

Metastability of Logit Dynamics for Coordination Games*

Vincenzo Auletta¹, Diodato Ferraioli¹, Francesco Pasquale², and Giuseppe Persiano¹

¹Università di Salerno

²Università di Roma “Tor Vergata”

Abstract

Logit Dynamics [Blume, Games and Economic Behavior, 1993] are randomized best response dynamics for strategic games: at every time step a player is selected uniformly at random and she chooses a new strategy according to a probability distribution biased toward strategies promising higher payoffs. This process defines an ergodic Markov chain, over the set of strategy profiles of the game, whose unique stationary distribution is the long-term equilibrium concept for the game. However, when the mixing time of the chain is large (e.g., exponential in the number of players), the stationary distribution loses its appeal as equilibrium concept, and the transient phase of the Markov chain becomes important. It can happen that the chain is “metastable”, i.e., on a time-scale shorter than the mixing time, it stays close to some probability distribution over the state space, while in a time-scale multiple of the mixing time it jumps from one distribution to another.

In this paper we give a quantitative definition of “metastable probability distributions” for a Markov chain and we study the metastability of the logit dynamics for some classes of coordination games. We first consider a pure n -player coordination game that highlights the distinctive features of our metastability notion based on distributions. Then, we study coordination games on the clique without a risk-dominant strategy (which are equivalent to the well-known Glauber dynamics for the Curie-Weiss model) and coordination games on a ring (both with and without risk-dominant strategy).

*A preliminary version of this paper appeared in the Proceedings of the Twenty-Third Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'12) [4]. Part of this work was done while the third author was at Università di Salerno

1 Introduction

Complex systems consist of a large number of components that interact according to simple rules at small scale and, despite of this, exhibit complex large scale behaviors. Complex systems are ubiquitous and some examples can be found in Economics (e.g., the market), Physics (e.g., ideal gases, spin systems), Biology (e.g., evolution of life) and Computer Science (e.g., Internet and social networks). Analyzing, understanding how such systems evolve, and predicting their long-term behaviour is a major research endeavor.

In this paper we focus on *selfish* systems in which the components (called the *players*) are selfish agents, each one with a set of possible actions or *strategies* trying to maximize her own payoff. The payoff obtained by each player depends not only on her decision but also on the decisions of the other players. We study specific dynamics, the *logit* dynamics (first studied by Blume [10]) and consider as solution concept their equilibrium states. Logit dynamics are a type of *noisy* best response dynamics that model in a clean and tractable way the limited knowledge (or bounded rationality) of the players in terms of a parameter β (in similar models studied in Physics, β is the inverse of the temperature). Intuitively, a low value of β (that is, high temperature and entropy) represents the situation where players choose their strategies “nearly at random”; a high value of β (that is, low temperature and entropy) represents the situation where players pick the strategies yielding high payoff with higher probability. It is well known that, for every strategic game, these dynamics induce a Markov chain with a *unique* stationary distribution (the Markov chain is ergodic). Thus no equilibrium selection problem arises. The drawback of using the stationary distribution to describe the system behavior is that the system may take too long to converge to that distribution, unless the chain is rapidly mixing. Logit dynamics for strategic games can be rapidly mixing or *not*, depending on the features of the underlying game, the temperature/noise, and the number of players [6, 7]. For this reason, in this work we focus on the *transient phase* of the logit dynamics and, in particular, we try to answer the following questions for games where logit dynamics have exponential mixing time: is the transient phase completely *chaotic* or can we spot some *regularities* even at time-scales shorter than mixing time? Ellison [19] showed that as long as the probability of suboptimal responses tends to 0 (i.e., $\beta \rightarrow \infty$), the logit dynamics (and many other similar noisy dynamics) usually evolve via a series of gradual steps between nearly stable states. However, how can we describe these stable states? And what happens for lower values of β ? Obviously, at a fine-grained level any Markov chain is perfectly described by the collections of probability distributions consisting of one distribution for each time step and each starting profile. However, this should be contrasted with the rapidly mixing case (i.e., a Markov chain with polynomial mixing time) in which one can approximately describe the state of the system (after the mixing time) using *one* distribution, i.e. the stationary distribution.

Our results show that there are games for which regularities can be observed even in the transient phase of the logit dynamics. In particular, we will show that it is often possible to identify a few probability distributions (the *metastable* distributions) such that, depending on the starting profile, the dynamics quickly reach one of those distributions and remain close to that one for a long time. We can describe our results also in terms of the quantity of information needed to predict the status of a system that evolves in time according to the logit dynamics. We know that the long-term behavior of the system can be compactly described in terms of a unique distribution but we have to wait a transient phase of length equal to the mixing time. Thus, if the system is rapidly mixing this description is significant after a short transient phase. However, when the mixing time is super-polynomial this description becomes significant only after a long time. Our results show that for a large class of n -player games whose logit dynamics are not rapidly mixing, the behaviour of the dynamics (the strategies played by the n players) can still be described with good approximation and for a super-polynomial number of steps by means of a small number of probability distributions. This comes at the price of sacrificing a short polynomial initial transient phase (so far we are on a par with the rapidly mixing case) and requires a few bits of information about the starting profile (this is not needed in the rapidly mixing case).

Our contribution. Our main contribution is the introduction of the notion of an (ε, T) -*metastable distribution* of a Markov chain and of the concept of *pseudo-mixing time* of a metastable distribution. Roughly speaking, a distribution μ is (ε, T) -metastable for a Markov chain if, starting from μ , the Markov chain stays at distance at most ε from μ for at least T steps. The pseudo-mixing time of μ starting from a state x is the number of steps needed by the Markov chain to get ε -close to μ when started from x .

In a rapidly-mixing Markov chain, after a “short time” and regardless of the starting state, the chain

converges rapidly to the stationary distribution and remains there. For the case of *non*-rapidly-mixing Markov chains, we replace the notions of “mixing time” and “stationary distribution” with those of “pseudo-mixing time” and “metastable distribution”. Intuitively speaking, we would like to say that, even when the mixing time is (prohibitively) high, there are “few” distributions which give us an accurate description of the chain over a “reasonable amount of time”. That is, the state space Ω can be partitioned into a small number of subsets $\Omega_1, \Omega_2, \dots$ of “equivalent” states so that if the chain starts in *any* of the states in Ω_i , then it will rapidly converge to a “metastable” distribution μ_i , where metastable denotes the fact that the chain remains there for “sufficiently” long.

The similarities with stationary distributions and mixing time are confirmed by some of the properties that are enjoyed by the metastable distribution and the pseudo-mixing time concepts. In Section 3 we highlight these properties and also the relationship of these concepts with other well-known Markov chain measures, such as *bottleneck ratio* and *hitting time*.

In order to familiarize with the concepts of metastable distribution and pseudo-mixing time, we first apply them to a simple three-state Markov chain and the Markov chain of the logit dynamics for 2-player coordination games. Then, we show the usefulness of these concepts in more involved examples. Specifically, in Section 4 we consider a pure n -player coordination game, in which all players would like to take the same action and each action is equivalently valued by any player. People forming teams, objects being categorized (e.g., movies into genres in a video store) and firms choosing their trading locations are some examples of real-world scenarios that can be modeled by this class of games [14]. The behavior of the logit dynamics for this game turns out to be very interesting. Indeed, it highlights a distinctive feature of our metastability notion based on distributions, namely that a metastable distribution exists even if every small subset of states in the support of that distribution is quickly left with high probability.

Finally, in Section 5, we obtain results about the metastability of the logit dynamics for different classes of graphical coordination games. These games are often used to model the spread of a new technology in a social network [40, 35, 34] with the strategy of maximum potential corresponding to adopting the new technology; players prefer to choose the same technology as their neighbors and the new technology is at least as preferable as the old one. Research along this direction has mainly focused on analyzing how the features of the social network affect the spread of innovations. This research line has been initiated by Ellison [18], that considers two extremal network topologies, the clique and the ring¹.

In this work we follow Ellison [18] and study the case in which social interaction between the players is described by the clique and ring topologies. Specifically, in Section 5.1, we analyze the metastable distributions for a special graphical coordination game embedding the Ising model on the complete graph, also known as the Curie-Weiss model. It has been studied in the context of population protocols [36] and used by physicists to model the interaction between magnets in a ferro-magnetic system [32]. Indeed, it is known [24, 34, 7] that this model can be seen as a game played by magnets and, in particular, the Glauber dynamics for the Gibbs measure on the Ising model are equivalent to the logit dynamics for this game where β is exactly the inverse of the temperature. The mixing time of these dynamics is known to be exponential for every $\beta > 1/n$. For this model, we show that distributions where all magnets have the same magnetization are $(1/n, t)$ -metastable for t greater than any polynomial when $\beta = \omega(\log n/n)$. Moreover we show that the pseudo-mixing time of these distributions is polynomial when the dynamics start from a profile where the difference in the number of positive and negative magnets is large.

In Section 5.2 we focus on the ring topology and we show that for every starting profile there is a metastable distribution and the dynamics approach it in a polynomial number of steps. The metastable distributions that arise in this case consist of a mixture of the extreme distributions that assign probability 1 to profiles in which all nodes adopt the same strategy. In the case of a symmetric coordination game (the two strategies are equally liked) the mixture that arises has coefficients that are almost equal to the fraction of nodes adopting each of the two strategies. If the game is asymmetric, it turns out that coefficients are biased towards the profile in which all nodes adopt the most preferred strategy. In particular, this coefficient gets very close to 1 even if in the initial profile there is a constant number (depending on the metastability parameter ε) of nodes with this strategy.

¹We point out that the dynamics studied by Ellison [18] are slightly different from the logit dynamics, that have been instead adopted by the later works [40, 35, 34].

Other equilibrium notions. We find illustrative to compare the concepts studied in this paper (the logit dynamics and their metastable distributions) with the notion of Nash equilibrium. Similar comparison can be made with other solution concepts from Game Theory (correlated equilibria, sink equilibria). The notion of a Nash equilibrium has been extensively studied as a solution concept for predicting the behavior of selfish players. Indeed, if the players happen to be in a (pure) Nash equilibrium any sequence of selfish best response (i.e., utility improving) moves keeps the players in the same state. Unfortunately, the theory of Nash equilibria does not explain how a Nash equilibrium is reached if players do not start from one and, in case multiple equilibria exist, does not say which equilibrium is selected (about this important issue see [25]). Even though it is not hard to see that sequences of best response moves may reach a pure Nash equilibrium (if it exists), recent hardness results regarding the computation of pure Nash equilibria [20] suggest that the best-response dynamics (or any other dynamics) might take super-polynomial time in the number of players to reach an equilibrium. Thus, even in case only one (pure) Nash equilibrium exists, the players might take very long to reach it and thus it cannot be taken to describe the state of the players (unless we are willing to ignore the super-polynomially long transient phase). In contrast, in the setting studied in this paper these drawbacks disappear: the solution concept is *defined* in terms of dynamics and for our specific dynamics we have a unique possible prediction. For rapidly mixing chains the equilibrium is quickly reached. The concept of metastable distribution and the results in this paper show that, even for the non-rapidly mixing case, the system can be often efficiently described after a short initial transient phase.

The Nash Equilibrium concept is based on the assumption that each player has complete information about the game and the strategies of his opponents and it is able to compute his best strategy with respect to the strategies played by the other players. However, in many complex systems, environmental factors can influence the way each agent selects her own strategy and limitations to the players' computational power can influence their behaviors. Logit dynamics are a clear and crispy way to model these settings.

1.1 Related Work

Logit dynamics. Logit dynamics were first studied by Blume [10] who showed that, for two-player two-strategy coordination games, the long-term behavior of the system is concentrated in the risk dominant equilibrium (see [25]). Ellison [18] initiated the analysis of how social networks affect the spread of innovations. Ellison considered dynamics that are slightly different from the logit dynamics and he focused on two extremal network topologies, the clique and the ring (the classes of games we study in Section 5.2). His results show that some large fraction of the players will eventually choose the strategy with maximum potential. Similar results were obtained by Peyton Young [35] for the logit dynamics and more general families of graphs. Montanari and Saberi [34] gave bounds, in terms of some graph theoretic properties of the underlying interaction network, on the time the logit dynamics take to hit the highest potential equilibrium. Asadpour and Saberi [3] studied the hitting time of the Nash equilibrium for the logit dynamics in a class of congestion games. Bounds on the convergence to specific equilibria have been given also in [27, 28].

The study of the mixing time of the logit dynamics for strategic games has been initiated in [6] (see also [7]), whose results highlight a separation between games where the mixing time can be bounded independently from the parameter β and games where the mixing time is necessarily exponential in β .

Several variants of these dynamics have been proposed for evaluating the robustness of known results [31, 1, 2, 5], for adapting them to different setting [38], or for speeding up their convergence to equilibria [11].

Metastability. The goal of metastability is to model processes showing the following typical behavior: starting from a given *state*, the system rather quickly visits a nearby *metastable state* and then it remains very close to such a state for a very long time; at some point, the system leaves the metastable state (and its neighborhood) and moves to some other metastable state; the process then is repeated. *Metastable states* are identified with subsets of profiles of the system, with the property that the probability to leave these subsets is exponentially small [13]. Similar concepts have been also introduced in the context of population protocols under the name of *locally stable states* (see, e.g., [36] and references therein). Usually, these states correspond to “deep” local energy minima (or local maximizers of the potential function, in the case of potential games), with absolute energy minima corresponding to stable states. Sometimes, metastable state are described through probability distributions over the profiles in the subset

[13, 9]. Anyway, the metastability property always refers to the subset and not to the distribution. We depart from this approach by considering metastable *distributions* and not just metastable subsets of profiles. Indeed this is necessary for our object of study (logit dynamics of strategic games) as we shall show that there exist games that admit quite natural metastable distributions even if in the support there is no strict local maximizers of the potential function (see Section 4).

The classical approach to metastability is concerned with evaluating the time the system takes to escape from a metastable state [26]. More recent approaches [23] based on large deviation theory try to describe also the typical trajectory that a system takes, that is the sequence of states bringing to the stable one. This approach has been recently applied also to the logit dynamics [37], under the hypothesis that both β and n tend to ∞ . Better result on the escape time have been achieved recently, through the so-called potential theoretical approach [12], that links hitting time theory with spectral properties of the transition matrix of the dynamics. In contrast our work not only tries to bound the time the system takes to escape from a metastable distribution, but also the time the system takes to enter this distribution from a specific starting profile.

Very recently, and independently from our work, Friedlin and Koralov [22] introduced a notion of a metastable distribution that resembles the one presented in this work. Unlike ours, the notion of metastable distribution in [22] does not require the chain to stay close to this distribution for a given number of steps. Moreover, in [22] it is not considered the time required to get close to a metastable distribution.

Quasi-stationarity and Censored dynamics. A *quasi-stationary distribution* [15] describes the limiting behavior of a system in a subset of states given that the system never leaves this subset. Quasi-stationary distributions are often associated with metastability (see, e.g., [9]). Anyway, we highlight that, even if quasi-stationary distributions can be metastable, the inverse is not necessarily true. Moreover, works about quasi-stationary distributions focus on bounding the time the dynamics take to converge to this distribution from the support. In this work we are also interested in the convergence time to metastable distributions from profiles that are not in the support.

The work on censored Glauber dynamics [30, 17, 16] is also related to ours: the mixing time in censored dynamics resembles the pseudo mixing time for the metastable distribution on a subset of states. However, we stress that the censored dynamics alter the original evolution of the Markov chain and the techniques developed do not seem useful to answer questions about the pseudo-mixing time.

1.2 Notations

We write $|S|$ for the size of a set S . We use bold symbols for vectors and the standard game theoretic notation (\mathbf{x}_{-i}, y) to denote the vector obtained from \mathbf{x} by replacing its i -th entry with y ; i.e., $(\mathbf{x}_{-i}, y) = (x_1, \dots, x_{i-1}, y, x_{i+1}, \dots, x_n)$. When we are given a probability distribution μ_x that assigns probability 1 to state x , we say that μ_x is *concentrated* in state x . We refer the reader to Appendix A for a review of useful facts about the concepts of *total variation distance* $\|\mu - \nu\|_{\text{TV}}$ between probability distributions, and of *mixing time* t_{mix} .

2 Logit dynamics

A *strategic game* \mathcal{G} is a triple $\mathcal{G} = ([n], \mathcal{S}, \mathcal{U})$. The set $[n] = \{1, \dots, n\}$ is a finite set of *players*; $\mathcal{S} = S_1 \times \dots \times S_n$ is the set of *strategy profiles* and the finite set S_i is the set of *strategies* for player i ; $\mathcal{U} = (u_1, \dots, u_n)$ is a tuple of *utility functions* (or *payoffs*) where $u_i: \mathcal{S} \rightarrow \mathbb{R}$ and $u_i(\mathbf{x})$ is the payoff of player i in strategy profile \mathbf{x} .

Dynamics for a strategic game are probabilistic rules by which at each step: (i) one or more players are selected; (ii) each selected player updates her strategy by sampling according to a probability distribution that may depend on the current strategy profile. For example, in the best response dynamics each selected player updates her strategy by selecting with probability 1 a strategy that maximizes her payoff given the strategies currently adopted by other players.

In this paper we study specific dynamics for strategic games, the logit dynamics, introduced by Blume [10] and referred to as the *log-linear model*. The *logit dynamics with parameter* $\beta \geq 0$ are randomized best-response dynamics where at each step one player $i \in [n]$ is selected uniformly at random

and she updates her strategy in profile $\mathbf{x} \in \mathcal{S}$ by choosing strategy $y \in S_i$ with probability $\sigma_i(y \mid \mathbf{x})$ defined as follows

$$\sigma_i(y \mid \mathbf{x}) = \frac{e^{\beta u_i(\mathbf{x}_{-i}, y)}}{T_i(\mathbf{x})}, \quad (1)$$

where $T_i(\mathbf{x}) = \sum_{z \in S_i} e^{\beta u_i(\mathbf{x}_{-i}, z)}$ is the normalizing factor. In other words, in the logit dynamics the logarithm of the ratio of the probability of choosing strategy y and y' is proportionally related through β to the difference of the respective utilities. Parameter β can be seen as descriptive of the *rationality level* of the players: for $\beta = 0$ (no rationality) player i updates her strategy by selecting a new strategy uniformly at random; for $\beta > 0$, the probability is biased towards strategies yielding higher payoffs; and for $\beta \rightarrow \infty$ (full rationality) player i chooses her best response strategy (if more than one best response is available, she chooses uniformly at random one of them).

For every $\beta \geq 0$, the logit dynamics with parameter β induce a Markov chain over the set \mathcal{S} of strategy profiles. More formally, we have the following definition.

Definition 2.1 (Markov chain of the logit dynamics [10]). *Let $\mathcal{G} = ([n], \mathcal{S}, \mathcal{U})$ be a strategic game and let $\beta \geq 0$. The Markov chain \mathcal{M}_β of the logit dynamics with parameter β for strategic game $\mathcal{G} = ([n], \mathcal{S}, \mathcal{U})$ is the Markov chain with state space \mathcal{S} and transition probability*

$$P(\mathbf{x}, \mathbf{y}) = \frac{1}{n} \cdot \begin{cases} \sum_{i=1}^n \sigma_i(y_i \mid \mathbf{y}), & \text{if } \mathbf{x} = \mathbf{y}; \\ \sigma_i(y_i \mid \mathbf{x}), & \text{if } \mathbf{x} \neq \mathbf{y} \text{ and } \mathbf{y}_{-i} = \mathbf{x}_{-i}; \\ 0, & \text{otherwise;} \end{cases} \quad (2)$$

where $\sigma_i(y_i \mid \mathbf{x})$ is defined in (1).

It is easy to see that the Markov chain is ergodic as $P^n(\mathbf{x}, \mathbf{y}) > 0$ for every pair of profiles \mathbf{x} and \mathbf{y} . Hence, a unique stationary distribution, i.e. a distribution π such that $\pi = \pi P$, exists and for every starting profile \mathbf{x} the distribution $P^t(\mathbf{x}, \cdot)$ approaches π as t tends to infinity.

Potential games. Function $\Phi: \mathcal{S} \rightarrow \mathbb{R}$ is a *potential* for strategic game $\mathcal{G} = ([n], \mathcal{S}, \mathcal{U})$ if for every player $i \in [n]$ and for profiles $\mathbf{x}, \mathbf{y} \in \mathcal{S}$ differing only for player i , it holds that $u_i(\mathbf{x}) - u_i(\mathbf{y}) = \Phi(\mathbf{x}) - \Phi(\mathbf{y})$. Games \mathcal{G} admitting a potential function are called *potential games* [33]. The stationary distribution of the Markov chain of the logit dynamics with parameter β of a potential game \mathcal{G} is the Gibbs measure π defined as follows

$$\pi(\mathbf{x}) = \frac{e^{\beta \Phi(\mathbf{x})}}{Z}, \quad (3)$$

where $Z = \sum_{\mathbf{y} \in \mathcal{S}} e^{\beta \Phi(\mathbf{y})}$ is the *partition function*. The Markov chain of the logit dynamics with parameter β for a potential game \mathcal{G} coincides with the well-studied *Glauber dynamics* with temperature $1/\beta$ for the Gibbs measure π .

Two-strategy coordination games. A two-strategy coordination game is a two-player strategic game where each player has two strategies $+1$ and -1 . Thus $S_1 = S_2 = \{\pm 1\}$ and $\mathcal{S} = \{\pm 1\}^2$. The utility function is described by the following table

	$+1$	-1
$+1$	a, a	c, d
-1	d, c	b, b

Thus, for example, $u_1(-1, +1) = d$ and $u_2(-1, +1) = c$. It is usually assumed that $a > d$ and $b > c$ to model the fact that each player prefers to choose the same strategy as the other player (hence the name of *coordination games*). This implies that profiles $(+1, +1)$ and $(-1, -1)$ are Nash equilibria of the two-strategy coordination games. We define $\Delta := a - d$ and $\delta := b - c$ and assume, w.l.o.g., that $\Delta \geq \delta$. If $\Delta > \delta$, then $+1$ is the *risk dominant strategy* ([25]). It is easy to see that the function Φ defined as follows

$$\Phi(+1, +1) = \Delta, \quad \Phi(-1, -1) = \delta, \quad \text{and} \quad \Phi(+1, -1) = \Phi(-1, +1) = 0$$

is a potential function for the two-strategy coordination game. The logit dynamics with parameter β for two-strategy coordination games have been first studied in [10] and have the following transition matrix

$$P = \left(\begin{array}{c|cccc} & (+1, +1) & (+1, -1) & (-1, +1) & (-1, -1) \\ \hline (+1, +1) & 1-p & p/2 & p/2 & 0 \\ (+1, -1) & (1-p)/2 & (p+q)/2 & 0 & (1-q)/2 \\ (-1, +1) & (1-p)/2 & 0 & (p+q)/2 & (1-q)/2 \\ (-1, -1) & 0 & q/2 & q/2 & 1-q \end{array} \right)$$

where $p = \frac{1}{1+e^{\Delta\beta}}$ and $q = \frac{1}{1+e^{\delta\beta}}$. The partition function is

$$Z = e^{\beta\Delta} + 1 + 1 + e^{\beta\delta} = \frac{1}{p} + \frac{1}{q}$$

and the stationary distribution for P is the Gibbs measure

$$\pi = \left(\frac{e^{\beta\Delta}}{Z}, \frac{1}{Z}, \frac{1}{Z}, \frac{e^{\beta\delta}}{Z} \right)$$

that can be re-written as

$$\pi = \left(\frac{q(1-p)}{p+q}, \frac{pq}{p+q}, \frac{pq}{p+q}, \frac{p(1-q)}{p+q} \right).$$

3 Metastability

In this section we give formal definitions of the concepts of *metastable distribution* (Section 3.1) and of *pseudo-mixing time* (Section 3.2) and we highlight some of their properties. Finally, we exemplify the new concepts by studying a simple three-state Markov chain and the Markov chain of the logit dynamics for two-player two-strategy coordination games (in Section 4 we will study an n -player pure coordination game and in Section 5 we study the metastability of coordination games on graphs).

3.1 Metastable distributions

Definition 3.1 (Metastable distribution). *Let P be a Markov chain with finite state space Ω . A probability distribution μ over Ω is (ε, T) -metastable for P if, for $0 \leq t \leq T$, it holds that*

$$\|\mu P^t - \mu\|_{\text{TV}} \leq \varepsilon.$$

The definition of a metastable distribution captures the idea of a distribution that behaves approximately like the stationary distribution; meaning that if we start from a metastable distribution and run the chain we stay close to that distribution (that is, within ε) for a long time (that is, for at least T time steps).

Let us now highlight some properties of metastable distributions, that can be easily derived from known results in Markov chain theory:

1. *Monotonicity:* If μ is (ε, T) -metastable for P then it is (ε', T') -metastable for every $\varepsilon' \geq \varepsilon$ and $T' \leq T$;
2. *Stationarity:* μ is stationary if and only if it is $(0, 1)$ -metastable.

A third property is given by the following easy and useful lemma.

Lemma 3.2 (Additivity). *If μ is $(\varepsilon, 1)$ -metastable for P then, for every integer $T > 0$, μ is $(\varepsilon T, T)$ -metastable for P .*

Proof. Since the total variation distance satisfies the triangle inequality (see Fact 1 in Appendix A) and μ is $(\varepsilon, 1)$ -metastable, we have

$$\|\mu P^T - \mu\|_{\text{TV}} \leq \|\mu P^T - \mu P\|_{\text{TV}} + \|\mu P - \mu\|_{\text{TV}} \leq \|\mu P^{T-1} - \mu\|_{\text{TV}} + \varepsilon.$$

The lemma then follows. \square

The next lemma states that the convex combination of two metastable distributions is metastable.

Lemma 3.3 (*Convexity*). *If μ_1 is (ε_1, T_1) -metastable for P and μ_2 is (ε_2, T_2) -metastable for P , then, for $0 \leq \alpha \leq 1$, $\mu = \alpha\mu_1 + (1-\alpha)\mu_2$ is (ε, T) -metastable for P with $\varepsilon = \max\{\varepsilon_1, \varepsilon_2\}$ and $T = \min\{T_1, T_2\}$.*

Proof. For any $t \leq T$, we have

$$\begin{aligned} \|\mu P^t - \mu\|_{\text{TV}} &= \|\alpha(\mu_1 P^t - \mu_1) + (1-\alpha)(\mu_2 P^t - \mu_2)\|_{\text{TV}} \\ &\leq \alpha \|\mu_1 P^t - \mu_1\|_{\text{TV}} + (1-\alpha) \|\mu_2 P^t - \mu_2\|_{\text{TV}} \leq \varepsilon. \end{aligned} \quad \square$$

Finally, we highlight a connection between metastability and *bottleneck ratio*. Given an ergodic Markov chain P with state space Ω and stationary distribution π , the bottleneck ratio $B(S)$ for a subset $S \subseteq \Omega$ of states, is defined as

$$B(S) = \frac{Q(S, \bar{S})}{\pi(S)},$$

where $Q(S, \bar{S}) = \sum_{x \in S} \sum_{y \in \Omega \setminus S} \pi(x) P(x, y)$. Let π_S be the stationary distribution conditioned on S ; i.e.,

$$\pi_S(x) = \begin{cases} \pi(x)/\pi(S), & \text{if } x \in S; \\ 0, & \text{otherwise.} \end{cases} \quad (4)$$

Lemma 3.4. *Let P be a Markov chain with finite state space Ω and let $S \subseteq \Omega$ be a subset of states. Then, π_S is $(B(S), 1)$ -metastable.*

Proof. The lemma follows from the fact that the bottleneck ratio $B(S)$ of S is equal to the total variation distance between π_S and $\pi_S P$; i.e., $\|\pi_S P - \pi_S\| = B(S)$. See, for example, Theorem 7.3 from [29]. \square

3.2 Pseudo-mixing time

Among all metastable distributions, we are interested in the ones that are quickly reached from a, possibly large, set of states. This motivates the following definition.

Definition 3.5 (Pseudo-mixing time). *Let P be a Markov chain with state space Ω , let $S \subseteq \Omega$ be a set of states and let μ be a probability distribution over Ω . We define $d_\mu^S(t)$ as*

$$d_\mu^S(t) = \max_{x \in S} \|P^t(x, \cdot) - \mu\|_{\text{TV}}.$$

Then the pseudo-mixing time $t_\mu^S(\varepsilon)$ of μ starting from S is

$$t_\mu^S(\varepsilon) = \inf \{t \in \mathbb{N} : d_\mu^S(t) \leq \varepsilon\}.$$

The pseudo-mixing time extends the concept of mixing time to metastable distributions and it coincides with the mixing time for a stationary distribution. Indeed, the stationary distribution π of an ergodic Markov chain is a $(0, 1)$ -metastable distribution that is reached within ε in time $t_{\text{mix}}(\varepsilon)$ from every state (see Appendix A). Thus, according to Definition 3.5, we have that $t_\pi^S(\varepsilon) = t_{\text{mix}}(\varepsilon)$.

We now highlight some interesting properties related to the pseudo-mixing time and metastable distribution concepts. The first property connects $d_\mu^S(t)$ for a metastable distribution μ to a quantity that does not depend on this distribution.

Lemma 3.6. *Let P be a Markov chain with finite state space Ω and let μ be an (ε, T) -metastable distribution supported over a subset S of the state space. Then, for $1 \leq t \leq T$, it holds that*

$$d_\mu^S(t) \leq \varepsilon + \max_{x, y \in S} \|P^t(x, \cdot) - P^t(y, \cdot)\|_{\text{TV}}.$$

Proof. From the triangle inequality, we have

$$\|P^t(x, \cdot) - \mu\|_{\text{TV}} \leq \|P^t(x, \cdot) - \mu P^t\|_{\text{TV}} + \|\mu P^t - \mu\|_{\text{TV}}.$$

Since μ is (ε, t) -metastable, then for every $t \leq T$ we have $\|\mu P^t - \mu\|_{\text{TV}} \leq \varepsilon$. Moreover, since $\mu(y) = 0$ for $y \notin S$, for every set of states $A \subseteq \Omega$ and for every t it holds that

$$\begin{aligned} |P^t(x, A) - \mu P^t(A)| &= \left| P^t(x, A) - \sum_{y \in S} \mu(y) P^t(y, A) \right| = \left| \sum_{y \in S} \mu(y) (P^t(x, A) - P^t(y, A)) \right| \\ &\leq \sum_{y \in S} \mu(y) |P^t(x, A) - P^t(y, A)| \leq \max_{y \in S} |P^t(x, A) - P^t(y, A)|. \end{aligned}$$

Thus, the total variation between $P^t(x, \cdot)$ and μP^t is

$$\begin{aligned} \|P^t(x, \cdot) - \mu P^t\|_{\text{TV}} &= \max_{A \subseteq \Omega} |P^t(x, A) - \mu P^t(A)| \\ &\leq \max_{A \subseteq \Omega} \max_{y \in S} |P^t(x, A) - P^t(y, A)| = \max_{y \in S} \|P^t(x, \cdot) - P^t(y, \cdot)\|_{\text{TV}}. \end{aligned} \quad \square$$

Notice that if μ is the stationary distribution then it is $(0, T)$ -metastable for every T and the above lemma gives

$$d_\mu^\Omega(t) \leq \max_{x, y \in \Omega} \|P^t(x, \cdot) - P^t(y, \cdot)\|_{\text{TV}}$$

for all t . This is a well-know inequality and it is widely used to bound the mixing time of Markov chains (see, for example, Lemma 4.11 in [29]).

Observe that, even if the pseudo-mixing time of a metastable distribution μ is finite, we cannot argue that the Markov chain converges to μ , but only that it converges to a distribution μ' that is *close* to μ . However, next lemma shows that when starting from μ' the Markov chain will behave approximatively as if it starts from the metastable distribution μ .

Lemma 3.7. *Let P be a Markov chain with finite state space Ω and let μ be an (ε, T) -metastable distribution supported over a subset S of the state space with $t_\mu^S(\varepsilon) < +\infty$. Then, for every $x \in S$ and $t_\mu^S(\varepsilon) \leq t \leq t_\mu^S(\varepsilon) + T$, it holds that*

$$\|P^t(x, \cdot) - \mu\|_{\text{TV}} \leq 2\varepsilon.$$

Proof. Let us name $\bar{t} = t - t_\mu^S(\varepsilon)$ for convenience sake. By using the triangle inequality for the total variation distance, the fact that $P^{\bar{t}}$ is a stochastic matrix, and the definitions of metastable distribution and pseudo-mixing, we have that

$$\begin{aligned} \|P^t(x, \cdot) - \mu\|_{\text{TV}} &= \|P^{t_\mu^S(\varepsilon)}(x, \cdot) P^{\bar{t}} - \mu\|_{\text{TV}} \\ &\leq \|P^{t_\mu^S(\varepsilon)}(x, \cdot) P^{\bar{t}} - \mu P^{\bar{t}}\|_{\text{TV}} + \|\mu P^{\bar{t}} - \mu\|_{\text{TV}} \\ &\leq \|P^{t_\mu^S(\varepsilon)}(x, \cdot) - \mu\|_{\text{TV}} + \|\mu P^{\bar{t}} - \mu\|_{\text{TV}} \leq 2\varepsilon. \end{aligned} \quad \square$$

Finally, we highlight some connections between pseudo-mixing time and *hitting time*. Given a Markov chain P with finite state space Ω , the hitting time τ_S of a subset $S \subset \Omega$ is the first time step in which the Markov chain reaches a state from S , i.e.,

$$\tau_S = \min\{t: X_t \in S\}.$$

When S consists of a single state x , we write τ_x rather than $\tau_{\{x\}}$. Then, we have the following lemma.

Lemma 3.8. *Let P be a Markov chain with finite state space Ω and let μ_y be an (ε, T) -metastable distribution concentrated in state y . Then for all $S \subseteq \Omega$ and $t \leq T$ we have*

$$d_{\mu_y}^S(t) \leq \varepsilon + \max_{x \in S} \mathbf{P}_x(\tau_y > t).$$

Proof. Since μ_y is concentrated in y , we have that

$$\begin{aligned}\|P^t(x, \cdot) - \mu_y\|_{\text{TV}} &= \mathbf{P}_x(X_t \neq y) = \mathbf{P}_x(X_t \neq y, \tau_y \leq t) + \mathbf{P}_x(X_t \neq y, \tau_y > t) \\ &= \mathbf{P}_x(X_t \neq y \mid \tau_y \leq t) \mathbf{P}_x(\tau_y \leq t) + \mathbf{P}_x(\tau_y > t).\end{aligned}$$

Moreover, observe that

$$\begin{aligned}\mathbf{P}_x(X_t \neq y \mid \tau_y \leq t) &= \sum_{k \leq t} \mathbf{P}_x(X_t \neq y \mid \tau_y = k) \mathbf{P}_x(\tau_y = k \mid \tau_y \leq t) \\ &= \sum_{k \leq t} \mathbf{P}_y(X_{t-k} \neq y) \mathbf{P}_x(\tau_y = k \mid \tau_y \leq t) \\ &= \sum_{k \leq t} \|\mu_y P^{t-k} - \mu_y\|_{\text{TV}} \mathbf{P}_x(\tau_y = k \mid \tau_y \leq t) \\ &\leq \varepsilon \sum_{k \leq t} \mathbf{P}_x(\tau_y = k \mid \tau_y \leq t) = \varepsilon.\end{aligned}$$

where in the inequality we used the metastability of μ_y . Hence,

$$\begin{aligned}\|P^t(x, \cdot) - \mu_y\|_{\text{TV}} &= \mathbf{P}_x(X_t \neq y \mid \tau_y \leq t) \mathbf{P}_x(\tau_y \leq t) + \mathbf{P}_x(\tau_y > t) \\ &\leq \varepsilon \mathbf{P}_x(\tau_y \leq t) + \mathbf{P}_x(\tau_y > t) \leq \varepsilon + \mathbf{P}_x(\tau_y > t).\end{aligned}\quad \square$$

The following lemma show another interesting relation between hitting time and the pseudo-mixing time of a convex combination of metastable distributions.

Lemma 3.9. *Let P be a Markov chain with finite state space Ω . Let μ_y be an (ε_y, T_y) -metastable distribution concentrated in state y and μ_z be an (ε_z, T_z) -metastable distribution concentrated in state z . Fix $\varepsilon = \max\{\varepsilon_y, \varepsilon_z\}$ and $T = \min\{T_y, T_z\}$. Then for all $x \in \Omega$ and $t \leq T$ we have*

$$d_{\mu_{x,t}}^{\{x\}}(t) \leq \varepsilon + \mathbf{P}_x(\tau_{\{y,z\}} > t),$$

where $\mu_{x,t} = \alpha_{x,t} \mu_y + (1 - \alpha_{x,t}) \mu_z$ with $\alpha_{x,t} = \mathbf{P}_x(\tau_y \leq \tau_z \mid \tau_{\{y,z\}} \leq t)$.

Proof. For sake of readability, let us denote by E the event $\tau_{\{y,z\}} > t$ and by \overline{E} its complement. Observe that

$$\begin{aligned}\|P^t(x, \cdot) - \mu_{x,t}\|_{\text{TV}} &= \max_{A \subset \Omega} |\mathbf{P}_x(X_t \in A) - \mu_{x,t}(A)| \\ &= \max_{A \subset \Omega} |\mathbf{P}_x(X_t \in A \mid \overline{E}) - \mu_{x,t}(A) + (\mathbf{P}_x(X_t \in A \mid E) - \mathbf{P}_x(X_t \in A \mid \overline{E})) \mathbf{P}_x(E)| \\ &\leq \mathbf{P}_x(E) + \max_{A \subset \Omega} |\mathbf{P}_x(X_t \in A \mid \overline{E}) - \mu_{x,t}(A)|.\end{aligned}$$

Since Markov chains are memoryless, we have

$$\mathbf{P}_x(X_t \in A \mid \overline{E}) = \sum_w \mathbf{P}_x(X_{\tau_{\{y,z\}}} = w \mid \overline{E}) \mathbf{P}_x(X_t \in A \mid X_{\tau_{\{y,z\}}} = w \wedge \overline{E}).$$

Observe that $\mathbf{P}_x(X_{\tau_{\{y,z\}}} = y \mid \overline{E}) = \alpha_{x,t}$, $\mathbf{P}_x(X_{\tau_{\{y,z\}}} = z \mid \overline{E}) = 1 - \alpha_{x,t}$, and $\mathbf{P}_x(X_{\tau_{\{y,z\}}} = w \mid \overline{E}) = 0$ for each state $w \neq y, z$. That is, $\mathbf{P}_x(X_{\tau_{\{y,z\}}} = w \mid \overline{E}) = \mu_{x,t}(w)$ for each $w \in \Omega$. Moreover, since $t - \tau_{\{y,z\}} \geq 0$ if the event \overline{E} holds, for each $w \in \Omega$ we have

$$\mathbf{P}_x(X_t \in A \mid X_{\tau_{\{y,z\}}} = w \wedge \overline{E}) = \mathbf{P}_w(X_{t-\tau_{\{y,z\}}} \in A) = P^{t-\tau_{\{y,z\}}}(w, A).$$

Thus,


$$\|P^t(x, \cdot) - \mu_{x,t}\|_{\text{TV}} \leq \mathbf{P}_x(E) + \max_{A \subset \Omega} |(\mu_{x,t} P^{t-\tau_{\{y,z\}}})(A) - \mu_{x,t}(A)| \leq \mathbf{P}_x(E) + \varepsilon,$$

where the last inequality follows from $\mu_{x,t}$ being (ε, T) -metastable from Lemma 3.3 and since $t - \tau_{\{y,z\}} \leq T$. \square

3.3 Examples

3.3.1 A simple three-state Markov chain

As a first example, let us consider the simplest Markov chain that highlights the concepts of *metastability* and *pseudo-mixing*,



The chain is ergodic with stationary distribution $\pi = (\varepsilon, (1-\varepsilon)/2, (1-\varepsilon)/2)$ and its mixing time is $t_{\text{mix}} = \Theta(1/\varepsilon)$. Hence the mixing time grows unbounded as ε tends to zero.

Now observe that, for every $\delta > \varepsilon$, distributions $\mu_1 = (0, 1, 0)$ and $\mu_2 = (0, 0, 1)$ are $(\delta, \Theta(\delta/\varepsilon))$ -metastable according to Definition 3.1. Moreover, if the chain starts from state 0, the state of the chain after one step is distributed as in the stationary distribution. Hence, even if the mixing time can be arbitrary large, for every ε and for every starting state x there is a $(\delta, \Theta(t_{\text{mix}}))$ -metastable distribution μ that is quickly (in constant, independent of ε , time) reached from x .

3.3.2 Two-strategy coordination games

In [6] it is proved that the mixing time of the logit dynamics with parameter β for a two-player two-strategy coordination game defined above is $t_{\text{mix}} = \Theta(1/q) = \Theta(1 + e^{\beta\delta})$ and thus it grows unbounded in β . We next describe distributions that are metastable for an amount of time of the same order of the mixing time and whose pseudo mixing time is independent from β . Specifically, consider distributions μ_+ and μ_- concentrated in states $(+1, +1)$ and $(-1, -1)$, respectively; i.e.,

$$\mu_+ = (1, 0, 0, 0) \quad \text{and} \quad \mu_- = (0, 0, 0, 1).$$

Observe that, if we start from μ_+ or μ_- , after one step of the chain we are respectively in distributions

$$\mu_+ P = (1-p, p/2, p/2, 0) \quad \text{and} \quad \mu_- P = (0, q/2, q/2, 1-q).$$

Then, since $\Delta \geq \delta$, we obtain

$$\|\mu_+ P - \mu_+\|_{\text{TV}} = p \leq q \quad \text{and} \quad \|\mu_- P - \mu_-\|_{\text{TV}} = q.$$

In other words, μ_- and μ_+ are $(q, 1)$ -metastable for the Markov chain of the logit dynamics with parameter β for the two-strategy coordination game. By Lemma 3.2, they are $(\varepsilon, \varepsilon/q)$ -metastable for any $\varepsilon > 0$ and, since each of the two distributions has a support set consisting of one element, the pseudo-mixing time starting from their respective supports is trivially 0.

Let us discuss what happens when the Markov chain starts from the states $(+1, -1)$ or $(-1, +1)$. The distributions after one step starting from the two states are

$$P((+1, -1), \cdot) = \left(\frac{1-p}{2}, \frac{p+q}{2}, 0, \frac{1-q}{2} \right), \quad \text{and} \quad P((-1, +1), \cdot) = \left(\frac{1-p}{2}, 0, \frac{p+q}{2}, \frac{1-q}{2} \right).$$

Thus, we compute the total variation distance from the stationary distribution π and obtain

$$\|P((+1, -1), \cdot) - \pi\|_{\text{TV}} = \|P((-1, +1), \cdot) - \pi\|_{\text{TV}} = \frac{p}{2} + \frac{1}{2} \cdot \frac{q-p}{q+p} < \frac{1}{2},$$

where we used that $p + q < 2$. Then, we obtain that the logit dynamics starting from $(+1, -1)$ or $(-1, +1)$ are ε -close to the stationary distribution after $\mathcal{O}(\log \varepsilon^{-1})$ steps for each $\varepsilon > 0$ (see Fact 3 in Appendix A).

We summarize the above discussion in the following theorem.

Theorem 3.10. *For every starting profile $\mathbf{x} \in \{\pm 1\}^2$ and $\varepsilon > 0$, the Markov chain of the logit dynamics with parameter β for the two-strategy coordination game admits a $(\varepsilon, t_{\text{mix}} \cdot \varepsilon)$ -metastable distribution $\mu_{\mathbf{x}}$ with pseudo-mixing time $t_{\mu_{\mathbf{x}}}^{\{\mathbf{x}\}} = \mathcal{O}(\log \varepsilon^{-1})$ (independent of β).*

We would like to remark how the above theorem, jointly with the result on the mixing time of [7], gives a complete picture of the behavior of a system that evolves according to the logit dynamics for the two-strategy coordination game: given the initial state, metastable distributions describe, with an error at most ε , the state of the system for every time t between $\log 1/\varepsilon$ and $(1 + e^{\beta\delta}) \cdot \varepsilon$, whereas the stationary distribution can be used for describing the system for larger values of t .

4 Metastability of Pure Coordination games

In this section, we study a simple potential game and describe three metastable distributions for its logit dynamics. One of these metastable distributions is particularly interesting: it gives positive probability to states that are not local maximum of the potential function and that are easy to leave. This shows that metastable distributions are not necessarily concentrated on states that are increasingly hard to leave as β grows.

For $n \geq 3$, the *pure coordination game* [14] (also known as *unanimity game*) is an n -player game where players have the same strategy set A and each player is *happy* when all players adopt the same strategy and *unhappy* otherwise. More formally, the utility $u_i(\mathbf{x})$ of player $i \in [n]$ in profile \mathbf{x} is $u_i(\mathbf{x}) = a$, if $\mathbf{x} = (s, \dots, s)$ for some $s \in A$; and $u_i(\mathbf{x}) = b < a$, otherwise. This class of games has several Nash equilibria (every profile except the ones in which a single player adopts a strategy different from the rest of the players) strategy is an equilibrium) and this makes it very difficult to predict the behavior of the system. These games are often used for modeling the division of people in teams (e.g., close friends deciding which soccer team to join), or the choice of firms' location in a space (e.g., linked industries deciding in which area of the country to produce their goods).

In order to keep the analysis of metastability for the logit dynamics as simple as possible, we consider a very simple setting in which the strategy set $A = \{+1, -1\}$ consists of only two strategies and $a = 1$ and $b = 0$. We let \mathbf{p} and \mathbf{m} denote, respectively, the profile in which all players play $+1$ and the profile in which all players play -1 . This game is a potential game with the following potential function Φ : $\Phi(\mathbf{p}) = \Phi(\mathbf{m}) = 1$ and $\Phi(\mathbf{x}) = 0$ for $\mathbf{x} \neq \mathbf{p}, \mathbf{m}$. The mixing time of its logit dynamics is roughly e^β , and hence super-polynomial for $\beta = \omega(\log n)$. To see this, recall the known relation between bottleneck ratio and mixing time (see, e.g., Chapter 7.2 in [29]) and observe that the probability of \mathbf{p} at stationarity is $\pi(\mathbf{p}) \leq 1/2$ and the bottleneck ratio for \mathbf{p} is $B(\mathbf{p}) = \frac{1}{1+e^\beta}$.

Let us now focus on the metastable distributions. We show that if we start the logit dynamics at a profile where at least one player is playing $+1$ and at least one is playing -1 , then after $\mathcal{O}(\log n)$ time steps the distribution is close to uniform on the profiles different from \mathbf{p} and \mathbf{m} , and it stays close to this distribution for exponential time. Hence, even if there is no small sub-set of the state space where the chain stays close for a long time, we can still say that the chain is metastable in the sense that the “distribution” of the chain stays close to some well-defined distribution for a long time.

Specifically, we identify three metastable distributions for the logit dynamics: π_+ , the distribution concentrated on the profile \mathbf{p} where every player plays $+1$; π_- , the distribution concentrated on the profile \mathbf{m} where every player plays -1 ; and the uniform distribution U over all remaining states. In addition, we prove that the pseudo-mixing time of U starting from a profile other than \mathbf{p} and \mathbf{m} is $\mathcal{O}(n \log n)$.

The idea behind this result is that the process according to which the number of players adopting a given strategy evolves is essentially an Ehrenfest urn (see Appendix C), the only difference being at extreme values, i.e., when this number is either 0 or n . Hence, as long as one of these extreme value is not hit (and this takes time that is exponential in n), the distribution on the profiles of the game is well-described by the stationary distribution of the Ehrenfest urn as an approximation of the future status.

The next lemma establishes the metastability of π_+ , π_- and U .

Lemma 4.1. *Consider the logit dynamics with parameter β for the n -player pure equilibrium game. For every $\varepsilon > 0$, the uniform distribution U over $R = \{+1, -1\}^n \setminus \{\mathbf{p}, \mathbf{m}\}$ is $(\varepsilon, \varepsilon \cdot 2^n)$ -metastable and π_+ and π_- are $(\varepsilon, \varepsilon \cdot e^\beta)$ -metastable.*

Proof. Let us denote by P the transition matrix of the logit dynamics and by π its stationary distribution. Then observe that $\pi(R) = \frac{2^n - 2}{Z}$ and $Z = 2e^\beta + 2^n - 2$. Moreover,

$$Q(R, \bar{R}) = \sum_{\mathbf{x} \in R} \pi(\mathbf{x}) \cdot (P(\mathbf{x}, \mathbf{p}) + P(\mathbf{x}, \mathbf{m})) = \frac{2}{Z} \cdot \frac{e^\beta}{1 + e^\beta}.$$

Therefore, the bottleneck ratio is

$$B(R) = \frac{Q(R, \bar{R})}{\pi(R)} = \frac{e^\beta}{1 + e^\beta} \cdot \frac{2}{2^n - 2} < \frac{1}{2^{n-1}}.$$

Similarly, it is immediate to see that $B(\mathbf{p}) = B(\mathbf{m}) = \frac{1}{1 + e^\beta}$. The lemma then follows from Lemma 3.4, Lemma 3.2 and by observing that the stationary distribution restricted to R is the uniform distribution U . \square

As for bounding the convergence to these metastable distributions, we first start with the following technical lemma.

Lemma 4.2. *Let $\tau_{\mathbf{p}, \mathbf{m}}$ denote the hitting time of either profile \mathbf{p} or profile \mathbf{m} for the logit dynamics with parameter β of the n -player pure coordination game. Then, for every profile \mathbf{x} with at least one player adopting action $+1$ and at least one player adopting action -1 ,*

$$\mathbf{P}_{\mathbf{x}} \left(\tau_{\mathbf{p}, \mathbf{m}} \leq n \log \frac{n}{\varepsilon} \right) \leq \frac{c}{n},$$

for a suitable constant $c = c(\varepsilon)$.

Proof. We start by proving that $\mathbf{P}_{\mathbf{x}}(\tau_{\mathbf{p}, \mathbf{m}} \leq t) \leq \mathbf{P}_k(\rho_{0,n} \leq t)$, where $\rho_{0,n}$ is the hitting time of state 0 or state n for the lazy Ehrenfest urn and k is the number of $+1$'s in \mathbf{x} .

We say that $\mathbf{x}, \mathbf{y} \in \Omega = \{+1, -1\}^n$ are equivalent if they have the same number of $+1$'s and we let $Z = \{Z_t\}$ be the *projection* of the logit dynamics starting from \mathbf{x} over the quotient space $\Omega_{\#} = \{0, 1, \dots, n\}$. The hitting time $\tau_{\mathbf{p}, \mathbf{m}}$ for the logit dynamics starting in \mathbf{x} coincide with the hitting time $\hat{\rho}_{0,n}$ of state 0, $n \in \Omega_{\#}$ for the projection Z starting in k . Z has the following transition matrix $P_{\#}$.

$$P_{\#}(i, i-1) = \frac{i}{2n}; \quad P_{\#}(i, i) = \frac{1}{2}; \quad P_{\#}(i, i+1) = \frac{n-i}{2n}; \quad \text{for } i = 2, \dots, n-2 \quad (5)$$

and

$$\begin{aligned} P_{\#}(1, 0) &= P_{\#}(n-1, n) = \frac{1}{n(1 + e^{-\beta})} \leq \frac{1}{n}; \\ P_{\#}(1, 1) &= P_{\#}(n-1, n-1) = \frac{n-1}{2n} + \frac{1}{n(1 + e^{\beta})}; \\ P_{\#}(1, 2) &= P_{\#}(n-1, n-2) = \frac{n-1}{2n}. \end{aligned}$$

$P_{\#}$ differs from the transition matrix of the lazy Ehrenfest urn only for states 1 and $n-1$. Observe though that the transitions from state 1 to state 0 and from state $n-1$ to state n in Z have smaller probability than in the Ehrenfest urn and this can only increase the hitting time of 0 and 1 in Z .

The lemma then follows from Lemmas C.1 and C.2. \square

In the next lemma we show that, if the chain starts from a state containing at least one 1 and at least one 0, then after $\mathcal{O}(n \log n)$ time steps the distribution of the chain is ε -close to the distribution U .

Lemma 4.3. *Let $R = \{+1, -1\}^n \setminus \{\mathbf{p}, \mathbf{m}\}$ and let U be the uniform distribution over R . Then for the logit dynamics for the n -player pure coordination game, it holds that*

$$t_U^R(\varepsilon) = \mathcal{O}\left(n \log \frac{n}{\varepsilon}\right).$$

Proof. Let $\{X_t\}$ be the Markov chain starting at $\mathbf{x} \in R$ and let $\{Y_t\}$ be a lazy random walk on the n -cube starting at the uniform distribution U^* on $\{+1, -1\}^n$, so that X_t is distributed according to $P^t(\mathbf{x}, \cdot)$ and Y_t is uniformly distributed over $\{+1, -1\}^n$. Consider the following *coupling* (X_t, Y_t) (for a formal definition and some useful facts about the coupling of Markov chains, we refer the reader to Appendix A): when chain $\{X_t\}$ is at state $\mathbf{y} \in \{+1, -1\}^n$ then choose a position $i \in [n]$ uniformly at random and, by denoting with $|\mathbf{y}|$ the number of players playing strategy $+1$ in \mathbf{y} ,

- If $2 \leq |\mathbf{y}| \leq n-2$, or $|\mathbf{y}| = 1$ and X_t has -1 in position i , or $|\mathbf{y}| = n-1$ and X_t has $+1$ in position i , then choose an action $a \in \{-1, 1\}$ uniformly at random and update both chains X_t and Y_t in position i with action a ;
- If $|\mathbf{y}| = 0$ or both $|\mathbf{y}| = 1$ and X_t has $+1$ in position i , then
 - update both chains at -1 in position i with probability $1/2$;
 - update both chains at $+1$ in position i with probability $1/(1 + e^\beta)$;
 - update chain X_t at -1 and chain Y_t at $+1$ in position i with probability $1/(1 + e^{-\beta}) - 1/2$.
- If $|\mathbf{y}| = n$ or both $|\mathbf{y}| = n-1$ and X_t has -1 in position i , then
 - update both chains at $+1$ in position i with probability $1/2$;
 - update both chains at -1 in position i with probability $1/(1 + e^\beta)$;
 - update chain X_t at $+1$ and chain Y_t at -1 in position i with probability $1/(1 + e^{-\beta}) - 1/2$.

By construction we have that (X_t, Y_t) is a coupling of $P^t(\mathbf{x}, \cdot)$ and U^* . Then

$$\|P^t(\mathbf{x}, \cdot) - U\|_{\text{TV}} \leq \|P^t(\mathbf{x}, \cdot) - U^*\|_{\text{TV}} + \frac{1}{2^{n-1}} \leq \mathbf{P}_{\mathbf{x}, U^*}(X_t \neq Y_t) + \frac{1}{2^{n-1}}, \quad (6)$$

where the last inequality follows from Theorem A.1. Moreover observe that, if at time t all players have been selected at least once and chain X_t has not yet hit profiles \mathbf{p} or \mathbf{m} , then the two random variables X_t and Y_t have the same value. Hence

$$\mathbf{P}_{\mathbf{x}, U^*}(X_t \neq Y_t) \leq \mathbf{P}_{\mathbf{x}}(\tau_{\mathbf{p}, \mathbf{m}} \leq t \cup \eta > t) \leq \mathbf{P}_{\mathbf{x}}(\tau_{\mathbf{p}, \mathbf{m}} \leq t) + \mathbf{P}_{\mathbf{x}}(\eta > t),$$

where $\tau_{\mathbf{p}, \mathbf{m}}$ is the hitting time of \mathbf{p} or \mathbf{m} for chain X_t , and η is the first time all players have been selected at least once.

From the coupon collector's argument (see, e.g., [29, Proposition 2.4]) it follows that for every $t \geq n \log(3n/\varepsilon)$

$$\mathbf{P}_{\mathbf{x}}(\eta > t) \leq \varepsilon/3. \quad (7)$$

Hence, for $t = n \log n + n \log(3/\varepsilon)$, by combining (6), (7) and Lemma 4.2 it holds that

$$\|P^t(\mathbf{x}, \cdot) - U\|_{\text{TV}} \leq \frac{1}{2^{n-1}} + \frac{\varepsilon}{3} + \frac{c}{n} \leq \frac{\varepsilon}{2},$$

for n sufficiently large. □

Finally, we have the following theorem.

Theorem 4.4. *For $\beta = \omega(\log n)$, for every $\varepsilon > 0$ and for every $\mathbf{x} \in \{+1, -1\}^n$, there exists function $T(n) = T_{\varepsilon, \mathbf{x}}(n)$ and distribution μ such that $T(n)$ is super-polynomial in n , μ is $(\varepsilon, T(n))$ -metastable and pseudo-mixing time $t_{\mu}^{\{\mathbf{x}\}}(\varepsilon)$ is polynomial in n and in $\log 1/\varepsilon$.*

Proof. For $\mathbf{x} \neq \mathbf{p}, \mathbf{m}$, the theorem follows from Lemma 4.3. Moreover, if the dynamics start from profile \mathbf{p} or \mathbf{m} it immediately reaches the metastable distributions π_+ or π_- . □

5 Metastability of Graphical Coordination games

Consider n players identified by the vertices of a graph $G = (V, E)$. For each edge of G , we have an instance of the two-strategy coordination game defined above, played by the endpoints of the edge. Each player picks a strategy from the set of strategy $\{\pm 1\}$ and uses it for each two-strategy coordination game in which she is involved. The utility of a player is the sum of the utilities for each two-strategy coordination game she plays. Graphical coordination games are often used to model the spread of a new technology in a social network [35, 34] with strategy $+1$ corresponding to adopting the new technology. It is easy to see that the graphical coordination game on a graph G is a potential game with potential function $\Phi(\mathbf{x}) = \sum_{e=(u,v) \in E} \Phi_e(\mathbf{x})$, where $\Phi_e(\mathbf{x})$ is the potential of the two-strategy coordination game associated with edge e .

We start by describing two distributions that turn out to be metastable for any graphical coordination game.

Lemma 5.1. *Consider the logit dynamics with parameter β for the coordination game on a graph G with minimum degree d_{\min} . Then, for any $\varepsilon > 0$, the distribution π_+ concentrated in profile $\mathbf{p} = (+1)^n$ is $(e^{-\beta \cdot \Delta \cdot d_{\min}}, 1)$ -metastable. Similarly, distribution π_- concentrated in profile $\mathbf{m} = (-1)^n$ is $(e^{-\beta \cdot \delta \cdot d_{\min}}, 1)$ -metastable.*

Proof. The bottleneck ratio of \mathbf{p} is

$$B(\mathbf{p}) = \frac{\sum_{i=1}^n \pi(\mathbf{p}) P(\mathbf{p}, (\mathbf{p}_{-i}, -1))}{\pi(\mathbf{p})} = \sum_{i=1}^n \frac{1}{n} \frac{1}{1 + e^{-\beta(\Phi(\mathbf{p}) - \Phi(\mathbf{p}_{-i}, -1))}}.$$

An edge adjacent to player i has potential Δ in \mathbf{p} , whereas in $(\mathbf{p}_{-i}, -1)$ it has potential 0. Therefore $\Phi(\mathbf{p}) - \Phi(\mathbf{p}_{-i}, -1) = \Delta \cdot d_i$, where d_i is the degree of i . Hence

$$B(\mathbf{p}) = \sum_{i=1}^n \frac{1}{n} \frac{1}{1 + e^{\beta \Delta \cdot d_i}} \leq \frac{1}{1 + e^{\beta \cdot \Delta \cdot d_{\min}}}.$$

Thus by Lemma 3.4, we have that π_+ is $((1 + e^{\beta \cdot \Delta \cdot d_{\min}})^{-1}, 1)$ -metastable and the theorem follows. The metastability of π_- can be proved in similar way. \square

The above result should not be surprising as it simply says that if players happen to be coordinated then they will stay so for exponentially long (in the parameter β). This however does not conclude the study of metastability of graphical coordination games as nothing is said about the behavior of the logit dynamics for different initial states. More specifically, the following question remains unanswered: do we have for each (class of) initial state a metastable distribution that is quickly reached? That is we look for (ε, T) -metastable distributions with small ε , large T and small pseudo-mixing time. We positively answer this question for the two extremal graphs in terms of connectivity: the complete graph and the ring (that are the graphical coordination games that have been firstly studied in [18]).

5.1 The Curie-Weiss game

The *Curie-Weiss game* is the game-theoretic formulation of the well-studied Curie-Weiss model (the Ising model on the complete graph), used by physicists to model the interaction between magnets in a ferro-magnetic system [32]. This game is also a *population game*. Hence, the metastable states of the logit dynamics for this game has been object of intense analysis by both physicists and economists [12, 36, 37]. Our approach complements these works by focusing on metastable distributions.

Specifically, the Curie-Weiss game is a special graphical coordination game where the graph G is a clique and each player is involved in a two-strategy coordination game with $a = b = 1$ and $c = d = -1$ with every other player. Note that for this game the utility of player i at profile $\mathbf{x} = (x_1, \dots, x_n) \in \{\pm 1\}^n$ can be written as $u_i(\mathbf{x}) = x_i \sum_{j \neq i} x_j$. Similarly, the potential function H can be written as $H(\mathbf{x}) = \sum_{(j,k)} x_j x_k$. Thus, the Markov chain of the logit dynamics for the Curie-Weiss game with parameter β has stationary distribution $\pi(\mathbf{x}) = e^{\beta H(\mathbf{x})} / Z$, where $Z = \sum_{\mathbf{y}} e^{\beta H(\mathbf{y})}$ is the partition function.

The *magnetization* of \mathbf{x} is defined as $S(\mathbf{x}) = \sum_{i=1}^n x_i$ and we observe that the potential of a profile \mathbf{x} depends only on its magnetization. More precisely, if $S(\mathbf{x}) = k$ then $H(\mathbf{x}) = H(k) := \frac{1}{2}(k^2 - n)$.

To see this, let us name p and m the number of $+1$ and -1 respectively, in profile \mathbf{x} , and observe that $p - m = S(\mathbf{x}) = k$ and $p + m = n$. Each pair of players with the same sign contributes for $+1$ to $H(\mathbf{x})$ and each pair of players with opposite signs contributes for -1 ; since there are $\binom{p}{2}$ pairs where both players play $+1$, $\binom{m}{2}$ pairs where both play -1 and $p \cdot m$ pairs where players play opposite strategies, we have that

$$H(\mathbf{x}) = \binom{p}{2} + \binom{m}{2} - p \cdot m = \frac{1}{2}((p - m)^2 - (p + m)).$$

The mixing time of the logit dynamics for the Curie-Weiss game is known to be exponential for every $\beta > 1/n$ (see Theorem 15.3 in [29]). In this section we show that, if the Markov chain starts from a profile where the number of $+1$ (respectively -1) is a sufficiently large majority and if β is large enough then, after a polynomial *pseudo-mixing* time, the distribution of the chain is close, in total variation distance, to π_+ (respectively, π_-). Note that from Lemma 5.1 and Lemma 3.2, since $\Delta = 2$ and $d_{\min} = n - 1$, it follows that π_+ and π_- are metastable for a time longer than any polynomial for any $\beta = \omega(\log n/n)$.

Let X_t be the logit dynamics for the Curie-Weiss game, and consider the magnetization process $S_t := S(X_t)$. Observe that S_t is itself a Markov chain, with state space $\Omega = \{-n, -n + 2, \dots, n - 4, n - 2, n\}$. When at state $k \in \Omega$, the probability to go right (to state $k + 2$) or left (to state $k - 2$) is respectively

$$\mathbf{P}_k(S_1 = k + 2) = p_k = \frac{n - k}{2n} \frac{1}{1 + e^{-2(k+1)\beta}}; \quad \mathbf{P}_k(S_1 = k - 2) = q_k = \frac{n + k}{2n} \frac{1}{1 + e^{2(k-1)\beta}}. \quad (8)$$

Indeed, let us evaluate the probability to jump from a profile \mathbf{x} with magnetization k to a profile with magnetization $k + 2$. If $S(\mathbf{x}) = k$ then there are $(n + k)/2$ players playing $+1$ and $(n - k)/2$ players playing -1 . The chain moves to a profile with magnetization $k + 2$ if a player playing -1 is selected, this happens with probability $(n - k)/2n$, and she updates her strategy to $+1$, this happens with probability

$$\frac{e^{\beta u_i(\mathbf{x}_{-i}, +1)}}{e^{\beta u_i(\mathbf{x}_{-i}, +1)} + e^{\beta u_i(\mathbf{x}_{-i}, -1)}} = \frac{1}{1 + e^{\beta[u_i(\mathbf{x}_{-i}, -1) - u_i(\mathbf{x}_{-i}, +1)]}}.$$

Finally observe that $u_i(\mathbf{x}_{-i}, -1) - u_i(\mathbf{x}_{-i}, +1) = -2 \sum_{j \neq i} x_j = -2(S(\mathbf{x}) - x_i) = -2(k + 1)$.

For $a, b \in [-n, n]$, with $a < b$, let $\tau_{a,b}$ be the random variable indicating the first time the magnetization chain reaches a state k with $k \leq a$ or $k \geq b$, that is

$$\tau_{a,b} = \min \{t \in \mathbb{N} : S_t \leq a \text{ or } S_t \geq b\}.$$

At time $\tau_{a,b}$, chain $S_{\tau_{a,b}}$ can be in one out of two states, namely the largest state smaller than or equal to a or the smallest state larger than or equal to b . The following technical lemmas give an upper bound on the probability that when the chain exits from interval (a, b) , it happens on the left side of the interval.

In the first lemma we show that, if the chain starts from a sufficiently large positive state k , and if $\beta k^2 \geq c \log n$ for a suitable constant c , then when chain S_t gets out of interval $(0, n/2)$, it happens on the $n/2$ side with high probability.

Lemma 5.2. *Let $k \in \Omega$ be the starting state with $4 \leq k \leq n/2$. If $\beta \geq 6/n$ and $\beta k^2 \geq 16 \log n$, then*

$$\mathbf{P}_k(S_{\tau_{0, n/2}} \leq 0) \leq 1/n.$$

Proof. According to (8), the ratio of q_h and p_h is

$$\frac{q_h}{p_h} = \frac{n + h}{n - h} \cdot \frac{1 + e^{-2(h+1)\beta}}{1 + e^{2(h-1)\beta}}.$$

Now observe that for all $h \geq 2$ it holds that

$$\frac{1 + e^{-2(h+1)\beta}}{1 + e^{2(h-1)\beta}} \leq e^{-2(h-1)\beta} \leq e^{-h\beta}, \quad (9)$$

and for all $h \leq n/2$ it holds that

$$\frac{n + h}{n - h} = \frac{1 + h/n}{1 - h/n} \leq e^{3h/n}.$$

Hence, for every $2 \leq h \leq n/2$ we can give the following upper bound

$$\frac{q_h}{p_h} \leq e^{3h/n} \cdot e^{-\beta h} = e^{-(\beta-3/n)h} \leq e^{-\frac{1}{2}\beta h}, \quad (10)$$

where in the last inequality we used $\beta \geq 6/n$.

Thus, for each state h of the chain with $k/2 \leq h \leq n/2$ we have that the ratio q_h/p_h is less than $e^{-\frac{1}{4}\beta k}$. If the chain starts at k , by applying Lemma B.3 it follows that the probability of reaching $k/2$ before reaching $n/2$ is less than $\left(e^{-\frac{1}{4}\beta k}\right)^\ell$, where ℓ is the number of states between $k/2$ and k , that is $\ell = k/4$. Hence, for every $4 \leq k \leq n/2$, if $\beta k^2 \geq 16 \log n$, the chain starting at k hits state $n/2$ before state $k/2$ with probability

$$\mathbf{P}_k(S_{\tau_{k/2, n/2}} \leq k/2) \leq e^{-\frac{1}{16}\beta k^2} \leq \frac{1}{n}.$$

The thesis follows by observing that $\mathbf{P}_k(S_{\tau_{0, n/2}} \leq 0) \leq \mathbf{P}_k(S_{\tau_{k/2, n/2}} \leq k/2)$. \square

In the next lemma we show that, if the chain starts from a state $k \geq n/2$, and if $\beta \geq c \log n/n$ for a suitable constant c , then when chain S_t reaches one of the endpoints of interval $(0, n)$ it is on the n side with probability exponentially close to 1.

Lemma 5.3. *Let $k \in \Omega$ be the starting state with $n/2 \leq k \leq n-1$. If $\beta \geq 8 \log n/n$, then*

$$\mathbf{P}_k(S_{\tau_{0, n}} \leq 0) \leq (2/n)^{n/8}.$$

Proof. Observe that for every $h \leq n-1$ it holds that $\frac{n+h}{n-h} \leq 2n$, and by using it together with (9) we have that $q_h/p_h \leq 2ne^{-\beta h}$ for every $2 \leq h \leq n-1$. Thus, for every $k/2 \leq h \leq n-1$ it holds that

$$q_h/p_h \leq 2ne^{-\frac{1}{2}\beta k} \leq 2/n, \quad (11)$$

where in the last inequality we used $k \geq n/2$ and $\beta \geq 8 \log n/n$. Hence, if the chain starts at k , by applying Lemma B.3 it follows that the probability of reaching $k/2$ before reaching n is less than $(2/n)^\ell$, where $\ell = k/4 \geq n/8$ is the number of states between $k/2$ and k . Hence,

$$\mathbf{P}_k(Y_{\tau_{0, n}} \leq 0) \leq \mathbf{P}_k(Y_{\tau_{k/2, n}} \leq k/2) \leq (2/n)^{n/8}. \quad \square$$

Finally, in the next lemma we show that for every starting state between 0 and n , the expected time the chain reaches 0 or n is at most $\mathcal{O}(n^3)$.

Lemma 5.4. *For every $k \in \Omega$ with $k \geq 0$ it holds that $\mathbf{E}_k[\tau_{0, n}] \leq n^3$.*

Proof. Consider the birth-and-death chain $\{Y_t^*\}$ on $\Omega_+ = \{l \in \Omega \mid l \geq 0\}$ and let τ_n^* the hitting time of n in this chain. It is obvious that $\mathbf{E}_k[\tau_{0, n}] \leq \mathbf{E}_k[\tau_n^*]$. It is well-known (see, for example, Section 2.5 in [29]) that

$$\mathbf{E}_k[\tau_n^*] = \sum_{l=\frac{k+2}{2}}^{n/2} \frac{1}{q_{2l}w_{2l}} \sum_{j=\frac{n \bmod 2}{2}}^{l-1} w_{2j},$$

where $w_{n \bmod 2} = 1$ and $w_{2j} = \prod_{i=1}^j p_{2(j-1)}/q_{2j}$. From simple computations, we obtain

$$\begin{aligned} \mathbf{E}_k[\tau_n^*] &= \sum_{l=\frac{k+2}{2}}^{n/2} \sum_{j=\frac{n \bmod 2}{2}}^{l-1} \frac{1}{p_{2j}} \prod_{i=j+1}^{l-1} \frac{q_{2i}}{p_{2i}} \\ &\leq \sum_{l=\frac{k+2}{2}}^{n/2} \sum_{j=\frac{n \bmod 2}{2}}^{l-1} \frac{1}{p_{2j}}, \end{aligned}$$

where the inequality follows from (10) and (11). Finally, the Lemma follows by observing that $p_{2j} \geq \frac{1}{2n}$ for every $j \geq 0$. \square

Now we can state and prove the main theorem of this section.

Theorem 5.5. *Let \mathbf{x} be a profile of magnetization $S(\mathbf{x}) = k$. If $\beta \geq 8 \log n / n$ and $k^2 > 16 \log n / \beta$ then $t_{\pi+}^{\{\mathbf{x}\}}(\varepsilon) \leq n^4$ if k is positive and $t_{\pi-}^{\{\mathbf{x}\}}(\varepsilon) \leq n^4$ if k is negative.*

Proof. Consider without loss of generality the case of starting state with positive magnetization, $S(\mathbf{x}) = k \geq 0$. Let τ_n be the first time the chain hits state with all +1 and recall that $\tau_{0,n}$ is the first time the magnetization of the chain is either n or less than or equal to 0. Since $\{\tau_n > t, \tau_{0,n} \leq t\}$ implies that the magnetization chain reaches 0 before reaching n we have

$$\begin{aligned} \mathbf{P}_{\mathbf{x}}(\tau_n > t) &= \mathbf{P}_{\mathbf{x}}(\tau_n > t, \tau_{0,n} > t) + \mathbf{P}_{\mathbf{x}}(\tau_n > t, \tau_{0,n} \leq t) \\ &\leq \mathbf{P}_{\mathbf{x}}(\tau_{0,n} > t) + \mathbf{P}_{\mathbf{x}}(S_{\tau_{0,n}} \leq 0) \\ &\leq \frac{\mathbf{E}_{\mathbf{x}}[\tau_{0,n}]}{t} + \mathbf{P}_{\mathbf{x}}(S_{\tau_{0,n}} \leq 0). \end{aligned}$$

As for the first term of the sum, from Lemma 5.4 it follows that $\mathbf{E}_{\mathbf{x}}[\tau_{0,n}]/t \leq 1/n$ for $t \geq n^4$. As for the second term, by conditioning on the position of the chain when it gets out of subinterval $(0, n/2)$ we have

$$\begin{aligned} \mathbf{P}_k(S_{\tau_{0,n}} \leq 0) &= \mathbf{P}_k(S_{\tau_{0,n}} \leq 0 \mid Y_{\tau_{0,n/2}} \leq 0) \mathbf{P}_k(S_{\tau_{0,n/2}} \leq 0) + \\ &\quad + \mathbf{P}_k(S_{\tau_{0,n}} \leq 0 \mid S_{\tau_{0,n/2}} \geq n/2) \mathbf{P}_k(S_{\tau_{0,n/2}} \geq n/2) \\ &\leq \mathbf{P}_k(S_{\tau_{0,n/2}} \leq 0) + \mathbf{P}_k(S_{\tau_{0,n}} \leq 0 \mid S_{\tau_{0,n/2}} \geq n/2). \end{aligned}$$

From Lemma 5.2 we have that $\mathbf{P}_k(S_{\tau_{0,n/2}} \leq 0) \leq 1/n$, and observe that

$$\mathbf{P}_k(S_{\tau_{0,n}} \leq 0 \mid S_{\tau_{0,n/2}} \geq n/2) \leq \mathbf{P}_{n/2}(S_{\tau_{0,n}} \leq 0) \leq (2/n^{n/8}),$$

where the last inequality follows from Lemma 5.3. Hence for every $t \geq n^4$ it holds that $\mathbf{P}_{\mathbf{x}}(\tau_n > t) \leq 3/n$. The thesis then follows. \square

5.2 Graphical Coordination games on rings

In this section we study graphical coordination games on a ring.

In [7] it is showed that the mixing time of the logit dynamics for this game is $\Omega(e^{2\delta\beta})$. If $\Delta > \delta$, techniques of [8] can be generalized to obtain an upper bound to the mixing time that is polynomial in n and exponential in β . If, instead, $\Delta = \delta$, an almost matching upper bound is given in [7]. These results show that the mixing time is polynomial in n for $\beta = \mathcal{O}(\log n)$ and greater than any polynomial in n , for $\beta = \omega(\log n)$. In this section we shall show that, for these large values of β , despite the large mixing time, from each starting point the logit dynamics quickly approach a metastable distribution. Specifically, our result is summarized by the following theorem.

Theorem 5.6. *Consider the logit dynamics with parameter β for an n -player coordination game on a ring. Then, for n sufficiently large, for $\beta = \omega(\log n)$, for every $\varepsilon > 0$ and for every $\mathbf{x} \in \{\pm 1\}^n$, there is a $(\varepsilon, T(n))$ -metastable distribution μ , with T super-polynomial in n , such that the pseudo-mixing time $t_{\mu}^{\{\mathbf{x}\}}(\varepsilon)$ is polynomial in n and in $1/\varepsilon$.*

We will show that, depending on the starting state \mathbf{x} , the Markov chain \mathcal{M}_{β} of the logit dynamics with parameter β for the n -player coordination game on the ring, rapidly approaches a $(\varepsilon, \varepsilon \cdot e^{2\beta\delta})$ -metastable distribution. Specifically, our results confirm the findings of [18, 19], that the logit dynamics for this game tend to reach in the long run the metastable states, namely, the profiles in which either all players adopt strategy -1 or they all adopt with strategy $+1$. Still our approach based on metastable distributions allows us to have a finer description of the system in the medium run. For example, our results for graphical coordination games on the ring without a risk-dominant strategy allow us to understand that the probability that one of these states is reached essentially depends only on the number of players that adopt a given strategy at beginning, and not on their location on the ring.

The proof of the theorem separately considers games with risk dominant strategies (i.e., $\Delta > \delta$) and games without risk dominant strategies (i.e., $\Delta = \delta$). As for the first ones, let us first introduce some

useful notation. We define $R \subseteq \{\pm 1\}^n$ as the set of profiles in which at least two *adjacent* players are playing +1. Moreover we define $S_d \subseteq \{\pm 1\}^n$ as the set of profiles where *exactly* d non-adjacent players are playing +1, $S_{\geq d} = \bigcup_{i=d}^n S_i$ and $S_d^* = S_{\geq d} \cup R$. We remind that \mathbf{p} denotes the profile where all players are playing +1 and \mathbf{m} denotes the profile where all players are playing -1.

Our proof is based on the following three technical lemmas whose proofs are postponed to Appendix D.1, Appendix D.2 and Appendix D.3, respectively. The first two lemmas provide an upper bound to the probability that \mathcal{M}_β takes too long to hit profile \mathbf{p} when starting from $\mathbf{x} \in R$ or from $\mathbf{x} \in S_d$. The last lemma, instead, provides an upper bound to the probability that \mathcal{M}_β takes too long to hit one of the two profiles \mathbf{p} and \mathbf{m} for any starting state.

More precisely, recalling that $\tau_{\mathbf{p}}$ denotes the hitting time of profile \mathbf{p} , and $\tau_{\{\mathbf{p}, \mathbf{m}\}}$ denotes the hitting time of the set of profiles $\{\mathbf{p}, \mathbf{m}\}$, we have that

Lemma 5.7. *For all $\mathbf{x} \in R$ and $\varepsilon > 0$, if $\Delta > \delta$, $\beta = \omega(\log n)$ and n is sufficiently large*

$$P_{\mathbf{x}} \left(\tau_{\mathbf{p}} > \frac{8 - \varepsilon}{\varepsilon} \cdot n^2 \right) \leq \frac{\varepsilon}{4}.$$

Lemma 5.8. *For all $d > 0$, $\mathbf{x} \in S_d$ and $\varepsilon > 0$, if $\Delta > \delta$, $\beta = \omega(\log n)$ and n is sufficiently large*

$$P_{\mathbf{x}} \left(\tau_{\mathbf{p}} > \frac{8}{\varepsilon} \cdot n^2 \right) \leq \frac{1}{2^d + 1} + \frac{\varepsilon}{4}.$$

Lemma 5.9. *For all $\mathbf{x} \in \{\pm 1\}^n$ and $\varepsilon > 0$, if $\Delta > \delta$, $\beta = \omega(\log n)$ and n is sufficiently large*

$$P_{\mathbf{x}} \left(\tau_{\{\mathbf{p}, \mathbf{m}\}} > \frac{8}{\varepsilon} \cdot n^2 \right) \leq \frac{\varepsilon}{2}.$$

As for the case that no risk dominant strategy exists, the proof is based on a technical lemma, whose proof is postponed to Appendix D.4, that shows that with high probability the hitting time of the set $\{\mathbf{p}, \mathbf{m}\}$ is polynomial in the number of players.

Lemma 5.10. *For all $\mathbf{x} \in \{\pm 1\}^n$ and $\varepsilon > 0$, if $\Delta = \delta$, $\beta = \omega(\log n)$ and n is sufficiently large*

$$P_{\mathbf{x}} (\tau_{\{\mathbf{p}, \mathbf{m}\}} > n^5) \leq o(1).$$

We are now ready to prove Theorem 5.6.

Proof (of Theorem 5.6). We start showing that for each possible initial state, the pseudo-mixing time of the Markov chain \mathcal{M}_β to a $(\varepsilon, \varepsilon \cdot T(n))$ -metastable distribution, for a function T super-polynomial in n , is polynomial in n and in $1/\varepsilon$. In particular, we will show that \mathcal{M}_β approaches a convex combination of π_+ and π_- . These two distributions, from Lemma 5.1 and since in a ring any vertex has degree 2, are $(e^{-2\delta\beta}, 1)$ -metastable. Then, from Lemma 3.2 it follows that they are $(\varepsilon, \varepsilon \cdot e^{2\delta\beta})$ -metastable. Consequently, by Lemma 3.3, any their convex combination is $(\varepsilon, \varepsilon \cdot e^{2\delta\beta})$ -metastable. Finally, since $\beta = \omega(\log n)$, it follows that $\varepsilon \cdot e^{2\delta\beta}$ is super-polynomial in n as desired.

For the pseudo-mixing time we distinguish three cases.

$\Delta > \delta$, $\mathbf{x} \in S_d^*$ and $d \geq \log_2 \left(\frac{4}{\varepsilon} - 1 \right)$: We prove that the pseudo-mixing time $t_{\pi_+}^{S_d^*}(\varepsilon)$ to π_+ from any $\mathbf{x} \in S_d^*$, with $d \geq \log_2 \left(\frac{4}{\varepsilon} - 1 \right)$, is polynomial in n and $1/\varepsilon$. Indeed, by Lemma 5.7 and Lemma 5.8 we have that

$$\max_{\mathbf{x} \in S_d^*} P_{\mathbf{x}} \left(\tau_{\mathbf{p}} > \frac{8}{\varepsilon} \cdot n^2 \right) \leq \frac{\varepsilon}{2}.$$

Moreover, from Lemma 5.1 and Lemma 3.2 it follows that π_+ is $(\varepsilon/2, \varepsilon/2 \cdot e^{2\delta\beta})$ -metastable. Thus, for sufficiently large n , since $\frac{8}{\varepsilon} \cdot n^2 \leq \frac{\varepsilon}{2} \cdot e^{2\delta\beta}$, π_+ is $(\frac{\varepsilon}{2}, \frac{8}{\varepsilon} \cdot n^2)$ -metastable. Then, by Lemma 3.8, we obtain that

$$d_{\pi_+}^{S_d^*} \left(\frac{8}{\varepsilon} \cdot n^2 \right) \leq \varepsilon.$$

Then, from the definition of pseudo-mixing time, we obtain that

$$t_{\pi_+}^{S_d^*}(\varepsilon) \leq \frac{8}{\varepsilon} \cdot n^2.$$

$\Delta > \delta$, $\mathbf{x} \in S_d$ **and** $d < \log_2 \left(\frac{4}{\varepsilon} - 1\right)$: In this case we prove that the pseudo-mixing time to the distribution

$$\mu_{\mathbf{x}} = \alpha_{\mathbf{x}}\pi_+ + (1 - \alpha_{\mathbf{x}})\pi_-,$$

where

$$\alpha_{\mathbf{x}} = \mathbf{P}_{\mathbf{x}} \left(\tau_{\mathbf{p}} \leq \tau_{\mathbf{m}} \mid \tau_{\mathbf{p}, \mathbf{m}} \leq \frac{8}{\varepsilon} \cdot n^2 \right) \quad (12)$$

is polynomial. Indeed from Lemma 5.9 we have that

$$P_{\mathbf{x}} \left(\tau_{\{\mathbf{p}, \mathbf{m}\}} > \frac{8}{\varepsilon} \cdot n^2 \right) \leq \frac{\varepsilon}{2}.$$

Moreover, from Lemma 5.1 and Lemma 3.2 it follows that π_+ and π_- are both $(\varepsilon/2, \varepsilon/2 \cdot e^{2\delta\beta})$ -metastable. Thus, for sufficiently large n , since $\frac{8}{\varepsilon} \cdot n^2 \leq \frac{\varepsilon}{2} \cdot e^{2\delta\beta}$, π_+ and π_- are both $(\frac{\varepsilon}{2}, \frac{8}{\varepsilon} \cdot n^2)$ -metastable. Then, from Lemma 3.9 we obtain that

$$d_{\mu_{\mathbf{x}}}^{\{\mathbf{x}\}} \left(\frac{8}{\varepsilon} \cdot n^2 \right) \leq \varepsilon.$$

Finally, from the definition of pseudo-mixing time, we have that

$$t_{\mu_{\mathbf{x}}}^{\{\mathbf{x}\}}(\varepsilon) \leq \frac{8}{\varepsilon} \cdot n^2.$$

$\Delta = \delta$: This case can be proved similarly to the previous case, by using Lemma 5.10 in place of Lemma 5.9. It is surprising that, if $\Delta = \delta$, the coefficient $\alpha_{\mathbf{x}}$ defined in (12), for any $\mathbf{x} \in \{\pm 1\}^n$, depends “almost” only on the number of players selecting strategy +1 in \mathbf{x} (see Lemma D.14 in Appendix D.5). Thus, for example, the metastable behavior of profiles with d adjacent players with strategy +1 is almost the same as profiles where the same players are far from each other. \square

6 Conclusions and open problems

Logit dynamics are clean and tractable dynamics that well model the behaviour of limited-rationality players in a strategic game. The stationary distribution of the induced Markov chain is the natural long-term equilibrium concept for games under logit dynamics. However, when the mixing time is long, the behavior of the Markov chain in the transient phase becomes important and it is worth looking for “regularities” at a time-scale shorter than mixing time. Such regularities have been previously explored, for some classes of Markov chains, by means of “metastable states”. We believe that a more general and useful concept is that of “metastable distributions”.

In this paper we defined a quantitative notion of metastable distribution and we analyzed the metastability properties of the logit dynamics for some classes of coordination games. We showed that, even when the mixing time is exponential, it is often possible to find some distributions that well-approximate the distribution of the chain for a time-window of super-polynomial size. Such metastable distributions can be found even in the case of the OR-game, where no partition of the state space in metastable *states* exists. A natural open question is whether the metastability properties for coordination games we observed in this paper hold in general for potential games.

In the case of the Curie-Weiss model on the complete graph, we showed that when $\beta > c \log n/n$ the two degenerate distributions are metastable for $\text{poly}(n)$ time and they are quickly reached from a large fraction of the state space. We note that the metastability properties when $1/n < \beta < \log n/n$ have been investigated in [21], where it is proved that, even in this case, for every starting profile the logit dynamics quickly converge to a distribution that is metastable for long time.

Acknowledgment. We wish to thank Paolo Penna for useful ideas, hints, and discussions. We also thank the anonymous referees for their comments and suggestions.

References

- [1] Carlos Alós-Ferrer and Nick Netzer. The logit-response dynamics. *Games and Economic Behavior*, 68(2):413–427, 2010.
- [2] Carlos Alós-Ferrer and Nick Netzer. Robust stochastic stability. *Economic Theory*, 58(1):31–57, 2015.
- [3] Arash Asadpour and Amin Saberi. On the inefficiency ratio of stable equilibria in congestion games. In *Proc. of the 5th Int. Workshop on Internet and Network Economics (WINE'09)*, volume 5929 of *Lecture Notes in Computer Science*, pages 545–552. Springer, 2009.
- [4] Vincenzo Auletta, Diodato Ferraioli, Francesco Pasquale, and Giuseppe Persiano. Metastability of logit dynamics for coordination games. In *Proceedings of the Twenty-Third Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2012, Kyoto, Japan, January 17-19, 2012*, pages 1006–1024, 2012.
- [5] Vincenzo Auletta, Diodato Ferraioli, Francesco Pasquale, Paolo Penna, and Giuseppe Persiano. Logit dynamics with concurrent updates for local interaction games. In *Algorithms - ESA 2013 - 21st Annual European Symposium, Sophia Antipolis, France, September 2-4, 2013. Proceedings*, pages 73–84, 2013.
- [6] Vincenzo Auletta, Diodato Ferraioli, Francesco Pasquale, and Giuseppe Persiano. Mixing time and stationary expected social welfare of logit dynamics. *Theory Comput. Syst.*, 53(1):3–40, 2013.
- [7] Vincenzo Auletta, Diodato Ferraioli, Francesco Pasquale, Paolo Penna, and Giuseppe Persiano. Convergence to equilibrium of logit dynamics for strategic games. *Algorithmica*, 76(1):110–142, 2016.
- [8] Noam Berger, Claire Kenyon, Elchanan Mossel, and Yuval Peres. Glauber dynamics on trees and hyperbolic graphs. *Probability Theory and Related Fields*, 131:311–340, 2005.
- [9] Alessandra Bianchi and Alexandre Gaudilliere. Metastable states, quasi-stationary and soft measures, mixing time asymptotics via variational principles. *arXiv preprint arXiv:1103.1143*, 2011.
- [10] Lawrence E. Blume. The statistical mechanics of strategic interaction. *Games and Economic Behavior*, 5:387–424, 1993.
- [11] Holly Borowski, Jason R. Marden, and Eric W. Frew. Fast convergence in semi-anonymous potential games. In *Decision and Control (CDC), 2013 IEEE 52nd Annual Conference on*, pages 2418–2423. IEEE, 2013.
- [12] Anton Bovier, Michael Eckhoff, Véronique Gayrard, and Markus Klein. Metastability in stochastic dynamics of disordered mean-field models. *Probability Theory and Related Fields*, 119:99–161, 2001.
- [13] Anton Bovier, Michael Eckhoff, Véronique Gayrard, and Markus Klein. Metastability and low lying spectra in reversible Markov chains. *Communications in Mathematical Physics*, 228:219–255, 2002.
- [14] Colin F. Camerer. *Behavioral Game Theory: Experiments in Strategic Interaction*. 2003.
- [15] Pierre Collet, Servet Martínez, and Jaime San Martín. *Quasi-Stationary Distributions: Markov Chains, Diffusions and Dynamical Systems*. Springer, 2012.
- [16] Jian Ding, Eyal Lubetzky, and Yuval Peres. The mixing time evolution of Glauber dynamics for the mean-field Ising model. *Communications in Mathematical Physics*, 289:725–764, 2009.
- [17] Jian Ding, Eyal Lubetzky, and Yuval Peres. Censored Glauber dynamics for the mean field Ising model. *Journal of Statistical Physics*, 137:407–458, 2009.
- [18] Glenn Ellison. Learning, local interaction, and coordination. *Econometrica*, 61(5):1047–1071, 1993.
- [19] Glenn Ellison. Basins of Attraction, Long-Run Stochastic Stability, and the Speed of Step-by-Step Evolution. *Review of Economic Studies*, 67(1):17–45, 2000.

- [20] Alex Fabrikant, Christos H. Papadimitriou, and Kunal Talwar. The complexity of pure nash equilibria. In *STOC*, pages 604–612, 2004.
- [21] Diodato Ferraioli and Carmine Ventre. Metastability of asymptotically well-behaved potential games. In *Mathematical Foundations of Computer Science 2015*, pages 311–323. Springer, 2015.
- [22] Mark Freidlin and Leonid Koralov. Metastable distributions of markov chains with rare transitions. *Journal of Statistical Physics*, pages 1–21, 2017.
- [23] Mark Iosifovich Freidlin and Alexander D. Wentzell. *Random Perturbations of Dynamical Systems*. Springer, 1984.
- [24] Serge Galam and Bernard Walliser. Ising model versus normal form game. *Physica A: Statistical Mechanics and its Applications*, 389(3):481 – 489, 2010.
- [25] John C. Harsanyi and Reinhard Selten. *A General Theory of Equilibrium Selection in Games*. MIT Press, 1988.
- [26] Julian Keilson. *Markov chain modelsrarity and exponentiality*, volume 28. Springer Science & Business Media, 2012.
- [27] Gabriel E. Kreindler and H. Peyton Young. Fast convergence in evolutionary equilibrium selection. *Games and Economic Behavior*, 80:39–67, 2013.
- [28] Gabriel E. Kreindler and H. Peyton Young. Rapid innovation diffusion in social networks. *Proceedings of the National Academy of Sciences*, 111(Supplement 3):10881–10888, 2014.
- [29] David Levin, Yuval Peres, and Elizabeth L. Wilmer. *Markov Chains and Mixing Times*. American Mathematical Society, 2008.
- [30] David Levin, Malwina Luczak, and Yuval Peres. Glauber dynamics for the mean-field Ising model: cut-off, critical power law, and metastability. *Probability Theory and Related Fields*, 146:223–265, 2010.
- [31] Jason R. Marden and Jeff S. Shamma. Revisiting log-linear learning: Asynchrony, completeness and payoff-based implementation. In *Communication, Control, and Computing (Allerton), 2010 48th Annual Allerton Conference on*, pages 1171–1172. IEEE, 2010.
- [32] Fabio Martinelli. Lectures on Glauber dynamics for discrete spin models. In *Lectures on Probability Theory and Statistics (Saint-Flour, 1997)*, volume 1717 of *Lecture Notes in Math.*, pages 93–191. Springer, 1999.
- [33] Dov Monderer and Lloyd S. Shapley. Potential games. *Games and Economic Behavior*, 14:124–143, 1996.
- [34] Andrea Montanari and Amin Saberi. Convergence to equilibrium in local interaction games. In *Proc. of the 50th Ann. Symp. on Foundations of Computer Science (FOCS’09)*. IEEE, 2009.
- [35] H. Peyton Young. The diffusion of innovations in social networks. Economics Working Paper Archive number 437, Johns Hopkins University, Department of Economics, 2000.
- [36] William H. Sandholm. *Population Games and Evolutionary Dynamics*. MIT Press, 2010.
- [37] William H. Sandholm and Mathias Staudigl. Large deviations and stochastic stability in the small noise double limit. *Theoretical Economics*, 11(1):279–355, 2016.
- [38] Devavrat Shah and Jinwoo Shin. Dynamics in congestion games. In *ACM SIGMETRICS Performance Evaluation Review*, volume 38, pages 107–118. ACM, 2010.
- [39] Howard M Taylor and Samuel Karlin. *An introduction to stochastic modeling*. Academic press, 2014.
- [40] H. Peyton Young. *Individual Strategy and Social Structure: An Evolutionary Theory of Institutions*. Princeton University Press, 1998.

Appendix

A Markov chain summary

In this section we recall some basic facts about Markov chains. For a more detailed treatment, we refer the reader to [29].

Total variation distance. The *total variation* distance $\|\mu - \nu\|_{\text{TV}}$ between two probability distributions μ and ν on Ω is defined as

$$\|\mu - \nu\|_{\text{TV}} := \max_{A \subseteq \Omega} |\mu(A) - \nu(A)| = \frac{1}{2} \sum_{x \in \Omega} |\mu(x) - \nu(x)|.$$

The total variation distance is actually a distance and, in particular, the triangle inequality holds. That is, the following simple fact holds.

Fact 1. *For distributions μ_1, μ_2 , and μ_3 , it holds that*

$$\|\mu_1 - \mu_3\|_{\text{TV}} \leq \|\mu_1 - \mu_2\|_{\text{TV}} + \|\mu_2 - \mu_3\|_{\text{TV}}.$$

A stochastic matrix P over Ω is a non-negative matrix in which rows and columns are indexed by elements of Ω and such that, for all $x \in \Omega$, it holds that

$$\sum_{y \in \Omega} P(x, y) = 1.$$

Fact 2. *For all distributions μ and ν on Ω and all stochastic matrices P it holds that*

$$\|\mu P - \nu P\|_{\text{TV}} \leq \|\mu - \nu\|_{\text{TV}}.$$

Mixing time. Consider a Markov chain $\mathcal{M} = \{X_t\}$ with *finite* state space Ω and transition matrix P . We stress that P is stochastic matrix over Ω and we often identify \mathcal{M} with P . It is a classical result that if \mathcal{M} is *irreducible* and *aperiodic*² (also called *ergodic*) there exists a unique *stationary distribution*; that is, a distribution π on Ω such that $\pi \cdot P = \pi$.

An ergodic Markov chain \mathcal{M} *converges* to its stationary distribution π ; specifically, there exist constants C and $0 < \alpha < 1$ such that

$$d(t) \leq C \cdot \alpha^t,$$

where

$$d(t) = \max_{x \in \Omega} \|P^t(x, \cdot) - \pi\|_{\text{TV}}$$

and $P^t(x, \cdot)$ is the distribution at time t of the Markov chain starting at x . For $0 < \varepsilon \leq 1$, the *mixing time* is defined as

$$t_{\text{mix}}(\varepsilon) = \min\{t \in \mathbb{N} : d(t) \leq \varepsilon\}.$$

It is usual to set $\varepsilon = 1/4$ and to write t_{mix} for $t_{\text{mix}}(1/4)$.

Fact 3. *For $0 < \varepsilon \leq 1$,*

$$t_{\text{mix}}(\varepsilon) \leq \lceil \log_2 \varepsilon^{-1} \rceil \cdot t_{\text{mix}}.$$

Coupling. A *coupling* of two probability distributions μ and ν on Ω is a pair of random variables (X, Y) defined on $\Omega \times \Omega$ such that the marginal distribution of X is μ and the marginal distribution of Y is ν . A well-known property of coupling is given by the following theorem (see, e.g., [29, Proposition 4.7]).

Theorem A.1. *Let μ and ν two probability distributions on Ω . Then*

$$\|\mu - \nu\|_{\text{TV}} = \inf \{\mathbf{P}(X \neq Y) : (X, Y) \text{ is a coupling of } \mu \text{ and } \nu\}$$

²Roughly speaking, a finite-state Markov chain is irreducible and aperiodic if there exists t such that, for all pairs of states x, y , the probability to be in y after t steps, starting from x , is positive.

B Biased birth-and-death chains

In this section we consider birth-and-death chains with state space $\Omega = \{0, 1, \dots, n\}$ (see Chapter 2.5 in [29] for a detailed description of such chains). For $k \in \{1, \dots, n-1\}$ let $p_k = \mathbf{P}_k(X_1 = k+1)$, $q_k = \mathbf{P}_k(X_1 = k-1)$, and $r_k = 1 - p_k - q_k = \mathbf{P}_k(X_1 = k)$. We will be interested in the probability that the chain starting at some state $h \in \Omega$ hits state n before state 0, namely $\mathbf{P}_h(X_{\tau_{0,n}} = n)$ where $\tau_{0,n} = \min\{t \in \mathbb{N} : X_t \in \{0, n\}\}$.

We start by giving an exact formula for such probability for the case when p_k and q_k do not depend on k .

Lemma B.1. *Suppose for all $k \in \{1, \dots, n-1\}$ it holds that $p_k = \varepsilon$ and $q_k = \delta$, for some ε and δ with $\varepsilon + \delta \leq 1$. Then the probability the chain hits state n before state 0 starting from state $h \in \Omega$ is*

$$\mathbf{P}_h(X_{\tau_{0,n}} = n) = \frac{1 - (\delta/\varepsilon)^h}{1 - (\delta/\varepsilon)^n}.$$

The proof of this result, that is showed below for sake of completeness, uses standard arguments (see, e.g., [39, p. 159]).

Proof. Let α_k be the probability to reach state n before state 0 starting from state k , i.e.

$$\alpha_k = \mathbf{P}_k(X_{\tau_{0,n}} = n).$$

Observe that for $k = 1, \dots, n-1$ we have

$$\alpha_k = \delta \cdot \alpha_{k-1} + \varepsilon \cdot \alpha_{k+1} + (1 - (\delta + \varepsilon)) \alpha_k. \quad (13)$$

Hence

$$\varepsilon \cdot \alpha_k - \delta \cdot \alpha_{k-1} = \varepsilon \cdot \alpha_{k+1} - \delta \cdot \alpha_k$$

with boundary conditions $\alpha_0 = 0$ and $\alpha_n = 1$. If we name $\Delta_k = \varepsilon \cdot \alpha_k - \delta \cdot \alpha_{k-1}$ we have $\Delta_k = \Delta_{k+1}$ for all k . By simple calculation and using that $\alpha_0 = 0$ it follows that

$$\alpha_k = \frac{\Delta}{\varepsilon} \sum_{i=0}^{k-1} \left(\frac{\delta}{\varepsilon}\right)^i = \frac{\Delta}{\varepsilon - \delta} (1 - (\delta/\varepsilon)^k).$$

From $\alpha_n = 1$ we get

$$\Delta = \frac{\varepsilon - \delta}{(1 - (\delta/\varepsilon)^n)}.$$

Hence

$$\alpha_k = \frac{1 - (\delta/\varepsilon)^k}{1 - (\delta/\varepsilon)^n}. \quad \square$$

Lemma B.2. *Suppose for all $k \in \{1, \dots, n-1\}$ it holds that $p_k \geq \varepsilon$ and $q_k \leq \delta$, for some ε and δ with $\varepsilon + \delta \leq 1$. Then the probability to hit state n before state 0 starting from state $h \in \Omega$ is*

$$\mathbf{P}_h(X_{\tau_{0,n}} = n) \geq \frac{1 - (\delta/\varepsilon)^h}{1 - (\delta/\varepsilon)^n}.$$

Proof. Let $\{Y_t\}$ be a birth-and-death chain with the same state space as $\{X_t\}$ but transition rates

$$\mathbf{P}_k(Y_1 = k-1) = \delta; \quad \mathbf{P}_k(Y_1 = k+1) = \varepsilon.$$

Consider the following coupling of X_t and Y_t : When (X_t, Y_t) is at state (k, h) , consider the two $[0, 1]$ intervals, each one partitioned in three subintervals as in Fig. 1. Let U be a uniform random variable over the interval $[0, 1]$ and choose the update for the two chains according to position of U in the two intervals. Observe that, since $p_k \geq \varepsilon$ and $q_k \leq \delta$, if the two chains start at the same state $h \in \Omega$, i.e. $(X_0, Y_0) = (h, h)$, then at every time t it holds that $X_t \geq Y_t$. Hence if chain Y_t hits state n before state




Figure 1: Partition for the coupling in Lemma B.2.

0, then chain X_t hits state n before state 0 as well. More formally, let $\tau_{0,n}$ and $\hat{\tau}_{0,n}$ be the random variables indicating the first time chains X_t and Y_t respectively hit state 0 or n ; then

$$\{Y_{\hat{\tau}_{0,n}} = n\} \Rightarrow \{X_{\tau_{0,n}} = n\}.$$

Thus

$$\mathbf{P}_h(X_{\tau_{0,n}} = n) \geq \mathbf{P}_h(Y_{\hat{\tau}_{0,n}} = n) \geq \frac{1 - (\delta/\varepsilon)^h}{1 - (\delta/\varepsilon)^n},$$

where in the last inequality we used Lemma B.1. \square

Lemma B.3. *Suppose for all $k \in \{1, \dots, n-1\}$ it holds that $q_k/p_k \leq \alpha$, for some $\alpha < 1$. Then the probability to hit state 0 before state n starting from state $h \in \Omega$ is*

$$\mathbf{P}_h(X_{\tau_{0,n}} = 0) \leq \alpha^h.$$

Proof. Let $\hat{p}_k = \frac{p_k}{p_k + q_k}$ and $\hat{q}_k = \frac{q_k}{p_k + q_k}$ and let $\{Y_t\}$ be the birth-and-death chain with transition rates \hat{p}_k and \hat{q}_k .

Let $\{U_t\}$ be an array of random variables such that $U_t = -1$ with probability q_{Y_t} , $U_t = +1$ with probability p_{Y_t} and $U_t = 0$ with remaining probability. We will use U_t to update chains X_t and Y_t at different time steps. Specifically, we denote with u the index of the first variables U_t not used for updating X_t (thus, at the beginning $u = 1$) and : set $Y_{t+1} = Y_t + U_t$; for chain X_t , we toss a coin that gives head with probability $p_{X_t} + q_{X_t}$ and if it gives tail we set $X_{t+1} = X_t$, otherwise we set $X_{t+1} = X_t + U_u$. Roughly speaking, we have that the chain X_t follows the path traced by chain Y_t : indeed, it is easy to see that, if they start at the same place, the sequence of states visited by the two chains is the same and in the same order. Hence chain X_t hits state 0 before state n if and only if chain Y_t hits state 0 before state n and thus $\mathbf{P}_k(X_{\tau_{0,n}^X} = 0) = \mathbf{P}_k(Y_{\tau_{0,n}^Y} = 0)$.

Finally, observe that $\frac{\hat{q}_k}{\hat{p}_k} = \frac{q_k}{p_k} \leq \alpha$ and $\hat{p}_k + \hat{q}_k = 1$. Hence, for every $k \in \{1, \dots, n-1\}$, we have that $\hat{p}_k \geq \frac{1}{1+\alpha}$ and $\hat{q}_k \leq \frac{\alpha}{1+\alpha}$. This implies, from Lemma B.2, that, from any state $h \in \Omega$,

$$\mathbf{P}_h(Y_{\tau_{0,n}^Y} = n) \geq \frac{1 - \alpha^h}{1 - \alpha^n} \geq 1 - \alpha^h.$$

The lemma follows. \square

C Ehrenfest urns

The *Ehrenfest urn* is the Markov chain with state space $\Omega = \{0, 1, \dots, n\}$ that, when at state k , moves to state $k-1$ or $k+1$ with probability k/n and $(n-k)/n$ respectively (see, for example, Section 2.3 in [29] for a detailed description). The next lemma gives an upper bound on the probability that the Ehrenfest urn starting at state k hits state 0 or n within time step t .

Lemma C.1. *Let $\{Z_t\}$ be the Ehrenfest urn over $\{0, 1, \dots, n\}$ and let $\tau_{0,n}$ be the first time the chain hits state 0 or state n . Then for every $k \geq 1$ it holds that*

$$\mathbf{P}_k(\tau_{0,n} < n \log n + cn) \leq \frac{c'}{n}.$$

for suitable positive constants c and c' .

Proof. First observe that for any $t \geq 3$ the probability of hitting 0 or n before time t for the chain starting at 1 is only $\mathcal{O}(1/n)$ larger than for the chain starting at 2, which in turn is only $\mathcal{O}(1/n)$ larger than for the chain starting at 3. Indeed, by conditioning on the first step of the chain, we have

$$\begin{aligned}\mathbf{P}_1(\tau_{0,n} < t) &= \mathbf{P}_1(\tau_{0,n} < t \mid Z_1 = 0) \mathbf{P}_1(Z_1 = 0) + \mathbf{P}_1(\tau_{0,n} < t \mid Z_1 = 2) \mathbf{P}_1(Z_1 = 2) \\ &= \frac{1}{n} + \frac{n-1}{n} \mathbf{P}_2(\tau_{0,n} < t-1) \leq \frac{1}{n} + \mathbf{P}_2(\tau_{0,n} < t)\end{aligned}$$

and

$$\begin{aligned}\mathbf{P}_2(\tau_{0,n} < t) &= \mathbf{P}_2(\tau_{0,n} < t \mid Z_1 = 1) \mathbf{P}_2(Z_1 = 1) + \mathbf{P}_2(\tau_{0,n} < t \mid Z_1 = 3) \mathbf{P}_2(Z_1 = 3) \\ &= \frac{2}{n} \mathbf{P}_1(\tau_{0,n} < t-1) + \frac{n-2}{n} \mathbf{P}_3(\tau_{0,n} < t-1) \\ &\leq \frac{2}{n} \left(\frac{1}{n} + \mathbf{P}_2(\tau_{0,n} < t) \right) + \frac{n-2}{n} \mathbf{P}_3(\tau_{0,n} < t).\end{aligned}$$

Hence,

$$\begin{aligned}\mathbf{P}_2(\tau_{0,n} < t) &\leq \frac{2}{n-2} + \mathbf{P}_3(\tau_{0,n} < t) \leq \frac{3}{n} + \mathbf{P}_3(\tau_{0,n} < t); \\ \mathbf{P}_1(\tau_{0,n} < t) &\leq \frac{4}{n} + \mathbf{P}_3(\tau_{0,n} < t).\end{aligned}$$

Moreover observe that the probability that the chain starting at k hits state 0 or n before time t is symmetric, i.e. $\mathbf{P}_k(\tau_{0,n} < t) = \mathbf{P}_{n-k}(\tau_{0,n} < t)$, and it is decreasing for k that goes from 0 to $n/2$. In particular, for every k such that $3 \leq k \leq n-3$ it holds that $\mathbf{P}_k(\tau_{0,n} < t) \leq \mathbf{P}_3(\tau_{0,n} < t)$. Now we show that $\mathbf{P}_3(\tau_{0,n} < n \log n + cn) = \mathcal{O}(1/n)$ and this will complete the proof.

First observe that

$$\mathbf{P}_3(\tau_{0,n} < t) = \mathbf{P}_3(\tau_{0,n} < t \wedge \tau_0 < \tau_n) + \mathbf{P}_3(\tau_{0,n} < t \wedge \tau_n < \tau_0),$$

where τ_0 and τ_n are the first time the chain hits state 0 and state n , respectively. Then

$$\mathbf{P}_3(\tau_{0,n} < t) \leq \mathbf{P}_3(\tau_0 < t) + \mathbf{P}_3(\tau_n < t) \leq \mathbf{P}_3(\tau_0 < t) + \mathbf{P}_{n-3}(\tau_n < t) = 2 \cdot \mathbf{P}_3(\tau_0 < t),$$

where the equality follows from the symmetry of the Ehrenfest urn Markov chain.

Let us now consider a path \mathcal{P} of length t starting at state 3 and ending at state 0. Observe that any such path must contain the sub-path going from state 3 to state 0 whose probability is $6/n^3$. Moreover, for all the other $t-3$ moves we have that if the chain crosses an edge $(i, i+1)$ from left to right then it must cross the same edge from right to left (and vice-versa). The probability for any such pair of moves is

$$\frac{n-i}{n} \cdot \frac{i+1}{n} \leq \frac{e^{2/n}}{4},$$

for every i . Hence, for any path \mathcal{P} of length t going from 3 to 0, the probability that the chain follows exactly path \mathcal{P} is³

$$\mathbf{P}_3((X_1, \dots, X_t) = \mathcal{P}) \leq \frac{6}{n^3} \cdot \left(\frac{e^{2/n}}{4} \right)^{(t-3)/2} = \frac{6}{n^3} \cdot \frac{2^3}{e^{3/n}} \cdot \left(\frac{e^{1/n}}{2} \right)^t \leq \frac{48}{n^3} \cdot \left(\frac{e^{1/n}}{2} \right)^t.$$

Let ℓ and r be the number of left and right moves respectively in path \mathcal{P} then $\ell + r = t$ and $\ell - r = 3$. Hence the total number of paths of length t going from 3 to 0 is less than

$$\binom{t}{\ell} = \binom{t}{\frac{t-3}{2}} \leq 2^t.$$

Thus, the probability that starting from 3 the chain hits 0 for the first time exactly at time t is

$$\mathbf{P}_3(\tau_0 = t) \leq \binom{t}{\frac{t-3}{2}} \frac{48}{n^3} \cdot \left(\frac{e^{1/n}}{2} \right)^t \leq \frac{48}{n^3} e^{t/n}.$$

³Notice that such probability is zero if $t-3$ is odd.

Finally, the probability that the hitting time of 0 is less than t is

$$\begin{aligned} \mathbf{P}_3(\tau_0 < t) &\leq \sum_{i=3}^{t-1} \mathbf{P}_3(\tau_0 = i) \\ &\leq \frac{48}{n^3} \sum_{i=3}^{t-1} e^{i/n} = \frac{48}{n^3} \cdot \frac{e^{t/n} - 1}{e^{1/n} - 1} \leq \frac{48e^c}{n}. \end{aligned}$$

In the last inequality we used that $e^{1/n} - 1 \geq 1/n$ and $t = n \log n + cn$. \square

In the proof of Lemma 4.2 we deal with the lazy version of the Ehrenfest urn. The next lemma, which is folklore, allows us to use the bound we achieved in Lemma C.1 for the non-lazy chain.

Lemma C.2. *Let $\{X_t\}$ be an irreducible Markov chain with finite state space Ω and transition matrix P and let $\{\hat{X}_t\}$ be its lazy version, i.e. the Markov chain with the same state space and transition matrix $\hat{P} = \frac{P+I}{2}$ where I is the $\Omega \times \Omega$ identity matrix. Let τ_A and $\hat{\tau}_A$ be the hitting time of the subset of states $A \subseteq \Omega$ in chains $\{X_t\}$ and $\{\hat{X}_t\}$ respectively. Then, for every starting state $b \in \Omega$ and for every time $t \in \mathbb{N}$ it holds that*

$$\mathbf{P}_b(\hat{\tau}_A \leq t) \leq \mathbf{P}_b(\tau_A \leq t).$$

D Proofs from Section 5.2

Throughout this section we denote the players with $0, 1, \dots, n-1$ so that the neighbors of i are $i \pm 1 \bmod n$.

D.1 Hitting \mathbf{p} starting from R when $\Delta > \delta$

In this section we prove Lemma 5.7 that provides an upper bound on the hitting time of state \mathbf{p} when starting from a state $\mathbf{x} \in R$.

Lemma 5.7. *For all $\mathbf{x} \in R$ and $\varepsilon > 0$, if $\Delta > \delta$, $\beta = \omega(\log n)$ and n is sufficiently large*

$$P_{\mathbf{x}}\left(\tau_{\mathbf{p}} > \frac{8-\varepsilon}{\varepsilon} \cdot n^2\right) \leq \frac{\varepsilon}{4}.$$

Remind that states in R have at least two neighbors both playing $+1$ and without loss of generality we assume that $x_0 = x_1 = +1$. Intuitively, for $\beta = \omega(\log n)$, each of player 0 and 1 changes her strategy with very low probability. Moreover, player 2, when selected for update, updates her strategy to $+1$ with high probability. Similarly, after player 2 has played $+1$, we have that each of player 0, 1 and 2 changes her strategy with very low probability and player 3, when selected for update, plays $+1$ with high probability. This process repeats until every player is playing $+1$. In the following, we estimate the number of steps sufficient to have all players playing strategy $+1$ with high probability.

For sake of compactness, we will denote the strategy of player i at time step t by X_t^i . We start with a simple observation that lower bounds the probability that a player picks strategy $+1$ when selected for update, given that at least one of their neighbors is playing $+1$.

Observation D.1. *If player i is selected for update at time t then, for $b \in \{-1, +1\}$*

$$\mathbf{P}(X_t^i = +1 \mid X_{t-1}^{i+b} = +1) \geq \left(1 - \frac{1}{1 + e^{(\Delta-\delta)\beta}}\right).$$

We start by lower bounding the probability that players $2, \dots, n-1$ are selected for update at least once in this order before a given number t of steps. Set $\rho_1 = 0$ and, for $i = 2, \dots, n-1$, let ρ_i be the first time player i is selected for update after time step ρ_{i-1} . Thus, at time ρ_i player i is selected for update and players $2, \dots, i-1$ have already been selected at least once in this order. In particular, ρ_{n-1} is the first time step at which every player i , $i \geq 2$, has been selected at least once after his left neighbor. Obviously, $\rho_i > \rho_{i-1}$ for $i = 2, \dots, n-1$. The next lemma lower bounds the probability that $\rho_{n-1} \leq t$.

Lemma D.2. For every $\mathbf{x} \in R$ and every $t > 0$, we have

$$\mathbf{P}_{\mathbf{x}}(\rho_{n-1} \leq t) \geq 1 - \frac{n^2}{t}.$$

Proof. Every player i has probability $\frac{1}{n}$ of being selected at any given time step. Therefore, $\mathbf{E}[\rho_2] = \mathbf{E}[\rho_2 - \rho_1] = n$ and $\mathbf{E}[\rho_i - \rho_{i-1}] = n$, for $i = 3, \dots, n-1$. Thus, by linearity of expectation,

$$\mathbf{E}[\rho_{n-1}] = \sum_{i=2}^{n-1} \mathbf{E}[\rho_i - \rho_{i-1}] \leq n^2.$$

The lemma follows from the Markov inequality. \square

The next lemma gives a lower bound to the probability that $X_t^i = +1$ for all players i and for $t \geq \rho_{n-1}$.

Lemma D.3. For every $\mathbf{x} \in R$, for every player i and for every $t > 0$, we have

$$\mathbf{P}_{\mathbf{x}}(X_t^i = +1 \mid \rho_{n-1} \leq t) \geq \left(1 - \frac{1}{1 + e^{(\Delta-\delta)\beta}}\right)^t.$$

We prove the lemma first for $i \geq 2$ and then we will deal with players 0 and 1. Fix player $i \geq 2$, time step t and set $s_{i+1} = t$. Starting from time step t and going backward to time step 0, we identify the sequence of time steps $s_i > s_{i-1} > \dots > s_2 > 0$ such that, for $j = i, i-1, \dots, 2$, s_j is the last time player j has been selected before time s_{j+1} . We remark that, since $t \geq \rho_{n-1} > \rho_i$ we have that players $2, \dots, i$ are selected at least once in this order and thus all the s_j are well defined. Strictly speaking, the sequence s_i, \dots, s_2 depends on i and t and thus a more precise, and more cumbersome, notation would have been $s_{i,j,t}$. Since player i and time step t will be clear from the context, we drop i and t .

In order to lower bound the probability that $X_t^i = +1$ for $i \geq 2$, we first bound it in terms of the probability that player 2 plays +1 at time s_2 and then we evaluate this last quantity. The next lemma is the first step of our proof.

Lemma D.4. For every $\mathbf{x} \in R$, every player $2 \leq i \leq n-1$ and every $t \geq \rho_{n-1}$, we have

$$\mathbf{P}_{\mathbf{x}}(X_t^i = +1 \mid \rho_{n-1} \leq t) \geq \left(1 - \frac{1}{1 + e^{(\Delta-\delta)\beta}}\right)^{i-2} \mathbf{P}_{\mathbf{x}}(X_{s_2}^2 = +1 \mid \rho_{n-1} \leq t).$$

Proof. For every i , s_i is the last time the player i is selected for update before t and thus $X_t^i = X_{s_i}^i$. Hence, for $i = 2$ the lemma obviously holds. For $i > 2$ and $j = 3, \dots, i$, we observe that, since $t \geq \rho_i$, s_{j-1} is the last time that player $j-1$ has been selected for update before time s_j and thus $X_{s_j}^{j-1} = X_{s_{j-1}}^{j-1}$. Then, from Observation D.1, we have

$$\begin{aligned} \mathbf{P}_{\mathbf{x}}(X_{s_j}^j = +1 \mid \rho_{n-1} \leq t) &\geq \mathbf{P}_{\mathbf{x}}(X_{s_j}^j = +1 \mid X_{s_{j-1}}^{j-1} = +1, \rho_{n-1} \leq t) \cdot \mathbf{P}_{\mathbf{x}}(X_{s_{j-1}}^{j-1} = +1 \mid \rho_{n-1} \leq t) \\ &\geq \left(1 - \frac{1}{1 + e^{(\Delta-\delta)\beta}}\right) \mathbf{P}_{\mathbf{x}}(X_{s_{j-1}}^{j-1} = +1 \mid \rho_{n-1} \leq t). \end{aligned}$$

We obtain the lemma by iteratively applying the same argument to $\mathbf{P}_{\mathbf{x}}(X_{s_{j-1}}^{j-1} = +1 \mid \rho_{n-1} \leq t)$, for $j = i-1, i-2, \dots, 3$. \square

We now bound the probability that player 2 plays +1 at time step s_2 . If player 1 has not been selected for update before time s_2 , then $X_{s_2-1}^1 = X_0^1 = +1$, and, from Observation D.1, we have

$$\begin{aligned} \mathbf{P}_{\mathbf{x}}(X_{s_2}^2 = +1 \mid \rho_{n-1} \leq t) &\geq \mathbf{P}_{\mathbf{x}}(X_{s_2}^2 = +1 \mid X_{s_2-1}^1 = +1, \rho_{n-1} \leq t) \\ &\geq \left(1 - \frac{1}{1 + e^{(\Delta-\delta)\beta}}\right). \end{aligned}$$

It remains to consider the case when player 1 has been selected for update at least once before time s_2 . For any fixed time step t , we define a new sequence of time steps $r_0 > r_1, \dots > 0$ in the following way.

We set $r_0 = s_2$ and let r_j , for $j > 0$, be the last time player $j \bmod 2$ has been selected before time r_{j-1} . For the last element in the sequence, r_k , it holds that player $(k+1) \bmod 2$ is not selected before time step r_k .

Since at time s_2 player 2 has been selected for update and since r_1 is the last time step player 1 has been selected for update before $r_0 = s_2$, we have $X_{s_2-1}^1 = X_{r_1}^1$ and, by Observation D.1,

$$\begin{aligned} \mathbf{P}_{\mathbf{x}}(X_{s_2}^2 = +1 \mid \rho_{n-1} \leq t) &\geq \mathbf{P}_{\mathbf{x}}(X_{s_2}^2 = +1 \mid X_{s_2-1}^1 = +1, \rho_{n-1} \leq t) \cdot \mathbf{P}_{\mathbf{x}}(X_{s_2-1}^1 = +1 \mid \rho_{n-1} \leq t) \\ &\geq \left(1 - \frac{1}{1 + e^{(\Delta-\delta)\beta}}\right) \cdot \mathbf{P}_{\mathbf{x}}(X_{r_1}^1 = +1 \mid \rho_{n-1} \leq t). \end{aligned} \quad (14)$$

Finally, we bound $\mathbf{P}_{\mathbf{x}}(X_{r_1}^1 = +1 \mid \rho_{n-1} \leq t)$.

Lemma D.5. *For every $\mathbf{x} \in R$, time step t , and player i , let r_0, \dots, r_k be defined as above. If $k > 0$, we have*

$$\mathbf{P}_{\mathbf{x}}(X_{r_1}^1 = +1 \mid \rho_{n-1} \leq t) \geq \left(1 - \frac{1}{1 + e^{(\Delta-\delta)\beta}}\right)^k.$$

Proof. For sake of compactness, in this proof we denote the parity of integer a with $p(a) = a \bmod 2$. Thus, the definition of sequence r_j gives that player $p(j)$ has been selected for update at time r_j and

$$\begin{aligned} \mathbf{P}_{\mathbf{x}}(X_{r_j}^{p(j)} = +1 \mid \rho_{n-1} \leq t) \\ \geq \mathbf{P}_{\mathbf{x}}(X_{r_j}^{p(j)} = +1 \mid X_{r_j-1}^{p(j+1)} = +1, \rho_{n-1} \leq t) \mathbf{P}_{\mathbf{x}}(X_{r_j-1}^{p(j+1)} = +1 \mid \rho_{n-1} \leq t). \end{aligned}$$

If $j \neq k$ player $p(j+1)$ has not been selected for update between time $r_{j+1} \leq r_j - 1$ and time r_j and, by Observation D.1, we obtain

$$\begin{aligned} \mathbf{P}_{\mathbf{x}}(X_{r_j}^{p(j)} = +1 \mid \rho_{n-1} \leq t) \\ \geq \mathbf{P}_{\mathbf{x}}(X_{r_j}^{p(j)} = +1 \mid X_{r_j-1}^{p(j+1)} = +1, \rho_{n-1} \leq t) \mathbf{P}_{\mathbf{x}}(X_{r_j-1}^{p(j+1)} = +1 \mid \rho_{n-1} \leq t) \\ \geq \left(1 - \frac{1}{1 + e^{(\Delta-\delta)\beta}}\right) \mathbf{P}_{\mathbf{x}}(X_{r_{j+1}}^{p(j+1)} = +1 \mid \rho_{n-1} \leq t). \end{aligned}$$

If $j = k$, instead, player $p(k+1)$ has not been selected for update before time r_k and thus $X_{r_k-1}^{p(k+1)} = X_0^{p(k+1)} = +1$. By Observation D.1, we have

$$\begin{aligned} \mathbf{P}_{\mathbf{x}}(X_{r_k}^{p(k)} = +1 \mid \rho_{n-1} \leq t) \\ \geq \mathbf{P}_{\mathbf{x}}(X_{r_k}^{p(k)} = +1 \mid X_{r_k-1}^{p(k+1)} = +1, \rho_{n-1} \leq t) \mathbf{P}_{\mathbf{x}}(X_{r_k-1}^{p(k+1)} = +1 \mid \rho_{n-1} \leq t) \\ \geq \left(1 - \frac{1}{1 + e^{(\Delta-\delta)\beta}}\right). \end{aligned} \quad \square$$

Proof (of Lemma D.3). For every player $i \geq 2$ and $t > 0$, we have

$$\begin{aligned} \mathbf{P}_{\mathbf{x}}(X_t^i = +1 \mid \rho_{n-1} \leq t) &\geq \left(1 - \frac{1}{1 + e^{(\Delta-\delta)\beta}}\right)^{i-2} \mathbf{P}_{\mathbf{x}}(X_{s_2}^2 = +1 \mid \rho_{n-1} \leq t) \quad (\text{from Lemma D.4}) \\ &\geq \left(1 - \frac{1}{1 + e^{(\Delta-\delta)\beta}}\right)^{i-1} \mathbf{P}_{\mathbf{x}}(X_{r_1}^1 = +1 \mid \rho_{n-1} \leq t) \quad (\text{from Equation 14}) \\ &\geq \left(1 - \frac{1}{1 + e^{(\Delta-\delta)\beta}}\right)^{i-1+k} \quad (\text{from Lemma D.5}) \\ &\geq \left(1 - \frac{1}{1 + e^{(\Delta-\delta)\beta}}\right)^t, \end{aligned}$$

where k is the index of the last term in the sequence r_0, r_1, \dots previously defined and where the last inequality follows from $i - 1 + k \leq t$, since the sequence of updates we are considering cannot be longer than t .

The lemma for players 0 and 1 can be proved in a similar way. Clearly, if player $i = 0, 1$ has never been selected for update before time t , we have that $X_t^i = +1$ with probability 1. If player i has been selected at least once we have to distinguish the cases $i = 0$ and $i = 1$. If $i = 1$, we define $r_0 = t + 1$ and we identify a sequence of time step $r_1 > r_2 > \dots > 0$ as above: we have that $X_t^1 = X_{r_1}^1$ and from Lemma D.5 follows that

$$\mathbf{P}_{\mathbf{x}}(X_t^i = +1 \mid \rho_{n-1} \leq t) \geq \left(1 - \frac{1}{1 + e^{(\Delta-\delta)\beta}}\right)^k \geq \left(1 - \frac{1}{1 + e^{(\Delta-\delta)\beta}}\right)^t,$$

where k is the last index of the sequence r_1, r_2, \dots . Finally, the probability that player 0 plays the strategy $+1$ at time t , given that she was selected for update at least once, can be handled similarly to the probability that player 2 plays the strategy $+1$ at time s_2 . \square

The following lemma gives the probability that the hitting time of the profile \mathbf{p} is less or equal to t , given that $\rho_{n-1} \leq t$.

Lemma D.6. *For every $\mathbf{x} \in R$ and every $t > 0$, we have*

$$\mathbf{P}_{\mathbf{x}}(\tau_{\mathbf{p}} \leq t \mid \rho_{n-1} \leq t) \geq \left(1 - \frac{1}{1 + e^{(\Delta-\delta)\beta}}\right)^{nt}.$$

Proof. Let f be the permutation that sorts the players according to the order of last selection for update; i.e., $f(0)$ is the last player that is selected for update, $f(1)$ is the next to last one, and so on. We have

$$\begin{aligned} \mathbf{P}_{\mathbf{x}}(\tau_{\mathbf{p}} \leq t \mid \rho_{n-1} \leq t) &\geq \mathbf{P}_{\mathbf{x}}\left(\bigwedge_{j=0}^{n-1} X_t^{f(j)} = +1 \mid \rho_{n-1} \leq t\right) \\ &= \prod_{j=0}^{n-1} \mathbf{P}_{\mathbf{x}}\left(X_t^{f(j)} = +1 \mid \bigwedge_{i=j+1}^{n-1} X_t^{f(i)} = +1, \rho_{n-1} \leq t\right) \\ &\geq \prod_{j=0}^{n-1} \mathbf{P}_{\mathbf{x}}\left(X_t^{f(j)} = +1 \mid \rho_{n-1} \leq t\right) \\ &\geq \left(1 - \frac{1}{1 + e^{(\Delta-\delta)\beta}}\right)^{nt}, \end{aligned}$$

where for the second inequality we used that the probability of selecting $+1$ can only increase when there are other players playing this strategy, whereas the last inequality follows from Lemma D.3. \square

Now we are ready to prove Lemma 5.7.

Proof (of Lemma 5.7). From Lemma D.2 and Lemma D.6, we have that for every $\mathbf{x} \in R$ and every $t > 0$

$$\begin{aligned} \mathbf{P}_{\mathbf{x}}(\tau_{\mathbf{p}} \leq t) &\geq \mathbf{P}_{\mathbf{x}}(\rho_{n-1} \leq t) \cdot \mathbf{P}_{\mathbf{x}}(\tau_{\mathbf{p}} \leq t \mid \rho_{n-1} \leq t) \\ &\geq \left(1 - \frac{n^2}{t}\right) \left(1 - \frac{1}{1 + e^{(\Delta-\delta)\beta}}\right)^{nt}. \end{aligned}$$

For sufficiently large n we have that, since $\beta = \omega(\log n)$,

$$e^{(\Delta-\delta)\beta} \geq \frac{8-\varepsilon}{\varepsilon} \cdot \frac{n^3}{\log \frac{8}{8-\varepsilon}}.$$

By setting $t = \frac{8-\varepsilon}{\varepsilon} \cdot n^2$, we have

$$\begin{aligned} \mathbf{P}_{\mathbf{x}}(\tau_{\mathbf{p}} \leq t) &\geq \left(1 - \frac{\varepsilon}{8-\varepsilon}\right) \left(1 - \frac{1}{1 + \frac{8-\varepsilon}{\varepsilon} \cdot \frac{n^3}{\log \frac{8}{8-\varepsilon}}}\right)^{\frac{8-\varepsilon}{\varepsilon} \cdot n^3} \\ &\geq \frac{8(1-\frac{\varepsilon}{4})}{8-\varepsilon} \frac{8-\varepsilon}{8} = 1 - \frac{\varepsilon}{4}, \end{aligned}$$

where the second inequality follows from the well known approximation $1 - a \geq e^{-\frac{a}{1-a}}$. \square

D.2 Hitting \mathbf{p} starting from S_d when $\Delta > \delta$

In this section we prove Lemma 5.8 that provides an upper bound on the hitting time of state \mathbf{p} when starting from a state $\mathbf{x} \in S_d$.

Lemma 5.8. *For all $d > 0$, $\mathbf{x} \in S_d$ and $\varepsilon > 0$, if $\Delta > \delta$, $\beta = \omega(\log n)$ and n is sufficiently large*

$$P_{\mathbf{x}}\left(\tau_{\mathbf{p}} > \frac{8}{\varepsilon} \cdot n^2\right) \leq \frac{1}{2^d + 1} + \frac{\varepsilon}{4}.$$

In fact, we show that from any $\mathbf{x} \in S_d$ the dynamics hit after a polynomial number of steps a profile in R with high probability. Then, the aimed result easily follows by Lemma 5.7.

Specifically, let us denote by $\text{negl}(n)$ a function in n that is smaller than the inverse of every polynomial in n . Then the following upper bound on the hitting time of R from the set S_d holds.

Lemma D.7. *For every $d > 0$ and $\mathbf{x} \in S_d$, if $\Delta > \delta$, $\beta = \omega(\log n)$ and n is sufficiently large*

$$\mathbf{P}_{\mathbf{x}}(\tau_R \leq n^2) \geq \frac{2^d}{2^d + 1} (1 - \text{negl}(n)).$$

Before proving this lemma, we show how Lemma 5.8 easily follows from Lemma D.7 and Lemma 5.7.

Proof (of Lemma 5.8). By Lemma 5.7 and Lemma D.7, we have

$$\begin{aligned} \mathbf{P}_{\mathbf{x}}\left(\tau_{\mathbf{p}} \leq \frac{8n^2}{\varepsilon}\right) &\geq \mathbf{P}_{\mathbf{x}}\left(\tau_{\mathbf{p}} \leq \frac{8n^2}{\varepsilon} \mid \tau_R \leq n^2\right) \mathbf{P}_{\mathbf{x}}(\tau_R \leq n^2) \\ &\geq \mathbf{P}_{X_{\tau_R}}\left(\tau_{\mathbf{p}} \leq \frac{(8-\varepsilon)n^2}{\varepsilon}\right) \mathbf{P}_{\mathbf{x}}(\tau_R \leq n^2) \\ &\geq \left(1 - \frac{\varepsilon}{4}\right) \frac{2^d}{2^d + 1} (1 - \text{negl}(n)) \geq \frac{2^d}{2^d + 1} - \frac{\varepsilon}{4}. \end{aligned} \quad \square$$

We now prove Lemma D.7. Specifically we will bound $\mathbf{P}_{\mathbf{x}}(\tau_R \leq t \wedge \tau_R \leq \tau_{\mathbf{m}})$ and will use the fact that

$$\mathbf{P}_{\mathbf{x}}(\tau_R \leq t) \geq \mathbf{P}_{\mathbf{x}}(\tau_R \leq t \wedge \tau_R \leq \tau_{\mathbf{m}}). \quad (15)$$

Let θ^* be the first time at which all players have been selected at least once. We define players playing +1 in profile \mathbf{x} as the *plus-players* of \mathbf{x} and their neighbors as *border-players*. Consider the event E that “a border-player is selected for update before at least one of her neighboring plus-players” and denote by \overline{E} its complement. Observe that

$$\begin{aligned} \mathbf{P}_{\mathbf{x}}(\tau_R \leq t \wedge \tau_R \leq \tau_{\mathbf{m}}) &\geq \mathbf{P}_{\mathbf{x}}(\tau_R \leq t \wedge \tau_R \leq \tau_{\mathbf{m}} \mid E \wedge \theta^* \leq t) \mathbf{P}_{\mathbf{x}}(E \wedge \theta^* \leq t) \\ &= \mathbf{P}_{\mathbf{x}}(\tau_R \leq t \wedge \tau_R \leq \tau_{\mathbf{m}} \mid E \wedge \theta^* \leq t) \mathbf{P}_{\mathbf{x}}(E \mid \theta^* \leq t) \mathbf{P}_{\mathbf{x}}(\theta^* \leq t). \end{aligned} \quad (16)$$

We next bound the three components appearing in the last line of (16). A bound on the probability that $\theta^* \leq t$ directly follows from the coupon collector argument; we include a proof of this bound for completeness.

Lemma D.8. *For every $t > 0$,*

$$\mathbf{P}_{\mathbf{x}}(\theta^* \leq t) \geq 1 - ne^{-t/n}.$$

Proof. The logit dynamics at each time step select a player for update uniformly and independently of the previous selections. Thus, the probability that i players are never selected for update in t steps is $(1 - \frac{i}{n})^t$ and

$$\mathbf{P}_{\mathbf{x}}(\theta^* > t) \leq \sum_{i=1}^{n-1} \left(1 - \frac{i}{n}\right)^t \leq \sum_{i=1}^{n-1} e^{-\frac{it}{n}} \leq ne^{-t/n}. \quad \square$$

Next we bound the probability of the event E given that all players have been selected at least once.

Lemma D.9. *For every $\mathbf{x} \in S_d$, $d > 0$ and every $t > 0$*

$$\mathbf{P}_{\mathbf{x}}(E \mid \theta^* \leq t) \geq \frac{2^d}{2^d + 1}.$$

Proof. Let us define $\ell(\mathbf{x})$ as the number of border-players in \mathbf{x} . Note that, since there are no adjacent plus-players in \mathbf{x} , $\ell(\mathbf{x}) \geq d$.

The proof proceeds by induction on d . Let $d = 1$ and denote by i the plus-player. Since we are conditioning on $\theta^* \leq t$, all players are selected at least once by time t and thus the probability that one of the two neighbors of i is selected for update before i is selected is $\frac{2}{3} = \frac{2^d}{2^d + 1}$.

Suppose now that the claim holds for $d - 1$ and consider $\mathbf{x} \in S_d$. Denote by $T_{\mathbf{x}}$ the set of all the plus-players in \mathbf{x} and their border-players and let i be the first player in $T_{\mathbf{x}}$ to be selected for update (notice that i is well defined since $\theta^* \leq t$). Observe that, if i is a border-player, then the event E occur and this happens with probability $\frac{\ell(\mathbf{x})}{\ell(\mathbf{x}) + d}$. If i is a plus-player, we consider the subset $\bar{T}_{\mathbf{x}} \subset T_{\mathbf{x}}$ of the remaining $d - 1$ plus-players and their border-players. The event E will occur if and only if at least one border-player in $\bar{T}_{\mathbf{x}}$ is selected before one of its neighboring plus-players. Notice though that $\bar{T}_{\mathbf{x}} = T_{\mathbf{y}}$, for $\mathbf{y} \in S_{d-1} \setminus R$ such that $y_i = -1$ and $\mathbf{y}_{-i} = \mathbf{x}_{-i}$. Thus, by inductive hypothesis, $\mathbf{P}_{\mathbf{y}}(E \mid \theta^* \leq t) \geq \frac{2^{d-1}}{2^{d-1} + 1} = \frac{2^d}{2^d + 2}$. Thus,

$$\begin{aligned} \mathbf{P}_{\mathbf{x}}(E \mid \theta^* \leq t) &= \frac{\ell(\mathbf{x})}{\ell(\mathbf{x}) + d} + \frac{d}{\ell(\mathbf{x}) + d} \cdot \mathbf{P}_{\mathbf{y}}(E \mid \theta^* \leq t) \\ &\geq \frac{\ell(\mathbf{x})}{\ell(\mathbf{x}) + d} + \frac{d}{\ell(\mathbf{x}) + d} \cdot \frac{2^d}{2^d + 2} \\ &= 1 - \frac{2d}{(\ell(\mathbf{x}) + d)(2^d + 2)} \\ &\geq 1 - \frac{1}{2^d + 2}, \end{aligned}$$

where the last inequality follows from $\ell(\mathbf{x}) \geq d$. The claim follows since $1 - \frac{1}{2^d + 2} \geq \frac{2^d}{2^d + 1}$. \square

Finally suppose that the event E occurs, i.e., there is a border-player i selected for update before than at least one of the neighboring plus-players, and $\theta^* \leq t$. Let τ be the time step at which E occurs. Note that, since all players have been selected at least once before t , then $\tau < t$. Moreover, at time τ , the player i has at least one neighbor playing strategy $+1$, and thus she will play this strategy with probability at least $\left(1 - \frac{1}{1 + e^{(\Delta - \delta)\beta}}\right)$. Then we have

$$\mathbf{P}_{\mathbf{x}}(\tau_R \leq t \wedge \tau_R \leq \tau_{\mathbf{m}} \mid E \wedge \theta^* \leq t) \geq \mathbf{P}_{\mathbf{x}}(X_{\tau}^i = +1 \mid E \wedge \theta^* \leq t) \geq \left(1 - \frac{1}{1 + e^{(\Delta - \delta)\beta}}\right), \quad (17)$$

where X_{τ}^i denotes the strategy of i in the profile X_{τ} .

Now we are ready to prove Lemma D.7.

Proof (of Lemma D.7). From (15), (16), (17), Lemma D.8 and Lemma D.9 we have

$$\mathbf{P}_{\mathbf{x}}(\tau_R \leq n^2) \geq \left(1 - \frac{1}{1 + e^{(\Delta - \delta)\beta}}\right) \cdot \frac{2d - 1}{2d} \cdot (1 - ne^{-n}) = \frac{2d - 1}{2d} \cdot (1 - \text{negl}(n)),$$

where the last equality holds for n sufficiently large since $\beta = \omega(\log n)$. \square

D.3 Hitting $\{\mathbf{p}, \mathbf{m}\}$ starting from S_d when $\Delta > \delta$

In this section we prove Lemma 5.9 that provides an upper bound on the hitting time of the set $\{\mathbf{p}, \mathbf{m}\}$ when starting from a state $\mathbf{x} \in S_d$.

Lemma 5.9. *For all $\mathbf{x} \in \{\pm 1\}^n$ and $\varepsilon > 0$, if $\Delta > \delta$, $\beta = \omega(\log n)$ and n is sufficiently large*

$$P_{\mathbf{x}} \left(\tau_{\{\mathbf{p}, \mathbf{m}\}} > \frac{8}{\varepsilon} \cdot n^2 \right) \leq \frac{\varepsilon}{2}.$$

The proof is similar to the proof of Lemma 5.8, except that now we need a bound on the probability to hit R or \mathbf{m} in polynomial time when starting from a state in S_d . Next lemma gives us such a bound.

Lemma D.10. *For all $\mathbf{x} \in \{\pm 1\}^n$, if $\Delta > \delta$, $\beta = \omega(\log n)$ and n is sufficiently large*

$$\mathbf{P}_{\mathbf{x}} (\tau_{R \cup \{\mathbf{m}\}} \leq n^2) \geq 1 - \text{negl}(n).$$

Lemma 5.9 then immediately follows.

Proof (of Lemma 5.9). By Lemma 5.7 and Lemma D.10, we have

$$\begin{aligned} \mathbf{P}_{\mathbf{x}} \left(\tau_{\{\mathbf{p}, \mathbf{m}\}} \leq \frac{8n^2}{\varepsilon} \right) &\geq \mathbf{P}_{\mathbf{x}} \left(\tau_{\{\mathbf{p}, \mathbf{m}\}} \leq \frac{8n^2}{\varepsilon} \mid \tau_{R \cup \{\mathbf{m}\}} \leq n^2 \right) \mathbf{P}_{\mathbf{x}} (\tau_{R \cup \{\mathbf{m}\}} \leq n^2) \\ &\geq \mathbf{P}_{X_{\tau_{R \cup \{\mathbf{m}\}}}} \left(\tau_{\{\mathbf{p}, \mathbf{m}\}} \leq \frac{(8-\varepsilon)n^2}{\varepsilon} \right) \mathbf{P}_{\mathbf{x}} (\tau_{R \cup \{\mathbf{m}\}} \leq n^2) \\ &\geq \left(1 - \frac{\varepsilon}{4} \right) (1 - \text{negl}(n)) \geq 1 - \frac{\varepsilon}{2}. \quad \square \end{aligned}$$

Let us now prove Lemma D.10. First of all, we note that

$$\mathbf{P}_{\mathbf{x}} (\tau_{R \cup \{\mathbf{m}\}} \leq t) = \mathbf{P}_{\mathbf{x}} (\tau_R \leq t \wedge \tau_R \leq \tau_{\mathbf{m}}) + \mathbf{P}_{\mathbf{x}} (\tau_{\mathbf{m}} \leq t \wedge \tau_{\mathbf{m}} \leq \tau_R). \quad (18)$$

A lower bound on the first term on the right hand side of the above equation was given in previous section. In fact, from (16), (17) and Lemma D.8, we have

$$\mathbf{P}_{\mathbf{x}} (\tau_R \leq t \wedge \tau_R \leq \tau_{\mathbf{m}}) \geq \left(1 - \frac{1}{1 + e^{(\Delta-\delta)\beta}} \right) \cdot (1 - ne^{-t/n}) \cdot \mathbf{P}_{\mathbf{x}} (E \mid \theta^* \leq t). \quad (19)$$

Similarly,

$$\begin{aligned} \mathbf{P}_{\mathbf{x}} (\tau_{\mathbf{m}} \leq t \wedge \tau_{\mathbf{m}} \leq \tau_R) &\geq \mathbf{P}_{\mathbf{x}} (\tau_{\mathbf{m}} \leq t \wedge \tau_{\mathbf{m}} \leq \tau_R \mid \overline{E} \wedge \theta^* \leq t) \mathbf{P}_{\mathbf{x}} (\overline{E} \wedge \theta^* \leq t) \\ &= \mathbf{P}_{\mathbf{x}} (\tau_{\mathbf{m}} \leq t \wedge \tau_{\mathbf{m}} \leq \tau_R \mid \overline{E} \wedge \theta^* \leq t) \mathbf{P}_{\mathbf{x}} (\overline{E} \mid \theta^* \leq t) \mathbf{P}_{\mathbf{x}} (\theta^* \leq t) \\ &\geq (1 - ne^{-t/n}) \cdot \mathbf{P}_{\mathbf{x}} (\tau_{\mathbf{m}} \leq t \wedge \tau_{\mathbf{m}} \leq \tau_R \mid \overline{E} \wedge \theta^* \leq t) \mathbf{P}_{\mathbf{x}} (\overline{E} \mid \theta^* \leq t) \end{aligned} \quad (20)$$

where the last inequality follows from Lemma D.8.

Next lemma bounds $\mathbf{P}_{\mathbf{x}} (\tau_{\mathbf{m}} \leq t \wedge \tau_{\mathbf{m}} \leq \tau_R \mid \overline{E} \wedge \theta^* \leq t)$.

Lemma D.11. *For every $d \geq 0$, for every $\mathbf{x} \in S_d$ and every $t > 0$, we have*

$$\mathbf{P}_{\mathbf{x}} (\tau_{\mathbf{m}} \leq t \wedge \tau_{\mathbf{m}} \leq \tau_R \mid \overline{E} \wedge \theta^* \leq t) \geq 1 - \frac{t}{e^{2\delta\beta}}.$$

Proof. Suppose that the event E does not occur, i.e., all plus-players are selected before their neighboring border-players, and $\theta^* \leq t$. Let τ be the time step at which the last plus-player is selected. Note that, since all players have been selected at least once before t , then $\tau < t$. Thus, if in the first τ time steps the selected player adopts strategy -1 , then $\tau_{\mathbf{m}} \leq \tau \leq t$ and $\tau_{\mathbf{m}} \leq \tau_R$, i.e., by denoting the player selected for update at time step j as $i(j)$,

$$\begin{aligned} \mathbf{P}_{\mathbf{x}} (\tau_{\mathbf{m}} \leq t \wedge \tau_{\mathbf{m}} \leq \tau_R \mid \overline{E} \wedge \theta^* \leq t) &\geq \mathbf{P}_{\mathbf{x}} \left(\bigwedge_{j=1}^{\tau} X_j^{i(j)} = -1 \mid \overline{E} \wedge \theta^* \leq t \right) \\ &= \prod_{j=1}^{\tau} \mathbf{P}_{\mathbf{x}} \left(X_j^{i(j)} = -1 \mid \overline{E} \wedge \theta^* \leq t \wedge \bigwedge_{k=1}^j X_k^{i(k)} = -1 \right). \end{aligned}$$

Given that no border-player is selected before the corresponding plus-player and that at each previous time step the selected player has adopted strategy -1 , we have that, for every j , player $i(j)$ has both neighbors playing -1 . Hence, $i(j)$ adopts strategy -1 with probability $\left(1 - \frac{1}{1+e^{2\delta\beta}}\right)$. Thus,

$$\mathbf{P}_{\mathbf{x}}(\tau_{\mathbf{m}} \leq t \wedge \tau_{\mathbf{m}} \leq \tau_R \mid \overline{E} \wedge \theta^* \leq t) \geq \left(1 - \frac{1}{1+e^{2\delta\beta}}\right)^\tau \geq \left(1 - \frac{1}{1+e^{2\delta\beta}}\right)^t \geq \left(1 - \frac{t}{e^{2\delta\beta}}\right),$$

where the last inequality follows from the approximations $1-a \leq e^{-a}$ and $1-a \geq e^{-\frac{a}{1-a}}$ for $0 \leq a \leq 1$. \square

Now we are ready to prove Lemma D.10.

Proof (of Lemma D.10). From (19), (20) and Lemma D.11, for n sufficiently large and $\beta = \omega(\log n)$ we have both

$$\mathbf{P}_{\mathbf{x}}(\tau_R \leq n^2 \wedge \tau_R \leq \tau_{\mathbf{m}}) \geq (1 - \text{negl}(n)) \cdot \mathbf{P}_{\mathbf{x}}(E \mid \theta^* \leq n^2).$$

and

$$\mathbf{P}_{\mathbf{x}}(\tau_{\mathbf{m}} \leq n^2 \wedge \tau_{\mathbf{m}} \leq \tau_R) \geq (1 - \text{negl}(n)) \cdot \mathbf{P}_{\mathbf{x}}(\overline{E} \mid \theta^* \leq n^2).$$

Then, from (18) it follows that

$$\mathbf{P}_{\mathbf{x}}(\tau_{R \cup \{\mathbf{m}\}} \leq n^2) \geq (1 - \text{negl}(n)) (\mathbf{P}_{\mathbf{x}}(E \mid \theta^* \leq n^2) + \mathbf{P}_{\mathbf{x}}(\overline{E} \mid \theta^* \leq n^2)) = 1 - \text{negl}(n). \quad \square$$

D.4 Hitting $\{\mathbf{p}, \mathbf{m}\}$ when $\Delta = \delta$

In this section we prove Lemma 5.10 that provides an upper bound on the hitting time of the set $\{\mathbf{p}, \mathbf{m}\}$ when $\Delta = \delta$.

Lemma 5.10. *For all $\mathbf{x} \in \{\pm 1\}^n$ and $\varepsilon > 0$, if $\Delta = \delta$, $\beta = \omega(\log n)$ and n is sufficiently large*

$$\mathbf{P}_{\mathbf{x}}(\tau_{\{\mathbf{p}, \mathbf{m}\}} > n^5) \leq o(1).$$

Let us start by introducing some useful notation. We say that the profile \mathbf{x} has a *plus-block* of size l starting at player i if $x_i = x_{i+1} = \dots = x_{i+l-1} = +1$ and $x_{i-1} = x_{i+l} = -1$ and players i and $i+l-1$ are the *border* players of the block. A similar definition is given for *minus-blocks*. Notice that every profile $\mathbf{x} \neq \mathbf{p}, \mathbf{m}$ has the same number of plus-blocks and minus-blocks and this number is called the *level* of \mathbf{x} and is denoted by $\ell(\mathbf{x})$. We set $\ell(\mathbf{p}) = \ell(\mathbf{m}) = 0$. Moreover, for a profile \mathbf{x} , we define $s_+(\mathbf{x})$ as the number of plus-blocks of size 1, $s_-(\mathbf{x})$ as the number of minus-blocks of size 1 and set $s(\mathbf{x}) = s_+(\mathbf{x}) + s_-(\mathbf{x})$.

We would like to study how long it takes to the logit dynamics to hit \mathbf{p} or \mathbf{m} . Let \mathbf{x} be the starting profile of the logit dynamics and assume that $\ell(\mathbf{x}) = \ell$. Since \mathbf{p} and \mathbf{m} have level 0, the dynamics have to go down ℓ levels before hitting one of these two target profiles. To reach a profile in a smaller level the dynamics have to reach a profile having a monochromatic block of size 1 is reached, select the unique player of this block for update and have this player changing her strategy to the same strategy as her neighbors. Let us denote with τ_i the hitting time of a profile at level i , i.e., $\tau_i = \tau_{\{\mathbf{x}: \ell(\mathbf{x})=i\}}$. By linearity of expectation, we have that

$$\mathbf{E}_{\mathbf{x}}[\tau_{\{\mathbf{p}, \mathbf{m}\}}] \leq \sum_{i=0}^{\ell-1} \max_{\mathbf{x}: \ell(\mathbf{x})=i+1} \mathbf{E}_{\mathbf{x}}[\tau_i]. \quad (21)$$

Thus, to bound the hitting time of $\{\mathbf{p}, \mathbf{m}\}$ we have only to compute $\mathbf{E}_{\mathbf{x}}[\tau_i]$ for any level $i = 0, 1, \dots, \ell-1$ and for any profile \mathbf{x} such that $\ell(\mathbf{x}) = i+1$.

Fix \mathbf{x} be a profile of level $i+1$ and number arbitrarily its $2(i+1)$ monochromatic blocks. We denote by $k_j(\mathbf{x})$ the size of the j -th monochromatic block. For each $1 \leq j \leq 2(i+1)$, let us define the quantity θ_j as follows. Suppose that the dynamics starting from \mathbf{x} hit for the first time a profile of level i after t steps and this happens because the j -th monochromatic block disappears; then, we set $\theta_j = t$ and $\theta_{j'} = +\infty$ for all $j' \neq j$. Given that the starting profile \mathbf{x} is at level $i+1$, we have $\tau_i = \min_j \theta_j$ and, thus,

$$\mathbf{E}_{\mathbf{x}}[\tau_i] = \mathbf{E}_{\mathbf{x}}\left[\min_j \theta_j\right] \leq \max_j \mathbf{E}_{\mathbf{x}}[\theta_j \mid \theta_j < \theta_{j'} \text{ for all } j' \neq j].$$

In order to have a more compact notation we define

$$\gamma_{i,l} = \max_{1 \leq j \leq 2(i+1)} \max_{\mathbf{x}: \ell(\mathbf{x})=i+1} \mathbf{E}_{\mathbf{x}} [\theta_j \mid \theta_j < \theta_{j'} \text{ for all } j' \neq j],$$

and set $\gamma_i = \max_l \gamma_{i,l}$. Observe that

$$\mathbf{E}_{\mathbf{x}} [\tau_i] \leq \gamma_i \quad (22)$$

and it is not hard to see that $\gamma_{i,l}$ is non-decreasing with l . Next lemma gives a bound γ_i in terms of γ_{i+1} . Using recursively this lemma we will be able to upper bound γ_i for any $i \geq 0$.

Lemma D.12. For $0 \leq i < \lfloor n/2 \rfloor - 1$

$$\gamma_i \leq n^3 \left(1 + \frac{1}{1 + e^{2\Delta\beta}} \gamma_{i+1} \right).$$

Moreover, $\gamma_{\lfloor n/2 \rfloor - 1} \leq n^3$.

Proof. We start by bounding $\gamma_{i,l}$ for any l . Let \mathbf{x} and j be the profile and the monochromatic block that attain the maximum $\gamma_{i,l}$, respectively. We remark that we are conditioning on the event that $\theta_{i,j} < \theta_{i,j'}$ for all $j' \neq j$ (i.e., the j -th block is the first one to disappear). To bound $\gamma_{i,l}$ we distinguish three cases, depending on the value of l .

$l = 1$: Let u be the unique player of the j -th monochromatic block and let v be the player selected for update. Consider all the possible selections of v :

- If $u = v$ and v changes her strategy the block disappears and thus $\theta_j = 1$. This event occurs with probability $\frac{1}{n} \cdot \left(1 - \frac{1}{1 + e^{2\Delta\beta}} \right)$.
- Suppose that v is a neighbor of u and she changes her strategy. Observe that v cannot belong to a monochromatic block of size 1 since we are assuming that j is the first block to disappear. Thus, the two neighbors of v are playing different strategies and when v has to change her strategy she selects at random. We can conclude that the probability that v is selected and she changes her strategy is at most $1/2 \cdot 2/n = 1/n$. Moreover, after this update the dynamics reach a profile of level $i + 1$ where the size of the j -th block increases to 2.
- Suppose that v is not a border player of any monochromatic block and she changes her strategy. The players that are not at the border of a block are $n - 4(i + 1) + s(\mathbf{x})$ and each of them changes her strategy with probability $\frac{1}{1 + e^{2\Delta\beta}}$. We can conclude that the probability of such an event is $\frac{n - 4(i + 1) + s(\mathbf{x})}{n(1 + e^{2\Delta\beta})}$ and in this case the dynamics reach a profile of level $i + 2$.
- for all the remaining choices for v both the level of the reached profile and the length of the j -th monochromatic block remain the same as in the starting profile.

Summing over all the possible choices for v and observing that $\gamma_{i,2} \geq \gamma_{i,1}$ we have

$$\begin{aligned} \gamma_{i,1} &\leq \frac{1}{n} \left(1 - \frac{1}{1 + e^{2\Delta\beta}} \right) + \frac{1}{n} (1 + \gamma_{i,2}) + \frac{n - 4(i + 1) + s(\mathbf{x})}{n} \frac{1}{1 + e^{2\Delta\beta}} (1 + \gamma_{i+1}) \\ &\quad + \left(\frac{n - 2}{n} - \frac{n - 4(i + 1) + s(\mathbf{x}) - 1}{n} \frac{1}{1 + e^{2\Delta\beta}} \right) (1 + \gamma_{i,1}). \end{aligned}$$

By simple calculations and using the fact that $n - 4(i + 1) + s(\mathbf{x}) \geq 0$, we obtain

$$\gamma_{i,1} \leq \left(\frac{1}{2} + \frac{1}{4e^{2\Delta\beta} + 2} \right) \left(n + \gamma_{i,2} + \frac{n - 4(i + 1) + s(\mathbf{x})}{1 + e^{2\Delta\beta}} \gamma_{i+1} \right).$$

Finally, since $\left(\frac{1}{2} + \frac{1}{4e^{2\Delta\beta} + 2} \right) \leq \frac{2}{3}$ for $\beta = \omega(\log n)$, we can conclude that

$$\gamma_{i,1} \leq \frac{2}{3} (\gamma_{i,2} + b_i), \quad (23)$$

where $b_i = n + \frac{\gamma_{i+1}}{1 + e^{2\Delta\beta}} \cdot \max_{\mathbf{y}: \ell(\mathbf{y})=i+1} (n - 4(i + 1) + s(\mathbf{y}))$.

$1 < l < n - 2i - 1$: Let v be the player selected for update. Consider all the possible selections of v :

- suppose v is one of the two players at the border of the j -th monochromatic block and she changes her strategy. Since the border players are two and their neighbors are playing different strategies, v selects her new strategy at random and thus the probability of this event is $1/n$. Moreover, the dynamics reach a profile \mathbf{y} of level $i+1$ where the length of the j -th monochromatic block decreases to $l-1$.
- suppose v does not belong to the j -th monochromatic block but she is a neighbor of its border players and she changes her strategy. Then the dynamics reach a profile \mathbf{y} of level $i+1$. Since v cannot belong to a monochromatic block of size 1 (otherwise her block would disappear before block j) we can state that the two neighbors of v are playing different strategies. Thus, when v updates her strategy she chooses each of the two alternatives with probability $1/2$. Since there are two players adjacent to the border players of block j , this case happens with probability at most $1/n$.
- suppose v is a player that is not a border player of any monochromatic block and she changes her strategy. Then, when v changes her strategy a monochromatic block is split and the reached profile has level $i+2$. Notice that there are $n-4(i+1)+s(\mathbf{x})$ such players and each of them has probability $\frac{1}{1+e^{2\Delta\beta}}$ to change her strategy.
- for all the remaining choices for v both the level of the reached profile and the length of the j -th monochromatic block remain the same as in the starting profile.

Hence,

$$\begin{aligned} \gamma_{i,l} \leq & \frac{1}{n}(1 + \gamma_{i,l-1}) + \frac{1}{n}(1 + \gamma_{i,l+1}) + \frac{n-4(i+1)+s(\mathbf{x})}{n} \frac{1}{1+e^{2\Delta\beta}}(1 + \gamma_{i+1}) \\ & + \left(\frac{n-2}{n} - \frac{n-4(i+1)+s(\mathbf{x})}{n} \frac{1}{1+e^{2\Delta\beta}} \right) (1 + \gamma_{i,l}). \end{aligned}$$

By simple calculations, similar to the ones for the case $l=1$, we obtain

$$\gamma_{i,l} \leq \frac{1}{2}(\gamma_{i,l-1} + \gamma_{i,l+1} + b_i).$$

From the previous inequality and Equation 23, a simple induction on l shows that, for every $1 \leq l < n-2i-1$, we have

$$\gamma_{i,l} \leq \frac{1}{l+2} \left((l+1)\gamma_{i,l+1} + \frac{l(l+3)}{2}b_i \right). \quad (24)$$

Moreover, from Equation 24, we can use a simple inductive argument to show that, for every $h \geq 1$,

$$\begin{aligned} \gamma_{i,l} & \leq \frac{l+1}{l+h+1} \gamma_{i,l+h} + \frac{l+1}{2} b_i \sum_{j=l}^{l+h-1} \frac{j(j+3)}{(j+1)(j+2)} \\ & \leq \frac{l+1}{l+h+1} \gamma_{i,l+h} + \frac{l+1}{2} h b_i. \end{aligned} \quad (25)$$

$l = n-2i-1$: in this case all blocks other than the j -th have size 1. Thus, every time one of these players is selected for update she doesn't change her strategy otherwise there would be a monochromatic block disappearing before block j . This means that the size of the j -th monochromatic block cannot increase. By using an argument similar to the one used in the previous cases, we obtain that

$$\gamma_{i,n-2i-1} \leq \gamma_{i,n-2i-2} + b_i.$$

By using Equation 24, we have

$$\gamma_{i,n-2i-1} \leq \frac{(n-2i-2)(n-2i+1) + 2(n-2i)}{2} b_i \leq \frac{n^2}{2} b_i.$$

Finally, for every $l \geq 1$, by using Equation 25 with $h = n-2i-1-l$, we have

$$\gamma_{i,l} \leq \frac{l+1}{n-2i} \gamma_{i,n-2i-1} + \frac{(l+1)(n-2i-1-l)}{2} b_i \leq n^2 b_i.$$

The lemma finally follows by observing that $\max_{\mathbf{y}: \ell(\mathbf{y})=i+1} (n - 4(i+1) + s(\mathbf{y})) \leq n$ for $0 \leq i < \lfloor n/2 \rfloor - 1$ and it is exactly 0 for $i = \lfloor n/2 \rfloor - 1$. \square

Corollary D.13. *If $\beta = \omega(\log n)$, then for every $i \geq 0$, $\gamma_i = \mathcal{O}(n^3)$.*

Proof. From Lemma D.12 we have $\gamma_{\lfloor n/2 \rfloor - 1} \leq n^3$. Instead, for $0 \leq i < \lfloor n/2 \rfloor - 1$, we have

$$\gamma_i \leq n^3 \left(1 + \frac{1}{1 + e^{2\Delta\beta}} \gamma_{i+1} \right) \leq n^3 \left(1 + \sum_{j=1}^{\lfloor n/2 \rfloor - i - 1} \left(\frac{n^3}{1 + e^{2\Delta\beta}} \right)^j \right). \quad (26)$$

The corollary follows by observing that, if $\beta = \omega\left(\frac{\log n}{\Delta}\right)$, then the summation in Equation 26 is $o(1)$. \square

The above corollary gives a polynomial bound to the time that the dynamics take to go from a profile at level $i+1$ to a profile at level i . Lemma 5.10 easily follows.

Proof (of Lemma 5.10). From (21) and (22), for every profile \mathbf{x} at level $1 \leq \ell \leq n/2$ we have

$$\mathbf{E}_{\mathbf{x}} [\tau_{\mathbf{p}, \mathbf{m}}] \leq \sum_{i=0}^{\ell-1} \gamma_i = \mathcal{O}(n^4),$$

where the last bound follows from Corollary D.13. The lemma then follows from the Markov inequality. \square

D.5 Bounding $\alpha_{\mathbf{x}}$ when $\Delta = \delta$

Here we prove the following lemma.

Lemma D.14. *For every $d \geq 0$, every profile \mathbf{x} with exactly d players playing +1 and $\beta = \omega(\log n)$*

$$\alpha_{\mathbf{x}} = \frac{d}{n} \pm o(1).$$

Proof. Trivially, $\alpha_{\mathbf{p}} = 1$ and $\alpha_{\mathbf{m}} = 0$. We next show that for $\beta = \omega(\log n)$ and \mathbf{x} with exactly d players playing +1, $\alpha_{\mathbf{x}} = \frac{d}{n} + \lambda_{\mathbf{x}}$, for some $\lambda_{\mathbf{x}} = o(1)$. By the definition of Markov chains we know that

$$\alpha_{\mathbf{x}} = P(\mathbf{x}, \mathbf{x}) \cdot \alpha_{\mathbf{x}} + \sum_{\mathbf{y} \in N(\mathbf{x})} P(\mathbf{x}, \mathbf{y}) \cdot p_{\mathbf{y}}.$$

We then partition the neighborhood $N(\mathbf{x})$ of profile \mathbf{x} of level i in 5 subsets, $N_1(\mathbf{x})$, $N_2(\mathbf{x})$, $N_3(\mathbf{x})$, $N_4(\mathbf{x})$, $N_5(\mathbf{x})$ such that, for two profiles $\mathbf{y}_1, \mathbf{y}_2$ in the same subsets it holds that $P(\mathbf{x}, \mathbf{y}_1) = P(\mathbf{x}, \mathbf{y}_2)$. Then

- $N_1(\mathbf{x})$ is the set of profiles \mathbf{y} obtained from \mathbf{x} by changing the strategy of a player of a plus-block of size 1. Observe that $|N_1(\mathbf{x})| = s_+(\mathbf{x})$. Moreover, for every $\mathbf{y} \in N_1(\mathbf{x})$, \mathbf{y} is at level $i-1$, has $d-1$ players playing +1 and $P(\mathbf{x}, \mathbf{y}) = \frac{1}{n} \cdot \left(1 - \frac{1}{1+e^{2\Delta\beta}}\right)$.
- $N_2(\mathbf{x})$ is the set of profiles \mathbf{y} obtained from \mathbf{x} by changing the strategy of a player of a minus-block of size 1. Observe that $|N_2(\mathbf{x})| = s_-(\mathbf{x})$. Moreover, for every $\mathbf{y} \in N_2(\mathbf{x})$, \mathbf{y} is at level $i-1$, has $d+1$ players playing +1 and $P(\mathbf{x}, \mathbf{y}) = \frac{1}{n} \cdot \left(1 - \frac{1}{1+e^{2\Delta\beta}}\right)$.
- $N_3(\mathbf{x})$ is the set of profiles \mathbf{y} obtained from \mathbf{x} by changing the strategy of a border player of a plus-block of size greater than 1. Observe that $|N_3(\mathbf{x})| = 2(i - s_+(\mathbf{x}))$. Moreover, for every $\mathbf{y} \in N_3(\mathbf{x})$, \mathbf{y} is at level i , and has $d-1$ players playing +1 and $P(\mathbf{x}, \mathbf{y}) = 1/2n$.
- $N_4(\mathbf{x})$ is the set of profiles \mathbf{y} obtained from \mathbf{x} by changing the strategy of a border player of a minus-block of size greater than 1. Observe that $|N_4(\mathbf{x})| = 2(i - s_-(\mathbf{x}))$. Moreover, for every $\mathbf{y} \in N_4(\mathbf{x})$, \mathbf{y} is at level i , and has $d+1$ players playing +1 and $P(\mathbf{x}, \mathbf{y}) = 1/2n$.
- $N_5(\mathbf{x})$ is the set of all the profiles $\mathbf{y} \in N(\mathbf{x})$ that do not belong to any of the previous 4 subsets. Observe that $|N_5(\mathbf{x})| = n - 4i + s(\mathbf{x})$. Moreover, for every $\mathbf{y} \in N_5(\mathbf{x})$, \mathbf{y} is at level $i+1$, and $P(\mathbf{x}, \mathbf{y}) = \frac{1}{n} \cdot \frac{1}{1+e^{2\Delta\beta}}$.

Moreover, we have that

$$P(\mathbf{x}, \mathbf{x}) = \frac{s(\mathbf{x})}{n} \frac{1}{1 + e^{2\Delta\beta}} + \frac{2i - s(\mathbf{x})}{n} + \frac{n - 4i + s(\mathbf{x})}{n} \left(1 - \frac{1}{1 + e^{2\Delta\beta}}\right).$$

Then, we have

$$\begin{aligned} \alpha_{\mathbf{x}} &= \frac{1}{n} \left(1 - \frac{1}{1 + e^{2\Delta\beta}}\right) \left(\sum_{\mathbf{y} \in N_1(\mathbf{x})} \alpha_{\mathbf{y}} + \sum_{\mathbf{y} \in N_2(\mathbf{x})} \alpha_{\mathbf{y}} \right) + \frac{1}{2n} \left(\sum_{\mathbf{y} \in N_3(\mathbf{x})} \alpha_{\mathbf{y}} + \sum_{\mathbf{y} \in N_4(\mathbf{x})} \alpha_{\mathbf{y}} \right) \\ &\quad + \frac{1}{n} \frac{1}{1 + e^{2\Delta\beta}} \sum_{\mathbf{y} \in N_5(\mathbf{x})} \alpha_{\mathbf{y}} \\ &\quad + \left(\frac{s(\mathbf{x})}{n} \frac{1}{1 + e^{2\Delta\beta}} + \frac{2i - s(\mathbf{x})}{n} + \frac{n - 4i + s(\mathbf{x})}{n} \left(1 - \frac{1}{1 + e^{2\Delta\beta}}\right) \right) \alpha_{\mathbf{x}} \\ &= \frac{1}{n} \left(\sum_{\mathbf{y} \in N_1(\mathbf{x})} \alpha_{\mathbf{y}} + \sum_{\mathbf{y} \in N_2(\mathbf{x})} \alpha_{\mathbf{y}} \right) + \frac{1}{2n} \left(\sum_{\mathbf{y} \in N_3(\mathbf{x})} \alpha_{\mathbf{y}} + \sum_{\mathbf{y} \in N_4(\mathbf{x})} \alpha_{\mathbf{y}} \right) + \frac{n - 2i}{n} \cdot \alpha_{\mathbf{x}} + \frac{c}{1 + e^{2\Delta\beta}}, \end{aligned} \tag{27}$$

where

$$c = \frac{1}{n} \left(\sum_{\mathbf{y} \in N_5(\mathbf{x})} \alpha_{\mathbf{y}} - \sum_{\mathbf{y} \in N_1(\mathbf{x}) \cup N_2(\mathbf{x})} \alpha_{\mathbf{y}} - (n - 4i) \alpha_{\mathbf{x}} \right).$$

We notice that, since $1 \leq i \leq n/2$ and $|N_1(\mathbf{x})| + |N_2(\mathbf{x})|, |N_5(\mathbf{x})| \leq n$, we have $|c| \leq 2$ and thus the last term in Equation 27 is negligible in n (since $\beta = \omega(\log n)$). Hence we have that the following condition holds for every level $i \geq 1$ and every profile \mathbf{x} at level i :

$$\alpha_{\mathbf{x}} = \frac{1}{2i} \left(\sum_{\mathbf{y} \in N_1(\mathbf{x})} \alpha_{\mathbf{y}} + \sum_{\mathbf{y} \in N_2(\mathbf{x})} \alpha_{\mathbf{y}} \right) + \frac{1}{4i} \left(\sum_{\mathbf{y} \in N_3(\mathbf{x})} \alpha_{\mathbf{y}} + \sum_{\mathbf{y} \in N_4(\mathbf{x})} \alpha_{\mathbf{y}} \right) + \eta_{\mathbf{x}},$$

where $\eta_{\mathbf{x}}$ is negligible in n . This gives us a linear system of equations in which the number of equations is the same that the number of variables.

Let us consider the *polished* version of this system, in which we omit the negligible part in every equation. We will compute a solution of the polished system. Finally, we argue the solution of original system should be close to the polished one.

Let us denote with $\alpha_{\mathbf{x}}^*$ the variable depending on \mathbf{x} in the polished system. Note that for each level i and every profile \mathbf{x} at level i , $\alpha_{\mathbf{x}}^*$ does not depend on profiles at higher level. This enable us to compute a solution for the polished system inductively on level i . Indeed, for every profile \mathbf{x} at level 0 (this is only possible for $d = 0$ or $d = n$), we have, as discussed above, $\alpha_{\mathbf{x}}^* = \frac{d}{n}$. Now, consider \mathbf{x} at level i with exactly d players playing +1. By assuming that, for every profile \mathbf{y} at level $i - 1$, $\alpha_{\mathbf{y}}^* = \frac{d_{\mathbf{y}}}{n}$, where $d_{\mathbf{y}}$ is the number of player playing +1 in \mathbf{y} , we have

$$\alpha_{\mathbf{x}}^* = \frac{s_+(\mathbf{x})}{2i} \cdot \frac{d - 1}{n} + \frac{s_-(\mathbf{x})}{2i} \cdot \frac{d + 1}{n} + \frac{1}{4i} \left(\sum_{\mathbf{y} \in N_3(\mathbf{x})} \alpha_{\mathbf{y}}^* + \sum_{\mathbf{y} \in N_4(\mathbf{x})} \alpha_{\mathbf{y}}^* \right). \tag{28}$$

Let us now consider Equation 28 for each \mathbf{x} at level i . These gives another linear system of equations. This system has a unique solution: indeed, the number of equations and the number of variables coincides and, moreover, the matrix of coefficients is a diagonally dominant matrix (since $|N_3(\mathbf{x})| + |N_4(\mathbf{x})| \leq 4i$) and thus it is non-singular. We will show this solution must set $\alpha_{\mathbf{x}}^* = \frac{d}{n}$ for every profile \mathbf{x} at level i with exactly d players playing +1. Indeed, with this assignment the right hand side of the Equation 28 becomes

$$\frac{s_+(\mathbf{x})}{2i} \frac{d - 1}{n} + \frac{s_-(\mathbf{x})}{2i} \frac{d + 1}{n} + \frac{i - s_+(\mathbf{x})}{2i} \frac{d - 1}{n} + \frac{i - s_-(\mathbf{x})}{2i} \frac{d + 1}{n} = \frac{d}{n}.$$

Thus, we can conclude $\alpha_{\mathbf{x}}^* = \frac{d}{n}$ for each profile \mathbf{x} with exactly d players playing +1.

Now, let $\lambda_{\mathbf{x}} = |\alpha_{\mathbf{x}} - \alpha_{\mathbf{x}}^*|$. As n grows unbounded, any original equation approaches the polished one, so we need $\alpha_{\mathbf{x}}$ to approach $\alpha_{\mathbf{x}}^*$. Then $\lambda_{\mathbf{x}} = o(1)$ for every profile \mathbf{x} . \square