

Mixing Time and Stationary Expected Social Welfare of Logit Dynamics*

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

May 24, 2011

Abstract

We study *logit dynamics* [3] for strategic games. This dynamics works as follows: at every stage of the game a player is selected uniformly at random and she plays according to a *noisy* best-response where the noise level is tuned by a parameter β . Such a dynamics defines a family of ergodic Markov chains, indexed by β , over the set of strategy profiles. We believe that the stationary distribution of these Markov chains gives a meaningful description of the long-term behavior for systems whose agents are not completely rational.

Our aim is twofold: On the one hand, we are interested in evaluating the performance of the game at equilibrium, i.e. the expected social welfare when the strategy profiles are random according to the stationary distribution. On the other hand, we want to estimate how long it takes, for a system starting at an arbitrary profile and running the logit dynamics, to get close to its stationary distribution; i.e., the *mixing time* of the chain.

In this paper we study the stationary expected social welfare for the 3-player CK game [6], for 2-player coordination games, and for two simple n -player games. For all these games, we also give almost tight upper and lower bounds on the mixing time of logit dynamics. Our results show two different behaviors: in some games the mixing time depends exponentially on β , while for other games it can be upper bounded by a function independent of β .

1 Introduction

The evolution of a system is determined by its dynamics and complex systems are often described by looking at the equilibrium states induced by their dynamics. Once the system reaches an equilibrium state it stays there, thus equilibrium states describe the long-term behavior of the system. In this paper we are mainly interested in systems whose individual components are *selfish* agents. The state of a selfish system is fully described by a vector of *strategies*, each controlled by one agent, and each state assigns a payoff to each agent. The agents are selfish in the sense that they pick their strategy so to maximize their payoff, given the strategies of the other agents. *Nash equilibrium* is the classical notion of equilibrium for selfish systems and it corresponds to the equilibrium induced by the *best-response* dynamics. The observation that selfish systems are described by their equilibrium states (that is, by the Nash equilibria) has motivated the notions of Price of Anarchy [15] and Price of Stability [1] and the analysis of efficiency of selfish systems based on such notions.

However, such analysis inherits some of the shortcomings of the concept of a Nash equilibrium. First of all, the best-response dynamics assumes that selfish agents have complete knowledge of the current state of the system; that is, they know the payoff associated with each of their possible choices and each of the strategies chosen by other agents. Instead, in most cases, agents have only approximate knowledge of the system state or they are not able to compute their best choice. Moreover, in presence of multiple

*Work partially supported by the PRIN 2008 research project COGENT (COmputational and Game-theoretic aspects of uncoordinated NeTworks), funded by the Italian Ministry of University and Research.

equilibria, it is not clear which one of them will be reached by the system, as it may depend on the initial state: Price of Anarchy considers the worst-case equilibrium, whereas Price of Stability focuses on the best-case equilibrium. Finally, Nash equilibria are hard to compute [7, 5] and thus for some systems it might take very long to reach a Nash equilibrium: in this case using equilibrium states to describe the system performance is not well justified. Rather, one would like to analyze the performance of a system by using a dynamics (and its related equilibrium notion) that has the following three properties:

- The dynamics takes into account the fact that the system components might have a perturbed or noisy knowledge of the system;
- For every system the equilibrium state exists and is unique;
- The system reaches the equilibrium very quickly regardless of the starting state.

In this paper, we consider *noisy best-response dynamics* in which the behavior of the agents is described by a parameter $\beta \geq 0$. The case $\beta = 0$ corresponds to agents picking their strategies completely at random (that is, the agents have no knowledge of the system) and the case $\beta = \infty$ corresponds to agents picking their strategies according to the best-response dynamics (in which the agents have full and complete knowledge of the system). The intermediate values of β correspond to agents that are roughly guided by the best-response dynamics but can make a sub-optimal choice due, for example, to *bounded rationality* of the agent or *limited knowledge* about the system: this sub-optimal behavior occurs with some probability that depends on β (and on the associated payoff).

We will study a specific noisy best-response dynamics for which the system evolves according to an ergodic Markov chain for all $\beta \geq 0$. For these systems, it is natural to look at the stationary distribution (which is the equilibrium state of the Markov chain) and to analyze the expected social welfare (the sum of utility functions) of the system at that distribution. We stress that the noisy best-response dynamics well models agents that only have approximate or noisy knowledge of the system and that for ergodic Markov chains (such as the ones arising in our study) the stationary distribution is known to exist and to be unique. Moreover, to justify the use of the stationary distribution for analyzing the performance of the system, we will study how fast the Markov chain converges to the stationary distribution.

Related Works and Our Results. Several dynamics, besides the best-response dynamics, and several notions of equilibrium, besides Nash equilibria, have been considered to describe the evolution of a selfish system and to analyze its performance. See, for example, [11, 21, 20].

Equilibrium concepts based on the best-response. When the game does not possess a Pure Nash equilibrium, the best-response dynamics will eventually cycle over a set of states (in a Nash equilibrium the set is a singleton). These states are called *sink equilibria* [12]. Sink equilibria exist for all games and, in some contexts, they seem a better approximation of the real setting than *mixed* Nash equilibria. Unfortunately, sink equilibria share two undesirable properties with Nash equilibria: a game can have more than one sink equilibrium and sink equilibria seem hard to compute [9].

Other notions of equilibrium state associated with best-response dynamics are the *unit-recall equilibria* and *component-wise unit-recall equilibria* (see [9]). However, we point out that the former does not always exist and that the latter imposes too strict limitations on the players.

No-Regret Dynamics. Another broadly explored set of dynamics are the no-regret dynamics (see, for example, [11]). The regret of an user is the difference between the long-term average cost and the average cost of the best strategy in hindsight. In the no-regret dynamics the regret of every player after t steps is $o(t)$ (sublinear with time). In [10, 14] it is showed that the no-regret dynamics converges to the set of *correlated equilibria*. Note that the convergence is to the *set* of correlated equilibria and not to a specific correlated equilibrium.

Our work. In this paper we consider a specific noisy best-response dynamics called the *logit* dynamics (see [3]) and we study its mixing time (that is, the time it takes to converge to the stationary distribution) and the stationary expected social welfare. Specifically,

- We start by analyzing the logit dynamics for a simple 3-player linear congestion game (the CK game [6]) which exhibits the worst Price of Anarchy among linear congestion games. We show that the mixing time of the logit dynamics is upper bounded by a constant independent of β . Moreover, we show that the stationary expected social welfare is larger than the social welfare of the worst Nash equilibrium for all β ;

- We then analyze the 2×2 coordination games studied in [3]. Here we show that, under some conditions, the stationary expected social welfare is larger than the social welfare of the worst Nash equilibrium. We give upper and lower bounds on the mixing time exponential in β . We also observe that the same bounds apply to anti-coordination games;
- Finally, we apply our analysis to two simple n -player games: the OR game and XOR game. We give upper and lower bounds on the mixing time: we show that the mixing time of the OR game can be upper bounded by a function independent of β , while the mixing time of the XOR game increases exponentially in β . We also prove that for $\beta = \mathcal{O}(\log n)$ the mixing time is polynomial in n for both games.

The *logit* dynamics has been first studied by Blume [3] who showed that, for 2×2 coordination games, the long-term behavior of the Markov chain is concentrated in the risk dominant equilibrium (see [13]) for sufficiently large β . Ellison [8] studied different noisy best-response dynamics for coordination games assuming that interaction among players were described by a graph; that is, the utility of a player is determined only by the strategies of the adjacent players. Specifically, Ellison [8] studied interaction modeled by rings and showed that some large fraction of the players will eventually choose the risk dominant strategy. Similar results were obtained by Peyton Young [22] for the logit dynamics and for more general families of graphs. Montanari and Saberi [17] gave bounds on the hitting time (the expected time that the logit dynamics takes to reach a specific state) of the risk dominant equilibrium state in terms of some graph theoretic properties of the underlying interaction network. Asadpour and Saberi [2] studied the hitting time for a broader class of congestion games. We notice that none of [3, 8, 22] gave any bound on the convergence time to the risk dominant equilibrium. Montanari and Saberi [17] were the first to do so but their study focuses on the hitting time of a specific configuration and not on the convergence time to the stationary distribution.

From a technical point of view, our work follows the lead of [3, 8, 22] and extends their technical findings by giving bounds on the mixing time of the Markov chain of the logit dynamics. We stress that previous results only proved that, for sufficiently large β , *eventually* the system concentrates around certain states without further quantifying the rate of convergence nor the asymptotic behaviour of the system for small values of β . Instead, we identify the stationary distribution of the logit dynamics as the *global* equilibrium and we evaluate the social welfare at stationarity and the time it takes the system to reach it (the mixing time) as explicit functions of β .

We choose to start our study from the class of coordination games considered in [3] and two simple n -player games (the OR game and the XOR game). We give nearly tight upper and lower bounds on the mixing time. Despite their *game-theoretic* simplicity, the analytical study of the mixing time of the logit dynamics for the two n -player games is far from trivial. We notice that the results in [17] cannot be used to derive upper bounds on the mixing time.

From a more conceptual point of view, our work tries (similarly to [12, 9, 19]) to introduce a solution concept that well models the behavior of selfish agents, is uniquely defined for any game, and is quickly reached from any starting state. We propose the stationary distribution induced by the logit dynamics as a possible solution concept and exemplify its use in the analysis of the performance of some 2×2 games (as the ones considered in [3]), of games used to obtain tight bounds on the Price of Anarchy, and of two simple multi-player games.

Organization of the paper. In Section 2 we summarize some Markov chain notions that we will use throughout the paper. In Section 3 we formally describe the logit dynamics for strategic games. We also describe the coupling we will repeatedly use in the proofs of the upper bounds on mixing times. In Sections 4, 5, 6 and 7 we study the stationary expected social welfare and the mixing time of the logit dynamics for CK game, coordination games, the OR game, and the XOR game, respectively. Finally, in Section 8 we present conclusions and some open problems.

Notation. We write \bar{S} for the complementary set of a set S ; we write $|S|$ for its size. We use bold symbols for vectors; when $\mathbf{x} = (x_1, \dots, x_n) \in \{0, 1\}^n$ we write $|\mathbf{x}|$ for the number of 1s in \mathbf{x} ; i.e., $|\mathbf{x}| = |\{i \in [n] : x_i = 1\}|$. For two vectors \mathbf{x}, \mathbf{y} let $H(\mathbf{x}, \mathbf{y}) = |\{i \in [n] : x_i \neq y_i\}|$ be their Hamming distance: we write $\mathbf{x} \sim \mathbf{y}$ if $H(\mathbf{x}, \mathbf{y}) = 1$. 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. $(\mathbf{x}_{-i}, y) = (x_1, \dots, x_{i-1}, y, x_{i+1}, \dots, x_n)$.

2 Markov chains summary and notation

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 5.2, 7.2, 12.2 and 14.2 of [16]. We refer the reader to [16] also for notational conventions).

Consider a Markov chain \mathcal{M} with *finite* state space Ω and transition matrix P . It is a classical result that for an *irreducible* and *aperiodic* Markov chain¹ there exists a unique *stationary distribution* π over Ω ; that is, a distribution π on Ω such that $\pi \cdot P = \pi$.

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

$$\|\mu - \nu\|_{\text{TV}} = \max_{A \subset \Omega} |\mu(A) - \nu(A)|.$$

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

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

where

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

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

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

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

Coupling. A *coupling* of two probability distributions μ and ν on Ω is a pair of random variables (X, Y) defined on $\Omega \times \Omega$ such that the marginal distribution of X is μ and the marginal distribution of Y is ν . A *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,

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

We will consider only couplings of Markov chains with the property that for $s \geq \tau_{\text{couple}}$, it holds $X_s = Y_s$. The following theorem can be used to give an upper bound on t_{mix} (see, for example, Corollary 5.3 in [16]).

Theorem 1 (Coupling) *Let \mathcal{M} be a Markov chain with 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*

$$d(t) \leq \max_{x,y \in \Omega} \mathbf{P}_{x,y}(\tau_{\text{couple}} > t).$$

Sometimes it is difficult to specify a coupling and to analyze the coupling time τ_{couple} for each pair of starting states x and y . The *Path Coupling* theorem says that it is sufficient to define a coupling only for pairs of Markov chains starting from *adjacent* states and an upper bound on the mixing time can be obtained if each of these couplings contracts their distance on average. More precisely, consider a Markov chain \mathcal{M} with state space Ω and transition matrix P ; let $G = (\Omega, E)$ be a connected graph and let $w : E \rightarrow \mathbb{R}$ be a function assigning weights to the edges such that $w(e) \geq 1$ for every edge $e \in E$; for $x, y \in \Omega$, we denote by $\rho(x, y)$ the weight of the shortest path in G between x and y . The following theorem holds.

¹Roughly 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.

Theorem 2 (Path Coupling [4]) Suppose that for every edge $\{x, y\} \in E$ a coupling (X_t, Y_t) of \mathcal{M} with $X_0 = x$ and $Y_0 = y$ exists such that $\mathbf{E}_{x,y}[\rho(X_1, Y_1)] \leq e^{-\alpha} \cdot w(\{x, y\})$ for some $\alpha > 0$. Then

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

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

Spectral techniques. A Markov chain \mathcal{M} with state space Ω and transition matrix P is said *reversible* if for all $x, y \in \Omega$, it holds that

$$\pi(x) \cdot P(x, y) = \pi(y) \cdot P(y, x).$$

The eigenvalues of the transition matrix P of a reversible Markov chain \mathcal{M} can be used to obtain upper and lower bounds on the mixing time. Observe that all the eigenvalues of any transition matrix P have absolute value at most 1, $\lambda = 1$ is an eigenvalue, and for irreducible and aperiodic chains, -1 is not an eigenvalue. The *relaxation time* t_{rel} of a Markov chain \mathcal{M} is defined as

$$t_{\text{rel}} = \frac{1}{1 - \lambda^*}$$

where λ^* is the largest absolute value among eigenvalues other than 1,

$$\lambda^* = \max\{|\lambda| : \lambda \text{ is an eigenvalue of } P, \lambda \neq 1\}.$$

Observe that, for \mathcal{M} reversible, irreducible and aperiodic, $0 \leq \lambda^* < 1$ and thus t_{rel} is positive and finite. We have the following theorem (see, for example, Theorems 12.3 and 12.4 in [16]).

Theorem 3 (Relaxation time) Let P be the transition matrix of a reversible, irreducible, and aperiodic Markov chain with state space Ω and stationary distribution π . Then

$$(t_{\text{rel}} - 1) \log\left(\frac{1}{2\varepsilon}\right) \leq t_{\text{mix}}(\varepsilon) \leq \log\left(\frac{1}{\varepsilon \pi_{\min}}\right) t_{\text{rel}}$$

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

Lower bound. We will use the following theorem to derive our lower bounds (see, for example, Theorem 7.3 in [16]).

Theorem 4 (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 $S \subseteq \Omega$ be any set with $\pi(S) \leq 1/2$. Then the mixing time is

$$t_{\text{mix}}(\varepsilon) \geq \frac{1 - 2\varepsilon}{2\Phi(S)}$$

where

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

3 The model and the problem

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

Consider the following *noisy best-response dynamics*, introduced in [3] and known as *logit dynamics*: at every time step

1. Select one player $i \in [n]$ uniformly at random;

2. Update the strategy of player i according to the following probability distribution over the set S_i of her strategies. For every $y \in S_i$

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

where $\mathbf{x} \in S_1 \times \dots \times S_n$ is the strategy profile played 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 $\beta \geq 0$.

Parameter β is called *inverse noise* of the system, indeed from (1) it is easy to see that, for $\beta = 0$ player i selects her strategy uniformly at random, for $\beta > 0$ the probability is biased toward strategies promising higher payoffs, and for $\beta \rightarrow \infty$ player i chooses her best response strategy (if more than one best response is available, she chooses uniformly at random one of them). Moreover observe that probability $\sigma_i(y | \mathbf{x})$ does not depend on the strategy x_i currently adopted by player i .

The above dynamics defines a Markov chain with the set of strategy profiles as state space, and where the transition probability from profile $\mathbf{x} = (x_1, \dots, x_n)$ to profile $\mathbf{y} = (y_1, \dots, y_n)$ is zero if the $H(\mathbf{x}, \mathbf{y}) \geq 2$ 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 5 (Logit dynamics [3]) Let $\mathcal{G} = ([n], \mathcal{S}, \mathcal{U})$ be a strategic game and let $\beta \geq 0$. The logit dynamics for \mathcal{G} is the Markov chain $\mathcal{M}_\beta = \{X_t : t \in \mathbb{N}\}$ with state space $\Omega = S_1 \times \dots \times S_n$ and transition matrix

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

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

Properties. Logit dynamics enjoys some interesting properties:

Ergodicity. It is easy to see that the logit dynamics is irreducible and aperiodic. Indeed, let $\mathbf{x} = (x_1, \dots, x_n)$ and $\mathbf{y} = (y_1, \dots, y_n)$ be two profiles and let $(\mathbf{z}^0, \dots, \mathbf{z}^n)$ be a *path* of profiles where $\mathbf{z}^0 = \mathbf{x}$, $\mathbf{z}^n = \mathbf{y}$ and $\mathbf{z}^i = (y_1, \dots, y_i, x_{i+1}, \dots, x_n)$ for $i = 1, \dots, n-1$. The probability that the chain starting at \mathbf{x} is in \mathbf{y} after n steps is

$$P^n(\mathbf{x}, \mathbf{y}) = P^n(\mathbf{z}^0, \mathbf{z}^n) \geq P^{n-1}(\mathbf{z}^0, \mathbf{z}^{n-1})P(\mathbf{z}^{n-1}, \mathbf{z}^n)$$

and recursively

$$P^n(\mathbf{x}, \mathbf{y}) \geq \prod_{i=1}^n P(\mathbf{z}^{i-1}, \mathbf{z}^i) > 0$$

where the last inequality follows from (2) because, for all $i = 1, \dots, n$, the Hamming distance between \mathbf{z}^{i-1} and \mathbf{z}^i is at most 1. Hence there is a unique stationary distribution π and, for every starting profile \mathbf{x} , the distribution of the chain $P^t(\mathbf{x}, \cdot)$ converges to π in total variation as t tends to infinity.

Invariance under utility translation. Let $\mathcal{G} = ([n], \mathcal{S}, \mathcal{U})$ be a game. If we change the utility functions by adding a constant c_i to all the utilities of player i , i.e. if we define a new family $\tilde{\mathcal{U}} = \{\tilde{u}_i : i \in [n]\}$ of utility functions as follows

$$\tilde{u}_i(\mathbf{x}) := u_i(\mathbf{x}) + c_i \quad \text{for all } \mathbf{x}$$

we get a new game $\tilde{\mathcal{G}} = ([n], \mathcal{S}, \tilde{\mathcal{U}})$ but the same logit dynamics. Indeed, according to (1), the probability player i chooses strategy y when the game is at profile \mathbf{x} is

$$\tilde{\sigma}_i(y | \mathbf{x}) = \frac{e^{\beta \tilde{u}_i(\mathbf{x}_{-i}, y)}}{\sum_{z \in S_i} e^{\beta \tilde{u}_i(\mathbf{x}_{-i}, z)}} = \frac{1}{\sum_{z \in S_i} e^{\beta [\tilde{u}_i(\mathbf{x}_{-i}, z) - \tilde{u}_i(\mathbf{x}_{-i}, y)]}} = \frac{1}{\sum_{z \in S_i} e^{\beta [u_i(\mathbf{x}_{-i}, z) - u_i(\mathbf{x}_{-i}, y)]}} = \sigma_i(y | \mathbf{x}).$$

Noise changes under utility rescaling. While translations of utilities do not affect logit dynamics, a *rescaling* of the utility functions for the same constant $\alpha > 0$ changes the inverse noise from β to $\alpha \cdot \beta$. Indeed, if for every player i and every profile \mathbf{x} we set

$$\tilde{u}_i(\mathbf{x}) := \alpha \cdot u_i(\mathbf{x}),$$

from (1) we have

$$\tilde{\sigma}_i(y | \mathbf{x}) = \frac{e^{\beta \tilde{u}_i(\mathbf{x}_{-i}, y)}}{\sum_{z \in S_i} e^{\beta \tilde{u}_i(\mathbf{x}_{-i}, z)}} = \frac{e^{\alpha \beta u_i(\mathbf{x}_{-i}, y)}}{\sum_{z \in S_i} e^{\alpha \beta u_i(\mathbf{x}_{-i}, z)}}.$$

Notice that, unlike translations constants, we here must have the same rescaling constant α for all utility functions.

Potential Games. A game $\mathcal{G} = ([n], \mathcal{S}, \mathcal{U})$ is said a (exact) *potential* game if a function $\Phi : S_1 \times \dots \times S_n \rightarrow \mathbb{R}$ 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 and its stationary distribution is the Gibbs measure

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

where $Z = \sum_{\mathbf{y} \in S_1 \times \dots \times S_n} e^{\beta \Phi(\mathbf{y})}$ is the normalizing constant. Except for the Matching Pennies example in Subsection 3.1, all the games we analyze in this paper are potential games.

Logit dynamics vs Glauber dynamics. When \mathcal{G} is a potential game, the logit dynamics is equivalent to the well-studied *Glauber dynamics*. For state space $\Omega = S_1 \times \dots \times S_n$ and probability distribution μ over Ω , the Glauber dynamics for μ proceeds as follows: From profile $\mathbf{x} \in \Omega$, pick a player $i \in [n]$ uniformly at random and update her strategy at $y \in S_i$ with probability μ conditioned on the other players being at \mathbf{x}_{-i} , i.e.

$$\mu(y | \mathbf{x}_{-i}) = \frac{\mu(\mathbf{x}_{-i}, y)}{\sum_{z \in S_i} \mu(\mathbf{x}_{-i}, z)}.$$

It is easy to see that the Markov chain defined by the Glauber dynamics is irreducible, aperiodic, and reversible with stationary distribution μ . When $\mathcal{G} = ([n], \mathcal{S}, \mathcal{U})$ is a potential game with potential function Φ , the logit dynamics defines the same Markov chain as the Glauber dynamics for the Gibbs distribution π in (3). Indeed, in that case we have

$$\begin{aligned} \sigma_i(y | \mathbf{x}) &= \frac{e^{\beta u_i(\mathbf{x}_{-i}, y)}}{\sum_{z \in S_i} e^{\beta u_i(\mathbf{x}_{-i}, z)}} = \frac{1}{\sum_{z \in S_i} e^{\beta(u_i(\mathbf{x}_{-i}, z) - u_i(\mathbf{x}_{-i}, y))}} \\ &= \frac{1}{\sum_{z \in S_i} e^{\beta(\Phi(\mathbf{x}_{-i}, z) - \Phi(\mathbf{x}_{-i}, y))}} = \frac{e^{\beta \Phi(\mathbf{x}_{-i}, y)}}{\sum_{z \in S_i} e^{\beta \Phi(\mathbf{x}_{-i}, z)}} = \frac{\pi(\mathbf{x}_{-i}, y)}{\sum_{z \in S_i} \pi(\mathbf{x}_{-i}, z)}. \end{aligned}$$

Hence, logit dynamics for potential games and Glauber dynamics for Gibbs distributions are two ways of looking at the same Markov chains: in the former case the dynamics is derived from the potential function, in the latter case from the stationary distribution. However, observe that, if \mathcal{G} is not a potential game and π is the stationary distribution of the logit dynamics for \mathcal{G} , in general the Glauber dynamics for π is different from the logit dynamics (see, for example, the Matching Pennies case in Subsection 3.1).

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).

Stationary expected social welfare and mixing time. Let $W : S_1 \times \dots \times S_n \rightarrow \mathbb{R}$ be a *social welfare function* (in this paper we assume that W is simply the sum of all the utility functions $W(\mathbf{x}) = \sum_{i=1}^n u_i(\mathbf{x})$, but clearly any other function of interest can be analysed). We study the *stationary expected social welfare*, i.e. the expectation of W when the strategy profiles are random according to the stationary distribution π of the Markov chain,

$$\mathbf{E}_\pi [W] = \sum_{\mathbf{x} \in S_1 \times \dots \times S_n} W(\mathbf{x}) \pi(\mathbf{x})$$

Since the Markov chain defined in (2) is irreducible and aperiodic, from every initial profile \mathbf{x} the distribution $P^t(\mathbf{x}, \cdot)$ of chain X_t starting at \mathbf{x} will eventually converge to π as t tends to infinity. We will

be interested in bounding how long it takes to get close to the stationary distribution, that is the *mixing time* of the Markov chain.

In the next subsection we illustrate the goals of our work with two simple examples.

3.1 Two simple examples: Matching Pennies and a *Stairs* game

Matching Pennies. Consider the classical *Matching Pennies* game. We write the utility functions in the standard bimatrix form.

	H	T	
H	+1, -1	-1, +1	(4)
T	-1, +1	+1, -1	

According to (1), the update probabilities for the logit dynamics are, for every $x \in \{H, T\}$

$$\begin{aligned} \sigma_1(H | (x, H)) &= \sigma_1(T | (x, T)) = \frac{1}{1+e^{-2\beta}} = \sigma_2(T | (H, x)) = \sigma_2(H | (T, x)), \\ \sigma_1(T | (x, H)) &= \sigma_1(H | (x, T)) = \frac{1}{1+e^{2\beta}} = \sigma_2(H | (H, x)) = \sigma_2(T | (T, x)). \end{aligned}$$

Hence the transition matrix (see (2)) is

$$P = \left(\begin{array}{c|cccc} & HH & HT & TH & TT \\ \hline HH & 1/2 & b/2 & (1-b)/2 & 0 \\ HT & (1-b)/2 & 1/2 & 0 & b/2 \\ TH & b/2 & 0 & 1/2 & (1-b)/2 \\ TT & 0 & (1-b)/2 & b/2 & 1/2 \end{array} \right)$$

where, for readability sake, we named $b = \frac{1}{1+e^{-2\beta}}$.

Since every column of the matrix adds up to 1, the uniform distribution π over the set of strategy profiles is the stationary distribution for the logit dynamics. The stationary expected social welfare is thus 0 for every inverse noise β .

As for the mixing time, it is easy to see that it is upper bounded by a constant independent of β . Indeed, a direct calculation shows that, for every $\mathbf{x} \in \{HH, HT, TH, TT\}$ and for every $\beta \geq 0$ it holds that

$$\|P^3(\mathbf{x}, \cdot) - \pi\|_{\text{TV}} \leq \frac{7}{16} < \frac{1}{2}.$$

A *stairs* game. One of the main techniques used to give upper bounds on the mixing time of Markov chains is the coupling technique (see Theorem 1). In the following example we use it to upper bound the mixing time of the logit dynamics for a simple game.

Let \mathcal{G} be a potential game where every player has two strategies, say *upstairs* (or 1) and *downstairs* (or 0), and the potential of a profile $\mathbf{x} \in \{0, 1\}^n$ is the number of players that are upstairs, i.e. $\Phi(\mathbf{x}) = |\mathbf{x}|$.

Notice that the logit dynamics (and thus the stationary distribution and the mixing time) is completely defined by the potential function, while if we wanted to evaluate the stationary expected social welfare we would need to specify the utility functions.

The partition function is

$$Z(\beta) = \sum_{\mathbf{x} \in \{0,1\}^n} e^{\beta|\mathbf{x}|} = \sum_{k=0}^n \binom{n}{k} e^{\beta k} = (1 + e^\beta)^n.$$

So the stationary distribution is

$$\pi(\mathbf{x}) = \frac{e^{\beta|\mathbf{x}|}}{(1 + e^\beta)^n}.$$

As for the mixing time, we can use the coupling technique as follows: observe that the probability of playing strategy 1 (or equivalently strategy 0), for the player selected for the update, is independent of

the current strategies of the other players. Indeed, according to (1), for every \mathbf{x} it holds that

$$\begin{aligned}\sigma_i(1 \mid \mathbf{x}) &= \frac{e^{\beta u_i(\mathbf{x}_{-i}, 1)}}{e^{\beta u_i(\mathbf{x}_{-i}, 1)} + e^{\beta u_i(\mathbf{x}_{-i}, 0)}} = \frac{1}{1 + e^{\beta(u_i(\mathbf{x}_{-i}, 0) - u_i(\mathbf{x}_{-i}, 1))}} \\ &= \frac{1}{1 + e^{\beta(\Phi(\mathbf{x}_{-i}, 0) - \Phi(\mathbf{x}_{-i}, 1))}} = \frac{1}{1 + e^{\beta(|\mathbf{x}_{-i}| - (|\mathbf{x}_{-i}| + 1))}} = \frac{1}{1 + e^{-\beta}}.\end{aligned}$$

We can define a coupling of two Markov chains starting at two different profiles as follows: choose $i \in [n]$ uniformly at random and perform the same update at player i in both chains². When every player has been chosen at least once the two chains have coalesced. From the coupon collector's argument, it takes $\mathcal{O}(n \log n)$ to have that, with probability at least $3/4$, all players have been chosen at least once. By applying Theorem 1 we have that the mixing time is $\mathcal{O}(n \log n)$.

In the above examples, it turned out that the mixing time of the logit dynamics can be upper bounded by functions that do not depend on the inverse noise β . As we shall see in the next sections, this is not always the case. Moreover, the analysis of the mixing time is usually far from trivial.

3.2 Description of the Coupling

Throughout the paper we will use the coupling and path-coupling techniques (see Theorem 1 and Theorem 2) to give upper bounds on mixing times. Since we will use the same coupling idea in several proofs, we describe it here and we will refer to this description when we will need it.

Consider an n -player 2-strategy game \mathcal{G} and let us rename 0 and 1 the strategies of every player. For every pair of strategy profiles $\mathbf{x} = (x_1, \dots, x_n), \mathbf{y} = (y_1, \dots, y_n) \in \{0, 1\}^n$ we define a coupling (X_1, Y_1) of two copies of the Markov chain with transition matrix P defined in (2) for which $X_0 = \mathbf{x}$ and $Y_0 = \mathbf{y}$.

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

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

Three easy observations are in order: if $\sigma_i(0 \mid \mathbf{x}) = \sigma_i(0 \mid \mathbf{y})$ and player i is chosen, then, after the update, we have $x_i = y_i$; for every player i , at most one of the updates $(x_i, y_i) = (0, 1)$ and $(x_i, y_i) = (1, 0)$ has positive probability; 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

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

We define $G = (\Omega, E)$ as the *Hamming graph* of the game, where $\Omega = \{0, 1\}^n$ is the set of strategy profiles, and two profiles $\mathbf{x} = (x_1, \dots, x_n), \mathbf{y} = (y_1, \dots, y_n) \in \Omega$ are adjacent if they differ only for the strategy of one player, i.e.

$$\{\mathbf{x}, \mathbf{y}\} \in E \iff \mathbf{x} \sim \mathbf{y}. \quad (5)$$

For the path coupling technique, the coupling described above is applied only to pairs of adjacent starting profiles.

²This is the same coupling used in the analysis of the lazy random walk on the hypercube (e.g. see Section 5.3.3 in [16]), the only difference being that the probability of choosing 0 or 1 is not $1/2, 1/2$ but $1/(1 + e^\beta), 1/(1 + e^{-\beta})$.

4 A 3-player congestion game

In this section we analyze the CK game, a simple 3-player linear congestion game introduced in [6]. This game is interesting because it highlights the weakness of the Price of Anarchy notion for the logit dynamics. Indeed, the CK game exhibits the worst Price of Anarchy with respect to the average social welfare among all linear congestion games with 3 or more players. But, as we shall see soon, the stationary expected social welfare of the logit dynamics is always larger than the social welfare of the worst Nash equilibrium and, for large enough β , players spend most of the time in the best Nash equilibrium. Moreover, we will show that the mixing time of the logit dynamics can be bounded independently from β : that is, the stationary distribution guarantees a good social welfare and it is quickly reached by the system.

Let us now describe the CK game. We have 3 players and 6 facilities divided into two sets: $G = \{g_1, g_2, g_3\}$ and $H = \{h_1, h_2, h_3\}$. Player $i \in \{0, 1, 2\}$ has two strategies: Strategy “0” consists in selecting facilities (g_i, h_i) ; Strategy “1” consists in selecting facilities $(g_{i+1}, h_{i-1}, h_{i+1})$ (index arithmetic is modulo 3). The cost of a facility is the number of players choosing such facility, and the welfare of a player is minus the sum of the costs of the facilities she selected. It is easy to see that this game has two pure Nash equilibria: the solution where every player plays strategy 0 (each player pays 2, which is optimal), and the solution where every player plays strategy 1 (each player pays 5). The game is a congestion game, and thus, by [18], it is also a potential game and its potential function is:

$$\Phi(\mathbf{x}) = \sum_{j \in G \cup H} \sum_{i=1}^{L_{\mathbf{x}}(j)} i,$$

where $L_{\mathbf{x}}(j)$ is the number of players using facility j in configuration \mathbf{x} .

Stationary expected social welfare. It is easy to see that the update probabilities given by the logit dynamics for this game (see Equation (1)) only depend on the number of players playing strategy 1 and not on which player is actually playing that strategy. In particular, we have that, from a profile \mathbf{x} , the player i , if selected for update, plays strategy 0 with the following probabilities:

$$\sigma_i(0 \mid |\mathbf{x}_{-i}| = 0) = \frac{1}{1 + e^{-4\beta}}, \quad \sigma_i(0 \mid |\mathbf{x}_{-i}| = 1) = \frac{1}{1 + e^{-2\beta}}, \quad \sigma_i(0 \mid |\mathbf{x}_{-i}| = 2) = \frac{1}{2}, \quad (6)$$

and strategy 1 with the remaining probabilities.

Next theorem evaluates the stationary expected social welfare for this game.

Theorem 6 (Expected social welfare) *The stationary expected social welfare $\mathbf{E}_{\pi}[W]$ of the logit dynamics for the CK game is*

$$\mathbf{E}_{\pi}[W] = -\frac{6 + 39e^{-4\beta} + 63e^{-6\beta}}{1 + 3e^{-4\beta} + 4e^{-6\beta}}.$$

Proof. We notice that two profiles with the same number of players playing strategy 1 have both the same potential (and, by Equation (3), the same stationary distribution) and the same social welfare. Thus, $\pi(\mathbf{x}) = \pi[k]$ and $W(\mathbf{x}) = W[k]$ for a profile \mathbf{x} such that $|\mathbf{x}| = k$, where

$$\pi[0] = \frac{e^{-6\beta}}{Z(\beta)}, \quad \pi[1] = \frac{e^{-10\beta}}{Z(\beta)}, \quad \pi[2] = \pi[3] = \frac{e^{-12\beta}}{Z(\beta)},$$

where $Z(\beta) = e^{-6\beta} + 3e^{-10\beta} + 4e^{-12\beta}$, and

$$W[0] = -6, \quad W[1] = -13, \quad W[2] = -16, \quad W[3] = -15.$$

Hence, the stationary expected social welfare is

$$\mathbf{E}_{\pi}[W] = -\frac{6 \cdot e^{-6\beta} + 3 \cdot 13 \cdot e^{-10\beta} + (3 \cdot 16 + 15) \cdot e^{-12\beta}}{e^{-6\beta} + 3e^{-10\beta} + 4e^{-12\beta}} = -\frac{6 + 39e^{-4\beta} + 63e^{-6\beta}}{1 + 3e^{-4\beta} + 4e^{-6\beta}}.$$

□

Notice that for $\beta = 0$ we have $\mathbf{E}_\pi[W] = -27/2$, which is better than the social welfare of the worst Nash equilibrium. This means that, even if each player selects her strategy at random, the logit dynamics drives the system to a random profile whose expectation to be better than the worst Nash equilibrium. We also observe that $\mathbf{E}_\pi[W]$ increases with β and thus the long-term behavior of the logit dynamics gives a better social welfare than the worst Nash equilibrium for any $\beta \geq 0$. Moreover, the stationary expected social welfare approaches the optimal social welfare as β tends to ∞ .

Mixing time. Now we study the mixing time of the logit dynamics for the CK game and we show that it is bounded by a constant for any $\beta \geq 0$. The proof will use Coupling Theorem (see Theorem 1).

Theorem 7 (Mixing time) *There exists a constant τ such that the mixing time t_{mix} of the logit dynamics of the CK game is upper bounded by τ for every $\beta \geq 0$.*

Proof. First, we notice that the update probabilities given in Equation (6) imply that

$$\forall i, \forall \mathbf{x}, \forall \beta, \quad \sigma_i(0 | \mathbf{x}) \geq 1/2. \quad (7)$$

Let X_t and Y_t be two copies of the logit dynamics for the CK game, starting in \mathbf{x} and \mathbf{y} respectively, coupled as described in Section 3.2. It is easy to check that, by Equation (7), the player selected for update, chooses strategy 0 in both chain with probability at least $1/2$.

Finally, we bound the probability that after three steps the two coupled chains coalesce: it is at least as large as the probability that we choose three different players and all of them play strategy 0 at their turn, i.e.

$$\mathbf{P}_{\mathbf{x}, \mathbf{y}}(X_3 = Y_3) \geq \frac{1}{2} \cdot \frac{1}{3} \cdot \frac{1}{6} = \frac{1}{36}.$$

Since this bound holds for every starting pair (\mathbf{x}, \mathbf{y}) , we have that the probability the two chains have not yet coalesced after $3t$ steps is

$$\mathbf{P}_{\mathbf{x}, \mathbf{y}}(X_{3t} \neq Y_{3t}) \leq \left(1 - \frac{1}{36}\right)^t \leq e^{-t/36}.$$

The thesis follows from the Theorem 1. □

5 Two player games

In this section we analyse the performance of the logit dynamics for 2×2 coordination games (the same class studied in [3]) and 2×2 anti-coordination games.

Coordination games. Coordination Games are two-player games in which the players have an advantage in selecting the same strategy. These games are often used to model the spread of a new technology [22]: two players have to decide whether to adopt or not a new technology. We assume that the players would prefer choosing the same technology as the other one and that choosing the new technology is at most as risky as choosing the old one.

We name 0 the NEW strategy and 1 the OLD strategy. The game is formally described by the following payoff matrix

	0	1	
0	(a, a)	(c, d)	(8)
1	(d, c)	(b, b)	

We assume that $a > d$ and $b > c$ (meaning that they prefer to coordinate) and that $a - d \geq b - c$ (meaning that for each player strategy 0 is at most as risky as strategy 1). Notice that we do not make any assumption on the relation between a and b . For convenience sake we name

$$\Delta := a - d \quad \text{and} \quad \delta := b - c.$$

It is easy to see that this game is a potential game and the following function is an exact potential for it:

$$\Phi(0, 0) = \Delta \quad \Phi(0, 1) = \Phi(1, 0) = 0 \quad \Phi(1, 1) = \delta.$$

This game has two pure Nash equilibria: $(0, 0)$, where each player has utility a , and $(1, 1)$, where each player has utility b . As $d + c < a + b$, the social welfare is maximized at one of the two equilibria.

We analyse the mixing time of the logit dynamics for 2×2 coordination games and compute its stationary expected social welfare as a function of β .

Stationary expected social welfare. The logit dynamics for the coordination game defined by the payoffs in Table 8 establishes that, from a profile \mathbf{x} , player i selected for update plays according to the following probability distribution (see Equation (1)):

$$\begin{aligned}\sigma_i(0 \mid \mathbf{x}_{-i} = 0) &= \frac{1}{1+e^{-\Delta\beta}}, & \sigma_i(1 \mid \mathbf{x}_{-i} = 0) &= \frac{1}{1+e^{\Delta\beta}}, \\ \sigma_i(0 \mid \mathbf{x}_{-i} = 1) &= \frac{1}{1+e^{\delta\beta}}, & \sigma_i(1 \mid \mathbf{x}_{-i} = 1) &= \frac{1}{1+e^{-\delta\beta}}.\end{aligned}$$

Next theorem bounds the stationary expected social welfare $\mathbf{E}_\pi[W]$ obtained by the logit dynamics and gives conditions for which $\mathbf{E}_\pi[W]$ is better than the social welfare SW_N of the worst Nash equilibrium.

Theorem 8 (Expected social welfare) *The stationary expected social welfare $\mathbf{E}_\pi[W]$ of the logit dynamics for the coordination game in Table 8 is*

$$\mathbf{E}_\pi[W] = 2 \cdot \frac{a + be^{-(\Delta-\delta)\beta} + (c+d)e^{-\Delta\beta}}{1 + e^{-(\Delta-\delta)\beta} + 2e^{-\Delta\beta}}.$$

Moreover, if $a \neq b$ then $\mathbf{E}_\pi[W] \geq \text{SW}_N$ for β sufficiently large.

Proof. The stationary distribution π of the logit dynamics is

$$\pi(0, 0) = \frac{e^{\Delta\beta}}{Z(\beta)} \quad \pi(1, 1) = \frac{e^{\delta\beta}}{Z(\beta)} \quad \pi(0, 1) = \pi(1, 0) = \frac{1}{Z(\beta)}$$

where $Z(\beta) = e^{\Delta\beta} + e^{\delta\beta} + 2$.

Since $\mathbf{E}_\pi[W] = 2 \cdot \mathbf{E}_\pi[u_i]$, we compute the expected utility $\mathbf{E}_\pi[u_i]$ of player i at the stationary distribution,

$$\begin{aligned}\mathbf{E}_\pi[u_i] &= \sum_{\mathbf{x} \in \{0,1\}^2} u_i(\mathbf{x})\pi(\mathbf{x}) \\ &= \frac{ae^{\Delta\beta} + be^{\delta\beta} + c + d}{e^{\Delta\beta} + e^{\delta\beta} + 2} \\ &= \frac{a + be^{-(\Delta-\delta)\beta} + (c+d)e^{-\Delta\beta}}{1 + e^{-(\Delta-\delta)\beta} + 2e^{-\Delta\beta}}.\end{aligned}$$

Thus, if $a > b$ and $\beta \geq \max\left\{0, \frac{1}{\Delta} \log \frac{2b-c-d}{a-b}\right\}$, we have

$$\mathbf{E}_\pi[W] - \text{SW}_N = 2 \cdot \frac{a + be^{-(\Delta-\delta)\beta} + (c+d)e^{-\Delta\beta}}{1 + e^{-(\Delta-\delta)\beta} + 2e^{-\Delta\beta}} - 2b = 2 \cdot \frac{(a-b) - (2b-c-d)e^{-\Delta\beta}}{1 + e^{-(\Delta-\delta)\beta} + 2e^{-\Delta\beta}} \geq 0.$$

Similarly, we obtain $\mathbf{E}_\pi[W] - \text{SW}_N \geq 0$ if $b > a$ and $\beta \geq \max\left\{0, \frac{1}{\delta} \log \frac{2a-c-d}{b-a}\right\}$. \square

Mixing time. Now we study the mixing time of the logit dynamics for coordination games and we show that it is exponential in β and in the minimum potential difference between adjacent profiles.

Theorem 9 (Mixing Time) *The mixing time of the logit dynamics for the coordination game of Table 8 is $\Theta(e^{\delta\beta})$ for every $\beta \geq 0$.*

Proof. Upper bound: We apply the Path Coupling technique (see Theorem 2) with the Hamming graph defined in (5) and all the edge-weights set to 1. Let \mathbf{x} and \mathbf{y} be two profiles differing only for the player j and consider the coupling defined in Section 3.2 for this pair of profiles. Now we bound the expected distance of the two coupled chains after one step.

We denote by $b_i(\mathbf{x}, \mathbf{y})$ the probability that both chains perform the same update given that player i has been selected for strategy update. Clearly, $b_i(\mathbf{x}, \mathbf{y}) = 1$ for $i = j$, while for $i \neq j$, we have

$$\begin{aligned} b_i(\mathbf{x}, \mathbf{y}) &= \min\{\sigma_i(0 \mid \mathbf{x}), \sigma_i(0 \mid \mathbf{y})\} + \min\{\sigma_i(1 \mid \mathbf{x}), \sigma_i(1 \mid \mathbf{y})\} \\ &= \frac{1}{1 + e^{\Delta\beta}} + \frac{1}{1 + e^{\delta\beta}}. \end{aligned}$$

For sake of readability we set

$$p = \frac{1}{1 + e^{\Delta\beta}} \quad \text{and} \quad q = \frac{1}{1 + e^{\delta\beta}}.$$

and thus $b_i(\mathbf{x}, \mathbf{y}) = p + q$. To compute $\mathbf{E}_{\mathbf{x}, \mathbf{y}}[\rho(X_1, Y_1)]$, we observe that the logit dynamics chooses player j with probability $1/2$. In this case, as $b_j(\mathbf{x}, \mathbf{y}) = 1$, the coupling updates both chains in the same way, resulting in $X_1 = Y_1$. Similarly, player $i \neq j$ is chosen for strategy update with probability $1/2$. In this case, with probability $b_i(\mathbf{x}, \mathbf{y})$ the coupling performs the same update in both chains resulting in $\rho(X_1, Y_1) = 1$. Instead with probability $1 - b_i(\mathbf{x}, \mathbf{y})$, the coupling performs different updates on the chains resulting in $\rho(X_1, Y_1) = 2$. Therefore we have,

$$\begin{aligned} \mathbf{E}_{\mathbf{x}, \mathbf{y}}[\rho(X_1, Y_1)] &= \frac{1}{2}b_i(\mathbf{x}, \mathbf{y}) + 2 \cdot \frac{1}{2}(1 - b_i(\mathbf{x}, \mathbf{y})) \\ &= 1 - \frac{1}{2}b_i(\mathbf{x}, \mathbf{y}) = 1 - \frac{1}{2}(p + q) \leq e^{-\frac{1}{2}(p+q)}. \end{aligned}$$

From Theorem 2, with $\alpha = \frac{1}{2}(p + q)$ and $\text{diam}(\Omega) = 2$, it follows that

$$t_{\text{mix}}(\varepsilon) \leq \frac{2(\log 2 + \log(1/\varepsilon))}{p + q} = \frac{1}{p + q} \log \frac{4}{\varepsilon^2}.$$

Lower bound: We use the *relaxation time* bound (see Theorem 3). The transition matrix of the logit dynamics is

$$P = \left(\begin{array}{c|cccc} & 00 & 01 & 10 & 11 \\ \hline 00 & 1-p & p/2 & p/2 & 0 \\ 01 & \frac{1-p}{2} & \frac{p+q}{2} & 0 & \frac{1-q}{2} \\ 10 & \frac{1-p}{2} & 0 & \frac{p+q}{2} & \frac{1-q}{2} \\ 11 & 0 & q/2 & q/2 & 1-q \end{array} \right)$$

It is easy to see that the second largest eigenvalue of P is $\lambda_* = \frac{(1-p)+(1-q)}{2}$, hence the relaxation time is $t_{\text{rel}} = 1/(1 - \lambda_*) = \frac{2}{p+q}$, and for the mixing time we have

$$\begin{aligned} t_{\text{mix}}(\varepsilon) &\geq (t_{\text{rel}} - 1) \log \frac{1}{2\varepsilon} = \frac{2 - (p + q)}{p + q} \log \frac{1}{2\varepsilon} \\ &\geq \frac{1}{p + q} \log \frac{1}{2\varepsilon}. \end{aligned} \tag{9}$$

In the last inequality we used that p and q are both smaller than $1/2$.

Finally, the theorem follows by observing that

$$\frac{1}{p + q} = \frac{1}{\frac{1}{1+e^{\Delta\beta}} + \frac{1}{1+e^{\delta\beta}}} = \Theta(e^{\delta\beta}).$$

□

Notice that, if we used the relaxation time to upper bound the mixing time (see Theorem 3) we would get a non-tight bound, hence in the above proof we had to resort to the path coupling for the upper bound.

Anti-coordination games. Very similar results can be obtained for anti-coordination games. These are two-player games in which the players have an advantage in selecting different strategies. They model many settings where there is a common and exclusive resource: two players have to decide whether to use the resource or to drop it. If they both try to use it, then a deadlock occurs and this is bad for both players. Usually, these games are described by a payoff matrix like the one in Table 8, where we assume that $d > a$ and $c > b$ and that $d - a \geq c - b$. Notice that Nash Equilibria of this game are *unfair*, as one player has utility $\max\{c, d\}$ and the other $\min\{c, d\}$.

For the logit dynamics, we have that, for all β , the stationary expected social welfare is worse than the one guaranteed by a Nash equilibrium. On the other hand, for sufficiently large β we have that the expected utility of a player is always better than $\min\{c, d\}$: that is, in the logit dynamics each player expects to gain more than in the worst Nash equilibrium. Moreover, the stationary distribution is a fair equilibrium, since every player has the same expected utility. As for the coordination games, the mixing time is exponential in β and in the minimum potential difference between adjacent profiles.

6 The OR game

In this section we consider the following simple n -player potential game that we here call *OR game*. 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 function of player $i \in [n]$ is

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

Notice that the OR game has $2^n - n$ Nash equilibria. The only profiles that are not Nash equilibria are the n profiles with exactly one player playing 1. Nash equilibrium $\mathbf{0}$ has social welfare 0, while all the others have social welfare $-n$.

In Theorem 10 we show that the stationary expected social welfare is always better than the social welfare of the worst Nash equilibrium, and it is *significantly* better for large β . Unfortunately, in Theorem 11 we show that if β is large enough to guarantee a *good* stationary expected social welfare, then the time needed to get close to the stationary distribution is exponential in n . Finally, in Theorem 12 we give upper bounds on the mixing time showing that if β is relatively small then the mixing time is polynomial in n , while for large β the upper bound is exponential in n and it is almost-tight with the lower bound. Despite the simplicity of the game, the analysis of the mixing time is far from trivial.

Theorem 10 (Expected social welfare) *The stationary expected social welfare of the logit dynamics for the OR game is $\mathbf{E}_\pi [W] = -\alpha n$ where $\alpha = \alpha(n, \beta) = \frac{(2^n - 1)e^{-\beta}}{1 + (2^n - 1)e^{-\beta}}$.*

Proof. Observe that the OR game is a potential game with exact potential Φ where $\Phi(\mathbf{0}) = 0$ and $\Phi(\mathbf{x}) = -1$ for every $\mathbf{x} \neq \mathbf{0}$. Hence the stationary distribution is

$$\pi(\mathbf{x}) = \begin{cases} 1/Z, & \text{if } \mathbf{x} = \mathbf{0}; \\ e^{-\beta}/Z, & \text{if } \mathbf{x} \neq \mathbf{0}; \end{cases}$$

where the normalizing factor is $Z = 1 + (2^n - 1)e^{-\beta}$. The expected social welfare is thus

$$\mathbf{E}_\pi [W] = \sum_{\mathbf{x} \in \{0,1\}^n} W(\mathbf{x})\pi(\mathbf{x}) = -n \cdot \frac{(2^n - 1)e^{-\beta}}{1 + (2^n - 1)e^{-\beta}}.$$

□

In the next theorem we show that the mixing time can be polynomial in n only if $\beta \leq c \log n$ for some constant c .

Theorem 11 (Lower bound on mixing time) *The mixing time of the logit dynamics for the OR game is*

1. $\Omega(e^\beta)$ if $\beta < \log(2^n - 1)$;

2. $\Omega(2^n)$ if $\beta > \log(2^n - 1)$.

Proof. Consider the set $S \subseteq \{0, 1\}^n$ containing only the state $\mathbf{0} = (0, \dots, 0)$ and observe that $\pi(\mathbf{0}) \leq 1/2$ for $\beta \leq \log(2^n - 1)$. The bottleneck ratio is

$$\Phi(\mathbf{0}) = \frac{1}{\pi(\mathbf{0})} \sum_{\mathbf{y} \in \{0,1\}^n} \pi(\mathbf{0})P(\mathbf{0}, \mathbf{y}) = \sum_{\mathbf{y} \in \{0,1\}^n : |\mathbf{y}|=1} P(\mathbf{0}, \mathbf{y}) = n \cdot \frac{1}{n} \frac{1}{1 + e^\beta}.$$

Hence, by applying Theorem 4, the mixing time is

$$t_{\text{mix}} \geq \frac{1}{\Phi(\mathbf{0})} = 1 + e^\beta.$$

If $\beta > \log(2^n - 1)$ instead we consider the set $R \subseteq \{0, 1\}^n$ containing all states except state $\mathbf{0}$, and observe that

$$\pi(R) = \frac{1}{Z}(2^n - 1)e^{-\beta} = \frac{(2^n - 1)e^{-\beta}}{1 + (2^n - 1)e^{-\beta}}.$$

and $\pi(R) \leq 1/2$ for $\beta > \log(2^n - 1)$. It holds that

$$Q(R, \bar{R}) = \sum_{\mathbf{x} \in R} \pi(\mathbf{x})P(\mathbf{x}, \mathbf{0}) = \sum_{\mathbf{x} \in \{0,1\}^n : |\mathbf{x}|=1} \pi(\mathbf{x})P(\mathbf{x}, \mathbf{0}) = n \frac{e^{-\beta}}{Z} \frac{1}{n} \frac{1}{1 + e^{-\beta}}.$$

The bottleneck ratio is

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

Hence, by applying Theorem 4, the mixing time is

$$t_{\text{mix}} \geq \frac{1}{\Phi(R)} > 2^n - 1.$$

□

In the next theorem we give upper bounds on the mixing time depending on the value of β . The theorem shows that, if $\beta \leq c \log n$ for some constant c , the mixing time is effectively polynomial in n with degree depending on c . The use of the path coupling technique in the proof of the theorem requires a careful choice of the edge-weights.

Theorem 12 (Upper bound on mixing time) *The mixing time of the logit dynamics for the OR game is $\mathcal{O}(n^{5/2}2^n)$ for every β . Moreover, for small values of β the mixing time is*

1. $\mathcal{O}(n \log n)$ if $\beta < (1 - \varepsilon) \log n$, for an arbitrary small constant $\varepsilon > 0$;
2. $\mathcal{O}(n^{c+3} \log n)$ if $\beta \leq c \log n$, where $c \geq 1$ is an arbitrary constant.

Proof. We apply the path coupling technique (see Theorem 2 in Section 2) with the Hamming graph defined in (5). Let $\mathbf{x}, \mathbf{y} \in \{0, 1\}^n$ be two profiles differing only at player $j \in [n]$ and, without loss of generality, let us assume $|\mathbf{x}| = k - 1$ and $|\mathbf{y}| = k$ for some $k = 1, \dots, n$. We set the weight of edge $\{\mathbf{x}, \mathbf{y}\}$ depending only on k , i.e. $\ell(\mathbf{x}, \mathbf{y}) = \delta_k$ where $\delta_k \geq 1$ will be chosen later. Consider the coupling defined in Subsection 3.2.

Now we evaluate the expected distance after one step $\mathbf{E}_{\mathbf{x}, \mathbf{y}}[\rho(X_1, Y_1)]$ of the two coupled chains (X_t, Y_t) starting at (\mathbf{x}, \mathbf{y}) . Let i be the player chosen for the update. Observe that if $i = j$, i.e. if we update the player where \mathbf{x} and \mathbf{y} are different (this holds with probability $1/n$), then the distance after one step is zero, otherwise we distinguish four cases depending on the value of k .

Case $k = 1$: In this case profile \mathbf{x} is all zeros and profile \mathbf{y} has only one 1 and the length of edge $\{\mathbf{x}, \mathbf{y}\}$ is $\ell(\mathbf{x}, \mathbf{y}) = \delta_1$. When choosing a player $i \neq j$ (this happens with probability $(n - 1)/n$), at the next step the two chains will be at distance δ_1 (if in both chains player i chooses strategy 0, and this holds with probability $\min\{\sigma_i(0 | \mathbf{x}), \sigma_i(0 | \mathbf{y})\}$), or at distance δ_2 (if in both chains player i chooses strategy 1, and

this holds with probability $\min\{\sigma_i(1|\mathbf{x}), \sigma_i(1|\mathbf{y})\}$, or at distance $\delta_1 + \delta_2$ (if player i chooses strategy 0 in chain X_1 and strategy 1 in chain Y_1 , and this holds with the remaining probability). Notice that, from the definition of the coupling, it will never happen that player i chooses strategy 1 in chain X_1 and strategy 0 in chain Y_1 , indeed we have that

$$\min\{\sigma_i(0|\mathbf{x}), \sigma_i(0|\mathbf{y})\} = \sigma_i(0|\mathbf{y}) = \frac{1}{2} \quad \text{and} \quad \min\{\sigma_i(1|\mathbf{x}), \sigma_i(1|\mathbf{y})\} = \sigma_i(1|\mathbf{x}) = \frac{1}{1+e^\beta}. \quad (10)$$

Hence the expected distance after one step is

$$\begin{aligned} \mathbf{E}_{\mathbf{x},\mathbf{y}}[\rho(X_1, Y_1)] &= \frac{n-1}{n} \left(\frac{1}{2}\delta_1 + \frac{1}{1+e^\beta}\delta_2 + \left(1 - \frac{1}{2} - \frac{1}{1+e^\beta}\right) (\delta_1 + \delta_2) \right) \\ &= \frac{n-1}{n} \left(\frac{\delta_1}{1+e^{-\beta}} + \frac{\delta_2}{2} \right). \end{aligned} \quad (11)$$

Case $k=2$: In this case we have $x_j = 0$ and $y_j = 1$, there is another player $h \in [n] \setminus \{j\}$ where $x_h = y_h = 1$, and for all the other players $i \in [n] \setminus \{j, h\}$ it holds $x_i = y_i = 0$. Hence the length of edge $\{\mathbf{x}, \mathbf{y}\}$ is $\ell(\mathbf{x}, \mathbf{y}) = \delta_2$.

When player h is chosen (this holds with probability $1/n$) we have that $\sigma_h(s|\mathbf{x})$ and $\sigma_h(s|\mathbf{y})$ for $s = 0, 1$ are the same as in (10). At the next step the two chains will be at distance δ_2 (if player h stays at strategy 1 in both chains), or at distance δ_1 (if player h chooses strategy 0 in both chains), or at distance $\delta_1 + \delta_2$ (if player h stays at strategy 0 in chain X_1 and chooses strategy 1 in chain Y_1).

When a player $i \notin \{h, j\}$ is chosen (this holds with probability $(n-2)/n$) we have that $\sigma_i(0, \mathbf{x}) = \sigma_i(1, \mathbf{x}) = \sigma_i(0, \mathbf{y}) = \sigma_i(1, \mathbf{y}) = 1/2$. Thus in this case the two coupled chains always perform the same choice at player i , and at the next step they will be at distance δ_2 (if player i stays at strategy 0 in both chains) or at distance δ_3 (if player i chooses strategy 1 in both chains).

Hence the expected distance after one step is

$$\begin{aligned} \mathbf{E}_{\mathbf{x},\mathbf{y}}[\rho(X_1, Y_1)] &= \frac{1}{n} \left(\frac{1}{2}\delta_1 + \frac{1}{1+e^\beta}\delta_2 + \left(1 - \frac{1}{2} - \frac{1}{1+e^\beta}\right) (\delta_1 + \delta_2) \right) + \frac{n-2}{n} \left(\frac{1}{2}\delta_2 + \frac{1}{2}\delta_3 \right) \\ &= \frac{1}{2n} \left(\frac{2}{1+e^{-\beta}}\delta_1 + (n-1)\delta_2 + (n-2)\delta_3 \right). \end{aligned} \quad (12)$$

Case $3 \leq k \leq n-1$: When a player $i \neq j$ is chosen such that $x_i = y_i = 1$ (this holds with probability $(k-1)/n$) then at the next step the two chains will be at distance δ_k (if i stays at strategy 1) or at distance δ_{k-1} (if i moves to strategy 0). When a player $i \neq j$ is chosen such that $x_i = y_i = 0$ (this holds with probability $(n-k)/n$) then at the next step the two chains will be at distance δ_k (if i chooses to stay at strategy 0) or at distance δ_{k+1} (if i chooses to move to strategy 0). Hence the expected distance after one step is

$$\begin{aligned} \mathbf{E}_{\mathbf{x},\mathbf{y}}[\rho(X_1, Y_1)] &= \frac{k-1}{n} \left(\frac{1}{2}\delta_k + \frac{1}{2}\delta_{k-1} \right) + \frac{n-k}{n} \left(\frac{1}{2}\delta_k + \frac{1}{2}\delta_{k+1} \right) \\ &= \frac{1}{2n} ((n-1)\delta_k + (k-1)\delta_{k-1} + (n-k)\delta_{k+1}). \end{aligned} \quad (13)$$

Case $k=n$: When a player $i \neq j$ is chosen, then at the next step the two chains will be at distance δ_n or at distance δ_{n-1} . Hence the expected distance after one step is

$$\mathbf{E}_{\mathbf{x},\mathbf{y}}[\rho(X_1, Y_1)] = \frac{n-1}{n} \left(\frac{1}{2}\delta_n + \frac{1}{2}\delta_{n-1} \right) = \frac{n-1}{2n} (\delta_n + \delta_{n-1}). \quad (14)$$

In order to apply Theorem 2 we now have to show that it is possible to choose the edge weights $\delta_1, \dots, \delta_n$ and a parameter $\alpha > 0$ such that

$$\begin{aligned} \frac{n-1}{n} \left(\frac{\delta_1}{1+e^{-\beta}} + \frac{\delta_2}{2} \right) &\leq \delta_1 e^{-\alpha}, \\ \frac{1}{2n} \left(\frac{2}{1+e^{-\beta}}\delta_1 + (n-1)\delta_2 + (n-2)\delta_3 \right) &\leq \delta_2 e^{-\alpha}, \\ \frac{1}{2n} ((n-1)\delta_k + (k-1)\delta_{k-1} + (n-k)\delta_{k+1}) &\leq \delta_k e^{-\alpha}, \quad \text{for } k = 3, \dots, n-1, \\ \frac{n-1}{2n} (\delta_n + \delta_{n-1}) &\leq \delta_n e^{-\alpha}. \end{aligned} \quad (15)$$

For different values of β , we make different choices for α and for the weights δ_k . For clarity's sake we split the proof in three different lemmas. We denote by δ^{\max} the largest δ_k .

In Lemma 13 we show that Inequalities (15) are satisfied for every value of β by choosing the weights as follows

$$\delta_k = \begin{cases} \frac{1}{2}[(n-1)\delta_2 + 1], & \text{if } k = 1; \\ \frac{n-k}{k}\delta_{k+1} + 1, & \text{if } 2 \leq k \leq n-1; \\ 1, & \text{if } k = n; \end{cases}$$

and by setting $\alpha = 1/(2n\delta^{\max})$. From Corollary 18, we have $\delta^{\max} = \mathcal{O}(\sqrt{n}2^n)$. Observe that the diameter of the Hamming graph is $\sum_{i=1}^n \delta_i \leq n\delta^{\max}$, hence from Theorem 2 we obtain $t_{\text{mix}} = \mathcal{O}(n^{5/2}2^n)$.

In Lemma 14 we show that, if $\beta < (1-\epsilon)\log n$ for an arbitrarily small constant $\epsilon > 0$, Inequalities (15) are satisfied, for sufficiently large n , by choosing weights $\delta_1 = n^{1-\epsilon}$, $\delta_2 = 4/3$, $\delta_3 = \dots = \delta_n = 1$, and $\alpha = 1/n$. In this case the diameter is $\mathcal{O}(n)$ and, by Theorem 2, $t_{\text{mix}} = \mathcal{O}(n \log n)$.

In Lemma 15 we show that, Inequalities (15) are satisfied by choosing weights as follows

$$\delta_k = \begin{cases} \frac{1+e^{-\beta}}{2} \left[\frac{a_1}{b_1} \delta_2 + 1 \right], & \text{if } k = 1; \\ \frac{a_k}{b_k} \delta_{k+1} + 1, & \text{if } 2 \leq k \leq n-1; \\ 1, & \text{if } k = n; \end{cases}$$

where $a_1 = n-1$ and $b_1 = ne^{-\beta} + 1$ and, for every $k = 2, \dots, n-1$

$$a_k = (n-k)b_{k-1} \quad b_k = (n+1)b_{k-1} - (k-1)a_{k-1};$$

and by setting $\alpha = 1/(2n\delta^{\max})$. From Corollary 21 it follows that, if $\beta \leq c \log n$ for a constant $c \in \mathbb{N}$, we have that $\delta^{\max} = \mathcal{O}(n^{c+2})$ and the diameter of the Hamming graph is $\mathcal{O}(n^{c+3})$. Thus, by Theorem 2 it follows that $t_{\text{mix}} = \mathcal{O}(n^{c+3} \log n)$. \square

6.1 Technical lemmas

In this section we prove the technical lemmas needed for completing the proof of Theorem 12.

Lemma 13 *Let $\delta_1, \dots, \delta_n$ be as follows*

$$\delta_k = \begin{cases} \frac{1}{2}[(n-1)\delta_2 + 1], & \text{if } k = 1; \\ \frac{n-k}{k}\delta_{k+1} + 1, & \text{if } 2 \leq k \leq n-1; \\ 1, & \text{if } k = n; \end{cases} \quad (16)$$

and let $\alpha = 1/(2n\delta^{\max})$ where $\delta^{\max} = \max\{\delta_k : k = 1, \dots, n\}$. Then Inequalities (15) are satisfied for every $\beta \geq 0$.

Proof. Observe that, for every $k = 1, \dots, n$, the right-hand side of the k -th inequality in (15) is

$$\delta_k e^{-\alpha} = \delta_k e^{-1/(2n\delta^{\max})} \geq \delta_k \left(1 - \frac{1}{2n\delta^{\max}} \right) = \delta_k - \frac{\delta_k}{2n\delta^{\max}} \geq \delta_k - \frac{1}{2n}. \quad (17)$$

Now we check that the left-hand side is at most $\delta_k - 1/(2n)$.

First inequality ($k = 1$): $\frac{n-1}{n} \left(\frac{\delta_1}{1+e^{-\beta}} + \frac{\delta_2}{2} \right) \leq \delta_1 e^{-\alpha}$.

From the definition of δ_1 in (16) we have that

$$\delta_2 = \frac{2\delta_1 - 1}{n-1}.$$

Hence the left-hand side is

$$\frac{n-1}{n} \left(\frac{\delta_1}{1+e^{-\beta}} + \frac{\delta_2}{2} \right) \leq \frac{n-1}{n} \left(\delta_1 + \frac{\delta_2}{2} \right) = \frac{n-1}{n} \left(\delta_1 + \frac{2\delta_1 - 1}{2(n-1)} \right) = \frac{1}{2n} (2n\delta_1 - 1) = \delta_1 - \frac{1}{2n}.$$

Second inequality ($k = 2$): $\frac{1}{2n} \left(\frac{2}{1+e^{-\beta}} \delta_1 + (n-1)\delta_2 + (n-2)\delta_3 \right) \leq \delta_2 e^{-\alpha}$.
From the definition of δ_2 in (16) we have that

$$\delta_3 = \frac{2}{n-2}(\delta_2 - 1).$$

Hence the left-hand side of the second inequality is

$$\begin{aligned} \frac{1}{2n} \left(\frac{2}{1+e^{-\beta}} \delta_1 + (n-1)\delta_2 + (n-2)\delta_3 \right) &\leq \frac{1}{2n} (2\delta_1 + (n-1)\delta_2 + (n-2)\delta_3) \\ &= \frac{1}{2n} ((n-1)\delta_2 + 1 + (n-1)\delta_2 + 2(\delta_2 - 1)) \\ &= \frac{1}{2n} (2n\delta_2 - 1) = \delta_2 - \frac{1}{2n}. \end{aligned}$$

Other inequalities ($k = 3, \dots, n-1$): $\frac{1}{2n} ((n-1)\delta_k + (k-1)\delta_{k-1} + (n-k)\delta_{k+1}) \leq \delta_k e^{-\alpha}$.
From the definition of δ_k in (16) we have that

$$\delta_{k+1} = \frac{k}{n-k}(\delta_k - 1).$$

Hence the left-hand side is

$$\begin{aligned} \frac{1}{2n} ((n-1)\delta_k + (k-1)\delta_{k-1} + (n-k)\delta_{k+1}) &= \frac{1}{2n} ((n-1)\delta_k + (n-k+1)\delta_k + (k-1) + k\delta_k - k) \\ &= \frac{1}{2n} (2n\delta_k - 1) = \delta_k - \frac{1}{2n}. \end{aligned}$$

Last inequality ($k = n$): $\frac{n-1}{2n}(\delta_n + \delta_{n-1}) \leq \delta_n e^{-\alpha}$.

Since $\delta_n = 1$ and $\delta_{n-1} = \frac{1}{n-1}\delta_n + 1 = \frac{n}{n-1}$, the left-hand side of the last inequality is

$$\frac{n-1}{2n}(\delta_n + \delta_{n-1}) = \frac{n-1}{2n} \left(1 + \frac{n}{n-1} \right) = 1 - \frac{1}{2n}.$$

□

Lemma 14 *Let $\delta_1, \dots, \delta_n$ be as follows*

$$\delta_1 = n^{1-\varepsilon}, \delta_2 = 4/3, \delta_3 = \dots = \delta_n = 1$$

where $\varepsilon > 0$ is an arbitrary small constant and let $\alpha = 1/n$. Then Inequalities (15) are satisfied for every $\beta \leq (1-\varepsilon) \log n$ and n sufficiently large.

Proof. We check that all the inequalities in (15) are satisfied.

First inequality ($k = 1$): $\frac{n-1}{n} \left(\frac{\delta_1}{1+e^{-\beta}} + \frac{\delta_2}{2} \right) \leq \delta_1 e^{-\alpha}$.

For the left-hand side we have

$$\begin{aligned} \frac{n-1}{n} \left(\frac{\delta_1}{1+e^{-\beta}} + \frac{\delta_2}{2} \right) &= \left(1 - \frac{1}{n} \right) \left(\frac{n^{1-\varepsilon}}{1+e^{-\beta}} + \frac{2}{3} \right) \\ &\leq \left(1 - \frac{1}{n} \right) \left(\frac{n^{1-\varepsilon}}{1 + \frac{1}{n^{1-\varepsilon}}} + \frac{2}{3} \right) = \left(1 - \frac{1}{n} \right) \left(\frac{n^{2(1-\varepsilon)}}{n^{1-\varepsilon} + 1} + \frac{2}{3} \right) \\ &= \left(1 - \frac{1}{n} \right) \left(\frac{(n^{1-\varepsilon} + 1)(n^{1-\varepsilon} - 1) + 1}{n^{1-\varepsilon} + 1} + \frac{2}{3} \right) \\ &= \left(1 - \frac{1}{n} \right) \left(n^{1-\varepsilon} + \frac{1}{n^{1-\varepsilon} + 1} - \frac{1}{3} \right). \end{aligned}$$

For the right-hand side we have

$$\delta_1 e^{-\alpha} = n^{1-\varepsilon} e^{-1/n} \geq n^{1-\varepsilon} \left(1 - \frac{1}{n} \right).$$

Hence the left-hand side is smaller than the right-hand one (for n sufficiently large).

Second inequality ($k = 2$): $\frac{1}{2n} \left(\frac{2}{1+e^{-\beta}} \delta_1 + (n-1)\delta_2 + (n-2)\delta_3 \right) \leq \delta_2 e^{-\alpha}$.

For the left-hand side we have

$$\begin{aligned} \frac{1}{2n} \left(\frac{2}{1+e^{-\beta}} \delta_1 + (n-1)\delta_2 + (n-2)\delta_3 \right) &= \frac{1}{2n} \left(\frac{2}{1+e^{-\beta}} n^{1-\varepsilon} + (n-1)\frac{4}{3} + (n-2) \right) \\ &\leq \frac{1}{2n} \left(2n^{1-\varepsilon} + \frac{7}{3}n \right) = \frac{7}{6} + \frac{1}{n^\varepsilon}. \end{aligned}$$

And for the right-hand side we have

$$\delta_2 e^{-\alpha} = \frac{4}{3} e^{-1/n} \geq \frac{4}{3} \left(1 - \frac{1}{n} \right) \geq \frac{4}{3} - \frac{1}{n}.$$

Hence the left-hand side is smaller than the right-hand one (for n sufficiently large).

Third inequality ($k = 3$): $\frac{1}{2n} ((n-1)\delta_3 + 2\delta_2 + (n-3)\delta_4) \leq \delta_3 e^{-\alpha}$.

For the left-hand side we have

$$\begin{aligned} \frac{1}{2n} ((n-1)\delta_3 + 2\delta_2 + (n-3)\delta_4) &= \frac{1}{2n} \left((n-1) + 2\frac{4}{3} + (n-3) \right) \\ &= \frac{1}{2n} (2n-3) \leq \left(1 - \frac{1}{n} \right). \end{aligned}$$

And for the right-hand side we have

$$\delta_3 e^{-\alpha} = e^{-1/n} \geq \left(1 - \frac{1}{n} \right).$$

Hence the left-hand side is smaller than the right-hand one.

Other inequalities ($k \geq 4$): $\frac{1}{2n} ((n-1)\delta_k + (k-1)\delta_{k-1} + (n-k)\delta_{k+1}) \leq \delta_k e^{-\alpha}$.

Since $\delta_k = \delta_{k-1} = \delta_{k+1} = 1$ the left-hand side is equal to $\frac{n-1}{n}$ and the right-hand side is $e^{-1/n} \geq \frac{n-1}{n}$. \square

Lemma 15 *Let $\delta_1, \dots, \delta_n$ be as follows*

$$\delta_k = \begin{cases} \frac{1+e^{-\beta}}{2} \left[\frac{a_1}{b_1} \delta_2 + 1 \right], & \text{if } k = 1; \\ \frac{a_k}{b_k} \delta_{k+1} + 1, & \text{if } 2 \leq k \leq n-1; \\ 1, & \text{if } k = n; \end{cases} \quad (18)$$

where $a_1 = n-1$ and $b_1 = ne^{-\beta} + 1$ and for every $k = 2, \dots, n-1$

$$a_k = (n-k)b_{k-1} \quad \text{and} \quad b_k = (n+1)b_{k-1} - (k-1)a_{k-1},$$

and let $\alpha = 1/(2n\delta^{\max})$ where $\delta^{\max} = \max\{\delta_k : k = 1, \dots, n\}$. Then Inequalities (15) are satisfied for every $\beta \geq 0$.

Before to prove the Lemma 15 we do the following observation.

Observation 16 *Let b_k defined as in the Lemma 15. Then, for every $k \geq 2$, it holds that $b_k \geq kb_{k-1}$.*

Proof. We proceed by induction on k . The base case $k = 2$ follows from

$$b_2 = (n+1)(ne^{-\beta} + 1) - (n-1) = (n+1)ne^{-\beta} + 2 > 2(ne^{-\beta} + 1) = 2b_1.$$

Now suppose the claim holds for $k-1$, that is $b_{k-1} \geq (k-1)b_{k-2}$. Then

$$\begin{aligned} b_k &= (n+1)b_{k-1} - (k-1)a_{k-1} \\ &= (n+1)b_{k-1} - (k-1)(n-k+1)b_{k-2} \\ &\geq [(n+1) - (n-k+1)]b_{k-1} = kb_{k-1}. \end{aligned}$$

□

Proof (Lemma 15). Observe that, as in Equation (17), for every $k = 1, \dots, n$, the right-hand side of the k -th inequality in (15) is

$$\delta_k e^{-\alpha} \geq \delta_k - \frac{1}{2n}.$$

Now we check that the left-hand side is at most $\delta_k - 1/(2n)$.

First inequality ($k = 1$): $\frac{n-1}{n} \left(\frac{\delta_1}{1+e^{-\beta}} + \frac{\delta_2}{2} \right) \leq \delta_1 e^{-\alpha}$.

From the definition of δ_1 in (18) we have that

$$\delta_2 = \frac{ne^{-\beta} + 1}{n-1} \left(\frac{2\delta_1}{1+e^{-\beta}} - 1 \right).$$

Hence the left-hand side is

$$\begin{aligned} \frac{n-1}{n} \left(\frac{\delta_1}{1+e^{-\beta}} + \frac{\delta_2}{2} \right) &= \frac{n-1}{n} \left[\frac{\delta_1}{1+e^{-\beta}} + \frac{ne^{-\beta} + 1}{n-1} \left(\frac{\delta_1}{1+e^{-\beta}} - \frac{1}{2} \right) \right] \\ &= \frac{n-1}{n} \frac{\delta_1}{1+e^{-\beta}} \left(1 + \frac{ne^{-\beta} + 1}{n-1} \right) - \frac{ne^{-\beta} + 1}{2n} \\ &\leq \delta_1 - \frac{1}{2n}. \end{aligned}$$

Second inequality ($k = 2$): $\frac{1}{2n} \left(\frac{2}{1+e^{-\beta}} \delta_1 + (n-1)\delta_2 + (n-2)\delta_3 \right) \leq \delta_2 e^{-\alpha}$.

From the definition of δ_2 in (18) we have that

$$\delta_3 = \frac{b_2}{a_2} (\delta_2 - 1) = \frac{(n+1)b_1 - a_1}{(n-2)b_1} (\delta_2 - 1).$$

Hence the left-hand side is

$$\begin{aligned} \frac{1}{2n} \left(\frac{2}{1+e^{-\beta}} \delta_1 + (n-1)\delta_2 + (n-2)\delta_3 \right) &= \frac{1}{2n} \left[\left(\frac{a_1}{b_1} \delta_2 + 1 \right) + (n-1)\delta_2 + \frac{(n+1)b_1 - a_1}{b_1} (\delta_2 - 1) \right] \\ &= \delta_2 - \frac{1}{2n} \frac{nb_1 - a_1}{b_1} = \delta_2 - \frac{1}{2n} \left(n - \frac{n-1}{ne^{-\beta} + 1} \right) \\ &\leq \delta_2 - \frac{1}{2n}. \end{aligned}$$

Other inequalities ($k = 3, \dots, n-1$): $\frac{1}{2n} ((n-1)\delta_k + (k-1)\delta_{k-1} + (n-k)\delta_{k+1}) \leq \delta_k e^{-\alpha}$.

From the definition of δ_k in (18) we have that

$$\delta_{k+1} = \frac{b_k}{a_k} (\delta_k - 1) = \frac{(n+1)b_{k-1} - (k-1)a_{k-1}}{(n-k)b_{k-1}} (\delta_k - 1).$$

Hence the left-hand side is

$$\begin{aligned} \frac{1}{2n} ((n-1)\delta_k + (k-1)\delta_{k-1} + (n-k)\delta_{k+1}) &= \frac{1}{2n} \left[(n-1)\delta_k + (k-1) \left(\frac{a_{k-1}}{b_{k-1}} \delta_k + 1 \right) \right. \\ &\quad \left. + \frac{(n+1)b_{k-1} - (k-1)a_{k-1}}{b_{k-1}} (\delta_k - 1) \right] \\ &= \delta_k - \frac{1}{2n} \frac{(n-k+2)b_{k-1} - (k-1)a_{k-1}}{b_{k-1}} \\ &= \delta_k - \frac{1}{2n} \left((n-k+2) - (k-1)(n-k+1) \frac{b_{k-2}}{b_{k-1}} \right) \\ &\leq \delta_k - \frac{1}{2n}. \end{aligned}$$

where the inequality follows from the Observation 16.

Last inequality ($k = n$): $\frac{n-1}{2n}(\delta_n + \delta_{n-1}) \leq \delta_n e^{-\alpha}$.

Since $\delta_n = 1$ and $\delta_{n-1} = \frac{a_{n-1}}{b_{n-1}}\delta_n + 1 = \frac{a_{n-1}}{b_{n-1}} + 1$, the left-hand side of the last inequality is

$$\begin{aligned} \frac{n-1}{2n}(\delta_n + \delta_{n-1}) &= \frac{n-1}{2n} \left(2 + \frac{a_{n-1}}{b_{n-1}} \right) = \frac{n-1}{2n} \left(2 + \frac{b_{n-2}}{b_{n-1}} \right) \\ &\leq \frac{n-1}{2n} \left(2 + \frac{1}{n-1} \right) = 1 - \frac{1}{2n}. \end{aligned}$$

where the inequality follows from the Observation 16. \square

In order to apply the path coupling theorem, we need to bound δ_{\max} : the next observation will represent the main tool to achieve this goal.

Observation 17 Let $\delta_1, \dots, \delta_n$ be defined recursively as follows: $\delta_n = 1$ and

$$\delta_k = \gamma_k \delta_{k+1} + 1,$$

where $\gamma_k > 0$ for every $k = 1, \dots, n-1$. Let $\delta^{\max} = \max\{\delta_k : k = 1, \dots, n\}$. Then

$$\delta^{\max} \leq n \max \left\{ \prod_{i=h}^j \gamma_i : 1 \leq h \leq j \leq n-1 \right\}.$$

Proof. The observation follows from the fact that, for $k = 1, \dots, n-1$, we have

$$\delta_k = 1 + \sum_{j=k}^{n-1} \prod_{i=k}^j \gamma_i.$$

\square

Corollary 18 Let $\delta_1, \dots, \delta_n$ be defined as in Lemma 13. Then $\delta^{\max} \leq c\sqrt{n}2^n$ for a suitable constant c .

Proof. From Observation 17 and the definition of $\delta_1, \dots, \delta_n$, it holds that

$$\begin{aligned} \delta^{\max} &\leq n \max \left\{ \prod_{i=h}^j \frac{n-i}{i} : 1 \leq h \leq j \leq n \right\} \\ &\leq n \prod_{i=1}^{\lfloor n/2 \rfloor} \frac{n-i}{i} \leq n \binom{n}{\lfloor n/2 \rfloor} \leq c\sqrt{n}2^n. \end{aligned}$$

for a suitable constant c . \square

In order to bound δ_{\max} when $\delta_1, \dots, \delta_n$ are defined as in Lemma 15 and $\beta \leq c \log n$ for a constant $c \in \mathbb{N}$, we define

$$\gamma_k = \frac{a_k}{b_k} = \frac{p_k e^{-\beta} + l_k}{q_k e^{-\beta} + r_k}. \quad (19)$$

You can check that $p_1 = 0$, $q_1 = n$ and

$$p_k = (n-k)q_{k-1} \quad q_k = (n+1)q_{k-1} - (k-1)p_{k-1};$$

we notice that $p_k = (n+1)q_{k-1} - (k+1)q_{k-1} \leq q_k$ for every k . We can also prove the following simple observation about q_k .

Observation 19 For every $k \geq 1$ constant, we have $q_k \geq 2^{-k}n^k$.

Proof. We proceed by induction on k , with the base $k = 1$ being obvious. Suppose the claim holds for $k-1$, that is $q_{k-1} \geq 2^{-(k-1)}n^{k-1}$, then

$$q_k = (n+1)q_{k-1} - (k-1)p_{k-1} \geq \frac{n}{2}q_{k-1} \geq 2^{-k}n^k.$$

□

Moreover, you can check that $l_1 = n - 1$, $r_1 = 1$ and

$$l_k = (n - k)r_{k-1} \quad r_k = (n + 1)r_{k-1} - (k - 1)l_{k-1};$$

we notice that above recursion gives $l_k = (n - k)(k - 1)!$ and $r_k = k!$. Next lemma bounds γ_k defined in Equation 19.

Lemma 20 *Let $\delta_1, \dots, \delta_n$ be defined as in Lemma 15, γ_k defined as in Equation (19) and $\beta \leq c \log n$ for a constant $c \in \mathbb{N}$. Then, for sufficiently large n , it holds that*

$$\begin{cases} \gamma_k < n & \forall k; \\ \gamma_k < 1 & \text{if } k > c + 2; \\ \gamma_{c+2} = \mathcal{O}(1). \end{cases}$$

Proof. Since $p_k \leq q_k$, then $(nq_k - p_k)e^{-\beta} > 0$; instead, $l_k - nr_k = (k - 1)!(n - k - nk) < 0$. Hence we have for every k

$$\gamma_k - n = \frac{p_k e^{-\beta} + l_k}{q_k e^{-\beta} + r_k} - n = \frac{(l_k - nr_k) - (nq_k - p_k)e^{-\beta}}{q_k e^{-\beta} + r_k} < 0.$$

Inductively, we show that for every $k \geq c + 3$, we have $\gamma_k < 1$. Set $k = c + 3$: c is a constant, thus Observation 19 holds for $k - 1$; hence and since $e^{-\beta} \geq n^{-c}$, we have that

$$(q_{c+3} - p_{c+3})e^{-\beta} = [(n + 1)q_{c+2} - (c + 2)p_{c+2} - (n - c - 3)q_{c+2}]e^{-\beta} \geq 2q_{c+2}e^{-\beta} \geq 2^{-(c+1)}n^2.$$

Instead, $l_{c+3} - r_{c+3} = (c + 2)!(n - 2c - 6) \leq (c + 2)! \cdot n$. Thus,

$$\gamma_{c+3} - 1 = \frac{(l_{c+3} - r_{c+3}) - (q_{c+3} - p_{c+3})e^{-\beta}}{q_{c+3}e^{-\beta} + r_{c+3}} \leq \frac{(c + 2)! \cdot n - 2^{-(c+1)}n^2}{q_{c+3}e^{-\beta} + r_{c+3}} < 0,$$

for n sufficiently large. Now, suppose that $\gamma_{k-1} < 1$; then, we have

$$\gamma_k - 1 = \frac{a_k - b_k}{b_k} = \frac{(k - 1)a_{k-1} - (k + 1)b_{k-1}}{b_k} < 0,$$

where $a_{k-1} < b_{k-1}$ is implied by the inductive hypothesis.

In order to complete the proof, we need to show that $\gamma_{c+2} = \mathcal{O}(1)$. Similarly to the case $k = c + 3$, we obtain $(q_{c+2} - p_{c+2})e^{-\beta} \geq 2^{-c}n$ and $l_{c+2} - r_{c+2} \leq (c + 1)! \cdot n$. Hence,

$$\gamma_{c+2} \leq \frac{p_{c+2} + r_{c+2} + (c + 1)! \cdot n}{p_{c+2} + r_{c+2} + 2^{-c}n} \leq (c + 1)! \cdot 2^c = \mathcal{O}(1).$$

□

Corollary 21 *Let $\delta_1, \dots, \delta_n$ and c be defined as in Lemma 15. Then $\delta^{\max} = \mathcal{O}(n^{c+2})$.*

Proof. From Observation 17, Lemma 20 and the definition of $\delta_1, \dots, \delta_n$ it follows that

$$\begin{aligned} \delta^{\max} &\leq n \max \left\{ \prod_{i=h}^j \frac{a_i}{b_i} : 1 \leq h \leq j \leq n \right\} \\ &\leq n \prod_{i=1}^{c+2} \frac{a_i}{b_i} = \mathcal{O}(n^{c+2}). \end{aligned}$$

□

7 The XOR game

In this section we analyze the logit dynamics for another simple n -player game, the *XOR game*. The XOR game is a symmetric n -player game in which each player has two strategies, denoted by 0 and 1, and each player pays the XOR of the strategies of all players (including herself). More formally, for each $i \in [n]$, the utility function $u_i(\cdot)$ is defined as follows

$$u_i(\mathbf{x}) = \begin{cases} -1, & \text{if } \mathbf{x} \text{ has an odd number of 1's;} \\ 0, & \text{if } \mathbf{x} \text{ has an even number of 1's.} \end{cases}$$

Notice that the XOR game has 2^{n-1} Nash equilibria, namely all profiles with an even number of players playing strategy 1. Nash equilibria have social welfare 0 and profiles not in equilibria have social welfare $-n$. Observe that the XOR game is a potential game with exact potential Φ where $\Phi(\mathbf{x}) = u_i(\mathbf{x})$ for every \mathbf{x} and every $i \in [n]$. Hence, the stationary distribution is

$$\pi(\mathbf{x}) = \begin{cases} e^{-\beta}/Z, & \text{if } \mathbf{x} \text{ has an odd number of 1's;} \\ 1/Z, & \text{if } \mathbf{x} \text{ has an even number of 1's;} \end{cases}$$

where the normalizing factor is $Z = 2^{n-1}(1 + e^{-\beta})$.

Even if this game looks similar to the OR game, it exhibits a different behavior. Theorem 22 gives the stationary expected social welfare of the XOR game and we can see that, as β increases, the expected social welfare tends from below to the social welfare at the Nash equilibria. In contrast the expected social welfare of the OR game is better than the worst Nash equilibrium for all values of β . Moreover, in Theorem 23 and Theorem 24 we show that the mixing time for the XOR game is polynomial in n and exponential in β , whereas the mixing time for the OR game can be bounded independently from β .

Theorem 22 (Expected social welfare) *The stationary expected social welfare of the logit dynamics for the XOR game is $\mathbf{E}_\pi[W] = -\frac{n}{1+e^\beta}$.*

Proof. The expected social welfare is

$$\mathbf{E}_\pi[W] = \sum_{\mathbf{x} \in \{0,1\}^n} W(\mathbf{x})\pi(\mathbf{x}) = -n \cdot \frac{2^{n-1}e^{-\beta}}{2^{n-1}(1+e^{-\beta})} = -\frac{n}{1+e^\beta}.$$

□

The next theorem shows that the mixing time is exponential in β for every $\beta > 0$.

Theorem 23 (Lower bound on mixing time) *The mixing time of the logit dynamics for the XOR game is $\Omega(e^\beta)$.*

Proof. Consider the set $S \subseteq \{0,1\}^n$ containing only the state $\mathbf{0} = (0, \dots, 0)$. Observe that $\pi(\mathbf{0}) \leq 1/2$. The bottleneck ratio is

$$B(\mathbf{0}) = \frac{1}{\pi(\mathbf{0})} \sum_{\mathbf{y} \in \{0,1\}^n} \pi(\mathbf{0})P(\mathbf{0}, \mathbf{y}) = \sum_{\mathbf{y} \in \{0,1\}^n : |\mathbf{y}|=1} P(\mathbf{0}, \mathbf{y}) = n \cdot \frac{1}{n} \cdot \frac{1}{1+e^\beta}.$$

Hence, by applying Theorem 4, the mixing time is

$$t_{\text{mix}} \geq \frac{1}{B(\mathbf{0})} = 1 + e^\beta.$$

□

Finally, in the next theorem we give an almost matching upper bound to the mixing time.

Theorem 24 (Upper bound on mixing time) *The mixing time of the logit dynamics for the OR game is $\mathcal{O}(n^3 e^\beta)$.*

The theorem is proved using coupling (see Theorem 1) and proof is presented in the next sections. Specifically, we use the coupling described in Section 3.2; in Section 7.1 we show that if the coupled chains are at even distance then distance does not increase after one step of the coupling; in Section 7.2 we show that if the coupled chains are at odd distance then they get closer distance with probability independent from β ; finally, in Section 7.3 we bound the expected time needed by the two chains to coalesce and use Theorem 1 to derive an upper bound for the mixing time.

7.1 Even Hamming distance

Let X_t and Y_t be two chains coupled as described in Section 3.2. Suppose that $X_t = \mathbf{x}$, $Y_t = \mathbf{y}$, and $H(\mathbf{x}, \mathbf{y}) = 2\ell$, for $\ell > 0$. In this case, $u_i(\mathbf{x}) = u_i(\mathbf{y}) = b$ for all $i \in [n]$ and some $b \in \{-1, 0\}$.

Let i be the index selected for update and let us distinguish two cases. In the first case $x_i = y_i$ and we have

$$u_i(\mathbf{x}_{-i}, 0) = u_i(\mathbf{y}_{-i}, 0) \quad \text{and} \quad u_i(\mathbf{x}_{-i}, 1) = u_i(\mathbf{y}_{-i}, 1)$$

and thus

$$\sigma_i(0 | \mathbf{x}) = \sigma_i(0 | \mathbf{y}) \quad \text{and} \quad \sigma_i(1 | \mathbf{x}) = \sigma_i(1 | \mathbf{y}).$$

Therefore the coupling always update the strategy of player i in the same way in the two chains and thus $H(X_{t+1}, Y_{t+1}) = 2\ell$.

In the second case we have $x_i \neq y_i$ and we assume, without loss of generality, that $x_i = 0$ and $y_i = 1$. We observe that, for $b \in \{-1, 0\}$,

$$u_i(\mathbf{x}_{-i}, 0) = u_i(\mathbf{y}_{-i}, 1) = b \quad \text{and} \quad u_i(\mathbf{y}_{-i}, 0) = u_i(\mathbf{x}_{-i}, 1) = -(1 + b).$$

Therefore we have

$$\sigma_i(0 | \mathbf{x}) = \sigma_i(1 | \mathbf{y}) = \frac{1}{1 + e^{-(1+2b)\beta}} \quad \text{and} \quad \sigma_i(1 | \mathbf{x}) = \sigma_i(0 | \mathbf{y}) = \frac{1}{1 + e^{(1+2b)\beta}}$$

and thus we have three possible updates for the strategy of player i :

1. both chains update to 0 (and thus $H(X_{t+1}, Y_{t+1}) = 2\ell - 1$) with probability

$$\min \left\{ \frac{1}{1 + e^{(1+2b)\beta}}, \frac{1}{1 + e^{-(1+2b)\beta}} \right\} = \frac{1}{1 + e^\beta};$$

2. both chains update to 1 (and thus $H(X_{t+1}, Y_{t+1}) = 2\ell - 1$) with probability

$$\min \left\{ \frac{1}{1 + e^{(1+2b)\beta}}, \frac{1}{1 + e^{-(1+2b)\beta}} \right\} = \frac{1}{1 + e^\beta};$$

3. chain X and Y choose two different strategies for updating the strategy of player i (and thus $H(X_{t+1}, Y_{t+1}) = 2\ell$) with probability

$$1 - \frac{2}{1 + e^\beta}.$$

The following lemma summarizes the above observations.

Lemma 25 *Suppose that $H(X_t, Y_t) = 2\ell$, for $\ell > 0$. Then*

$$H(X_{t+1}, Y_{t+1}) = \begin{cases} 2\ell - 1, & \text{with probability } \frac{2\ell}{n} \cdot \frac{2}{1 + e^\beta}; \\ 2\ell, & \text{with probability } 1 - \frac{2\ell}{n} \cdot \frac{2}{1 + e^\beta}. \end{cases}$$

7.2 Odd Hamming distance

Let X_t and Y_t be two chains coupled as described in Section 3.2. Suppose that $X_t = \mathbf{x}$, $Y_t = \mathbf{y}$, and $H(\mathbf{x}, \mathbf{y}) = 2\ell - 1$, for $\ell > 0$. In this case we have $u_i(\mathbf{x}) = b$ and $u_i(\mathbf{y}) = -(1 + b)$ for some $b \in \{-1, 0\}$. Let i be the index selected for update and let us distinguish two cases.

In the case in which $x_i = y_i = c$ for some $c \in \{0, 1\}$, we have

$$u_i(\mathbf{x}_{-i}, c) = u_i(\mathbf{y}_{-i}, 1 - c) = b \quad \text{and} \quad u_i(\mathbf{x}_{-i}, 1 - c) = u_i(\mathbf{y}_{-i}, c) = -(1 + b).$$

Therefore

$$\sigma_i(c | \mathbf{x}) = \sigma_i(1 - c | \mathbf{y}) = \frac{1}{1 + e^{-(1+2b)\beta}} \quad \text{and} \quad \sigma_i(1 - c | \mathbf{x}) = \sigma_i(c | \mathbf{y}) = \frac{1}{1 + e^{(1+2b)\beta}}$$

and thus we have three possible updates:

1. both chains update to c (and thus $H(X_{t+1}, Y_{t+1}) = 2\ell - 1$) with probability

$$\min \left\{ \frac{1}{1 + e^{-(1+2b)\beta}}, \frac{1}{1 + e^{(1+2b)\beta}} \right\} = \frac{1}{1 + e^\beta};$$

2. both chains update to $1 - c$ (and thus $H(X_{t+1}, Y_{t+1}) = 2\ell - 1$) with probability

$$\min \left\{ \frac{1}{1 + e^{-(1+2b)\beta}}, \frac{1}{1 + e^{(1+2b)\beta}} \right\} = \frac{1}{1 + e^\beta};$$

3. chains X and Y choose two different strategies for updating the strategy player i (and thus $H(X_{t+1}, Y_{t+1}) = 2\ell$) with probability $1 - \frac{2}{1+e^\beta}$.

In the second case we have $x_i \neq y_i$ and we assume, without loss of generality, that $x_i = 0$ and $y_i = 1$. We observe that

$$u_i(\mathbf{x}_{-i}, 0) = u_i(\mathbf{y}_{-i}, 0) = b \quad \text{and} \quad u_i(\mathbf{x}_{-i}, 1) = u_i(\mathbf{y}_{-i}, 1) = -(1 + b).$$

Therefore we have

$$\sigma_i(0 | \mathbf{x}) = \sigma_i(0 | \mathbf{y}) \quad \text{and} \quad \sigma_i(1 | \mathbf{x}) = \sigma_i(1 | \mathbf{y})$$

and thus in this case $H(X_{t+1}, Y_{t+1}) = 2\ell - 2$.

The following lemma summarizes the above observations.

Lemma 26 *Suppose that $H(X_t, Y_t) = 2\ell - 1$, for $\ell > 0$. Then*

$$H(X_{t+1}, Y_{t+1}) = \begin{cases} 2\ell - 2, & \text{with probability } \frac{2\ell-1}{n}; \\ 2\ell - 1, & \text{with probability } \frac{n-2\ell+1}{n} \cdot \frac{2}{1+e^\beta}; \\ 2\ell, & \text{with probability } \frac{n-2\ell+1}{n} \left(1 - \frac{2}{1+e^\beta}\right). \end{cases}$$

7.3 Time to coalesce

We denote with τ_k the random variable indicating the first time at which the two coupled chains have distance k . More precisely,

$$\tau_k = \min\{t : H(X_t, Y_t) = k\}.$$

Therefore, $\tau_{\text{couple}} = \tau_0$ is the time needed for the two chains to coalesce. We next give a bound on the expected time $\mathbf{E}_{\mathbf{x}, \mathbf{y}}[\tau_{\text{couple}}]$ for the two chains to coalesce starting from \mathbf{x} and \mathbf{y} . If \mathbf{x} and \mathbf{y} have distance 2ℓ , we denote by μ_ℓ the expected time to reach distance $2\ell - 2$. That is,

$$\mu_\ell = \mathbf{E}_{\mathbf{x}, \mathbf{y}}[\tau_{2\ell-2}].$$

Similarly, if \mathbf{x} and \mathbf{y} have distance $2\ell - 1$, we denote by ν_ℓ the expected time to reach distance $2\ell - 2$. That is,

$$\nu_\ell = \mathbf{E}_{\mathbf{x}, \mathbf{y}}[\tau_{2\ell-2}].$$

Notice that, if $H(\mathbf{x}, \mathbf{y}) = H(\mathbf{x}', \mathbf{y}')$ then

$$\mathbf{E}_{\mathbf{x}, \mathbf{y}}[\tau_k] = \mathbf{E}_{\mathbf{x}', \mathbf{y}'}[\tau_k]$$

for all k , and thus the μ_ℓ and ν_ℓ are well defined.

From Lemma 25 and Lemma 26, we have the following relations

$$\begin{aligned} \mu_\ell &= 1 + \mu_\ell \cdot \left(1 - \frac{2\ell}{n} \cdot \frac{2}{1+e^\beta}\right) + \nu_\ell \cdot \frac{2\ell}{n} \cdot \frac{2}{1+e^\beta} \\ \nu_\ell &= 1 + \nu_\ell \cdot \frac{n-2\ell+1}{n} \cdot \frac{2}{1+e^\beta} + \mu_\ell \cdot \frac{n-2\ell+1}{n} \cdot \left(1 - \frac{2}{1+e^\beta}\right). \end{aligned}$$

Simple algebraic manipulations give

$$\nu_\ell = \frac{n}{2\ell-1} \left(1 + \frac{n-2\ell+1}{2\ell} \cdot \frac{e^\beta-1}{2} \right)$$

and

$$\begin{aligned} \mu_\ell &= \nu_\ell + \frac{n}{2\ell} \cdot \frac{1+e^\beta}{2} \\ &= \frac{n}{2\ell-1} + \frac{n}{2\ell} \left(\frac{n}{2\ell-1} \cdot \frac{e^\beta-1}{2} + 1 \right) \\ &\leq \frac{n}{2\ell-1} \left(\frac{n}{2\ell-1} \cdot \frac{e^\beta-1}{2} + 2 \right) \\ &\leq n \left(n \cdot \frac{e^\beta-1}{2} + 2 \right). \end{aligned}$$

Hence,

$$\mathbf{E}_{\mathbf{x},\mathbf{y}} [\tau_{\text{couple}}] \leq 1 + \sum_{\substack{2 \leq \ell \leq n \\ \ell \text{ even}}} \mu_\ell \leq \frac{n^2}{2} \left(n \cdot \frac{e^\beta-1}{2} + 2 \right) + 1 = \mathcal{O}(n^3 e^\beta).$$

From Markov inequality we have that

$$\mathbf{P}_{\mathbf{x},\mathbf{y}} (\tau_{\text{couple}} > t) \leq \frac{\mathbf{E}_{\mathbf{x},\mathbf{y}} [\tau_{\text{couple}}]}{t}$$

and thus, by taking $t_0 = 4\mathbf{E}_{\mathbf{x},\mathbf{y}} [\tau_{\text{couple}}]$, we have $d(t_0) \leq 1/4$. Therefore, by using Theorem 1, we have that

$$t_{\text{mix}} = \mathcal{O}(n^3 e^\beta).$$

8 Conclusions and open problems

In this paper we studied strategic games where at every run a player is selected uniformly at random and she is assumed to choose her strategy for the next run according to the *logit dynamics*: a noisy best-response dynamics where the noise level is tuned by a parameter β . Such dynamics defines a family of ergodic Markov chains, indexed by β , over the set of strategy profiles.

We proposed the stationary distribution of these Markov chains as solution concept for games where players have bounded rationality or limited knowledge about the system. Since this solution concept does not assume full rationality of agents, it avoids one of the main drawbacks of many classical equilibria concepts. Moreover, the stationary distribution of an ergodic Markov chain always exists, it is unique, and the chain converges to such a distribution from any starting state.

In order to evaluate the long-term performance of the system, on the one hand we analyzed the expected social welfare when the strategy profiles are random according to the stationary distribution, on the other hand we studied the *mixing time*, i.e. how long it takes, for a chain starting at an arbitrary profile, to get close to its stationary distribution.

In this paper we applied this approach to some simple but well-studied games with a constant number of players: the CK game, that obtains the worst Price of Anarchy bound between linear congestion games, and the 2×2 coordination games considered in the seminal paper about logit dynamics [3]. We also considered two simple n -player games, the OR game and the XOR game: the analysis of the mixing time turned out to be far from trivial even for such simple games. The above games highlight a twofold behavior: for some games, namely CK game and OR game, the mixing time can be upper bounded by a function independent of β , whereas the mixing time for the other games depends exponentially on the noise parameter β .

The main goal of our line of research is to investigate logit dynamics for notable classes of n -player games. It would also be interesting to consider variations of the logit dynamics where players update their strategies simultaneously or where the noise is not uniform between players.

We have seen that, for some games and for some values of β , the mixing time can be exponential in the number of players. When it takes such a long time to reach the stationary distribution, it would be interesting to investigate the evolution of the system in the *transient* phase of the logit dynamics.

References

- [1] Elliot Anshelevich, Anirban Dasgupta, Jon M. Kleinberg, Éva Tardos, Tom Wexler, and Tim Roughgarden. The price of stability for network design with fair cost allocation. *SIAM Journal on Computing*, 38(4):1602–1623, 2008. Preliminary version in FOCS’04.
- [2] 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 *Lecture Notes in Computer Science*, pages 545–552. Springer, 2009.
- [3] Lawrence E. Blume. The statistical mechanics of strategic interaction. *Games and Economic Behavior*, 5:387–424, 1993.
- [4] R. Bubley and M. Dyer. Path Coupling: A technique for proving rapid mixing in Markov chains. In *Proceedings of the 38th Annual Symposium on Foundations of Computer Science*, pages 223–, Washington, DC, USA, 1997. IEEE Computer Society.
- [5] Xi Chen, Xiaotie Deng, and Shang-Hua Teng. Settling the complexity of computing two-player Nash equilibria. *Journal of the ACM*, 56(3), 2009.
- [6] George Christodoulou and Elias Koutsoupias. The price of anarchy of finite congestion games. In *Proc. of the 37th Annual ACM Symposium on Theory of Computing (STOC’05)*, pages 67–73. ACM, 2005.
- [7] Constantinos Daskalakis, Paul W. Goldberg, and Christos H. Papadimitriou. The complexity of computing a Nash equilibrium. *SIAM Journal on Computing*, 39(1):195–259, 2009.
- [8] Glenn Ellison. Learning, local interaction, and coordination. *Econometrica*, 61(5):1047–1071, 1993.
- [9] Alex Fabrikant and Christos H. Papadimitriou. The complexity of game dynamics: Bgp oscillations, sink equilibria, and beyond. In *Proc. of the 19th annual ACM-SIAM symposium on Discrete algorithms (SODA’08)*, pages 844–853. ACM, 2008.
- [10] Dean P. Foster and Rakesh V. Vohra. Calibrated learning and correlated equilibrium. *Games and Economic Behavior*, 21(1-2):40–55, 1997.
- [11] Drew Fudenberg and David K. Levine. *The Theory of Learning in Games*. MIT, 1998.
- [12] Michel Goemans, Vahab Mirrokni, and Adrian Vetta. Sink equilibria and convergence. In *Proc. of the 46th Annual IEEE Symposium on Foundations of Computer Science (FOCS’05)*, pages 142–154. IEEE, 2005.
- [13] John C. Harsanyi and Reinhard Selton. *A General Theory of Equilibrium Selection in Games*. MIT Press, 1988.
- [14] Sergiu Hart and Andreu Mas-Colell. A simple adaptive procedure leading to correlated equilibrium. *Econometrica*, 68(5):1127–1150, 2000.
- [15] Elias Koutsoupias and Christos H. Papadimitriou. Worst-case equilibria. *Computer Science Review*, 3(2):65–69, 2009. Preliminary version in STACS 1999.
- [16] David Levin, Yuval Peres, and Elizabeth L. Wilmer. *Markov Chains and Mixing Times*. American Mathematical Society, 2008.

- [17] 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'09)*. IEEE, 2009.
- [18] Robert W. Rosenthal. A class of games possessing pure-strategy nash equilibria. *International Journal of Game Theory*, 2(1):6567, 1973.
- [19] Aaron Roth, Nina Balcan, Adam Kalai, and Yishay Mansour. On the equilibria of alternating move games. In *Proc. of ACM-SIAM Symposium on Discrete Algorithms (SODA'10)*, 2010.
- [20] William H. Sandholm. *Population Games and Evolutionary Dynamics*. MIT Press, 2010.
- [21] H. Peyton Young. *Individual Strategy and Social Structure: An Evolutionary Theory of Institutions*. Princeton University Press, 1998.
- [22] H. Peyton Young. The diffusion of innovations in social networks. Economics Working Paper Archive number 437, Johns Hopkins University, Department of Economics, 2000.