ and $\mathcal{K}_{i_1}$, $\mathcal{K}_{i_1}$ and $\mathcal{K}_{i_2}$, …, $\mathcal{K}_{i_{n-1}}$ and $\mathcal{K}_j$ for some $i_1,\dots,i_{n-1} \in \mathcal{I}$ that do not enter any other equivalence class in between. Therefore, the inequality ($\geq$) follows from the definition of $B_{i,j}$ and $\widetilde{B}_{i_k,i_{k+1}}$.

We now focus on ($\leq$).

Fix ϵ. Take $n \geq 1$, $i_0 = i$, $i_{n-1} = j$, $i_k \in \mathcal{I}$ for $k = 1,\dots,n-2$. There are paths $\xi^0,\dots,\xi^{n-2}$ of lengths $T_0,\dots,T_{n-2}$ such that $\xi_0^k \in \mathcal{K}_{i_k}$, $\xi_{T_k-1}^k \in \mathcal{K}_{i_{k+1}}$, $\mathcal{A}_{T_k}(\xi^k) \leq \widetilde{B}_{i_k,i_{k+1}} + \epsilon/n$ for $k = 0,\dots,n-2$.

By Lemma E.25, for all $k = 0,\dots,n-2$, $\xi_{T_k-1}^k$ and ξ_0^{k+1} can be connected by path of cost at most ϵ/n. Therefore concatenating all these paths yield ζ of length T with $\zeta_0 \in \mathcal{K}_i$, $\zeta_{T-1} \in \mathcal{K}_j$ and $\mathcal{A}_T(\zeta) \leq \sum_{k=0}^{n-2} \widetilde{B}_{i_k,i_{k+1}} + 2\epsilon$. Since $\mathcal{A}_T(\zeta) \geq B_{i,j}$, we obtain the desired result. ∎

Notation 1. We will write, for non-decreasing $f : \mathbb{R} \to \mathbb{R}$,

$$a \asymp f(b \pm c) \iff f(b-c) \leq a \leq f(b+c)\,. \tag{F.93}$$

Proposition F.1. *For any $\epsilon > 0$ and any small enough neighborhoods $\mathcal{U}_i$ of $\mathcal{K}_i$, $i = 1,\dots,K$, there exists $T \geq 0$, $\eta_0 > 0$ such that for all i, j, $x \in \mathcal{U}_i$, $0 < \eta < \eta_0$,*

$$\mathbb{Q}_{\mathcal{U}}^T(x, \mathcal{U}_j) \asymp \exp\left(-\frac{B_{i,j}}{\eta} \pm \frac{\epsilon}{\eta}\right). \tag{F.94}$$

Proof. Let us first start with ($\geq$).

Let T satisfy the conditions of Lemma F.12 and define $T' := (K-1)T$. For any $i, j \in \mathcal{I}$, by Lemma F.13, there exist $i_0 = i$, $i_{K-1} = j$, $i_k \in \mathcal{I}$ for $k = 1,\dots,K-2$ such that

$$B_{i,j} = \sum_{k=0}^{K-2} \widetilde{B}_{i_k,i_{k+1}}\,. \tag{F.95}$$

Therefore, we have that the probability of reaching $\mathcal{U}_j$ from $x \in \mathcal{U}_i$ in $T' = (K-1)T$ steps is greater than the probability of reaching sequentially the $\mathcal{U}_{i_{k+1}}$ from any point $x \in \mathcal{U}_{i_k}$, i.e.,

$$\begin{aligned}\inf_{x\in\mathcal{U}_i} \mathbb{Q}_{\mathcal{U}}^{T'}(x, \mathcal{U}_j) &\geq \prod_{k=0}^{K-2} \inf_{x\in\mathcal{U}_{i_k}} \mathbb{Q}_{\mathcal{U}}^{T}(x, \mathcal{U}_{i_{k+1}}) \\ &\geq \prod_{k=0}^{K-2} \exp\left(-\frac{\widetilde{B}_{i_k,i_{k+1}}}{\eta} - \frac{\epsilon}{\eta}\right) = \exp\left(-\frac{B_{i,j}}{\eta} - \frac{(K-1)\epsilon}{\eta}\right)\end{aligned} \tag{F.96}$$

where we used Lemma F.12 to get the second inequality.

Now for the reverse inequality ($\leq$). Take $i, j \in \mathcal{I}$ and denote $i_0 = i$, $i_{T'-1} = j$. By Lemma F.11, we have that

$$\begin{aligned}\sup_{x\in\mathcal{U}_i} \mathbb{Q}_{\mathcal{U}}^{T'}(x, \mathcal{U}_j) &\leq \sum_{i_1,\dots,i_{K-2}\in\mathcal{I}} \prod_{k=0}^{K-2} \sup_{x\in\mathcal{U}_{i_k}} \mathbb{Q}_{\mathcal{U}}^{T}(x, \mathcal{U}_{i_{k+1}}) \\ &\leq \sum_{i_1,\dots,i_{K-2}\in\mathcal{I}} \prod_{k=0}^{K-2} \exp\left(-\frac{\widetilde{B}_{i_k,i_{k+1}}}{\eta} + \frac{\epsilon}{\eta}\right) \\ &\leq \sum_{i_1,\dots,i_{K-2}\in\mathcal{I}} \exp\left(-\frac{\sum_{k=0}^{T'} \widetilde{B}_{i_k,i_{k+1}}}{\eta} + \frac{(K-1)\epsilon}{\eta}\right).\end{aligned} \tag{F.97}$$

Using Lemma F.13, we obtain that

$$\sup_{x\in\mathcal{U}_i} \mathbb{Q}_{\mathcal{U}}^{T'}(x,\mathcal{U}_j) \leq \sum_{i_1,\dots,i_{K-2}\in\mathcal{I}} \exp\left(-\frac{B_{i,j}}{\eta}+\frac{(K-1)\epsilon}{\eta}\right)$$
$$= K^{K-2}\exp\left(-\frac{B_{i,j}}{\eta}+\frac{(K-1)\epsilon}{\eta}\right) \tag{F.98}$$

which concludes the proof. ∎

G Long-run distribution

G.1 Convergence and stability 69
G.2 First estimates of the invariant measures 74
G.3 Characterization of the long-run distribution 79

G.1. Convergence and stability. This section is devoted to the study of the properties of stable and unstable sets and their influence on the long run behavior of the process. Compared to Azizian et al. [9], the main difficulty is that the process does not admit a Lyapunov function anymore. The principal results of this section thus require much more intricate proofs, see in particular the proofs of Lemma G.3 and Lemma G.10. We note that only the proofs of Lemmas G.4–G.6 are very minor adaptations of the proofs of Azizian et al. [9], but they are reproduced here for the sake of completeness and readability.

Let us begin by showing the omega-limit sets of the flow are each contained in one of the $\mathcal{K}_i$.

Lemma G.1. *For every $x \in \mathcal{V}$, if $y, z \in \omega(x)$, then $y \leftrightsquigarrow z$*

Proof. Fix $x \in \mathcal{V}$ and $y, z \in \omega(x)$. We first prove that $y \rightsquigarrow z$. Fix $\varepsilon > 0$ and $T > 0$. By the continuity of the flow, there exists $\delta > 0$ such that

$$\|x' - y\| < \delta \quad \Longrightarrow \quad \|\Phi_T(x') - \Phi_T(y)\|^2 < \frac{\varepsilon}{2}\,. \tag{G.1}$$

Since $y \in \omega(x)$, there exists $s \geq 0$ such that

$$\|\Phi_s(x) - y\| < \delta\,. \tag{G.2}$$

It follows from the previous display that

$$\|\Phi_T(\Phi_s(x)) - \Phi_T(y)\|^2 < \frac{\varepsilon}{2}\,. \tag{G.3}$$

Using the semiflow property,

$$\Phi_T(\Phi_s(x)) = \Phi_{s+T}(x)\,. \tag{G.4}$$

Next, since $z \in \omega(x)$, there exists $t \geq s + 2T$ such that

$$\|\Phi_t(x) - z\|^2 < \frac{\varepsilon}{2}\,. \tag{G.5}$$

We now define the chain

$$x_0 = y, \qquad x_1 = \Phi_{s+T}(x), \qquad x_2 = z, \tag{G.6}$$

with times

$$t_1 = T, \qquad t_2 = t - (s + T)\,. \tag{G.7}$$

Since $t \geq s + 2T$, we have $t_1, t_2 \geq T$. Moreover,

$$\begin{aligned}\sum_{i=1}^{2} \|\Phi_{t_i}(x_{i-1}) - x_i\|^2 &= \|\Phi_T(y) - \Phi_{s+T}(x)\|^2 + \|\Phi_{t-(s+T)}(\Phi_{s+T}(x)) - z\|^2 \\ &= \|\Phi_T(y) - \Phi_{s+T}(x)\|^2 + \|\Phi_t(x) - z\|^2 \\ &< \varepsilon\,.\end{aligned} \tag{G.8}$$

Thus $y \rightsquigarrow z$. By symmetry, $z \rightsquigarrow y$, and hence $y \leftrightsquigarrow z$. ∎

Theorem E.1 provides the following corollary.

Corollary G.1. *For every $x \in \mathcal{V}$, there is $i \in \mathcal{I}$ such that $\omega(x) \subset \mathcal{K}_i$.*

Proof. This is an immediate consequence of Theorem E.1 and Lemma G.1. ∎

Now, since the vector field is dissipative, the flow does not escape at infinity and therefore it must converge to one of the $\mathcal{K}_i$.

Lemma G.2. *For every $x \in \mathcal{V}$, there is $i \in \mathcal{I}$ such that $\Phi_t(x)$ converges to $\mathcal{K}_i$ as $t \to +\infty$ i.e.,*

$$\lim_{t \to +\infty} d(\Phi_t(x), \mathcal{K}_i) = 0\,. \tag{G.9}$$

Proof. Fix $x \in \mathcal{V}$. By Corollary G.1, there is $i \in \mathcal{I}$ such that $\omega(x) \subset \mathcal{K}_i$. Now, by Lemma E.5, there exists $\mathcal{K}$ a positively invariant compact set which contains x. As a consequence, for every $t \geq 0$, $\Phi_t(x) \in \mathcal{K}$. By contradiction, if $\Phi_t(x)$ did not converge to $\mathcal{K}_i$, there would be a subsequence that would remain outside of a neighborhood of $\mathcal{K}_i$. Since it would remain in $\mathcal{K}$, it would have an accumulation point that would be outside of $\mathcal{K}_i$, which is a contradiction. ∎

We now define the notions of stable component and attractors, and show that they coincide in our context. We then show that non-stable components are unstable in the sense of Freidlin & Wentzell [36, Chap. 6,§4] and that they have negligible weight in the invariant measure.

Definition G.1 (Stable component)**.** For any $i \in \mathcal{I}$, we say that $\mathcal{K}_i$ is *Lyapunov stable* or *stable* if for every $\delta > 0$, there exists $\delta' > 0$ such that,

$$\Phi_t(\mathcal{U}_{\delta'}(\mathcal{K}_i)) \subset \mathcal{U}_\delta(\mathcal{K}_i) \quad \text{for all } t \geq 0\,. \tag{G.10}$$

$\mathcal{K}_i$ is said to be *unstable* if it is not stable.

Lemma G.3. *For any $i \in \mathcal{I}$, $\mathcal{K}_i$ is stable if and only if there exists $\delta > 0$ such that*

$$\Phi_t(x) \to \mathcal{K}_i \quad \text{as } t \to +\infty \text{ for all } x \in \mathcal{U}_\delta(\mathcal{K}_i)\,. \tag{G.11}$$

Proof. ($\Rightarrow$) Assume that $\mathcal{K}_i$ is stable. Since they are only a finite number of equivalence classes and they are compact (Lemma E.10), there exists $\delta > 0$ such that $\mathcal{U}_\delta(\mathcal{K}_i)$ is disjoint from the other $\mathcal{K}_j$ for $j \neq i$. By Definition G.1, there exists $\delta' > 0$ such that $\Phi_t(\mathcal{U}_{\delta'}(\mathcal{K}_i)) \subset \mathcal{U}_\delta(\mathcal{K}_i)$ for all $t \geq 0$ and, by Lemma G.2, for any $x \in \mathcal{U}_{\delta'}(\mathcal{K}_i)$, $\Phi_t(x)$ must converges to some $\mathcal{K}_j$ as $t \to +\infty$. Since $\mathcal{U}_\delta(\mathcal{K}_i)$ is disjoint from the other $\mathcal{K}_j$, $\Phi_t(x)$ must converge to $\mathcal{K}_i$.

($\Leftarrow$) Assume that there exists $\delta > 0$ such that, for any $x \in \mathcal{U}_\delta(\mathcal{K}_i)$, $\Phi_t(x)$ converges to $\mathcal{K}_i$ as $t \to +\infty$. We show that $\mathcal{K}_i$ is stable by contradiction. Therefore, assume that there exists $\delta > 0$ and a sequence $x_k \in \mathcal{U}_{1/k}(\mathcal{K}_i)$ such that $\Phi_t(x_k)$ exits $\mathcal{U}_\delta(\mathcal{K}_i)$ for some $t \geq 0$ i.e., there exists $t_k \geq 0$ such that $y_k := \Phi_{t_k}(x_k) \notin \mathcal{U}_\delta(\mathcal{K}_i)$. But, by assumption, for k large enough, $\Phi_t(x_k)$ converges to $\mathcal{K}_i$ as $t \to +\infty$ i.e., $\omega(x_k) \subset \mathcal{K}_i$. Take $z_k \in \omega(x_k) \subset \mathcal{K}_i$. By definition, there exists s_k such that

$$\|\Phi_{s_k}(y_k) - z_k\| < 1/k\,. \tag{G.12}$$

Moreover, by Lemma E.5, since all $(x_k)_k$ are in a compact set, there exists a compact set containing all the trajectories of the flow starting from x_k for $k \geq 1$. Hence, there exists a compact set containing the three sequences $(x_k)_k$, $(y_k)_k$ and $(z_k)_k$. Without loss of generality, we can thus assume they converge to, respectively, $x_\infty, y_\infty, z_\infty$ with x_∞ and z_∞ in $\mathcal{K}_i$ and y_∞ outside of $\mathcal{U}_\delta(\mathcal{K}_i)$. We now show that $x_\infty \rightsquigarrow y_\infty \rightsquigarrow z_\infty$.

- $x_\infty \rightsquigarrow y_\infty$: there are two cases.
  - If t_k is unbounded: this means that it has a subsequence that goes to $+\infty$. Without loss of generality, assume that t_k itself goes to $+\infty$. Fix $T > 0$, $\varepsilon > 0$. For k large enough, we thus have $t_k \geq T$ and
  $$\|x_\infty - x_k\|^2 + \|\Phi_{t_k}(x_k) - y_\infty\|^2 < \varepsilon\,. \tag{G.13}$$
  Hence, Lemma E.6 implies that $x_\infty \rightsquigarrow y_\infty$.
  - If t_k is bounded: then, by Assumption 3(*a*), there exists $C \geq 1$ such that $\Phi_{t_k}(\cdot)$ are uniformly C-Lipschitz for $k = 0, 1, \ldots$. Fix $T > 0$, $\varepsilon > 0$. By Theorem E.1, $\mathcal{K}_i$ is an equivalence class for $\rightsquigarrow$ and therefore $x_\infty \rightsquigarrow x_\infty$: there exists $x'_0 = x_\infty, x'_1, \ldots, x'_\ell = x_\infty$ and $\tau'_1, \ldots, \tau_k \geq T$ such that
  $$\sum_{i=1}^{\ell} \|\Phi_{\tau'_i}(x'_{i-1}) - x'_i\|^2 < \varepsilon / C^2\,. \tag{G.14}$$
  For k large enough, we have
  $$\|\Phi_{t_k}(x_k) - y_\infty\|^2 = \|y_k - y_\infty\|^2 < \varepsilon/C^2 \quad \text{and} \quad \|x_k - x_\infty\|^2 < \varepsilon/C^2\,. \tag{G.15}$$

Therefore, by the triangle inequality and the C-Lipschitzness of $\Phi_{t_k}(\cdot)$, we have

$$\|\Phi_{\tau_\ell + t_k}(x'_{\ell-1}) - y_\infty\| \tag{G.16}$$
$$\leq \|\Phi_{\tau_\ell + t_k}(x'_{\ell-1}) - \Phi_{t_k}(x_\infty)\| + \|\Phi_{t_k}(x_\infty) - \Phi_{t_k}(x_k)\| + \|\Phi_{t_k}(x_k) - y_\infty\| \tag{G.17}$$
$$\leq C\|\Phi_{\tau_\ell}(x'_{\ell-1}) - x_\infty\| + C\|x_k - x_\infty\| + \|\Phi_{t_k}(x_k) - y_\infty\| \tag{G.18}$$
$$< 3\sqrt{\varepsilon}\,, \tag{G.19}$$

by combining Eq. (G.14) and Eq. (G.15). Thus considering the chain

$$x_0 = x_\infty, x_1, \ldots, x_{\ell-1}, y_\infty \quad \text{with } \tau_1, \ldots, \tau_{\ell-1}, \tau_\ell + t_k\,, \tag{G.20}$$

yields $x_\infty \rightsquigarrow y_\infty$.

- $y_\infty \rightsquigarrow z_\infty$: this is shown by the same argument as above, using that now z_∞ is in $\mathcal{K}_i$.

We thus have shown that $x_\infty \rightsquigarrow y_\infty \rightsquigarrow z_\infty$ with x_∞ and z_∞ in $\mathcal{K}_i$. By Theorem E.1, this means that y_∞ is in $\mathcal{K}_i$, which is a contradiction since y_∞ is outside of $\mathcal{U}_\delta(\mathcal{K}_i)$. ∎

Remark G.1. With this lemma in hand, we could now show that stability implies the following stronger condition: there exists $\delta > 0$ such that, $\Phi_t(x)$ converges to $\mathcal{K}_i$ as $t \to +\infty$ uniformly for all $x \in \mathcal{U}_\delta(\mathcal{K}_i)$.

The next lemma shows that if $\mathcal{K}_i$ is not stable, then it is unstable in the sense of Freidlin & Wentzell [36, Chap. 6,§4].

Lemma G.4. *If $\mathcal{K}_i$ is not stable, then there exists $j \in \mathcal{I}$ such that $B_{i,j} = 0$.*

Proof. If $\mathcal{K}_i$ is not stable, then, by Lemma G.3, for every $n \geq 1$, there exists $x_n \in \mathcal{U}_{1/n}(\mathcal{K}_i)$ such that $\Phi_t(x_n)$ does not converge to $\mathcal{K}_i$. By Lemma G.2, there exists $j_n \in \mathcal{I}$ such that $(\Phi_t(x_n))_t$ converges to $\mathcal{K}_{j_n}$. Since $\mathcal{I}$ is finite, there exists $j \in \mathcal{I}$ such that $j_n = j$ for infinitely many n. Replacing $(x_n)_{n\geq 1}$ by a subsequence, we can assume that $(\Phi_t(x_n))_t$ converges to $\mathcal{K}_j$ for all $n \geq 1$.

We now show that $B_{i,j} = 0$.

Fix $\epsilon > 0$ and consider $\mathcal{W}_\epsilon(\mathcal{K}_i), \mathcal{W}_\epsilon(\mathcal{K}_j)$ which are neighborhoods of $\mathcal{K}_i, \mathcal{K}_j$ by Corollary E.4. Since $(x_n)_{n\geq 1}$ converges to $\mathcal{K}_i$, there exists n such that $x_n \in \mathcal{W}_\epsilon(\mathcal{K}_i)$. In turn, since $\Phi_t(x_n)$ converges to $\mathcal{K}_j$, there exists $T > 0$, which we can choose integer, such that, $\Phi_T(x_n) \in \mathcal{W}_\epsilon(\mathcal{K}_j)$. Therefore, there exists $x \in \mathcal{K}_i$ and $y \in \mathcal{K}_j$ such that $\rho(x, x_n) < \epsilon$ and $\rho(\Phi_T(x_n), y) < \epsilon$. Consider the discrete path $\xi \in \mathcal{V}^{T+3}$ defined by $\xi_0 = x$, $\xi_n = \Phi_{n-1}(x)$ for all $1 \leq n \leq T+1$ and $\xi_{T+2} = y$. By Lemma E.2, we have that

$$\mathcal{A}_{T+3}(\xi) = \rho(x, x_n) + \rho(\Phi_T(x), y) < 2\epsilon\,, \tag{G.21}$$

and therefore, since $\xi_0 \in \mathcal{K}_i$ and $\xi_{T+2} \in \mathcal{K}_j$, $B_{i,j} < 2\epsilon$. ∎

The following lemma is now a straightforward adaptation of Freidlin & Wentzell [36, Chap. 6, Lemma 4.2].

Lemma G.5. *If $\mathcal{K}_i$ is not stable, then there exists $j \in \mathcal{I}$ such that $B_{i,j} = 0$ and such that, for any $k \in \mathcal{I}$, $B_{j,k} > 0$.*

Proof. For the sake of contradiction, assume that such a j does not exist. Then we build an infinite sequence $j_0 = i, j_1, \ldots$ such that $B_{j_n, j_{n+1}} = 0$ for all $n \geq 0$. By definition of equivalence classes, any j cannot appear twice in this sequence. But since $\mathcal{I}$ is finite, this is a contradiction. ∎

We now reach our final result on unstable equivalence classes: they have negligible weight in the invariant measure.

Lemma G.6. *If $\mathcal{K}_i$ is not stable, or, equivalently, non-minimizing, then there exists $j \in \mathcal{I}$ such that $\mathcal{K}_j$ is stable, $B_{i,j} = 0$ and,*

$$E_j < E_i \tag{G.22}$$

Proof. By Lemma G.5, there exists $j \in \mathcal{I}$ such that $B_{i,j} = 0$ and such that, for any $k \in \mathcal{I}$, $B_{j,k} > 0$. Lemma G.4 implies in particular that $\mathcal{K}_j$ must be stable

It remains to show that $E_j < E_i$. Take $T \in \mathcal{T}_i$ such that

$$E_i = \sum_{(k\to l)\in T} B_{k,l}\,. \tag{G.23}$$

j has an outgoing edge in T and denote by j' the other end of that edge. Now, consider the tree $T' \in \mathcal{T}_j$ obtained from T by removing the outgoing edge from j to i' and adding an edge from i to j. Then, by definition of E_j,

$$E_j \le \sum_{(k\to l)\in T'} B_{k,l} = \sum_{(k\to l)\in T} B_{k,l} - B_{j,j'} + B_{i,j}\,. \tag{G.24}$$

But, by definition of j, $B_{i,j} = 0$ and $B_{j,j'} > 0$. Therefore, $E_j < E_i$. ∎

We now show the second part of our main result: the invariant measure concentrates on the ground states, which are stable by Lemma G.6.

Before obtaining the main results of this section, we require the following technical lemma. It is a continuous-time reformulation of the first item of Lemma F.7.

We will require the following technical lemma.

Lemma G.7. *Fix $T > 0$. Define, for $\gamma \in \mathcal{C}_T(\mathcal{V})$,*

$$\mathcal{R}_T[\gamma] := \begin{cases} \frac{1}{2}\int_0^T \|\dot\gamma(t) - F(\gamma(t))\|^2\,dt & \text{if } \gamma \text{ is absolutely continuous} \\ +\infty & \text{otherwise}\,. \end{cases} \tag{G.25}$$

For any $\mathcal{K} \subset \mathcal{V}$ compact, $s \ge 0$, the set

$$\{\gamma \in \mathcal{C}_T(\mathcal{V}) : \gamma(0) \in \mathcal{K}, \mathcal{R}_T[\gamma] \le s\} \tag{G.26}$$

is compact and $\mathcal{R}_T[\cdot]$ is l.s.c. on $\mathcal{C}_T(\mathcal{V})$.

Proof. Let us first check that $\mathcal{R}_T[\cdot]$ is l.s.c. on $\mathcal{C}_T(\mathcal{V})$ by applying Ioffe et al. [59, §9.1.4, Thm. 3]: $(t,x,w) \mapsto \|w - F(x)\|^2$ is a normal integrand since $(x,w) \mapsto \|w - F(x)\|^2$ is continuous, quasiregular since $w \mapsto \|w - F(x)\|^2$ is convex for any $x \in \mathcal{V}$ and satisfies the growth condition because it is non-negative. Hence, $\mathcal{R}_T[\cdot]$ is l.s.c. on $\mathcal{C}_T(\mathcal{V})$.

The compactness of $\Gamma_{[0,T]}(\mathcal{K}; s)$ follows from the idea of proof as Freidlin & Wentzell [36, Chap. 3, Lem. 2.1; Chap. 4, Thm 1.1] but with the added difficulty that the vector field is not bounded. Compared to Azizian et al. [9], the action functional in this lemma does not imply a bound on the derivative of the curve.

Take $\gamma \in \mathcal{C}_T(\mathcal{V})$ such that $\gamma(0) \in \mathcal{K}$ and $\mathcal{R}_T[\gamma] \le s < +\infty$. We first show that γ is bounded. Fix $t \in [0,T]$. We have, by definition of absolutely continuous functions and by the triangle inequality,

$$\|\gamma(t)\| = \left\|\gamma(0) + \int_0^t \dot\gamma(s)\mathrm{s}\right\| \tag{G.27}$$

$$\le \|\gamma(0)\| + \int_0^t \|\dot\gamma(s) - F(\gamma(s))\|\,ds + \int_0^t \|F(\gamma(s))\|\,ds\,. \tag{G.28}$$

By Eq. (D.4), the last term is bounded as

$$\int_0^t \|F(\gamma(s))\|\,ds \le M\left(T + \int_0^t \|\gamma(s)\|\,ds\right). \tag{G.29}$$

By Cauchy-Schwarz inequality and the definition of $\mathcal{R}_T[\gamma]$, the second term in Eq. (G.28) is bounded as

$$\int_0^t \|\dot\gamma(s) - F(\gamma(s))\|\,ds \le \sqrt{T}\sqrt{\int_0^T \|\dot\gamma(s) - F(\gamma(s))\|^2\,ds} \tag{G.30}$$

$$= \sqrt{2T\mathcal{R}_T[\gamma]} \tag{G.31}$$

$$\leq \sqrt{2Ts}\,. \tag{G.32}$$

Plugging Eq. (G.29) and Eq. (G.32) into Eq. (G.28), we have

$$\|\gamma(t)\| \leq \|\gamma(0)\| + \sqrt{2Ts} + MT + M\int_0^t \|\gamma(s)\|\,ds\,. \tag{G.33}$$

We can now apply Grönwall's lemma (since γ is continuous) to obtain that, for any $t \in [0,T]$,

$$\|\gamma(t)\| \leq e^{MT}\left(\sup_{x\in\mathcal{K}}\|x\| + \sqrt{2Ts} + MT\right). \tag{G.34}$$

Hence, γ is bounded uniformly in $\gamma \in \mathcal{C}_T(\mathcal{V})$ such that $\gamma(0) \in \mathcal{K}$ and $\mathcal{R}_T[\gamma] \leq s$. We now show that this family of curves is equicontinuous. Combining Eq. (D.4) and Eq. (G.34), we have that there exists $C > 0$ such that, for any such curve γ and any $t \in [0,T]$,

$$\|F(\gamma(t))\| \leq C\,. \tag{G.35}$$

For the equicontinuity, take $0 \leq t < s \leq T$. We have, by the triangle inequality,

$$\|\gamma(s) - \gamma(t)\| = \left\|\int_t^s \dot\gamma(\tau)\,d\tau\right\| \tag{G.36}$$

$$\leq \int_t^s \|\dot\gamma(\tau) - F(\gamma(\tau))\|\,d\tau + \int_t^s \|F(\gamma(\tau))\|\,d\tau\,. \tag{G.37}$$

By Cauchy-Schwarz inequality, the definition of $\mathcal{R}_T[\gamma]$ and, Eq. (G.35), we have

$$\|\gamma(s) - \gamma(t)\| \leq \sqrt{s-t}\sqrt{\int_0^T \|\dot\gamma(\tau) - F(\gamma(\tau))\|^2\,d\tau} + C(s-t) \tag{G.38}$$

$$\leq \sqrt{2s}\sqrt{s-t} + C(s-t)\,. \tag{G.39}$$

Hence, the family of curves γ such that $\gamma(0) \in \mathcal{K}$ and $\mathcal{R}_T[\gamma] \leq s$ is equicontinuous and uniformly bounded. Moreover, it is closed by l.s.c. of $\mathcal{R}_T[\cdot]$, so that, by the Arzelà-Ascoli theorem, it is compact. ∎

We will require the following technical lemma, which is a consequence Kifer [63, Lem. 5.3].

Lemma G.8. *Consider $\mathcal{K} \subset \mathcal{V}$ a compact set which does not intersect any of the $\mathcal{K}_i$, $i \in \mathcal{I}$. Then, there exists $a, b > 0$ such that, for any $T > 0$, $\gamma \in \mathcal{C}_T(\mathcal{K})$,*

$$\mathcal{R}_T[\gamma] \geq aT - b\,. \tag{G.40}$$

Proof. The proof follows Kifer [63, Lem. 5.3] by setting

$$\rho(x,x') \coloneqq \inf\{\mathcal{R}_T[\gamma] : T > 0, \gamma \in \mathcal{C}_T(\mathcal{V}), \gamma(0) = x, \gamma(T) = x'\}\,. \tag{G.41}$$

and the fact that, for a curve $\gamma \in \mathcal{C}_T(\mathcal{V})$, we have

$$\mathcal{R}_T[\gamma] \geq \sum_{n=0}^{\lfloor T\rfloor - 1} \rho(\gamma(n), \gamma(n+1))\,. \tag{G.42}$$

∎

We now prove the key technical result:

Lemma G.9. *If $\mathcal{K}_i$ is stable, then there exists $\delta > 0$ such that*

$$\inf\{\mathcal{R}_T[\gamma] : T > 0, \gamma \in \mathcal{C}_T(\mathcal{V}), \gamma(0) \in \mathcal{K}_i, \gamma(T) \notin \mathcal{U}_\delta(\mathcal{K}_i)\} \geq 0\,. \tag{G.43}$$

Proof. For this, we prove the contrapositive. Assume that, for any $\delta > 0$, there exists $\delta' \leq \delta$ such that

$$\inf\{\mathcal{R}_T[\gamma] : T \geq 0, \gamma \in \mathcal{C}_T(\mathcal{V}), \gamma(0) \in \mathcal{K}_i, \gamma(T) \notin \mathcal{U}_\delta(\mathcal{K}_i)\} = 0\,. \tag{G.44}$$

. We show that $\mathcal{K}_i$ is not stable. To simplify notations, in this proof, we write $\mathcal{U}_\delta$ for $\mathcal{U}_\delta(\mathcal{K}_i)$.

By assumption, there is $\delta_n \to 0$ such that

$$\inf\left\{\mathcal{R}_T[\gamma] : T \geq 0, \gamma \in \mathcal{C}_T(\mathcal{V}), \gamma(0) \in \mathcal{K}_i, \gamma(T) \notin \mathcal{U}_{\delta_n}\right\} \tag{G.45}$$

for all $n \geq 1$. Take $\delta > 0$ such that $\mathrm{cl}\,\mathcal{U}_{2\delta}$ is disjoint from the other $\mathcal{K}_j$ for $j \neq i$. We will show that, for any $\delta' < \delta$, $\Phi_t(\mathcal{U}_{\delta'})$ is not included in $\mathcal{U}_\delta$ for some $t \geq 0$, which will imply that $\mathcal{K}_i$ is not stable (Definition G.1). Fix $\delta' < \delta$ and, without loss of generality, assume that $\delta_n < \delta'/2$ for all $n \geq 1$. Define $\mathcal{K} := (\mathrm{cl}\,\mathcal{U}_{2\delta}) \setminus \mathcal{U}_{\delta'/4}$. Note, by definition of δ, $\mathcal{K}$ does not intersect any of the $\mathcal{K}_j$ for $j \neq i$. Eq. (G.45) implies that, for any $n \geq 1$, there exists $\gamma_n(\cdot) \in \mathcal{C}_{T_n}(\mathcal{K})$ for some $T_n \geq 0$ such that

$$\gamma_n(0) \in \mathrm{cl}\,\mathcal{U}_{\delta'/2} \setminus \mathcal{U}_{\delta'/4}\,, \quad \gamma_n(T_n) \in \mathrm{cl}\,\mathcal{U}_{2\delta} \setminus \mathcal{U}_\delta\,, \quad \text{and} \quad \mathcal{R}_{T_n}[\gamma_n(\cdot)] \leq 1/n\,. \tag{G.46}$$

By Lemma G.8 applied to $\mathcal{K}$, T_n is bounded by some constant $T > 0$ for all $n \geq 1$. Extend all the paths $\gamma_n(\cdot)$ to $[0,T]$ by setting $\gamma_n(t), t \in [T_n, T]$ to be a trajectory of the flow starting from $\gamma_n(T_n)$. Eq. (G.46) still holds, also with $\mathcal{S}_T[\gamma_n(\cdot)] \leq 1/n$ since the flow has a zero action cost. By Lemma G.7, $\gamma_n(\cdot)$,$n \geq 1$ is contained in a compact subset of $\mathcal{C}_T(\mathcal{V})$. Hence, there exists a subsequence of $\gamma_n(\cdot)$ that converges uniformly to some $\gamma_\infty(\cdot) \in \mathcal{C}_T(\mathcal{V})$. Without loss of generality, we can assume that $\gamma_n(\cdot)$ itself converges to $\gamma_\infty(\cdot)$. By l.s.c. of the action (Lemma G.7), $\mathcal{R}_T[\gamma_\infty(\cdot)] = 0$ and therefore $\gamma_\infty(\cdot)$ is a trajectory of the flow. Moreover, we thus also have that the family $\gamma_n(\cdot)$, $n \geq 1$ is equicontinuous (Arzelà-Ascoli theorem) and therefore

$$\gamma_\infty(0) = \lim_{n\to+\infty} \gamma_n(0) \in \mathrm{cl}\,\mathcal{U}_{\delta'/2} \setminus \mathcal{U}_{\delta'/4}\,, \quad \text{and} \quad \gamma_\infty(T) = \lim_{n\to+\infty} \gamma_n(T) \in \mathrm{cl}\,\mathcal{U}_{2\delta} \setminus \mathcal{U}_\delta\,. \tag{G.47}$$

Thus, $\gamma_\infty(\cdot)$ is a trajectory of the flow that starts in $\mathrm{cl}\,\mathcal{U}_{\delta'/2} \subset \mathcal{U}_{\delta'}$ and ends outside of $\mathcal{U}_\delta$: this shows that $\mathcal{K}_i$ is not stable. ∎

Lemma G.10. *For any $i \in \mathcal{I}$ such that $\mathcal{K}_i$ is stable, there exists $\delta > 0$ such that, for any $0 < \delta' \leq \delta$,*

$$B(\mathcal{K}_i, \mathcal{V} \setminus \mathcal{U}'_\delta(\mathcal{K}_i)) > 0\,. \tag{G.48}$$

Proof. Consider the set $\mathcal{K}\,\mathrm{cl}\,\mathcal{U}_{\delta'}(\mathcal{K}_i)$, which is compact. Hence, on $\mathcal{K}$,

$$C := \sup_{x\in\mathcal{K}} \sigma^2_\infty(\|x\|^2) < +\infty\,. \tag{G.49}$$

By definition, we have that

$$B(\mathcal{K}_i, \mathcal{V} \setminus \mathcal{U}'_\delta(\mathcal{K}_i)) = \inf\left\{\mathcal{S}_T[\gamma] : T \in \mathbb{N}, \gamma \in \mathcal{C}_T(\mathcal{V}), \gamma(0) \in \mathcal{K}_i, \gamma(T) \notin \mathcal{U}'_\delta(\mathcal{K}_i)\right\}. \tag{G.50}$$

Take any such T and γ. There exists $t_\gamma \in [0,T]$ such that $\gamma(t_\gamma) \in \mathrm{bd}\,\mathcal{U}'_\delta(\mathcal{K}_i)$ and $\gamma(t) \in \mathcal{K}$ for all $t \in [0, t_\gamma]$. We thus have that

$$\mathcal{S}_T[\gamma] \geq \mathcal{S}_{t_\gamma}[\gamma] \geq \frac{1}{2C} \int_0^{t_\gamma} \|\dot\gamma(t) - F(\gamma(t))\|^2\, dt\,, \tag{G.51}$$

by Lemma D.4. Therefore,

$$\mathcal{S}_T[\gamma] \geq \frac{1}{C} \inf\left\{\mathcal{R}_{t_\gamma}[\gamma] : t_\gamma \in [0,T], \gamma \in \mathcal{C}_{t_\gamma}(\mathcal{V}), \gamma(0) \in \mathcal{K}_i, \gamma(t_\gamma) \in \mathrm{bd}\,\mathcal{U}'_\delta(\mathcal{K}_i)\right\}. \tag{G.52}$$

Taking the infimum over T and γ and invoking Lemma G.9, we obtain that $B(\mathcal{K}_i, \mathcal{V}\setminus\mathcal{U}'_\delta(\mathcal{K}_i)) > 0$. ∎

G.2. First estimates of the invariant measures. The results of this section are exactly the same as the ones of Azizian et al. [9] but with a more general setting. Though the proofs are the same, we include them for completeness and to make the paper self-contained.

Let us now give our first estimates on the invariant measure.

Definition G.2 ([36, Chap. 6,§4])**.** Define, for every $i \in \mathcal{I}$,

$$E_i = E(\mathcal{K}_i) := \min_{T\in\mathcal{T}_i} \sum_{(j\to k)\in T} B_{j,k}\,, \tag{G.53}$$

where $\mathcal{T}_i$ denotes the set of trees rooted at i in the complete graph on $\mathcal{I}$.

Proposition G.1. *For any $\mu_{\mathcal{U}}$ invariant probability measure for $\mathbb{Q}_{\mathcal{U}}$, the induced chain on $\mathcal{U}$, in the setting of Proposition F.1, for any $i \in \mathcal{I}$,*

$$\mu_{\mathcal{U}}(\mathcal{U}_i) \asymp \exp\left(-\frac{E(\mathcal{K}_i) - \min_{j\in\mathcal{I}} E(\mathcal{K}_j)}{\eta} \pm \frac{\epsilon}{\eta}\right). \tag{G.54}$$

Proof. If $\mu_{\mathcal{U}}$ is an invariant measure of $\mathbb{Q}_{\mathcal{U}}$, then it is an invariant measure of $\mathbb{Q}_{\mathcal{U}}^T$ for any T given by Proposition F.1.

Freidlin & Wentzell [36, Chap. 6, Lem. 3.1-3.2] combined with Proposition F.1 then give

$$\begin{aligned}
&\exp(-(2\#\mathcal{I}+2)\epsilon/\eta)\frac{\exp(-E(\mathcal{K}_i)/\eta)}{\sum_{j\in\mathcal{I}}\exp(-E(\mathcal{K}_j)/\eta)} \\
&\quad\leq \mu_{\mathcal{U}}(\mathcal{U}_i) \\
&\quad\leq \exp((2\#\mathcal{I}+2)\epsilon/\eta)\frac{\exp(-E(\mathcal{K}_i)/\eta)}{\sum_{j\in\mathcal{I}}\exp(-E(\mathcal{K}_j)/\eta)}\,.
\end{aligned} \tag{G.55}$$

For η small enough, it holds that

$$\sum_{j\in\mathcal{I}}\exp(-E(\mathcal{K}_j)/\eta) \asymp \exp\left(-\min_{j\in\mathcal{I}}E(\mathcal{K}_j)/\eta \pm \epsilon/\eta\right), \tag{G.56}$$

which concludes the proof. ∎

We now state a result that links invariant measures of $(x_n^\eta)_n$ and $\mathbb{Q}_{\mathcal{U}}$. It is a consequence of Douc et al. [33, Thm. 3.6.5].

Lemma G.11. *There is $\eta_0 > 0$ such that, for $0 < \eta \leq \eta_0$, if $(x_n^\eta)_n$ has an invariant probability measure μ_∞, then, for any $\mathcal{U}$ measurable neighborhood of $\bigcup_{i\in\mathcal{I}}\mathcal{K}_i$, we have that $\mu_\infty(\mathcal{U}) > 0$ and $\mu_{\mathcal{U}}$, the restriction of $\mu_\infty/\mu_\infty(\mathcal{U})$ to $\mathcal{U}$ is an invariant measure for the induced chain on $\mathcal{U}$ and, for any measurable set $E \subset \mathcal{X}$,*

$$\frac{\mu_\infty(E)}{\mu_\infty(\mathcal{U})} = \int_{\mathcal{U}} d\mu_{\mathcal{U}}(x)\,\mathbb{E}_x\left[\sum_{n=0}^{\sigma_{\mathcal{U}}-1}\mathbb{1}\{x_n^\eta \in E\}\right]. \tag{G.57}$$

Proof. We invoke the first item of Lemma F.8: there exists $\mathcal{D} \subset \mathcal{V}$ a compact set, $\eta_0 > 0$, such that for any $\mathcal{D}' \subset \mathcal{V}$ compact set such that $\mathcal{D} \subset \mathcal{D}'$, there exists $\lambda_0 > 0$ such that,

$$\forall \eta \leq \eta_0,\, x \in \mathcal{D}', \quad \mathbb{E}_x\left[e^{\frac{\lambda_0\sigma_{\mathcal{D}}}{\eta}}\right] < +\infty\,. \tag{G.58}$$

Without loss of generality, at the potential expense of expanding $\mathcal{D}$, assume that $\mathcal{U} \subset \mathcal{D}$.

By applying Lemma F.10 to $\mathcal{U} \leftarrow \mathcal{U}$ and $\mathcal{D} \leftarrow \mathcal{D}$, we get that there is some $\eta_0, \lambda_0, > 0$ such that, for any $\eta \leq \eta_0$, $x \in \mathcal{D}$,

$$\mathbb{E}_x\left[e^{\frac{\lambda\sigma_{\mathcal{U}}}{\eta}}\right] < +\infty\,. \tag{G.59}$$

In particular, we have that $\mathbb{P}_x(\sigma_{\mathcal{U}} < \infty) = 1$ for any $x \in \mathcal{D}$, and *a fortiori* for any $x \in \mathcal{U}$.

Let us now show that, for any $x \in \mathcal{V}$, $\mathbb{P}_x(\sigma_{\mathcal{U}} < \infty) = 1$. Fix $x \in \mathcal{V}$. By choosing $\mathcal{D}'$ large enough to contain both $\mathcal{D}$ and x, Eq. (G.58) implies that, with $x_0 = x$, $\sigma_{\mathcal{D}} < \infty$ almost surely. Therefore, by the strong Markov property, it holds that,

$$\mathbb{P}_x(\sigma_{\mathcal{U}} < \infty) \geq \inf_{y\in\mathcal{D}}\mathbb{P}_y(\sigma_{\mathcal{U}} < \infty) = 1\,. \tag{G.60}$$

Therefore the assumptions of Douc et al. [33, Thm. 3.6.5] are satisfied as well as its first item which yields the result. ∎

We reach the next proposition, which is the first part of our main result. It is an adaptation of Freidlin & Wentzell [36, Thm. 4.1] to the discrete time setting.

Proposition G.2. *For any $\epsilon > 0$, for any $\mathcal{U}_1, \ldots, \mathcal{U}_K$ measurable neighborhoods of $\mathcal{K}_1, \ldots, \mathcal{K}_K$ small enough, there exists $\eta_0 > 0$ such that for any $0 < \eta < \eta_0$, μ_∞ invariant probability measure for $(x_n^\eta)_n$, for any $i \in \mathcal{I}$,*

$$\mu_\infty(\mathcal{U}_i) \asymp \exp\left(-\frac{E(\mathcal{K}_i) - \min_{j\in\mathcal{I}}E(\mathcal{K}_j)}{\eta} \pm \frac{\epsilon}{\eta}\right). \tag{G.61}$$

Proof. Let us first provide estimates for the unnormalized measure defined by, for any measurable set E,

$$\nu(E) \coloneqq \int_E d\mu_{\mathcal{U}}(x)\, \mathbb{E}_x\left[\sum_{n=0}^{\sigma_{\mathcal{U}}-1} \mathbb{1}\{x_n^\eta \in E\}\right]. \tag{G.62}$$

By definition of $\sigma_{\mathcal{U}}$, in the sequence of points $x_0^\eta, \dots, x_{\sigma_{\mathcal{U}}-1}^\eta$, only x_0^η can be in $\mathcal{U}$. Therefore, for any $i, j \in \mathcal{I}$, $x \in \mathcal{U}_j$,

$$\mathbb{E}_x\left[\sum_{n=0}^{\sigma_{\mathcal{U}}-1} \mathbb{1}\{x_n^\eta \in \mathcal{U}_i\}\right] = \begin{cases} 1 & \text{if } i = j\,, \\ 0 & \text{otherwise}\,. \end{cases} \tag{G.63}$$

Thus, we have that,

$$\begin{aligned} \nu(\mathcal{U}_i) &= \mu_{\mathcal{U}}(\mathcal{U}_i) \\ &\asymp \exp\left(-\frac{E(\mathcal{K}_i) - \min_{j\in\mathcal{I}} E(\mathcal{K}_j)}{\eta} \pm \frac{\epsilon}{\eta}\right), \end{aligned} \tag{G.64}$$

by Proposition G.1. It now remains to estimate the normalization constant $\nu(\mathcal{V})$. On the one hand, we have that

$$\begin{aligned} \nu(\mathcal{V}) &\geq \max_{i\in\mathcal{I}} \nu(\mathcal{U}_i) \\ &\geq \max_{i\in\mathcal{I}} \exp\left(-\frac{E(\mathcal{K}_i) - \min_{j\in\mathcal{I}} E(\mathcal{K}_j)}{\eta} - \frac{\epsilon}{\eta}\right) = \exp\left(-\frac{\epsilon}{\eta}\right). \end{aligned} \tag{G.65}$$

On the other hand, by Lemma F.10 applied with $\mathcal{U} \leftarrow \mathcal{U}$ and $\mathcal{D} \leftarrow \operatorname{cl}\mathcal{U}$ (choosing $\mathcal{U}$ small enough so that is bounded), the quantity

$$c \coloneqq \sup\{\mathbb{E}_x[\sigma_{\mathcal{U}}] : x \in \mathcal{U},\, 0 < \eta \leq \eta_0\} \tag{G.66}$$

is finite. Therefore, we have that

$$\begin{aligned} \nu(\mathcal{V}) &= \int_{\mathcal{V}} d\mu_{\mathcal{U}}(x)\, \mathbb{E}_x[\sigma_{\mathcal{U}}] \\ &\leq c \int_{\mathcal{V}} d\mu_{\mathcal{U}}(x) = c\,, \end{aligned} \tag{G.67}$$

which, along with choosing η_0 small enough so that $c \leq \exp(\epsilon/\eta)$, concludes the proof. ∎

We will need the following lemma to prove the second part of our main result.

Lemma G.12. *For any $\epsilon > 0$, for any $\mathcal{U}_1, \dots, \mathcal{U}_K$ measurable neighborhoods of $\mathcal{K}_1, \dots, \mathcal{K}_K$ small enough, $\mathcal{D}$ measurable set, $\delta_{\mathcal{D}} > 0$, there exists $\eta_0 > 0$ such that, for any $0 < \eta < \eta_0$, for any $i \in \mathcal{I}$, $x \in \mathcal{U}_i$,*

$$\mathbb{P}_x\left(\sigma_{\mathcal{D}_{-\delta_{\mathcal{D}}}} < \sigma_{\mathcal{U}}\right) \leq \exp\left(-\frac{B(\mathcal{K}_i, \mathcal{D}) - \epsilon}{\eta}\right), \tag{G.68}$$

where

$$\mathcal{D}_{-\delta_{\mathcal{D}}} \coloneqq \{y \in \mathcal{D} : d(y, \mathcal{V} \setminus \mathcal{D}) \geq \delta_{\mathcal{D}}\}\,. \tag{G.69}$$

Proof. The requirement on the $\mathcal{U}_i$ are the same in the proof of Lemma F.11 but we restate them here for completeness.

Assume that, without loss of generality, ϵ is small enough so that Corollary E.4 with $r \leftarrow \epsilon$ can be applied to every $\mathcal{K}_i$, $i = 1, \dots, K$. Denote by $\mathcal{W}_i$, $i = 1, \dots, K$ the corresponding neighborhoods of $\mathcal{K}_i$. Since these $\mathcal{W}_i$'s are neighborhoods of the $\mathcal{K}_i$'s, there exists $\delta > 0$ such that $\mathcal{U}_{2\delta}(\mathcal{K}_i) \subset \mathcal{W}_i$ for all $i = 1, \dots, K$. Require then that $\mathcal{U}_i$ be contained in $\mathcal{U}_\delta(\mathcal{K}_i)$ so that $\mathcal{U}_\delta(\operatorname{cl}\mathcal{U}_i) \subset \mathcal{W}_i$ for all $i = 1, \dots, K$. Moreover, assume that $\delta > 0$ is small enough so that the neighborhoods $\mathcal{U}_\delta(\mathcal{K}_i)$, $i = 1, \dots, K$ are pairwise disjoint and that $\delta \leq \delta_{\mathcal{D}}$. Define $0 < \delta' \leq \delta$ such that $\mathcal{U}'_\delta(\mathcal{K}_i)$ is contained in $\mathcal{U}_i$ for all $i = 1, \dots, K$.

Fix $i \in \mathcal{I}$ and consider $\xi \in (\mathcal{V})^T$ such that $\xi_0 \in \mathcal{U}'_\delta(\operatorname{cl}\mathcal{U}_i)$, $\xi_{T-1} \in \mathcal{U}'_\delta(\mathcal{D}_{-\delta_{\mathcal{D}}})$. By construction, $\xi_0 \in \mathcal{U}_\delta(\mathcal{U}_i) \subset \mathcal{W}_i$. Therefore, there exists $x \in \mathcal{K}_i$ such that $\rho(x, \xi_0) < \epsilon$. Moreover, since $\delta' \leq \delta \leq \delta_{\mathcal{D}}$, $\mathcal{U}'_\delta(\mathcal{D}_{-\delta_{\mathcal{D}}}) \subset \mathcal{D}$ so that $\xi_{T-1} \in \mathcal{D}$.

Define $\zeta := (x, \xi_0, \dots, \xi_{T-1})$ which is a path from $\mathcal{K}_i$ to $\mathcal{D}$ so that

$$\mathcal{A}_T(\xi) \geq \mathcal{A}_{T+1}(\zeta) - \epsilon \geq B(\mathcal{K}_i, \mathcal{D}) - \epsilon \,. \tag{G.70}$$

We now follow the same outline as for the proof of Lemma F.11. Fix $x \in \mathcal{U}_i$. For any $T \geq 0$, we have that

$$\mathbb{P}_x\left(\sigma_{\mathcal{D}_{-\delta_\mathcal{D}}} < \sigma_\mathcal{U}\right) \leq \mathbb{P}_x\left(\sigma_{\mathcal{D}_{-\delta_\mathcal{D}}} < \sigma_\mathcal{U}\,, \sigma_\mathcal{U} < T\right) + \mathbb{P}_x(\sigma_\mathcal{U} \geq T)\,. \tag{G.71}$$

For some T large enough, Lemma F.10 yields

$$\mathbb{P}_x(\sigma_\mathcal{U} \geq T) \leq \exp\left(-\frac{B(\mathcal{K}_i, \mathcal{D}) - \epsilon}{\eta}\right). \tag{G.72}$$

We now bound $\mathbb{P}_x\left(\sigma_{\mathcal{D}_{-\delta_\mathcal{D}}} < \sigma_\mathcal{U}\,, \sigma_\mathcal{U} < T\right)$ for this choice of T. For this, it suffices to note that $\sigma_{\mathcal{D}_{-\delta_\mathcal{D}}} < \sigma_\mathcal{U}\,, \sigma_\mathcal{U} < T$ implies that $\mathrm{dist}_T\left(x^\eta, \Gamma_T^{\{x\}}(\xi - 2\epsilon)\right) > \frac{\delta'}{2}$. Then, applying Corollary D.2 yields

$$\mathbb{P}_x\left(\sigma_{\mathcal{D}_{-\delta_\mathcal{D}}} < \sigma_\mathcal{U}\,, \sigma_\mathcal{U} < T\right) \leq \exp\left(-\frac{B(\mathcal{K}_i, \mathcal{D}) - 3\epsilon}{\eta}\right), \tag{G.73}$$

which concludes the proof. ∎

We can use this lemma to upper-bound the $\mu_\infty(\mathcal{D})$.

Lemma G.13. *For any $\epsilon > 0$, any measurable set $\mathcal{D}$ with $\mathrm{bd}\,\mathcal{D}$ bounded, there exists $\eta_0 > 0$ such that for any $0 < \eta < \eta_0$, μ_∞ invariant probability measure for $(x_n^\eta)_n$, for any $i \in \mathcal{I}$,*

$$\mu_\infty(\mathcal{D}_{-\delta_\mathcal{D}}) \leq \exp\left(-\frac{\min_{i\in\mathcal{I}}\{E(\mathcal{K}_i) + B(\mathcal{K}_i, \mathcal{D})\} - \min_{i\in\mathcal{I}} E(\mathcal{K}_i)}{\eta} + \frac{\epsilon}{\eta}\right), \tag{G.74}$$

where

$$\mathcal{D}_{-\delta_\mathcal{D}} := \{y \in \mathcal{V} : d(y, \mathcal{V} \setminus \mathcal{D}) \geq \delta_\mathcal{D}\}\,. \tag{G.75}$$

Proof. Using $\mathcal{U}_1, \dots, \mathcal{U}_K$ measurable neighborhoods of $\mathcal{K}_1, \dots, \mathcal{K}_K$ small enough given by Lemma G.12, we provide an estimate for the weight of $\mathcal{D}$ for the unnormalized measure ν defined in the proof of Proposition G.2, i.e., for

$$\nu(\mathcal{D}_{-\delta_\mathcal{D}}) := \int_\mathcal{U} d\mu_\mathcal{U}(x)\, \mathbb{E}_x\left[\sum_{n=0}^{\sigma_\mathcal{U}-1} \mathbb{1}\{x_n^\eta \in \mathcal{D}_{-\delta_\mathcal{D}}\}\right], \tag{G.76}$$

and the result will follow from the estimate on the normalization constant $\nu(\mathcal{V})$ obtained in the proof of Proposition G.2.

Since $\mathrm{bd}\,\mathcal{D}$ is bounded, $\mathrm{bd}\,\mathcal{D}_{-\delta_\mathcal{D}}$ is also bounded. Therefore, by Lemma D.1, there exists $\mathcal{D}$ compact set such that, almost surely, if $x_{n-1}^\eta \notin \mathcal{D}_{-\delta_\mathcal{D}}$ and $x_n^\eta \in \mathcal{D}_{-\delta_\mathcal{D}}$, then $x_n^\eta \in \mathcal{D}$. Without loss of generality, we can assume that $\mathcal{D}$ contains $\mathrm{cl}\,\mathcal{U}$ (choosing $\mathcal{U}$ small enough so that it is bounded if necessary).

By Lemma F.10 applied with $\mathcal{U} \leftarrow \mathcal{U}$ and $\mathcal{D} \leftarrow \mathcal{D}$, the quantity

$$c := \sup\{\mathbb{E}_x[\sigma_\mathcal{U}] : x \in \mathcal{D},\ 0 < \eta \leq \eta_0\} \tag{G.77}$$

is finite.

Fix $i \in \mathcal{I}$ and $x \in \mathcal{U}_i$. If $\sigma_{\mathcal{D}_{-\delta_\mathcal{D}}} \geq \sigma_\mathcal{U}$, then $\sum_{n=0}^{\sigma_\mathcal{U}-1} \mathbb{1}\{x_n^\eta \in \mathcal{D}_{-\delta_\mathcal{D}}\}$ would be 0 so, we have that

$$\begin{aligned}
\mathbb{E}_x\left[\sum_{n=0}^{\sigma_\mathcal{U}-1} \mathbb{1}\{x_n^\eta \in \mathcal{D}_{-\delta_\mathcal{D}}\}\right] &= \mathbb{E}_x\left[\mathbb{1}\{\sigma_{\mathcal{D}_{-\delta_\mathcal{D}}} < \sigma_\mathcal{U}\} \sum_{n=0}^{\sigma_\mathcal{U}-1} \mathbb{1}\{x_n^\eta \in \mathcal{D}_{-\delta_\mathcal{D}}\}\right] \\
&\leq \mathbb{E}_x\left[\mathbb{1}\{\sigma_{\mathcal{D}_{-\delta_\mathcal{D}}} < \sigma_\mathcal{U}\}\sigma_\mathcal{U}\right] \\
&= \mathbb{E}_x\left[\mathbb{1}\{\sigma_{\mathcal{D}_{-\delta_\mathcal{D}}} < \sigma_\mathcal{U}\}\left(\sigma_{\mathcal{D}_{-\delta_\mathcal{D}}} + \mathbb{E}_{x^\eta_{\sigma_{\mathcal{D}_{-\delta_\mathcal{D}}}}}[\sigma_\mathcal{U}]\right)\right],
\end{aligned} \tag{G.78}$$

where we used the strong Markov property. Bounding $\sigma_{\mathcal{D}_{-\delta_\mathcal{D}}}$ by $\sigma_\mathcal{U}$ on the event $\{\sigma_{\mathcal{D}_{-\delta_\mathcal{D}}} < \sigma_\mathcal{U}\}$, we obtain that

$$\begin{aligned}\mathbb{E}_x\left[\sum_{n=0}^{\sigma_\mathcal{U}-1} \mathbb{1}\{x_n^\eta \in \mathcal{D}_{-\delta_\mathcal{D}}\}\right] &\le \mathbb{E}_x\left[\mathbb{1}\{\sigma_{\mathcal{D}_{-\delta_\mathcal{D}}} < \sigma_\mathcal{U}\}\left(\sigma_\mathcal{U} + \mathbb{E}_{x^\eta_{\sigma_{\mathcal{D}_{-\delta_\mathcal{D}}}}}[\sigma_\mathcal{U}]\right)\right] \\ &\le 2c\,\mathbb{P}_x\left(\sigma_{\mathcal{D}_{-\delta_\mathcal{D}}} < \sigma_\mathcal{U}\right) \\ &\le 2c\exp\left(-\frac{B(\mathcal{K}_i,\mathcal{D})-\epsilon}{\eta}\right), \end{aligned} \tag{G.79}$$

where we used that $x^\eta_{\sigma_{\mathcal{D}_{-\delta_\mathcal{D}}}} \in \mathcal{D}$ almost surely on the event $\{\sigma_{\mathcal{D}_{-\delta_\mathcal{D}}} < \sigma_\mathcal{U}\}$ and the definition of c for the second inequality and where we invoked Lemma G.12 for the last inequality.

Combining this bound and Proposition G.1, we obtain that

$$\begin{aligned}\nu(\mathcal{D}_{-\delta_\mathcal{D}}) &\le 2c\sum_{i\in\mathcal{I}}\mu_\mathcal{U}(\mathcal{U}_i)\exp\left(-\frac{B(\mathcal{K}_i,\mathcal{D})-\epsilon}{\eta}\right) \\ &\le 2c\,\#\mathcal{I}\exp\left(-\frac{\min_{i\in\mathcal{I}}\{E(\mathcal{K}_i)+B(\mathcal{K}_i,\mathcal{D})\}-\min_{i\in\mathcal{I}}E(\mathcal{K}_i)}{\eta}+\frac{2\epsilon}{\eta}\right), \end{aligned} \tag{G.80}$$

which concludes the proof. ∎

The next proposition shows that the invariant measure concentrates exponentially on states that are stable (and contain the ground states), leveraging crucial results from Appendix G.1.

Proposition G.3. *Consider $J \subset \mathcal{I}$ such that, for all $i \in J$, $\mathcal{K}_i$ is stable and, such that, J contains* $\arg\min_{i\in\mathcal{I}} E_i$. *Consider $\mathcal{U}_i$ small enough neighborhoods of $\mathcal{K}_i$ for $i \in J$. Then, there exists $c > 0$, $\eta_0 > 0$ such that, for any $\eta \le \eta_0$, for any μ_∞ invariant measure of $(x_n^\eta)_{n\ge0}$,*

$$\mu_\infty\left(\mathcal{V}\setminus\bigcup_{i\in J}\mathcal{U}_i\right) \le e^{-\frac{c}{\eta}}\,. \tag{G.81}$$

Proof. Take $\delta > 0$ small enough so that, for any $i \in J$, $\mathcal{U}_\delta(\mathcal{K}_i) \subset \mathcal{U}_i$ and

$$d(\mathcal{K}_i, \mathcal{V}\setminus\mathcal{U}_\delta(\mathcal{K}_i)) > 0\,. \tag{G.82}$$

This is possible by Lemma G.10.

Define

$$\mathcal{D} := \mathcal{V}\setminus\bigcup_{i\in J}\mathcal{U}_{\delta/2}(\mathcal{K}_i)\,. \tag{G.83}$$

With the notations of Lemma G.13, we show that

$$\mathcal{V}\setminus\bigcup_{i\in J}\mathcal{U}_i \subset \mathcal{D}_{-\delta/2}\,. \tag{G.84}$$

Indeed, take $x \in \mathcal{V}\setminus\cup_{i\in J}\mathcal{U}_i$. Then, for any $i \in J$, it holds that $d(x,\mathcal{K}_i) \ge \delta$ and so, we have

$$d(x, \mathcal{V}\setminus\mathcal{D}) = d\left(x, \mathcal{V}\setminus\bigcup_{i\in J}\mathcal{U}_{\delta/2}(\mathcal{K}_i)\right) \ge \frac{\delta}{2}\,. \tag{G.85}$$

Hence $x \in \mathcal{D}_{-\delta/2}$ and, it suffices to bound $\mu_\infty(\mathcal{D}_{-\delta/2})$ to show the result.

Moreover, for any $i \in J$,

$$B(\mathcal{K}_i,\mathcal{D}) \ge B(\mathcal{K}_i, \mathcal{V}\setminus\mathcal{U}_{\delta/2}(\mathcal{K}_i)) > 0\,, \tag{G.86}$$

so that the quantity

$$c := \min\left(\min_{i\notin J}E_i - \min_{i\in\mathcal{I}}E_i, \min_{i\in J}B(\mathcal{K}_i,\mathcal{D})\right) \tag{G.87}$$

is positive.

Apply Lemma G.13 with $\delta_{\mathcal{D}} \leftarrow \delta/2$, $\epsilon \leftarrow c/2$ to get that, for any $\eta \le \eta_0$,

$$\mu_\infty(\mathcal{D}_{-\delta/2}) \le \exp\left(-\frac{\min_{i\in\mathcal{I}}\{E(\mathcal{K}_i) + B(\mathcal{K}_i, \mathcal{D})\} - \min_{j\in\mathcal{I}} E(\mathcal{K}_j)}{\eta} + \frac{c}{2\eta}\right). \tag{G.88}$$

But, for any $i \in \mathcal{I}$, the exponent can be estimated as

$$\{E(\mathcal{K}_i) + B(\mathcal{K}_i, \mathcal{D})\} - \min_{j\in\mathcal{I}} E(\mathcal{K}_i) \ge \begin{cases} E(\mathcal{K}_i) - \min_{j\in\mathcal{I}} E(\mathcal{K}_j) & \text{if } i \notin J \\ B(\mathcal{K}_i, \mathcal{D}) & \text{if } i \in J\,, \end{cases} \tag{G.89}$$

which is always positive, even in the first case, since $\arg\min_{j\in\mathcal{I}} E(\mathcal{K}_j) \subset J$. Therefore, it holds that

$$\min_{i\in\mathcal{I}}\{E(\mathcal{K}_i) + B(\mathcal{K}_i, \mathcal{D})\} - \min_{j\in\mathcal{I}} E(\mathcal{K}_j) - \epsilon \ge c - \epsilon = \frac{c}{2}\,, \tag{G.90}$$

which concludes the proof. ∎

G.3. Characterization of the long-run distribution.

G.3.1. **The invariant measure μ_∞.** In this section, we restate the main results Theorems G.1–G.4 and provide their proofs. They are now mostly corollaries of the results of Appendices G.1 and G.2.

Theorem G.1. *For any $\epsilon > 0$, for any $\mathcal{U}_1, \dots, \mathcal{U}_K$ measurable neighborhoods of $\mathcal{K}_1, \dots, \mathcal{K}_K$ small enough, there exists $\eta_0 > 0$ such that, for any $0 < \eta < \eta_0$, for any μ_∞ invariant probability measure for $(x_n^\eta)_n$, and for any $i, j \in \mathcal{I}$, we have*

$$|\eta \log \mu_\infty(\mathcal{U}_i) + E(\mathcal{K}_i) - \min_{k\in\mathcal{I}} E(\mathcal{K}_k)| \le \epsilon \tag{G.91}$$

and

$$\left|\eta \log \frac{\mu_\infty(\mathcal{U}_i)}{\mu_\infty(\mathcal{U}_j)} + E(\mathcal{K}_i) - E(\mathcal{K}_j)\right| \le \epsilon. \tag{G.92}$$

More compactly, with notation as above, we have:

$$\mu_\infty(\mathcal{U}_i) \propto \exp\left(-\frac{E(\mathcal{K}_i) + \mathcal{O}(\epsilon)}{\eta}\right). \tag{G.93}$$

Proof. This result is then a direct consequence of Proposition G.2 (Appendix G.2). ∎

Theorem G.2. *Let $i \in \mathcal{I}$ be such that $\mathcal{K}_i$ is unstable. Then there exists $j \in \mathcal{I}$ such that $\mathcal{K}_j$ is stable and a positive constant $c \equiv c(i, j) > 0$ such that, for any $\epsilon > 0$, for any $\mathcal{U}_i$ and $\mathcal{U}_j$ measurable neighborhoods of $\mathcal{K}_i$ and $\mathcal{K}_j$ small enough, there exists $\eta_0 > 0$ such that, for any $0 < \eta < \eta_0$ and for any μ_∞ invariant probability measure for $(x_n^\eta)_n$,*

$$\frac{\mu_\infty(\mathcal{U}_i)}{\mu_\infty(\mathcal{U}_j)} \le \exp\left(-\frac{c(i, j) - \epsilon}{\eta}\right) \tag{G.94}$$

In particular, in the limit $\eta \to 0$, we have $\mu_\infty(\mathcal{U}_i) \to 0$.

Proof. By Lemma G.6, there exists $j \in \mathcal{I}$ such that $\mathcal{K}_j$ is stable and $E(\mathcal{K}_j) < E(\mathcal{K}_i)$. The statement then follows from Theorem G.1. ∎

Theorem G.3. *For any measurable neighborhood $\mathcal{U}$ of $\bigcup_{i\in\mathcal{I}} \mathcal{K}_i$ small enough, there exist constants $c \equiv c_{\mathcal{U}} > 0$ and $\eta_0 > 0$ such that, for any $0 < \eta < \eta_0$ and for any μ_∞ invariant probability measure for $(x_n^\eta)_n$, we have:*

$$\mu_\infty(\mathcal{U}) \ge 1 - e^{-c/\eta}. \tag{G.95}$$

Proof. Define $J \coloneqq \{i \in \mathcal{I} : \mathcal{K}_i \text{ is stable}\}$. Since $\mathcal{U}$ is a neighborhood of $\bigcup_{i\in\mathcal{I}} \mathcal{K}_i$, we can choose neighborhoods $\mathcal{U}_i$ of $\mathcal{K}_i$, for $i \in J$, small enough so that $\bigcup_{i\in J} \mathcal{U}_i \subseteq \mathcal{U}$. This is then a consequence of Proposition G.3 (Appendix G.1) with $J \leftarrow J$. Note J contains the ground states since they are stable by Lemma G.6. ∎

Theorem G.4. *Let $J \coloneqq \arg\min_{i\in\mathcal{I}} E(\mathcal{K}_i)$ be the set of ground states. For any measurable neighborhood $\mathcal{U}_J$ of $\bigcup_{i\in J} \mathcal{K}_i$ small enough, there exist constants $c \equiv c_{\mathcal{U}_J} > 0$ and $\eta_0 > 0$ such that, for any $0 < \eta < \eta_0$ and for any μ_∞ invariant probability measure for $(x_n^\eta)_n$, we have:*

$$\mu_\infty(\mathcal{U}_J) \ge 1 - e^{-c/\eta}. \tag{G.96}$$

Proof. It suffices to apply Proposition G.3 (Appendix G.1) with J as the set of ground states $\arg\min_{i\in\mathcal{I}} E(\mathcal{K}_i)$. They are necessarily stable by Lemma G.6. ∎

G.3.2. **The mean occupation measures μ_n.** We now state and prove analogues of Theorems G.1–G.4 for the mean occupation measure μ_n.

For this, we begin by showing that, in our setting, the sequence of mean occupation measures $(\mu_n)_{n\geq 0}$ is tight, so that it admits a weak accumulation point which is an invariant probability measure for the process.

Lemma G.14. *The sequence of mean occupation measures $(\mu_n)_{n\geq 0}$ is tight: for any $\varepsilon > 0$, there exists a compact set $\mathcal{K} \subset \mathcal{V}$ such that, for all $n \geq 0$, we have $\mu_n(\mathcal{K}) \geq 1 - \varepsilon$.*

Proof. We begin by showing that the sequence of random variables $x_n, n = 0, 1, \ldots$ has a uniformly bounded second moment, i.e.,

$$\sup_{n\geq 0} \mathbb{E}\left[\|x_n\|^2\right] < \infty. \tag{G.97}$$

Fix $n \geq 0$. We have that

$$\|x_{n+1}\|2 = \|x_n\|^2 + 2\eta\langle x_n, F(x_n) + U(x_n, \omega_n)\rangle + \eta^2\|F(x_n) + U(x_n, \omega_n)\|^2 . \tag{G.98}$$

Taking the expectation conditionally on x_n yields that, since $U(x_n, \omega_n)$ is zero-mean conditionally on x_n by Assumption 3,

$$\mathbb{E}\left[\|x_{n+1}\|^2 \mid x_n\right] = \|x_n\|^2 + 2\eta\langle x_n, F(x_n)\rangle + \eta^2\,\mathbb{E}\left[\|F(x_n) + U(x_n, \omega_n)\|^2 \mid x_n\right] \tag{G.99}$$

$$\leq \|x_n\|^2 - 2\eta(\kappa\|x_n\|^2 - C) + 8\eta^2 M^2(1 + \|x_n\|^2) , \tag{G.100}$$

by dissipativity (Assumption 4(a)) and the growth condition (Eq. (D.4)). Hence, setting

$$\eta_0 \coloneqq \frac{\kappa}{8M^2}, \qquad a \coloneqq \kappa, \qquad b \coloneqq 2C + 8M^2\eta_0, \tag{G.101}$$

we have, for all $\eta \leq \eta_0$,

$$\mathbb{E}\left[\|x_{n+1}\|^2 \mid x_n\right] \leq \left(1 - 2\kappa\eta + 8M^2\eta^2\right)\|x_n\|^2 + 2C\eta + 8M^2\eta^2 \tag{G.102}$$

$$\leq (1 - a\eta)\|x_n\|^2 + b\eta. \tag{G.103}$$

Taking expectations and writing

$$m_n \coloneqq \mathbb{E}\left[\|x_n\|^2\right], \tag{G.104}$$

we obtain

$$m_{n+1} \leq (1 - a\eta)m_n + b\eta. \tag{G.105}$$

Applying this inequality recursively gives, for all $n \geq 0$,

$$m_n \leq (1 - a\eta)^n m_0 + b\eta\sum_{j=0}^{n-1}(1 - a\eta)^j \leq m_0 + \frac{b}{a}. \tag{G.106}$$

Therefore, with

$$C \coloneqq \mathbb{E}\left[\|x_0\|^2\right] + \frac{b}{a}, \tag{G.107}$$

we have, for all $n \geq 0$,

$$\mathbb{E}\left[\|x_n\|^2\right] \leq C. \tag{G.108}$$

Hence, there exists $\eta_0, a, b > 0$ such that, for all $\eta \leq \eta_0$,

$$\mathbb{E}\left[\|x_{n+1}\|^2 \mid x_n\right] \leq (1 - a\eta)\|x_n\|^2 + b\,. \tag{G.109}$$

Taking the expectation and applying this inequality recursively yields that there exists $C = C(x_0, \eta) > 0$ such that, for all $n \geq 0$,

$$\mathbb{E}\left[\|x_n\|^2\right] \leq C\,. \tag{G.110}$$

As a consequence, the mean occupation measure μ_n also have a uniformly bounded second moment: for all $n \geq 0$,

$$\int \|x\|^2\, d\mu_n = \frac{1}{n}\sum_{k=0}^{n-1}\mathbb{E}\left[\|x_k\|^2\right] \leq C\,. \tag{G.111}$$

Markov inequality now yields that for any $\varepsilon > 0$, we have, for all $n \geq 0$,

$$\mu_n\Big(\{x : \|x\|^2 > C/\varepsilon\}\Big) \leq \varepsilon\,, \tag{G.112}$$

which shows that the sequence of mean occupation measures $(\mu_n)_{n\geq 0}$ is tight. ∎

Theorem G.5 (Occupation variant of Theorem G.1)**.** *For any $\epsilon > 0$, for any $\mathcal{U}_1, \dots, \mathcal{U}_K$ measurable neighborhoods of $\mathcal{K}_1, \dots, \mathcal{K}_K$ small enough, there exists $\eta_0 > 0$ such that, for any $0 < \eta < \eta_0$ and all n large enough, for any $i, j \in \mathcal{I}$, we have*

$$|\eta \log \mu_n(\mathcal{U}_i) + E(\mathcal{K}_i) - \min_{k\in\mathcal{I}} E(\mathcal{K}_k)| \le \epsilon \tag{G.113}$$

and

$$\left|\eta \log \frac{\mu_n(\mathcal{U}_i)}{\mu_n(\mathcal{U}_j)} + E(\mathcal{K}_i) - E(\mathcal{K}_j)\right| \le \epsilon. \tag{G.114}$$

More compactly, with notation as above, we have:

$$\mu_n(\mathcal{U}_i) \propto \exp\left(-\frac{E(\mathcal{K}_i) + \mathcal{O}(\epsilon)}{\eta}\right). \tag{G.115}$$

Proof. We show that, for sufficiently small neighborhoods and sufficiently small $\eta > 0$, for any $i \in \mathcal{I}$, we have that

$$\begin{aligned}\exp\left(-\frac{E(\mathcal{K}_i) - \min_{j\in\mathcal{I}} E(\mathcal{K}_j) + \epsilon}{\eta}\right) &\le \liminf_{n\to\infty} \mu_n(\mathcal{U}_i)\\ &\le \limsup_{n\to\infty} \mu_n(\mathcal{U}_i)\\ &\le \exp\left(-\frac{E(\mathcal{K}_i) - \min_{j\in\mathcal{I}} E(\mathcal{K}_j) - \epsilon}{\eta}\right)\end{aligned} \tag{G.116}$$

and the results in the statement will follow with 2ϵ in place of ϵ.

We apply Proposition G.2 (Appendix G.2): take the neighborhoods small enough so Proposition G.2 can be applied with both the open neighborhoods $\mathcal{U}_1, \dots, \mathcal{U}_K$ and $\operatorname{cl}\mathcal{U}_1, \dots, \operatorname{cl}\mathcal{U}_K$. One then obtains $\eta_0 > 0$ such that, for all $0 < \eta < \eta_0$ and any μ_∞ invariant probability measure for $(x_n^\eta)_n$, for any $i \in \mathcal{I}$,

$$\begin{aligned}\mu_\infty(\mathcal{U}_i) &\ge \exp\left(-\frac{E(\mathcal{K}_i) - \min_{j\in\mathcal{I}} E(\mathcal{K}_j)}{\eta} - \frac{\epsilon}{\eta}\right)\\ \mu_\infty(\operatorname{cl}\mathcal{U}_i) &\le \exp\left(-\frac{E(\mathcal{K}_i) - \min_{j\in\mathcal{I}} E(\mathcal{K}_j)}{\eta} + \frac{\epsilon}{\eta}\right).\end{aligned} \tag{G.117}$$

We now prove that Eq. (G.116) holds. Fix $i \in \mathcal{I}$. By Lemma G.14, the sequence of mean occupation measures $(\mu_n)_{n\ge0}$ is tight, so that, by Prohorov theorem [61, Thm. 23.2], it is sequentially compact for the weak topology, or, in other terms, for the convergence in distribution. Therefore, every subsequence of $(\mu_n)_{n\ge0}$ admits a weak accumulation point which is a probability distribution and that we denote by ν. Applying Portmanteau theorem [61, Thm. 5.25] to the open set $\mathcal{U}_i$ and the closed set $\operatorname{cl}\mathcal{U}_i$ yields that

$$\nu(\mathcal{U}_i) \le \liminf_{n\to\infty} \mu_n(\mathcal{U}_i) \le \limsup_{n\to\infty} \mu_n(\operatorname{cl}\mathcal{U}_i) \le \nu(\operatorname{cl}\mathcal{U}_i)\,. \tag{G.118}$$

Since $(x_n)_{n\ge0}$ is (weak) Feller by Lemma F.4, ν is actually invariant for $(x_n)_{n\ge0}$, by, e.g., Douc et al. [33, Prop. 12.3.1], and *a fortiori* invariant for the accelerated process $(x_n^\eta)_{n\ge0}$. Combining Eq. (G.117) with Eq. (G.118) gives the result Eq. (G.116). ∎

Theorem G.6 (Occupation variant of Theorem G.2)**.** *Let $i \in \mathcal{I}$ be such that $\mathcal{K}_i$ is unstable. Then there exists $j \in \mathcal{I}$ such that $\mathcal{K}_j$ is stable and a positive constant $c \equiv c(i,j) > 0$ such that, for any $\epsilon > 0$, for any $\mathcal{U}_i$ and $\mathcal{U}_j$ measurable neighborhoods of $\mathcal{K}_i$ and $\mathcal{K}_j$ small enough, there exists $\eta_0 > 0$ such that, for any $0 < \eta < \eta_0$ and all n large enough,*

$$\frac{\mu_n(\mathcal{U}_i)}{\mu_n(\mathcal{U}_j)} \le \exp\left(-\frac{c(i,j) - \epsilon}{\eta}\right) \tag{G.119}$$

In particular, in the limit $\eta \to 0$, we have $\mu_n(\mathcal{U}_i) \to 0$.

Proof. By Lemma G.6, there exists $j \in \mathcal{I}$ such that $\mathcal{K}_j$ is stable and

$$E(\mathcal{K}_j) < E(\mathcal{K}_i)\,. \tag{G.120}$$

Set $c(i,j) \coloneqq E(\mathcal{K}_i) - E(\mathcal{K}_j) > 0$. The result then follows from the ratio estimate in Theorem G.5. ∎

Theorem G.7 (Occupation variant of Theorem G.3)**.** *For any measurable neighborhood $\mathcal{U}$ of $\bigcup_{i\in\mathcal{I}} \mathcal{K}_i$ small enough, there exist constants $c \equiv c_{\mathcal{U}} > 0$ and $\eta_0 > 0$ such that, for any $0 < \eta < \eta_0$ and all n large enough, we have:*

$$\mu_n(\mathcal{U}) \geq 1 - e^{-c/\eta}. \tag{G.121}$$

Proof. We follow the proof of Theorem G.5. Let $a > 0$ and $\eta_0 > 0$ be the constants given by Theorem G.3 for the neighborhood $\mathcal{U}$. Fix $0 < \eta < \eta_0$. By Lemma G.14, the sequence of mean occupation measures $(\mu_n)_{n\geq 0}$ is tight, so that, by Prohorov theorem [61, Thm. 23.2], it is sequentially compact for the weak topology. Therefore, every subsequence of $(\mu_n)_{n\geq 0}$ admits a weak accumulation point which is a probability distribution and that we denote by ν. Applying Portmanteau theorem [61, Thm. 5.25] to the open set $\mathcal{U}$ yields that

$$\nu(\mathcal{U}) \leq \liminf_{n\to\infty} \mu_n(\mathcal{U})\,. \tag{G.122}$$

Since $(x_n)_{n\geq 0}$ is (weak) Feller by Lemma F.4, ν is invariant for $(x_n)_{n\geq 0}$, by, e.g., Douc et al. [33, Prop. 12.3.1], and *a fortiori* invariant for the accelerated process $(x_n^\eta)_{n\geq 0}$. Applying Theorem G.3 to ν gives

$$\liminf_{n\to\infty} \mu_n(\mathcal{U}) \geq 1 - e^{-a/\eta}\,. \tag{G.123}$$

The statement follows by taking $c \coloneqq a/2$, and then n large enough. ∎

Theorem G.8 (Occupation variant of Theorem G.4)**.** *Let $J \coloneqq \arg\min_{i\in\mathcal{I}} E(\mathcal{K}_i)$ be the set of ground states. For any measurable neighborhood $\mathcal{U}_J$ of $\bigcup_{i\in J} \mathcal{K}_i$ small enough, there exist constants $c \equiv c_{\mathcal{U}_J} > 0$ and $\eta_0 > 0$ such that, for any $0 < \eta < \eta_0$ and all n large enough, we have:*

$$\mu_n(\mathcal{U}_J) \geq 1 - e^{-c/\eta}. \tag{G.124}$$

Proof. We follow the proof of Theorem G.5. Let $a > 0$ and $\eta_0 > 0$ be the constants given by Theorem G.4 for the neighborhood $\mathcal{U}_J$. Fix $0 < \eta < \eta_0$. By Lemma G.14, the sequence of mean occupation measures $(\mu_n)_{n\geq 0}$ is tight, so that, by Prohorov theorem [61, Thm. 23.2], it is sequentially compact for the weak topology. Therefore, every subsequence of $(\mu_n)_{n\geq 0}$ admits a weak accumulation point which is a probability distribution and that we denote by ν. Applying Portmanteau theorem [61, Thm. 5.25] to the open set $\mathcal{U}_J$ yields that

$$\nu(\mathcal{U}_J) \leq \liminf_{n\to\infty} \mu_n(\mathcal{U}_J)\,. \tag{G.125}$$

Since $(x_n)_{n\geq 0}$ is (weak) Feller by Lemma F.4, ν is invariant for $(x_n)_{n\geq 0}$, by, e.g., Douc et al. [33, Prop. 12.3.1], and *a fortiori* invariant for the accelerated process $(x_n^\eta)_{n\geq 0}$. Applying Theorem G.4 to ν gives

$$\liminf_{n\to\infty} \mu_n(\mathcal{U}_J) \geq 1 - e^{-a/\eta}\,. \tag{G.126}$$

The statement follows by taking $c \coloneqq a/2$, and then n large enough. ∎

G.3.3. **The lifted mean occupation measures π_n.** We now work in the setting of Appendix E.6. The standard continuous dynamics are

$$\dot{x}_t = F(x_t), \qquad x_t = \Phi_t(x_0), \tag{G.127}$$

while the lifted continuous dynamics are

$$\dot{y}_t = \tilde{F}(y_t), \qquad y_t = \Psi_t(y_0). \tag{G.128}$$

Recall that $Q\colon \mathcal{V} \to \operatorname{int}\mathcal{X}$ is a diffeomorphism. By Lemma E.28, these two flows are conjugate:

$$\Psi_t = Q \circ \Phi_t \circ Q^{-1}. \tag{G.129}$$

We write

$$\widetilde{\mathcal{K}}_i \coloneqq Q(\mathcal{K}_i), \qquad i \in \mathcal{I}. \tag{G.130}$$

We use the same Lyapunov stability notion as in Definition G.1, with Ψ and neighborhoods taken in $\operatorname{int}\mathcal{X}$, for the lifted continuous dynamics.

Lemma G.15. *For any $i \in \mathcal{I}$, $\widetilde{\mathcal{K}}_i = Q(\mathcal{K}_i)$ is stable for the lifted continuous dynamics Eq. (G.128) if and only if $\mathcal{K}_i$ is stable for the standard continuous dynamics Eq. (G.127).*

Proof. By Lemma E.28, the diffeomorphism Q is a topological conjugacy between Φ and Ψ. Using the topological neighborhood formulation of Lyapunov stability, if $\widetilde{\mathcal{U}}$ is a neighborhood of $\widetilde{\mathcal{K}}_i$, then $\mathcal{U} := Q^{-1}(\widetilde{\mathcal{U}})$ is a neighborhood of $\mathcal{K}_i$. Stability of $\mathcal{K}_i$ gives a neighborhood $\mathcal{U}_0$ of $\mathcal{K}_i$ such that $\Phi_t(\mathcal{U}_0) \subseteq \mathcal{U}$ for all $t \geq 0$. Then $\widetilde{\mathcal{U}}_0 := Q(\mathcal{U}_0)$ is a neighborhood of $\widetilde{\mathcal{K}}_i$ and, by the conjugacy relation,

$$\Psi_t(\widetilde{\mathcal{U}}_0) = Q(\Phi_t(\mathcal{U}_0)) \subseteq \widetilde{\mathcal{U}}, \qquad t \geq 0. \tag{G.131}$$

Hence $\widetilde{\mathcal{K}}_i$ is stable for the lifted continuous dynamics. The converse follows by applying the same argument to Q^{-1}. ∎

For readability, we define the lifted sequence by

$$y_n := Q(x_n), \qquad n \geq 0. \tag{G.132}$$

The mean occupation measure of the lifted sequence is defined by

$$\pi_n(\mathcal{B}) := \mathbb{E}\left[\frac{1}{n}\sum_{k=0}^{n-1} \mathbb{1}\{y_k \in \mathcal{B}\}\right], \qquad \mathcal{B} \subseteq \operatorname{int}\mathcal{X} \text{ Borel}. \tag{G.133}$$

Equivalently, for every Borel set $\mathcal{B} \subseteq \operatorname{int}\mathcal{X}$,

$$\pi_n(\mathcal{B}) = \mu_n(Q^{-1}(\mathcal{B})). \tag{G.134}$$

Theorem G.9 (Lifted variant of Theorem G.5)**.** *For any $\epsilon > 0$, for any $\widetilde{\mathcal{U}}_1, \ldots, \widetilde{\mathcal{U}}_K$ measurable neighborhoods of $\widetilde{\mathcal{K}}_1, \ldots, \widetilde{\mathcal{K}}_K$ small enough in* $\operatorname{int}\mathcal{X}$*, there exists $\eta_0 > 0$ such that, for any $0 < \eta < \eta_0$ and all n large enough, for any $i, j \in \mathcal{I}$, we have*

$$\left|\eta \log \pi_n(\widetilde{\mathcal{U}}_i) + E(\mathcal{K}_i) - \min_{k\in\mathcal{I}} E(\mathcal{K}_k)\right| \leq \epsilon \tag{G.135}$$

and

$$\left|\eta \log \frac{\pi_n(\widetilde{\mathcal{U}}_i)}{\pi_n(\widetilde{\mathcal{U}}_j)} + E(\mathcal{K}_i) - E(\mathcal{K}_j)\right| \leq \epsilon. \tag{G.136}$$

More compactly, with notation as above, we have:

$$\pi_n(\widetilde{\mathcal{U}}_i) \propto \exp\left(-\frac{E(\mathcal{K}_i) + \mathcal{O}(\epsilon)}{\eta}\right). \tag{G.137}$$

Proof. Set $\mathcal{U}_i := Q^{-1}(\widetilde{\mathcal{U}}_i)$ for each $i \in \mathcal{I}$. Since Q is a diffeomorphism, each $\mathcal{U}_i$ is a measurable neighborhood of $\mathcal{K}_i$, and these neighborhoods are small enough whenever $\widetilde{\mathcal{U}}_i$ is small enough. Moreover, by definition of π_n,

$$\pi_n(\widetilde{\mathcal{U}}_i) = \mu_n(\mathcal{U}_i) \tag{G.138}$$

for every $i \in \mathcal{I}$. The result follows from Theorem G.5. ∎

Theorem G.10 (Lifted variant of Theorem G.6)**.** *Let $i \in \mathcal{I}$ be such that $\widetilde{\mathcal{K}}_i$ is unstable. Then there exists $j \in \mathcal{I}$ such that $\widetilde{\mathcal{K}}_j$ is stable and a positive constant $c \equiv c(i,j) > 0$ such that, for any $\epsilon > 0$, for any $\widetilde{\mathcal{U}}_i$ and $\widetilde{\mathcal{U}}_j$ measurable neighborhoods of $\widetilde{\mathcal{K}}_i$ and $\widetilde{\mathcal{K}}_j$ small enough in* $\operatorname{int}\mathcal{X}$*, there exists $\eta_0 > 0$ such that, for any $0 < \eta < \eta_0$ and all n large enough,*

$$\frac{\pi_n(\widetilde{\mathcal{U}}_i)}{\pi_n(\widetilde{\mathcal{U}}_j)} \leq \exp\left(-\frac{c(i,j) - \epsilon}{\eta}\right) \tag{G.139}$$

In particular, in the limit $\eta \to 0$, we have $\pi_n(\widetilde{\mathcal{U}}_i) \to 0$.

Proof. Since $\widetilde{\mathcal{K}}_i$ is unstable for the lifted dynamics, Lemma G.15 implies that $\mathcal{K}_i$ is unstable. By Lemma G.6, there exists $j \in \mathcal{I}$ such that $\mathcal{K}_j$ is stable and

$$E(\mathcal{K}_j) < E(\mathcal{K}_i)\,. \tag{G.140}$$

Again by Lemma G.15, $\widetilde{\mathcal{K}}_j$ is stable for the lifted dynamics. Set $c(i,j) := E(\mathcal{K}_i) - E(\mathcal{K}_j) > 0$. The result then follows from the ratio estimate in Theorem G.9. ∎

Theorem G.11 (Lifted variant of Theorem G.7)**.** *For any measurable neighborhood $\mathcal{U}$ of $\bigcup_{i\in\mathcal{I}} \widetilde{\mathcal{K}}_i$ small enough in* $\operatorname{int}\mathcal{X}$*, there exist constants $c \equiv c_{\mathcal{U}} > 0$ and $\eta_0 > 0$ such that, for any $0 < \eta < \eta_0$ and all n large enough, we have:*

$$\pi_n(\mathcal{U}) \geq 1 - e^{-c/\eta}. \tag{G.141}$$

Proof. Set $\mathcal{U} \coloneqq Q^{-1}(\mathcal{U})$. Since Q is a diffeomorphism, $\mathcal{U}$ is a measurable neighborhood of $\bigcup_{i\in\mathcal{I}} \mathcal{K}_i$, and it is small enough whenever $\mathcal{U}$ is small enough. Moreover, by definition of π_n,

$$\pi_n(\mathcal{U}) = \mu_n(\mathcal{U}). \tag{G.142}$$

The result follows from Theorem G.7. ∎

Theorem G.12 (Lifted variant of Theorem G.8)**.** *Let $J \coloneqq \arg\min_{i\in\mathcal{I}} E(\mathcal{K}_i)$ be the set of ground states. For any measurable neighborhood $\mathcal{U}_J$ of $\bigcup_{i\in J} \widetilde{\mathcal{K}}_i$ small enough in* $\operatorname{int}\mathcal{X}$*, there exist constants $c \equiv c_{\mathcal{U}_J} > 0$ and $\eta_0 > 0$ such that, for any $0 < \eta < \eta_0$ and all n large enough, we have:*

$$\pi_n(\mathcal{U}_J) \geq 1 - e^{-c/\eta}. \tag{G.143}$$

Proof. Set $\mathcal{U}_J \coloneqq Q^{-1}(\mathcal{U}_J)$. Since Q is a diffeomorphism, $\mathcal{U}_J$ is a measurable neighborhood of $\bigcup_{i\in J} \mathcal{K}_i$, and it is small enough whenever $\mathcal{U}_J$ is small enough. Moreover, by definition of π_n,

$$\pi_n(\mathcal{U}_J) = \mu_n(\mathcal{U}_J). \tag{G.144}$$

The result follows from Theorem G.8. ∎

H Potential games with Gaussian noise

H.1 Setup: potential games 85
H.2 Informal computations 85
H.3 Technical preliminaries 87
H.4 Potential games with truncated Gaussian noise 90

H.1. Setup: potential games. In this part, we consider potential games: we assume that there exists $f : \mathcal{V} \to \mathbb{R}$ and $Q : \mathcal{V} \to \mathcal{V}$ diffeomorphism such that, for every $x \in \mathcal{V}$,

$$F(x) = \nabla f(Q(x)) \,. \tag{H.1}$$

Assume in addition that Q is the gradient of a strictly convex function. This ensures that the ICT sets of the flow are exactly the critical points of $f \circ Q$ and the equivalence classes are exactly the connected components of the critical points of $f \circ Q$.

Assumption 6 (Essential connectivity of the critical points)**.** For any connected component $\mathcal{K}$ of crit f, and for any two points $x, x' \in \mathcal{K}$, there exists (a) a continuous, almost everywhere differentiable curve $\gamma\colon [0,1] \to \mathcal{K}$ such that $\gamma(0) = x$, $\gamma(1) = x'$; and (b) a partition $0 = t_0 < \cdots < t_n < \cdots < t_T = 1$ of $[0,1]$ such that γ is integrable on every closed interval of (t_{n-1}, t_n) for all $n = 1, \dots, T$.

This assumption is introduced in Azizian et al. [9]. In particular, it is satisfied whenever the connected components of crit(f) are smoothly connected (e.g., smooth manifolds), reduced to isolated critical points or, more generally, when they are definable. Indeed, if f is definable, then crit(f) is definable as well and therefore can be connected by piecewise smooth paths [25, 103].

Lemma H.1. *In the setting above and with Assumption 6, the equivalence classes of $\leftrightsquigarrow$ are exactly the connected components of the critical points of $f \circ Q$.*

Proof. The proof relies on three ingredients.

1. Any connected component of crit$(f \circ Q)$ is included in an equivalence class. This is a consequence of Assumption 6 and is done by Azizian et al. [9, D.11].
2. Any equivalence class is connected: any equivalence class is an ICT set (Theorem 2) and any ICT set is connected; see Benaïm [13, Prop. 5.3].
3. ICT sets are included in crit$(f \circ Q)$: this is a consequence of the fact that $f \circ Q$ is a Lyapunov function for the flow (since Q is the gradient of a strictly convex function) and that any ICT set must be included in the set of points where the Lyapunov function is constant, which is exactly crit$(f \circ Q)$ [13, Prop. 6.]. Therefore, any equivalence class is included in crit$(f \circ Q)$ (Theorem 2).

All together, these three ingredients show that the equivalence classes are exactly the connected components of crit$(f \circ Q)$.

■

In this section, we thus characterize the asymptotic distribution, and more precisely the energies (Appendix G.2) of the different components of crit$(f \circ Q)$ when the noise is Gaussian. The general idea follows Azizian et al. [9, §4.1] but with significant additional challenges: the fact that f is composed with Q changes the geometry and requires to consider the Jacobian of Q in the noise covariance matrix and thus nontrivial covariance geometry.

H.2. Informal computations. Even for potential games, the formula for the invariant measure cannot be simplified in general. In particular, the ground state might not even be the global minimizers of f. However, with careful choices of the noise distribution, it is possible to get a more explicit expression for the invariant measure.

Though it does not formally fit into our setting, we thus consider Gaussian noise. Since it is unbounded, our assumptions are not satisfied and our theorems describing the invariant measure do not apply. However, all the objects we consider are still well-defined, and, in that case, it is possible

to compute the E_i explicitly. Moreover, this section serves as a blueprint for the truncated Gaussian case of the next section.

Assume that, for every $x \in \mathcal{V}$, $U(x, \omega)$ follows a centered Gaussian distribution with positive definite covariance matrix $\Sigma(x)$. For the invariant measure to match the Gibbs measure with potential f, it is required that $\Sigma(x)$ be of a specific form. If Q was the identity, then we would need Σ to be proportional to the identity. For a general Q, we require $\Sigma(x)$ to be proportional to $(\operatorname{Jac} Q(x))^{-1}$. Formally, assume that there exists a continuous function $\sigma^2 : \mathbb{R} \to (0, +\infty)$ such that, for every $x \in \mathcal{V}$,

$$\Sigma(x) = \sigma^2 \circ f \circ Q(x) (\operatorname{Jac} Q(x))^{-1} . \tag{H.2}$$

As we will see with the formal version of this result, this equality is actually only required on a large enough compact set.

Define, for convenience, the norm induced by $\Sigma(x)$: for every $x, w \in \mathcal{V}$,

$$\|w\|_x := \sqrt{\langle w, \Sigma(x)^{-\top} w \rangle} . \tag{H.3}$$

Akin to g_∞ in Appendix D, a key role is played by the function $g : \mathcal{V} \to \mathbb{R}$ defined by

$$g(x) := 2\alpha \circ f \circ Q(x) \quad \text{with} \quad \alpha : \mathbb{R} \to \mathbb{R} \quad \text{a primitive of} \quad 1/\sigma^2 . \tag{H.4}$$

Since the noise is Gaussian, the Lagrangian and Hamiltonian have explicit expressions: for every $x, p, w \in \mathcal{V}$,

$$\mathcal{H}(x, p) = \langle \nabla f(Q(x)), p \rangle + \frac{1}{2} \langle p, \Sigma(x) p \rangle , \qquad \mathcal{L}(x, w) = \frac{1}{2} \|w - \nabla f(Q(x))\|_x^2 . \tag{H.5a}$$

This expression of $\mathcal{L}$ makes it clear that the action function penalizes the deviation of a path from the flow: for a path $\gamma \in \mathcal{C}([0, T])$ for some $T > 0$,

$$\mathcal{S}_T(\gamma) = \int_0^T \frac{1}{2} \|\dot{\gamma}_t - F(\gamma_t)\|_{\gamma_t}^2 \, dt . \tag{H.6}$$

Moreover, with our choice of Σ, we have, for every $x \in \mathcal{V}$,

$$\begin{aligned} \nabla g(x) &= 2\alpha'(f \circ Q(x)) (\operatorname{Jac} Q(x))^\top \nabla f(Q(x)) \\ &= \frac{2}{\sigma^2(f \circ Q(x))} (\operatorname{Jac} Q(x))^\top \nabla f(Q(x)) \\ &= 2\Sigma(x)^{-\top} F(x) . \end{aligned} \tag{H.7}$$

The computation of the E_i relies on the following observation. Take a path $\gamma \in \mathcal{C}([0, T])$ for some $T > 0$ and consider φ defined by $\varphi_t = \gamma_{T-t}$ for $t \in [0, T]$. Then, the action cost of φ is given by,

$$\begin{aligned} \mathcal{S}_T(\varphi) &= \int_0^T \frac{1}{2} \|-\dot{\gamma}_t - F(\gamma_t)\|_{\gamma_t}^2 \, dt \\ &= \int_0^T \frac{1}{2} \|\dot{\gamma}_t - F(\gamma_t)\|_{\gamma_t}^2 \, dt + \int_0^T 2\langle \dot{\gamma}_t, \Sigma(\gamma_t)^{-\top} F(\gamma_t) \rangle \, dt \\ &= \mathcal{S}_T(\gamma) + \int_0^T \langle \dot{\gamma}_t, \nabla g(\gamma_t) \rangle \, dt . \end{aligned} \tag{H.8}$$

Therefore, we get that

$$\mathcal{S}_T(\varphi) = \mathcal{S}_T(\gamma) + (g(\gamma_T) - g(\gamma_0)) . \tag{H.9}$$

Take $i, j \in \mathcal{I}$. This equality then translates to a relation between $B_{i,j}$ and $B_{j,i}$: considering $\gamma \in \mathcal{C}([0, T])$ such that $\gamma_0 \in \mathcal{K}_i$, $\gamma_T \in \mathcal{K}_j$ and taking the infimum over all such paths, we get that

$$B_{j,i} \le B_{i,j} + (g_j - g_i) . \tag{H.10}$$

where, since g is constant on $\mathcal{K}_i$ and $\mathcal{K}_j$, we denote by g_i and g_j the values of g on $\mathcal{K}_i$ and $\mathcal{K}_j$ respectively.

Reversing the roles of i and j and applying the same argument shows that this inequality is an equality:

$$B_{j,i} - g_j = B_{i,j} - g_i \,. \tag{H.11}$$

Denote by $C_{i,j}$ this common value. Crucially, $C_{i,j}$ is symmetric in i and j.

Consider now T^* a minimum weight spanning tree on the complete but now undirected graph on $\mathcal{I}$ with weights $(C_{i,j})_{i,j}$. We show that the minima in the E_i are attained T^*, or more precisely, a directed version of it.

Fix i and consider $T \in \mathcal{T}_i$ a spanning tree rooted at i. We have that, since any node $j \neq i$ is the origin of exactly one edge in T,

$$\begin{aligned}\sum_{(j\to k)\in T} B_{j,k} - \sum_{j\in\mathcal{I}} g_j &= \sum_{(j\to k)\in T} \big(B_{j,k} - g_j\big) - g_i \\ &= \sum_{(j\to k)\in T} C_{j,k} - g_i \,. \end{aligned} \tag{H.12}$$

But by definition of T^*, this sum is at least greater than $\sum_{(j\leftrightarrow k)\in T^*} C_{j,k}$ so that we have

$$\sum_{(j\to k)\in T} B_{j,k} - \sum_{j\in\mathcal{I}} g_j \geq \sum_{(j\leftrightarrow k)\in T^*} C_{j,k} - g_i \,, \tag{H.13}$$

and taking T_i an oriented version of T^* rooted at i, the equality is attained. Therefore, we have that

$$E_i = \sum_{(j\to k)\in T_i} B_{j,k} = \sum_{(j\leftrightarrow k)\in T^*} C_{j,k} + \sum_{j\in\mathcal{I}} g_j - g_i \,, \tag{H.14}$$

or, in short, $E_i = -g_i + c$ where $c > 0$ is independent of i.

Therefore, the mass distribution over critical points is governed by a Gibbs measure with potential $-g$.

Let us now mention two particular cases.

- If σ^2 is constant, then $g = \frac{2f\circ Q}{\sigma^2}$, and

$$E_i = -\frac{2f\circ Q(\mathcal{K}_i)}{\sigma^2} + c \,. \tag{H.15}$$

- It σ^2 is linear, i.e., of the form

$$\sigma^2(f\circ Q(x)) = \sigma_1^2\Big(f\circ Q(x) + \sigma_0^2\Big) \,, \tag{H.16}$$

then

$$g = \frac{2}{\sigma_1^2}\log\Big(f\circ Q + \sigma_0^2\Big) \,, \quad \text{and} \quad E_i = -\frac{2}{\sigma_1^2}\log\Big(f\circ Q + \sigma_0^2\Big)(\mathcal{K}_i) + c \,. \tag{H.17}$$

H.3. Technical preliminaries. To make the computations of the previous section formal, we need to generalize the technical lemmas from Azizian et al. [9, App. F] to the case of Gaussian distributions with arbitrary covariance. We start by generalizing Azizian et al. [9, Lem. F.2]

Lemma H.2. *Consider $X \sim N(0,\Sigma)$ a multivariate Gaussian distribution with covariance matrix Σ positive definite. For $R > 0$, define the truncated Gaussian random variable (r.v.) X_R by conditioning X to the Euclidean ball $\bar{\mathbb{B}}(0,R)$. Define*

$$\bar{\mathcal{H}}(p) \coloneqq \log \mathbb{E}\big[e^{\langle p, X_R\rangle}\big] \tag{H.18a}$$

$$\bar{\mathcal{L}}(w) \coloneqq \bar{\mathcal{H}}^*(w) \,, \tag{H.18b}$$

and

$$E(\Sigma,R) \coloneqq \frac{e^{-\frac{R^2}{16\|\Sigma\|}}(\operatorname{tr}\Sigma + \|\Sigma\|)2^{d+3}}{\lambda_{\min}(\Sigma)} \,. \tag{H.19}$$

Assume that $R > 0$ satisfies

$$R \geq 4\sqrt{\|\Sigma\|\left[(d+5)\log 2 + \log\left(\frac{\operatorname{tr}\Sigma + \|\Sigma\|}{\lambda_{\min}(\Sigma)}\right)\right]} \,. \tag{H.20}$$

Then, for any $p \in \mathcal{V}$ such that

$$\|\Sigma p\| \le \frac{R}{2}\,, \tag{H.21}$$

it holds that

$$(1 - E(\Sigma, R))\frac{\langle p, \Sigma p\rangle}{2} \le \bar{\mathcal{H}}(p) \le (1 + E(\Sigma, R))\frac{\langle p, \Sigma p\rangle}{2}\,. \tag{H.22}$$

Moreover, for any $w \in \mathcal{V}$ such that

$$\|w\| \le \frac{R}{4}\,, \tag{H.23}$$

it holds that

$$(1 - 2E(\Sigma, R))\frac{\langle w, \Sigma^{-1} w\rangle}{2} \le \bar{\mathcal{L}}(w) \le (1 + 2E(\Sigma, R))\frac{\langle w, \Sigma^{-1} w\rangle}{2}\,. \tag{H.24}$$

Proof. We apply Azizian et al. [9, Lem. F.1]with $\widetilde{R} \leftarrow \frac{R}{2}$. First, note that our choice of R implies that

$$\frac{R}{2} \ge \sqrt{\|\Sigma\|(2d+4)\log 2}\,, \tag{H.25}$$

so that the condition of Azizian et al. [9, Lem. F.1]is satisfied. Moreover, our choice of R also implies that

$$\begin{aligned} E(\Sigma, R) &= \frac{e^{-\frac{R^2}{16\|\Sigma\|}}(\operatorname{tr}\Sigma + \|\Sigma\|)2^{d+3}}{\lambda_{\min}(\Sigma)} \\ &\le 2^{-(d+5)}\,\frac{\lambda_{\min}(\Sigma)}{\operatorname{tr}\Sigma + \|\Sigma\|}\,\frac{(\operatorname{tr}\Sigma + \|\Sigma\|)2^{d+3}}{\lambda_{\min}(\Sigma)} \\ &= \frac{1}{4}\,. \end{aligned} \tag{H.26}$$

Let $p \in \mathcal{V}$ be such that

$$\|\Sigma p\| \le \frac{R}{2}\,. \tag{H.27}$$

By Azizian et al. [9, Lem. F.1], we have

$$\begin{aligned} \left|\bar{\mathcal{H}}(p) - \frac{1}{2}\langle p, \Sigma p\rangle\right| &\le \frac{1}{2}e^{-\frac{R^2}{16\|\Sigma\|}}(\operatorname{tr}\Sigma + \|\Sigma\|)2^{d+3}\|p\|^2 \\ &= \frac{1}{2}E(\Sigma, R)\lambda_{\min}(\Sigma)\|p\|^2 \\ &\le \frac{1}{2}E(\Sigma, R)\langle p, \Sigma p\rangle\,. \end{aligned} \tag{H.28}$$

This gives

$$(1 - E(\Sigma, R))\frac{\langle p, \Sigma p\rangle}{2} \le \bar{\mathcal{H}}(p) \le (1 + E(\Sigma, R))\frac{\langle p, \Sigma p\rangle}{2}\,. \tag{H.29}$$

We now prove the bound on $\bar{\mathcal{L}}$. Set

$$\mathcal{C}_R := \left\{p \in \mathcal{V} : \|\Sigma p\| \le \frac{R}{2}\right\}. \tag{H.30}$$

We first show that, for any $w \in \mathcal{V}$ such that $\|w\| \le \frac{R}{4}$, the supremum defining $\bar{\mathcal{L}}(w)$ is attained inside $\mathcal{C}_R$.

Consider the maximization problem

$$\sup_{p \in \mathcal{C}_R} \langle w, p\rangle - \bar{\mathcal{H}}(p)\,. \tag{H.31}$$

Since $\mathcal{C}_R$ is compact and $p \mapsto \langle w, p\rangle - \bar{\mathcal{H}}(p)$ is continuous, there exists a maximizer $p_R \in \mathcal{C}_R$.

We claim that p_R cannot belong to the boundary of $\mathcal{C}_R$. Assume by contradiction that

$$\|\Sigma p_R\| = \frac{R}{2}\,. \tag{H.32}$$

Then, by the first-order optimality condition for the constrained maximization problem, there exists $\lambda \geq 0$ such that

$$w - \nabla \bar{\mathcal{H}}(p_R) = 2\lambda \Sigma^2 p_R \,. \tag{H.33}$$

Therefore,

$$\langle \Sigma p_R, \nabla \bar{\mathcal{H}}(p_R) - w \rangle = -2\lambda p_R^\top \Sigma^3 p_R \leq 0 \,. \tag{H.34}$$

On the other hand, using again Azizian et al. [9, Lem. F.1], we have

$$\begin{aligned} \left\| \nabla \bar{\mathcal{H}}(p_R) - \Sigma p_R \right\| &\leq e^{-\frac{R^2}{16\|\Sigma\|}} (\operatorname{tr} \Sigma + \|\Sigma\|) 2^{d+3} \|p_R\| \\ &= E(\Sigma, R) \lambda_{\min}(\Sigma) \|p_R\| \\ &\leq E(\Sigma, R) \|\Sigma p_R\| \,. \end{aligned} \tag{H.35}$$

Hence,

$$\begin{aligned} \langle \Sigma p_R, \nabla \bar{\mathcal{H}}(p_R) - w \rangle &= \|\Sigma p_R\|^2 + \langle \Sigma p_R, \nabla \bar{\mathcal{H}}(p_R) - \Sigma p_R \rangle - \langle \Sigma p_R, w \rangle \\ &\geq \|\Sigma p_R\|^2 - E(\Sigma, R) \|\Sigma p_R\|^2 - \|\Sigma p_R\| \|w\| \\ &= \|\Sigma p_R\| ((1 - E(\Sigma, R)) \|\Sigma p_R\| - \|w\|) \\ &\geq \frac{R}{2} \left(\frac{3R}{8} - \frac{R}{4} \right) > 0 \,, \end{aligned} \tag{H.36}$$

where we used that $E(\Sigma, R) \leq \frac{1}{4}$ and $\|w\| \leq \frac{R}{4}$. This is a contradiction. Therefore, p_R belongs to the interior of $\mathcal{C}_R$, and so

$$\nabla \bar{\mathcal{H}}(p_R) = w \,. \tag{H.37}$$

Since $\bar{\mathcal{H}}$ is convex, p_R is also a global maximizer of $p \mapsto \langle w, p \rangle - \bar{\mathcal{H}}(p)$. Therefore,

$$\bar{\mathcal{L}}(w) = \sup_{p \in \mathcal{C}_R} \langle w, p \rangle - \bar{\mathcal{H}}(p) \,. \tag{H.38}$$

We can now use the previous bound on $\bar{\mathcal{H}}$. For the lower bound, set

$$p_0 \coloneqq \frac{1}{1 + E(\Sigma, R)} \Sigma^{-1} w \,. \tag{H.39}$$

Then

$$\|\Sigma p_0\| = \frac{\|w\|}{1 + E(\Sigma, R)} \leq \frac{R}{4} < \frac{R}{2} \,, \tag{H.40}$$

so that $p_0 \in \mathcal{C}_R$. Hence,

$$\begin{aligned} \bar{\mathcal{L}}(w) &\geq \langle w, p_0 \rangle - \bar{\mathcal{H}}(p_0) \\ &\geq \langle w, p_0 \rangle - (1 + E(\Sigma, R)) \frac{\langle p_0, \Sigma p_0 \rangle}{2} \\ &= \frac{1}{2(1 + E(\Sigma, R))} \langle w, \Sigma^{-1} w \rangle \,. \end{aligned} \tag{H.41}$$

For the upper bound, we use that

$$\begin{aligned} \bar{\mathcal{L}}(w) &= \sup_{p \in \mathcal{C}_R} \langle w, p \rangle - \bar{\mathcal{H}}(p) \\ &\leq \sup_{p \in \mathcal{C}_R} \langle w, p \rangle - (1 - E(\Sigma, R)) \frac{\langle p, \Sigma p \rangle}{2} \\ &\leq \sup_{p \in \mathcal{V}} \langle w, p \rangle - (1 - E(\Sigma, R)) \frac{\langle p, \Sigma p \rangle}{2} \\ &= \frac{1}{2(1 - E(\Sigma, R))} \langle w, \Sigma^{-1} w \rangle \,. \end{aligned} \tag{H.42}$$

Therefore,

$$\frac{1}{1 + E(\Sigma, R)} \frac{\langle w, \Sigma^{-1} w \rangle}{2} \leq \bar{\mathcal{L}}(w) \leq \frac{1}{1 - E(\Sigma, R)} \frac{\langle w, \Sigma^{-1} w \rangle}{2} \,. \tag{H.43}$$

Since, for $x \in [0, 1/2]$, both $\frac{1}{1+x} \geq 1 - 2x$ and $\frac{1}{1-x} \leq 1 + 2x$ hold, and since $E(\Sigma, R) \leq \frac{1}{4}$, we obtain

$$(1 - 2E(\Sigma, R))\frac{\langle w, \Sigma^{-1} w\rangle}{2} \leq \bar{\mathcal{L}}(w) \leq (1 + 2E(\Sigma, R))\frac{\langle w, \Sigma^{-1} w\rangle}{2}\,. \tag{H.44}$$

This concludes the proof. ∎

We now generalize Azizian et al. [9, Lem. F.3].

Lemma H.3. *Consider $w, w \in \mathcal{V}$ such that $w \neq 0$ and $w \neq 0$. Let Σ be a symmetric positive definite matrix. Define*

$$\lambda := \sqrt{\frac{\langle w, \Sigma^{-1} w\rangle}{\langle w, \Sigma^{-1} w\rangle}}\,, \tag{H.45}$$

and assume that

$$\left\|\Sigma^{-1}\left(\frac{w}{\lambda} + w\right)\right\| \leq R. \tag{H.46}$$

Define

$$f(u) = \sup_{p \in \mathcal{V}: \|p\| \leq R} \langle p, u\rangle - \frac{1}{2}\langle p, \Sigma p\rangle\,. \tag{H.47}$$

Then, we have

$$\lambda f\left(\frac{w}{\lambda} + w\right) \leq f(w + w)\,. \tag{H.48}$$

Proof. Define

$$p := \Sigma^{-1}\left(\frac{w}{\lambda} + w\right), \tag{H.49}$$

which has norm at most R by assumption. Then, we have that

$$\begin{aligned}\lambda f\left(\frac{w}{\lambda} + w\right) - f(w + w) &\leq \lambda\left(\left\langle p, \frac{w}{\lambda} + w\right\rangle - \frac{1}{2}\langle p, \Sigma p\rangle\right) - \left(\langle p, w + w\rangle - \frac{1}{2}\langle p, \Sigma p\rangle\right) \\ &= (\lambda - 1)\left(\langle p, w\rangle - \frac{1}{2}\langle p, \Sigma p\rangle\right).\end{aligned} \tag{H.50}$$

Moreover, since

$$\Sigma p = \frac{w}{\lambda} + w, \tag{H.51}$$

we have that

$$\begin{aligned}\langle p, w\rangle - \frac{1}{2}\langle p, \Sigma p\rangle &= \left\langle \Sigma^{-1}\left(\frac{w}{\lambda} + w\right), w\right\rangle - \frac{1}{2}\left\langle \Sigma^{-1}\left(\frac{w}{\lambda} + w\right), \frac{w}{\lambda} + w\right\rangle \\ &= \frac{1}{2}\left(\langle w, \Sigma^{-1} w\rangle - \frac{1}{\lambda^2}\langle w, \Sigma^{-1} w\rangle\right) \\ &= 0,\end{aligned} \tag{H.52}$$

by definition of λ. Therefore,

$$\lambda f\left(\frac{w}{\lambda} + w\right) - f(w + w) \leq 0, \tag{H.53}$$

and our proof is complete. ∎

H.4. Potential games with truncated Gaussian noise. To fit into our theoretical framework, we consider truncated Gaussian noise instead. The general outline of the proof but with added steps to handle the error due to the truncation and to allow for covariance matrices that do not necessarily match the Jacobian of Q.

Assumption 7 (Truncated Gaussian noise)**.** Assume that $U(x, \omega)$ follows a centered Gaussian distribution with covariance $\Sigma(x)$ conditioned on being in $\mathbb{B}(0, R(x))$ for some $R(x)$. Moreover, $\Sigma(x)$ is a continuous function of x and is always positive-definite and that $R(x) > 0$ is a continuous function of x that is always positive.

We first begin by refining Lemma E.17.

Lemma H.4. *Define $\mathcal{C} \coloneqq \operatorname{conv}\{\mathcal{K}_i : i \in \mathcal{I}\}$. Then, there exists a compact set $\mathcal{K} \subset \mathcal{V}$ that depends only on*

$$\max_{x\in\mathcal{C}}\|x\|, \max_{\mathcal{C}}\|\Sigma\|, \min_{\mathcal{C}}\lambda_{\min}(\Sigma), \min_{\mathcal{C}} R, M, \overline{\sigma}^2, s, \kappa, C \tag{H.54}$$

such that, for every $i, j \in \mathcal{I}$,

$$B_{i,j} = B_{\mathcal{K}}(\mathcal{K}_i, \mathcal{K}_j)\,. \tag{H.55}$$

Proof. This result is a consequence of the construction in Lemma E.17 and more precisely of Eq. (E.105). To show the result, it suffices to show that $B(\mathcal{K}_i, \mathcal{K}_j)$ can be bounded by a quantity that depends only on the aforementioned parameters, so that we can choose $\mathcal{K}$ as the compact set constructed in Lemma E.17 using Eq. (E.105) with this bound as $B(\mathcal{K}_i, \mathcal{K}_j)$.

By continuity of Σ and R, there exists $C < +\infty$ such that,

$$\max_{x\in\mathcal{C}}\|x\| \le C, \quad \max_{\mathcal{C}}\|\Sigma\| \le C, \quad \min_{\mathcal{C}}\lambda_{\min}(\Sigma) \ge C^{-1}, \quad \min_{\mathcal{C}} R \ge C^{-1}\,. \tag{H.56}$$

Take $i \neq j$ and $x_i \in \mathcal{K}_i$, $x_j \in \mathcal{K}_j$. Consider the path $\gamma \in \mathcal{C}_T(\mathcal{C})$ defined by, for any $t \in [0, T]$,

$$\gamma(t) \coloneqq x_i + \frac{t}{T}(x_j - x_i) \in \mathcal{C}\,, \tag{H.57}$$

with $T \coloneqq \|x_j - x_i\|/(4C)$. In particular, for any $t \in [0, T]$, $\|\dot\gamma(t)\| \le \frac{R(\gamma(t))}{4}$. We can thus apply the second set of inequalities of Lemma H.2 to get that, there exists C' depending only on C such that, for any $t \in [0, T]$,

$$\mathcal{L}(\gamma(t), \dot\gamma(t)) = \bar{\mathcal{L}}(\gamma(t), \dot\gamma(t) - F(\gamma(t))) \tag{H.58}$$

$$\le \frac{C'}{2}\|\dot\gamma(t) - F(\gamma(t))\|^2, \qquad \le C'(\|\dot\gamma(t)\|^2 + \|F(\gamma(t))\|^2)\,, \le C'\Big(\frac{1}{16C^2} + M^2(1 + C^2)\Big)\,, \tag{H.59}$$

where we used Young's inequality, the choices of T and C and the growth assumption on F (Eq. (D.4)). Integrating over $t \in [0, T]$ yields that $\mathcal{S}_T[\gamma] \le C''$ for some C'' depending only on C and M, so that $B_{i,j} \le C''$. ∎

In particular, one must show that, without loss of generality, we can only consider paths whose derivative has the same $\|\bullet\|_\bullet$-norm as the vector field. This is done with Freidlin & Wentzell [36, Chap. 4, Lem. 3.1] that we adapt to our setting.

Assume that there exists $f : \mathcal{V} \to \mathbb{R}$ and $Q : \mathcal{V} \to \mathcal{V}$ diffeomorphism such that, for every $x \in \mathcal{V}$,

$$F(x) = \nabla f(Q(x))\,. \tag{H.60}$$

Equivalently, compared to the gradient case, the objective is $-f \circ Q$. Assume moreover that Q is the gradient of a strongly convex function and let $\sigma : \mathbb{R} \to (0, +\infty)$ be a continuous function, with $\sigma^2 = \sigma^2$. We do not assume that the covariance matrix is exactly of the form appearing in Appendix H.2; instead, the defect will be measured on the compact set $\mathcal{K}$ given by Lemma H.4. As in Appendix H.2, for every $x, w \in \mathcal{V}$, we write

$$\|w\|_x \coloneqq \sqrt{\langle w, \Sigma(x)^{-1} w\rangle}\,. \tag{H.61}$$

As in Definition D.1, we define

$$g(x) = 2\alpha \circ f \circ Q(x) \quad \text{with} \quad \alpha' = \sigma^{-2}\,, \tag{H.62}$$

and denote by g_i the value taken by g on $\mathcal{K}_i$. With this definition, we have, for every $x \in \mathcal{K}$,

$$\begin{aligned}\nabla g(x) &= 2\alpha'(f \circ Q(x))\operatorname{Jac} Q(x)\,\nabla f(Q(x)) \\ &= 2\sigma^{-2}(f \circ Q(x))\operatorname{Jac} Q(x)\,\nabla f(Q(x)) \\ &= 2\Big(\sigma^{-2}(f \circ Q(x))\operatorname{Jac} Q(x)\Big)F(x)\,.\end{aligned} \tag{H.63}$$

Consider

$$\delta \coloneqq \sup_{x\in\mathcal{K}} \frac{e^{-\frac{R(x)^2}{16\|\Sigma(x)\|}}(\operatorname{tr}\Sigma(x) + \|\Sigma(x)\|)2^{d+3}}{\lambda_{\min}(\Sigma(x))} \tag{H.64}$$

and assume that $0 < \delta \le \frac{1}{2}$, so that the error term $2E(\Sigma(x), R(x))$ in Lemma H.2 is bounded by δ on $\mathcal{K}$. Moreover, assume that, for any $x \in \mathcal{K}$,

$$\sqrt{\|\Sigma(x)\|\|\Sigma(x)^{-1}\|}\|F(x)\| \le \frac{R(x)}{8}\,. \tag{H.65}$$

Lemma H.5. *Consider $\gamma \in \mathcal{C}_T(\mathcal{K})$. Then, there exists $\widetilde{\gamma} \in \mathcal{C}_S(\mathcal{K})$ a reparametrization of γ such that, for any $s \in [0, S]$,*

$$\|\dot{\widetilde{\gamma}}_s\|_{\widetilde{\gamma}_s} = \|F(\widetilde{\gamma}_s)\|_{\widetilde{\gamma}_s}\,. \tag{H.66}$$

and

$$\mathcal{S}_T(\gamma) \ge \int_0^S \frac{1}{2(1+\delta)}\|\dot{\widetilde{\gamma}}_s - F(\widetilde{\gamma}_s)\|^2_{\widetilde{\gamma}_s}\,ds\,. \tag{H.67}$$

Proof. By the proof Freidlin & Wentzell [36, Chap. 4, Lem. 3.1], there exists $t(s)$ change of time such that, with $\widetilde{\gamma}_s = \gamma_{t(s)}$,

$$\|\dot{\widetilde{\gamma}}_s\|_{\widetilde{\gamma}_s} = \|F(\widetilde{\gamma}_s)\|_{\widetilde{\gamma}_s}\,. \tag{H.68}$$

We have that

$$\mathcal{S}_T(\gamma) = \int_0^{t^{-1}(T)} \dot{t}(s)\mathcal{L}(\widetilde{\gamma}_s, (\dot{t}(s))^{-1}\dot{\widetilde{\gamma}}_s)\,ds\,, \tag{H.69}$$

so it suffices to bound $\mathcal{L}(\widetilde{\gamma}_s, (\dot{t}(s))^{-1}\dot{\widetilde{\gamma}}_s)$ from below. By definition, we have

$$\begin{aligned}\mathcal{L}(\widetilde{\gamma}_s, (\dot{t}(s))^{-1}\dot{\widetilde{\gamma}}_s) &\ge \sup\left\{\langle p, (\dot{t}(s))^{-1}\dot{\widetilde{\gamma}}_s - F(\widetilde{\gamma}_s)\rangle - \bar{\mathcal{H}}(\widetilde{\gamma}_s, p) : \|\Sigma(\widetilde{\gamma}_s)p\| \le \frac{R(\widetilde{\gamma}_s)}{2}\right\}\\ &\ge \sup\left\{\langle p, (\dot{t}(s))^{-1}\dot{\widetilde{\gamma}}_s - F(\widetilde{\gamma}_s)\rangle - (1+\delta)\frac{1}{2}\langle p, \Sigma(\widetilde{\gamma}_s)p\rangle : \|\Sigma(\widetilde{\gamma}_s)p\| \le \frac{R(\widetilde{\gamma}_s)}{2}\right\},\end{aligned} \tag{H.70}$$

by Lemma H.2. Applying Lemma H.3, after the change of variables $p \leftarrow \Sigma(\widetilde{\gamma}_s)p$, with

$$w \leftarrow (\dot{t}(s))^{-1}\dot{\widetilde{\gamma}}_s, \quad w \leftarrow -F(\widetilde{\gamma}_s), \quad \lambda \leftarrow (\dot{t}(s))^{-1}, \tag{H.71}$$

now exactly yields, for almost all s,

$$\begin{aligned}\mathcal{L}(\widetilde{\gamma}_s, (\dot{t}(s))^{-1}\dot{\widetilde{\gamma}}_s) &\ge (\dot{t}(s))^{-1}\sup\left\{\langle p, \dot{\widetilde{\gamma}}_s - F(\widetilde{\gamma}_s)\rangle - (1+\delta)\frac{1}{2}\langle p, \Sigma(\widetilde{\gamma}_s)p\rangle : \|\Sigma(\widetilde{\gamma}_s)p\| \le \frac{R(\widetilde{\gamma}_s)}{2}\right\}\\ &= (\dot{t}(s))^{-1}\frac{1}{2(1+\delta)}\|\dot{\widetilde{\gamma}}_s - F(\widetilde{\gamma}_s)\|^2_{\widetilde{\gamma}_s}\,,\end{aligned} \tag{H.72}$$

since

$$\|\dot{\widetilde{\gamma}}_s - F(\widetilde{\gamma}_s)\| \le \frac{R(\widetilde{\gamma}_s)}{4}\,. \tag{H.73}$$

Indeed, by the choice of the change of time and the previous assumption,

$$\|\dot{\widetilde{\gamma}}_s\| \le \sqrt{\|\Sigma(\widetilde{\gamma}_s)\|\|\Sigma(\widetilde{\gamma}_s)^{-1}\|}\|F(\widetilde{\gamma}_s)\| \le \frac{R(\widetilde{\gamma}_s)}{8}\,, \tag{H.74}$$

and the same bound holds for $\|F(\widetilde{\gamma}_s)\|$. This concludes the proof. ∎

Theorem H.1. *In the setting of Lemma H.4, define*

$$\delta := \sup_{x\in\mathcal{K}} \frac{e^{-\frac{R(x)^2}{16\|\Sigma(x)\|}}(\operatorname{tr}\Sigma(x) + \|\Sigma(x)\|)2^{d+3}}{\lambda_{\min}(\Sigma(x))}\,, \tag{H.75}$$

$$\zeta := \sup_{x\in\mathcal{K}}\left\|I - \Sigma^{1/2}(x)\sigma^{-2}(f\circ Q(x))\operatorname{Jac}Q(x)\Sigma^{1/2}(x)\right\|. \tag{H.76}$$

Assume that $0 < \delta \le \frac{1}{2}$, that $0 \le \zeta < 1$, and that, for any $x \in \mathcal{K}$,

$$\sqrt{\|\Sigma(x)\|\|\Sigma(x)^{-1}\|}\|F(x)\| \le \frac{R(x)}{8}\,. \tag{H.77}$$

Finally, define

$$Q_{\mathcal{K}} \coloneqq \max_{i,j\in\mathcal{I}} B_{i,j}, \quad M_{\mathcal{K}} \coloneqq \max_{i,j\in\mathcal{I}} |g_i - g_j| \tag{H.78}$$

and

$$\Delta_{\mathcal{K}} \coloneqq \frac{2(\delta+\zeta)Q_{\mathcal{K}} + (\delta+\zeta-\delta\zeta)M_{\mathcal{K}}}{(1-\delta)(1-\zeta)}\,. \tag{H.79}$$

Then, for every $i \in \mathcal{I}$,

$$E_i = -g_i + \epsilon_i c_{\mathcal{K}}\,, \quad \text{for some } \epsilon_i \in [0, (|\mathcal{I}|-1)\Delta_{\mathcal{K}}]\,, \tag{H.80}$$

and where $c_{\mathcal{K}}$ is a constant independent of i.

Proof. Consider $\gamma \in \mathcal{C}_T(\mathcal{K})$ such that $\gamma_0 \in \mathcal{K}_i$, $\gamma_T \in \mathcal{K}_j$ and $\mathcal{S}_T(\gamma) < +\infty$. Then, by Lemma H.5, there exists $\widetilde{\gamma} \in \mathcal{C}_S(\mathcal{K})$ a reparametrization of γ such that, for any $s \in [0, S]$,

$$\begin{aligned}
\mathcal{S}_T(\gamma) &\geq \int_0^S \frac{1}{2(1+\delta)} \|\dot{\widetilde{\gamma}}_s - F(\widetilde{\gamma}_s)\|^2_{\widetilde{\gamma}_s}\, ds \\
&= \int_0^S \frac{1}{2(1+\delta)} \|-\dot{\widetilde{\gamma}}_s - F(\widetilde{\gamma}_s)\|^2_{\widetilde{\gamma}_s}\, ds - \frac{g(\gamma_T) - g(\gamma_0)}{1+\delta} \\
&\quad - \frac{2}{1+\delta} \int_0^S \langle \dot{\widetilde{\gamma}}_s, \Big(\Sigma(\widetilde{\gamma}_s)^{-1} - \sigma^{-2}(f\circ Q(\widetilde{\gamma}_s))\operatorname{Jac} Q(\widetilde{\gamma}_s)\Big) F(\widetilde{\gamma}_s)\rangle\, ds\,.
\end{aligned} \tag{H.81}$$

Indeed, this is the same computation as above in Appendix H.2, except that Σ^{-1} is now not necessarily proportional to $\operatorname{Jac} Q$ and that we have an additional error term due to the truncation of the Gaussian noise. By construction of the change of time, we have

$$\|\dot{\widetilde{\gamma}}_s\|_{\widetilde{\gamma}_s} = \|F(\widetilde{\gamma}_s)\|_{\widetilde{\gamma}_s}\,. \tag{H.82}$$

Therefore, for almost every s,

$$\begin{aligned}
&2\Big|\langle \dot{\widetilde{\gamma}}_s, \Big(\Sigma(\widetilde{\gamma}_s)^{-1} - \sigma^{-2}(f\circ Q(\widetilde{\gamma}_s))\operatorname{Jac} Q(\widetilde{\gamma}_s)\Big) F(\widetilde{\gamma}_s)\rangle\Big| \\
&= 2\Big|\langle \dot{\widetilde{\gamma}}_s, \Sigma^{-1/2}(\widetilde{\gamma}_s)\Big(I - \Sigma^{1/2}(\widetilde{\gamma}_s)\sigma^{-2}(f\circ Q(\widetilde{\gamma}_s))\operatorname{Jac} Q(\widetilde{\gamma}_s)\Sigma^{1/2}(\widetilde{\gamma}_s)\Big)\Sigma^{-1/2}(\widetilde{\gamma}_s) F(\widetilde{\gamma}_s)\rangle\Big| \\
&\leq 2\zeta \|\dot{\widetilde{\gamma}}_s\|_{\widetilde{\gamma}_s} \|F(\widetilde{\gamma}_s)\|_{\widetilde{\gamma}_s} \\
&= \zeta(\|\dot{\widetilde{\gamma}}_s\|^2_{\widetilde{\gamma}_s} + \|F(\widetilde{\gamma}_s)\|^2_{\widetilde{\gamma}_s}) \\
&= \frac{\zeta}{2}\Big(\|-\dot{\widetilde{\gamma}}_s - F(\widetilde{\gamma}_s)\|^2_{\widetilde{\gamma}_s} + \|\dot{\widetilde{\gamma}}_s - F(\widetilde{\gamma}_s)\|^2_{\widetilde{\gamma}_s}\Big)\,.
\end{aligned} \tag{H.83}$$

Therefore, for every $s \in [0, S]$, writing

$$\theta_s = \langle \dot{\widetilde{\gamma}}_s, \Big(\Sigma(\widetilde{\gamma}_s)^{-1} - \sigma^{-2}(f\circ Q(\widetilde{\gamma}_s))\operatorname{Jac} Q(\widetilde{\gamma}_s)\Big) F(\widetilde{\gamma}_s)\rangle, \tag{H.84}$$

we have

$$\begin{aligned}
\frac{1}{2}\|\dot{\widetilde{\gamma}}_s - F(\widetilde{\gamma}_s)\|^2_{\widetilde{\gamma}_s} &= \frac{1}{2}\|-\dot{\widetilde{\gamma}}_s - F(\widetilde{\gamma}_s)\|^2_{\widetilde{\gamma}_s} - 2\langle \dot{\widetilde{\gamma}}_s, \Sigma(\widetilde{\gamma}_s)^{-1} F(\widetilde{\gamma}_s)\rangle \\
&= \frac{1}{2}\|-\dot{\widetilde{\gamma}}_s - F(\widetilde{\gamma}_s)\|^2_{\widetilde{\gamma}_s} - \frac{d}{ds} g(\widetilde{\gamma}_s) - 2\theta_s\,,
\end{aligned} \tag{H.85}$$

where we used the definition of g for the second equality. Hence, by the previous estimate on θ_s,

$$\begin{aligned}
\frac{1}{2}\|\dot{\widetilde{\gamma}}_s - F(\widetilde{\gamma}_s)\|^2_{\widetilde{\gamma}_s} \geq \frac{1-\zeta}{2}\|-\dot{\widetilde{\gamma}}_s - F(\widetilde{\gamma}_s)\|^2_{\widetilde{\gamma}_s} - \frac{d}{ds} g(\widetilde{\gamma}_s) \\
- \frac{\zeta}{2}\|\dot{\widetilde{\gamma}}_s - F(\widetilde{\gamma}_s)\|^2_{\widetilde{\gamma}_s}\,.
\end{aligned} \tag{H.86}$$

Integrating over $[0, S]$ gives

$$\int_0^S \frac{1}{2}\|\dot{\widetilde{\gamma}}_s - F(\widetilde{\gamma}_s)\|^2_{\widetilde{\gamma}_s}\, ds \geq \frac{1-\zeta}{2}\int_0^S \|-\dot{\widetilde{\gamma}}_s - F(\widetilde{\gamma}_s)\|^2_{\widetilde{\gamma}_s}\, ds$$

$$- (g(\gamma_T) - g(\gamma_0)) \\ - \frac{\zeta}{2} \int_0^S \|\dot{\widetilde{\gamma}}_s - F(\widetilde{\gamma}_s)\|^2_{\widetilde{\gamma}_s} \, ds \,. \tag{H.87}$$

Hence, using again the lower-bound from Lemma H.5,

$$(1+\zeta)\mathcal{S}_T(\gamma) \geq \frac{1-\zeta}{2(1+\delta)} \int_0^S \|-\dot{\widetilde{\gamma}}_s - F(\widetilde{\gamma}_s)\|^2_{\widetilde{\gamma}_s} \, ds \\ - \frac{g(\gamma_T) - g(\gamma_0)}{1+\delta} \,. \tag{H.88}$$

Considering the path $(\widetilde{\gamma}_{S-s})_{s\in[0,S]}$ and invoking the upper-bound on the Lagrangian from Lemma H.2 with $-\dot{\widetilde{\gamma}}_s - F(\widetilde{\gamma}_s)$ which still has norm less than $R(\widetilde{\gamma}_s)/4$, we get that

$$\mathcal{S}_T(\gamma) \geq \frac{1-\zeta}{1+\zeta}\frac{1-\delta}{1+\delta} B_{j,i} - \frac{g(\gamma_T) - g(\gamma_0)}{(1+\delta)(1+\zeta)} \,. \tag{H.89}$$

Taking the infimum over all such paths, we obtain

$$B_{i,j} \geq \alpha_{\mathcal{K}} B_{j,i} - \beta_{\mathcal{K}} \big(g_j - g_i\big) \,, \tag{H.90}$$

where

$$\alpha_{\mathcal{K}} \coloneqq \frac{1-\zeta}{1+\zeta}\frac{1-\delta}{1+\delta} \quad \text{and} \quad \beta_{\mathcal{K}} \coloneqq \frac{1}{(1+\delta)(1+\zeta)} \,. \tag{H.91}$$

where, since g is constant on $\mathcal{K}_i$ and $\mathcal{K}_j$, we denote by g_i and g_j the values of g on $\mathcal{K}_i$ and $\mathcal{K}_j$ respectively.

Reversing the roles of i and j and applying the same argument shows that the quantities

$$B_{i,j} - g_i \quad \text{and} \quad B_{j,i} - g_j \tag{H.92}$$

are not exactly equal anymore, but they are close. Indeed, from the previous inequality, we have

$$B_{j,i} - g_j \leq B_{i,j} - g_i + \Big(\alpha_{\mathcal{K}}^{-1} - 1\Big) B_{i,j} + \left(\frac{\beta_{\mathcal{K}}}{\alpha_{\mathcal{K}}} - 1\right)\big(g_j - g_i\big) \\ \leq B_{i,j} - g_i + \Delta_{\mathcal{K}} \,. \tag{H.93}$$

In the last inequality, we used that

$$\alpha_{\mathcal{K}}^{-1} - 1 = \frac{2(\delta+\zeta)}{(1-\delta)(1-\zeta)} \quad \text{and} \quad \frac{\beta_{\mathcal{K}}}{\alpha_{\mathcal{K}}} - 1 = \frac{\delta + \zeta - \delta\zeta}{(1-\delta)(1-\zeta)} \,. \tag{H.94}$$

Reversing the roles of i and j gives the reverse inequality. Thus, for every $i, j \in \mathcal{I}$,

$$\big|\big(B_{i,j} - g_i\big) - \big(B_{j,i} - g_j\big)\big| \leq \Delta_{\mathcal{K}} \,. \tag{H.95}$$

Define, for every $i, j \in \mathcal{I}$,

$$C_{i,j} \coloneqq \min\big\{B_{i,j} - g_i, B_{j,i} - g_j\big\} \,. \tag{H.96}$$

Consider now T^* a minimum weight spanning tree on the complete undirected graph on $\mathcal{I}$ with weights $(C_{i,j})_{i,j}$. We show that the minima in the E_i are attained by a directed version of it, up to the explicit error above.

Fix i and consider $T \in \mathcal{T}_i$ a spanning tree rooted at i. We have that, since any node $j \neq i$ is the origin of exactly one edge in T,

$$\sum_{(j\to k)\in T} B_{j,k} - \sum_{j\in\mathcal{I}} g_j = \sum_{(j\to k)\in T} \big(B_{j,k} - g_j\big) - g_i \\ \geq \sum_{(j\to k)\in T} C_{j,k} - g_i \,. \tag{H.97}$$

But by definition of T^*, this sum is at least greater than $\sum_{(j\leftrightarrow k)\in T^*} C_{j,k}$ so that we have

$$\sum_{(j\to k)\in T} B_{j,k} - \sum_{j\in\mathcal{I}} g_j \geq \sum_{(j\leftrightarrow k)\in T^*} C_{j,k} - g_i \,. \tag{H.98}$$

Conversely, taking T_i an oriented version of T^* rooted at i, the previous bound on the asymmetry yields

$$\sum_{(j\to k)\in T_i} B_{j,k} - \sum_{j\in\mathcal{I}} g_j \le \sum_{(j\leftrightarrow k)\in T^*} C_{j,k} - g_i + (|\mathcal{I}|-1)\Delta_{\mathcal{K}}\,. \tag{H.99}$$

Therefore, we have that, for some $r_i \in [0, (|\mathcal{I}|-1)\Delta_{\mathcal{K}}]$,

$$E_i = c_{\mathcal{K}} - g_i + r_i\,, \tag{H.100}$$

where

$$c_{\mathcal{K}} \coloneqq \sum_{(j\leftrightarrow k)\in T^*} C_{j,k} + \sum_{j\in\mathcal{I}} g_j \tag{H.101}$$

is independent of i. This concludes the proof. ∎

Let us now mention two particular cases.

- If σ^2 is constant, then $g = \frac{2f\circ Q}{\sigma^2}$, and

$$E_i = -\frac{2f\circ Q(\mathcal{K}_i)}{\sigma^2} + c_{\mathcal{K}} + r_i\,, \quad 0 \le r_i \le (|\mathcal{I}|-1)\Delta_{\mathcal{K}}\,. \tag{H.102}$$

- It σ^2 is linear, i.e., of the form

$$\sigma^2(f\circ Q(x)) = \sigma_1^2\Big(f\circ Q(x) + \sigma_0^2\Big)\,, \tag{H.103}$$

  then

$$g = \frac{2}{\sigma_1^2}\log\Big(f\circ Q + \sigma_0^2\Big)\,, \quad \text{and} \quad E_i = -\frac{2}{\sigma_1^2}\log\Big(f\circ Q + \sigma_0^2\Big)(\mathcal{K}_i) + c_{\mathcal{K}} + r_i\,, \quad 0 \le r_i \le (|\mathcal{I}|-1)\Delta_{\mathcal{K}}\,. \tag{H.104}$$

I Proofs of the large deviation principles

I.1 Preliminaries 96

I.2 Large deviation principle for interpolated trajectories 97

I.3 Large deviation principle for discrete trajectories 99

The goal of this section is to prove the large deviation principles for continuous and discrete trajectories that are stated in Appendix D.5. To facilitate the reading and avoid going back and forth between these parts, we will recall all necessary notation so that the present section can be read independently.

I.1. Preliminaries. From our dynamics of interest (SD), we define the sequences $(x_n)_{n\geq 0}$ and $(\omega_n)_{n\geq 0}$, we define three sequences: one discrete (in lowercase) and two continuous (in uppercase):

(a) The discrete "rescaled" trajectory

$$x_n^\eta \coloneqq x_{n\lfloor 1/\eta \rfloor}\,. \tag{I.1a}$$

(b) The continuous "interpolated" trajectory, defined for any $n \geq 0$, $t \in [\eta n, \eta(n+1)]$ by

$$X_t = x_n + \left(\frac{t}{\eta} - n\right)(x_{n+1} - x_n) \tag{I.1b}$$

(c) The continuous "discretized noise" trajectory, defined by $Z_0 = x_0$ and for any $t > 0$

$$\dot{Z}_t = F(Z_t) + U(Z_t, \omega_{\lfloor t/\eta \rfloor}) \tag{I.1c}$$

Note that all three sequences are "accelerated" by a factor $1/\eta$ compared to the original sequences appearing in (SD), as classically done in the large deviation literature (see [36, Chap. 7.1]).

Our final goal in this part is to establish a large deviation principle for the discrete rescaled sequence $(x_n^\eta)_{n\geq 0}$, i.e., Corollary D.2. To do so, we build upon Freidlin & Wentzell [36, Chap. 7] to obtain a large deviations principle on $(Z_t)_{t\geq 0}$ and then transfer it to the continuous interpolated $(X_t)_{t\geq 0}$ and finally to the discrete rescaled trajectory $(x_n^\eta)_{n\geq 0}$.

We equip the space of continuous functions $\mathcal{C}_T \coloneqq \mathcal{C}_T(\mathcal{V})$ with the distance induced by the uniform norm

$$\mathrm{dist}_{[0,T]}(\gamma, \gamma') \coloneqq \sup_{t\in[0,T]} \|\gamma(t) - \gamma'_t\|\,. \tag{I.2}$$

In order to use it as a proxy later, we will now bound the distance between the (continuous) "interpolated" trajectory and the "discretized noise" trajectory. To do so, we will first bound the latter.

Lemma I.1 (Growth of the trajectory)**.** *For any $x_0 \in \mathcal{V}$, $\eta > 0$, $t \geq 0$, we have $\|Z_t\| \leq e^{2Mt}(\|x_0\| + 2Mt)$.*

Proof. Using the definition of Z_t in (I.1c) and the growth condition (D.4) on F and U (see Assumption 3), we have that

$$\|\dot{Z}_t\| = \big\|F(Z_t) + U(Z_t, \omega_{\lfloor t/\eta \rfloor})\big\| \leq 2M(1 + \|Z_t\|)\,. \tag{I.3}$$

Hence, for any $t \geq 0$,

$$\|Z_t\| \leq \|y_0\| + \int_0^t 2M(1 + \|Z_s\|)\,ds = \|y_0\| + 2Mt + \int_0^t 2M\|Z_s\|\,ds\,. \tag{I.4}$$

Invoking Grönwall's lemma then yields the result. ∎

The following lemma states that the distance between the "interpolated" and "discretized noise" trajectories is bounded by a factor proportional to the stepsize η.

Lemma I.2. *Fix $\mathcal{K} \subset \mathcal{V}$ compact, $\eta_0 > 0$, $T > 0$. Then, there exists some constant $c = c(\mathcal{K}, \eta_0, T, F, U, \Omega) < +\infty$ such that, for any $x_0 \in \mathcal{K}$, $\eta \in (0, \eta_0]$,*

$$\mathrm{dist}_{[0,T]}(X, Z) \leq c\,\eta\,. \tag{I.5}$$

Proof. Before starting, notice that Lemmas D.1 and I.1 imply that there is some compact set $\mathcal{K}'$ that depends on $\mathcal{K}$, T, η_0 and M such that, for any $x_0 \in \mathcal{K}$, $\eta \in (0, \eta_0]$, $t \in [0,T]$, X_t and Z_t belong to $\mathcal{K}'$. In particular, U is therefore Lipschitz-continuous on $\mathcal{K}' \times \Omega$ so that, for any $\omega \in \Omega$, $x \mapsto F(x) + U(x,\omega)$ is Lipschitz-continuous on $\mathcal{K}'$ with constant L and bounded with constant B.

Let us now estimate the derivative of the interpolated trajectory X_t, which is piecewise differentiable by definition (see Eq. (I.1b)). For any $t \in [0,T]$ such that $t \in (\eta n, \eta(n+1))$ for some $n \geq 0$, we have that

$$\begin{aligned}
&\left\|\dot{X}_t - \big(F(X_t) + U(X_t, \omega_{\lfloor t/\eta \rfloor})\big)\right\| \\
&\quad = \left\|\frac{x_{n+1} - x_n}{\eta} - \big(F(X_t) + U(X_t, \omega_{\lfloor t/\eta \rfloor})\big)\right\| \\
&\quad = \left\|\big(-F(X_{\eta n}) + U(X_{\eta n}, \omega_n)\big) - \big(F(X_t) + U(X_t, \omega_{\lfloor t/\eta \rfloor})\big)\right\| \\
&\quad \leq L\|X_{\eta n} - X_t\|
\end{aligned} \tag{I.6}$$

where we used the L-Lipschitz-continuity of $x \mapsto F(x) + U(x,\omega)$ on $\mathcal{K}'$ uniformly in $\omega \in \Omega$. Since this map is also bounded by B, we have that

$$\left\|\dot{X}_t - \big(F(X_t) + U(X_t, \omega_{\lfloor t/\eta \rfloor})\big)\right\| \leq L\|X_{\eta n} - X_t\| \leq L\|x_n - x_{n+1}\| \leq LB\eta\,. \tag{I.7}$$

Moreover, the L-Lipschitz-continuity of $x \mapsto F(x) + U(x,\omega)$ also gives us that

$$\begin{aligned}
&\|\dot{Z}_t - \big(F(X_t) + U(X_t, \omega_{\lfloor t/\eta \rfloor})\big)\| \\
&\quad = \|-F(Z_t) + U(Z_t, \omega_{\lfloor t/\eta \rfloor}) - \big(F(X_t) + U(X_t, \omega_{\lfloor t/\eta \rfloor})\big)\| \\
&\quad \leq L\|Z_t - X_t\|\,.
\end{aligned} \tag{I.8}$$

Putting everything together and integrating $X_t - Z_t$ from 0 to t yields (since X_t is absolutely continuous), for any $t \in [0,T]$,

$$\|X_t - Z_t\| \leq \|X_0 - y_0\| + \int_0^t \|\dot{X}_s - \dot{Z}_s\|\,ds \leq LB\eta + L\int_0^t \|X_s - Z_s\|\,ds \tag{I.9}$$

where we used that $y_0 = X_0 = x_0$ by definition of the trajectories. Finally, Grönwall's lemma then yields the result. ∎

I.2. Large deviation principle for interpolated trajectories. From the Lagrangian defined in Eq. (D.19), we define, on $\mathcal{C}_T = \mathcal{C}_T(\mathcal{V})$, the normalized action functional $\mathcal{S}_{[0,T]}$ as

$$\mathcal{S}_{[0,T]}(\gamma) = \begin{cases} \int_0^T \mathcal{L}(\gamma(t), \dot{\gamma}(t))\,dt & \text{if } \gamma \text{ is absolutely continuous,} \\ \infty & \text{otherwise,} \end{cases} \tag{I.10}$$

following Freidlin & Wentzell [36, Chap. 3.2], as a manner to quantify how "probable" a trajectory is.

We first show that the set

$$\Gamma_{[0,T]}(\mathcal{K}; s) := \{\gamma \in \mathcal{C}_T(\mathcal{V}) : \gamma(0) \in \mathcal{K}, \mathcal{S}_{[0,T]}(\gamma) \leq s\} \tag{I.11}$$

of trajectories with bounded action functional is compact and $\mathcal{S}_{[0,T]}$ is lower semi-continuous.

Lemma I.3. *Fix $T > 0$. For any $\mathcal{K} \subset \mathcal{V}$ compact, $s \geq 0$, the set $\Gamma_{[0,T]}(\mathcal{K}; s)$ is compact and $\mathcal{S}_{[0,T]}$ is l.s.c. on $\mathcal{C}_T(\mathcal{V})$.*

Proof. Let us first check that $\mathcal{S}_{[0,T]}$ is l.s.c. on $\mathcal{C}_T(\mathcal{V})$ by applying Ioffe et al. [59, §9.1.4, Thm. 3]: $(t,x,w) \mapsto \mathcal{L}(x,w)$ is a normal integrand since $(x,w) \mapsto \mathcal{L}(x,w)$ is l.s.c. by construction, quasiregular since $\mathcal{L}(x,\cdot)$ is convex for any $x \in \mathcal{V}$ and satisfies the growth condition because it is non-negative (see Lemma D.3). Hence, $\mathcal{S}_{[0,T]}$ is l.s.c. on $\mathcal{C}_T(\mathcal{V})$.

The compactness of $\Gamma_{[0,T]}(\mathcal{K}; s)$ follows from the idea of proof as Freidlin & Wentzell [36, Chap. 7, Lemma 4.2] but with the added difficulty that the vector field F and the noise U are not uniformly bounded. Take $\gamma \in \Gamma_{[0,T]}(\mathcal{K}; s)$. Since $\mathcal{S}_{[0,T]}(\gamma) \leq s < +\infty$, it means that $\mathcal{L}(\gamma(t), \dot{\gamma}(t)) < +\infty$ for almost every t so that by Lemma D.3, almost everywhere,

$$\|\dot{\gamma}(t)\| \leq 2M(1 + \|\gamma(t)\|)\,. \tag{I.12}$$

Grönwall's lemma then yields that, for any $t \in [0,T]$,

$$\|\gamma(t)\| \leq e^{2MT}(\|\gamma(0)\| + 2MT) \tag{I.13}$$

which is bounded uniformly in $\gamma \in \Gamma_{[0,T]}(\mathcal{K}; s)$ since $\gamma(0)$ is in $\mathcal{K}$ compact. Hence, $\|\dot{\gamma}(t)\|$ is also uniformly bounded by Eq. (I.12). Therefore, the functions in $\Gamma_{[0,T]}(\mathcal{K}; s)$ are equicontinuous and uniformly bounded and, $\Gamma_{[0,T]}(\mathcal{K}; s)$ is closed by l.s.c. of $\mathcal{S}_{[0,T]}$, so that, by the Arzelà-Ascoli theorem, $\Gamma_{[0,T]}(\mathcal{K}; s)$ is compact. ∎

The following result establishes the fact that the functional $\eta^{-1}\mathcal{S}_{[0,T]}$ is the action functional in $\mathcal{C}([0,T], \mathcal{V})$ of the interpolated process $(X_t)_{t\in[0,T]}$ of the algorithm started at x_0, uniformly with respect to the initial point x_0 in any compact set $\mathcal{K} \subset \mathcal{V}$, as $\eta \to 0$. This enables to build on Freidlin & Wentzell [36, Chap. 7, Thm. 4.1′] to provide a large deviation principle for the interpolated trajectory $(X_t)_{t\in[0,T]}$, meaning that for η small enough, it will be (a) close to probable trajectories with a probability exponentially big in their action functional; and (b) far from the most probable trajectories with a probability exponentially small in their action value.

Proposition I.1. *Fix $T > 0$. For any $s, \delta, \epsilon > 0$, $\mathcal{K} \subset \mathcal{V}$ compact, there exists $\eta_0 > 0$ such that, for any $\eta \in (0, \eta_0]$, for any $x_0 \in \mathcal{K}$, we have that*

$$\mathbb{P}_{x_0}\left(\mathrm{dist}_{[0,T]}(X, \gamma) < \delta\right) \geq \exp\left(-\frac{\mathcal{S}_{[0,T]}(\gamma) + \epsilon}{\eta}\right) \tag{I.14a}$$

$$\mathbb{P}_{x_0}\left(\mathrm{dist}_{[0,T]}(X, \Gamma_{[0,T]}(\{x_0\}; s)) > \delta\right) \leq \exp\left(-\frac{s - \epsilon}{\eta}\right) \tag{I.14b}$$

for all $\gamma \in \Gamma_{[0,T]}(\{x_0\}; s)$.

Proof. Our strategy is to apply Freidlin & Wentzell [36, Chap. 7, Thm. 4.1′] to the process $(Z_t)_{t\in[0,T]}$ and use Lemma I.2 to transfer the result to $(X_t)_{t\in[0,T]}$ (both starting from x_0). However, this theorem requires $b(x,\omega) \coloneqq F(x) + U(x,\omega)$ to be uniformly bounded as well as its derivative. To avoid this issue, note that, for a fixed compact set $\mathcal{K}$, $s > 0$, and all $\eta \leq 1$, the trajectories of $\Gamma_{[0,T]}(\mathcal{K}; s)$ and of the process $(Z_t)_{t\in[0,T]}$ are contained in a compact set $\mathcal{K}'$ by the compactness of $\Gamma_{[0,T]}(\mathcal{K}; s)$ and Lemma I.1.

Now, consider $b' \colon \mathcal{V} \times \mathcal{V}^* \to \mathcal{V}$ twice differentiable that coincides with b on $\mathcal{K}' \times \Omega$ but that is uniformly bounded along with its derivative. Define $\mathcal{H}'(x,w) = \log \mathbb{E}[\exp(\langle w, b'(x,\omega)\rangle)]$ which is still differentiable and satisfies "Condition F" of Freidlin & Wentzell [36, Chap. 7.4] by Freidlin & Wentzell [36, Chap. 7, Lem. 4.3] and i.i.d. assumption. Hence, Freidlin & Wentzell [36, Chap. 7, Thm. 4.1′] yields that, with $\mathcal{K}$ and s fixed above, for any $\delta, \epsilon > 0$, there exists $\eta_0 > 0$ such that, for any $\eta \in (0, \eta_0]$, for any $x_0 \in \mathcal{K}$, $\gamma \in \Gamma_{[0,T]}(\{x_0\}; s)$,

$$\mathbb{P}_{x_0}\left(\mathrm{dist}_{[0,T]}(Z, \gamma) < \delta\right) \geq \exp\left(-\frac{\mathcal{S}_{[0,T]}(\gamma) + \epsilon}{\eta}\right) \tag{I.15a}$$

$$\mathbb{P}_{x_0}\left(\mathrm{dist}_{[0,T]}(Z, \Gamma_{[0,T]}(\{x_0\}; s)) > \delta\right) \leq \exp\left(-\frac{s - \epsilon}{\eta}\right). \tag{I.15b}$$

To obtain the result for the process $(X_t)_{t\in[0,T]}$, fix $\delta > 0$. By Lemma I.2, there exists $\eta_0 > 0$ such that, for any $\eta \in (0, \eta_0]$, for any initial point $x_0 \in \mathcal{K}$, $\mathrm{dist}_{[0,T]}(X, Z) < \delta$. Combining this with the previous result on $(Z_t)_{t\in[0,T]}$ yields that, for any $\delta, \epsilon > 0$, there exists $\eta_0 > 0$ such that, for any $\eta \in (0, \eta_0]$, for any $x_0 \in \mathcal{K}$, $\gamma \in \Gamma_{[0,T]}(\{x_0\}; s)$,

$$\mathbb{P}_{x_0}\left(\mathrm{dist}_{[0,T]}(X, \gamma) < 2\delta\right) \geq \exp\left(-\frac{\mathcal{S}_{[0,T]}(\gamma) + \epsilon}{\eta}\right) \tag{I.16a}$$

and

$$\mathbb{P}_{x_0}\left(\mathrm{dist}_{[0,T]}(X, \Gamma_{[0,T]}(\{x_0\}; s)) > 2\delta\right) \leq \exp\left(-\frac{s - \epsilon}{\eta}\right) \tag{I.16b}$$

which concludes the proof. ∎

I.3. Large deviation principle for discrete trajectories. We now leverage the previous section to show a LDP for the discrete rescaled trajectories $(x_n^\eta)_{n\geq 0}$ defined in Eq. (I.1a). To do so, we will use the results of the previous section by considering a finite number of points $(X_n)_{n\geq 0}$ from the continuous interpolation.

For some $T > 0$, we will first equip $\mathcal{V}^T$ with the distance

$$\mathrm{dist}_T(\xi, \zeta) = \max_{0\leq n\leq T-1} \|\xi_n - \zeta_n\| \tag{I.17}$$

and bound the difference between the discrete rescaled trajectory and the continuous interpolation.

Lemma I.4. *Fix $\mathcal{K} \subset \mathcal{V}$ compact, $\eta_0 > 0$, $T \geq 1$. Then, there exists some constant $c = c(\mathcal{K}, \eta_0, T, F, U, \Omega) < +\infty$ such that, for any $x_0 \in \mathcal{K}$, $\eta \in (0, \eta_0]$,*

$$\mathrm{dist}_T(x^\eta, (X_n)_{0\leq n\leq T-1}) = \max_{0\leq n\leq T-1} \|x_n^\eta - X_n\| \leq c\eta\,. \tag{I.18}$$

Proof. By Lemma D.1, there is some compact set $\mathcal{K}'$ that depends on $\mathcal{K}$, T, η_0 and M such that, for any $x_0 \in \mathcal{K}$, $\eta \in (0, \eta_0]$, $n \leq \lceil 1/\eta \rceil T$, x_n belongs to $\mathcal{K}'$. In particular, for almost every $t \in [0, T-1]$, we have that

$$\dot{X}_t = \frac{x_{\lfloor t/\eta \rfloor + 1} - x_{\lfloor t/\eta \rfloor}}{\eta} = F(x_{\lfloor t/\eta \rfloor}) + U(x_{\lfloor t/\eta \rfloor}, \omega_{\lfloor t/\eta \rfloor}) \tag{I.19}$$

with $x_{\lfloor t/\eta \rfloor}$ belonging to $\mathcal{K}'$ since $\left\lfloor \frac{t}{\eta} \right\rfloor \leq \left\lfloor \frac{T-1}{\eta} \right\rfloor \leq \left\lceil \frac{T}{\eta} \right\rceil$. Hence, the norm of $\dot{X}_t$ is bounded by some constant B for almost every $t \in [0, T-1]$ uniformly in $x_0 \in \mathcal{K}$, $\eta \in (0, \eta_0]$ by the growth condition of Eq. (D.4).

Therefore, for any $0 \leq n \leq T-1$, since $x_n^\eta = x_{n\lfloor 1/\eta \rfloor} = X_{n\eta\lfloor 1/\eta \rfloor}$, we have that

$$\|x_n^\eta - X_n\| \leq Bn|1 - \eta\lfloor 1/\eta \rfloor| \leq BT\eta \tag{I.20}$$

which concludes the proof. ∎

Now, for $T \geq 1$, $\xi = (\xi_0, \dots, \xi_{T-1}) \in \mathcal{V}^T$, let us define the normalized discrete action functional

$$\mathcal{A}_T(\xi) \coloneqq \sum_{n=0}^{T-2} \rho(\xi_n, \xi_{n+1}), \quad \mathcal{A}_1(\xi) \coloneqq 0, \tag{I.21}$$

where the cost of moving from one iteration to the next is defined for any $x, x' \in \mathcal{V}$ from the previous continuous normalized action functional (cf. Eq. (I.10)) with horizon 1 as

$$\rho(x, x') \coloneqq \inf\{\mathcal{S}_{0,1}(\gamma) : \gamma \in \mathcal{C}_1(\mathcal{V}), \gamma(0) = x, \gamma(1) = x'\}\,. \tag{I.22}$$

Again, we show that the set of discrete trajectories with low action functional

$$\Gamma_T(\mathcal{K}; s) \coloneqq \{\xi \in \mathcal{V}^T : \xi_0 \in \mathcal{K}, \mathcal{A}_T(\xi) \leq s\} \tag{I.23}$$

is compact and $\mathcal{A}_T$ is l.s.c.

Lemma I.5. *Fix $T \geq 0$. For any $\mathcal{K} \subset \mathcal{V}$ compact, $s \geq 0$, the set $\Gamma_T(\mathcal{K}; s)$ is compact and $\mathcal{A}_T$ is l.s.c. on $\mathcal{V}^T$.*

Proof. First, let us show that, for $s \geq 0$, $\mathcal{K} \subset \mathcal{V}$, the set $\{(x, x') \in \mathcal{K} \times \mathcal{V} : \rho(x, x') \leq s\}$ is compact. Let $(x^k, x'^k)_{k\geq 0}$ be a sequence in $\mathcal{K} \times \mathcal{V}$ such that $\rho(x^k, x'^k) \leq s$ for all $k \geq 0$. Then, for any $k \geq 0$, there exists $\gamma_k \in \mathcal{C}_1(\mathcal{V})$ such that $\gamma_k(0) = x^k$, $\gamma_k(1) = x'^k$ and $\mathcal{S}_{0,1}(\gamma_k) \leq \rho(x^k, x'^k) + (1+k)^{-1}$. In particular, for any $k \geq 0$, γ_k belongs to $\{\gamma \in \mathcal{C}_1(\mathcal{V}) : \gamma(0) \in \mathcal{K}, \mathcal{S}_{0,1}(\gamma) \leq s + 1\}$ which is compact by Lemma I.3. Hence, there exists a subsequence that converges uniformly to some $\gamma \in \mathcal{C}_1(\mathcal{V})$. Without loss of generality, assume that $\gamma_k \to \gamma$ as $k \to \infty$. In particular, (x^k, x'^k) converges to $(\gamma(0), \gamma(1))$ as $k \to \infty$ with $\gamma(0)$ belonging to $\mathcal{K}$. Then, by l.s.c. of $\mathcal{S}_{0,1}$, we have that $\mathcal{S}_{0,1}(\gamma) \leq \liminf_{k\to\infty} \rho(x^k, x'^k) \leq s$ so that $\rho(\gamma(0), \gamma(1)) \leq s$. Therefore, we have shown that $\{(x, x') \in \mathcal{K} \times \mathcal{V} : \rho(x, x') \leq s\}$ is compact.

As a consequence ρ is l.s.c. on $\mathcal{V} \times \mathcal{V}$: indeed, for any $s \geq 0$, any convergent sequence of points of $\{(x, x') \in \mathcal{V} \times \mathcal{V} : \rho(x, x') \leq s\}$ must be included in $\{(x, x') \in \mathcal{K} \times \mathcal{V} : \rho(x, x') \leq s\}$ for some

$\mathcal{K} \subset \mathcal{V}$ compact, which is closed and therefore contains the limit point of any such sequences. $\mathcal{A}_T$ is then immediately l.s.c. on $\mathcal{V}^T$.

Finally, the compactness of $\Gamma_T(\mathcal{K}; s)$ follows by induction on T. For $T = 1$, $\Gamma_1(\mathcal{K}; s) = \{(x, x') \in \mathcal{V} \times \mathcal{V} : \rho(x, x') \leq s\}$ which is compact by the previous argument for any compact set $\mathcal{K}$. Now, assume that $\Gamma_T(\mathcal{K}; s)$ is compact for some $T \geq 1$. As a consequence, the set

$$\mathcal{K}' := \{x \in \mathcal{V} : \exists (\xi_0, \dots, \xi_{T-2}) \in \mathcal{V}^{T-1}, (\xi_0, \dots, \xi_{T-2}, x) \in \Gamma_T(\mathcal{K}; s)\} \tag{I.24}$$

is compact as well. Hence, $\Gamma_{T+1}(\mathcal{K}; s)$ is included in the product of compact sets $\Gamma_{T-1}(\mathcal{K}; s) \times \Gamma_1(\mathcal{K}'; s)$ and is therefore bounded. Moreover, $\Gamma_{T+1}(\mathcal{K}; s)$ is closed by l.s.c. of $\mathcal{A}_T$ and therefore compact. This concludes the proof by induction and the proof of the lemma. ∎

We will first provide a discrete analogue of Proposition I.1 for the interpolated trajectory at times $n = 0, \dots, T-1$, with $\eta^{-1}\mathcal{A}_T$ the action functional in $\mathcal{V}^T$ of the process $(X_n)_{0 \leq n \leq T-1}$, uniformly with respect to the starting point x_0 in any compact set $\mathcal{K} \subset \mathcal{V}$, as $\eta \to 0$.

Proposition I.2. *Fix $T \geq 0$. For any $s, \delta, \epsilon > 0$, $\mathcal{K} \subset \mathcal{V}$ compact, there exists $\eta_0 > 0$ such that, for any $\eta \in (0, \eta_0]$, for any $x_0 \in \mathcal{K}$, $\xi \in \Gamma_T(\{x_0\}; s)$, we have that*

$$\mathbb{P}_{x_0}(\mathrm{dist}_T((X_n)_{0 \leq n \leq T-1}, \xi) < \delta) \geq \exp\left(-\frac{\mathcal{A}_T(\xi) + \epsilon}{\eta}\right) \tag{I.25a}$$

and

$$\mathbb{P}_{x_0}(\mathrm{dist}_T((X_n)_{0 \leq n \leq T-1}, \Gamma_T(\{x_0\}; s)) > \delta) \leq \exp\left(-\frac{s - \epsilon}{\eta}\right). \tag{I.25b}$$

for all $\xi \in \Gamma_T(\{x_0\}; s)$

Proof. Invoke the first part of Proposition I.1 with $T \leftarrow T - 1$ and $s \leftarrow s + \epsilon$. There exists $\eta_0 > 0$ such that, for any $\eta \in (0, \eta_0]$, for any $x_0 \in \mathcal{K}$, $\gamma \in \Gamma_{[0,T]}(\{x_0\}; s + \epsilon)$,

$$\mathbb{P}_{x_0}\left(\mathrm{dist}_{0,T-1}(X, \gamma) < \delta\right) \geq \exp\left(-\frac{\mathcal{S}_{0,T-1}(\gamma) + \epsilon}{\eta}\right). \tag{I.26}$$

Take $\eta \in (0, \eta_0]$, $x_0 \in \mathcal{K}$, $\xi \in \Gamma_T(\{x_0\}; s)$. Then, there exists $\gamma \in \Gamma_{[0,T]}(\{x_0\}; s + \epsilon)$, for any $0 \leq n \leq T - 1$, $\gamma(n) = \xi_n$, cf. Eq. (I.22). Hence,

$$\begin{aligned}
\mathbb{P}_{x_0}(\mathrm{dist}_T((X_n)_{0 \leq n \leq T-1}, \xi) < \delta) &\geq \mathbb{P}_{x_0}\left(\mathrm{dist}_{0,T-1}(X, \gamma) < \delta\right) \\
&\geq \exp\left(-\frac{\mathcal{S}_{0,T-1}(\gamma) + \epsilon}{\eta}\right) \\
&\geq \exp\left(-\frac{\mathcal{A}_T(\xi) + 2\epsilon}{\eta}\right)
\end{aligned} \tag{I.27}$$

which proves the first part of the result.

For the second part, we have similarly from Proposition I.1 with $T \leftarrow T - 1$ and $\delta \leftarrow \delta/2$ that for any $\eta \in (0, \eta_0]$, for any $x_0 \in \mathcal{K}$,

$$\mathbb{P}_{x_0}\left(\mathrm{dist}_{0,T-1}(X, \Gamma_{[0,T-1]}(\{x_0\}; s)) > \delta/2\right) \leq \exp\left(-\frac{s - \epsilon}{\eta}\right). \tag{I.28}$$

Now, note that if $\mathrm{dist}_{0,T-1}(X, \Gamma_{[0,T-1]}(\{x_0\}; s)) < \delta$, then there must exist $\gamma \in \Gamma_{[0,T-1]}(\{x_0\}; s)$ such that $\mathrm{dist}_{0,T-1}(X, \gamma) \leq \delta$. Consider the discrete path $\xi \in \mathcal{V}^T$ defined by $\xi_n = \gamma(n)$ for $0 \leq n \leq T - 1$. Then, by construction, $\mathcal{A}_T(\xi) \leq \mathcal{S}_{0,T-1}(\gamma) \leq s$ and $\mathrm{dist}_T((X_n)_{0 \leq n \leq T-1}, \xi) \leq \delta$. Thus,

$$\mathrm{dist}_{0,T-1}(X, \Gamma_{[0,T-1]}(\{x_0\}; s)) < \delta \implies \mathrm{dist}_T((X_n)_{0 \leq n \leq T-1}, \Gamma_T(\{x_0\}; s)) \leq \delta\,. \tag{I.29}$$

Putting all together, we have that

$$\begin{aligned}
\mathbb{P}_{x_0}(\mathrm{dist}_T((X_n)_{0 \leq n \leq T-1}, \Gamma_T(\{x_0\}; s)) > \delta) &\leq \mathbb{P}_{x_0}\left(\mathrm{dist}_{0,T-1}(X, \Gamma_{[0,T-1]}(\{x_0\}; s) \geq \delta)\right) \\
&\leq \mathbb{P}_{x_0}\left(\mathrm{dist}_{0,T-1}(X, \Gamma_{[0,T-1]}(\{x_0\}; s) > \delta/2)\right) \\
&\leq \exp\left(-\frac{s - \epsilon}{\eta}\right)
\end{aligned} \tag{I.30}$$

which concludes our proof. ∎

Finally, we end up with a large deviation principle on the discrete rescaled iterates $(x_n^\eta)_{0\le n\le T-1} = (x_{n\lfloor \eta^{-1}\rfloor})_{0\le n}$ by leveraging Lemma I.4. In the following result, the functional $\eta^{-1}\mathcal{A}_T$ is thus the action functional in $\mathcal{V}^T$ of the process $(x_n^\eta)_{0\le n\le T-1}$ uniformly with respect to the starting point x_0 in any compact set $\mathcal{K}\subset\mathcal{V}$, as $\eta\to 0$.

Corollary I.1. *Fix $T\ge 0$. For any $s,\delta,\epsilon>0$, $\mathcal{K}\subset\mathcal{V}$ compact, there exists $\eta_0>0$ such that, for any $\eta\in(0,\eta_0]$, for any $x_0\in\mathcal{K}$, $\xi\in\Gamma_T(\{x_0\};s)$, we have that*

$$\mathbb{P}_{x_0}(\mathrm{dist}_T(x^\eta,\xi)<\delta)\ge\exp\left(-\frac{\mathcal{A}_T(\xi)+\epsilon}{\eta}\right) \tag{I.31a}$$

and

$$\mathbb{P}_{x_0}(\mathrm{dist}_T(x^\eta,\Gamma_T(\{x_0\};s))>\delta)\le\exp\left(-\frac{s-\epsilon}{\eta}\right). \tag{I.31b}$$

for all $\xi\in\Gamma_T(\{x_0\};s)$.

Proof. Fix $\delta>0$. Choose η_0 such that both Proposition I.2 and Lemma I.4 hold with $c\eta\le\delta$. Then for any $\eta\in(0,\eta_0]$, for any $x_0\in\mathcal{K}$, $\xi\in\Gamma_T(\{x_0\};s)$, we have that

$$\mathbb{P}_{x_0}(\mathrm{dist}_T((X_n)_{0\le n\le T-1},\xi)<\delta)\ge\exp\left(-\frac{\mathcal{A}_T(\xi)+\epsilon}{\eta}\right). \tag{I.32}$$

Now, if $\mathrm{dist}_T((X_n)_{0\le n\le T-1},\xi)<\delta$ and $\mathrm{dist}_T(x^\eta,(X_n)_{0\le n\le T-1})\le\delta$, then $\mathrm{dist}_T(x^\eta,\xi)<2\delta$. Thus,

$$\mathbb{P}_{x_0}(\mathrm{dist}_T(x^\eta,\xi)<2\delta)\ge\mathbb{P}_{x_0}(\mathrm{dist}_T((X_n)_{0\le n\le T-1},\xi)<\delta)\ge\exp\left(-\frac{\mathcal{A}_T(\xi)+\epsilon}{\eta}\right). \tag{I.33}$$

The second part can be obtained similarly. ∎

J Numerical illustrations

In this appendix, we provide some details about the setup of Figs. 1 and 2 in the main body of the paper.

The games we run were for Fig. 1 where 2-player 1-dimensional games defined as follows:

- Fig. 1, left: An "almost bilinear" zero-sum game with a spurious attractor and min-max gradient field given by

$$v(x) = \left(x_2,\ -x_1 + \varepsilon x_2 - \varepsilon x_2^3\right) \tag{J.1}$$

- Fig. 1, right: An "almost bilinear" zero-sum game with a forsaken equilibrium and min-max gradient field given by

$$v(x) = \left(x_2 - \frac{\varepsilon}{2}x_1 + 2\varepsilon x_1^3 - \varepsilon x_1^5,\ -x_1 - \frac{\varepsilon}{2}x_2 + 2\varepsilon x_2^3 - \varepsilon x_2^5\right) \tag{J.2}$$

Both games were taken from Hsieh et al. [56], and our experiments were run with $\varepsilon = 0.1$. The regularization weight was set to $\lambda = 0.05$, and we used a step-size of $\eta = 0.1$, and ran the algorithm for 10^3 iterations over 10^4 randomly drawn initializations.

For Fig. 2, we ran a standard game of Chicken and a coordination game with risk-dominant payoff of 1 and payoff-dominant payoff of 1.1, and the same parameters as above.

NeurIPS Paper Checklist

1. **Claims**

   Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

   Answer: [Yes]

   Justification: The claims made in the abstract and introduction correspond to the main results of the paper, presented in Section 4.

   Guidelines:

   - The answer [N/A] means that the abstract and introduction do not include the claims made in the paper.
   - The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A [No] or [N/A] answer to this question will not be perceived well by the reviewers.
   - The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
   - It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. **Limitations**

   Question: Does the paper discuss the limitations of the work performed by the authors?

   Answer: [Yes]

   Justification: The limitations of our assumptions are discussed in the main text and in Appendix D.

   Guidelines:

   - The answer [N/A] means that the paper has no limitation while the answer [No] means that the paper has limitations, but those are not discussed in the paper.
   - The authors are encouraged to create a separate "Limitations" section in their paper.
   - The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
   - The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
   - The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
   - The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
   - If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
   - While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. **Theory assumptions and proofs**

   Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

   Answer: [Yes]

   Justification: The full set of assumptions can be found in the main (Sections 2 and 3) while the proofs are in the appendices, see Appendix A for an overview of the structure of the appendices.

Guidelines:

- The answer [N/A] means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. **Experimental result reproducibility**

   Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

   Answer: [N/A]

   Justification: This paper does not include experiments, only illustrative examples.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - If the paper includes experiments, a [No] answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
   - If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
   - Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general. releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
   - While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
     (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
     (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
     (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
     (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. **Open access to data and code**

   Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

   Answer: [N/A]

   Justification: This paper does not include experiments, only illustrative examples.

   Guidelines:

   - The answer [N/A] means that paper does not include experiments requiring code.

- Please see the NeurIPS code and data submission guidelines (https://neurips.cc/public/guides/CodeSubmissionPolicy) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so [No] is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (https://neurips.cc/public/guides/CodeSubmissionPolicy) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. **Experimental setting/details**

   Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer) necessary to understand the results?

   Answer: [N/A]

   Justification: This paper does not include experiments, only illustrative examples.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
   - The full details can be provided either with the code, in appendix, or as supplemental material.

7. **Experiment statistical significance**

   Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

   Answer: [N/A]

   Justification: This paper does not include experiments, only illustrative examples.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - The authors should answer [Yes] if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
   - The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
   - The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
   - The assumptions made should be given (e.g., Normally distributed errors).
   - It should be clear whether the error bar is the standard deviation or the standard error of the mean.
   - It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
   - For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g., negative error rates).
   - If error bars are reported in tables or plots, the authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. **Experiments compute resources**

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [N/A]

Justification: This paper does not include experiments, only illustrative examples.

Guidelines:

- The answer [N/A] means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. **Code of ethics**

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics https://neurips.cc/public/EthicsGuidelines?

Answer: [Yes]

Justification: This research conforms with the NeurIPS Code of Ethics.

Guidelines:

- The answer [N/A] means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer [No], they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. **Broader impacts**

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [N/A]

Justification: This work is foundational research and does not have any direct societal impacts.

Guidelines:

- The answer [N/A] means that there is no societal impact of the work performed.
- If the authors answer [N/A] or [No], they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate Deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. **Safeguards**

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pre-trained language models, image generators, or scraped datasets)?

Answer: [N/A]

Justification: This research does not involve the release of data or models.

Guidelines:

- The answer [N/A] means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. **Licenses for existing assets**

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [N/A]

Justification: This research does not use existing assets.

Guidelines:

- The answer [N/A] means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, `paperswithcode.com/datasets` has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. **New assets**

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [N/A]

Justification: This research does not introduce new assets.

Guidelines:

- The answer [N/A] means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and research with human subjects**

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [N/A]

Justification: This research does not involve crowdsourcing nor research with human subjects.

Guidelines:

- The answer [N/A] means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional review board (IRB) approvals or equivalent for research with human subjects**

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [N/A]

Justification: This research does not involve crowdsourcing nor research with human subjects.

Guidelines:

- The answer [N/A] means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. **Declaration of LLM usage**

Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does *not* impact the core methodology, scientific rigor, or originality of the research, declaration is not required.

Answer: [N/A]

Justification: This research does not involve LLMs as any important, original, or non-standard components.

Guidelines:

- The answer [N/A] means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
- Please refer to our LLM policy in the NeurIPS handbook for what should or should not be described.