Metastability of Logit Dynamics for Coordination Games

Vincenzo Auletta, Diodato Ferraioli, Francesco Pasquale, and Giuseppe Persiano

Dipartimento di Informatica "Renato M. Capocelli"
Università di Salerno
auletta,ferraioli,pasquale,giuper@dia.unisa.it

Abstract

Logit Dynamics [Blume, Games and Economic Behavior, 1993] is a 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. In several cases it happens that on a time-scale shorter than mixing time the chain is "quasi-stationary", meaning that it stays close to some small set of the state space, while in a time-scale multiple of the mixing time it jumps from one quasi-stationary configuration to another; this phenomenon is usually called "metastability".

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. In particular, we study no-risk-dominant coordination games on the clique (which is equivalent to the well-known Glauber dynamics for the Ising model) and coordination games on a ring (both the risk-dominant and no-risk-dominant case). We also describe a simple "artificial" game that highlights the distinctive features of our metastability notion based on distributions.

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 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 future states 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 a specific dynamics, the *logit* dynamics [6], and consider as solution concept its equilibrium states. Logit dynamics is a type of *noisy* best response dynamics that models 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 "almost surely" play their best response; that is, they pick the strategies yielding high payoff with higher probability. It is well known that this dynamics induces, for every strategic game, a Markov chain which has 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 behaviour is that the system may take too long to reach it, 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 [3, 2]. 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: when the mixing time is exponential, the transient phase is completely chaotic, or we can still spot some regularities? Are we able to say something about the behavior of the chain before it reaches the stationary distribution? Obviously, such a chain is perfectly described by the collections of probability distributions consisting of one distribution for each time step and each starting profile. 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 (that is, the stationary distribution).

Our results show games for which regularities can be observed even in the transient phase. In particular, we will show that, depending on the starting profile, the dynamics rapidly reaches a distribution and remains close to this distribution for a sufficiently long time (we call such a distribution *metastable*). 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 logit dynamics is not rapidly mixing but the profile (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 results. In this work, we obtain results about the metastability of the logit dynamics for different classes of coordination games.

- We start in Section 3 by introducing the notion of an (ε, T)-metastable distribution μ and of its pseudo-mixing time. Roughly speaking, 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, $t_\mu^x(\varepsilon)$, 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" by that of "pseudo-mixing time" and that of "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". Roughly speaking, the state space Ω can be partitioned into a small number of subsets $\Omega_1, \Omega_2, \ldots$ of "equivalent" states; that is, 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.

- In Section 4, we analyze the metastable distributions of the Ising model on the complete graph, also known as the Curie-Weiss model, and used by physicists to model the interaction between magnets in a ferro-magnetic system [26]. The model describes also the process of opinion formation and the spread of new technologies in fully connected societies (see, for example [29]). It is known [18, 27, 2] 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 is equivalent to the logit dynamics for this game where β is exactly the inverse of the temperature. The mixing time of this 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 any $t = poly(n)$ when $\beta = \Omega(\log n/n)$. Moreover we show that the pseudo-mixing time of these distributions is polynomial when the dynamics starts from a profile where the difference in the number of positive and negative magnets is large.

- Graphical coordination games are often used to model the spread of a new technology in a social network [30, 27] 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. In Section 5 we follow [15] and study the case in which social interaction between the players is described by the ring topology. We show that for every starting profile there is a metastable distribution and the dynamics approaches it in a polynomial number of steps.

Finally, we consider the OR-game, an artificial game defined in [3], that highlights the distinctive features of our metastability notion based on distributions.

Other equilibrium notions. Let us compare the solution concept studied in this paper (the stationary distribution of the logit dynamics) 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 [19]). Even tough 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 Nash equilibria [12, 11] suggest that Best Response, 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 a dynamics and for each dynamics we have a unique equilibrium. For rapidly mixing chains the equilibrium is quickly reached. The results in this paper show that, even for the non-rapidly mixing case, the system can be described if we discount the 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 is a clear and crispy way to model these settings.

Related works. Logit dynamics has been first studied by Blume [6] who showed that, for 2×2 coordination games, the long-term behavior of the system is concentrated in the risk dominant equilibrium (see [19]). Ellison [15] studied logit dynamics for graphical coordination games on rings (the class of games we study in Section 5) and showed that some large fraction of the players will eventually choose the strategy with maximum potential. Similar results were obtained by Peyton Young [30] for more general families of graphs. Montanari and Saberi [27] gave bounds on the hitting time of the highest potential equilibrium for the logit dynamics in terms of some graph theoretic properties of the underlying interaction network. Asadpour and Saberi [1] studied the hitting time of the Nash equilibrium for a class of congestion games.

The study of the mixing time of logit dynamics for strategic games has been initiated in [3] (see also [2]), 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. The Ising model is a very well-studied topic and we refer the reader to the survey of Martinelli [26] and to Chapter 15 of [25].

In Physics, Chemistry or Biology, metastability is a phenomenon related to the evolution of systems under noisy dynamics. In particular, metastability concerns the transition between different regions of the state space and the existence of multiple, well-separated time scales: at short time scales the system appears to be in a *quasi-equilibrium*, and it explores only a confined region of the available space state, while, at larger time scales, it undergoes transitions between such different regions. Examples of metastability can be found in Biology, Climatology, Economics, Materials Science and Physics.

Metastability appears for the first time around 1935 with the work of Eyring [16] and Kramers [22] on diffusion in potential wells, but the mathematically rigorous analysis of metastability phenomena in the context of randomly perturbed dynamical systems start in the early 1970's with the work of Freidlin and Wentzell [17]. Since then, metastability is a very well studied topic in Physics and several monographs on this subject are available (see, for example [20, 28, 7, 21]). The goal of metastability is to model processes showing the following typical behavior: starting from a given profile, the system will rather quickly visit the nearby maximum of the potential function (a metastable state); the dynamics stays very close to such a state for a very long time, avoiding visits to other local maxima; at some point, the system leaves the metastable state (and its neighborhood) and moves to some other local maximum, usually better than the previous one; the process then is repeated. Research in Physics about metastability aims at expressing typical features of a metastable state and to evaluate the transition time between metastable states; the main approaches used to this analysis are based on large deviation theory [17] or on potential theory [8]. Our approach is closest to the one of Bovier et al. [9]. They define the notion of a *metastable point* as a state that is quickly reached and difficult to leave. For every metastable point x they define the *local valley* of x as the set of states for which x is the metastable point with the smallest hitting time and the associated metastable distribution associated with x is the stationary distribution restricted to the local valley. In [10], Bovier and Manzo apply the approach of [9] in the context of zero temperature limit of Glauber dynamics of spin systems in finite volume and show that the transition times can be expressed in terms of properties of the potential function.

Metastability was analysed not only for discrete dynamics, but also for continuous Markov processes. In [23] Larralde et al. define a metastable state by two components: spectral feature of a state (namely, isolated eigenstate of the master operator of the Markov Process having an exceptionally low eigenvalue) and the technical condition meaning that the probability of being in a metastable state at equilibrium is vanishingly small. These conditions partition the state space in two disjoint set: the metastable states and the equilibrium states. They show that for any starting profile $\mathbf{x}$, the dynamics quickly reach, with a probability $p_{\mathbf{x}}$, a state which is fully in the metastable region and, with probability $1 - p_{\mathbf{x}}$, the equilibrium. Further, if the dynamics start from the metastable region, then the probability of leaving it in short time is very low. Moreover, they consider a *restricted dynamics* in which the process is reflected each time it attempts to leave the metastable region, whose equilibrium is described by the restriction of the stationary distribution to the metastable region: they show that these restricted dynamics well mimic the process when the starting point is in the metastable region.

Very recently, Beltran and Landim [4] describe for the continuous time Markov process of the Ising model all metastable behaviors, defining time scales at which they occur, the metastable set associated to each time scale, and the asymptotic dynamics which specifies at which rate the process jumps from one metastable state to another.

The work on censored Glauber dynamics [24, 13, 14] is also related to ours: the mixing time in a censored dynamics resemble the pseudo mixing time for the metastable distribution on a subset of states. However, we stress that the censored dynamics alters the original evolution of the Markov chain and the techniques developed do not seem useful to answer questions about the pseudo-mixing time.

Notations. We write $|S|$ for the size of a set S. We use bold symbols for vectors; when $\mathbf{x} = (x_1, \ldots, x_n) \in \{0,1\}^n$ we write $|\mathbf{x}|_a$ for the number of occurrence of a in $\mathbf{x}$; i.e., $|\mathbf{x}|_a = |\{i \in [n] \,:\, x_i = a\}|$. We use the standard game theoretic notation $(\mathbf{x}_{-i}, y)$ to mean the vector obtained from $\mathbf{x}$ by replacing the i-th entry with y; i.e., $(\mathbf{x}_{-i}, y) = (x_1, \ldots, x_{i-1}, y, x_{i+1}, \ldots, x_n)$. We denote by $\mathsf{negl}\,(n)$ a function in n that is smaller than the inverse of every polynomial in n. For distributions μ and ν on the same space Ω, the *total variation distance* between μ and ν is defined as $\|\mu - \nu\|_{\mathrm{TV}} := \frac{1}{2}\sum_{x \in \Omega} |\mu(x) - \nu(x)|$.

2 Logit dynamics

A *strategic game* $\mathcal{G}$ is a triple $\mathcal{G} = ([n], \mathcal{S}, \mathcal{U})$, where $[n] = \{1, \ldots, n\}$ is a finite set of *players*, $\mathcal{S} = \{S_1, \ldots, S_n\}$ is a family of non-empty finite sets (S_i is the set of *strategies* for player i), and $\mathcal{U} = \{u_1, \ldots, u_n\}$ is a family of *utility functions* (or *payoffs*), where $u_i : S_1 \times \cdots \times S_n \to \mathbb{R}$ is the utility function of player i.

Logit dynamics is a *randomized* best-response dynamics where at every time step a players $i \in [n]$ is selected uniformly at random, and she updates her strategy according to the following probability distribution over the set S_i. For every $y \in S_i$

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

where $\mathbf{x} \in S_1 \times \cdots \times S_n$ is the strategy profile played by the players at the current time step, $T_i(\mathbf{x}) = \sum_{z \in S_i} e^{\beta u_i(\mathbf{x}_{-i}, z)}$ is the normalizing factor, and $\beta \geqslant 0$.

Parameter β is a *rationality level* for the system. From (1) it is easy to see that, for $\beta = 0$ (no rationality) player i selects her strategy uniformly at random, for $\beta > 0$ the probability is biased toward strategies promising higher payoffs, and for $\beta \to \infty$ (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 $\beta \geqslant 0$, the above dynamics defines a Markov chain with the set of strategy profiles as state space, where the transition probability from profile $\mathbf{x} = (x_1, \ldots, x_n)$ to profile $\mathbf{y} = (y_1, \ldots, y_n)$ is zero if the $\mathbf{x}$ and $\mathbf{y}$ differ at more than one player, and it is $\frac{1}{n}\sigma_i(y_i \,|\, \mathbf{x})$ if the two profiles differ exactly at player i. More formally, we can define the logit dynamics as follows.

Definition 2.1 (Logit dynamics [6]) *Let $\mathcal{G} = ([n], \mathcal{S}, \mathcal{U})$ be a strategic game and let $\beta \geqslant 0$. The* logit dynamics *for $\mathcal{G}$ is the Markov chain $\mathcal{M}_\beta = \{X_t \,:\, t \in \mathbb{N}\}$ with state space $\Omega = S_1 \times \cdots \times S_n$ and transition matrix*

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

where $\sigma_i(y_i \,|\, \mathbf{x})$ 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 π exists and for every starting state $\mathbf{x}$ the distribution of the chain $P^t(\mathbf{x}, \cdot)$ at time t approaches π as t tends to infinity.

Potential games. A game $\mathcal{G} = ([n], \mathcal{S}, \mathcal{U})$ is said a (exact) *potential* game if a function $\Phi : S_1 \times \cdots \times S_n \to \mathbb{R}$ exists such that, for every player i and for every pair of profiles $\mathbf{x}$ and $\mathbf{y}$ that differ only at position i, it holds that $u_i(\mathbf{x}) - u_i(\mathbf{y}) = \Phi(\mathbf{x}) - \Phi(\mathbf{y})$. It is easy to see that, if $\mathcal{G} = ([n], \mathcal{S}, \mathcal{U})$ is a potential game with potential function Φ, then the Markov chain given by (2) is *reversible* (i.e. $\pi(\mathbf{x})P(\mathbf{x}, \mathbf{y}) = \pi(\mathbf{y})P(\mathbf{y}, \mathbf{x})$ for every $\mathbf{x}$ and $\mathbf{y}$) and its stationary distribution is the Gibbs measure

$$\pi(\mathbf{x}) = \frac{1}{Z} e^{\beta \Phi(\mathbf{x})} \tag{3}$$

where $Z = \sum_{\mathbf{y} \in S_1 \times \cdots \times S_n} e^{\beta \Phi(\mathbf{y})}$ is the normalizing constant.

When $\mathcal{G}$ is a potential game, the logit dynamics is equivalent to the well-studied *Glauber dynamics*, meaning that logit dynamics for $\mathcal{G}$ and Glauber dynamics for the Gibbs distribution π in (3) define the same Markov chain. Due to the analogies between logit and Glauber dynamics, we will sometimes adopt the terminology used by physicists to indicate the quantities involved; in particular we will call parameter β the inverse noise or *inverse temperature* and we will call *partition function* the normalizing constant Z of the Gibbs distribution (3). All games we consider in this paper are potential games.

3 Metastability

In this section we give formal definitions of *metastable distributions* and *pseudo-mixing time*. As a simple example we analyze metastability for the logit dynamics for 2-player coordination games and we also highlight some connections between metastability and the bottleneck ratio.

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 *if for every $0 \leqslant t \leqslant T$ it holds that*

$$\left\|\mu P^t - \mu\right\|_{\mathrm{TV}} \leqslant \varepsilon$$

Here are two obvious property of metastable distributions.

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

Proof. By using the triangle inequality, we have

$$\left\|\mu P^T - \mu\right\|_{\mathrm{TV}} \leqslant \left\|\mu P^t - \mu P\right\|_{\mathrm{TV}} + \left\|\mu P - \mu\right\|_{\mathrm{TV}} \leqslant \left\|\mu P^{t-1} - \mu\right\|_{\mathrm{TV}} + \varepsilon\,,$$

where the last inequality follows from the $(\varepsilon,1)$-metastability of μ and from the fact that if μ and ν are two probability distributions and P is a stochastic matrix then $\|\mu P - \nu P\|_{\mathrm{TV}} \leqslant \|\mu - \nu\|_{\mathrm{TV}}$. □

The definition of metastable distribution captures the idea of a distribution that behaves approximately like the stationary distribution, meaning that if we start from such distribution and run the chain we stay close to it for a long 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 3.3 (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 the* pseudo-mixing time $t_\mu^S(\varepsilon)$ *as*

$$t_\mu^S(\varepsilon) = \inf\{t \in \mathbb{N} \,:\, \left\|P^t(x,\cdot) - \mu\right\|_{\mathrm{TV}} \leqslant \varepsilon \text{ for all } x \in S\}$$

We observe that the stationary distribution π of an ergodic Markov chain is reached within ε in time $t_{\mathrm{mix}}(\varepsilon)$ from every state (see Appendix A). Thus, according to Definition 3.3, we have that $t_\pi^\Omega(\varepsilon) = t_{\mathrm{mix}}(\varepsilon)$.

The following simple lemma connects metastability and pseudo-mixing.

Lemma 3.4 *Let μ be a (ε, T)-metastable distribution and let $S \subseteq \Omega$ be a set of states such that $t_\mu^S(\varepsilon) < +\infty$. Then for every $x \in S$ it holds that*

$$\left\|P^t(x,\cdot) - \mu\right\|_{\mathrm{TV}} \leqslant 2\varepsilon \quad \text{for every } t_\mu^S(\varepsilon) \leqslant t \leqslant t_\mu^S(\varepsilon) + T$$

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

$$\begin{aligned}
\left\|P^t(x,\cdot) - \mu\right\|_{\mathrm{TV}} &= \left\|P^{t_\mu^S(\varepsilon)}(x,\cdot)\, P^{\bar{t}} - \mu\right\|_{\mathrm{TV}} \\
&\leqslant \left\|P^{t_\mu^S(\varepsilon)}(x,\cdot)\, P^{\bar{t}} - \mu P^{\bar{t}}\right\|_{\mathrm{TV}} + \left\|\mu P^{\bar{t}} - \mu\right\|_{\mathrm{TV}} \\
&\leqslant \left\|P^{t_\mu^S(\varepsilon)}(x,\cdot) - \mu\right\|_{\mathrm{TV}} + \left\|\mu P^{\bar{t}} - \mu\right\|_{\mathrm{TV}} \leqslant 2\varepsilon
\end{aligned}$$

□

3.1 Example: A simple three-state Markov chain

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

$$P = \begin{pmatrix} \varepsilon & \frac{1-\varepsilon}{2} & \frac{1-\varepsilon}{2} \\ \varepsilon & 1-\varepsilon & 0 \\ \varepsilon & 0 & 1-\varepsilon \end{pmatrix}$$

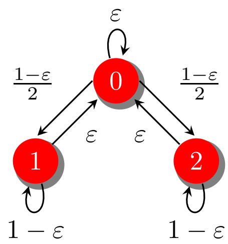

The chain is ergodic with stationary distribution $\pi = (\varepsilon, (1-\varepsilon)/2, (1-\varepsilon)/2)$, and its mixing time is $t_{\text{mix}} = \Theta(1/\varepsilon)$. Hence the mixing time increases as ε tends to zero.

Now observe that, for every $\delta > \varepsilon$, degenerate[1] distributions $\mu_1 = (0,1,0)$ and $\mu_2 = (0,0,1)$ are $(\delta, \Theta(\delta/\varepsilon))$-metastable according to Definition 3.1. If we start from the first state (i.e. from degenerate distribution $\nu = (1,0,0)$), after one step we are in the stationary distribution.

Hence, even if the mixing time can be arbitrary large, for every starting state x there is a $(1/4, \Theta(t_{\text{mix}}))$-metastable distribution μ that is quickly (in constant time, independent of ε) reached from x.

3.2 Example: Two-player coordination games

Coordination games[2] with two players are examples of games where the mixing time is a function increasing exponentially in β. The transition matrix of the logit dynamics for such games is

$$P = \begin{pmatrix} \begin{array}{c|cccc} & (0,0) & (0,1) & (1,0) & (1,1) \\ \hline (0,0) & 1-\varepsilon & \varepsilon/2 & \varepsilon/2 & 0 \\ (0,1) & (1-\varepsilon)/2 & (\varepsilon+\delta)/2 & 0 & (1-\delta)/2 \\ (1,0) & (1-\varepsilon)/2 & 0 & (\varepsilon+\delta)/2 & (1-\delta)/2 \\ (1,1) & 0 & \delta/2 & \delta/2 & 1-\delta \end{array} \end{pmatrix}$$

where

$$\varepsilon = \frac{1}{1+e^{(a-d)\beta}}\,; \qquad \delta = \frac{1}{1+e^{(b-c)\beta}}$$

and a, b, c, d are the parameters of the coordination game with $a > d$, $b > c$, and $a - d > b - c$ (see Section 4.1 in [3] for further details). The stationary distribution for P is

$$\pi = \frac{1}{\varepsilon+\delta}\left[\delta(1-\varepsilon),\, \varepsilon\delta,\, \varepsilon\delta,\, \varepsilon(1-\delta)\right]$$

In [3] we proved that the mixing time for such games is $t_{\text{mix}} = \Theta(1/\delta)$.
Now consider the special case when $\varepsilon = \delta$, hence the stationary distribution is

$$\pi = ((1-\varepsilon)/2,\, \varepsilon/2,\, \varepsilon/2,\, (1-\varepsilon)/2)$$

Let $\mu_{(0,0)}$ and $\mu_{(1,1)}$ be the two distributions concentrated in states $(0,0)$ and $(1,1)$ respectively, i.e.

$$\mu_{(0,0)} = [1,0,0,0] \quad \text{and } \mu_{(1,1)} = [0,0,0,1]$$

Observe that, if we start from $\mu_{(0,0)}$ or $\mu_{(1,1)}$, after one step of the chain we are respectively in distributions

$$\mu_{(0,0)}P = [1-\varepsilon,\, \varepsilon/2,\, \varepsilon/2,\, 0] \quad \text{and} \quad \mu_{(1,1)}P = [0,\, \varepsilon/2,\, \varepsilon/2,\, 1-\varepsilon]$$

Hence

$$\left\|\mu_{(0,0)}P - \mu_{(0,0)}\right\|_{\text{TV}} = \left\|\mu_{(1,1)}P - \mu_{(1,1)}\right\|_{\text{TV}} = \varepsilon$$

[1] A probability distribution is *degenerate* if it is concentrated in one single element

[2] See Section 5 for the definition of coordination games.

By using Lemma 3.4, we have that $\mu_{(0,0)}$ and $\mu_{(1,1)}$ are $(1/4, \Theta(1/\varepsilon))$-metastable according to Definition 3.1. Moreover, if the chain starts from state $(0,1)$ or from state $(1,0)$, after 1 step of the chain we are ε-close to the stationary distribution π, indeed

$$(0,1,0,0)P = \left[\frac{1-\varepsilon}{2},\, \varepsilon,\, 0,\, \frac{1-\varepsilon}{2}\right], \quad (0,0,1,0)P = \left[\frac{1-\varepsilon}{2},\, 0,\, \varepsilon,\, \frac{1-\varepsilon}{2}\right]$$

and

$$\|(0,1,0,0)P - \pi\|_{\mathrm{TV}} = \|(0,1,0,0)P - \pi\|_{\mathrm{TV}} = \varepsilon\,.$$

We can summarize what we have just shown in the following theorem.

Theorem 3.5 *Let P be the transition matrix of the logit dynamics for a 2-player coordination game. For every starting profile $\mathbf{x} \in \Omega$ there is a $(1/4, \Theta(t_{mix}))$-metastable distribution $\mu_{\mathbf{x}}$ such that $t_{\mu_{\mathbf{x}}}^{\{\mathbf{x}\}} = \Theta(1)$.*

3.3 Metastability and the bottleneck ratio

Consider an ergodic Markov chain P with state space Ω and stationary distribution π. For a subset S of states, the *bottleneck ratio* $B(S)$ is defined as

$$B(S) = \frac{Q(S,\bar{S})}{\pi(S)} \qquad \text{where } \; Q(S,\bar{S}) = \sum_{x\in S}\sum_{y\in\Omega\setminus S}\pi(x)P(x,y)\,.$$

Let π_S be the stationary distribution conditioned on S, i.e.

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

It is well-known (see e.g. Theorem 7.3 in [25]) that the bottleneck ratio at set S equals the total variation distance between π_S and $\pi_S P$, i.e., $\|\pi_S P - \pi_S\| = B(S)$. Hence, the following lemma about the metastability of π_S holds.

Lemma 3.6 *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.*

3.4 Pseudo-mixing time tools

In order to upper bound the mixing time of an ergodic chain, it is often used the fact that, for every starting state $x \in \Omega$ the total variation distance between the distribution of the chain at time t and the stationary distribution π is upper bounded by the maximum, over all states $y \in \Omega$, of the total variation between the chain starting at x and the chain starting at y (see Lemma 4.11 in [25]), i.e.

$$\left\|P^t(x,\cdot) - \pi\right\|_{\mathrm{TV}} \leqslant \max_{y\in\Omega}\left\|P^t(x,\cdot) - P^t(y,\cdot)\right\|_{\mathrm{TV}}$$

In the following lemma we formalize and prove an analogous statement for metastable distributions.

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 $S \subseteq \Omega$ of the state space. Then for every $x \in S$ and every $1 \leqslant t \leqslant T$, it holds that*

$$\left\|P^t(x,\cdot) - \mu\right\|_{\mathrm{TV}} \leqslant \varepsilon + \max_{y\in S}\left\|P^t(x,\cdot) - P^t(y,\cdot)\right\|_{\mathrm{TV}}\,.$$

Proof. From triangle inequality we have

$$\left\|P^t(x,\cdot) - \mu\right\|_{\mathrm{TV}} \leqslant \left\|P^t(x,\cdot) - \mu P^t\right\|_{\mathrm{TV}} + \left\|\mu P^t - \mu\right\|_{\mathrm{TV}}$$

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

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

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

$$
\begin{aligned}
\left\| P^t(x, \cdot) - \mu P^t \right\|_{\mathrm{TV}} &= \max_{A \subseteq \Omega} |P^t(x, A) - \mu P^t(A)| \\
&\leqslant \max_{A \subseteq \Omega} \max_{y \in S} \left| P^t(x, A) - P^t(y, A) \right| = \max_{y \in S} \left\| P^t(x, \cdot) - P^t(y, \cdot) \right\|_{\mathrm{TV}} .
\end{aligned}
$$

□

In some cases metastable distributions are concentrated in one single state. The following lemma shows that, in those cases, the hitting time of such a state can be used to establish the pseudo-mixing time of the metastable distribution.

Lemma 3.8 *Let P be a Markov chain with finite state space Ω and let μ be an (ε, T)-metastable distribution concentrated on a single state y. Let τ_y be the hitting time of this state. Then for all $x \in \Omega$ and $1 \leqslant t \leqslant T$, we have*

$$
\left\| P^t(x, \cdot) - \mu \right\|_{\mathrm{TV}} \leqslant \varepsilon + (1 - \varepsilon) \mathbf{P}_x \left(\tau_y > t \right) .
$$

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

$$
\begin{aligned}
\left\| P^t(x, \cdot) - \mu \right\|_{\mathrm{TV}} &= \mathbf{P}_x \left(X_t \neq y \right) = \mathbf{P}_x \left(X_t \neq y, \tau_y \leqslant t \right) + \mathbf{P}_x \left(X_t \neq y, \tau_y > t \right) \\
&= \mathbf{P}_x \left(X_t \neq y \,|\, \tau_y \leqslant t \right) \mathbf{P}_x \left(\tau_y \leqslant t \right) + \mathbf{P}_x \left(\tau_y > t \right)
\end{aligned}
$$

Moreover, observe that

$$
\begin{aligned}
\mathbf{P}_x \left(X_t \neq y \,|\, \tau_y \leqslant t \right) &= \sum_{k \leqslant t} \mathbf{P}_x \left(X_t \neq y \,|\, \tau_y = k \right) \mathbf{P}_x \left(\tau_y = k \,|\, \tau_y \leqslant t \right) \\
&= \sum_{k \leqslant t} \mathbf{P}_y \left(X_{t-k} \neq y \right) \mathbf{P}_x \left(\tau_y = k \,|\, \tau_y \leqslant t \right) \\
&= \sum_{k \leqslant t} \left\| \mu P^{t-k} - \mu \right\|_{\mathrm{TV}} \mathbf{P}_x \left(\tau_y = k \,|\, \tau_y \leqslant t \right) \\
&\leqslant \varepsilon \sum_{k \leqslant t} \mathbf{P}_x \left(\tau_y = k \,|\, \tau_y \leqslant t \right) = \varepsilon
\end{aligned}
$$

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

$$
\begin{aligned}
\left\| P^t(x, \cdot) - \mu \right\|_{\mathrm{TV}} &= \mathbf{P}_x \left(X_t \neq y \,|\, \tau_y \leqslant t \right) \mathbf{P}_x \left(\tau_y \leqslant t \right) + \mathbf{P}_x \left(\tau_y > t \right) \\
&\leqslant \varepsilon \mathbf{P}_x \left(\tau_y \leqslant t \right) + \mathbf{P}_x \left(\tau_y > t \right) = \varepsilon + (1 - \varepsilon) \mathbf{P}_x \left(\tau_y > t \right)
\end{aligned}
$$

□

4 Ising model on the complete graph

Consider the following game, that here we call *Ising game*: each one of n players has two strategies, -1 and $+1$, and the utility of player i at profile $\mathbf{x} = (x_1, \ldots, x_n) \in \{-1, +1\}^n$ is $u_i(\mathbf{x}) = x_i \sum_{j \neq i} x_j$. This is the "game-theoretic" formulation of the well-studied Curie-Weiss model (the Ising model on the complete graph).

Observe that for every player i it holds that $u_i(\mathbf{x}_{-i}, +1) - u_i(\mathbf{x}_{-i}, -1) = H(\mathbf{x}_{-i}, +1) - H(\mathbf{x}_{-i}, -1)$ where $H(\mathbf{x}) = \sum_{\{j,k\} \in \binom{[n]}{2}} x_j x_k$, hence the Ising game is a potential game with potential function H and, as observed in Section 2 the logit dynamics has stationary distribution $\pi(\mathbf{x}) = e^{\beta H(\mathbf{x})}/Z$.

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

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

In this section we study the metastability properties of the logit dynamics for the Ising game (equivalently the Glauber dynamics for the Curie-Weiss model) from our quantitative point of view. Namely, we show that, if we start from a profile where the number of $+1$ (respectively -1) is a sufficiently large majority, and if β is large enough then, after an initial *pseudo-mixing* phase, the distribution of the chain at time t is close, in total variation distance, to the degenerate distribution concentrated in the profile with all $+1$'s (respectively all -1's) for all $t = \mathcal{O}(\mathrm{poly}(n))$.

Let π_+ and π_- be the two degenerate distributions concentrated in the states with all $+1$ and all -1, respectively. The next lemma shows that, for $\beta = \Omega(\log n/n)$, π_+ and π_- are metastable for a polynomially-long time.

Lemma 4.1 *If $\beta > c \log n/n$ then π_+ and π_- are $\left(1/n, n^{c-2}\right)$-metastable.*

Proof. We prove the result for π_+, exactly the same proof (by swapping minuses and pluses) works for π_-. Since $\pi_+(\mathbf{x}) = 0$ for all $\mathbf{x} \neq +\mathbf{1}$ and

$$\pi_+ P(\mathbf{x}) = \begin{cases} 0 & \text{if } S(\mathbf{x}) < (n-2)/n \ \text{ (i.e if } \mathbf{x} \text{ contains more than one “} -1\text{”)} \\ \frac{1}{n} \cdot \frac{1}{1+e^{\beta(n-2)}} & \text{if } S(\mathbf{x}) = (n-2)/n \ \text{ (i.e. if } \mathbf{x} \text{ contains exactly one “} -1\text{”)} \\ \frac{1}{1+e^{-\beta(n-2)}} & \text{if } S(\mathbf{x}) = 1 \ \text{ (i.e. if } \mathbf{x} = +\mathbf{1}) \end{cases}$$

the total variation distance between $\pi_+ P$ and π_+ is

$$\|\pi_+ P - \pi_+\| = \frac{1}{2} \sum_{\mathbf{x} \in \{-1,+1\}^n} |\pi_+ P(\mathbf{x}) - \pi_+(\mathbf{x})| = \frac{1}{1 + e^{\beta(n-2)}} \leqslant e^{-\beta(n-2)} \leqslant n^{c-1} \,.$$

In the last inequality we used $\beta \geqslant c \log n/n$. Hence π_+ is $(n^{-(c-1)}, 1)$-metastable. The thesis follows from Lemma 3.2. $\square$

In order to give an upper bound on the pseudo-mixing time we need some preliminary results about birth-and-death chains, which we show in the next subsection.

4.1 Biased birth-and-death chains

In this section we consider birth-and-death chains with state space $\Omega = \{0, 1, \ldots, n\}$ (see Chapter 2.5 in [25] for a detailed description of such chains). For $k \in \{1, \ldots, n-1\}$ let $p_k = \mathbf{P}_k\left(X_1 = k+1\right)$, $q_k = \mathbf{P}_k\left(X_1 = k-1\right)$, and $r_k = 1 - p_k - q_k = \mathbf{P}_k\left(X_1 = k\right)$. We will be interested in the probability that the chain starting at some state $h \in \Omega$ hits state 0 before state n, namely $\mathbf{P}_k\left(X_{\tau_{0,n}} = n\right)$ 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 4.2 *Suppose for all $k \in \{1, \ldots, n-1\}$ it holds that $p_k = \varepsilon$ and $q_k = \delta$, for some ε and δ with $\varepsilon + \delta \leqslant 1$. Then the probability the chain hits state n before state 0 starting from state $h \in \Omega$ is*

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

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

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

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

$$\alpha_k = \delta \cdot \alpha_{k-1} + \varepsilon \cdot \alpha_{k+1} + (1 - (\delta + \varepsilon))\, \alpha_k . \tag{5}$$

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

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

From $\alpha_n = 1$ we get

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

Hence

$$\alpha_k = \frac{1 - (\delta/\varepsilon)^k}{1 - (\delta/\varepsilon)^n} . \tag{6}$$

□

Lemma 4.3 *Suppose for all $k \in \{1, \ldots, n-1\}$ it holds that $p_k \geqslant \varepsilon$ and $q_k \leqslant \delta$, for some ε and δ with $\varepsilon + \delta \leqslant 1$. Then the probability to hit state n before state 0 starting from state $h \in \Omega$ is*

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

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

$$\mathbf{P}_k\left(Y_1 = k-1\right) = \delta \qquad \mathbf{P}_k\left(Y_1 = k+1\right) = \varepsilon$$

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. 4.1. Let U be a uniform r.v. over the interval $[0,1]$ and chose the update for the two chains according to position of U in the two intervals.

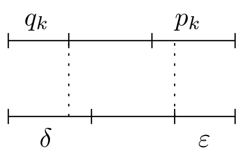

Figure 1: Coupling

Observe that, since $p_k \geqslant \varepsilon$ and $q_k \leqslant \delta$, 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 $X_t \geqslant Y_t$. Hence if chain Y_t hits state n before state 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, hence

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

Thus

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

In the last inequality we used Lemma 4.2. □

Lemma 4.4 *Suppose for all $k \in \{1, \dots, n-1\}$ it holds that $q_k/p_k \leqslant \alpha$, for some $\alpha < 1$. Then the probability to hit state 0 before state n starting from state $h \in \Omega$ is*

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

Idea of the 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 $\hat{q}_k$. Observe that $\mathbf{P}_k \left(X_{\tau^X_{0,n}} = 0 \right) = \mathbf{P}_k \left(Y_{\tau^Y_{0,n}} = 0 \right)$, indeed it is easy to define a coupling of the two chains where chain X_t *follows the path* traced by chain Y_t: Let $\{U_t\}$ be i.i.d. random variables uniformly distributed over $[0,1]$ and use such variables to run chain Y_t. As for chain X_t, when at state h toss a coin that gives head with probability $p_h + q_h$; if it gives tail than don't move, if it gives head than move to one of the two adjacent states by using the first r.v. U_i that has not been previously used in process X_t. In this way 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.

Finally observe that $\hat{q}_k/\hat{p}_k = q_k/p_k \leqslant \alpha$, and since $\hat{p}_k + \hat{q}_k = 1$ it implies that $\hat{p}_k \geqslant 1/(1+\alpha)$ and $\hat{q}_k \leqslant \alpha/(1+\alpha)$. The thesis then follows from Lemma 4.3. □

4.2 Convergence time at low temperature

If X_t is the logit dynamics for the Ising game, the magnetization process $S_t := S(X_t)$ is itself a Markov chain, with state space $\Omega = \{-n, -n+2, \cdots, n-4, n-2, n\}$. When at state $k \in \Omega$, the probability to go right (to state $k+2$) or left (to state $k-2$) is respectively

$$\mathbf{P}_k (S_1 = k+2) = p_k = \frac{n-k}{2n} \frac{1}{1+e^{-2(k+1)\beta}} ; \qquad \mathbf{P}_k (S_1 = k-2) = q_k = \frac{n+k}{2n} \frac{1}{1+e^{2(k-1)\beta}} . \tag{7}$$

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

For $a, b \in [-n, n]$, with $a < b$, let $\tau_{a,b}$ be the random variable indicating the first time the chain reaches a state x with $x \leqslant a$ or $x \geqslant b$,

$$\tau_{a,b} = \inf\{t \in \mathbb{N} \,:\, Y_t \leqslant a \text{ or } Y_t \geqslant b\}$$

At time $\tau_{a,b}$, chain $Y_{\tau_{a,b}}$ can be in one out of two states, namely the larger state smaller than a or the smallest state larger than b. We need to 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 next lemma we show that, if the chain starts from a sufficiently large positive state k, and if $\beta k^2 \geqslant c \log n$ for a suitable constant c, then when chain Y_t gets out of interval $(0, n/2)$, it happens on the $n/2$ side w.h.p.

Lemma 4.5 *Let $k \in \Omega$ be the starting state with $4 \leqslant k \leqslant n/2$. If $\beta \geqslant 6/n$ and $\beta k^2 \geqslant 16 \log n$, then*

$$\mathbf{P}_k \left(Y_{\tau_{0,n/2}} \leqslant 0 \right) \leqslant 1/n$$

Proof. According to (7), 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 $h \geqslant 2$ it holds that

$$\frac{1 + e^{-2(h+1)\beta}}{1 + e^{2(h-1)\beta}} \leqslant e^{-2(h-1)\beta} \leqslant e^{-h\beta} \tag{8}$$

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

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

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

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

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

Thus, for each state h of the chain with $k/2 \leqslant h \leqslant n/2$ 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 4.4 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 $4 \leqslant k \leqslant n/2$, if $\beta k^2 \geqslant 16 \log n$, the chain starting at k hits state $n/2$ before state $k/2$ w.h.p.

$$\mathbf{P}_k \left(Y_{\tau_{k/2,n/2}} \leqslant k/2 \right) \leqslant e^{\frac{1}{16}\beta k^2} \leqslant \frac{1}{n}$$

The thesis follows by observing that $\mathbf{P}_k \left(Y_{\tau_{0,n/2}} \leqslant 0 \right) \leqslant \mathbf{P}_k \left(Y_{\tau_{k/2,n/2}} \leqslant k/2 \right)$. □

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

Lemma 4.6 *Let $k \in \Omega$ be the starting state with $n/2 \leqslant k \leqslant n-1$. If $\beta \geqslant 8 \log n / n$, then*

$$\mathbf{P}_k \left(Y_{\tau_{0,n}} \leqslant 0 \right) \leqslant (2/n)^{n/8}$$

Proof. Observe that for every $h \leqslant n-1$ it holds that $\frac{n+h}{n-h} \leqslant 2n$, and by using it together with (8) we have that $q_h/p_h \leqslant 2ne^{-\beta h}$ for every $2 \leqslant h \leqslant n-1$. Thus, for every $k/2 \leqslant h \leqslant n-1$ it holds that $q_h/p_h \leqslant 2ne^{-\frac{1}{2}\beta k} \leqslant 2/n$, where in the last inequality we used $k \geqslant n/2$ and $\beta \geqslant 8 \log n / n$. Hence, if the chain starts at k, by applying Lemma 4.4 it follows that the probability of reaching $k/2$ before reaching n is less than $(2/n)^{\ell}$, where $\ell = k/4 \geqslant n/8$ is the number of states between $k/2$ and k. Hence,

$$\mathbf{P}_k \left(Y_{\tau_{0,n}} \leqslant 0 \right) \leqslant \mathbf{P}_k \left(Y_{\tau_{k/2,n}} \leqslant k/2 \right) \leqslant (2/n)^{n/8}$$

□

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 4.7 *For every $k \in \Omega$ with $k \geqslant 0$ it holds that $\mathbf{E}_k [\tau_{0,n}] \leqslant n^3$.*

Idea of the proof. For the *unbiased* birth-and-death chain, i.e. when $p_k = q_k = 1/2$ for all k, it is well-known that the expectation is $\mathcal{O}(n^2)$. If the chain has a drift toward one of the end points, e.g. $p_k = p > 1/2$, $q_k = 1 - p$ for all k, then such expected time cannot be larger. In the magnetization chain we have $q_k < p_k$ for all positive k, but since $p_k + q_k < 1$ we have also to consider the time the chain spends *not moving*. At state k the chain moves with probability $(p_k + q_k)$, hence the expected time the chain stays in a state before moving is at most $\max_k \{1/(p_k + q_k)\}$. Since $p_k + q_k \geqslant 1/n$ for all $k \in \{1, \ldots, n-1\}$ it follows that n^3 is a (rough) upper bound on the expected time for our chain. □

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

Theorem 4.8 *Let* $\mathbf{x}$ *be a profile whose magnetization* $S(\mathbf{x})$ *has absolute value* $|S(\mathbf{x})| = k$. *If* $\beta \geqslant c\log n/n$ *and* $k^2 > c\log n/\beta$ *then*

$$\|P^t(\mathbf{x},\cdot) - \pi_+\| = \mathcal{O}(1/n)$$

for every $n^4 \leqslant t \leqslant n^{c-2}$.

Proof. Consider w.l.o.g. the case of starting state with positive magnetization, $S(\mathbf{x}) = k$. Let τ_n be the first time the chain hits state with all $+1$ and let $\tau_{0,n}$ be the first time the magnetization of the chain is either n or less than or equal to 0,

$$\tau_{0,n} = \min\{t \in \mathbb{N} \,:\, S(X_t) = n \text{ or } S(X_t) \leqslant 0\}$$

Since $\{\tau_n > t,\, \tau_{0,n} \leqslant t\}$ implies that the magnetization chain reaches 0 before reaching n we have

$$\begin{aligned}
\mathbf{P}_{\mathbf{x}}\left(\tau_n > t\right) &= \mathbf{P}_{\mathbf{x}}\left(\tau_n > t,\, \tau_{0,n} > t\right) + \mathbf{P}_{\mathbf{x}}\left(\tau_n > t,\, \tau_{0,n} \leqslant t\right) \\
&\leqslant \mathbf{P}_{\mathbf{x}}\left(\tau_{0,n} > t\right) + \mathbf{P}_{\mathbf{x}}\left(S(X_{\tau_{0,n}}) \leqslant 0\right) \\
&\leqslant \frac{\mathbf{E}_{\mathbf{x}}\left[\tau_{0,n}\right]}{t} + \mathbf{P}_{\mathbf{x}}\left(Y_{\tau_{0,n}} \leqslant 0\right)
\end{aligned}$$

Where Y_t is the birth-and-death chain with state space Ω and transition rates as in (7). As for the first term of the sum, from Lemma 4.7 it follows that $\mathbf{E}_{\mathbf{x}}\left[\tau_{0,n}\right]/t \leqslant 1/n$ for $t \geqslant n^4$. As for the second term, by conditioning on the position of the chain when it gets out of subinterval $(0, n/2)$ we have

$$\begin{aligned}
\mathbf{P}_k\left(Y_{\tau_{0,n}} \leqslant 0\right) &= \mathbf{P}_k\left(Y_{\tau_{0,n}} \leqslant 0 \,|\, Y_{\tau_{0,n/2}} \leqslant 0\right) \mathbf{P}_k\left(Y_{\tau_{0,n/2}} \leqslant 0\right) + \\
&\quad + \mathbf{P}_k\left(Y_{\tau_{0,n}} \leqslant 0 \,|\, Y_{\tau_{0,n/2}} \geqslant n/2\right) \mathbf{P}_k\left(Y_{\tau_{0,n/2}} \geqslant n/2\right) \\
&\leqslant \mathbf{P}_k\left(Y_{\tau_{0,n/2}} \leqslant 0\right) + \mathbf{P}_k\left(Y_{\tau_{0,n}} \leqslant 0 \,|\, Y_{\tau_{0,n/2}} \geqslant n/2\right)
\end{aligned}$$

From Lemma 4.5 we have that $\mathbf{P}_k\left(Y_{\tau_{0,n/2}} \leqslant 0\right) \leqslant 1/n$, and observe that

$$\mathbf{P}_k\left(Y_{\tau_{0,n}} \leqslant 0 \,|\, Y_{\tau_{0,n/2}} \geqslant n/2\right) \leqslant \mathbf{P}_{n/2}\left(Y_{\tau_{0,n}} \leqslant 0\right) \leqslant (2/n^{n/8})$$

where the last inequality follows from Lemma 4.6. Hence for every $t \geqslant n^4$ it holds that $\mathbf{P}_{\mathbf{x}}\left(\tau_n > t\right) \leqslant 3/n$. Since π_+ is $(1/n, n^{c-2})$-metastable when $\beta \geqslant c\log n/n$, the thesis follows from Lemma 3.4. $\square$

5 Graphical Coordination games on rings

In this section we study coordination games on a ring [15]. Consider n players identified with the integers $\{0, \ldots, n-1\}$; the two neighbors of player i are $(i-1) \bmod n$ (the *left* neighbor) and $(i+1) \bmod n$ (the *right* neighbor). Every player can pick one of two strategies, 0 and 1, and use it in the two instances of the following basic coordination game, played with each of her two neighbors:

	0	1
0	a, a	c, d
1	d, c	b, b

We assume that $a > d$ and $b > c$ which implies that players have an advantage in selecting the same strategy of their neighbors and that profiles $(0,0)$ and $(1,1)$ are Nash equilibria of the basic coordination game. The utility of a player is the sum of the utilities gained in each basic coordination game she plays.

The set of profiles of the game is $\Omega = \{0,1\}^n$. We define $S_d \subseteq \Omega$ as the set of profiles where *exactly* d players are playing 0 and $R \subseteq \Omega$ as the set of profiles in which at least two *adjacent* players are playing 0. We also set $\tilde{S}_d = \bigcup_{i=d}^{n} S_i$. Further, we denote by $\mathbf{0}$ the profile where all players are playing 0 and by $\mathbf{1}$ the profile where all players are playing 1.

We define $\Delta := a - d$ and $\delta := b - c$ and assume, w.l.o.g., that $\Delta \geqslant \delta$. It is easy to see that the graphical coordination game on a ring is a potential game with potential function $\Phi(\mathbf{x}) = \sum_{i=0}^{n-1} \Phi_i(\mathbf{x})$, where

$$\Phi_i(\mathbf{x}) = \begin{cases} 0, & \text{if } x_i \neq x_{i+1}\,; \\ \delta, & \text{if } x_i = x_{i+1} = 1\,; \\ \Delta, & \text{if } x_i = x_{i+1} = 0\,. \end{cases}$$

In [2] it is showed that the mixing time of the logit dynamics for this game is $\Omega\left(e^{2\delta\beta}\right)$. If $\Delta > \delta$, then 0 is called the *risk dominant* strategy [19] and, in this case, the techniques of [5] can be generalized to obtain an upper bound to the mixing time that is polynomial in n and exponential in $\Delta + \delta$ and β. If instead $\Delta = \delta$ (and thus no risk dominant strategy exists) an almost matching upper bound is given in [2]. 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 $\beta = \omega(\log n)$.

5.1 Games with risk dominant strategies

In this section we study the case $\Delta > \delta$ and prove that for $\beta = \omega(\log n)$, the logit dynamics reaches in polynomial time a metastable distribution and remains close to it for super-polynomial time. On the other hand, we know that for $\beta = \mathcal{O}(\log n)$ the logit dynamics is rapidly mixing (and thus reaches in polynomial time the stationary distribution and stays close to it forever).

Theorem 5.1 *For $\beta = \omega(\log n)$, for every $\varepsilon > 0$ and for every $\mathbf{x} \in \Omega$, the dynamics starting from $\mathbf{x}$ approaches in polynomial time a distribution that is $(\varepsilon, T(n))$, for a super-polynomial function T of n.*

Throughout this section, we set $S_d^\star = R \cup \tilde{S}_d$.

5.1.1 Metastable distributions

The next theorem identifies three classes of metastable distributions. We remind the reader that π_S is the stationary distribution (that is the Gibbs measure defined in Eq. 3) restricted to the set S of profiles (see Eq. 4).

Theorem 5.2 *For every $\varepsilon > 0$, $\pi_{\mathbf{1}}$ is $(\varepsilon, \varepsilon \cdot e^{2\delta\beta})$-metastable and $\pi_{\mathbf{0}}$ is $(\varepsilon, \varepsilon \cdot e^{2\Delta\beta})$-metastable. Moreover, for $\beta = \omega(\log n)$ and constant $d > 0$, $\pi_{S_d^\star}$ is $(\varepsilon, e^{\Omega(n \log n)})$-metastable.*

Proof. The bottleneck ratio of $\mathbf{1}$ is

$$B(\mathbf{1}) = \frac{\sum_{\mathbf{x} \neq 1} \pi(\mathbf{1}) P(\mathbf{1}, \mathbf{x})}{\pi(\mathbf{1})} = e^{-2\delta\beta}\,.$$

Thus by Lemma 3.6, we have that $\pi_{\mathbf{1}}$ is $(e^{-2\delta\beta}, 1)$-metastable. By applying Lemma 3.2, we obtain that $\pi_{\mathbf{1}}$ is $(\varepsilon, \varepsilon \cdot e^{2\delta\beta})$-metastable for every $\varepsilon > 0$.

Similarly, the bottleneck ratio of $\mathbf{0}$ is

$$B(\mathbf{0}) = \frac{\sum_{\mathbf{x} \neq \mathbf{0}} \pi(\mathbf{0}) P(\mathbf{0}, \mathbf{x})}{\pi(\mathbf{0})} = e^{-2\Delta\beta}$$

and thus, by applying Lemma 3.6 and Lemma 3.2, we have that $\pi_{\mathbf{0}}$ is $(\varepsilon, \varepsilon \cdot e^{2\Delta\beta})$-metastable for every ε.

Finally, the bottleneck ratio of $S_d^\star$ is

$$\begin{aligned}
B(S_d^\star) &= \frac{\sum_{\mathbf{x} \in \partial S_d^\star} \pi(\mathbf{x}) \sum_{\mathbf{y} \in \Omega \setminus S_d^\star} P(\mathbf{x}, \mathbf{y})}{\sum_{\mathbf{x} \in S_d^\star} \pi(\mathbf{x})} \\
&\leqslant \frac{\sum_{\mathbf{x} \in \partial S_d^\star} \pi(\mathbf{x})}{\sum_{\mathbf{x} \in S_d^\star} \pi(\mathbf{x})} && \left(\textstyle\sum_{\mathbf{y} \in \Omega \setminus S_d^\star} P(\mathbf{x}, \mathbf{y}) \leqslant 1\right) \\
&\leqslant \frac{n^{d+1} \max_{\mathbf{x} \in \partial S_d^\star} \pi(\mathbf{x})}{\max_{\mathbf{x} \in S_d^\star} \pi(\mathbf{x})} && (|\partial S_d^\star| \leqslant n^{d+1}) \\
&\leqslant \frac{n^{d+1} e^{[(n-d-1)\delta + (d-1)\Delta]\beta}}{e^{n\Delta\beta}} && (\text{max when } d \text{ 0's are adjacent}) \\
&\leqslant n^{d+1} e^{-(n-d-1)(\Delta-\delta)\beta} \\
&\leqslant n^{-n+2(d+1)} && (\beta = \omega(\log n) \text{ and } (\Delta - \delta) \text{ is a positive constant}) \\
&= e^{-\Omega(n \log n)}, && (\text{constant } d)
\end{aligned}$$

where $\partial S_d^\star$ is the set of all profiles in $S_d^\star$ with at least a neighbor in $\Omega \setminus S_d^\star$. Thus, by applying Lemma 3.6 and Lemma 3.2, we have that $\pi_{S_d^\star}$ is $(\varepsilon, e^{\Omega(n \log n)})$-metastable for every ε. $\square$

5.1.2 Pseudo-mixing time

In this subsection we look at the the pseudo-mixing time of the metastable distributions described in Theorem 5.2 and we show that, for every starting profile, the dynamics rapidly approaches one of them. We remind the reader that the interesting case is $\beta = \omega(\log n)$ as for $\beta = \mathcal{O}(\log n)$ the mixing time of the logit dynamics is polynomial in n.

Not to overburden our notation, we will denote distribution $\pi_{S_d^\star}$ by π_d.

Our proof distinguishes cases depending on the starting profile $\mathbf{x}$ of the chain. We start by considering $\mathbf{x} \in S_d^\star$, for constant d, and show (see Theorem 5.7) that the pseudo-mixing time of π_d is polynomial. Finally, in Theorem 5.9, we show that if the dynamics starts from one of the remaining profiles, in polynomial time it hits either $S_d^\star$ or $\mathbf{1}$ with high probability.

Starting from $S_d^\star$. From the definition of pseudo mixing time and by Lemma 3.7, for any S, we can bound $t_{\pi_S}^S(\gamma)$ in terms of the total variation distance of two copies of the same Markov chain starting in different states. Next lemma relates this quantity, for the logit dynamics we are studying, to the hitting time $\tau_{\mathbf{0}}$ of profile $\mathbf{0}$.

Lemma 5.3 *For every $S \subseteq \Omega$ and for $\mathbf{x}, \mathbf{y} \in S$ we have*

$$\left\| P^t(\mathbf{x}, \cdot) - P^t(\mathbf{y}, \cdot) \right\|_{\mathrm{TV}} \leqslant 2 \cdot \max_{\mathbf{z} \in S} \mathbf{P}_{\mathbf{z}} \left(\tau_{\mathbf{0}} \geqslant t \right) .$$

Thus, by the previous lemma and Lemma 3.7, in order to bound $t_{\pi_d}^{S_d^\star}(\gamma)$, it is sufficient to give an upper bound on $\mathbf{P}_{\mathbf{x}}\left(\tau_{\mathbf{0}} \geqslant t\right)$, for $\mathbf{x} \in S_d^\star$. The next lemma bounds the hitting time of $\mathbf{0}$ when starting from R.

Lemma 5.4 *For $\beta = \omega(\log n)$, for every $\lambda > 0$ and every $\mathbf{x} \in R$, we have*

$$\mathbf{P}_{\mathbf{x}} \left(\tau_{\mathbf{0}} > \frac{(8-\lambda)n^2}{\lambda} \right) \leqslant \frac{\lambda}{4} .$$

Next we show that, when starting from $\mathbf{x} \in S_d$, the dynamics hits R in polynomial time.

Lemma 5.5 *For $\beta = \omega(\log n)$, for every $d > 0$ and for every profile $\mathbf{x} \in S_d$,*

$$\mathbf{P}_{\mathbf{x}} \left(\tau_R \leqslant n^2 \right) \geqslant \frac{2d}{2d+1} \left(1 - \mathsf{negl}(n) \right) .$$

Finally, we have

Lemma 5.6 *For $\beta = \omega(\log n)$, for every $d > 0$, for $\mathbf{x} \in S_d^\star$ and for every $\lambda > 0$,*

$$\mathbf{P}_{\mathbf{x}} \left(\tau_{\mathbf{0}} > \frac{8n^2}{\lambda} \right) \leqslant \frac{1}{2d+1} + \frac{\lambda}{4} .$$

Proof. We need to consider only $x \in S_d^\star \setminus R$. By Lemma 5.4 and Lemma 5.5, we have

$$\begin{aligned}
\mathbf{P}_{\mathbf{x}} \left(\tau_{\mathbf{0}} \leqslant \frac{8n^2}{\lambda} \right) &\geqslant \mathbf{P}_{\mathbf{x}} \left(\tau_{\mathbf{0}} \leqslant \frac{8n^2}{\lambda} \mid \tau_R \leqslant n^2 \right) \mathbf{P}_{\mathbf{x}} \left(\tau_R \leqslant n^2 \right) \\
&\geqslant \mathbf{P}_{X_{\tau_R}} \left(\tau_{\mathbf{0}} \leqslant \frac{(8-\lambda)n^2}{\lambda} \right) \mathbf{P}_{\mathbf{x}} \left(\tau_R \leqslant n^2 \right) \\
&\geqslant \left(1 - \frac{\lambda}{4} \right) \frac{2d}{2d+1} \left(1 - \mathsf{negl}(n) \right) \geqslant \frac{2d}{2d+1} - \frac{\lambda}{4}.
\end{aligned}$$

□

We are now ready to prove an upper bound on the pseudo-mixing time of π_d, when starting from a profile in $S_d^\star$.

Theorem 5.7 *For $\beta = \omega(\log n)$, for constant $d > 1$ and for every $\lambda > 0$,*

$$t_{\pi_d}^{S_d^\star}(\gamma) \leqslant \frac{8n^2}{\lambda} ,$$

where $\gamma = \frac{2}{2d+1} + \lambda$.

Proof. By Theorem 5.2, we have that π_d is $(\lambda/2, 8n^2/\lambda)$-metastable for every $\lambda > 0$ and sufficiently large n. Therefore, for every $\mathbf{x} \in S_d^\star$ and $t \geqslant 8n^2/\lambda$, we have

$$\begin{aligned}
\left\|P^t(\mathbf{x},\cdot) - \pi_d\right\|_{\mathrm{TV}} &\leqslant \tfrac{\lambda}{2} + \max_{\mathbf{y}\in S_d^\star} \left\|P^t(\mathbf{x},\cdot) - P^t(\mathbf{y},\cdot)\right\|_{\mathrm{TV}} && \text{(by Lemma 3.7)}\\
&\leqslant \tfrac{\lambda}{2} + \max_{\mathbf{z}\in S_d^\star} \mathbf{P}_{\mathbf{z}}\left(\tau_{\mathbf{0}} \geqslant t\right) && \text{(by Lemma 5.3)}\\
&\leqslant \tfrac{2}{2d+1} + \lambda\,. && \text{(by Lemma 5.4 and Corollary 5.6)}
\end{aligned}$$

□

Starting from outside $S_d^\star$. Observe that when $\mathbf{x} = \mathbf{1}$, metastable distribution $\pi_{\mathbf{1}}$ is trivially reached immediately. Thus, it only remains to analyze $\mathbf{x} \notin S_d^\star \cup \{\mathbf{1}\}$. For this, it is enough to prove that for such an $\mathbf{x}$ the hitting time of $S_d^\star \cup \{\mathbf{1}\}$ is polynomial.

Lemma 5.8 *For $\beta = \omega(\log n)$, for every $d > 0$ and for every $\mathbf{x} \in S_d$, we have*

$$\mathbf{P}_{\mathbf{x}}\left(\tau_{\mathbf{1}} \leqslant n^2\right) \geqslant \frac{1}{3^d}\left(1 - \mathsf{negl}\left(n\right)\right) .$$

We can now state the following theorem.

Theorem 5.9 *For $\beta = \omega(\log n)$, for every $d > z > 0$ and for every profile $\mathbf{x} \in S_z$, we have*

$$\mathbf{P}_{\mathbf{x}}\left(\tau_{S_d^\star \cup \{\mathbf{1}\}} \leqslant n^2\right) \geqslant \left(\frac{2z}{2z+1} + \frac{1}{3^z}\right)\left(1 - \mathsf{negl}\left(n\right)\right).$$

The discussion above concludes the proof of the Theorem 5.1.

Staying arbitrarly close. We observe that, in Theorem 5.7, the distance between the dynamics and the metastable distribution cannot be made arbitrarily small. We can achieve this at the cost of slightly reducing the set of starting states from which convergence is proved. Specifically, the next theorem shows that, for $d = \omega(1)$ and arbitrarily small $\gamma > 0$, the logit dynamics starting from $S_d^\star$ is within distance γ from $\pi_{\mathbf{0}}$ in a number of steps that is polynomial in n and in $1/\gamma$.

Theorem 5.10 *For $\beta = \omega(\log n)$, $d = \omega(1)$ and every $\gamma > 0$,*

$$t_{\pi_{\mathbf{0}}}^{S_d^\star}(\gamma) \leqslant \frac{8n^2}{\gamma}.$$

Proof. Since $\beta = \omega(\log n)$, Theorem 5.2 implies that $\pi_{\mathbf{0}}$ is $(\gamma/2, 8n^2/\gamma)$-metastable for every $\gamma > 0$ and sufficiently large n. Therefore, for every $\mathbf{x} \in S_d^\star$ and $t \geqslant 8n^2/\gamma$, we have

$$\begin{aligned}
\|P^t(\mathbf{x},\cdot) - \pi_{\mathbf{0}}\|_{\mathrm{TV}} &\leqslant \tfrac{\gamma}{2} + \mathbf{P}_{\mathbf{x}}\left(\tau_{\mathbf{0}} > t\right) && \text{(by Lemma 3.8)}\\
&\leqslant \tfrac{1}{2d+1} + \tfrac{3}{4}\gamma && \text{(by Lemma 5.4 and Corollary 5.6)}\\
&\leqslant \gamma\,. && \text{(since } d = \omega(1))
\end{aligned}$$

□

5.2 Games without risk dominant strategies

In this section we study the case of graphical coordination games without risk dominant strategies (that is, $\Delta = \delta$) played on a ring by n players. Next theorem identifies a class of metastable distributions.

Theorem 5.11 *For every $\varepsilon > 0$ and for every $0 \leqslant d \leqslant n$, distribution $\mu_d = \frac{d}{n}\pi_{\mathbf{0}} + \left(1 - \frac{d}{n}\right)\pi_{\mathbf{1}}$ is $(\varepsilon, \varepsilon e^{2\Delta\beta})$-metastable.*

Proof. We notice that

$$\begin{aligned}
\|\mu_d P - \mu_d\|_{\mathrm{TV}} &= \frac{1}{2}\sum_{\mathbf{x}} \left|(\mu_d P)(\mathbf{x}) - \mu_d(\mathbf{x})\right| = \frac{1}{2}\sum_{\mathbf{x}} \left|\sum_{\mathbf{y}} \mu_d(\mathbf{y})P(\mathbf{y},\mathbf{x}) - \mu_d(\mathbf{x})\right|\\
&= d\sum_{\mathbf{x}\in\Omega_{n-1}} P(\mathbf{0},\mathbf{x}) + (1-d)\sum_{\mathbf{x}\in\Omega_1} P(\mathbf{1},\mathbf{x}) = \frac{1}{1+e^{2\Delta\beta}}\,.
\end{aligned}$$

Thus μ is $\left(\frac{1}{1+e^{2\Delta\beta}}, 1\right)$-metastable. The Theorem follows from Lemma 3.2. $\square$

The main and quite surprising result in this section is that for every starting profile $\mathbf{x} \in S_d$ the dynamics starting in $\mathbf{x}$ converges in polynomial time to μ_d, for $d = 1, \ldots, n$. In order to prove this result, we define $\tau_{\mathbf{0},\mathbf{1}} = \min\{\tau_{\mathbf{0}}, \tau_{\mathbf{1}}\}$ and prove that this quantity is polynomial in n with very high probability; then we show that the dynamics starting at $\mathbf{x} \in S_d$ after $\tau_{\mathbf{0},\mathbf{1}}$ steps is distributed as a metastable distribution very close to μ_d. We formalize these arguments in two technical lemmas.

Lemma 5.12 *If $\beta = \omega(\log n)$, then for every $\mathbf{x} \in \Omega$, $\mathbf{P}_{\mathbf{x}}\left(\tau_{\mathbf{0},\mathbf{1}} \leqslant n^5\right) \geqslant 1 - o(1)$.*

Lemma 5.13 *For every d, $\mathbf{x} \in S_d$ and $\beta = \omega(\log n)$, the random variable $X_{\tau_{\mathbf{0},\mathbf{1}}}$ given that $X_0 = \mathbf{x}$, has distribution $\nu_x = \left(\frac{d}{n} + \lambda_{\mathbf{x}}\right)\pi_{\mathbf{0}} + \left(1 - \frac{d}{n} - \lambda_{\mathbf{x}}\right)\pi_{\mathbf{1}}$, with $|\lambda_{\mathbf{x}}| = o(1)$.*

The pseudo mixing time of distributions μ_d, for $d = 0, 1, \ldots, n$, is given by the next Theorem.

Theorem 5.14 *If $\beta = \omega(\log n)$, for every d and every $\gamma > 0$*

$$t_{\mu_d}^{S_d}(\gamma) \leqslant n^5.$$

Proof. From Lemma 5.12, for every $\mathbf{x} \in \Omega$ we have

$$\begin{aligned}
\left\| P^{n^5}(\mathbf{x}, \cdot) - \mu_d \right\|_{\mathrm{TV}} &= \max_{A \subset \Omega} \left| \mathbf{P}_{\mathbf{x}}\left(X_{n^5} \in A\right) - \mu_d(A) \right| \\
&= o(1) + \max_{A \subset \Omega} \left| \mathbf{P}_{\mathbf{x}}\left(X_{n^5} \in A \mid \tau_{\mathbf{0},\mathbf{1}} \leqslant n^5\right) - \mu_d(A) \right| \\
&= o(1) + \left\| \mathbf{P}_{\mathbf{x}}\left(X_{n^5} \mid \tau_{\mathbf{0},\mathbf{1}} \leqslant n^5\right) - \mu_d \right\|_{\mathrm{TV}} \\
&\leqslant o(1) + \left\| \mathbf{P}_{\mathbf{x}}\left(X_{n^5} \mid \tau_{\mathbf{0},\mathbf{1}} \leqslant n^5\right) - \nu_{\mathbf{x}} \right\|_{\mathrm{TV}} + \left\| \nu_{\mathbf{x}} - \mu_d \right\|_{\mathrm{TV}},
\end{aligned}$$

where the last inequality follows from the triangle inequality of the total variation distance. Moreover, from Lemma 5.13, for every $\mathbf{x} \in S_d$, we have

$$\left\| \nu_{\mathbf{x}} - \mu_d \right\|_{\mathrm{TV}} = \left\| P^{\tau_{\mathbf{0},\mathbf{1}}}(\mathbf{x}, \cdot) - \mu_d \right\|_{\mathrm{TV}} = o(1).$$

Finally, from Theorem 5.11 we have that $\mu_{\mathbf{x}}$ is $(o(1), n^5)$-metastable and thus,

$$\begin{aligned}
\left\| \mathbf{P}_{\mathbf{x}}\left(X_{n^5} \mid \tau_{\mathbf{0},\mathbf{1}} \leqslant n^5\right) - \nu_{\mathbf{x}} \right\|_{\mathrm{TV}} &= \left\| P^{\tau_{\mathbf{0},\mathbf{1}}}(\mathbf{x}, \cdot) P^{n^5 - \tau_{\mathbf{0},\mathbf{1}}} - \nu_{\mathbf{x}} \right\|_{\mathrm{TV}} \\
&= \left\| \nu_{\mathbf{x}} P^{n^5 - \tau_{\mathbf{0},\mathbf{1}}} - \nu_{\mathbf{x}} \right\|_{\mathrm{TV}} = o(1).
\end{aligned}$$

$\square$

5.3 Proofs from Section 5.1.2

5.3.1 Proof of Lemma 5.3

Proof of Lemma 5.3. Consider the following partial order $\preceq$ over Ω: for profiles $\mathbf{x}, \mathbf{y} \in \Omega$, $\mathbf{x} \preceq \mathbf{y}$ iff for every $0 \leqslant i \leqslant n-1$, we have that $x_i \geqslant y_i$. That is, if $\mathbf{x} \preceq \mathbf{y}$ then $\mathbf{x}$ can be obtained from $\mathbf{y}$ by changing 0's into 1's. We note that, in according to this order, $\mathbf{0}$ is the unique maximum.

We next show that for every two profiles $\mathbf{x}, \mathbf{y} \in \Omega$ there exists a monotone (w.r.t. $(\Omega, \preceq)$) coupling (X_1, Y_1) of two copies of the logit dynamics for the graphical coordination game on the ring for which $X_0 = \mathbf{x}$ and $Y_0 = \mathbf{y}$. The Lemma then follows from Theorem A.1 and Lemma A.2.

The coupling proceeds as follows: first, pick a player i uniformly at random; then, update the strategies x_i and y_i of player i in the two chains, by setting

$$(x_i, y_i) = \begin{cases}
(0,0), & \text{with probability } \min\{\sigma_i(0 \mid \mathbf{x}), \sigma_i(0 \mid \mathbf{y})\}; \\
(1,1), & \text{with probability } \min\{\sigma_i(1 \mid \mathbf{x}), \sigma_i(1 \mid \mathbf{y})\}; \\
(0,1), & \text{with probability } \sigma_i(0 \mid \mathbf{x}) - \min\{\sigma_i(0 \mid \mathbf{x}), \sigma_i(0 \mid \mathbf{y})\}; \\
(1,0), & \text{with probability } \sigma_i(1 \mid \mathbf{x}) - \min\{\sigma_i(1 \mid \mathbf{x}), \sigma_i(1 \mid \mathbf{y})\}.
\end{cases}$$

See Equation 1, for the definition of σ_i. Next we observe that if i is chosen for update, then the marginal distributions of x_i and y_i agree with $\sigma_i(\cdot \mid \mathbf{x})$ and $\sigma_i(\cdot \mid \mathbf{y})$, respectively. Indeed, for $b \in \{0,1\}$, the probability that $x_i = b$ is

$$\min\{\sigma_i(b \mid \mathbf{x}), \sigma_i(b \mid \mathbf{y})\} + \sigma_i(b \mid \mathbf{x}) - \min\{\sigma_i(b \mid \mathbf{x}), \sigma_i(b \mid \mathbf{y})\} = \sigma_i(b \mid \mathbf{x}),$$

and the probability that $y_i = b$ is

$$\begin{aligned}&\min\{\sigma_i(b \mid \mathbf{x}), \sigma_i(b \mid \mathbf{y})\} + \sigma_i(1-b \mid \mathbf{x}) - \min\{\sigma_i(1-b \mid \mathbf{x}), \sigma_i(1-b \mid \mathbf{y})\} \\ &= \min\{\sigma_i(b \mid \mathbf{x}), \sigma_i(b \mid \mathbf{y}) \ + (1-\sigma_i(b \mid \mathbf{x})) - (1-\max\{\sigma_i(b \mid \mathbf{x}), \sigma_i(b \mid \mathbf{y})\}) = \sigma_i(b \mid \mathbf{y}).\end{aligned}$$

The coupling described above is monotone w.r.t. $(\Omega, \preceq)$. Indeed, suppose $\mathbf{x} \preceq \mathbf{y}$ and that the player i was selected for update. Since $\mathbf{x} \preceq \mathbf{y}$, the number of neighbors of i playing 0 in $\mathbf{x}$ is less or equal than in $\mathbf{y}$ and thus $\sigma_i(0 \mid \mathbf{x}) \leqslant \sigma_i(0 \mid \mathbf{y})$. Thus, the coupling either sets $x_i = y_i$ or $x_i = 1$ and $y_i = 0$. In both cases, $X_1 \preceq Y_1$. □

5.3.2 Proof of Lemma 5.4

Lemma 5.4 gives an upper bound on the hitting time, $\tau_{\mathbf{0}}$, of $\mathbf{0}$, for a dynamics starting from a profile $\mathbf{x} \in R$ (profiles in R are those in which at least two adjacent players play 0). For convenience, we rename players so that $x_0 = x_1 = 0$. 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, plays 0 with high probability. Similarly, after player 2 has played 0, 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 0 with high probability. This process repeats until every player is playing 0. In the following, we estimate the number of steps sufficient to have all players playing strategy 0 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 0 when selected for update, given that at least one of their neighbors is playing 0.

Observation 5.15 *For every player i, if i is selected for update at time t, then, for $b \in \{-1, 1\}$*

$$\mathbf{P}\left(X_t^i = 0 \mid X_t^{i+b} = 0\right) \geqslant \left(1 - \frac{1}{1+e^{(\Delta-\delta)\beta}}\right).$$

We start by evaluating the probability that the dynamics selects players $2, \ldots, n-1$ at least once in this order before time t. To this aim, we set $\rho_1 = 0$ and, for $i = 2, \ldots, n-1$, we define ρ_i as 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 $2, \ldots, i-1$ have been selected at least once in this order. In particular, ρ_{n-1} is the first time step at which every player i, $i \geqslant 3$, has been selected at least once after his left neighbor. Obviously, $\rho_i > \rho_{i-1}$ for $i = 2, \ldots, n-1$. The next lemma lower bounds the probability that $\rho_{n-1} \leqslant t$.

Lemma 5.16 *For every $\mathbf{x} \in R$ and every $t > 0$, we have*

$$\mathbf{P}_{\mathbf{x}}\left(\rho_{n-1} \leqslant t\right) \geqslant 1 - \frac{n^2}{t}.$$

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 $i = 3, \ldots, n-1$. Thus, by linearity of expectation,

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

The lemma follows from the Markov inequality. □

Suppose now that $t \geqslant \rho_{n-1}$. The next lemma shows that, for all players i, the probability that $X_t^i = 0$ is high.

Lemma 5.17 *For every starting profile $\mathbf{x} \in R$, for every player i and for every time step $t > 0$, we have*

$$\mathbf{P}_{\mathbf{x}}\left(X_t^i = 0 \mid \rho_{n-1} \leqslant t\right) \geqslant \left(1 - \frac{1}{1+e^{(\Delta-\delta)\beta}}\right)^t.$$

We prove the lemma first for $i \geqslant 2$. Then we deal with players 0 and 1.

Fix player $i \geqslant 2$, 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 $s_i > s_{i-1} > \ldots > s_2 > 0$ such that, for $j = i, i-1, \ldots, 2$, s_j is the last time player j has been selected before time s_{j+1}. We remark that, since $t \geqslant \rho_{n-1} > \rho_i$ we have that players $2, \ldots, i$ are selected at least once in this order and thus all the s_j^i are well defined. Strictly speaking, the sequence $s_i, \ldots, s_2$ 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 $X_t^i = 0$ for $i \geqslant 2$, we first bound it in terms of the the probability that player 2 plays 0 at time s_2 and then we evaluate this last quantity. The next lemma is the first step.

Lemma 5.18 *For every $\mathbf{x} \in R$, every player $2 \leqslant i \leqslant n-1$ and every time step t, we have*

$$\mathbf{P}_{\mathbf{x}}\left(X_t^i = 0 \mid \rho_{n-1} \leqslant t\right) \geqslant \left(1 - \frac{1}{1+e^{(\Delta-\delta)\beta}}\right)^{i-2} \mathbf{P}_{\mathbf{x}}\left(X_{s_2}^2 = 0 \mid \rho_{n-1} \leqslant t\right).$$

Proof. For every fixed 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 $j = 2, \ldots, i$, we observe that, since $t \geqslant \rho_{n-1} > \rho_i$, player j has been selected for update at time s_j and s_j is the last time that player j is selected for update before time s_{j+1} and thus $X_{s_{j+1}}^j = X_{s_j}^j$.

From Observation 5.15, we have

$$\begin{aligned}\mathbf{P}_{\mathbf{x}}\left(X_{s_j}^j = 0 \mid \rho_{n-1} \leqslant t\right) &\geqslant \mathbf{P}_{\mathbf{x}}\left(X_{s_j}^j = 0 \mid X_{s_j}^{j-1} = 0, \rho_{n-1} \leqslant t\right) \cdot \mathbf{P}_{\mathbf{x}}\left(X_{s_j}^{j-1} = 0 \mid \rho_{n-1} \leqslant t\right) \\ &\geqslant \left(1 - \frac{1}{1+e^{(\Delta-\delta)\beta}}\right) \mathbf{P}_{\mathbf{x}}\left(X_{s_{j-1}}^{j-1} = 0 \mid \rho_{n-1} \leqslant t\right),\end{aligned}$$

and the lemma follows. □

We now bound the probability that player 2 plays 0 at time step s_2. If player 1 has not been selected for update before time s_2, then $X_{s_2}^1 = X_0^1 = 0$, and, from Observation 5.15, we have

$$\begin{aligned}\mathbf{P}_{\mathbf{x}}\left(X_{s_2}^2 = 0 \mid \rho_{n-1} \leqslant t\right) &\geqslant \mathbf{P}_{\mathbf{x}}\left(X_{s_2}^2 = 0 \mid X_{s_2}^1 = 0, \rho_{n-1} \leqslant t\right) \\ &\geqslant \left(1 - \frac{1}{1+e^{(\Delta-\delta)\beta}}\right).\end{aligned}$$

It remains to consider the case when player 1 has been selected for update at least once before time s_2. For fixed player i and time step t we define a new sequence of time steps $r_0 > r_1, \ldots > 0$ in the following way. We set $r_0 = s_2$, and, starting from time step s_2 and going backward to time step 0, r_j, for $j > 0$, is 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 r_1 is the last time player 1 has been selected for update before $r_0 = s_2$, we have $X_{s_2}^1 = X_{r_1}^1$ and, by Observation 5.15,

$$\begin{aligned}\mathbf{P}_{\mathbf{x}}\left(X_{s_2}^2 = 0 \mid \rho_{n-1} \leqslant t\right) &\geqslant \mathbf{P}_{\mathbf{x}}\left(X_{s_2}^2 = 0 \mid X_{s_2}^1 = 0, \rho_{n-1} \leqslant t\right) \cdot \mathbf{P}_{\mathbf{x}}\left(X_{s_2}^1 = 0 \mid \rho_{n-1} \leqslant t\right) \\ &\geqslant \left(1 - \frac{1}{1+e^{(\Delta-\delta)\beta}}\right) \cdot \mathbf{P}_{\mathbf{x}}\left(X_{r_1}^1 = 0 \mid \rho_{n-1} \leqslant t\right).\end{aligned} \tag{9}$$

Finally, we bound $\mathbf{P}_{\mathbf{x}}\left(X_{r_1}^1 = 0 \mid \rho_{n-1} \leqslant t\right)$.

Lemma 5.19 *For every starting profile $\mathbf{x} \in R$ and for every time step t, for every fixed player i, let $r_0, \ldots, r_k$ be defined as above. If $k > 0$, we have*

$$\mathbf{P}_{\mathbf{x}}\left(X_{r_1}^1 = 0 \mid \rho_{n-1} \leqslant t\right) \geqslant \left(1 - \frac{1}{1+e^{(\Delta-\delta)\beta}}\right)^k.$$

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

$$\mathbf{P_x}\left(X_{r_j}^{\mathcal{P}(j)} = 0 \mid \rho_{n-1} \leqslant t\right) \geqslant \mathbf{P_x}\left(X_{r_j}^{\mathcal{P}(j)} = 0 \mid X_{r_j}^{\mathcal{P}(j+1)} = 0, \rho_{n-1} \leqslant t\right) \cdot \mathbf{P_x}\left(X_{r_j}^{\mathcal{P}(j+1)} = 0 \mid \rho_{n-1} \leqslant t\right).$$

If $j \neq k$ player $\mathcal{P}(j+1)$ has not been selected for update between time r_{j+1} and time r_j and by Observation 5.15

$$\begin{aligned}\mathbf{P_x}\left(X_{r_j}^{\mathcal{P}(j)} = 0 \mid \rho_{n-1} \leqslant t\right) &\geqslant \mathbf{P_x}\left(X_{r_j}^{\mathcal{P}(j)} = 0 \mid X_{r_j}^{\mathcal{P}(j+1)} = 0, \rho_{n-1} \leqslant t\right) \cdot \mathbf{P_x}\left(X_{r_{j+1}}^{\mathcal{P}(j+1)} = 0 \mid \rho_{n-1} \leqslant t\right) \\ &\geqslant \left(1 - \frac{1}{1+e^{(\Delta-\delta)\beta}}\right) \mathbf{P_x}\left(X_{r_{j+1}}^{\mathcal{P}(j+1)} = 0 \mid \rho_{n-1} \leqslant t\right).\end{aligned}$$

If $j = k$, instead, player $\mathcal{P}(k+1)$ has not been selected for update before time r_k and thus $X_{r_k}^{\mathcal{P}(k+1)} = X_0^{\mathcal{P}(k+1)} = 0$. By Observation 5.15, we have

$$\begin{aligned}\mathbf{P_x}\left(X_{r_k}^{\mathcal{P}(k)} = 0 \mid \rho_{n-1} \leqslant t\right) &\geqslant \mathbf{P_x}\left(X_{r_k}^{\mathcal{P}(k)} = 0 \mid X_{r_k}^{\mathcal{P}(k+1)} = 0, \rho_{n-1} \leqslant t\right) \\ &\geqslant \left(1 - \frac{1}{1+e^{(\Delta-\delta)\beta}}\right).\end{aligned}$$

□

Thus, for every player $i \geqslant 2$ and for every time step $t > 0$, we have

$$\begin{aligned}\mathbf{P_x}\left(X_t^i = 0 \mid \rho_{n-1} \leqslant t\right) &\geqslant \left(1 - \tfrac{1}{1+e^{(\Delta-\delta)\beta}}\right)^{i-2} \mathbf{P_x}\left(X_{s_2}^2 = 0 \mid \rho_{n-1} \leqslant t\right) && \text{(from Lemma 5.18)} \\ &\geqslant \left(1 - \tfrac{1}{1+e^{(\Delta-\delta)\beta}}\right)^{i-1} \mathbf{P_x}\left(X_{r_1}^1 = 0 \mid \rho_{n-1} \leqslant t\right) && \text{(from Equation 9)} \\ &\geqslant \left(1 - \tfrac{1}{1+e^{(\Delta-\delta)\beta}}\right)^{i-1+k} && \text{(from Equation 5.19)} \\ &\geqslant \left(1 - \tfrac{1}{1+e^{(\Delta-\delta)\beta}}\right)^{t}, && \text{(since } i-1+k \leqslant t\text{)}\end{aligned}$$

where k is the index of the last term in the sequence $r_0, r_1, \ldots$ previously defined.

This ends the proof of Lemma 5.17 for player $i \geqslant 2$. The cases $i = 0, 1$ can be proved in similar way. Clearly, if player i has never been selected for update before time t, we have that $X_t^i = 0$ 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 $r_1 > r_2 > \ldots > 0$ as above: we have that $X_t^1 = X_{r_1}^1$ and from Lemma 5.19 follows that

$$\mathbf{P_x}\left(X_t^i = 0 \mid \rho_{n-1} \leqslant t\right) \geqslant \left(1 - \frac{1}{1+e^{(\Delta-\delta)\beta}}\right)^k \geqslant \left(1 - \frac{1}{1+e^{(\Delta-\delta)\beta}}\right)^t,$$

where k is the last index of the sequence $r_1, r_2, \ldots$. Finally, the probability that player 0 plays the strategy 0 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 0 at time s_2. This concludes the proof of Lemma 5.17.

The following lemma gives the probability that the hitting time of the profile $\mathbf{0}$ is less or equal to t, given that $\rho_{n-1} \leqslant t$.

Lemma 5.20 *For every $\mathbf{x} \in R$ and every $t > 0$, we have*

$$\mathbf{P_x}\left(\tau_{\mathbf{0}} \leqslant t \mid \rho_{n-1} \leqslant t\right) \geqslant \left(1 - \frac{1}{1+e^{(\Delta-\delta)\beta}}\right)^{nt}.$$

Proof. To prove our lemma we will show a bound on the probability that, conditioned on $\rho_{n-1} \leqslant t$, all players are playing 0 at time t.

Let f be the permutation that sort players in 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

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

where the last inequality follows from Lemma 5.17. □

Now we are ready to prove Lemma 5.4.

Proof of Lemma 5.4. From Lemma 5.16 and Lemma 5.20, we have that for every $\mathbf{x} \in R$ and every $t > 0$

$$\begin{aligned}
\mathbf{P}_{\mathbf{x}}\left(\tau_{\mathbf{0}} \leqslant t\right) &\geqslant \mathbf{P}_{\mathbf{x}}\left(\rho_{n-1} \leqslant t\right) \mathbf{P}_{\mathbf{x}}\left(\tau_{\mathbf{0}} \leqslant t \mid \rho_{n-1} \leqslant t\right) \\
&\geqslant \left(1 - \frac{n^2}{t}\right)\left(1 - \frac{1}{1 + e^{(\Delta-\delta)\beta}}\right)^{nt}.
\end{aligned}$$

Thus for every $\lambda > 0$, we have for $t = \frac{(8-\lambda)n^2}{\lambda}$

$$\begin{aligned}
\mathbf{P}_{\mathbf{x}}\left(\tau_{\mathbf{0}} \leqslant t\right) &\geqslant \left(1 - \frac{\lambda}{8-\lambda}\right)\left(1 - \frac{1}{1 + \frac{(8-\lambda)n^3}{\lambda \log\left(\frac{8}{8-\lambda}\right)}}\right)^{\frac{(8-\lambda)n^3}{\lambda}} \\
&\geqslant \frac{8\left(1 - \frac{\lambda}{4}\right)}{8-\lambda} \frac{8-\lambda}{8} = 1 - \frac{\lambda}{4},
\end{aligned}$$

where the first inequality follows from the fact that, since $\beta = \omega(\log n)$, for n large enough, we have $\beta \geqslant \frac{\log\left(\frac{(8-\lambda)n^3}{\lambda}\right) - \log\log\left(\frac{8}{8-\lambda}\right)}{\Delta-\delta}$, whereas the second inequality follows from the well known approximation $1 - a \geqslant e^{-\frac{a}{1-a}}$. □

5.3.3 Proof of Lemma 5.5

Let $\theta^\star$ be the first time at which all players have been selected at least once. The following lemma directly follows from coupon collector argument; we include a proof for completeness.

Lemma 5.21 *For every $t > 0$,*

$$\mathbf{P}_{\mathbf{x}}\left(\theta^\star \leqslant t\right) \geqslant 1 - ne^{-t/n}.$$

Proof. The logit dynamics at each time step selects 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 $\left(1 - \frac{i}{n}\right)^t$ and

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

□

We define players playing 0 in profile $\mathbf{x}$ as the *zero-players* of $\mathbf{x}$ and their neighbors as *border-players*; we also define $l(\mathbf{x}) \geqslant d+1$ as the number of border-players in $\mathbf{x}$.

Let $\tau^\star$ be the first time step at which a border-player is selected for update before one of its neighboring zero-players; if this event does not occur then $\tau^\star = +\infty$. The next lemma bounds the probability that $\tau^\star$ is finite given that all players have been selected at least once within time t.

Lemma 5.22 *For every* $\mathbf{x} \in S_d \setminus R$ *and every* $t > 0$

$$\mathbf{P_x}\left(\tau^\star \leqslant t \mid \theta^\star \leqslant t\right) \geqslant \frac{2d}{2d+1}.$$

Proof. Observe that if $\theta^\star \leqslant t$, then $\tau^\star > t$ is equivalent to say that $\tau^\star$ is infinite: thus we will consider $\mathbf{P_x}\left(\tau^\star \text{ is finite} \mid \theta^\star \leqslant t\right)$.

The proof proceeds by induction on d. Let $d = 1$ and denote by i the one zero-player. Notice that $\tau^\star$ is finite if and only if one of the two neighbors of i is selected for update before i is selected. Since we are conditioning on $\theta^\star \leqslant t$, all players are selected at least once by time t and thus the probability of this event is $\frac{2}{3} = \frac{2d}{2d+1}$.

Suppose now that the claim holds for $d-1$ and let $\mathbf{x} \in S_d \setminus R$. Denote by $T_\mathbf{x}$ the set of all the zero-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 $\theta^\star \leqslant t$). Observe that, if i is a border-player, then $\tau^\star$ is finite and this happens with probability $\frac{l(\mathbf{x})}{l(\mathbf{x})+d}$. If i is a zero-player, we consider the subset $\overline{T}_\mathbf{x} \subset T_\mathbf{x}$ of the remaining $d-1$ zero-players and their border-players. $\tau^\star$ is finite only if at least one border-player in $\overline{T}_\mathbf{x}$ is selected before one of its neighboring zero-players. Notice though that $\overline{T}_\mathbf{x} = T_\mathbf{y}$, 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, $\mathbf{P_y}\left(\tau^\star \text{ is finite} \mid \theta^\star \leqslant t\right) \geqslant \frac{2d-2}{2j-1}$. Finally,

$$\begin{aligned}
\mathbf{P_x}\left(\tau^\star \text{ is finite} \mid \theta^\star \leqslant t\right) &= \frac{l(\mathbf{x})}{l(\mathbf{x})+d} + \frac{d}{l(\mathbf{x})+d} \cdot \mathbf{P_y}\left(\tau^\star \text{ is finite} \mid \theta^\star \leqslant t\right) \\
&\geqslant \frac{l(\mathbf{x})}{l(\mathbf{x})+d} + \frac{d}{l(\mathbf{x})+d} \cdot \frac{2d-2}{2d-1} \\
&= 1 - \frac{d}{(l(\mathbf{x})+d)(2d-1)} \\
&\geqslant 1 - \frac{1}{2d+1},
\end{aligned}$$

where the last inequality follows from $l(\mathbf{x}) \geqslant d+1$. □

Now we are ready to prove Lemma 5.5.

Proof of Lemma 5.5. Suppose $\tau^\star$ is finite and let i be the border-player selected for update at time $\tau^\star$. Then, at time $\tau^\star$, i has at least one neighbor playing 0 and thus i plays 0 with probability

$$\mathbf{P_x}\left(X^i_{\tau^\star} = 0 \mid \tau^\star \leqslant t\right) \geqslant \left(1 - \frac{1}{1+e^{(\Delta-\delta)\beta}}\right).$$

Moreover, if i plays strategy 0, then at time $\tau^\star$ the dynamics hits a profile in R. Thus, for a finite $t > 0$, we have

$$\begin{aligned}
\mathbf{P_x}\left(\tau_R \leqslant t\right) &\geqslant \mathbf{P_x}\left(X^i_{\tau^\star} = 0 \wedge \tau^\star \leqslant t\right) \\
&= \mathbf{P_x}\left(X^i_{\tau^\star} = 0 \mid \tau^\star \leqslant t\right) \mathbf{P_x}\left(\tau^\star \leqslant t\right) \\
&\geqslant \mathbf{P_x}\left(X^i_{\tau^\star} = 0 \mid \tau^\star \leqslant t\right) \mathbf{P_x}\left(\tau^\star \leqslant t \mid \theta^\star \leqslant t\right) \mathbf{P_x}\left(\theta^\star \leqslant t\right) \\
&\geqslant \left(1 - \frac{1}{1+e^{(\Delta-\delta)\beta}}\right) \frac{2d}{2d+1} \left(1 - ne^{-t/n}\right),
\end{aligned}$$

where the last inequality follows from Lemma 5.21 and Lemma 5.22. Finally, the lemma follows since $\beta = \omega(\log n)$ and by taking $t = n^2$. □

5.3.4 Proof of Lemma 5.8

This proof is very similar to the proof of Lemma 5.5: in particular, we refer to notation defined in Section 5.3.3. If all zero-players are selected before both neighboring border-players, we set $\tau^\star$ be the time step at which the last zero-players is selected, otherwise we set $\tau^\star$ be infinity.

Lemma 5.23 *For every* $\mathbf{x} \in S_d \setminus R$

$$\mathbf{P_x}\left(\tau^\star \leqslant t \mid \theta^\star \leqslant t\right) \geqslant \frac{1}{3^d}.$$

Proof. Observe that if $\theta^\star \leqslant t$, then $\tau^\star > t$ is equivalent to say that $\tau^\star$ is infinite: thus we will consider $\mathbf{P_x}\left(\tau^\star \text{ is finite} \mid \theta^\star \leqslant t\right)$.

The proof proceeds by induction on d. For the base case, we denote by i the only zero-player and $\tau^\star$ is finite if and only if i is selected for update before her neighbors. Since we are conditioning on $\theta^\star \leqslant t$, we know that all players have been selected at least once and thus the probability of this event is $\frac{1}{3}$.

Suppose that the claim holds for $d-1$ and let $\mathbf{x} \in S_d \setminus R$. Denote by $T_\mathbf{x}$ the set of all the zero-players in $\mathbf{x}$ along with 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 $\theta^\star \leqslant t$ and thus all players has been selected at least once). Observe that, if i is a border-player, then $\tau^\star$ is *infinite* and this happens with probability $\frac{l(\mathbf{x})}{l(\mathbf{x})+j}$. Otherwise, we consider the subset $\overline{T}_\mathbf{x} \subset T_\mathbf{x}$ of the remaining $d-1$ zero-players and their border-players. $\tau^\star$ will be finite only if all zero-players in $\overline{T}_\mathbf{x}$ are selected before their border-players. However, $\overline{T}_\mathbf{x} = T_\mathbf{y}$, where $\mathbf{y} \in S_{d-1} \setminus R$ is the profile obtained from $\mathbf{x}$ by setting $y_i = 1$. By inductive hypothesis the probability $\mathbf{P_y}\left(\tau^\star \text{ is finite} \mid \theta^\star \leqslant t\right) \geqslant \frac{1}{3^{d-1}}$. Thus, we have

$$\begin{aligned}
\mathbf{P_x}\left(\tau^\star \text{ is finite} \mid \theta^\star \leqslant t\right) &= \frac{d}{l(\mathbf{x})+d} \cdot \mathbf{P_y}\left(\tau^\star \text{ is finite} \mid \theta^\star \leqslant t\right) \\
&\geqslant \frac{1}{3} \cdot \frac{1}{3^{d-1}} \,.
\end{aligned}$$

□

Now we prove Lemma 5.8.

Proof of Lemma 5.8. Notice that if $\tau^\star$ is finite, every time a player is selected for update she have both neighbors that are playing 0 and thus

$$\mathbf{P_x}\left(X_{\tau^\star} = \mathbf{1} \mid \tau^\star \leqslant t\right) \geqslant \left(1 - \frac{1}{1+e^{2\delta\beta}}\right)^t \geqslant \left(1 - \frac{t}{e^{2\delta\beta}}\right),$$

where the last inequality follows from the approximations $1 - a \leqslant e^{-x}$ and $1 - a \geqslant e^{-\frac{x}{1-x}}$.

Obviously, if $X_{\tau^\star} = \mathbf{1}$, then $\tau_\mathbf{1} \leqslant \tau^\star$. Thus, for a finite $t > 0$, we have

$$\begin{aligned}
\mathbf{P_x}\left(\tau_\mathbf{1} \leqslant t\right) &\geqslant \mathbf{P_x}\left(X_{\tau^\star} = \mathbf{1} \wedge \tau^\star \leqslant t\right) \\
&= \mathbf{P_x}\left(X_{\tau^\star} = \mathbf{1} \mid \tau^\star \leqslant t\right) \mathbf{P_x}\left(\tau^\star \leqslant t\right) \\
&\geqslant \mathbf{P_x}\left(X_{\tau^\star} = \mathbf{1} \mid \tau^\star \leqslant t\right) \mathbf{P_x}\left(\tau^\star \leqslant t \mid \theta^\star \leqslant t\right) \mathbf{P_x}\left(\theta^\star \leqslant t\right) \\
&\geqslant \left(1 - \frac{t}{e^{2\delta\beta}}\right) \frac{1}{3^z} \left(1 - ne^{-t/n}\right),
\end{aligned}$$

where the last inequality follows from Lemma 5.21 and Lemma 5.23. Finally, the lemma follows since $\beta = \omega(\log n)$ and by taking $t = n^2$. □

5.4 Proofs from Section 5.2

We say that a profile $\mathbf{x}$ has a *zero-block* of size l starting at player i if $x_i = x_{i+1} = \ldots = x_{i+l-1} = 0$ and $x_{i-1} = x_{i+l} = 1$. Players i and $i+h-1$ are the *border* players of the block. A similar definition is given for *one-blocks*. Notice that every profile $\mathbf{x} \neq \mathbf{0}, \mathbf{1}$ has the same number of zero-blocks and one-blocks and this number is called the *level* of $\mathbf{x}$ and is denoted by $\ell(\mathbf{x})$. We set $\ell(\mathbf{0}) = \ell(\mathbf{1}) = 0$.

The following observation gives the *level structure* of the potential function (note that we are studying the case $\Delta = \delta$).

Observation 5.24 *For every profile* $\mathbf{x}$*, the potential of* $\mathbf{x}$ *is* $\Phi(\mathbf{x}) = (n - 2\ell(\mathbf{x}))\Delta$*, regardless of the sizes of the zero-blocks and one-blocks.*

Moreover, for a profile $\mathbf{x}$, we defines $s_0(\mathbf{x})$ as the number of zero-blocks of size 1, $s_1(\mathbf{x})$ as the number of one-blocks of size 1 and set $s(\mathbf{x}) = s_0(\mathbf{x}) + s_1(\mathbf{x})$.

5.4.1 Proof of Lemma 5.12

We would like to study how long it takes for the logit dynamics to reach $\mathbf{0}$ or $\mathbf{1}$. Starting from profile $\mathbf{x}$ at level $i \geqslant 1$, the logit dynamics needs to go down i levels to hit a profile at level 0; and to go down one level, it is necessary for one monochromatic block (that is, either a zero-block or a one-block) to disappear. We next show that we do not have to wait too long for this to happen.

Our first step bounds the time τ_i needed to go from level $i+1$ to level i. Consider a profile $\mathbf{x}$ at level $i+1$ and number arbitrarily the $2(i+1)$ monochromatic blocks of $\mathbf{x}$ and denote by $k_j(\mathbf{x})$ the size of the j-th monochromatic block. We define $\tau_{i,j}$ in the following way. Suppose that the dynamics reaches level i for the first time after t steps and suppose that this happens because the j-th monochromatic block disappears. Then we set $\tau_{i,j} = t$ and $\tau_{i,j'} = +\infty$ for all $j' \neq j$. Obviously, for any starting profile $\mathbf{x}$ we have that

$$\mathbf{E}_{\mathbf{x}}\left[\tau_i\right] = \mathbf{E}_{\mathbf{x}}\left[\min_j \tau_{i,j}\right] \leqslant \max_j \mathbf{E}_{\mathbf{x}}\left[\tau_{i,j} \mid \tau_{i,j} < \tau_{i,j'} \text{ for all } j' \neq j\right].$$

For sake of compactness of notation we define

$$\gamma_{i,l} = \max_{1 \leqslant j \leqslant 2(i+1)} \max_{\substack{\mathbf{x}\colon \ell(\mathbf{x})=i+1 \\ k_j(\mathbf{x})=l}} \mathbf{E}_{\mathbf{x}}\left[\tau_{i,j} \mid \tau_{i,j} < \tau_{i,j'} \text{ for all } j' \neq j\right],$$

set $\gamma_i = \max_l \gamma_{i,l}$ and observe that $\mathbf{E}_{\mathbf{x}}\left[\tau_i\right] \leqslant \gamma_i$. It is also easy to see that $\gamma_{i,l}$ is non-decreasing with l. Next we bound γ_i in terms of γ_{i+1}.

Lemma 5.25 *For every $i \geqslant 0$*

$$\gamma_i \leqslant n^2 b_i$$

where $b_i = n + \frac{n-4(i+1)+s(\mathbf{x}^\star)}{1+e^{2\Delta\beta}}\gamma_{i+1}$.

Proof.

We next bound $\gamma_{i,l}$ by distinguishing cases depending on the size l. For each case, we let $\mathbf{x}$ and j be the profile and the monochromatic block that attains the maximum $\gamma_{i,l}$.

<u>$l = 1$:</u>

- if the unique player of the j-th monochromatic block is selected for update and she changes her strategy then $\tau_{i,j} = 1$. This happens with probability $\frac{1}{n} \cdot \left(1 - \frac{1}{1+e^{2\Delta\beta}}\right)$.
- if a neighbor v of the unique player of the j-th monochromatic block is selected for update and she changes her strategy then the dynamics reaches a profile $\mathbf{y}$ at same level $i+1$ and the size of the j-th block increases to 2. If v belongs to a monochromatic block of size 1, this has probability 0 (we are conditioning on $\tau_{i,j} < \tau_{i,j'}$ for all $j' \neq j$); otherwise, the probability is at most $1/2 \cdot 2/n = 1/n$.
- if we select for update a player that is not at the borders of a monochromatic block and she changes her strategy, then the dynamics reaches a profile $\mathbf{y}$ at level $i+2$. This has probability $\frac{n-4(i+1)+s(\mathbf{x})}{1+e^{2\Delta\beta}}$ of occurring.
- in the remaining cases neither the level nor the length of the j-th monochromatic block changes.

Hence, by observing that $\gamma_{i,2} \geqslant \gamma_{i,1}$ we have

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

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

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

Since $\left(\frac{1}{2} + \frac{1}{4e^{2\Delta\beta}+2}\right) \leqslant \frac{2}{3}$ for $\beta = \omega(\log n)$, we have

$$\gamma_{i,1} \leqslant \frac{2}{3}(\gamma_{i,2} + b_i)\,. \tag{10}$$

<u>$1 < l < n - 2i - 1$:</u>

- if a player at the borders of the j-th monochromatic block is selected for update (there are two of these players) and she changes her strategy (this happens with probability $1/2$), then the dynamics reaches a profile $\mathbf{y}$ at same level $i + 1$ and the length of the j-th monochromatic block decreases to $l - 1$;

- if a neighbor v of the border players of the j-th monochromatic block is selected for update and she changes her strategy, then the number of monochromatic blocks does not change (and thus we are at still at level $i + 1$) but the j-th monochromatic block increases in size.

  Notice that, in this case, player v does not belong to a monochromatic block of size 1, since we are conditioning on the fact that the j-th monochromatic block is the first to disappear ($\tau_{i,j} < \tau_{i,j'}\ \forall j' \neq j$). Therefore the two neighbors of v are playing two different strategies and thus v adopts any of the two with probability $1/2$. Since there are two players adjacent to the border players of block j, this case happens with probability at most $1/n$.

- if a player v that is not at the borders of a monochromatic block is selected for update and she changes her strategy then the two new adjacent monochromatic blocks are created and the level increases 1. Notice that there are $n - 4(i+1) + s(\mathbf{x}^\star)$ such player v and each has probability $\frac{1}{1+e^{2\Delta\beta}}$ of changing her strategy.

- in the remaining cases neither the level nor the length of the j-th monochromatic block changes.

Hence,

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

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

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

From the previous inequality and Equation 10, a simple induction on l shows that, for every $1 \leqslant l < n - 2i - 1$, we have

$$\gamma_{i,l} \leqslant \frac{1}{l+2}\left((l+1)\gamma_{i,l+1} + \frac{l(l+3)}{2}b_i\right). \tag{11}$$

Moreover, from Equation 11, a simple inductive argument shows that, for every $h \geqslant 1$,

$$\begin{aligned}\gamma_{i,l} &\leqslant \frac{l+1}{l+h+1}\gamma_{i,l+h} + \frac{l+1}{2}b_i \sum_{j=l}^{l+h-1} \frac{j(j+3)}{(j+1)(j+2)} \\ &\leqslant \frac{l+1}{l+h+1}\gamma_{i,l+h} + \frac{l+1}{2}hb_i\,.\end{aligned} \tag{12}$$

<u>$l = n - 2i - 1$:</u> in this case all blocks other than the j-th have size 1 and thus every time we select one of these players, she changes her strategy with probability 0 (we are conditioning on j being the first monochromatic block to disappear). This means that that the size of the j-th monochromatic block cannot increase. Reasoning similar to the ones used in the previous cases, we obtain that

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

By using Equation 11, we have

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

Finally, for every $l \geqslant 1$, by using Equation 12 with $h = n-2i-1-l$, we have

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

□

Corollary 5.26 *If $\beta = \omega(\log n)$, then for every $i \geqslant 0$, $\gamma_i = \mathcal{O}(n^3)$.*

Proof. Note that $b_{\lfloor n/2 \rfloor -1} = n$ and thus, by using Lemma 5.25, $\gamma_{\lfloor n/2 \rfloor -1} \leqslant n^3$. Moreover, for $0 \leqslant i < \lfloor n/2 \rfloor - 1$, since $n - 4(i+1) + s(\mathbf{x}) \leqslant n$, we have

$$\gamma_i \leqslant n^3 \left(1 + \frac{1}{1+e^{2\Delta\beta}}\gamma_{i+1}\right) \leqslant n^3 \left(1 + \sum_{j=1}^{\lfloor n/2 \rfloor - i - 1} \left(\frac{n^3}{1+e^{2\Delta\beta}}\right)^j\right) . \tag{13}$$

The corollary follows by observing that, if $\beta = \omega\left(\frac{\log n}{\Delta}\right)$, then the summation in Equation 13 is $o(1)$. □

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 5.12 easily follows.

Proof of Lemma 5.12. Obviously, for every $\mathbf{x}$ at level $1 \leqslant k \leqslant n/2$,

$$\mathbf{E}_{\mathbf{x}}\left[\tau_{\mathbf{0},\mathbf{1}}\right] \leqslant \sum_{i=0}^{k-1} \max_{\mathbf{y}\colon \ell(\mathbf{y}) = i+1} \mathbf{E}_{\mathbf{y}}\left[\tau_i\right] \leqslant \sum_{i=0}^{k-1} \gamma_i = \mathcal{O}(n^4) \,.$$

The lemma follows from the Markov inequality. □

5.4.2 Proof of Lemma 5.13

Proof of Lemma 5.13. For a profile $\mathbf{x}$, we denote by $p_{\mathbf{x}}$ the probability that the logit dynamics starting from $\mathbf{x}$ at step $\tau_{\mathbf{0},\mathbf{1}}$ is in profile $\mathbf{0}$; in other words, $p_{\mathbf{x}} = \mathbf{P}_{\mathbf{x}}\left(\tau_{\mathbf{0}} < \tau_{\mathbf{1}}\right)$. Trivially, $p_{\mathbf{0}} = 1$ and $p_{\mathbf{1}} = 0$.

Clearly, at time $\tau_{\mathbf{0},\mathbf{1}}$ the dynamics is either in the state $\mathbf{0}$ (this happens with probability $p_{\mathbf{x}}$) or in the state $\mathbf{1}$ (this happens with probability 1 - $p_{\mathbf{x}}$). Thus, the state of the dynamics at time $\tau_{\mathbf{0},\mathbf{1}}$ when starting from profile x is distributed according to the probability distribution

$$\nu_{\mathbf{x}} = p_{\mathbf{x}}\pi_{\mathbf{0}} + (1-p_{\mathbf{x}})\pi_{\mathbf{1}} .$$

We next show that for $\beta = \omega(\log n)$ and $x \in S_d$, $p_{\mathbf{x}} = \frac{d}{n} + \lambda_{\mathbf{x}}$, for some $\lambda_{\mathbf{x}} = o(1)$. By the definition of Markov chains we know that

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

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 $P(\mathbf{x},\mathbf{y}_1) = P(\mathbf{x},\mathbf{y}_2)$.

- $N_1(\mathbf{x})$ is the set of profiles $\mathbf{y}$ obtained from $\mathbf{x}$ by changing the strategy of a player of a zero-block of size 1. Observe that $|N_1(\mathbf{x})| = s_0(\mathbf{x})$. Moreover, for every $\mathbf{y} \in N_1(\mathbf{x})$, $\mathbf{y}$ is at level $i-1$, has $|\mathbf{x}|_0 - 1$ players playing 0 and $P(\mathbf{x},\mathbf{y}) = \frac{1}{n} \cdot (1 - \frac{1}{1+e^{2\Delta\beta}})$.
- $N_2(\mathbf{x})$ is the set of profiles $\mathbf{y}$ obtained from $\mathbf{x}$ by changing the strategy of a player of a one-block of size 1. Observe that $|N_2(\mathbf{x})| = s_1(\mathbf{x})$. Moreover, for every $\mathbf{y} \in N_2(\mathbf{x})$, $\mathbf{y}$ is at level $i-1$, has $|\mathbf{x}|_0 + 1$ players playing 0 and $P(\mathbf{x},\mathbf{y}) = \frac{1}{n} \cdot (1 - \frac{1}{1+e^{2\Delta\beta}})$.
- $N_3(\mathbf{x})$ is the set of profiles $\mathbf{y}$ obtained from $\mathbf{x}$ by changing the strategy of a border player of a zero-block of size greater than 1. Observe that $|N_3(\mathbf{x})| = 2(i - s_0(\mathbf{x}))$. Moreover, for every $\mathbf{y} \in N_3(\mathbf{x})$, $\mathbf{y}$ is at level i, and has $|\mathbf{x}|_0 - 1$ players playing 0 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 one-block of size greater than 1. Observe that $|N_4(\mathbf{x})| = 2(i - s_1(\mathbf{x}))$. Moreover, for every $\mathbf{y} \in N_4(\mathbf{x})$, $\mathbf{y}$ is at level i, and has $|\mathbf{x}|_0 + 1$ players playing 0 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

$$\begin{aligned}
p_{\mathbf{x}} &= \frac{1}{n}\left(1 - \frac{1}{1+e^{2\Delta\beta}}\right)\left(\sum_{\mathbf{y}\in N_1(\mathbf{x})} p_{\mathbf{y}} + \sum_{\mathbf{y}\in N_2(\mathbf{x})} p_{\mathbf{y}}\right) + \frac{1}{2n}\left(\sum_{\mathbf{y}\in N_3(\mathbf{x})} p_{\mathbf{y}} + \sum_{\mathbf{y}\in N_4(\mathbf{x})} p_{\mathbf{y}}\right) \\
&\quad + \frac{1}{n}\frac{1}{1+e^{2\Delta\beta}} \sum_{\mathbf{y}\in N_5(\mathbf{x})} p_{\mathbf{y}} \\
&\quad + \left(\frac{s(\mathbf{x})}{n}\frac{1}{1+e^{2\Delta\beta}} + \frac{2i - s(\mathbf{x})}{n} + \frac{n-4i+s(\mathbf{x})}{n}\left(1 - \frac{1}{1+e^{2\Delta\beta}}\right)\right) p_{\mathbf{x}} \\
&= \frac{1}{n}\left(\sum_{\mathbf{y}\in N_1(\mathbf{x})} p_{\mathbf{y}} + \sum_{\mathbf{y}\in N_2(\mathbf{x})} p_{\mathbf{y}}\right) + \frac{1}{2n}\left(\sum_{\mathbf{y}\in N_3(\mathbf{x})} p_{\mathbf{y}} + \sum_{\mathbf{y}\in N_4(\mathbf{x})} p_{\mathbf{y}}\right) + \frac{n-2i}{n}\cdot p_{\mathbf{x}} + \frac{c}{1+e^{2\Delta\beta}}, \qquad (14)
\end{aligned}$$

where

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

We notice that, since $1 \leqslant i \leqslant n/2$ and $|N_1(\mathbf{x})| + |N_2(\mathbf{x})|, |N_5(\mathbf{x})| \leqslant n$, we have $|c| \leqslant 2$ and thus the last term in Equation 14 is negligible in n (since $\beta = \omega(\log n)$). Hence we have that the following condition holds for every level $i \geqslant 1$ and every profile $\mathbf{x}$ at level i:

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

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.

Next we find a solution to a "modified" version of above system, where we omit the negligible part in every equation, and then we show that this solution cannot be very different from the solution of the "original" system.

We build the solution for the "modified" system inductively on the level i: for every profile $\mathbf{x} \in S_d$ at level 0 (this is only possible for $d = 0$ or $d = n$), we have, as discussed above, $p_{\mathbf{x}} = \frac{d}{n}$. Now, we assume that for every profile $\mathbf{x} \in S_d$ at level $i-1$, $p_{\mathbf{x}} = \frac{d}{n}$ is a solution for the system. For $\mathbf{x} \in S_d$ at level i, we can rewrite the "modified" condition as follows:

$$p_{\mathbf{x}} = \frac{s_0(\mathbf{x})}{2i}\cdot\frac{d-1}{n} + \frac{s_1(\mathbf{x})}{2i}\cdot\frac{d+1}{n} + \frac{1}{4i}\left(\sum_{\mathbf{y}\in N_3(\mathbf{x})} p_{\mathbf{y}} + \sum_{\mathbf{y}\in N_4(\mathbf{x})} p_{\mathbf{y}}\right). \qquad (15)$$

Equation 15 gives another linear system of equations. This system has a unique solution: indeed, it has the same number of equations and variables and the matrix of coefficients is a diagonally dominant matrix (since $|N_3(\mathbf{x}) \cup N_4(\mathbf{x})| \leqslant 4i$) and thus it is nonsingular. Moreover, if we set, for every profile $\mathbf{x}$ at level i, $p_{\mathbf{x}} = \frac{|\mathbf{x}|_0}{n}$, then the right hand side of the Equation 15 becomes

$$\frac{s_0(\mathbf{x})}{2i}\frac{d-1}{n} + \frac{s_1(\mathbf{x})}{2i}\frac{d+1}{n} + \frac{i - s_0(\mathbf{x})}{2i}\frac{d-1}{n} + \frac{i - s_1(\mathbf{x})}{2i}\frac{d+1}{n} = \frac{d}{n},$$

and hence the system is satisfied by this assignment. Summarizing, we have found that the "modified" system has a unique solution $p_{\mathbf{x}} = \frac{|\mathbf{x}|_0}{n}$ for every profile $\mathbf{x}$. Now, let $p^\star_{\mathbf{x}} = p_{\mathbf{x}} + \lambda_{\mathbf{x}}$ be the assignment that satisfies all "original" conditions: since, as n grows unbounded, these conditions approach the "modified" ones, we have that $p^\star_{\mathbf{x}}$ has to approach to $p_{\mathbf{x}}$ and thus we have that $|\lambda_{\mathbf{x}}| = o(1)$ for every profile $\mathbf{x}$. □

6 The OR game

The *OR-game* is a toy n-player game where every player has two strategies, say $\{0, 1\}$, and each player *pays* the *OR* of the strategies of all players (including herself). More formally, the utility of player $i \in [n]$ is $u_i(\mathbf{0}) = 0$ and $u_i(\mathbf{x}) = -1$ for every $\mathbf{x} \neq \mathbf{0}$.

In [3] we showed that the mixing time of the logit dynamics for the OR-game is roughly e^β for $\beta = \mathcal{O}(\log n)$ and it is roughly 2^n for larger β. Here we study the metastability properties of the OR game to highlight the distinguishing features of our quantitative notion of metastability based on distributions. Namely, we show that if we start the logit dynamics at a profile where at least one player is playing 1, then after $\mathcal{O}(\log n)$ time steps the distribution of the chain is close to uniform, and it stays close to uniform for exponential time. Hence, even if there is no small set of the state space where the chain stays close for a long time, we can still say that the chain is "metastable" meaning that the "distribution" of the chain stays close to some well-defined distribution for a long time.

6.1 Ehrenfest urns

We first need two simple lemmas that will be used in the proof of Theorem 6.4.

The *Ehrenfest urn* is the Markov chain with state space $\Omega = \{0, 1, \ldots, n\}$ 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 [25] 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 within time step t.

Lemma 6.1 *Let $\{Z_t\}$ be the Ehrenfest urn over $\{0, 1, \ldots, n\}$ and let τ_0 be the first time the chain hits state 0. Then for every $k \geqslant 1$ it holds that*

$$\mathbf{P}_k\left(\tau_0 < n \log n + cn\right) \leqslant \frac{c'}{n}$$

for suitable positive constants c and c'.

Proof. First observe that for any $t \geqslant 3$ the probability of hitting 0 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

$$\begin{aligned}
\mathbf{P}_1\left(\tau_0 < t\right) &= \mathbf{P}_1\left(\tau_0 < t \,|\, Z_1 = 0\right) \mathbf{P}_1\left(Z_1 = 0\right) + \mathbf{P}_1\left(\tau_0 < t \,|\, Z_1 = 2\right) \mathbf{P}_1\left(Z_1 = 2\right) \\
&= \frac{1}{n} + \frac{n-1}{n} \mathbf{P}_2\left(\tau_0 < t - 1\right) \leqslant \frac{1}{n} + \mathbf{P}_2\left(\tau_0 < t\right) \\
\mathbf{P}_2\left(\tau_0 < t\right) &= \mathbf{P}_2\left(\tau_0 < t \,|\, Z_1 = 1\right) \mathbf{P}_2\left(Z_1 = 1\right) + \mathbf{P}_2\left(\tau_0 < t \,|\, Z_1 = 3\right) \mathbf{P}_2\left(Z_1 = 3\right) \\
&= \frac{2}{n} \mathbf{P}_1\left(\tau_0 < t - 1\right) + \frac{n-2}{n} \mathbf{P}_3\left(\tau_0 < t - 1\right) \\
&\leqslant \frac{2}{n}\left(\frac{1}{n} + \mathbf{P}_2\left(\tau_0 < t\right)\right) + \frac{n-2}{n} \mathbf{P}_3\left(\tau_0 < t\right)
\end{aligned}$$

Hence,

$$\begin{aligned}
\mathbf{P}_2\left(\tau_0 < t\right) &\leqslant \frac{2}{n-2} + \mathbf{P}_3\left(\tau_0 < t\right) \leqslant \frac{3}{n} + \mathbf{P}_3\left(\tau_0 < t\right) \\
\mathbf{P}_1\left(\tau_0 < t\right) &\leqslant \frac{4}{n} + \mathbf{P}_3\left(\tau_0 < t\right)
\end{aligned}$$

Moreover observe that the probability that the chain starting at k hits state 0 before time t is decreasing in k, in particular, for every $k \geqslant 3$ it holds that $\mathbf{P}_k(\tau_0 < t) \leqslant \mathbf{P}_3(\tau_0 < t)$. Now we show that $\mathbf{P}_3(\tau_0 < n\log n + cn) = \mathcal{O}(1/n)$ and this will complete the proof.

Let us 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 viceversa). The probability for any such pair of moves is

$$\frac{n-i}{n} \cdot \frac{i+1}{n} \leqslant \frac{e^{1/n}}{4}$$

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 $\mathcal{P}$ is[3]

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

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

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

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

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

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

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

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

In the proof of Theorem 6.4 we will be dealing with the lazy version of the Ehrenfest urn. The next lemma, which is folklore, allows us to use the bound we achieved in Lemma 6.1 for the non-lazy chain.

Lemma 6.2 *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 state $a \in \Omega$ 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*

$$\mathbf{P}_b(\hat{\tau}_a \leqslant t) \leqslant \mathbf{P}_b(\tau_a \leqslant t)$$

6.2 OR game metastability

The next lemma shows that, if we start from the uniform distribution, the distribution of the logit dynamics stays ε-close to uniform for $\varepsilon 2^n$ time steps.

Lemma 6.3 *Let P be the transition matrix of the logit dynamics for the n-player OR-game, let U be the uniform distribution over $\{0,1\}^n$. Then U is $(\varepsilon, \varepsilon 2^n)$-metastable.*

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

Proof. Observe that, by starting from the stationary distribution, the probability of being in $\mathbf{y} \in \{0,1\}^n$ after one step of the chain is

$$UP(\mathbf{y}) = \sum_{\mathbf{x} \in \{0,1\}^n} U(\mathbf{x})P(\mathbf{x},\mathbf{y}) = \frac{1}{2^n} \sum_{\mathbf{x} \in \{0,1\}^n} P(\mathbf{x},\mathbf{y}) = \begin{cases} 2^{-n} & \text{if } |\mathbf{y}| \geqslant 2 \\ 2^{-n}\left(\frac{n-1}{n} + \frac{1}{n}\frac{2}{1+e^{\beta}}\right) & \text{if } |\mathbf{y}| = 1 \\ 2^{-n}\left(\frac{2}{1+e^{-\beta}}\right) & \text{if } \mathbf{y} = \mathbf{0} \end{cases}$$

Hence, the total variation distance between the uniform distribution and the distribution of the chain after one step is

$$\|UP - U\| = \frac{1}{2} \sum_{\mathbf{y} \in \{0,1\}^n} |UP(\mathbf{y}) - U(\mathbf{y})| = 2^{-n}\frac{e^{\beta}-1}{e^{\beta}+1} \leqslant 2^{-n}$$

Thus, the uniform distribution is $(2^{-n}, 1)$-metastable and the thesis follows from Lemma 3.2. □

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

Theorem 6.4 *Let P be the transition matrix of the logit dynamics for the n-player OR-game, let U be the uniform distribution over $\{0,1\}^n$, let $\mathbf{x} \in \{0,1\}^n$ with $|\mathbf{x}| = k \geqslant 1$ be the starting state of the chain, and let $\varepsilon > 0$, then it holds that*

$$\|P^t(\mathbf{x},\cdot) - U\| \leqslant \varepsilon$$

for every time t such that $n\log(3n/\varepsilon) \leqslant t \leqslant \varepsilon 2^{n-1}$.

Proof. Let $\{X_t\}$ be the Markov chain starting at $\mathbf{x}$ and let $\{Y_t\}$ be a lazy random walk on the n-cube starting at the uniform distribution, so that X_t is distributed according to $P^t(\mathbf{x},\cdot)$ and Y_t is uniformly distributed over $\{0,1\}^n$. Consider the following coupling (X_t, Y_t): when chain $\{X_t\}$ is at state $\mathbf{y} \in \{0,1\}^n$ then choose a position $i \in [n]$ u.a.r. and

- If $|\mathbf{y}| \geqslant 2$ then[4], choose an action $a \in \{0,1\}$ u.a.r. and update both chains X_t and Y_t in position i with action a;
- If $|\mathbf{y}| = 1$ then
  - if X_t has 0 in position i than proceed as in the previous case;
  - if X_t has 1 in position i then
    - update both chains at 0 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 0 and chain Y_t at 1 in position i with probability $1/(1+e^{-\beta}) - 1/2$.
- If $|\mathbf{y}| = 0$ then
  - update both chains at 0 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 0 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, hence $\|P^t(\mathbf{x},\cdot) - U\| \leqslant \mathbf{P}_{\mathbf{x},U}(X_t \neq Y_t)$. Moreover observe that, if at time t all players have been selected at least once and chain X_t has not yet hit profile $\mathbf{0} = (0,\cdots,0) \in \{0,1\}^n$, then the two random variables X_t and Y_t have the same value. Hence

$$\begin{aligned} \|P^t(\mathbf{x},\cdot) - U\| &\leqslant \mathbf{P}_{\mathbf{x},U}(X_t \neq Y_t) \\ &\leqslant \mathbf{P}_{\mathbf{x},U}(\tau_{\mathbf{0}} \leqslant t \cup \eta < t) \\ &\leqslant \mathbf{P}_{\mathbf{x},U}(\tau_{\mathbf{0}} \leqslant t) + \mathbf{P}_{\mathbf{x},U}(\eta < t) \end{aligned}$$

[4]With $|\mathbf{y}|$ we mean the number of 1's in profile $\mathbf{y}$

where $\tau_{\mathbf{0}}$ is the hitting time of $\mathbf{0}$ for chain X_t, and η is the first time all players have been selected at least once.

From the coupon collector's argument it follows that for every $t \geqslant n \log(3n/\varepsilon)$

$$\mathbf{P}_{\mathbf{x},U}\left(\eta < t\right) \leqslant \varepsilon/3 \tag{16}$$

As for the second term observe that $\mathbf{P}_{\mathbf{x},U}\left(\tau_{\mathbf{0}} \leqslant t\right) \leqslant \mathbf{P}_k\left(\rho_0 \leqslant t\right)$ where ρ_0 is the hitting time of state 0 for the lazy Ehrenfest urn. More formally, consider the equivalence relation over $\Omega = \{0,1\}^n$ such that two profiles $\mathbf{x}$ and $\mathbf{y}$ are equivalent if they have the same number of 1's and let $\{Z_t\}$ be the *projection* of chain $\{X_t\}$ over the quotient space $\Omega_{\#} = \{0,1,\ldots,n\}$ of such equivalence relation. Then $\{Z_t\}$ is a Markov chain with state space $\Omega_{\#}$ and transition matrix

$$P_{\#}(i,i-1) = \frac{i}{2n}; \qquad P_{\#}(i,i) = \frac{1}{2}; \qquad P_{\#}(i,i+1) = \frac{n-i}{2n}; \qquad \text{for } i = 2,\ldots,n \tag{17}$$

and

$$P_{\#}(1,0) = \frac{1}{n(1+e^{-\beta})} \leqslant \frac{1}{n} \qquad P_{\#}(1,1) = \frac{n-1}{2n} + \frac{1}{n(1+e^{\beta})} \qquad P_{\#}(1,2) = \frac{n-1}{2n}$$

The hitting time $\tau_{\mathbf{0}}$ of state $\mathbf{0} \in \Omega$ for chain $\{X_t\}$ coincide with the hitting time $\hat{\rho}_0$ of state $0 \in \Omega_{\#}$ for the projection Z_t.

Observe that, from the transition probabilities in (17), chain $\{Z_t\}$ is *almost* the lazy Ehrenfest urn, the only difference being at states 1 and 0. Moreover, the transition from state 1 to state 0 in the Z_t holds with probability smaller than the probability of the same transition in the Ehrenfest urn. From Lemmas 6.1 and 6.2 it follows that

$$\mathbf{P}_{\mathbf{x},U}\left(\tau_{\mathbf{0}} \leqslant n\log n + n\log(3/\varepsilon)\right) \leqslant c/n \tag{18}$$

for a suitable constant $c = c(\varepsilon)$. Hence, for $t = n\log n + n\log(3/\varepsilon)$, by combining (16) and (18) it holds that

$$\|P^t(\mathbf{x},\cdot) - U\| \leqslant \frac{\varepsilon}{3} + \frac{c}{n} \leqslant \frac{\varepsilon}{2}$$

for n sufficiently large.

Since from Lemma 6.3 we have that the uniform distribution is $(\varepsilon/2, \varepsilon 2^{n-1})$-metastable, the thesis follows by applying Lemma 3.4. □

7 Conclusions and open problems

Logit dynamics is a clean and tractable dynamics that well models 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 possible to find some distributions that well-approximate the distribution of the chain for a time-window of 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 Ising model on the complete graph, we showed that when $\beta > c\log n/n$ the two degenerate distributions are metastable for $poly(n)$ time and they are quickly reached from a large fraction of the state space. It would be interesting to investigate the metastability properties when $1/n < \beta < \log n/n$. Indeed, in that range the mixing time is exponential but the distributions concentrated in the two extremal states are not metastable.

Acknowledgement. We wish to thank Paolo Penna for useful ideas, hints, and discussions.

References

[1] 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.

[2] Vincenzo Auletta, Diodato Ferraioli, Francesco Pasquale, Paolo Penna, and Giuseppe Persiano. Convergence to equilibrium of logit dynamics for strategic games. In *Proc. of the 23rd ACM Symp. on Parallelism in Algorithms and Architectures (SPAA'11)*, pages 197–206. ACM, 2011.

[3] Vincenzo Auletta, Diodato Ferraioli, Francesco Pasquale, and Giuseppe Persiano. Mixing time and stationary expected social welfare of logit dynamics. In *Proc. of the 3rd Int. Symp. on Algorithmic Game Theory (SAGT'10)*, volume 6386 of *LNCS*, pages 54–65. Springer, 2010.

[4] Joel Beltran and Claudio Landim. Metastability of reversible finite state Markov processes. *Stochastic Processes and their Applications*, 121:1633–1677, 2011.

[5] 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.

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

[7] Anton Bovier. Metastability: a potential theoretic approach. In *Proc. of the International Congress of Mathematicians*, volume III, pages 499–518. European Mathematical Society, 2006.

[8] 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.

[9] 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.

[10] Anton Bovier and Francesco Manzo. Metastability in Glauber dynamics in the low-temperature limit: Beyond exponential asymptotics. *Journal of Statistical Physics*, 107:757–779, 2002.

[11] Xi Chen, Xiaotie Deng, and Shang-Hua Teng. Settling the complexity of computing two-player Nash equilibria. *J. ACM*, 56(3), 2009.

[12] Constantinos Daskalakis, Paul W. Goldberg, and Christos H. Papadimitriou. The complexity of computing a Nash equilibrium. *SIAM J. Comput.*, 39(1):195–259, 2009.

[13] Jian Ding, Eyal Lubetzky, and Yuval Peres. Censored Glauber dynamics for the mean field Ising model. *Journal of Statistical Physics*, 137:407–458, 2009.

[14] 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.

[15] Glenn Ellison. Learning, local interaction, and coordination. *Econometrica*, 61(5):1047–1071, 1993.

[16] Henry Eyring. The activated complex in chemical reactions. *Journal of Chemical Physics*, 3:107–115, 1935.

[17] Mark Iosifovich Freidlin and Alexander D. Wentzell. *Random Perturbations of Dynamical Systems*. Springer, 1984.

[18] Serge Galam and Bernard Walliser. Ising model versus normal form game. *Physica A: Statistical Mechanics and its Applications*, 389(3):481 – 489, 2010.

[19] John C. Harsanyi and Reinhard Selten. *A General Theory of Equilibrium Selection in Games*. MIT Press, 1988.

[20] Frank den Hollander. Metastability under stochastic dynamics. *Stochastic Processes and their Applications*, 114(1):1–26, 2004.

[21] Frank den Hollander. Three lectures on metastability under stochastic dynamics. In *Methods of Contemporary Mathematical Statistical Physics*, volume 1970 of *Lecture Notes in Mathematics*, pages 1–24. Springer Berlin / Heidelberg, 2009.

[22] Hendrik Anthony Kramers. Brownian motion in a field of force and the diffusion model of chemical reactions. *Physica*, 7(4):284–304, 1940.

[23] Hernán Larralde, Francois Leyvraz, and David P. Sanders. Metastability in Markov processes. *Journal of Statistical Mechanics: Theory and Experiment*, 2006(8):1–29, 2006.

[24] 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.

[25] David Levin, Yuval Peres, and Elizabeth L. Wilmer. *Markov Chains and Mixing Times*. American Mathematical Society, 2008.

[26] 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.

[27] 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.

[28] Enzo Olivieri and Maria Eulália Vares. *Large deviation and metastability*. Cambridge University Press, 2005.

[29] H. Peyton Young. *Individual Strategy and Social Structure: An Evolutionary Theory of Institutions*. Princeton University Press, 1998.

[30] H. Peyton Young. The diffusion of innovations in social networks. Economics Working Paper Archive number 437, Johns Hopkins University, Department of Economics, 2000.

Appendix

A Markov chain summary

In this section we recall some basic facts about Markov chains. Fore a more detailed description and for notation conventions we refer the reader to [25].

Consider a Markov chain $\mathcal{M}$ with *finite* state space Ω and transition matrix P. It is a classical result that if $\mathcal{M}$ is *irreducible* and *aperiodic*[5] (i.e., *ergodic*) there exists an unique *stationary distribution*; that is, a distribution π on Ω such that $\pi \cdot P = \pi$.
The *total variation* distance $\|\mu - \nu\|_{\mathrm{TV}}$ between two probability distributions μ and ν on Ω is defined as

$$\|\mu - \nu\|_{\mathrm{TV}} = \max_{A \subset \Omega} |\mu(A) - \nu(A)| = \frac{1}{2} \sum_{x \in \Omega} |\mu(x) - \nu(x)| \, .$$

An irreducible and aperiodic Markov chain $\mathcal{M}$ *converges* to its stationary distribution π; specifically, there exists $0 < \alpha < 1$ such that

$$d(t) \leqslant \alpha^t,$$

where

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

and $P^t(x, \cdot)$ is the distribution at time t of the Markov chain starting at x. For $0 < \varepsilon < 1/2$, the *mixing time* is defined as

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

It is usual to set $\varepsilon = 1/4$ or $\varepsilon = 1/2e$. If not explicitly specified, when we write t_{mix} we mean $t_{\mathrm{mix}}(1/4)$. Observe that $t_{\mathrm{mix}}(\varepsilon) \leqslant \lceil \log_2 \varepsilon^{-1} \rceil t_{\mathrm{mix}}$.

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 *coupling of a Markov chain* $\mathcal{M}$ with transition matrix P is a process $(X_t, Y_t)_{t=0}^{\infty}$ with the property that both X_t and Y_t are Markov chains with transition matrix P. When the two coupled chains start at $(X_0, Y_0) = (x, y)$, we write $\mathbf{P}_{x,y}(\cdot)$ and $\mathbf{E}_{x,y}[\cdot]$ for the probability and the expectation on the space where the two chains are both defined.
We denote by τ_{couple} the first time the two chains meet; that is,

$$\tau_{\mathrm{couple}} = \min\{t : X_t = Y_t\} \, .$$

We will consider only couplings of Markov chains with the property that for $s \geqslant \tau_{\mathrm{couple}}$, it holds $X_s = Y_s$. The following theorem establish the importance of this tool (see, for example, Theorem 5.2 in [25]).

Theorem A.1 (Coupling) *Let $\mathcal{M}$ be a Markov chain with finite state space Ω and transition matrix P. For each pair of states $x, y \in \Omega$ consider a coupling (X_t, Y_t) of $\mathcal{M}$ with starting states $X_0 = x$ and $Y_0 = y$. Then*

$$\left\| P^t(x, \cdot) - P^t(y, \cdot) \right\|_{\mathrm{TV}} \leqslant \mathbf{P}_{x,y}\left(\tau_{couple} > t\right) \, .$$

Consider a partial order $\preceq$ over the states in Ω. A coupling of a Markov chain is said to be *monotone* w.r.t. $(\Omega, \preceq)$ if, for every $t \geqslant 0$, $X_t \preceq Y_t \Rightarrow X_{t+1} \preceq Y_{t+1}$. For a state $z \in \Omega$, the *hitting time* τ_z of z is the first time the chain is in state z, $\tau_z = \inf\{t : X_t = z\}$. The following theorem holds.

Lemma A.2 *Let $\mathcal{M}$ be a Markov chain with finite state space Ω and transition matrix P. Let $\preceq$ be a partial order over Ω. For each pair of states $x, y \in \Omega$ consider a coupling (X_t, Y_t) of $\mathcal{M}$ with starting states $X_0 = x$ and $Y_0 = y$ that is monotone w.r.t. $(\Omega, \preceq)$. Moreover, suppose the ordered set $(\Omega, \preceq)$ has an unique maximum at z. Then*

$$\mathbf{P}_{x,y}\left(\tau_{couple} > t\right) \leqslant 2 \cdot \max\{\mathbf{P}_x\left(\tau_z > t\right), \mathbf{P}_y\left(\tau_z > t\right)\} \, .$$

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