

Mixing Time and Stationary Expected Social Welfare of Logit Dynamics*

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

Dipartimento di Informatica e Applicazioni “Renato M. Capocelli”
{auletta,ferraioli,pasquale,giuper}@dia.unisa.it

Abstract. We study *logit dynamics* [Blume, *Games and Economic Behavior*, 1993] for strategic games. At every stage of the game a player is selected uniformly at random and she plays according to a *noisy* best-response dynamics 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. Our aim is twofold: On the one hand, we are interested in the expected social welfare when the strategy profiles are random according to the stationary distribution of the Markov chain, because we believe it gives a meaningful description of the long-term behavior of the system. 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 the stationary distribution; i.e., the *mixing time* of the chain.

In this paper we study the stationary expected social welfare for the 3-player congestion game that exhibits the worst Price of Anarchy [Christodoulou and Koutsoupias, *STOC'05*], for 2-player coordination games (the same class of games studied by Blume), and for a simple n -player game. For all these games, we give almost-tight upper and lower bounds on the mixing time of logit dynamics.

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 enters an equilibrium state, it stays there and thus it can be rightly said that an equilibrium state describes the long-term behavior of the system. In this paper we are mainly interested in *selfish* 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. The notion of a Nash equilibrium is the classical notion of equilibrium for selfish systems and it corresponds to the equilibrium induced by the *best-response* dynamics. The

* Work partially supported by EU project FRONTS and by MIUR PRIN project COGENT.

observation that selfish systems are described by their equilibrium states (that is, by the Nash equilibria) has motivated the notion of Price of Anarchy [15] (and Price of Stability [1]) and the efficiency analysis of selfish systems based on such notions.

The analysis based on Nash equilibria inherits some of the shortcomings of the concept of a Nash equilibrium. First of all, the best-response dynamics assumes that the selfish agents have complete knowledge of the current state of the system; that is, of the payoff associated with each possible choice and of the strategies chosen by the other agents. Instead, in most cases, agents have only an approximate knowledge of the system state. Moreover, in presence of multiple equilibria, it is not clear which equilibrium will be reached by the system as it may depend on the initial state of the system. The notion of Price of Anarchy solves this problem by considering 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 system it might take very long to enter 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; the equilibrium state exists and is unique for every system; independently from the starting state; the system enters the equilibrium very quickly.

In this paper, we consider *noisy best-response dynamics* in which the behavior of the agents is described by a parameter $\beta \geq 0$ (β is sometimes called the *inverse temperature*). 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 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 response 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 performance of the system at the stationary 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, [11,20,19].

Equilibrium Concepts Based on the Best-Response. In case 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 context, 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]). We point out though 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 a user is the difference between the long term average cost and average cost of the best strategy in hindsight. In the no-regret dynamics the regret of every player after t step 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 [4]) and we study its mixing time (that is, the time it takes to converge to the stationary distribution) for various games. 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 convergence time to stationarity of the logit dynamics is upper bounded by a constant independent of β . Moreover, we show that the expected social cost at stationarity is smaller than the cost of the worst Nash equilibrium for all β .
- We then study the 2×2 coordination games studied by [4]. Here we show that, under some conditions, the expected social welfare at stationarity is better than the social welfare of the worst Nash equilibrium. We give exponential in β upper and lower bounds on the convergence time to stationarity for all values of β .
- Finally, we apply our analysis to a simple n -player game, the OR-game, and give upper and lower bound on the convergence time to stationarity. In particular, we prove that for $\beta = \mathcal{O}(\log n)$ the convergence time is polynomial in n .

The *logit* dynamics has been first studied by Blume [4] who showed that, for 2×2 coordination games, the long-term behaviour 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 2×2 games and assumed 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 [21] for the logit dynamics and for more general families of graphs. Montanari and Saberi [17] gave bounds on the hitting time of the risk dominant equilibrium states for the logit dynamics in terms of some graph theoretic properties of the underlying interaction network. Asadpour and Saberi [2] studied the hitting time for a class of congestion games. We notice that none of [4,8,21] 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.

From a technical point of view, our work follows the lead of [4,8,21] 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 the inverse temperature β of the system. For $\beta \rightarrow \infty$, the *logit* dynamics tends to the best response dynamics. It should come to no surprise that for large β the mixing time could be super-polynomial.

We choose to start our study from the class of coordination games considered in [4] for which we give tight upper and lower bound on the mixing time and then look also at other 2-player games and a simple n -player game (the OR-game). Despite its game-theoretic simplicity, the analytical study of the mixing time of the Markov chain associated with the OR-game as a function of β is far from trivial. Also we notice that the results of [17] cannot be used to derive upper bounds on the mixing time as in [17] the authors give a tight estimation of the hitting time only for a specific state of the Markov chain. The mixing time instead is upper bounded by the *maximum* hitting time.

From a more conceptual point of view, our work tries (similarly to [12,9,18]) to introduce a solution concept that well models the behaviour of selfish agents, is uniquely defined for any game and is quickly reached by the game. 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 [4,8,21]), in games used to obtain tight bounds on the Price of Anarchy and on a simple multiplayer game.

Organization of the Paper. In Section 2 we formally describe the logit dynamics Markov chain for a strategic game. In Sections 3, 4, and 5 we study the stationary expected social welfare and the mixing time of the logit dynamics for CK game, coordination games, and the OR-game respectively. Due to lack of space, the proofs are omitted and are available in the full version [3]. Finally, in Section 6 we present conclusions and some open problems.

Notation. We write \bar{S} for the complementary set of a set S and $|S|$ 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., $|\mathbf{x}| = |\{i \in [n] : x_i = 1\}|$. 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. $(\mathbf{x}_{-i}, y) = (x_1, \dots, x_{i-1}, y, x_{i+1}, \dots, x_n)$. We use standard Markov chain terminology (see [16]).

2 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 [4] 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 $\beta \geq 0$ is the *inverse noise*.

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 an ergodic finite 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 two profiles differ at more than one player and it is $\frac{1}{n} \sigma_i(y_i | \mathbf{x})$ if the two profiles differ exactly at player i . More formally, we have the following definition.

Definition 1 (Logit dynamics [4]). Let $\mathcal{G} = ([n], \mathcal{S}, \mathcal{U})$ be a strategic game and let $\beta \geq 0$ be the inverse noise. 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

$$P(\mathbf{x}, \mathbf{y}) = \frac{1}{n} \sum_{i=1}^n \frac{e^{\beta u_i(\mathbf{x}_{-i}, y_i)}}{T_i(\mathbf{x})} \mathbb{I}_{\{y_j = x_j \text{ for every } j \neq i\}}. \quad (2)$$

It is easy to see that, if $([n], \mathcal{S}, \mathcal{U})$ is a potential game with exact potential 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 (the *partition function* in physicists' language). Except for the Matching Pennies example in Subsection 2.1, all the games we analyse in this paper are potential games.

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 the *mixing time* t_{mix} of the chain, i.e. the time needed to have that $P^t(\mathbf{x}, \cdot)$ is *close* to π for every initial configuration \mathbf{x} . More formally, we define

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

where $\|P^t(\mathbf{x}, \cdot) - \pi\|_{\text{TV}} = \frac{1}{2} \sum_{\mathbf{y} \in \Omega} |P^t(\mathbf{x}, \mathbf{y}) - \pi(\mathbf{y})|$ is the *total variation distance*, and we set $t_{\text{mix}} = t_{\text{mix}}(1/4)$.

2.1 An Example: Matching Pennies

As an example consider the classical *Matching Pennies* game:

		<i>H</i>	<i>T</i>	
<i>H</i>	+1, -1	-1, +1		(4)
<i>T</i>	-1, +1	+1, -1		

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

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

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

So the transition matrix (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 we named $b = \frac{1}{1+e^{-2\beta}}$ for readability sake.

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 expected stationary 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

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

3 Warm Up: A 3-Player Congestion Game

In this section we study the CK game, a simple 3-player linear congestion game introduced in [6] that exhibits the worst Price of Anarchy of the average social welfare among linear congestion games with 3 or more players. This game has two equilibria: one with social welfare -6 (which is also optimal) and one with social welfare -15 . As we shall see briefly, 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 is bounded by a constant independent of 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: $G = \{g_0, g_1, g_2\}$ and $H = \{h_0, h_1, h_2\}$. 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 cost of a player is the sum of the costs of the facilities she selected. It is easy to see that this game has two pure Nash equilibria: when every player plays strategy 0 (each player pays 2, which is optimal), and when every player plays strategy 1 (each player pays 5). The game is a congestion game, and thus a potential game with following potential function:

$$\Phi(\mathbf{x}) = \sum_{j \in GUH} \sum_{i=1}^{L_{\mathbf{x}}(j)} i$$

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

Stationary Expected Social Welfare and Mixing Time. The logit dynamics for the CK game gives the following update probabilities (see Equation (1))

$$\begin{aligned} \sigma_i(0 \mid |\mathbf{x}_{-i}| = 0) &= \frac{1}{1+e^{-4\beta}} & \sigma_i(1 \mid |\mathbf{x}_{-i}| = 0) &= \frac{1}{1+e^{4\beta}} \\ \sigma_i(0 \mid |\mathbf{x}_{-i}| = 1) &= \frac{1}{1+e^{-2\beta}} & \sigma_i(1 \mid |\mathbf{x}_{-i}| = 1) &= \frac{1}{1+e^{2\beta}} \\ \sigma_i(0 \mid |\mathbf{x}_{-i}| = 2) &= \frac{1}{2} & \sigma_i(1 \mid |\mathbf{x}_{-i}| = 2) &= \frac{1}{2}. \end{aligned}$$

It is easy to check that the following distribution is stationary for the logit dynamics:

$$\begin{aligned}\pi[(0, 0, 0)] &= \frac{e^{-6\beta}}{Z(\beta)} \\ \pi[(0, 0, 1)] &= \pi[(0, 1, 0)] = \pi[(1, 0, 0)] = \frac{e^{-10\beta}}{Z(\beta)} \\ \pi[(0, 1, 1)] &= \pi[(1, 1, 0)] = \pi[(1, 0, 1)] = \pi[(1, 1, 1)] = \frac{e^{-12\beta}}{Z(\beta)}\end{aligned}$$

where $Z(\beta) = e^{-6\beta} + 3e^{-10\beta} + 4e^{-12\beta}$. Let k be the number of players playing strategy 1; the social welfare is -6 when $k = 0$, it is -13 if $k = 1$, it is -16 if $k = 2$, and -15 when $k = 3$. Thus the stationary expected social welfare is

$$\mathbf{E}_\pi[W] = -\frac{6e^{-6\beta} + 39e^{-10\beta} + (48 + 15)e^{-12\beta}}{e^{-6\beta} + 3e^{-10\beta} + 4e^{-12\beta}} = -\frac{3[2 + 13e^{-4\beta} + 21e^{-6\beta}]}{1 + 3e^{-4\beta} + 4e^{-6\beta}}.$$

For $\beta = 0$, we have $\mathbf{E}_\pi[W] = -27/2$ which is better than the social welfare of the worst Nash equilibrium. As β tends to ∞ , $\mathbf{E}_\pi[W]$ approaches the optimal social welfare. Furthermore, we observe that $\mathbf{E}_\pi[W]$ increases with β and thus we can conclude that the long-term behavior of the logit dynamics gives a better social welfare than the worst Nash equilibrium for any $\beta \geq 0$.

Theorem 1 (Mixing time of CK game). *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 $\beta \geq 0$.*

4 Coordination Games

Coordination Games are two-player games in which the players have an advantage in selecting the same strategy. They are often used to model the spread of a new technology [21]: two players have to decide whether to adopt or not a new technology. We assume that the players would prefer choosing the same technology and that choosing the new technology is risk dominant.

We analyse the mixing time of the logit dynamics for 2×2 coordination games and compute the stationary expected social welfare of the game as a function of β . We show that, for large enough β , players will spend most of the time in the risk dominant equilibrium and the expected utility is better than the one associated with the worst Nash equilibrium. Similar results can be obtained for anti coordination games (see [3]).

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

	0	1	
0	(a, a)	(c, d)	(5)
1	(d, c)	(b, b)	

We assume that $a > d$ and $b > c$ (meaning that they prefer to coordinate) and that $a - d > b - c$ (meaning that strategy 0 is the risk dominant strategy [11])

for each player). Notice that we do not make any assumption on the relation between a and b . It is easy to see that this game is a potential game and the following function is an exact potential for it:

$$\Phi(0,0) = a - d \quad \Phi(0,1) = \Phi(1,0) = 0 \quad \Phi(1,1) = b - c.$$

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 in correspondence of one of the two equilibria and the Price of Anarchy is equal to $\max\{b/a, a/b\}$.

Stationary Expected Social Welfare and Mixing Time. The logit dynamics for the game defined by the payoffs in Table 5 gives the following update probabilities for any strategy $x \in \{0,1\}$ (see Equation (1))

$$\begin{aligned} \sigma_1(0|(x,0)) = \sigma_2(0|(0,x)) &= \frac{1}{1+e^{-(a-d)\beta}} & \sigma_1(1|(x,0)) = \sigma_2(1|(0,x)) &= \frac{1}{1+e^{(a-d)\beta}} \\ \sigma_1(0|(x,1)) = \sigma_2(0|(1,x)) &= \frac{1}{1+e^{(b-c)\beta}} & \sigma_1(1|(x,1)) = \sigma_2(1|(1,x)) &= \frac{1}{1+e^{-(b-c)\beta}}. \end{aligned}$$

Theorem 2 (Expected social welfare). *The stationary expected social welfare $\mathbf{E}_\pi[W]$ of the logit dynamics for the coordination game is*

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

The following observation gives conditions on β and the players' utility for which the expected social welfare $\mathbf{E}_\pi[W]$ obtained by the logit dynamics is better than the social welfare SW_N of the worst Nash Equilibrium.

Observation 2. *For the coordination game described in Table 5, we have*

- if $a > b$ and $b \leq \max\{\frac{a+c+d}{3}, \frac{c+d}{2}\}$ then $\mathbf{E}_\pi[W] > \text{SW}_N$ for all β ;
- if $a > b$ and $b > \max\{\frac{a+c+d}{3}, \frac{c+d}{2}\}$ then $\mathbf{E}_\pi[W] > \text{SW}_N$ for all sufficiently large β ;
- if $a < b$ and $a \leq \max\{\frac{b+c+d}{3}, \frac{c+d}{2}\}$ then $\mathbf{E}_\pi[W] > \text{SW}_N$ for all β ;
- if $a < b$ and $a > \max\{\frac{b+c+d}{3}, \frac{c+d}{2}\}$ then $\mathbf{E}_\pi[W] > \text{SW}_N$ for all sufficiently large β ;
- if $a = b$ then $\mathbf{E}_\pi[W] < \text{SW}_N$ for any β, a, c and d .

Theorem 3 (Mixing Time of Coordination Games). *The mixing time of the logit dynamics with parameter β for the coordination game of Table 5 is $\Theta(e^{(b-c)\beta})$.*

5 A Simple n -Player Game: OR-Game

In this section we consider the following simple n -player potential game that we here call *OR-game*. For the upper bound we use the path coupling technique on the Hamming graph with carefully chosen edge weights. 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 $-n$. Despite its simplicity, the analysis of the mixing time is far from trivial (see full version [3]).

In Theorem 4 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 5 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 6 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.

Theorem 4 (Expected social welfare). *The stationary expected social welfare of the logit dynamics for the OR-game is $\mathbf{E}_\pi[W] = -\alpha n$ where $\alpha = \alpha(n, \beta) = \frac{(2^n - 1)e^{-\beta}}{1 + (2^n - 1)e^{-\beta}}$.*

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

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 6 (Upper bound on mixing time). *The mixing time of the logit dynamics for the OR-game 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.

Moreover the mixing time is $\mathcal{O}(n^{5/2} 2^n)$ for every β .

6 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 a *noisy best-response*, 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 study the long-term behavior of the system by analysing the expected social welfare when the strategy profiles are random according to the stationary distribution of such chains, and we compare it with the social welfare at Nash equilibria.

In order for such analysis to be meaningful we are also interested in the *mixing time* of the chains, i.e. how long it takes, for a chain starting at an arbitrary profile, to get close to its stationary distribution. The analysis of the mixing time is usually far from trivial even for very simple games.

We study several examples of applications of this approach to games with two and three players and to a simple n -players game. We started by showing that the social welfare at stationarity for the 3-player linear congestion game that attains the maximum Price of Anarchy is larger than the social welfare of the worst Nash equilibrium. This result is made significant by the fact that, for all β , the logit dynamics converges at the stationary distribution in constant time. For 2-player coordination games the mixing time turns out to be exponential in β and we give conditions for the expected social welfare at stationarity to be smaller than the social welfare of the worst Nash equilibrium. In the n -player OR-game, the mixing time is $\mathcal{O}(n \log n)$ for β up to $\log n$; if $\beta < c \log n$ with $c > 1$ constant, the mixing time is polynomial in n with the degree depending on the constant c ; finally, for large β the mixing time is exponential in n .

We leave several questions for further investigation. For example, we would like to close gaps between upper and lower bounds for the mixing time of the OR-game. Moreover, we would like to investigate logit dynamics for notable classes of n -player games.

Acknowledgements

We thank Paolo Penna and Carmine Ventre for helpful discussions and pointers to the literature.

References

1. Anshelevich, E., Dasgupta, A., Kleinberg, J.M., Tardos, É., Wexler, T., Roughgarden, T.: The price of stability for network design with fair cost allocation. *SIAM Journal on Computing* 38(4), 1602–1623 (2008)
2. Asadpour, A., Saberi, A.: On the inefficiency ratio of stable equilibria in congestion games. In: Leonardi, S. (ed.) *WINE 2009*. LNCS, vol. 5929, pp. 545–552. Springer, Heidelberg (2009)
3. Auletta, V., Ferraioli, D., Pasquale, F., Persiano, G.: Mixing time and stationary expected social welfare of logit dynamics. *CoRR* abs/1002.3474 (2010)

4. Blume, L.E.: The statistical mechanics of strategic interaction. *Games and Economic Behavior* 5, 387–424 (1993)
5. Chen, X., Deng, X., Teng, S.-H.: Settling the complexity of computing two-player Nash equilibria. *Journal of the ACM* 56(3) (2009)
6. Christodoulou, G., Koutsoupias, E.: The price of anarchy of finite congestion games. In: *Proc. of the 37th Annual ACM Symposium on Theory of Computing (STOC 2005)*, pp. 67–73. ACM, New York (2005)
7. Daskalakis, C., Goldberg, P.W., Papadimitriou, C.H.: The complexity of computing a Nash equilibrium. *SIAM Journal on Computing* 39(1), 195–259 (2009)
8. Ellison, G.: Learning, local interaction, and coordination. *Econometrica* 61(5), 1047–1071 (1993)
9. Fabrikant, A., Papadimitriou, C.H.: The complexity of game dynamics: BGP oscillations, sink equilibria, and beyond. In: *Proc. of the 19th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2008)*, pp. 844–853. ACM, New York (2008)
10. Foster, D.P., Vohra, R.V.: Calibrated learning and correlated equilibrium. *Games and Economic Behavior* 21(1-2), 40–55 (1997)
11. Fudenberg, D., Levine, D.K.: *The Theory of Learning in Games*. MIT, Cambridge (1998)
12. Goemans, M., Mirrokni, V., Vetta, A.: Sink equilibria and convergence. In: *Proc. of the 46th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2005)*, pp. 142–154. IEEE, Los Alamitos (2005)
13. Harsanyi, J.C., Selten, R.: *A General Theory of Equilibrium Selection in Games*. MIT Press, Cambridge (1988)
14. Hart, S., Mas-Colell, A.: A simple adaptive procedure leading to correlated equilibrium. *Econometrica* 68(5), 1127–1150 (2000)
15. Koutsoupias, E., Papadimitriou, C.H.: Worst-case equilibria. *Computer Science Review* 3(2), 65–69 (2009); Preliminary version in STACS 1999
16. Levin, D., Peres, Y., Wilmer, E.L.: *Markov Chains and Mixing Times*. American Mathematical Society, Providence (2008)
17. Montanari, A., Saberi, A.: Convergence to equilibrium in local interaction games. In: *Proc. of the 50th Annual Symposium on Foundations of Computer Science (FOCS 2009)*. IEEE, Los Alamitos (2009)
18. Roth, A., Balcan, N., Kalai, A., Mansour, Y.: On the equilibria of alternating move games. In: *Proc. of ACM-SIAM Symposium on Discrete Algorithms, SODA 2010* (2010)
19. Sandholm, W.H.: *Population Games and Evolutionary Dynamics*. MIT Press, Cambridge (2010)
20. Peyton Young, H.: *Individual Strategy and Social Structure: An Evolutionary Theory of Institutions*. Princeton University Press, Princeton (1998)
21. Peyton Young, H.: The diffusion of innovations in social networks. *Economics Working Paper Archive* number 437, Johns Hopkins University, Department of Economics (2000)