

Convergence to Equilibrium of Logit Dynamics for Strategic Games*

Vincenzo Auletta
Dipartimento di Informatica
Università di Salerno
84084 - Fisciano (SA), Italy
auletta@dia.unisa.it

Diodato Ferraioli
Dipartimento di Informatica
Università di Salerno
84084 - Fisciano (SA), Italy
ferraioli@dia.unisa.it

Francesco Pasquale
Dipartimento di Informatica
Università di Salerno
84084 - Fisciano (SA), Italy
pasquale@dia.unisa.it

Paolo Penna
Dipartimento di Informatica
Università di Salerno
84084 - Fisciano (SA), Italy
penna@dia.unisa.it

Giuseppe Persiano
Dipartimento di Informatica
Università di Salerno
84084 - Fisciano (SA), Italy
giuper@dia.unisa.it

ABSTRACT

We present the first general bounds on the mixing time of logit dynamics for wide classes of strategic games. The logit dynamics describes the behaviour of a complex system whose individual components act *selfishly* and keep responding according to some partial (“noisy”) knowledge of the system. In particular, we prove nearly tight bounds for potential games and games with dominant strategies. Our results show that, for potential games, the mixing time is upper and lower bounded by an *exponential* in the inverse of the noise and in the maximum potential difference. Instead, for games with dominant strategies, the mixing time cannot grow arbitrarily with the inverse of the noise. Finally, we refine our analysis for a subclass of potential games called *graphical* coordination games and we give evidence that the mixing time strongly depends on the structure of the underlying graph. Games in this class have been previously studied in Physics and, more recently, in Computer Science in the context of diffusion of new technologies.

Categories and Subject Descriptors

G.3 [Probability and Statistics]: Markov processes; J.4 [Social and Behavioral Sciences]: Economics

General Terms

Algorithms, Theory

Keywords

Game Theory, Markov Chains, Mixing Time

*Partially supported by MIUR PRIN project COGENT.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

SPAA '11, June 4–6, 2011, San Jose, California, USA.

Copyright 2011 ACM 978-1-4503-0743-7/11/06 ...\$10.00.

1. INTRODUCTION

Complex systems are often studied by looking at their dynamics and the equilibria induced by these dynamics. In this paper we concentrate on specific complex systems arising from *strategic games*. Here we have a set of selfish agents or *players*, each with a set of possible actions or *strategies*. An agent continuously evaluates her utility or *payoff*, that depends on her own strategy and on the strategies played by the other agents. A dynamics specifies the rule used by the players to update their strategies. In its most general form an equilibrium is a distribution over the set of states that has the property of being invariant with respect to the dynamics. For example, a very well-studied dynamics for strategic games is the *best response dynamics* whose associated equilibria are the Nash equilibria.

There are several characteristics of a dynamics and of the associated equilibrium concept that concur to make the dynamics descriptive of a system. First of all, it is desirable that the dynamics gives only one equilibrium state or, in case a system admits more than one equilibrium for a given dynamics, that the equilibria look similar. For example, this is not the case for Nash equilibria as a game can admit more than one Nash equilibrium and sometimes the equilibria have strikingly different characteristics. In addition, the dynamics must be descriptive of the way individual agents behave. For example, the best response dynamics is well-tailored for modeling players that have a complete knowledge of the global state of the system and of their payoffs. Finally, if a dynamics takes very long time to reach an equilibrium then the system spends most of its life outside of the equilibrium and thus knowledge gained from the study of the equilibrium is not very relevant.

In this work we study a specific *noisy* best-response dynamics, the *logit dynamics* (defined in [4]) in which, at each time step, a player is randomly selected for strategy update and the update is performed with respect to a “noisy” knowledge of the game and of the state of the system, that is, the strategies currently played by the players. Intuitively, “high noise” represents the situation where players choose their strategies “nearly at random” because they have a limited knowledge of the system; instead, “low noise” represents the situation where players “almost surely” play the best re-

sponse; that is, they pick the strategies yielding high payoff with “much higher” probability. After a sufficiently large number of steps, the probability that the system is found in a specific profile remains essentially unchanged and we say that the logit dynamics has converged to a *stationary distribution*, that is unique and independent of the starting state. We believe that this makes the logit dynamics the elective choice of a dynamics for large and complex systems in which agents have limited knowledge. However, one more step is needed to complete the picture. How long does the logit dynamics take to converge to the stationary distribution? This is the main technical focus of this paper. Specifically, we study the *mixing time* of the logit dynamics, that is, the time needed to get close to the stationary distribution. This depends on the underlying game and on the noise of the system (roughly speaking, the payoffs and how much players care about them). Since previous work has shown that the mixing time can vary a lot (from linear to exponential [2]) it is natural to ask the following questions: (1) How do the *noise* level and the *structure* of the game affect the mixing time? (2) Can the mixing time grow *arbitrarily*? We give general bounds on the mixing time for wide classes of games, including potential games and games with dominant strategies, and coordination games played between neighboring nodes of a given network.

We prove that, for all potential games, the mixing time of the logit dynamics is upper-bounded by a *polynomial* in the number of players and by an *exponential* in the inverse of the noise and in the maximum potential difference, an important structural property of the game.

We complement the upper bound with a lower bound showing that there exist potential games with mixing time exponential in the inverse of the noise and in the maximum potential difference. Thus the mixing time can grow indefinitely in potential games as noise decreases. We also study a special class of potential games, the *graphical coordination games*, in which players are connected by a network and adjacent players play a two-player coordination game. By looking at two extreme cases, we give evidence that the mixing time depends on the connectedness of the underlying network. Specifically, we prove that, as a function of the number of players, the mixing time is exponential on the clique and polynomial on the ring, for a class of coordination games.

Going to the second question, we show that for games with dominant strategies (not necessarily potential games) the mixing time cannot exceed some *absolute bound* T which depends uniquely on the number of players n and on the number of strategies m . Though $T = T(n, m)$ is of the form $\mathcal{O}(m^n)$, it is independent of the noise and we show that, in general, such an exponential growth is the best possible.

Our results suggest that the structural properties of the game are important for the mixing time. For small noise, players tend to play best response and for those games that have more than one pure Nash equilibrium (PNE) with similar potential the system is likely to remain in a PNE for a long time, whereas the stationary distribution gives each PNE approximately the same weight. This happens for (certain) potential games, whence the exponential growth of mixing time with respect to the noise. On the contrary, for games with dominant strategies there is a PNE (a dominant profile) with high stationary probability *and* players

are guaranteed to play that profile with non-vanishing probability (regardless of the noise).

Related works. The logit dynamics was first studied by Blume [4] who showed that, for 2×2 coordination games, the long-term behavior of the system is concentrated in the risk dominant equilibrium (see [6]). The study of the mixing time of the logit dynamics for strategic games has been initiated in [2], where, among others, bounds were given for the class of 2×2 coordination games studied in [4]. Before the work reported in [2], the rate of convergence was studied only for the hitting time of specific profiles; see for example the work by Asadpour and Saberi [1] who studied the hitting time of the Nash equilibrium for a class of congestion games.

Graphical coordination games are often used to model the spread of a new technology in a social network [14] 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. Ellison [5] studied the logit dynamics for graphical coordination games on rings and showed that some large fraction of the players will eventually choose the strategy with maximum potential. Similar results were obtained by Peyton Young [14] for the logit dynamics and for more general families of graphs. Montanari and Saberi [11] 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. We notice that none of [4, 5, 14] gave bounds on the convergence rate of the dynamics, while Montanari and Saberi [11] studied the convergence time of a specific configuration, namely the hitting time of the highest potential equilibrium.

Our work is also strictly related to the well-studied Glauber dynamics on the Ising model (see, for example, [9] and Chapter 15 of [7]). Indeed, the Ising model can be seen as a special graphical coordination game without risk dominant equilibria, and the Glauber dynamics on the Ising model is equivalent to the logit dynamics. In particular, we note that Berger et al. [3] relate the mixing time of the Ising model to the cutwidth of the underlying graph. Their results can be specialized to derive upper bounds on the mixing time of graphical coordination games without risk dominant equilibria. However the bounds we present in Section 5 are tighter.

Even if the logit dynamics has attracted a lot of attention in different scientific communities, many other promising dynamics that deal with partial or noise-corrupted knowledge of the game have been proposed (see, for example, the recent work of Marden et al. [8] and of Mertikopoulos and Moustakas [10] and references in [13]).

Paper organization. We give formal definitions of logit dynamics and some of the used techniques in Section 2. The upper bounds for potential games, for games with dominant strategies, and for graphical coordination games are given in Section 3, Section 4, and Section 5, respectively. For the sake of completeness, in Appendix A we present some known facts about Markov chains that we use throughout the paper.

2. PRELIMINARIES

In this section we review the background on strategic games, introduce the logit dynamics and describe the proof techniques for deriving our bounds.

Games. In a *strategic game* we are given a finite set of players $\{1, \dots, n\}$, with each player i having a finite set of *strategies* S_i and a *utility function* $u_i : S_1 \times \dots \times S_n \rightarrow \mathbb{R}$. Each player can choose a strategy $x_i \in S_i$ and the resulting *strategy profile* is the vector $\mathbf{x} = (x_1, \dots, x_n)$. Given a profile \mathbf{x} , the utility (or payoff) for player i is $u_i(\mathbf{x})$. Throughout the paper we adopt the standard game theoretic notation and write (a, \mathbf{x}_{-i}) to denote the vector $(x_1, \dots, x_{i-1}, a, x_{i+1}, \dots, x_n)$. We also let $S := S_1 \times \dots \times S_n$ denote the set of all strategy profiles.

Logit dynamics. In this paper we consider the logit dynamics (see [4]). In the *logit dynamics with inverse noise* $\beta \geq 0$ for an n -player strategic game $\mathcal{G} = (S_1, \dots, S_n, u_1, \dots, u_n)$ at every time step a player i is selected uniformly at random and her strategy is updated to strategy $y \in S_i$ with probability $\sigma_i(y | \mathbf{x})$ defined as

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

where $\mathbf{x} \in S$ is the current strategy profile and $T_i(\mathbf{x}) = \sum_{z \in S_i} e^{\beta u_i(z, \mathbf{x}_{-i})}$ is the normalizing factor.

The logit dynamics for \mathcal{G} naturally defines a Markov chain $\mathcal{M}_\beta^\mathcal{G} = \{X_t : t \in \mathbb{N}\}$ with state space $\Omega = S$ and transition probabilities

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

We will find convenient to identify the logit dynamics for \mathcal{G} with the Markov chain $\mathcal{M}_\beta^\mathcal{G}$. It is easy to see that $\mathcal{M}_\beta^\mathcal{G}$ is irreducible and aperiodic. Therefore, there exists a unique *stationary distribution* π , such that, for every initial profile \mathbf{x} , the distribution $P^t(\mathbf{x}, \cdot)$ of the position of the chain after t steps converges to π as t tends to infinity. We are interested in the *mixing time* of the chain; i.e., the time needed for $P^t(\mathbf{x}, \cdot)$ to be close to π for every initial configuration \mathbf{x} :

$$t_{\text{mix}}(\varepsilon) = \min \{t \in \mathbb{N} : \|P^t(\mathbf{x}, \cdot) - \pi\|_{\text{TV}} \leq \varepsilon \text{ for all } \mathbf{x} \in \Omega\}$$

where $\|P^t(\mathbf{x}, \cdot) - \pi\|_{\text{TV}} = \frac{1}{2} \sum_{\mathbf{y} \in \Omega} |P^t(\mathbf{x}, \mathbf{y}) - \pi(\mathbf{y})|$ is the *total variation distance*. We will use the standard convention of setting $t_{\text{mix}} = t_{\text{mix}}(1/4)$ and observe that $t_{\text{mix}}(\varepsilon) \leq t_{\text{mix}}(1/4) \cdot \log 1/\varepsilon$.

For an irreducible and aperiodic Markov chain over finite state space Ω with transition matrix P and stationary distribution π , we will call *edge stationary distribution* the probability distribution Q over the set $\Omega \times \Omega$ given by $Q(\mathbf{x}, \mathbf{y}) = \pi(\mathbf{x}) P(\mathbf{x}, \mathbf{y})$.

Potential games. A strategic game is a *potential game* if there exists a function $\Phi : S \rightarrow \mathbb{R}$ such that for every player i , every profile $\mathbf{x} \in S$, and every pair of strategies $a, b \in S_i$, it holds that

$$u_i(a, \mathbf{x}_{-i}) - u_i(b, \mathbf{x}_{-i}) = \Phi(a, \mathbf{x}_{-i}) - \Phi(b, \mathbf{x}_{-i}).$$

For every potential game \mathcal{G} with potential function Φ , let π be the so-called Gibbs measure

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

where $Z = \sum_{\mathbf{y} \in S} e^{\beta \Phi(\mathbf{y})}$ is the normalizing constant, also called the *partition function*. We write Z_β and π_β when we want to stress the dependence on the inverse noise β .

It is easy to see that $\mathcal{M}_\beta^\mathcal{G}$ is reversible (i.e., $Q(\mathbf{x}, \mathbf{y}) = Q(\mathbf{y}, \mathbf{x})$ for all states \mathbf{x}, \mathbf{y}) and the Gibbs measure π in (3) is its stationary distribution.

Further notation. We use bold symbols for vectors. We write $|\mathbf{x}|_a$ for the number of occurrences of a in \mathbf{x} , i.e., $|\mathbf{x}|_a = |\{i \in [n] : x_i = a\}|$. We write $\mathbf{x} \sim \mathbf{y}$ to denote the fact that \mathbf{x} differs from \mathbf{y} in exactly one coordinate.

2.1 Proof techniques for the upper bounds

To derive our upper bounds, we employ two techniques: Markov chain coupling and Markov chain comparison. Coupling is a well-established technique for bounding the mixing time and it is summarized in Theorems 17 and 18 in Appendix A. We next review the comparison technique.

Let P be the transition matrix of a Markov chain with finite state space Ω and let us label the eigenvalues of P in non-increasing order

$$\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_{|\Omega|}.$$

It is well-known (see, for example, Lemma 12.1 in [7]) that, if P is irreducible and aperiodic, then $\lambda_2 < 1$ and $\lambda_{|\Omega|} > -1$. For irreducible and aperiodic chains the *relaxation time* t_{rel} is defined as

$$t_{\text{rel}} = \max \left\{ \frac{1}{1 - \lambda_2}, \frac{1}{1 + \lambda_{|\Omega|}} \right\},$$

and for reversible Markov chains (and thus also for the logit dynamics of potential games) it is related to the mixing time by the following theorem (see, for example, Theorem 12.3 in [7]).

THEOREM 1. *Let P the transition matrix of a reversible, irreducible, and aperiodic Markov chain with state space Ω and stationary distribution π . Then it holds that*

$$t_{\text{mix}} \leq \log \left(\frac{1}{\pi_{\min}} \right) t_{\text{rel}}$$

where $\pi_{\min} = \min_{\mathbf{x} \in \Omega} \pi(\mathbf{x})$.

The following theorem allows us to relate the relaxation times of two chains by comparing their stationary and edge-stationary distributions. We will use it in Theorem 5 to relate the logit dynamics with a random walk on a generalized hypercube whose relaxation time is known.

THEOREM 2 (COMPARISON THEOREM). *Let P and \hat{P} be the transition matrices of two reversible, irreducible, and aperiodic Markov chains with the same state space Ω , stationary distributions π and $\hat{\pi}$ respectively, and edge stationary distributions Q and \hat{Q} respectively. Suppose that two constants α, γ exist such that, for all $x, y \in \Omega$,*

$$\hat{Q}(x, y) \leq \alpha \cdot Q(x, y) \quad (4)$$

$$\pi(x) \leq \gamma \cdot \hat{\pi}(x). \quad (5)$$

Then the relaxation time t_{rel} of P and the relaxation time \hat{t}_{rel} of \hat{P} satisfy $t_{\text{rel}} \leq \alpha \cdot \gamma \cdot \hat{t}_{\text{rel}}$.

When $t_{\text{rel}} = \frac{1}{1 - \lambda_2}$ Theorem 2 can be derived from Lemma 13.22 in [7]. For completeness sake, we give a full proof for the general case in Appendix B.

2.2 Proof technique for lower bounds

To derive our lower bounds we will use the Bottleneck Ratio Theorem (see Theorem 19 in Appendix A) and a refinement of it for the logit dynamics of potential games (see Theorem 3 below).

Let $\mathbf{x} \in S$ be a profile of a potential game and let $M \subseteq S \setminus \{\mathbf{x}\}$ be a set of profiles different from \mathbf{x} . We define $R_{\mathbf{x},M}$ as the set of profiles in the connected component of the Hamming graph¹ with vertex set $S \setminus M$ that contains \mathbf{x} and define

$$\partial R_{\mathbf{x},M} := \{\mathbf{y} \in R_{\mathbf{x},M} : \mathbf{y} \sim \mathbf{z} \text{ for some } \mathbf{z} \in M\}.$$

In other words, $\partial R_{\mathbf{x},M}$ consists exactly of those profiles in $R_{\mathbf{x},M}$ that have a neighbor in M . We have the following theorem.

THEOREM 3. *For any potential game \mathcal{G} in which each player has exactly 2 strategies, for any profile $\mathbf{x} \in S$ and for any $M \subset S \setminus \{\mathbf{x}\}$, if $R = R_{\mathbf{x},M}$ satisfies $\pi(R) \leq 1/2$ then the mixing time of the logit dynamics with inverse noise β for \mathcal{G} satisfies*

$$t_{\text{mix}} = \Omega\left(\frac{e^{\beta(\Phi^R - \Phi^M)}}{|\partial R|}\right),$$

where Φ^R and Φ^M are the maximum potential among profiles in R and M , respectively.

PROOF. Observe that for every pair \mathbf{y}, \mathbf{z} of adjacent profiles it holds that

$$\pi(\mathbf{y})P(\mathbf{y}, \mathbf{z}) = \frac{e^{\beta\Phi(\mathbf{y})}}{Z} \cdot \frac{1}{n} \cdot \frac{e^{\beta\Phi(\mathbf{z})}}{e^{\beta\Phi(\mathbf{y})} + e^{\beta\Phi(\mathbf{z})}} \leq \frac{e^{\beta\Phi(\mathbf{z})}}{nZ}.$$

Note that for every $\mathbf{y} \in \partial R$ there are at most n neighbors outside R and all of them belong to M by definition, thus

$$\begin{aligned} Q(R, \bar{R}) &= \sum_{\substack{\mathbf{y} \in R \\ \mathbf{z} \in \bar{R}}} \pi(\mathbf{y})P(\mathbf{y}, \mathbf{z}) = \sum_{\substack{\mathbf{y} \in \partial R \\ \mathbf{z} \in M}} \pi(\mathbf{y})P(\mathbf{y}, \mathbf{z}) \\ &\leq \sum_{\substack{\mathbf{y} \in \partial R \\ \mathbf{z} \in M}} \frac{e^{\beta\Phi(\mathbf{z})}}{nZ} \leq |\partial R| \frac{e^{\beta\Phi^M}}{Z}. \end{aligned}$$

Let $\mathbf{x}^+ \in R$ be a profile with the highest potential in R ; that is, $\Phi(\mathbf{x}^+) = \Phi^R$. Obviously

$$\pi(R) \geq \pi(\mathbf{x}^+) = \frac{e^{\beta\Phi^R}}{Z}.$$

These two inequalities yield

$$\frac{Q(R, \bar{R})}{\pi(R)} \leq \frac{|\partial R|}{e^{\beta(\Phi^R - \Phi^M)}}$$

and since $\pi(R) \leq 1/2$ the thesis follows from the Bottleneck Ratio Theorem (Theorem 19 in the appendix). \square

The above theorem gives good lower bounds when we choose \mathbf{x} and M such that all profiles in M have low potential, the resulting set $R = R_{\mathbf{x},M}$ contains at least one profile of high potential (and thus $\Phi^R - \Phi^M$ is large) and the boundary ∂R is small.

¹In the Hamming graph with vertex set $S' \subseteq S$, two profiles \mathbf{x} and \mathbf{y} are adjacent if and only if they differ in exactly one component.

3. POTENTIAL GAMES

For a function $\Phi : S \rightarrow \mathbb{R}$ over a finite set S , let us name $\Delta\Phi$ the difference between the maximum and minimum values of Φ and L its Lipschitz constant, i.e.

$$\begin{aligned} \Delta\Phi &= \Phi_{\max} - \Phi_{\min} = \max\{\Phi(\mathbf{x}) - \Phi(\mathbf{y}) : \mathbf{x}, \mathbf{y} \in S\} \\ L &= \max\{\Phi(\mathbf{x}) - \Phi(\mathbf{y}) : \mathbf{x}, \mathbf{y} \in S, \mathbf{x} \sim \mathbf{y}\}. \end{aligned}$$

In this section we shall see that it is possible to give upper bounds on the mixing time of the logit dynamics for potential games depending only on those two quantities. Moreover we will show that such bounds are nearly tight by providing examples of games whose logit dynamics mixing time is close to the given upper bound.

Upper bound. In order to give the upper bound on the mixing time, we first give an upper bound on the relaxation time and then use Theorem 1.

In the proof of Theorem 5 we obtain the upper bound on the relaxation time by comparing the logit dynamics with inverse noise β for a potential game \mathcal{G} and the logit dynamics with inverse noise 0 for the same game. When the inverse noise is zero, the logit dynamics is a random walk on a generalized hypercube. Next lemma evaluates the relaxation time of such a chain. The proof is a simple generalization of the proof for the relaxation time of the lazy random walk on the hypercube and is omitted.

LEMMA 4. *For every n -player game the relaxation time of the logit dynamics with inverse noise $\beta = 0$ is $t_{\text{rel}} = n$.*

The following theorem is the main result of this section.

THEOREM 5. *Let \mathcal{G} be a n -player potential game with potential function Φ . The relaxation time of the logit dynamics for \mathcal{G} with inverse noise β is $t_{\text{rel}} = \mathcal{O}\left(n \cdot e^{\beta(\Delta\Phi + L)}\right)$.*

PROOF. Remember that the stationary distribution is

$$\pi_{\beta}(\mathbf{x}) = \frac{e^{\beta\Phi(\mathbf{x})}}{Z_{\beta}} \leq \frac{e^{\beta\Phi_{\max}}}{Z_{\beta}} \quad \text{for all profiles } \mathbf{x} \in S$$

where $Z_{\beta} = \sum_{\mathbf{y} \in S} e^{\beta\Phi(\mathbf{y})}$ is the partition function. As for the edge-stationary distribution, for two adjacent profiles $\mathbf{x} \sim \mathbf{y}$ that differ at player $i \in [n]$ we have

$$Q_{\beta}(\mathbf{x}, \mathbf{y}) = \frac{e^{\beta\Phi(\mathbf{x})}}{Z_{\beta}} \frac{1}{n} \frac{e^{\beta\Phi(\mathbf{y})}}{T_i(\mathbf{x})} \geq \frac{e^{\beta\Phi_{\min}}}{Z_{\beta}} \frac{1}{n} \frac{1}{|S_i| \cdot e^{\beta L}}, \quad (6)$$

where we used that

$$\begin{aligned} \frac{e^{\beta\Phi(\mathbf{y})}}{T_i(\mathbf{x})} &= \frac{e^{\beta\Phi(\mathbf{y})}}{\sum_{z \in S_i} e^{\beta\Phi(\mathbf{x}_{-i}, z)}} \\ &= \frac{1}{\sum_{z \in S_i} e^{\beta[\Phi(\mathbf{x}_{-i}, z) - \Phi(\mathbf{y})]}} \geq \frac{1}{|S_i| \cdot e^{\beta L}}. \end{aligned}$$

Moreover, for any profile \mathbf{x} we have

$$Q_{\beta}(\mathbf{x}, \mathbf{x}) = \frac{e^{\beta\Phi(\mathbf{x})}}{Z_{\beta}} \frac{1}{n} \sum_{i=1}^n \frac{e^{\beta\Phi(\mathbf{x})}}{T_i(\mathbf{x})} \geq \frac{e^{\beta\Phi_{\min}}}{Z_{\beta}} \frac{1}{n} \frac{1}{e^{\beta L}} \sum_{i=1}^n \frac{1}{|S_i|}.$$

Hence, for all $\mathbf{x}, \mathbf{y} \in S$ it holds that

$$\begin{aligned} \pi_{\beta}(\mathbf{x}) &\leq \frac{Z_0}{Z_{\beta}} e^{\beta\Phi_{\max}} \pi_0(\mathbf{x}) \\ Q_{\beta}(\mathbf{x}, \mathbf{y}) &\geq \frac{Z_0}{Z_{\beta}} \frac{e^{\beta\Phi_{\min}}}{e^{\beta L}} Q_0(\mathbf{x}, \mathbf{y}). \end{aligned}$$

Since from Lemma 4 it holds that for $\beta = 0$ the relaxation time is $\mathcal{O}(n)$, the thesis follows by applying the comparison theorem (Theorem 2) with

$$\alpha = \frac{Z_\beta}{Z_0} \frac{e^{\beta L}}{e^{\beta \Phi_{\min}}} \quad \text{and} \quad \gamma = \frac{Z_0}{Z_\beta} e^{\beta \Phi_{\max}}.$$

□

A slightly better upper bound holds when the players have two strategies.

COROLLARY 6. *If every player has only two strategies then the relaxation time is $t_{\text{rel}} = \mathcal{O}(n \cdot e^{\beta \Delta \Phi})$.*

PROOF. Observe that, when every player has two strategies, in Equation (6) we have that

$$\frac{e^{\beta \Phi(\mathbf{x})} e^{\beta \Phi(\mathbf{y})}}{T_i(\mathbf{x})} = \frac{e^{\beta \Phi(\mathbf{x})} e^{\beta \Phi(\mathbf{y})}}{e^{\beta \Phi(\mathbf{x})} + e^{\beta \Phi(\mathbf{y})}} \geq \frac{e^{\beta \min\{\Phi(\mathbf{x}), \Phi(\mathbf{y})\}}}{2}.$$

Hence, we obtain

$$Q_\beta(\mathbf{x}, \mathbf{y}) \geq \frac{e^{\beta \Phi_{\min}}}{Z_\beta} \frac{1}{2n} \quad \text{and} \quad Q_\beta(\mathbf{x}, \mathbf{x}) \geq \frac{e^{\beta \Phi_{\min}}}{Z_\beta} \frac{1}{2}$$

and we can apply the comparison theorem with

$$\alpha = \frac{Z_\beta}{Z_0} \frac{1}{e^{\beta \Phi_{\min}}} \quad \text{and} \quad \gamma = \frac{Z_0}{Z_\beta} e^{\beta \Phi_{\max}}.$$

□

Finally, we can obtain the bounds on the mixing time by using Theorem 1 and the fact that $\pi_{\min} \geq 1/(e^{\beta \Delta \Phi} |S|)$.

COROLLARY 7. *For every potential game the mixing time of the logit dynamics is*

$$t_{\text{mix}} = \mathcal{O}\left(n \cdot e^{\beta(\Delta \Phi + L)} (\beta \Delta \Phi + \log |S|)\right),$$

where S is the set of strategy profiles.

For potential games with two strategies per player the mixing time is $\mathcal{O}(n \cdot e^{\beta \Delta \Phi} (\beta \Delta \Phi + \log |S|))$.

Lower bound. It is easy to find potential games whose logit dynamics mixing time is $\Omega(e^{\beta \Delta \Phi})$ when $\Delta \Phi = L$; e.g., games whose potential function Φ has only two values and at least two non-adjacent maxima (see, for example, 2-player coordination games studied in [2]). One naturally wonders whether a similar lower bound can be achieved for games where the Lipschitz constant L is small compared to $\Delta \Phi$. The following theorem shows that the term $e^{\beta \Delta \Phi}$ in the upper bound in Corollary 7 cannot be essentially improved for L smaller than $\Delta \Phi$.

THEOREM 8. *For every $0 < \delta < 1$ and for every $L = \omega(\log n)$ a family of potential games with two strategies per player exists such that the potential function Φ has Lipschitz constant L , it satisfies $\Delta \Phi/L > n^\delta$ and the mixing time of the logit dynamics is $\Omega(e^{(\beta - o(1)) \Delta \Phi})$.*

PROOF. Consider the game with n players in which every player has strategies 0 and 1, and whose potential function is

$$\Phi(\mathbf{x}) = \Phi(|\mathbf{x}|_1) = \min\{c - |\mathbf{x}|_1, 0\} \cdot L$$

where $c = \lceil n^\delta \rceil$. Note that the maximum of the potential is $\Phi(\mathbf{0}) = \Delta \Phi = cL$, while the minimum is zero and is attained at all states in the set $M = \{\mathbf{x} \in S : |\mathbf{x}|_1 = c\}$.

Consider the set $R_{\mathbf{0}, M}$ (see Section 2) and observe that

$$\begin{aligned} R_{\mathbf{0}, M} &= \{\mathbf{x} \in S : |\mathbf{x}|_1 < c\} \\ \partial R_{\mathbf{0}, M} &= \{\mathbf{x} \in S : |\mathbf{x}|_1 = c - 1\}. \end{aligned}$$

By the symmetry of the potential function, the stationary probability of $R_{\mathbf{0}, M}$ is $\pi(R_{\mathbf{0}, M}) \leq \frac{1}{2}$ and the size of its boundary is

$$|\partial R_{\mathbf{0}, M}| \leq \binom{n}{c} \leq e^{c \log n} = e^{(\Delta \Phi / L) \log n}.$$

Thus, from Theorem 3 we have that the mixing time of the logit dynamics is

$$t_{\text{mix}} = \Omega\left(e^{\beta \Delta \Phi - (\Delta \Phi / L) \log n}\right)$$

and since $L = \omega(\log n)$ the thesis follows. □

4. DOMINANT STRATEGIES GAMES

In the previous section, we analyzed potential games and derived upper and lower bounds on the mixing time for the logit dynamics that are exponential in β . In this section we prove that, for the class of games with *dominant strategies*, it is possible to give upper bounds that are independent of β . In other words, the mixing time of the logit dynamics for games with dominant strategies does not grow arbitrarily as β tends to infinity.

A strategy $z \in S_i$ is *dominant* for player i if it yields the maximum payoff regardless of the strategies of the other players; that is, $u_i(z, \mathbf{x}_{-i}) \geq u_i(z', \mathbf{x}_{-i})$ for every $z' \in S_i$ and every $\mathbf{x}_{-i} \in S_{-i}$. In a game with *dominant strategies* every player has a dominant strategy. Let us name 0 a dominant strategy for all players and consider the profile $\mathbf{0} = (0, \dots, 0)$. It is easy to see that the following observation holds for the logit dynamics of a game with dominant strategies.

OBSERVATION 9. *In every profile and for every β , if player i is selected then her strategy is updated to the dominant strategy with probability at least $1/|S_i|$. That is, for all \mathbf{x} , β and i , $\sigma_i(0 | \mathbf{x}) \geq 1/|S_i|$.*

We are now ready to derive an upper bound on the mixing time of the logit dynamics for dominant strategy games.

THEOREM 10. *For n -player games with dominant strategies where each player has at most m strategies, the mixing time is $t_{\text{mix}} = \mathcal{O}(m^n n \log n)$.*

PROOF. We apply the coupling technique (see Theorem 17 in Appendix A). Let $\{X_t\}$ and $\{Y_t\}$ be two instances of the logit dynamics starting at \mathbf{x} and \mathbf{y} respectively, and consider a coupling with the following properties: at every step the same player in both chains is chosen for the update, the probability that the strategy of the chosen player is updated to 0 in both chains is at least $1/|S_i| \geq 1/m$ (notice that this is possible because of Observation 9), and once the two chains coalesce they stay coupled for all the following time steps. An example of such a coupling can be found in Appendix C.

Let τ be the first time such that all the players have been selected at least once and let $t^* = 2n \log n$. Observe that for all starting profiles \mathbf{z} and \mathbf{w} , it holds that

$$\mathbf{P}_{\mathbf{z},\mathbf{w}}(X_{t^*} = \mathbf{0} \text{ and } Y_{t^*} = \mathbf{0} \mid \tau \leq t^*) \geq \frac{1}{m^n}. \quad (7)$$

Indeed, given that all players have been selected at least once within time t^* , both chains are in profile $\mathbf{0}$ at time t^* if and only if every player chose strategy 0 in both chains the last time she played before time t^* . From the construction of the coupling it follows that such event holds with probability at least $1/m^n$.

Hence, for all starting profiles \mathbf{z} and \mathbf{w} , we have that

$$\begin{aligned} \mathbf{P}_{\mathbf{z},\mathbf{w}}(X_{t^*} = Y_{t^*}) &\geq \mathbf{P}_{\mathbf{z},\mathbf{w}}(X_{t^*} = \mathbf{0} \text{ and } Y_{t^*} = \mathbf{0}) \\ &\geq \mathbf{P}_{\mathbf{z},\mathbf{w}}(X_{t^*} = \mathbf{0} \text{ and } Y_{t^*} = \mathbf{0} \mid \tau \leq t^*) \mathbf{P}_{\mathbf{z},\mathbf{w}}(\tau \leq t^*) \\ &\geq \frac{1}{m^n} \cdot \frac{1}{2} \end{aligned} \quad (8)$$

where in the last inequality we used (7) and the Coupon Collector's argument.

Therefore, by considering k phases each one lasting t^* time steps, since the bound in (8) holds for every starting states of the Markov chain, we have that the probability that the two chains have not yet coupled after kt^* time steps is

$$\mathbf{P}_{\mathbf{x},\mathbf{y}}(X_{kt^*} \neq Y_{kt^*}) \leq \left(1 - \frac{1}{2m^n}\right)^k \leq e^{-k/2m^n}$$

which is less than $1/4$, for $k = \mathcal{O}(m^n)$. By applying the Coupling Theorem (see Theorem 17 in the Appendix) we have that $t_{\text{mix}} = \mathcal{O}(m^n n \log n)$. \square

In [2] a n -player game with two strategies per player is shown whose logit dynamics mixing time is $\Omega(2^n)$ for large values of β . We next prove that, for every $m \geq 2$, there are n -player games with m strategies per player whose logit dynamics mixing time is $\Omega(m^{n-1})$. Thus the m^n factor in the upper bound given by Theorem 10 cannot be essentially improved.

THEOREM 11. *For every $m \geq 2$ and $n \geq 2$, there exists a n -player potential game with dominant strategies where each player has m strategies and such that, for sufficiently large β , $t_{\text{mix}} = \Omega(m^{n-1})$.*

PROOF. Consider the game with n players, each of them having strategies $\{0, \dots, m-1\}$, such that for every player i :

$$u_i(\mathbf{x}) = \begin{cases} 0, & \text{if } \mathbf{x} = \mathbf{0}; \\ -1, & \text{otherwise.} \end{cases}$$

Note that 0 is a dominant strategy. This is a potential game with potential $\Phi(\mathbf{x}) = u_i(\mathbf{x})$ and thus the stationary distribution is given by the Gibbs measure. We apply the bottleneck ratio (see Theorem 19 in Appendix A) with $R = \{0, \dots, m-1\}^n \setminus \{\mathbf{0}\}$, for which we have

$$\pi(R) = \frac{e^{-\beta}}{Z}(m^n - 1)$$

with $Z = 1 + e^{-\beta}(m^n - 1)$. It is easy to see that $\pi(R) < 1/2$ for $\beta > \log(m^n - 1)$ and furthermore

$$\begin{aligned} Q(R, \bar{R}) &= \sum_{\mathbf{x} \in R} \pi(\mathbf{x}) P(\mathbf{x}, \mathbf{0}) \\ &= \frac{e^{-\beta}}{Z} \sum_{\mathbf{x} \in R} P(\mathbf{x}, \mathbf{0}) = \frac{e^{-\beta}}{Z} \sum_{\mathbf{x} \in R_1} P(\mathbf{x}, \mathbf{0}), \end{aligned}$$

where R_1 is the subset of R containing all states with exactly one non-zero entry. Notice that, for every $\mathbf{x} \in R_1$, we have

$$P(\mathbf{x}, \mathbf{0}) = \frac{1}{n} \cdot \frac{1}{1 + (m-1)e^{-\beta}}.$$

As $|R_1| = n(m-1)$, we have

$$\begin{aligned} Q(R, \bar{R}) &= \frac{e^{-\beta}}{Z} |R_1| \frac{1}{n} \cdot \frac{1}{1 + (m-1)e^{-\beta}} \\ &= \frac{e^{-\beta}}{Z} \cdot \frac{m-1}{1 + (m-1)e^{-\beta}} \end{aligned}$$

whence

$$\begin{aligned} t_{\text{mix}} &\geq \frac{1}{4} \cdot \frac{\pi(R)}{Q(R, \bar{R})} \\ &\geq \frac{1}{4} \cdot (m^n - 1) \cdot \frac{1 + (m-1)e^{-\beta}}{m-1} > \frac{1}{4} \cdot \frac{m^n - 1}{m-1}. \end{aligned}$$

\square

Extensions. Observe that, by using the same techniques exploited in this section, it is possible to prove an upper bound *independent of β* for *max-solvable* games [12], a class which contains games with dominant strategies as a special case, albeit with an upper bound that is much larger than $\mathcal{O}(m^n n \log n)$.

5. GRAPHICAL COORDINATION GAMES

Consider the following basic two-player coordination game

$$\begin{array}{c|cc} & 0 & 1 \\ \hline 0 & a, a & c, d \\ \hline 1 & d, c & b, b \end{array} \quad (9)$$

We assume that $a > d$ and $b > c$ which implies that players have an advantage in selecting the same strategy and that $(0, 0)$ and $(1, 1)$ are Nash equilibria. If $a - d > b - c$ then equilibrium $(0, 0)$ is said to be *risk dominant* and, analogously, if $a - d < b - c$ then equilibrium $(1, 1)$ is said to be *risk dominant* [6]. Tight bounds for the mixing time of the basic coordination games have been given in [2].

In this section we consider *graphical coordination games* in which n players are connected by a network G (encoding, for example, social relationships) and every player plays the basic coordination game (9) with each of the adjacent players. Specifically, when a player selects her strategy, such a strategy is played against each one of her adjacent players. The payoff of a player is given by the sum of the payoffs gained from each instance of the basic coordination game. We focus on two network topologies: the clique (Section 5.1), where the mixing time dependence on $e^{\beta \Delta \Phi}$ showed in Corollary 7 cannot be improved, and the ring (Section 5.2), where a more local interaction implies a faster convergence to the stationary distribution.

In the rest of this section we will assume w.l.o.g. that $a - d \geq b - c$.

5.1 Coordination games on the clique

We study the mixing time of graphical coordination games on the clique; that is, every player plays the basic coordination game (9) with every other player. We give upper and lower bounds on the mixing time. As we shall see, both such bounds turn out to be exponential in n , even for $\beta = \Theta(1)$.

We first observe that the game is a potential game. This will allow us to use Corollary 7 to derive an upper bound on the mixing time and to use Theorem 3 to get a lower bound.

It is not difficult to see that $\Phi(\mathbf{x}) = \phi(|\mathbf{x}|_0)$ is a potential function for the graphical coordination game on the clique, where

$$\phi(k) := (k^* - k) \left(\frac{2n - k^* - k - 1}{2} (b - c) - \frac{k^* + k - 1}{2} (a - d) \right)$$

and $k^* = \left\lceil (n - 1) \frac{b - c}{(a - d) + (b - c)} \right\rceil$. Notice that the minimum of the potential is attained when k^* players are playing 0 and, since $\phi(k^*) = 0$, we have that $\Delta\Phi = \max_k \phi(k)$. Moreover, it is easy to check that $\phi(k)$ monotonically decreases as k goes from 0 to k^* and then monotonically increases as k goes from k^* to n . Therefore, $\Delta\Phi = \max\{\phi(0), \phi(n)\}$.

Notice that, since $a - d \geq b - c$, then $\phi(k) \leq \phi(n - k)$ for $k < k^*$, $\Delta\Phi = \phi(n)$ and

$$\sum_{k=0}^{k^*} \Phi(k) \leq \sum_{k=n-k^*}^n \Phi(k). \quad (10)$$

Since $\Delta\Phi = \phi(n)$, by applying our general result on the mixing time of the logit dynamics of potential games (see Corollary 7) we get $t_{\text{mix}} = \mathcal{O}\left(n \cdot e^{\beta\phi(n)} \cdot (\beta\phi(n) + n)\right)$. We next state a lower bound on the mixing time for coordination games on a clique.

LEMMA 12. *For coordination games on a clique the mixing time is $t_{\text{mix}} = \Omega\left(e^{(\beta - o(1))\phi(0)}\right)$.*

PROOF. We obtain our lower bound by applying Theorem 3 with configuration $\mathbf{x}^* = (1, \dots, 1)$ and set $M = \{\mathbf{x} \in S : |\mathbf{x}|_0 = k^*\}$.

The connected component R of $S \setminus M$ that contains \mathbf{x}^* is

$$R = \{\mathbf{x} \in S : |\mathbf{x}|_0 < k^*\}$$

From (10) it follows that $\pi(R) \leq \frac{1}{2}$. Finally, notice that

$$\begin{aligned} |\partial R| &\leq |\{\mathbf{x} \in \Omega : |\mathbf{x}|_0 = k^* - 1\}| \\ &= \binom{n}{k^* - 1} \leq n^{k^*} \leq n^{\frac{2}{b-c} \frac{\phi(0)}{n-1}}. \end{aligned}$$

The lemma follows by applying Theorem 3 and by observing that the maximum potential among profiles in R and M are $\Phi^R = \phi(0)$ and $\Phi^M = 0$, respectively. \square

We stress that when the basic coordination game has no risk dominant strategy (that is the case $a - d = b - c$), $\phi(0) = \phi(n)$ and thus the exponents of the upper and lower bound coincide up to a $o(1)$ term. In general, by observing that $\phi(0), \phi(n) = \Theta(n^2)$, we can say that the mixing time is exponential in n^2 and β . More precisely, we obtain the following theorem.

THEOREM 13. *For every graphical coordination game on a clique there exist two constants C and D such that $C^{\beta n^2} \leq t_{\text{mix}} \leq D^{\beta n^2}$.*

5.2 Coordination games on the ring

In this section we give upper and lower bounds on the mixing time for graphical coordination games on the ring when there is no risk dominant strategy. Unlike the clique, the ring encodes a very local type of interaction between the players which is more likely to occur in a social context. Our results show that the mixing time is polynomial in the number of players n and e^β .

Let us name $\delta := a - d = b - c$. It is not difficult to see that $\Phi(\mathbf{x}) = \sum_{i=1}^n \Phi_i(\mathbf{x})$ is a potential for the coordination game on the ring, where

$$\Phi_i(\mathbf{x}) = \begin{cases} \delta, & \text{if } x_{i-1} = x_i = x_{i+1}; \\ \frac{\delta}{2}, & \text{if } x_{i-1} \neq x_{i+1}; \\ 0, & \text{if } x_i \neq x_{i-1} = x_{i+1}. \end{cases}$$

Observe that $\Phi(\mathbf{1}) = \Phi(\mathbf{0}) = n\delta$. Moreover, if n is even, the configuration \mathbf{x} where every player selects a strategy different from the one selected by her neighbors has potential $\Phi(\mathbf{x}) = 0$: thus, there are graphical coordination games on the ring where $\Delta\Phi = n\delta$. If we used Corollary 7, we would get an upper bound exponential in n . Instead we here show an upper bound that is polynomial in n .

The proof of the upper bound uses the path coupling technique (see Theorem 18) and can be seen as a generalization of the upper bound on the mixing time for the Ising model on the ring (see Chapter 15 of [7]).

THEOREM 14. *For graphical coordination games with no risk-dominant strategy ($a - d = b - c = \delta$) on a ring with n players the mixing time is $t_{\text{mix}} = \mathcal{O}(n \log n \cdot e^{2\beta\delta})$.*

PROOF. We identify the n players with the integers in $\{0, \dots, n-1\}$ and assume that every player i plays the basic coordination game with her two adjacent players, $(i-1) \bmod n$ and $(i+1) \bmod n$. Let $S = \{0, 1\}^n$ be the set of profiles for n players playing the graphical coordination game on the ring and consider the Hamming graph G over S where profiles \mathbf{x} and \mathbf{y} are adjacent if and only if they differ in exactly one position.

Let us consider two adjacent configurations \mathbf{x} and \mathbf{y} . Denote by j the position in which they differ and assume, without loss of generality, that $x_j = 1$ and $y_j = 0$. We consider the following coupling for two chains X and Y starting respectively from $X_0 = \mathbf{x}$ and $Y_0 = \mathbf{y}$: Pick $i \in \{0, \dots, n-1\}$ and $U \in [0, 1]$ independently and uniformly at random and update position i of \mathbf{x} and \mathbf{y} by setting

$$x_i = \begin{cases} 0, & \text{if } U \leq \sigma_i(0 | \mathbf{x}); \\ 1, & \text{if } U > \sigma_i(0 | \mathbf{x}); \end{cases} \quad y_i = \begin{cases} 0, & \text{if } U \leq \sigma_i(0 | \mathbf{y}); \\ 1, & \text{if } U > \sigma_i(0 | \mathbf{y}). \end{cases}$$

We next compute the expected distance between X_1 and Y_1 after one step of the coupling. Notice that $\sigma_i(0 | \mathbf{x})$ only depends on x_{i-1} and x_{i+1} and $\sigma_i(0 | \mathbf{y})$ only on y_{i-1} and y_{i+1} . Therefore, since \mathbf{x} and \mathbf{y} only differ at position j , $\sigma_i(0 | \mathbf{x}) = \sigma_i(0 | \mathbf{y})$ for $i \neq j-1, j+1$.

We start by observing that if position j is chosen for update (this happens with probability $1/n$) then, by the observation above, both chains perform the same update. Since \mathbf{x} and \mathbf{y} differ only for player j , we have that the two chains are coupled (and thus at distance 0). Similarly, if $i \neq j-1, j, j+1$ (which happens with probability $(n-3)/n$) we have that both chains perform the same update and thus remain at distance 1. Finally, let us consider the case in

which $i \in \{j-1, j+1\}$. In this case, since $x_j = 1$ and $y_j = 0$, we have that $\sigma_i(0|\mathbf{x}) \leq \sigma_i(0|\mathbf{y})$. Therefore, with probability $\sigma_i(0|\mathbf{x})$ both chains update position i to 0 and thus remain at distance 1; with probability $1 - \sigma_i(0|\mathbf{y})$ both chains update position i to 1 and thus remain at distance 1; and with probability $\sigma_i(0|\mathbf{y}) - \sigma_i(0|\mathbf{x})$ chain X updates position i to 1 and chain Y updates position i to 0 and thus the two chains go to distance 2. By summing up, we have that the expected distance $E[\rho(X_1, Y_1)]$ after one step of coupling of the two chains is

$$\begin{aligned} E[\rho(X_1, Y_1)] &= \\ &= \frac{n-3}{n} + \frac{1}{n} \sum_{i \in \{j-1, j+1\}} [\sigma_i(0|\mathbf{x}) + 1 - \sigma_i(0|\mathbf{y}) \\ &\quad + 2 \cdot (\sigma_i(0|\mathbf{y}) - \sigma_i(0|\mathbf{x}))] \\ &= \frac{n-3}{n} + \frac{1}{n} \cdot \sum_{i \in \{j-1, j+1\}} (1 + \sigma_i(0|\mathbf{y}) - \sigma_i(0|\mathbf{x})) \\ &= \frac{n-1}{n} + \frac{1}{n} \cdot \sum_{i \in \{j-1, j+1\}} (\sigma_i(0|\mathbf{y}) - \sigma_i(0|\mathbf{x})) \end{aligned}$$

Let us now evaluate the difference $\sigma_i(0|\mathbf{y}) - \sigma_i(0|\mathbf{x})$ for $i = j-1$ (the same computation holds for $i = j+1$). We distinguish two cases depending on the strategies of player $j-2$ and start with the case $x_{j-2} = y_{j-2} = 1$. In this case we have that

$$\sigma_{j-1}(0|\mathbf{x}) = \frac{1}{1+e^{2\beta\delta}} \quad \text{and} \quad \sigma_{j-1}(0|\mathbf{y}) = \frac{1}{2}.$$

Thus,

$$\sigma_{j-1}(0|\mathbf{y}) - \sigma_{j-1}(0|\mathbf{x}) = \frac{1}{2} - \frac{1}{1+e^{2\beta\delta}}.$$

If instead $x_{j-2} = y_{j-2} = 0$, we have

$$\sigma_{j-1}(0|\mathbf{x}) = \frac{1}{2} \quad \text{and} \quad \sigma_{j-1}(0|\mathbf{y}) = \frac{1}{1+e^{-2\beta\delta}}.$$

Thus

$$\begin{aligned} \sigma_{j-1}(0|\mathbf{y}) - \sigma_{j-1}(0|\mathbf{x}) &= \frac{1}{1+e^{-2\beta\delta}} - \frac{1}{2} \\ &= 1 - \frac{1}{1+e^{2\beta\delta}} - \frac{1}{2} = \frac{1}{2} - \frac{1}{1+e^{2\beta\delta}}. \end{aligned}$$

We can conclude that the expected distance after one step of the chain is

$$\begin{aligned} E[\rho(X_1, Y_1)] &= \frac{n-1}{n} + \frac{1}{n} \left(1 - \frac{2}{1+e^{2\beta\delta}} \right) \\ &= 1 - \frac{2}{n(1+e^{2\beta\delta})} \leq e^{-\frac{2}{n(1+e^{2\beta\delta})}}. \end{aligned}$$

Since the diameter of G is $\text{diam}(G) = n$, by applying Theorem 18 with $\alpha = \frac{2}{n(1+e^{2\beta\delta})}$, we obtain the theorem. \square

The upper bound in Theorem 14 is nearly tight (up to the $n \log n$ factor). Indeed, a lower bound can be obtained by applying the Bottleneck Ratio technique (see Theorem 19 in Appendix A) to the set $R = \{\mathbf{1}\}$. Notice that $\pi(R) \leq \frac{1}{2}$ since profile $\mathbf{0}$ has the same potential as $\mathbf{1}$. Thus set R satisfies

the hypothesis of Theorem 19. Simple computations show that the bottleneck ratio is

$$B(R) = \sum_{\mathbf{y} \neq \mathbf{1}} P(\mathbf{1}, \mathbf{y}) = \frac{1}{1+e^{2\beta\delta}}.$$

Thus, by applying Theorem 19, we obtain the following bound.

THEOREM 15. *For graphical coordination games with no risk-dominant strategy on a ring with n players the mixing time is $t_{\text{mix}} = \Omega(e^{2\beta\delta})$.*

6. CONCLUSIONS AND OPEN PROBLEMS

In this work we give bounds on the mixing time of the logit dynamics for wide classes of games, highlighting how the noise level of the logit dynamics and the structural properties of the game affect the convergence rate to stationarity. In fact, we show that the mixing time for potential games depends polynomially on the number of players and exponentially on the inverse noise and the maximum potential difference $\Delta\Phi$: this dependence shows both in the upper and the lower bound, even if they are not completely matching; thus, it is natural to ask if it is possible to close the gap.

On the other hand, we show that there exists a class of games, namely dominant strategy games, such that the mixing time of the logit dynamics does not grow indefinitely with the inverse noise.

Finally, we consider coordination games on the clique and the ring, a subset of potential games, where we give evidence that the mixing time is affected also by other structural properties as the connectedness of the network: it might be interesting to investigate other graph structures to highlight other properties influencing the mixing time (e.g., degree of the graph, expansion, etc.)

The main goal of this line of research is to give general bounds on the logit dynamics mixing time for any game, highlighting the features of the game that distinguish between polynomial and exponential mixing time. We stress that, when the game is not a potential game, in general there is not a simple closed form for the stationary distribution like Equation (3).

At every step of the logit dynamics one single player is selected to update her strategy. It would be interesting to consider variations of such dynamics where players are allowed to update their strategies simultaneously. The special case of parallel best response (that is $\beta = \infty$) has been studied in [12]. Another interesting variant of the logit dynamics is the one in which the value of β is not fixed, but varies according to some *learning process* by which players acquire more information on the game as time progresses.

When the mixing time of the logit dynamics is polynomial, we know that the stationary distribution gives good predictions of the state of the system after a polynomial number of time steps. When the mixing time is exponential, it would be interesting to analyze the *transient* phase of the logit dynamics, in order to investigate what kind of predictions can be made about the state of the system in such a phase.

7. REFERENCES

- [1] Arash Asadpour and Amin Saberi. On the inefficiency ratio of stable equilibria in congestion games. In *Proc. of the 5th International Workshop on Internet and*

Network Economics (WINE'09), volume 5929 of *LNCS*, pages 545–552. Springer, 2009.

- [2] Vincenzo Auletta, Diodato Ferraioli, Francesco Pasquale, and Giuseppe Persiano. Mixing time and stationary expected social welfare of logit dynamics. In *3rd Int. Symp. on Algorithmic Game Theory (SAGT)*, volume 6386 of *LNCS*, pages 54–65. Springer, 2010.
- [3] 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.
- [4] Lawrence E. Blume. The statistical mechanics of strategic interaction. *Games and Economic Behavior*, 5:387–424, 1993.
- [5] Glenn Ellison. Learning, local interaction, and coordination. *Econometrica*, 61(5):1047–1071, 1993.
- [6] John C. Harsanyi and Reinhard Selten. *A General Theory of Equilibrium Selection in Games*. MIT Press, 1988.
- [7] David Levin, Yuval Peres, and Elizabeth L. Wilmer. *Markov Chains and Mixing Times*. American Mathematical Society, 2008.
- [8] Jason R. Marden, H. Peyton Young, Gürdal Arslan, and Jeff S. Shamma. Payoff-based dynamics for multiplayer weakly acyclic games. *SIAM Journal on Control and Optimization*, 48(1):373–396, 2009.
- [9] 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.
- [10] Panayotis Mertikopoulos and Aris L. Moustakas. The emergence of rational behavior in the presence of stochastic perturbations. *The Annals of Applied Probability*, 20(4):1359–1388, 2010.
- [11] Andrea Montanari and Amin Saberi. Convergence to equilibrium in local interaction games. In *Proc. of the 50th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 303–312. IEEE, 2009.
- [12] Noam Nisan, Michael Schapira, and Aviv Zohar. Asynchronous best-reply dynamics. In *4th International Workshop on Internet and Network Economics (WINE)*, volume 5385 of *LNCS*, pages 531–538. Springer, 2008.
- [13] H. Peyton Young. Adaptive learning in systems of interacting agents. In *Internet and Network Economics*, volume 5929 of *LNCS*, pages 13–16. Springer Berlin / Heidelberg, 2009.
- [14] H. Peyton Young. *The diffusion of innovations in social networks*, chapter in “The Economy as a Complex Evolving System”, vol. III, Lawrence E. Blume and Steven N. Durlauf, eds. Oxford University Press, 2003.

APPENDIX

A. MARKOV CHAINS’ SUMMARY

We summarize the main tools we use to bound the mixing time of Markov chains (for a complete description of such tools see, for example, Chapters 4.2, 5.2, 7.2 and 14.2 in [7]).

Definition 16. A coupling of two probability distributions μ and ν is a pair of random variables (X, Y) defined on the

same probability space such that the marginal distribution of X is μ and the marginal distribution of Y is ν . That is, a coupling (X, Y) of μ and ν satisfies $\mathbf{P}(X = x) = \mu(x)$ and $\mathbf{P}(Y = y) = \nu(y)$.

THEOREM 17 (COUPLING). Let \mathcal{M} be an irreducible and aperiodic Markov chain with finite state space Ω and transition matrix P . For two states x and y , let (X_t, Y_t) be a coupling of $P^t(x, \cdot)$ and $P^t(y, \cdot)$ such that $X_s = Y_s$ implies $X_t = Y_t$ for $t \geq s$. Let τ_{couple} be the first time the chains meet, $\tau_{\text{couple}} := \min\{t : X_t = Y_t\}$. Then, the mixing time is

$$t_{\text{mix}}(\varepsilon) \leq \min \left\{ t \in \mathbb{N} : \max_{x, y \in \Omega} \mathbf{P}_{x, y}(\tau_{\text{couple}} > t) \leq \varepsilon \right\}.$$

THEOREM 18 (PATH COUPLING). Let $\mathcal{M} = \{X_t : t \in \mathbb{N}\}$ be an irreducible and aperiodic Markov chain with finite state space Ω and transition matrix P . Let $G = (\Omega, E)$ be a connected graph, let $\ell : E \rightarrow \mathbb{R}$ be a function assigning weights to edges such that $\ell(e) \geq 1$ for every edge $e \in E$, and let $\rho : \Omega \times \Omega \rightarrow \mathbb{R}$ be the corresponding path distance, i.e. $\rho(x, y)$ is the length of the (weighted) shortest path in G between x and y .

Suppose that for every edge $\{x, y\} \in E$ a coupling (X, Y) of distributions $P(x, \cdot)$ and $P(y, \cdot)$ exists such that

$$\mathbf{E}_{x, y}[\rho(X, Y)] \leq \ell(\{x, y\})e^{-\alpha}$$

for some $\alpha > 0$, then the mixing time of \mathcal{M} is

$$t_{\text{mix}}(\varepsilon) \leq \frac{\log(\text{diam}(G)) + \log(1/\varepsilon)}{\alpha}$$

where $\text{diam}(G)$ is the (weighted) diameter of G .

THEOREM 19 (BOTTLENECK RATIO). Let $\mathcal{M} = \{X_t : t \in \mathbb{N}\}$ be an irreducible and aperiodic Markov chain with finite state space Ω , transition matrix P and stationary distribution π . Let $R \subseteq \Omega$ be any set with $\pi(R) \leq 1/2$. Then the mixing time is

$$t_{\text{mix}}(\varepsilon) \geq \frac{1 - 2\varepsilon}{2B(R)}$$

where

$$B(R) = \frac{Q(R, \bar{R})}{\pi(R)} \quad \text{and} \quad Q(R, \bar{R}) = \sum_{x \in R, y \in \bar{R}} \pi(x)P(x, y).$$

B. PROOF OF COMPARISON THEOREM

Let P be the transition matrix of an irreducible, aperiodic, and reversible Markov chain with finite state space Ω and stationary distribution π .

For a function $f : \Omega \rightarrow \mathbb{R}$ let $\mathcal{E}_\pi(f)$ be its Dirichlet form, i.e. $\mathcal{E}_\pi(f) = \langle (I - P)f, f \rangle_\pi$, where I is the identity matrix of size $|\Omega|$ and $\langle \cdot, \cdot \rangle_\pi$ is the inner product defined by

$$\langle f, g \rangle_\pi = \sum_{x \in \Omega} \pi(x)f(x)g(x) \quad \text{for } f, g : \Omega \rightarrow \mathbb{R}$$

The Dirichlet form of a function f can be written as (see Lemma 13.11 in [7])

$$\mathcal{E}_\pi(f) = \frac{1}{2} \sum_{x, y \in \Omega} Q(x, y) (f(x) - f(y))^2 \quad (11)$$

The following lemma states that by comparing the Dirichlet forms and the stationary distributions of two chains over the same state space it is possible to compare their spectral gaps.

LEMMA 20 (SEE LEMMA 13.22 IN [7]). *Let P and \hat{P} be irreducible and reversible transition matrices over the same state space Ω and stationary distributions π and $\hat{\pi}$, respectively. If $\mathcal{E}_{\hat{\pi}}(f) \leq \alpha \mathcal{E}_{\pi}(f)$ for any function f , then*

$$1 - \hat{\lambda}_2 \leq \left[\max_{x \in \Omega} \frac{\pi(x)}{\hat{\pi}(x)} \right] \alpha (1 - \lambda_2)$$

In order to compare the relaxation times of two chains, we still need to compare their last eigenvalues. To this aim, consider the form $\mathcal{E}_{\pi}^+(f) = \langle (I + P)f, f \rangle_{\pi}$.

OBSERVATION 21.

$$\mathcal{E}_{\pi}^+(f) = \frac{1}{2} \sum_{x, y \in \Omega} Q(x, y) (f(x) + f(y))^2 \quad (12)$$

PROOF.

$$\begin{aligned} & \sum_{x, y \in \Omega} Q(x, y) (f(x) + f(y))^2 = \\ &= 2 \sum_{x \in \Omega} \pi(x) f(x)^2 + 2 \sum_{x \in \Omega} \pi(x) f(x) (Pf)(x) \\ &= 2 \langle f, f \rangle_{\pi} + 2 \langle Pf, f \rangle_{\pi} = 2 \langle (I + P)f, f \rangle_{\pi} \end{aligned}$$

□

The next observation shows that, just like the Dirichlet form is related to the spectral gap, \mathcal{E}^+ is related to the smallest eigenvalue of the transition matrix.

OBSERVATION 22.

$$1 + \lambda_{|\Omega|} = \min_{f \neq 0} \frac{\mathcal{E}_{\pi}^+(f)}{\langle f, f \rangle_{\pi}}$$

PROOF. Since P is irreducible, aperiodic, and reversible there is a basis of \mathbb{R}^{Ω} formed by eigenvectors of P that are orthonormal w.r.t. the inner product $\langle \cdot, \cdot \rangle_{\pi}$ (see e.g. Lemma 12.2 in [7]). Let $f_1, \dots, f_{|\Omega|}$ be such a basis where f_i is the eigenvector with eigenvalue λ_i . Let f be any function, then it can be written as a linear combination of eigenvectors $f = \sum_{i=1}^{|\Omega|} \alpha_i f_i$. Hence $Pf = \sum_{i=1}^{|\Omega|} \alpha_i P f_i = \sum_{i=1}^{|\Omega|} \alpha_i \lambda_i f_i$. Since $f_1, \dots, f_{|\Omega|}$ are orthonormal w.r.t. $\langle \cdot, \cdot \rangle_{\pi}$ it holds that

$$\begin{aligned} \langle f, f \rangle_{\pi} &= \sum_i \alpha_i^2 \\ \langle Pf, f \rangle_{\pi} &= \sum_{i=1}^{|\Omega|} \sum_{j=1}^{|\Omega|} \lambda_i \alpha_i \alpha_j \langle f_i, f_j \rangle_{\pi} \\ &= \sum_{i=1}^{|\Omega|} \lambda_i \alpha_i^2 \geq \lambda_{|\Omega|} \langle f, f \rangle_{\pi} \end{aligned}$$

Thus, for every function $f \neq 0$ we have that

$$\frac{\mathcal{E}_{\pi}^+(f)}{\langle f, f \rangle_{\pi}} = \frac{\langle f, f \rangle_{\pi} + \langle Pf, f \rangle_{\pi}}{\langle f, f \rangle_{\pi}} \geq 1 + \lambda_{|\Omega|}$$

And by taking the eigenvector $f_{|\Omega|}$ we have

$$\mathcal{E}_{\pi}^+(f_{|\Omega|}) / \langle f_{|\Omega|}, f_{|\Omega|} \rangle_{\pi} = 1 + \lambda_{|\Omega|}.$$

□

By using the form \mathcal{E}^+ an analogous of Lemma 20 can be shown for the last eigenvalue.

LEMMA 23. *Let P and \hat{P} be irreducible and reversible transition matrices over the same state space Ω and stationary distributions π and $\hat{\pi}$, respectively. If $\mathcal{E}_{\hat{\pi}}^+(f) \leq \alpha \mathcal{E}_{\pi}^+(f)$, then*

$$1 + \hat{\lambda}_{|\Omega|} \leq \left[\max_{x \in \Omega} \frac{\pi(x)}{\hat{\pi}(x)} \right] \alpha (1 + \lambda_{|\Omega|})$$

PROOF. Let $c(\pi, \hat{\pi}) = \max\{\pi(x)/\hat{\pi}(x) : x \in \Omega\}$ be the maximum ratio between π and $\hat{\pi}$, then for every function f , the π -norm squared is at most $c(\pi, \hat{\pi})$ times the $\hat{\pi}$ -norm squared, i.e.

$$\langle f, f \rangle_{\pi} = \sum_{x \in \Omega} f(x)^2 \pi(x) = \sum_{x \in \Omega} f(x)^2 \frac{\pi(x)}{\hat{\pi}(x)} \hat{\pi}(x) \leq c(\pi, \hat{\pi}) \langle f, f \rangle_{\hat{\pi}}$$

Hence, by using the hypothesis $\mathcal{E}_{\hat{\pi}}^+(f) \leq \alpha \mathcal{E}_{\pi}^+(f)$, for every function $f \neq 0$ we have that

$$\frac{\mathcal{E}_{\hat{\pi}}^+(f)}{\langle f, f \rangle_{\hat{\pi}}} \leq \alpha c(\pi, \hat{\pi}) \frac{\mathcal{E}_{\pi}^+(f)}{\langle f, f \rangle_{\pi}} \quad (13)$$

And the thesis follows from Observation 22 by taking the minimum over all $f \neq 0$ on both sides of (13). □

Finally, we can prove the Comparison Theorem as stated in Section 2.

Proof of Theorem 2. Since $\hat{Q}(x, y) \leq \alpha Q(x, y)$ for all $x, y \in \Omega$, from (11) and (12) it follows that $\mathcal{E}_{\hat{\pi}}(f) \leq \alpha \mathcal{E}_{\pi}(f)$ and $\mathcal{E}_{\hat{\pi}}^+(f) \leq \alpha \mathcal{E}_{\pi}^+(f)$ for every function f . Since $\pi(x) \leq \gamma \hat{\pi}(x)$, from Lemmas 20 and 23 it follows that

$$1 - \hat{\lambda}_2 \leq \alpha \gamma (1 - \lambda_2) \quad \text{and} \quad 1 + \hat{\lambda}_{|\Omega|} \leq \alpha \gamma (1 + \lambda_{|\Omega|})$$

The thesis follows from the definition of relaxation time. □

C. THE COUPLING

In this section, we describe, for each $\mathbf{x}, \mathbf{y} \in S$, a coupling of $P(\mathbf{x}, \cdot)$ and $P(\mathbf{y}, \cdot)$ for the Markov chain \mathcal{M}_{β}^S whose transition matrix P is given by Equation (2). We will then show that the coupling described has the properties required by the proof of Theorem 10.

For each player i , we partition two copies of the interval $[0, 1]$, called $I_{X,i}$ and $I_{Y,i}$, in sub-intervals each labeled with a strategy from the set $S_i = \{z_1, \dots, z_{|S_i|}\}$ of strategies of player i . The sub-intervals are constructed as follows. For $k = 1, \dots, |S_i|$, we take the leftmost not yet labeled interval of length $l_k = \min\{\sigma_i(z_k | \mathbf{x}), \sigma_i(z_k | \mathbf{y})\}$ of both $I_{X,i}$ and $I_{Y,i}$ and label it with strategy z_k . In addition, we take the rightmost non yet labeled interval of length $\sigma_i(z_k | \mathbf{x}) - l_k$ of I_X and the rightmost non yet labeled interval of length $\sigma_i(z_k | \mathbf{y}) - l_k$ of I_Y and label both with z_k . Notice that at least one of these two intervals has length 0. Define functions $h_{X,i}: I_{X,i} \rightarrow S_i$ and $h_{Y,i}: I_{Y,i} \rightarrow S_i$ that for $s \in [0, 1]$ return the labels $h_{X,i}(s)$ and $h_{Y,i}(s)$ of the sub-intervals containing s .

Given the above partitions of $I_{X,i}$ and $I_{Y,i}$ for each i , the coupling can be described as follows: pick $i \in [n]$ and $U \in [0, 1]$ uniformly at random and update X and Y by setting $X_i = h_{X,i}(U)$ and $Y_i = h_{Y,i}(U)$. By construction we have that (X, Y) is a coupling of $P(\mathbf{x}, \cdot)$ and $P(\mathbf{y}, \cdot)$.

We finish by observing that, if player i is selected, the probability that both chains choose strategy z for player i is exactly $\min\{\sigma_i(z | \mathbf{x}), \sigma_i(z | \mathbf{y})\}$. If z is dominant for player i , we have that $\sigma_i(z | \mathbf{x}), \sigma_i(z | \mathbf{y}) \geq 1/|S_i|$ and thus the probability that the coupling updates to z is at least $1/|S_i|$.