

Logit Dynamics with Concurrent Updates for Local Interaction Potential Games

Vincenzo Auletta · Diodato Ferraioli ·
Francesco Pasquale · Paolo Penna · Giuseppe Persiano

Received: 30 December 2013 / Accepted: 11 November 2014 / Published online: 10 December 2014
© Springer Science+Business Media New York 2014

Abstract *Logit choice* dynamics constitute a family of randomized best response dynamics based on the *logit choice function* (McFadden in *Frontiers in econometrics*. Academic Press, New York, 1974) that models players with limited rationality and knowledge. In this paper we study the *all-logit dynamics* [also known as *simultaneous learning* (Alós-Ferrer and Netzer in *Games Econ Behav* 68(2):413–427, 2010)], where at *each* time step *all* players *concurrently* update their strategies according to the logit choice function. In the well studied (*one*-)logit dynamics (Blume in *Games Econ Behav* 5(3):387–424, 1993) instead at each step *only one* randomly chosen player is allowed to update. We study properties of the all-logit dynamics in the context of *local interaction potential games*, a class of games that has been used to model complex

Vincenzo Auletta and Giuseppe Persiano are supported by Italian MIUR under the PRIN 2010–2011 project *ARS TechnoMedia—Algorithmics for Social Technological Networks*. Diodato Ferraioli and Francesco Pasquale are supported by EU FET project MULTIPLEX 317532.

V. Auletta · G. Persiano
Università di Salerno, Salerno, Italy
e-mail: auletta@unisa.it

G. Persiano
e-mail: giuper@gmail.com

D. Ferraioli (✉) · F. Pasquale
“Sapienza” Università di Roma, Rome, Italy
e-mail: ferraioli@dis.uniroma1.it

F. Pasquale
e-mail: pasquale@di.uniroma1.it

P. Penna
LIAFA, University Paris Diderot, Paris, France
e-mail: penna@liafa.univ-paris-diderot.fr

social phenomena (Montanari and Saberi 2009; Peyton in The economy as a complex evolving system. Oxford University Press, Oxford, 2003) and physical systems (Levin et al. in Probab Theory Relat Fields 146(1–2):223–265, 2010; Martinelli in Lectures on probability theory and statistics. Springer, Berlin, 1999). In a local interaction potential game players are the vertices of a *social graph* whose edges are two-player potential games. Each player picks one strategy to be played for all the games she is involved in and the payoff of the player is the sum of the payoffs from each of the games. We prove that local interaction potential games characterize the class of games for which the all-logit dynamics is reversible. We then compare the stationary behavior of one-logit and all-logit dynamics. Specifically, we look at the expected value of a notable class of observables, that we call *decomposable* observables. We prove that the difference between the expected values of the observables at stationarity for the two dynamics depends only on the *rationality* level β and on the distance of the social graph from a bipartite graph. In particular, if the social graph is bipartite then decomposable observables have the same expected value. Finally, we show that the mixing time of the all-logit dynamics has the same twofold behavior that has been highlighted in the case of the one-logit: for some games it exponentially depends on the rationality level β , whereas for other games it can be upper bounded by a function independent from β .

Keywords Logit dynamics · Markov chain · Game dynamics · Concurrent updates · Noisy dynamics · Reversibility

1 Introduction

In the last decade, we have observed an increasing interest in understanding phenomena occurring in complex systems consisting of a large number of simple networked components that operate autonomously guided by their own objectives and influenced by the behavior of the neighbors. Even though (online) social networks are a primary example of such systems, other remarkable typical instances can be found in Economics (e.g., markets), Physics (e.g., Ising model and spin systems) and Biology (e.g., evolution of life). A common feature of these systems is that the behavior of each component depends only on the interactions with a limited number of other components (its neighbors) and these interactions are usually very simple.

Game Theory is the main tool used to model the behavior of agents that are guided by their own objective in contexts where their gains depend also on the choices made by neighboring agents. Game theoretic approaches have been often proposed for modeling phenomena in a complex social network, such as the formation of the social network itself [4, 10, 16–18, 21, 28], the formation of opinions [14, 22, 30] and the spread of innovation [38, 40] as well as in physical systems [33, 34]. Many of these models are based on *local interaction games* [39], where agents are represented as vertices on a *social graph* and the relationship between two agents is represented by a simple two-player game played on the edge joining the corresponding vertices.

We are interested in the *dynamics* that govern such phenomena and several dynamics have been studied in the literature like, for example, the best response dynamics [24],

the logit dynamics [15], fictitious play [23] or no-regret dynamics [27]. Any such dynamics can be seen as made of two components:

- *Selection rule*: by which the set of players that update their state (strategy) is determined;
- *Update rule*: by which the selected players update their strategy.

For example, the classical best response dynamics compose the *best response* update rule with a selection rule that selects one player at the time. In the best response update rule, the selected player picks the strategy that, given the current strategies of the other players, guarantees the highest utility. The Cournot dynamics [19] instead combine the best response update rule with the selection rule that selects all players. Other dynamics in which all players concurrently update their strategy are fictitious play [23] and the no-regret dynamics [27].

In this paper, we study a specific class of randomized update rules called the *logit choice function* [15,35,45] which is a type of noisy best response that models in a clean and tractable way the limited knowledge (or bounded rationality) of the players in terms of a parameter β called *inverse noise*. In similar models studied in Physics, β is the inverse of the temperature. Intuitively, a low value of β (that is, high temperature) models a noisy scenario in which players choose their strategies “nearly at random”; a high value of β (that is, low temperature) models a scenario with little noise in which players pick the strategies yielding higher payoffs with higher probability.

The logit choice function can be coupled with different selection rules so to give different dynamics. For example, in the *logit* dynamics [15] at every time step a single player is selected uniformly at random and the selected player updates her strategy according to the logit choice function. The remaining players are not allowed to revise their strategies in this time step. One of the appealing features of the logit dynamics is that it naturally describes an ergodic Markov chain. This means that the underlying Markov chain admits a *unique stationary distribution* which we take as solution concept. This distribution describes the long-run behavior of the system (which states appear more frequently over a long run). The interplay between the noise and the underlying game naturally determines the system behavior: (i) As the noise becomes “very large” the equilibrium point is “approximately” the uniform distribution; (ii) As the noise vanishes the stationary distribution concentrates on so called stochastically stable states¹ which, for certain classes of games, correspond to pure Nash equilibria [1, 15].

While the logit choice function is a very natural behavioral model for approximately rational agents, the specific selection rule that selects one single player per time step avoids any form of concurrency. Therefore a natural question arises:

What happens if *concurrent* updates are allowed?

Alos-Ferrer and Netzer [1] addressed this question by characterizing the states that the dynamics selects when the noise vanishes (more details on the results of [1] are given in the related work section).

¹ A stochastically stable state is a state that has non-zero probability as β goes to infinity [42].

Our contributions In this paper, we study how the logit choice function behaves in an extreme case of concurrency. Specifically, we couple the logit choice function with a selection rule by which *all* players update their strategies at every time step. We call such dynamics *all-logit*, as opposed to the classical (*one*-)logit dynamics in which only one player at a time is allowed to move. Roughly speaking, the all-logit are to the one-logit what the Cournot dynamics are to the best response dynamics.

We study the all-logit dynamics for local interaction potential games [9, 39, 44]. Here players are vertices of a graph, called the *social graph*, and each edge is a two-player (exact) potential game. We remark that games played on different edges by a player may be different but, nonetheless, they have the same strategy set for the player. Each player picks one strategy that is used for all of her edges and the payoff is a (weighted) sum of the payoffs obtained from each game. This class of games includes coordination games on a network [20] (see also [44] for a survey of more recent works about these games) that have been used to model the spread of innovation and of new technology in social networks [40], and the Ising model [33, 34], a model for ferromagnetism. Other examples of local interaction potential games arise when geographical interactions are embedded in classical potential games, such as routing games [41] or Cournot oligopolies [37] (see also [9] and references therein). In particular, we study the all-logit dynamics on local interaction potential games for every possible value of the inverse noise β and we are interested on properties of the original one-logit dynamics that are preserved by the all-logit.

As a warm-up, we discuss two classical two-player games (these are trivial local interaction potential games played on a graph with two vertices and one edge): the coordination game and the prisoner's dilemma. Even though for both games the stationary distribution of the one-logit and of the all-logit are quite different, we identify three similarities. First, for both games, both Markov chains are reversible. Moreover, for both games, the expected number of players playing a certain strategy at the stationarity of the all-logit is exactly the same as if the expectation was taken on the stationary distribution of the one-logit. Finally, for these games the mixing time is asymptotically the same regardless of the selection rule. In this paper we will show that none of these findings is accidental.

We first study the *reversibility* of the all-logit dynamics, an important property of stochastic processes that is useful also to obtain explicit formulas for the stationary distribution. We *characterize* the class of games for which the all-logit dynamics (that is, the Markov chain resulting from the all-logit dynamics) are reversible and it turns out that this class coincides with the class of local interaction potential games. This implies that the all-logit dynamics of all two-player potential games are reversible; whereas not all potential games have reversible all-logit dynamics. This is to be compared with the well-known result saying that one-logit dynamics of every potential game are reversible.

The tools developed for this characterization yield a closed formula for the stationary distribution of reversible all-logit dynamics. In particular, we show that the stationary distribution of the all-logit dynamics of local interaction potential games resembles the Gibbs measure, that describes the stationary distribution for the one-logit [15].

Then, we focus on the *observables* of local interaction potential games. An observable is a function of the strategy profile (that is the sequence of strategies adopted by the players) and we are interested in its expected values at stationarity for both the one-logit and the all-logit. A prominent example of observable is the difference between the number of players adopting two given strategies in a game, that we name *Diff*. For example, in a local interaction potential game modeling the spread of innovation on a social network, this observable counts the difference between the number of adopters of the new and old technology. We show that there exists a class of observables whose expectation at stationarity of the all-logit is the same as the expectation at stationarity of the one-logit as long as the social network underlying the local interaction potential game is bipartite (and thus trivially for all two-player games). This class of observables includes the *Diff* observable. We extend this result by showing that for general graphs, the extent at which the expectations of these observables differ can be upper and lower bounded by a function of β and of the distance of the social graph from a bipartite graph.

Two applications highlight the significance of this result. First, observe that in the Ising model the *Diff* observable defines the magnetic field of a magnet. It is interesting to note that the Ising game has been mainly studied for bipartite graphs (e.g., the two-dimensional and the three-dimensional lattice). Then, our result implies that, for the Ising model, the all-logit dynamics are compatible with the observations and it is arguably more natural than the one-logit (that postulate that at any given time step only one particle updates its status and that the update strategy is instantaneously propagated). Second, it has been showed that the behavior predicted by the one-logit dynamics often matches the one taken by real agents in experimental testing [3, 25]. Thus, by having that some observables of one-logit and all-logit are often comparable in magnitude, the same experiments can highlight a good predictive power also for the all-logit.

Finally, we give the first bounds on the mixing time of the all-logit. We start by giving a *general* upper bound on the mixing time of the all-logit in terms of the *cumulative utility* of the game. We then look at two specific classes of games: graphical coordination games and games with a dominant profile. For *graphical coordination games* we prove an upper bound to the mixing time that exponentially depends on β . Note that it is known [6] that the one-logit also take a time exponential in β for converging to the stationary distribution. For *games with a dominant profile* we instead prove that the mixing time can be bounded by a function independent from β . Thus, also for these games the mixing time of the all-logit has the same behavior of the one-logit mixing time [6].

Related works on logit dynamics Some previous works on the logit dynamics have focused on the identification of the *stochastically stable* states. These are states that are visited with positive probability when β grows unboundedly. Blume [15] showed that, for 2×2 coordination games, the risk dominant equilibria (see [26]) are stochastically stable. Baron et al. [11, 12] characterized the stochastically stable states for larger classes of games, among which the local interaction potential games. Alos-Ferrer and Netzer [1] extended this investigation to general selection rules (including the one-logit, termed *asynchronous learning*, and the all-logit, termed *simultaneous learning*) for general strategic games. Specifically, in [1] the authors investigate which selection rule guarantees that only Nash equilibria are stochastically stable. They proved that if

the selection rule is *regular*, that is every agent has positive probability to be the only one selected for update, then the set of stochastically stable states is a subset of the set of Nash equilibria. On the other side, they show that for non-regular selection rules, such as the all-logit, this is not the case.

Works described above look at the logit dynamics as a learning dynamics, adopted by the players to learn a Nash equilibrium (or a refinement thereof). In this work, the focus is on the logit dynamics as model to understand the spread of innovations in a social network and, more generally, to study potential games played on a network (of which spread of innovations is a special case). In this setting, the focus is on the stationary distribution of the Markov Chain induced by the dynamics for the whole range of values of β .

Ellison [20] was the first to consider dynamics to model spread of innovations in a network and Peyton Young [40] was the first to study the one-logit in this context. A general upper bound on the mixing time of the one-logit dynamics for these games is given by Berger et al. [13]. Montanari and Saberi [38] instead studied the hitting time of the highest potential configuration and relate this quantity to a connectivity property of the underlying network. Asadpour and Saberi [5] considered the same problem for congestion games. The mixing time and the metastability of the one-logit dynamics for strategic games have been studied in [6–8].

2 Definitions

In this section, we formally define the local interaction potential games and the Markov chain induced by the all-logit dynamics.

Strategic games Let $\mathcal{G} = ([n], S_1, \dots, S_n, u_1, \dots, u_n)$ be a finite normal-form strategic game. The set $[n] = \{1, \dots, n\}$ is the player set, S_i is the set of *strategies* for player $i \in [n]$, $S = S_1 \times S_2 \times \dots \times S_n$ is the set of *strategy profiles* and $u_i: S \rightarrow \mathbb{R}$ is the *utility* function of player $i \in [n]$.

We adopt the standard game-theoretic notation and denote by S_{-i} the set $S_{-i} = S_1 \times \dots \times S_{i-1} \times S_{i+1} \times \dots \times S_n$ and, for $\mathbf{x}_{-i} = (x_1, \dots, x_{i-1}, x_{i+1}, \dots, x_n) \in S_{-i}$ and $y \in S_i$, we denote by (\mathbf{x}_{-i}, y) the strategy profile $(x_1, \dots, x_{i-1}, y, x_{i+1}, \dots, x_n) \in S$. Also, for a subset $L \subseteq [n]$ and strategy profile \mathbf{x} , we denote by \mathbf{x}_L the components of \mathbf{x} corresponding to players in L .

Potential games We say that function $\Phi: S \rightarrow \mathbb{R}$ is an *exact potential* (or simply a *potential*) for game \mathcal{G} if for every $i \in [n]$ and every $\mathbf{x}_{-i} \in S_{-i}$

$$u_i(\mathbf{x}_{-i}, y) - u_i(\mathbf{x}_{-i}, z) = \Phi(\mathbf{x}_{-i}, z) - \Phi(\mathbf{x}_{-i}, y)$$

for all $y, z \in S_i$. A game \mathcal{G} that admits a potential is called a *potential game* [37].

The following is an important characterization of potential games in terms of the utilities. A *circuit* $\omega = \langle \mathbf{x}_0, \dots, \mathbf{x}_\ell \rangle$ of length ℓ is a sequence of strategy profiles such that $\mathbf{x}_0 = \mathbf{x}_\ell$, $\mathbf{x}_h \neq \mathbf{x}_k$ for $1 \leq h \neq k \leq \ell$ and, for $k = 1, \dots, \ell$, there exists player

i_k such that \mathbf{x}_{k-1} and \mathbf{x}_k differ only for player i_k . For such a circuit ω we define the utility improvement $I(\omega)$ as

$$I(\omega) = \sum_{k=1}^{\ell} [u_{i_k}(\mathbf{x}_k) - u_{i_k}(\mathbf{x}_{k-1})].$$

The following theorem then holds.

Theorem 2.1 ([37, Thm 2.8]) *A game \mathcal{G} is a potential game if and only if $I(\omega) = 0$ for all circuits of length 4.*

Local interaction potential games In a local interaction potential game \mathcal{G} , each player i , with strategy set S_i , is represented by a vertex of a graph $G = (V, E)$ (called *social graph*). For every edge $e = (i, j) \in E$ there is a two-players game \mathcal{G}_e with potential function Φ_e in which the set of strategies of endpoints are exactly S_i and S_j . We denote with u_i^e the utility function of player i in the game \mathcal{G}_e . Given a strategy profile \mathbf{x} , the utility function of player i in the local interaction potential game \mathcal{G} sets

$$u_i(\mathbf{x}) = \sum_{e=(i,j)} u_i^e(x_i, x_j).$$

It is easy to check that the function $\Phi = \sum_e \Phi_e$ is a potential function for the local interaction potential game \mathcal{G} . Note that we assume that the graph G is unweighted. However, it is immediate to see that weights do not give any modeling power.

Logit choice function We study the interaction of n players of a strategic game \mathcal{G} that update their strategy according to the *logit choice function* [15,35,45] described as follows: from profile $\mathbf{x} \in S$ player $i \in [n]$ updates her strategy to $y \in S_i$ with probability

$$\sigma_i(y | \mathbf{x}) = \frac{e^{\beta u_i(\mathbf{x}_{-i}, y)}}{\sum_{z \in S_i} e^{\beta u_i(\mathbf{x}_{-i}, z)}}. \tag{1}$$

In other words, the logit choice function leans towards strategies promising higher utility. The parameter $\beta \geq 0$ is a measure of how much the utility influences the choice of the player.

All-logit In this paper we consider the *all-logit* dynamics, by which *all* players *concurrently* update their strategy using the logit choice function. Most of the previous works have focused on dynamics where at each step *one* player is chosen uniformly at random and she updates her strategy by following the logit choice function. We call those dynamics *one-logit*, to distinguish them from the *all-logit*.

The all-logit dynamics induce a Markov chain over the set of strategy profiles whose transition probability $P(\mathbf{x}, \mathbf{y})$ from profile $\mathbf{x} = (x_1, \dots, x_n)$ to profile $\mathbf{y} = (y_1, \dots, y_n)$ is

$$P(\mathbf{x}, \mathbf{y}) = \prod_{i=1}^n \sigma_i(y_i | \mathbf{x}) = \frac{e^{\beta \sum_{i=1}^n u_i(\mathbf{x}_{-i}, y_i)}}{\prod_{i=1}^n \sum_{z \in S_i} e^{\beta u_i(\mathbf{x}_{-i}, z)}}. \quad (2)$$

Sometimes it is useful to write the transition probability from \mathbf{x} to \mathbf{y} in terms of the *cumulative utility* of \mathbf{x} with respect to \mathbf{y} defined as $U(\mathbf{x}, \mathbf{y}) = \sum_i u_i(\mathbf{x}_{-i}, y_i)$. Indeed, from the distributive property of the multiplication, it follows that

$$\prod_{i=1}^n \sum_{z \in S_i} e^{\beta u_i(\mathbf{x}_{-i}, z)} = \sum_{\mathbf{z} \in S} \prod_{i=1}^n e^{\beta u_i(\mathbf{x}_{-i}, z_i)}.$$

Hence, we can rewrite (2) as

$$P(\mathbf{x}, \mathbf{y}) = \frac{e^{\beta U(\mathbf{x}, \mathbf{y})}}{T(\mathbf{x})}, \quad (3)$$

where $T(\mathbf{x}) = \sum_{\mathbf{z} \in S} e^{\beta U(\mathbf{x}, \mathbf{z})}$. For a potential game \mathcal{G} with potential Φ , we define for each pair of profiles (\mathbf{x}, \mathbf{y}) the quantity

$$K(\mathbf{x}, \mathbf{y}) = \sum_i \Phi(\mathbf{x}_{-i}, y_i) - (n-2)\Phi(\mathbf{x}) = 2\Phi(\mathbf{x}) + \sum_i (\Phi(\mathbf{x}_{-i}, y_i) - \Phi(\mathbf{x})). \quad (4)$$

Simple algebraic manipulations show that, for a potential game, we can rewrite the transition probabilities in (3) as

$$P(\mathbf{x}, \mathbf{y}) = \frac{e^{-\beta K(\mathbf{x}, \mathbf{y})}}{\gamma_A(\mathbf{x})},$$

where $\gamma_A(\mathbf{x}) = \sum_{\mathbf{z} \in S} e^{-\beta K(\mathbf{x}, \mathbf{z})}$.

It is easy to see that a Markov chain with transition matrix (2) is ergodic. Indeed, for example, ergodicity follows from the fact that all entries of the transition matrix are strictly positive.

Reversibility, Observables, Mixing time In this work we focus on three features of the all-logit dynamics, that we formally define here.

Let \mathcal{M} be a Markov chain with transition matrix P and state set Ω . \mathcal{M} is *reversible* with respect to a distribution π if, for every pair of states $s, r \in \Omega$, the following *detailed balance condition* holds

$$\pi(s)P(s, r) = \pi(r)P(r, s). \quad (5)$$

It is easy to see that if \mathcal{M} is reversible with respect to π then π is also stationary, i.e., $\pi P = \pi$.

An *observable* O is a function $O: \Omega \rightarrow \mathbb{R}$, i.e. it is a function that assigns a value to each state of the Markov chain.

An ergodic Markov chain has a unique stationary distribution π and for every starting state s the distribution $P^t(s, \cdot)$ of the chain at time t converges to π as t goes to infinity. The *mixing time* is a measure of how long it takes to get close to the stationary distribution from the *worst-case* starting state, and it is defined as

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

where $\|P^t(s, \cdot) - \pi\|_{\text{TV}} = \frac{1}{2} \sum_{r \in \Omega} |P^t(s, r) - \pi(r)|$ is the *total variation distance*. We will usually use t_{mix} for $t_{\text{mix}}(1/4)$. We refer the reader to [32] for a more detailed description of notational conventions about Markov chains and mixing times.

3 Warm-Up: Two-Player Games

In this section, we compare the behavior of the one-logit and the all-logit dynamics for two simple two-player potential games (thus two simple local information potential games): a *coordination game* and the *Prisoner’s Dilemma*. The analysis of these games highlights that the stationary distribution of the two dynamics can significantly differ. However, it turns out that for both games the Markov chain induced by the all-logit is reversible, just as for the one-logit dynamics. More surprisingly, we see that the expected number of players taking a certain action in each one of these games is exactly the same regardless whether the expectation is taken according the stationary distribution of the all-logit or of the one-logit. Finally, we observe that the mixing time of the all-logit dynamics is asymptotically the same as the mixing time of the one-logit. Next sections will show that these results are not accidental.

Two-player coordination games These are 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 [40]: two players have to decide whether to adopt or not a new technology. Each player prefers to adopt the same technology as the other player. We denote by -1 the strategy of adopting the new technology and by $+1$ the strategy of adopting the old technology. The game is formally described by the following payoff matrix

	-	+	
-	a, a	c, d	
+	d, c	b, b	(6)

We assume that $a > d$ and $b > c$ (meaning that players prefer to coordinate) and that $a - d = b - c = \Delta$ (meaning that there is not a risk dominant strategy [23]). It is easy to see that this game is a potential game. It is well known that the stationary distribution of the one-logit of a potential game is the Gibbs distribution, that assigns to $\mathbf{x} \in S$ probability $e^{-\beta\Phi(\mathbf{x})}/Z$, where $Z = \sum_{\mathbf{x} \in S} e^{-\beta\Phi(\mathbf{x})}$ is the *partition function*.

The transition matrix of the Markov chain induced by the all-logit dynamics is

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

where $p = 1/(1 + e^{\Delta\beta})$. Observe that this transition matrix is doubly-stochastic, implying that the stationary distribution of the all-logit is uniform (and hence very different from the one-logit case). However, it is easy to check that the chain is reversible and the mixing time is $\Theta(e^{\Delta\beta})$ (as in the one-logit case). Moreover, the expected number of players adopting the new strategy at stationarity is one, both when considering the one-logit and the all-logit dynamics.

Prisoner’s Dilemma The Prisoner’s Dilemma game is described by the payoff matrix given in (6), where with -1 we denote the strategy Confess and with $+1$ the strategy Defect. Moreover, payoffs satisfy the following conditions: (i) $a > d$ (so that $--$ is a Nash equilibrium); (ii) $b < c$ (so that $++$ is not a Nash equilibrium); (iii) $2a < c + d < 2b$ (so that $++$ is the social optimum and $--$ is the worst social profile). It is easy to check that the game is a potential game.

The transition matrix of the Markov chain induced by the all-logit dynamics is

$$P = \begin{pmatrix} & \begin{matrix} -- & -+ & +- & ++ \end{matrix} \\ \begin{matrix} -- \\ -+ \\ +1 \\ ++ \end{matrix} & \begin{matrix} (1-p)^2 & p(1-p) & p(1-p) & p^2 \\ (1-p)(1-q) & p(1-q) & q(1-p) & pq \\ (1-p)(1-q) & q(1-p) & p(1-q) & pq \\ (1-q)^2 & q(1-q) & q(1-q) & q^2 \end{matrix} \end{pmatrix}$$

where we let $p = 1/(1 + e^{(a-d)\beta})$ be the probability that a player does not confess given the other player is currently confessing and $q = 1/(1 + e^{(c-b)\beta})$ be the probability that a player does not confess given the other player is currently not confessing. Note that both p and q go to 0 as β goes to infinity.

It is easy to check that the transition matrix is reversible (as for the one-logit). The stationary distribution is

$$\begin{aligned} \pi(--)&= \frac{(1-q)^2}{(1+p-q)^2} & \pi(++)&= \frac{p^2}{(1+p-q)^2} \\ \pi(-+)&= \pi(+-) = \frac{p(1-q)}{(1+p-q)^2}. \end{aligned}$$

Moreover, we can see that the mixing time is upper bounded by a constant independent of β (as for the one-logit). You may also check that the expected number of confessing prisoners is exactly the same in the stationary distribution of the one-logit and of the all-logit. Indeed,

$$2\pi(--)+\pi(-+)+\pi(+)=2\cdot\frac{1-q}{1-p-q}=2\cdot\frac{e^{\beta(a-d)}+1}{e^{\beta(a-d)}+e^{\beta(b-c)}+2},$$

and the last term turns out to be exactly the expected number of confessing prisoners according to the one-logit.

4 Reversibility and Stationary Distribution

Reversibility is an important property of Markov chains and, in general, of stochastic processes. Roughly speaking, for a reversible Markov chain the stationary frequency of transitions from a state s to a state r is equal to the stationary frequency of transitions from r to s . It is easy to see that the one-logit for a game \mathcal{G} are reversible if and only if \mathcal{G} is a potential game. This does not hold for the all-logit. Indeed, we will prove that the class of games for which the all-logit are reversible is exactly the class of local interaction potential games. The tools developed in this proof will also enable us to show that the stationary distribution π_A of the all-logit dynamics of local interaction potential games has a form that resembles the Gibbs measure, that describes the stationary distribution of the one-logit. Specifically, Corollary 4.11 shows that for any profile $\mathbf{x} \in S$

$$\pi_A(\mathbf{x}) = \frac{\sum_{\mathbf{y} \in S} e^{-\beta K(\mathbf{x}, \mathbf{y})}}{Z_A},$$

where $Z_A = \sum_{\mathbf{x}, \mathbf{y} \in S} e^{-\beta K(\mathbf{x}, \mathbf{y})}$.

4.1 Reversibility Criteria

As previously stated, a Markov chain \mathcal{M} is reversible if there exists a distribution π such that the detailed balance condition (5) is satisfied. The Kolmogorov reversibility criterion allows us to establish the reversibility of a process directly from the transition probabilities. Before stating the criterion, we introduce the following notation. A *directed path* Γ from state $s \in \Omega$ to state $r \in \Omega$ is a sequence of states $\langle s_0, s_1, \dots, s_\ell \rangle$ such that $s_0 = s$ and $s_\ell = r$. The probability $\mathbf{P}(\Gamma)$ of path Γ is defined as $\mathbf{P}(\Gamma) = \prod_{j=1}^\ell P(s_{j-1}, s_j)$. The *inverse of path* $\Gamma = \langle s_0, s_1, \dots, s_\ell \rangle$ is the path $\Gamma^{-1} = \langle s_\ell, s_{\ell-1}, \dots, s_0 \rangle$. Finally, a cycle C is simply a path from a state s to itself. We are now ready to state Kolmogorov’s reversibility criterion (see, for example, [29]).

Theorem 4.1 (Kolmogorov’s Reversibility Criterion) *An irreducible Markov chain \mathcal{M} with state space Ω and transition matrix P is reversible if and only if for every cycle C it holds that*

$$\mathbf{P}(C) = \mathbf{P}(C^{-1}).$$

The following lemma will be very useful for proving reversibility conditions for the all-logit dynamics and for stating a closed expression for its stationary distribution.

Lemma 4.2 *Let \mathcal{M} be an irreducible Markov chain with transition probability P and state space Ω . \mathcal{M} is reversible if and only if for every pair of states $s, r \in \Omega$, there exists a constant $c_{s,r}$ such that for all paths Γ from s to r , it holds that*

$$\frac{\mathbf{P}(\Gamma)}{\mathbf{P}(\Gamma^{-1})} = c_{s,r}.$$

Proof Fix $s, r \in \Omega$ and consider two paths, Γ_1 and Γ_2 , from s to r . Let C_1 and C_2 be the cycles $C_1 = \Gamma_1 \circ \Gamma_2^{-1}$ and $C_2 = \Gamma_2 \circ \Gamma_1^{-1}$, where \circ denotes the concatenation of paths. If \mathcal{M} is reversible then, by the Kolmogorov Reversibility Criterion, $\mathbf{P}(C_1) = \mathbf{P}(C_2)$. On the other hand,

$$\mathbf{P}(C_1) = \mathbf{P}(\Gamma_1) \cdot \mathbf{P}(\Gamma_2^{-1}) \quad \text{and} \quad \mathbf{P}(C_2) = \mathbf{P}(\Gamma_2) \cdot \mathbf{P}(\Gamma_1^{-1}).$$

Thus

$$\frac{\mathbf{P}(\Gamma_1)}{\mathbf{P}(\Gamma_1^{-1})} = \frac{\mathbf{P}(\Gamma_2)}{\mathbf{P}(\Gamma_2^{-1})}.$$

For the other direction, fix $s^* \in \Omega$ and, for all $s \in \Omega$, set $\tilde{\pi}(s) = c_{s^*,s}/Z$, where $Z = \sum_r c_{s^*,r}$ is the normalizing constant. Now consider any two states $s, r \in \Omega$ of \mathcal{M} , let Γ_1 be any path from s^* to s and set $\Gamma_2 = \Gamma_1 \circ \langle s, r \rangle$ (that is, Γ_2 is Γ_1 concatenated with the edge (s, r)). We have that

$$\begin{aligned} \frac{\tilde{\pi}(s)}{\tilde{\pi}(r)} &= \frac{c_{s^*,s}}{c_{s^*,r}} \\ &= \frac{\mathbf{P}(\Gamma_1)}{\mathbf{P}(\Gamma_1^{-1})} \cdot \frac{\mathbf{P}(\Gamma_2^{-1})}{\mathbf{P}(\Gamma_2)} \\ &= \frac{\mathbf{P}(\Gamma_1)}{\mathbf{P}(\Gamma_1^{-1})} \cdot \frac{\mathbf{P}(\Gamma_1^{-1}) \cdot P(r, s)}{\mathbf{P}(\Gamma_1) \cdot P(s, r)} \\ &= \frac{P(r, s)}{P(s, r)} \end{aligned}$$

and therefore \mathcal{M} is reversible with respect to $\tilde{\pi}$. □

4.2 All-Logit Reversibility Implies Potential Games

In this section, we prove that if the all-logit for a game \mathcal{G} are reversible then \mathcal{G} is a potential game.

The following lemma shows a condition on the cumulative utility of a game \mathcal{G} that is necessary and sufficient for the reversibility of the all-logit of \mathcal{G} .

Lemma 4.3 *The all-logit for game \mathcal{G} are reversible if and only if the following property holds for every $\mathbf{x}, \mathbf{y}, \mathbf{z} \in S$:*

$$U(\mathbf{x}, \mathbf{y}) - U(\mathbf{y}, \mathbf{x}) = \left(U(\mathbf{x}, \mathbf{z}) + U(\mathbf{z}, \mathbf{y}) \right) - \left(U(\mathbf{y}, \mathbf{z}) + U(\mathbf{z}, \mathbf{x}) \right). \tag{7}$$

Proof To prove the only if part, pick any three $\mathbf{x}, \mathbf{y}, \mathbf{z} \in S$ and consider paths $\Gamma_1 = \langle \mathbf{x}, \mathbf{y} \rangle$ and $\Gamma_2 = \langle \mathbf{x}, \mathbf{z}, \mathbf{y} \rangle$. From Lemma 4.2 we have that reversibility implies

$$\frac{\mathbf{P}(\Gamma_1)}{\mathbf{P}(\Gamma_1^{-1})} = \frac{\mathbf{P}(\Gamma_2)}{\mathbf{P}(\Gamma_2^{-1})}.$$

Then, by definition of $\mathbf{P}(\Gamma_i)$ and (3),

$$\frac{e^{\beta U(\mathbf{x}, \mathbf{y})}}{T(\mathbf{x})} \frac{T(\mathbf{y})}{e^{\beta U(\mathbf{y}, \mathbf{x})}} = \frac{e^{\beta U(\mathbf{x}, \mathbf{z})}}{T(\mathbf{x})} \frac{e^{\beta U(\mathbf{z}, \mathbf{y})}}{T(\mathbf{z})} \frac{T(\mathbf{y})}{e^{\beta U(\mathbf{y}, \mathbf{z})}} \frac{T(\mathbf{z})}{e^{\beta U(\mathbf{z}, \mathbf{x})}},$$

which in turn implies (7).

As for the if part, let us fix state $\mathbf{z} \in S$ and define $\tilde{\pi}(\mathbf{x}) = \frac{P(\mathbf{z}, \mathbf{x})}{Z \cdot P(\mathbf{x}, \mathbf{z})}$, where Z is the normalizing constant. For any $\mathbf{x}, \mathbf{y} \in S$, we have

$$\frac{\tilde{\pi}(\mathbf{x})}{\tilde{\pi}(\mathbf{y})} = \frac{P(\mathbf{z}, \mathbf{x})}{P(\mathbf{x}, \mathbf{z})} \cdot \frac{P(\mathbf{y}, \mathbf{z})}{P(\mathbf{z}, \mathbf{y})} = \frac{e^{\beta U(\mathbf{z}, \mathbf{x})}}{e^{\beta U(\mathbf{x}, \mathbf{z})}} \cdot \frac{e^{\beta U(\mathbf{y}, \mathbf{z})}}{e^{\beta U(\mathbf{z}, \mathbf{y})}} \cdot \frac{T(\mathbf{x})}{T(\mathbf{y})} = \frac{e^{\beta U(\mathbf{y}, \mathbf{x})}}{e^{\beta U(\mathbf{x}, \mathbf{y})}} \cdot \frac{T(\mathbf{x})}{T(\mathbf{y})} = \frac{P(\mathbf{y}, \mathbf{x})}{P(\mathbf{x}, \mathbf{y})},$$

where the first equality follows from the definition of $\tilde{\pi}$, the second and the fourth follow from (3) and the third follows from (7). Therefore, the detailed balance equation holds for $\tilde{\pi}$ and thus the Markov chain is reversible. \square

We are now ready to prove that the all-logit are reversible only for potential games.

Proposition 4.4 *If the all-logit for game \mathcal{G} are reversible then \mathcal{G} is a potential game.*

Proof We show that if the all-logit are reversible then the utility improvement $I(\omega)$ over any circuit ω of length 4 is 0. The theorem then follows from Theorem 2.1.

Consider circuit $\omega = \langle \mathbf{x}, \mathbf{z}, \mathbf{y}, \mathbf{w}, \mathbf{x} \rangle$ and let i be the player in which \mathbf{x} and \mathbf{z} differ and let j be the player in which \mathbf{z} and \mathbf{y} differ. Then \mathbf{y} and \mathbf{w} differ in player i and \mathbf{w} and \mathbf{x} differ in player j . In other words, $\mathbf{z} = (\mathbf{x}_{-i}, y_i) = (\mathbf{y}_{-j}, x_j)$ and $\mathbf{w} = (\mathbf{x}_{-i}, y_j) = (\mathbf{y}_{-i}, x_i)$. Therefore we have that

$$\begin{aligned} U(\mathbf{x}, \mathbf{y}) &= \sum_{k \neq i, j} u_k(\mathbf{x}) + u_i(\mathbf{z}) + u_j(\mathbf{w}) & U(\mathbf{y}, \mathbf{x}) &= \sum_{k \neq i, j} u_k(\mathbf{y}) + u_i(\mathbf{w}) + u_j(\mathbf{z}) \\ U(\mathbf{x}, \mathbf{z}) &= \sum_{k \neq i, j} u_k(\mathbf{x}) + u_i(\mathbf{z}) + u_j(\mathbf{x}) & U(\mathbf{z}, \mathbf{y}) &= \sum_{k \neq i, j} u_k(\mathbf{z}) + u_i(\mathbf{z}) + u_j(\mathbf{y}) \\ U(\mathbf{y}, \mathbf{z}) &= \sum_{k \neq i, j} u_k(\mathbf{y}) + u_i(\mathbf{y}) + u_j(\mathbf{z}) & U(\mathbf{z}, \mathbf{x}) &= \sum_{k \neq i, j} u_k(\mathbf{z}) + u_i(\mathbf{x}) + u_j(\mathbf{z}) \end{aligned}$$

By plugging the above expressions into (7) and rearranging terms, we obtain

$$\left(u_i(\mathbf{z}) - u_i(\mathbf{x}) \right) + \left(u_j(\mathbf{y}) - u_j(\mathbf{z}) \right) + \left(u_i(\mathbf{w}) - u_i(\mathbf{y}) \right) + \left(u_j(\mathbf{x}) - u_j(\mathbf{w}) \right) = 0$$

which shows $I(\omega) = 0$. \square

4.3 A Necessary and Sufficient Condition for All-Logit Reversibility

In the previous section we have established that the all-logit are reversible only for potential games and therefore, from now on, we only consider potential games \mathcal{G} with potential function Φ . Proposition 4.5 below gives a necessary and sufficient condition for reversibility that involves only the potential function. The condition will then be used in the next section to prove that local interaction potential games are exactly the games whose all-logit are reversible.

Proposition 4.5 *The all-logit for a game \mathcal{G} with potential Φ are reversible if and only if, for all strategy profiles $\mathbf{x}, \mathbf{y} \in S$,*

$$K(\mathbf{x}, \mathbf{y}) = K(\mathbf{y}, \mathbf{x}), \quad (8)$$

where K is as defined in (4).

Proof If $K(\mathbf{x}, \mathbf{y}) = K(\mathbf{y}, \mathbf{x})$, then

$$\sum_i \left(\Phi(\mathbf{y}_{-i}, x_i) - \Phi(\mathbf{y}) \right) - \sum_i \left(\Phi(\mathbf{x}_{-i}, y_i) - \Phi(\mathbf{x}) \right) = 2 \left(\Phi(\mathbf{x}) - \Phi(\mathbf{y}) \right).$$

Hence, for any pair of strategy profiles \mathbf{x}, \mathbf{y} we have

$$\begin{aligned} U(\mathbf{x}, \mathbf{y}) - U(\mathbf{y}, \mathbf{x}) &= n \left(\Phi(\mathbf{x}) - \Phi(\mathbf{y}) \right) + \sum_i \left(u_i(\mathbf{x}_{-i}, y_i) - u_i(\mathbf{x}) \right) \\ &\quad - \sum_i \left(u_i(\mathbf{y}_{-i}, x_i) - u_i(\mathbf{y}) \right) \\ &= n \left(\Phi(\mathbf{x}) - \Phi(\mathbf{y}) \right) + \sum_i \left(\Phi(\mathbf{y}_{-i}, x_i) - \Phi(\mathbf{y}) \right) \\ &\quad - \sum_i \left(\Phi(\mathbf{x}_{-i}, y_i) - \Phi(\mathbf{x}) \right) \\ &= (n + 2) \left(\Phi(\mathbf{x}) - \Phi(\mathbf{y}) \right). \end{aligned}$$

It is then immediate to check that (7) holds.

As for the other direction, we proceed by induction on the Hamming distance between \mathbf{x} and \mathbf{y} . Let \mathbf{x} and \mathbf{y} be two profiles at Hamming distance 1; that is, \mathbf{x} and \mathbf{y} differ in only one player, say j . This implies that $(y_j, \mathbf{x}_{-j}) = \mathbf{y}$ and $(x_j, \mathbf{y}_{-j}) = \mathbf{x}$. Moreover, for $i \neq j$, $(y_i, \mathbf{x}_{-i}) = \mathbf{x}$ and $(x_i, \mathbf{y}_{-i}) = \mathbf{y}$. Thus,

$$\begin{aligned} K(\mathbf{x}, \mathbf{y}) - K(\mathbf{y}, \mathbf{x}) &= \sum_i \left(\Phi(y_i, \mathbf{x}_{-i}) - \Phi(x_i, \mathbf{y}_{-i}) \right) - (n - 2) \left(\Phi(\mathbf{x}) - \Phi(\mathbf{y}) \right) \\ &= \left(\Phi(y_j, \mathbf{x}_{-j}) - \Phi(x_j, \mathbf{y}_{-j}) \right) + \sum_{i \neq j} \left(\Phi(y_i, \mathbf{x}_{-i}) - \Phi(x_i, \mathbf{y}_{-i}) \right) \end{aligned}$$

$$\begin{aligned}
 & - (n - 2) \left(\Phi(\mathbf{x}) - \Phi(\mathbf{y}) \right) \\
 = & \left(\Phi(\mathbf{y}) - \Phi(\mathbf{x}) \right) + (n - 1) \left(\Phi(\mathbf{x}) - \Phi(\mathbf{y}) \right) \\
 & - (n - 2) \left(\Phi(\mathbf{x}) - \Phi(\mathbf{y}) \right) = 0.
 \end{aligned}$$

Now assume that the claim holds for any pair of profiles at Hamming distance $k < n$ and let \mathbf{x} and \mathbf{y} be two profiles at distance $k + 1$. Let j be any player such that $x_j \neq y_j$ and let $\mathbf{z} = (y_j, \mathbf{x}_{-j})$: \mathbf{z} is at distance at most k from \mathbf{x} and from \mathbf{y} . Consider paths $\Gamma_1 = \langle \mathbf{x}, \mathbf{y} \rangle$ and $\Gamma_2 = \langle \mathbf{x}, \mathbf{z}, \mathbf{y} \rangle$. From Lemma 4.2 we have that reversibility implies

$$\frac{e^{\beta K(\mathbf{x}, \mathbf{y})}}{\gamma_A(\mathbf{x})} \frac{\gamma_A(\mathbf{y})}{e^{\beta K(\mathbf{y}, \mathbf{x})}} = \frac{e^{\beta K(\mathbf{x}, \mathbf{z})}}{\gamma_A(\mathbf{x})} \frac{e^{\beta K(\mathbf{z}, \mathbf{y})}}{\gamma_A(\mathbf{z})} \frac{\gamma_A(\mathbf{y})}{e^{\beta K(\mathbf{y}, \mathbf{z})}} \frac{\gamma_A(\mathbf{z})}{e^{\beta K(\mathbf{z}, \mathbf{x})}}.$$

Hence $K(\mathbf{x}, \mathbf{y}) - K(\mathbf{y}, \mathbf{x}) = \left(K(\mathbf{x}, \mathbf{z}) - K(\mathbf{z}, \mathbf{x}) \right) + \left(K(\mathbf{y}, \mathbf{z}) - K(\mathbf{z}, \mathbf{y}) \right)$ and the thesis follows from the inductive hypothesis. □

4.4 Reversibility and Local Interaction Potential Games

Here we prove that the games whose all-logit are reversible are exactly the local interaction potential games.

A potential $\Phi : S_1 \times \dots \times S_n \rightarrow \mathbb{R}$ is a *two-player potential* if there exist $u, v \in [n]$ such that, for any $\mathbf{x}, \mathbf{y} \in S$ with $x_u = y_u$ and $x_v = y_v$ we have $\Phi(\mathbf{x}) = \Phi(\mathbf{y})$. In other words, Φ is a function of only its u -th and v -th argument. An interesting fact about two-player potential games is given by the following lemma.

Lemma 4.6 *Any two-player potential satisfies (8).*

Proof Let Φ be a two-player potential and let u and v be its two players. Then we have that for $w \neq u, v$, $\Phi(y_w, \mathbf{x}_{-w}) = \Phi(\mathbf{x})$ and that $\Phi(y_u, \mathbf{x}_{-u}) = \Phi(x_v, \mathbf{y}_{-v})$ and $\Phi(y_v, \mathbf{x}_{-v}) = \Phi(x_u, \mathbf{y}_{-u})$. Thus

$$K(\mathbf{x}, \mathbf{y}) = \Phi(y_u, \mathbf{x}_{-u}) + \Phi(y_v, \mathbf{x}_{-v})$$

and $K(\mathbf{y}, \mathbf{x}) = \Phi(x_v, \mathbf{y}_{-v}) + \Phi(x_u, \mathbf{y}_{-u}) = \Phi(y_u, \mathbf{x}_{-u}) + \Phi(y_v, \mathbf{x}_{-v})$. □

We say that a potential Φ is the sum of two-player potentials if there exist N two-player potentials Φ_1, \dots, Φ_N such that $\Phi = \Phi_1 + \dots + \Phi_N$. It is easy to see that generality is not lost by further requiring that $1 \leq l \neq l' \leq N$ implies $(u_l, v_l) \neq (u_{l'}, v_{l'})$, where u_l and v_l are the two players of potential Φ_l . At every game \mathcal{G} whose potential is the sum of two-player potentials, i.e., $\Phi = \Phi_1 + \dots + \Phi_N$, we can associate a *social graph* G that has a vertex for each player of \mathcal{G} and has edge (u, v) if and only if there exists l such that potential Φ_l depends on players u and v . In other words, each game whose potential is the sum of two-player potentials is a local interaction potential game.

Observe that the sum of two potentials satisfying (8) also satisfies (8). Hence we have the following proposition.

Proposition 4.7 *The all-logit dynamics for a local interaction potential game are reversible.*

Next we prove that if an n -player potential Φ satisfies (8) then it can be written as the sum of at most $N = \binom{n}{2}$ two-player potentials, Φ_1, \dots, Φ_N and thus it represents a local interaction potential game. We do so by describing an effective procedure that constructs the N two-player potentials.

Let us fix a strategy w_i^* for each player i and denote as \mathbf{w}^* the strategy profile (w_1^*, \dots, w_n^*) . Moreover, we fix an arbitrary ordering $(u_1, v_1), \dots, (u_N, v_N)$ of the N unordered pairs of players. For a potential Φ we define the sequence $\vartheta_0, \dots, \vartheta_N$ of potentials as follows: $\vartheta_0 = \Phi$ and, for $i = 1, \dots, N$, set

$$\vartheta_i = \vartheta_{i-1} - \Phi_i \tag{9}$$

where, for $\mathbf{x} \in S$, $\Phi_i(\mathbf{x})$ is defined as

$$\Phi_i(\mathbf{x}) = \vartheta_{i-1}(x_{u_i}, x_{v_i}, \mathbf{w}_{-u_i v_i}^*).$$

Observe that, for $i = 1, \dots, N$, Φ_i is a two-player potential and its players are u_i and v_i . From Lemma 4.6, Φ_i satisfies (8). Hence, if Φ satisfies (8), then also ϑ_i , for $i = 1, \dots, N$, satisfies (8).

By summing for $i = 1, \dots, N$ in (9) we obtain

$$\sum_{i=1}^N \vartheta_i = \sum_{i=0}^{N-1} \vartheta_i - \sum_{i=1}^N \Phi_i.$$

Thus

$$\Phi - \vartheta_N = \sum_{i=1}^N \Phi_i.$$

The next two lemmas prove that, if Φ satisfies (8), then ϑ_N is identically zero. This implies that Φ is the sum of at most N non-zero two-player potentials and thus a local interaction potential game.

A ball $B(r, \mathbf{x})$ of radius $r \leq n$ centered in $\mathbf{x} \in S$ is the subset of S containing all profiles \mathbf{y} that differ from \mathbf{x} in at most r coordinates.

Lemma 4.8 *For any n -player potential function Φ and for any ordering of the pairs of players, $\vartheta_N(\mathbf{x}) = 0$ for every $\mathbf{x} \in B(2, \mathbf{w}^*)$.*

Proof We distinguish three cases based on the distance of \mathbf{x} from \mathbf{w}^* .
 $\mathbf{x} = \mathbf{w}^*$: for every $i \geq 1$, we have

$$\vartheta_i(\mathbf{w}^*) = \vartheta_{i-1}(\mathbf{w}^*) - \Phi_i(\mathbf{w}^*) = \vartheta_{i-1}(\mathbf{w}^*) - \vartheta_{i-1}(\mathbf{w}^*) = 0.$$

\mathbf{x} is at distance 1 from \mathbf{w}^* : That is, there exists $u \in [n]$ such that $\mathbf{x} = (x_u, \mathbf{w}_{-u}^*)$, with $x_u \neq w_u^*$. Let us denote by $t(u)$ the smallest t such that the t -th pair contains u . We next show that for $i \geq t(u)$, $\vartheta_i(\mathbf{x}) = 0$. Indeed, we have that if u is a component of the i -th pair then

$$\vartheta_i(\mathbf{x}) = \vartheta_{i-1}(\mathbf{x}) - \Phi_i(\mathbf{x}) = \vartheta_{i-1}(\mathbf{x}) - \vartheta_{i-1}(\mathbf{x}) = 0;$$

On the other hand, if u is not a component of the i -th pair then

$$\vartheta_i(\mathbf{x}) = \vartheta_{i-1}(\mathbf{x}) - \Phi_i(\mathbf{x}) = \vartheta_{i-1}(\mathbf{x}) - \vartheta_{i-1}(\mathbf{w}^*) = \vartheta_{i-1}(\mathbf{x});$$

\mathbf{x} is at distance 2 from \mathbf{w}^* : That is, there exist u and v such that $\mathbf{x} = (x_u, x_v, \mathbf{w}_{-uv}^*)$, with $x_u \neq w_u^*$ and $x_v \neq w_v^*$. Let t be the index of the pair (u, v) . Notice that $t \geq t(u), t(v)$. We show that $\vartheta_t(\mathbf{x}) = 0$ and that this value does not change for all $i > t$. Indeed, we have

$$\vartheta_t(\mathbf{x}) = \vartheta_{t-1}(\mathbf{x}) - \Phi_t(\mathbf{x}) = \vartheta_{t-1}(\mathbf{x}) - \vartheta_{t-1}(\mathbf{x}) = 0;$$

If instead neither of u and v belongs to the i -th pair, with $i > t$, then we have

$$\vartheta_i(\mathbf{x}) = \vartheta_{i-1}(\mathbf{x}) - \Phi_i(\mathbf{x}) = \vartheta_{i-1}(\mathbf{x}) - \vartheta_{i-1}(\mathbf{w}^*) = \vartheta_{i-1}(\mathbf{x});$$

Finally, suppose that the i -th pair, for $i > t$, contains exactly one of u and v , say u . Then we have

$$\vartheta_i(\mathbf{x}) = \vartheta_{i-1}(\mathbf{x}) - \Phi_i(\mathbf{x}) = \vartheta_{i-1}(\mathbf{x}) - \vartheta_{i-1}(x_u, \mathbf{w}_{-u}^*).$$

We conclude the proof by observing that $t(u) \leq t \leq i - 1$ and thus, by the previous case, $\vartheta_{i-1}(x_u, \mathbf{w}_{-u}^*) = 0$. □

The next lemma shows that if a potential ϑ satisfies (8) and is constant in a ball of radius 2, then it is constant everywhere.

Lemma 4.9 *Let ϑ be a function that satisfies (8). If there exist $\mathbf{x} \in S$ and $c \in \mathbb{R}$ such that $\vartheta(\mathbf{y}) = c$ for every $\mathbf{y} \in B(2, \mathbf{x})$, then $\vartheta(\mathbf{y}) = c$ for every $\mathbf{y} \in S$.*

Proof Fix $h > 2$ and suppose that $\vartheta(\mathbf{z}) = c$ for every $\mathbf{z} \in B(h - 1, \mathbf{x})$. Consider $\mathbf{y} \in B(h, \mathbf{x}) \setminus B(h - 1, \mathbf{x})$ and observe that $(y_i, \mathbf{x}_{-i}) \in B(h - 1, \mathbf{x})$ and $(x_i, \mathbf{y}_{-i}) \in B(h - 1, \mathbf{x})$ for every i such that $x_i \neq y_i$. Then, since ϑ satisfies (8), we have

$$(h - 2) (\vartheta(\mathbf{x}) - \vartheta(\mathbf{y})) = \sum_{i: x_i \neq y_i} \left(\vartheta(y_i, \mathbf{x}_{-i}) - \vartheta(x_i, \mathbf{y}_{-i}) \right) = 0,$$

that implies $\vartheta(\mathbf{y}) = \vartheta(\mathbf{x}) = c$. \square

We can thus conclude that if the all-logit of a potential game \mathcal{G} are reversible then \mathcal{G} is a local interaction potential game. By combining this result with Proposition 4.4 and Proposition 4.7, we obtain

Theorem 4.10 *The all-logit dynamics of game \mathcal{G} are reversible if and only if \mathcal{G} is a local interaction potential game.*

As a corollary of this theorem we have a closed form for the stationary distribution of the all-logit for local interaction potential games.

Corollary 4.11 (Stationary distribution) *Let \mathcal{G} be a local interaction potential game with potential function Φ . Then the stationary distribution of the all-logit for \mathcal{G} is*

$$\pi_A(\mathbf{x}) \propto \sum_{\mathbf{y} \in S} e^{-\beta K(\mathbf{x}, \mathbf{y})}. \quad (10)$$

Proof Fix any profile \mathbf{y} . The detailed balance equation gives for every $\mathbf{x} \in S$

$$\frac{\pi_A(\mathbf{x})}{\pi_A(\mathbf{y})} = \frac{P(\mathbf{y}, \mathbf{x})}{P(\mathbf{x}, \mathbf{y})} = e^{\beta(K(\mathbf{x}, \mathbf{y}) - K(\mathbf{y}, \mathbf{x}))} \frac{\gamma_A(\mathbf{x})}{\gamma_A(\mathbf{y})}.$$

By Proposition 4.5 we have

$$\pi_A(\mathbf{x}) = \gamma_A(\mathbf{x}) \cdot \frac{\pi_A(\mathbf{y})}{\gamma_A(\mathbf{y})}.$$

Since the term $\frac{\pi_A(\mathbf{y})}{\gamma_A(\mathbf{y})}$ is constant for each profile \mathbf{x} , the claim follows. \square

Note that for a local interaction potential game \mathcal{G} with potential function Φ , we write $\pi_1(\mathbf{x})$, the stationary distribution of the one-logit of \mathcal{G} , as $\pi_1(\mathbf{x}) = \gamma_1(\mathbf{x})/Z_1$ where $\gamma_1(\mathbf{x}) = e^{-\beta\Phi(\mathbf{x})}$ is the Boltzmann factor and $Z_1 = \sum_{\mathbf{x}} \gamma_1(\mathbf{x})$ is the partition function. From Corollary 4.11, we derive that $\pi_A(\mathbf{x})$, the stationary distribution of the all-logit of \mathcal{G} , can be written in similar way; that is, $\pi_A(\mathbf{x}) = \frac{\gamma_A(\mathbf{x})}{Z_A}$, where $\gamma_A(\mathbf{x}) = \sum_{\mathbf{y}} e^{-\beta K(\mathbf{x}, \mathbf{y})}$ and

$$Z_A = \sum_{\mathbf{x} \in S} \gamma_A(\mathbf{x}) = \sum_{\mathbf{x}, \mathbf{y} \in S} e^{-\beta K(\mathbf{x}, \mathbf{y})}.$$

The Z_A factor can thus be considered as the partition function of the all-logit.

5 Observables of Local Information Potential Games

In this section, we study observables of local interaction potential games and we focus on the relation between the expected value $\langle O, \pi_1 \rangle$ of an observable O at the stationarity of the one-logit and its expected value $\langle O, \pi_A \rangle$ at the stationarity of the all-logit dynamics. We start by studying invariant observables, that is, observables for which the two expected values coincide. In Theorem 5.6, we give a sufficient condition for an observable to be invariant. The sufficient condition is related to the existence of a *decomposition* of the set $S \times S$ that splits quantity K appearing in the expression for the stationary distribution of the all-logit of the local interaction potential game \mathcal{G} (see Eq. 10) into a sum of two potentials. In Theorem 5.6 we show that if \mathcal{G} admits such a decomposition μ and, in addition, observable O is also decomposed by μ (see Definition 5.2) then O has the same expected value at the stationarity of the one-logit and of the all-logit. We then go on to show that all local interaction potential games on *bipartite* social graphs admit a decomposition permutation (see Theorem 5.4) and give examples of invariant observables.

We then look at local interaction potential games \mathcal{G} on general social graphs G and show that the expected values of a decomposable observable O with respect to the stationary distributions of the one-logit and of the all-logit differ by a quantity that depends on β and on how far away the social graph G is from being bipartite (which in turn is related to the smallest eigenvalue of G [43]). This proves that the all-logit can be often taken as a close approximation of the one-logit and, hence, of those settings in which the latter gives good predictions [3, 25].

The above findings follow from a relation between the partition functions of the one-logit and of the all-logit that might be of independent interest. More precisely, in Theorem 5.1 we show that if the game \mathcal{G} admits a decomposition then the partition function of the all-logit is the square of the partition function of the one-logit. The partition function of the one-logit is easily seen to be equal to the partition function of the canonical ensemble used in Statistical Mechanics (see for example [31]). It is well known that a partition function of a canonical ensemble that is the union of two independent canonical ensembles is the product of the two partition functions. Thus Theorem 5.1 (and Corollary 5.5) can be seen as a further confirmation that the all-logit can be decomposed into two independent one-logit dynamics.

5.1 Decomposable Observables for Bipartite Social Graphs

We start by introducing the concept of a *decomposition* and we prove that for all local interaction potential games on a bipartite social graph there exists a decomposition. Then we define the concept of a *decomposable* observable and prove that a decomposable observable has the same expectation at stationarity for the one-logit and the all-logit.

Definition 5.1 A permutation

$$\mu : (\mathbf{x}, \mathbf{y}) \mapsto (\mu_1(\mathbf{x}, \mathbf{y}), \mu_2(\mathbf{x}, \mathbf{y}))$$

of $S \times S$ is a *decomposition* for a local interaction potential game \mathcal{G} with potential Φ if, for all (\mathbf{x}, \mathbf{y}) , we have that

$$K(\mathbf{x}, \mathbf{y}) = \Phi(\mu_1(\mathbf{x}, \mathbf{y})) + \Phi(\mu_2(\mathbf{x}, \mathbf{y})),$$

$$\mu_1(\mathbf{x}, \mathbf{y}) = \mu_2(\mathbf{y}, \mathbf{x}) \text{ and } \mu_2(\mathbf{x}, \mathbf{y}) = \mu_1(\mathbf{y}, \mathbf{x}).$$

Roughly speaking, a decomposition is a tool that allows us to see K as the sum of two potentials and in this way it relates the stationary distribution of the all-logit and of the one-logit. Indeed, if μ decomposes local interaction potential game \mathcal{G} then

$$\pi_A(\mathbf{x}) = \sum_{\mathbf{y}} \pi_1(\mu_1(\mathbf{x}, \mathbf{y})) \cdot \pi_1(\mu_2(\mathbf{x}, \mathbf{y})).$$

In turn, this will help us in proving the relationship between the expectation of some observables according to these stationary distributions.

We first show a relation between the partition functions of the one-logit and of the all-logit that might be of independent interest.

Theorem 5.1 *If a local interaction potential game \mathcal{G} admits a decomposition μ , then $Z_A = Z_1^2$.*

Proof From (10) and from the fact that μ is a permutation of $S \times S$, we have

$$Z_A = \sum_{\mathbf{x}, \mathbf{y}} e^{-\beta K(\mathbf{x}, \mathbf{y})} = \sum_{\mathbf{x}, \mathbf{y}} e^{-\beta[\Phi(\mu_1(\mathbf{x}, \mathbf{y})) + \Phi(\mu_2(\mathbf{x}, \mathbf{y}))]} = \sum_{\mathbf{x}, \mathbf{y}} e^{-\beta[\Phi(\mathbf{x}) + \Phi(\mathbf{y})]} = Z_1^2.$$

□

We next prove that for all local interaction potential games on a bipartite social graph there exists a decomposition. We start by showing that we can decompose $K(\mathbf{x}, \mathbf{y})$ in the contributions of each edge of the social graph G of the local interaction potential game \mathcal{G} . Specifically, for strategy profiles \mathbf{x} and \mathbf{y} and edge $e = (u, v)$ of G we define $K_e(\mathbf{x}, \mathbf{y})$ as

$$K_e(\mathbf{x}, \mathbf{y}) = \Phi_e(x_u, y_v) + \Phi_e(y_u, x_v). \tag{11}$$

Then we have the following lemma that will be useful for giving a sufficient condition for having a decomposition.

Lemma 5.2

$$K(\mathbf{x}, \mathbf{y}) = \sum_e K_e(\mathbf{x}, \mathbf{y}).$$

Proof By definition $K(\mathbf{x}, \mathbf{y}) = \sum_i \Phi(\mathbf{x}_{-i}, y_i) - (n - 2)\Phi(\mathbf{x})$. Then by expressing Φ as sum of the potential over the edges we have

$$K(\mathbf{x}, \mathbf{y}) = \sum_i \sum_e \Phi_e(\mathbf{x}_{-i}, y_i) - (n - 2)\Phi(\mathbf{x}).$$

Then observe that edge $e = (u, v)$ and each of the $(n - 2)$ vertices $i \neq u, v$ contribute $\Phi_e(x_u, x_v)$ to the sum. On the other hand, the total contribution for $e = (u, v)$ and $i = u, v$ is $K_e(\mathbf{x}, \mathbf{y}) = \Phi_e(y_u, x_v) + \Phi_e(x_u, y_v)$. Therefore we obtain

$$\begin{aligned} K(\mathbf{x}, \mathbf{y}) &= \sum_i \sum_{e=(u,v)} \Phi_e(\mathbf{x}_{-i}, y_i) - (n - 2)\Phi(\mathbf{x}) \\ &= \sum_{e=(u,v)} [(n - 2)\Phi_e(x_u, x_v) + K_e(\mathbf{x}, \mathbf{y})] - (n - 2)\Phi(\mathbf{x}) \\ &= \sum_e K_e(\mathbf{x}, \mathbf{y}). \end{aligned}$$

□

From Lemma 5.2, we then achieve the following sufficient condition for a permutation to be a decomposition.

Lemma 5.3 *Let \mathcal{G} be a local interaction potential game with potential Φ on a graph G . Consider a permutation $(\mathbf{x}, \mathbf{y}) \mapsto (\tilde{\mathbf{x}}, \tilde{\mathbf{y}})$ such that for all $\mathbf{x}, \mathbf{y} \in S$ and for each edge $e = (u, v)$ of G at least one of the following equalities holds*

$$(\tilde{x}_u, \tilde{x}_v, \tilde{y}_u, \tilde{y}_v) = (x_u, y_v, y_u, x_v), \tag{12}$$

$$(\tilde{x}_u, \tilde{x}_v, \tilde{y}_u, \tilde{y}_v) = (y_u, x_v, x_u, y_v). \tag{13}$$

Then, $K(\mathbf{x}, \mathbf{y}) = \phi(\tilde{\mathbf{x}}) + \phi(\tilde{\mathbf{y}})$.

Proof Observe that, if one of Equation (12) and (13) holds,

$$\Phi_e(\tilde{x}_u, \tilde{x}_v) + \Phi_e(\tilde{y}_u, \tilde{y}_v) = \Phi_e(x_u, y_v) + \Phi_e(y_u, x_v) = K_e(\mathbf{x}, \mathbf{y}).$$

The corollary then follows by summing over all edges e . □

We are now ready for the main result of the section.

Theorem 5.4 *Let \mathcal{G} be a local interaction potential game on a bipartite graph G . Then \mathcal{G} admits a decomposition.*

Proof Let (L, R) be the sets of vertices in which G is bipartite. For each $(\mathbf{x}, \mathbf{y}) \in S \times S$ define

$$\tilde{\mathbf{x}} = \mu_1(\mathbf{x}, \mathbf{y}) = (\mathbf{x}_L, \mathbf{y}_R) \quad \text{and} \quad \tilde{\mathbf{y}} = \mu_2(\mathbf{x}, \mathbf{y}) = (\mathbf{y}_L, \mathbf{x}_R). \tag{14}$$

First of all, observe that the mapping is an involution and thus it is also a permutation and that $\mu_1(\mathbf{x}, \mathbf{y}) = \mu_2(\mathbf{y}, \mathbf{x})$ and $\mu_2(\mathbf{x}, \mathbf{y}) = \mu_1(\mathbf{y}, \mathbf{x})$. Since G is bipartite, for every edge (u, v) exactly one endpoint is in L and exactly one is in R . If $u \in L$, then we have that $(\tilde{x}_u, \tilde{x}_v, \tilde{y}_u, \tilde{y}_v) = (x_u, y_v, y_u, x_v)$ and thus (12) is satisfied. If instead $u \in R$, then we have that $(\tilde{x}_u, \tilde{x}_v, \tilde{y}_u, \tilde{y}_v) = (y_u, x_v, x_u, y_v)$ and thus (13) is satisfied. Therefore for each edge one of (12) and (13) is satisfied. By Lemma 5.3, we can conclude that the mapping is a decomposition. □

Consider a local interaction potential game on a bipartite graph $G = (L, R, E)$ and let us denote by $L(\mathbf{x})$ (respectively, $R(\mathbf{x})$) the set of profiles agreeing with \mathbf{x} for every vertex of L (respectively, of R). That is, $L(\mathbf{x}) = \{\mathbf{y} : \mathbf{y}_L = \mathbf{x}_L\}$ and $R(\mathbf{x}) = \{\mathbf{y} : \mathbf{y}_R = \mathbf{x}_R\}$. The following corollary of Theorem 5.1 and Theorem 5.4 proves an interesting characterization of the stationary distribution of the all-logit dynamics for local interaction potential games on bipartite graphs that might be of independent interest.

Corollary 5.5 *Let \mathcal{G} be a local interaction potential game on a bipartite graph $G = (L, R, E)$. For every profile \mathbf{x} , we have*

$$\pi_A(\mathbf{x}) = \pi_1(L(\mathbf{x})) \cdot \pi_1(R(\mathbf{x})).$$

Proof Observe that

$$\pi_1(L(\mathbf{x})) \cdot \pi_1(R(\mathbf{x})) = \frac{1}{Z_1^2} \sum_{(\mathbf{y}, \mathbf{z}) \in L(\mathbf{x}) \times R(\mathbf{x})} e^{-\beta(\Phi(\mathbf{y}) + \Phi(\mathbf{z}))}.$$

For each pair $(\mathbf{y}, \mathbf{z}) \in L(\mathbf{x}) \times R(\mathbf{x})$, consider the profile $\mathbf{w}_{\mathbf{y}, \mathbf{z}} = (\mathbf{z}_L, \mathbf{y}_R)$. Then, $(\mathbf{y}, \mathbf{z}) = \mu(\mathbf{x}, \mathbf{w}_{\mathbf{y}, \mathbf{z}})$, where μ is the decomposition (14). Note that the correspondence between pairs $(\mathbf{y}, \mathbf{z}) \in L(\mathbf{x}) \times R(\mathbf{x})$ and profiles \mathbf{w} is actually a bijection. Indeed, for each profile $\mathbf{w} \in S$, the pair $(\mu_1(\mathbf{x}, \mathbf{w}), \mu_2(\mathbf{x}, \mathbf{w}))$ belongs to $L(\mathbf{x}) \times R(\mathbf{x})$. Hence and from Theorem 5.1, it follows that

$$\begin{aligned} \pi_1(L(\mathbf{x})) \cdot \pi_1(R(\mathbf{x})) &= \frac{1}{Z_A} \sum_{\mathbf{w}} e^{-\beta(\Phi(\mu_1(\mathbf{x}, \mathbf{w})) + \Phi(\mu_2(\mathbf{x}, \mathbf{w})))} \\ &= \frac{1}{Z_A} \sum_{\mathbf{w}} e^{-\beta K(\mathbf{x}, \mathbf{w})} = \pi_A(\mathbf{x}). \end{aligned}$$

□

We now define the concept of a decomposable observable.

Definition 5.2 An observable O is *decomposable* for local interaction potential game \mathcal{G} if there exists a decomposition μ of \mathcal{G} such that, for all (\mathbf{x}, \mathbf{y}) , we have that

$$O(\mathbf{x}) + O(\mathbf{y}) = O(\mu_1(\mathbf{x}, \mathbf{y})) + O(\mu_2(\mathbf{x}, \mathbf{y})).$$

We next prove that a decomposable observable has the same expectation at stationarity of the one-logit and the all-logit.

Theorem 5.6 *If observable O is decomposable then*

$$\langle O, \pi_1 \rangle = \langle O, \pi_A \rangle.$$

Proof Suppose that O is decomposed by μ . Then we have that, for all $\mathbf{x} \in S, \pi_A(\mathbf{x}) = \sum_{\mathbf{y}} \pi_1(\mu_1(\mathbf{x}, \mathbf{y})) \cdot \pi_1(\mu_2(\mathbf{x}, \mathbf{y}))$ and thus

$$\begin{aligned} \langle O, \pi_A \rangle &= \sum_{\mathbf{x}} O(\mathbf{x}) \cdot \pi_A(\mathbf{x}) \\ &= \sum_{\mathbf{x}, \mathbf{y}} O(\mathbf{x}) \cdot \pi_1(\mu_1(\mathbf{x}, \mathbf{y})) \cdot \pi_1(\mu_2(\mathbf{x}, \mathbf{y})) \\ &= \frac{1}{2} \sum_{\mathbf{x}, \mathbf{y}} [O(\mathbf{x}) + O(\mathbf{y})] \cdot \pi_1(\mu_1(\mathbf{x}, \mathbf{y})) \cdot \pi_1(\mu_2(\mathbf{x}, \mathbf{y})) \end{aligned}$$

In the last equality we have used that $\mu_1(\mathbf{x}, \mathbf{y}) = \mu_2(\mathbf{y}, \mathbf{x})$ and $\mu_2(\mathbf{x}, \mathbf{y}) = \mu_1(\mathbf{y}, \mathbf{x})$ which implies that

$$\sum_{\mathbf{x}, \mathbf{y}} O(\mathbf{x}) \cdot \pi_1(\mu_1(\mathbf{x}, \mathbf{y})) \cdot \pi_1(\mu_2(\mathbf{x}, \mathbf{y})) = \sum_{\mathbf{x}, \mathbf{y}} O(\mathbf{y}) \cdot \pi_1(\mu_1(\mathbf{x}, \mathbf{y})) \cdot \pi_1(\mu_2(\mathbf{x}, \mathbf{y})).$$

Now, since O is decomposable we have that $O(\mathbf{x}) + O(\mathbf{y}) = O(\mu_1(\mathbf{x}, \mathbf{y})) + O(\mu_2(\mathbf{x}, \mathbf{y}))$ and thus we can write

$$\begin{aligned} \langle O, \pi_A \rangle &= \frac{1}{2} \sum_{\mathbf{x}, \mathbf{y}} [O(\mu_1(\mathbf{x}, \mathbf{y})) + O(\mu_2(\mathbf{x}, \mathbf{y}))] \cdot \pi_1(\mu_1(\mathbf{x}, \mathbf{y})) \cdot \pi_1(\mu_2(\mathbf{x}, \mathbf{y})) \\ &= \frac{1}{2} \sum_{\mathbf{x}, \mathbf{y}} [O(\mathbf{x}) + O(\mathbf{y})] \cdot \pi_1(\mathbf{x}) \cdot \pi_1(\mathbf{y}) \\ &= \sum_{\mathbf{x}, \mathbf{y}} O(\mathbf{x}) \cdot \pi_1(\mathbf{x}) \cdot \pi_1(\mathbf{y}) \\ &= \sum_{\mathbf{x}} O(\mathbf{x}) \cdot \pi_1(\mathbf{x}) \cdot \sum_{\mathbf{y}} \pi_1(\mathbf{y}) \\ &= \langle O, \pi_1 \rangle. \end{aligned}$$

□

We now give examples of decomposable observables.

The Diff observable For a local interaction potential game in which players have only two strategies, namely -1 and $+1$, we define the observable **Diff** that returns the (signed) difference between the number of vertices adopting the strategy -1 and the number of vertices adopting strategy $+1$. That is, $\text{Diff}(\mathbf{x}) = \sum_u x_u$. In local interaction potential games used to model the diffusion of new technology in a social network [40], **Diff** is a measure of how widely the new technology has been adopted. The **Diff** observable is also meaningful in the Ising model for ferromagnetism (see, for example, [34]) as it coincides with the measured magnetism.

To prove that **Diff** is decomposable we consider the mapping defined in (14) and observe that, for every vertex u and for every $(\mathbf{x}, \mathbf{y}) \in S \times S$, we have $x_u + y_u = \tilde{x}_u + \tilde{y}_u$. Whence we conclude that $\text{Diff}(\mathbf{x}) + \text{Diff}(\mathbf{y}) = \text{Diff}(\tilde{\mathbf{x}}) + \text{Diff}(\tilde{\mathbf{y}})$.

The MonoC observable Another interesting decomposable observable is the signed difference **MonoC** between the number of “−1”-monochromatic edges of the social graph (that is, edges in which both endpoints play −1) and the number of “+1”-monochromatic edges. A monochromatic edge is a simple clique in which all vertices adopt the same strategy (notice that bipartite graphs do not have cliques larger than two). That is, $\text{MonoC}(\mathbf{x}) = \frac{1}{2} \sum_{(u,v) \in E} (x_u + x_v)$. Again, we consider the mapping defined in (14) and the decomposability of **MonoC** follows from the property that, for every $(\mathbf{x}, \mathbf{y}) \in S \times S$, we have $x_u + y_u = \tilde{x}_u + \tilde{y}_u$.

Corollary 5.7 *Observables Diff and MonoC are decomposable and thus, for local interaction potential games on bipartite social graphs,*

$$\langle \text{Diff}, \pi_1 \rangle = \langle \text{Diff}, \pi_A \rangle \quad \text{and} \quad \langle \text{MonoC}, \pi_1 \rangle = \langle \text{MonoC}, \pi_A \rangle.$$

5.2 General Graphs

Let us start by slightly generalizing concepts of decomposition and decomposable observable.

Definition 5.3 A permutation

$$\mu : (\mathbf{x}, \mathbf{y}) \mapsto (\mu_1(\mathbf{x}, \mathbf{y}), \mu_2(\mathbf{x}, \mathbf{y}))$$

of $S \times S$ is an α -decomposition for a local interaction potential game \mathcal{G} with potential Φ if, for all (\mathbf{x}, \mathbf{y}) , we have that

$$|K(\mathbf{x}, \mathbf{y}) - \Phi(\mu_1(\mathbf{x}, \mathbf{y})) - \Phi(\mu_2(\mathbf{x}, \mathbf{y}))| \leq \alpha,$$

$$\mu_1(\mathbf{x}, \mathbf{y}) = \mu_2(\mathbf{y}, \mathbf{x}) \text{ and } \mu_2(\mathbf{x}, \mathbf{y}) = \mu_1(\mathbf{y}, \mathbf{x}).$$

Note that a decomposition is actually a 0-decomposition (see Definition 5.1).

Definition 5.4 An observable O is α -decomposable if it is decomposed by an α -decomposition.

We prove that for all local interaction potential games there exists an α -decomposition with α depending only on how far away the social graph G is from being bipartite. Specifically, for each edge e of the social graph we define the *weight* $w_e = \max_{\mathbf{x}, \mathbf{y} \in \mathcal{G}_e} (\Phi_e(\mathbf{x}) - \Phi_e(\mathbf{y}))$, i.e., w_e is the maximum difference in the potential Φ_e of the two-player game \mathcal{G}_e played on edge e . We say that a subset of edges of G is *bipartiting* if its removal makes the graph bipartite. We denote with $B(G)$ the bipartiting subset of minimum weight and with $b(G)$ its weight. Note that quantity $b(G)$ is related to the *bipartiteness ratio* of G which in turn is related to the smallest eigenvalue of G [43].

We have the following theorem.

Theorem 5.8 *Let \mathcal{G} be a social interaction game on a graph G . Then \mathcal{G} admits an α -decomposition for any $\alpha \geq 2 \cdot b(G)$.*

Proof Let us name as $G' = (L, R, E')$ the bipartite graph obtained by deleting from G the edges of $B(G)$ and consider the mapping (14). We know this mapping is actually a permutation and $\mu_1(\mathbf{x}, \mathbf{y}) = \mu_2(\mathbf{y}, \mathbf{x})$ and $\mu_2(\mathbf{x}, \mathbf{y}) = \mu_1(\mathbf{y}, \mathbf{x})$. We will show that, for every \mathbf{x}, \mathbf{y}

$$|K(\mathbf{x}, \mathbf{y}) - \Phi(\tilde{\mathbf{x}}) - \Phi(\tilde{\mathbf{y}})| \leq 2 \cdot b(G), \tag{15}$$

where $\tilde{\mathbf{x}} = \mu_1(\mathbf{x}, \mathbf{y})$ and $\tilde{\mathbf{y}} = \mu_2(\mathbf{x}, \mathbf{y})$.

Observe that $K(\mathbf{x}, \mathbf{y}) = \sum_{e \in E'} K_e(\mathbf{x}, \mathbf{y}) + \sum_{e \in E \setminus E'} K_e(\mathbf{x}, \mathbf{y})$. From Theorem 5.4, for each edge $e = (u, v) \in E'$ we have $K_e(\mathbf{x}, \mathbf{y}) = \Phi_e(\tilde{x}_u, \tilde{x}_v) + \Phi_e(\tilde{y}_u, \tilde{y}_v)$. As for each edge $e = (u, v) \in E \setminus E'$ we have that the endpoints are either both in L or both in R . In both cases, it turns out that

$$\Phi_e(\tilde{x}_u, \tilde{x}_v) + \Phi_e(\tilde{y}_u, \tilde{y}_v) = \Phi_e(x_u, x_v) + \Phi_e(y_u, y_v).$$

Then we distinguish four cases:

1. $y_u = x_u$ and $y_v = x_v$. In this case $K_e(\mathbf{x}, \mathbf{y}) = 2 \cdot \Phi_e(x_u, x_v)$ and thus $K_e(\mathbf{x}, \mathbf{y}) = \Phi_e(\tilde{x}_u, \tilde{x}_v) + \Phi_e(\tilde{y}_u, \tilde{y}_v)$.
2. $y_u \neq x_u$ and $y_v = x_v$. In this case $K_e(\mathbf{x}, \mathbf{y}) = \Phi_e(x_u, x_v) + \Phi_e(y_u, x_v)$ and thus $K_e(\mathbf{x}, \mathbf{y}) = \Phi_e(\tilde{x}_u, \tilde{x}_v) + \Phi_e(\tilde{y}_u, \tilde{y}_v)$.
3. $y_u = x_u$ and $y_v \neq x_v$. In this case $K_e(\mathbf{x}, \mathbf{y}) = \Phi_e(x_u, x_v) + \Phi_e(x_u, y_v)$ and thus $K_e(\mathbf{x}, \mathbf{y}) = \Phi_e(\tilde{x}_u, \tilde{x}_v) + \Phi_e(\tilde{y}_u, \tilde{y}_v)$.
4. $y_u \neq x_u$ and $y_v \neq x_v$. In this case $K_e(\mathbf{x}, \mathbf{y}) = \Phi_e(y_u, x_v) + \Phi_e(x_u, y_v)$. Since $|\Phi_e(x_u, x_v) - \Phi_e(y_u, x_v)| \leq w_e$ and $|\Phi_e(y_u, y_v) - \Phi_e(x_u, y_v)| \leq w_e$, then

$$|K_e(\mathbf{x}, \mathbf{y}) - \Phi_e(\tilde{x}_u, \tilde{x}_v) - \Phi_e(\tilde{y}_u, \tilde{y}_v)| \leq 2w_e.$$

By summing the contribution of every edge we achieve (15). □

Finally, we next prove that for an α -decomposable observable the extent at which the expectations at stationarity for the one-logit and the all-logit differ depends only on α and β .

Theorem 5.9 *If observable O is α -decomposable then*

$$e^{-2\alpha\beta} \cdot \langle O, \pi_1 \rangle \leq \langle O, \pi_A \rangle \leq e^{2\alpha\beta} \cdot \langle O, \pi_1 \rangle.$$

Proof By mimicking the proof of Theorem 5.1, we have $e^{-\alpha\beta} Z_1^2 \leq Z_A \leq e^{\alpha\beta} Z_1^2$ and

$$e^{-\alpha\beta} \sum_{\mathbf{y}} \gamma_1(\mu_1(\mathbf{x}, \mathbf{y})) \cdot \gamma_1(\mu_2(\mathbf{x}, \mathbf{y})) \leq \gamma_A(\mathbf{x}) \leq e^{\alpha\beta} \sum_{\mathbf{y}} \gamma_1(\mu_1(\mathbf{x}, \mathbf{y})) \cdot \gamma_1(\mu_2(\mathbf{x}, \mathbf{y})).$$

The theorem then follows by the same arguments given in the proof of Theorem 5.6. □

6 Mixing Time

The all-logit dynamics for a strategic game have the property that, for every pair of profiles \mathbf{x}, \mathbf{y} and for every value of β , the transition probability from \mathbf{x} to \mathbf{y} is strictly positive. In order to give upper bounds on the mixing time, we will use the following simple well-known lemma (see e.g. Theorem 11.5 in [36]).

Lemma 6.1 *Let P be the transition matrix of an ergodic Markov chain with state space Ω . For every $r \in \Omega$ let us name $\alpha_r = \min\{P(s, r) : s \in \Omega\}$ and $\alpha = \sum_{r \in \Omega} \alpha_r$. Then the mixing time of P is $t_{\text{mix}} = \mathcal{O}(1/\alpha)$.*

We now give an upper bound holding for *every game*. Recall that for a strategic game \mathcal{G} , in Sect. 2 we defined the cumulative utility function for the ordered pair of profiles (\mathbf{x}, \mathbf{y}) as $U(\mathbf{x}, \mathbf{y}) = \sum_{i=1}^n u_i(\mathbf{x}_{-i}, y_i)$. Let us name ΔU the size of the range of U ,

$$\Delta U = \max\{U(\mathbf{x}, \mathbf{y}) : \mathbf{x}, \mathbf{y} \in S\} - \min\{U(\mathbf{x}, \mathbf{y}) : \mathbf{x}, \mathbf{y} \in S\}.$$

By using Lemma 6.1 we can give a simple upper bound on the mixing time of the all-logit dynamics for \mathcal{G} as a function of β and ΔU .

Theorem 6.2 (General upper bound) *For any strategic game \mathcal{G} the mixing time of the all-logit dynamics for \mathcal{G} is $\mathcal{O}(e^{\beta \Delta U})$.*

Proof Let P be the transition matrix of the all-logit dynamics for \mathcal{G} and let $\mathbf{x}, \mathbf{y} \in S$ be two profiles. From (3) we have that

$$P(\mathbf{x}, \mathbf{y}) = \frac{e^{\beta U(\mathbf{x}, \mathbf{y})}}{\sum_{\mathbf{z} \in S} e^{\beta U(\mathbf{x}, \mathbf{z})}} = \frac{1}{\sum_{\mathbf{z} \in S} e^{\beta(U(\mathbf{x}, \mathbf{z}) - U(\mathbf{x}, \mathbf{y}))}} \geq \frac{1}{|S| e^{\beta \Delta U}}.$$

Hence for every $\mathbf{y} \in S$ it holds that

$$\alpha_{\mathbf{y}} \geq \frac{e^{-\beta \Delta U}}{|S|}$$

and $\alpha = \sum_{\mathbf{y} \in S} \alpha_{\mathbf{y}} \geq e^{-\beta \Delta U}$. The thesis then follows from Lemma 6.1. \square

Next sections will give specific bounds for two specific classes of games (that contain the games analyzed in the Sect. 3), namely graphical coordination games and games with a dominant profile. These results show that the mixing time of the all-logit dynamics has the same twofold behavior that has been highlighted in the case of the one-logit: for some games it depends exponentially on β , whereas for other games it can be upper-bounded by a function independent from β .

6.1 Graphical Coordination Games

A *graphical coordination game* is a local interaction potential game in which all edges play the coordination game described in Eq. (6). By applying Theorem 6.2 we obtain the following upper bound.

Theorem 6.3 *The mixing time of the all-logit for a graphical coordination game on a graph $G = (V, E)$ is*

$$t_{\text{mix}} = \mathcal{O}\left(e^{2\beta(\max\{a,b\}-\min\{c,d\})|E|}\right).$$

Proof Suppose that $a \geq b$. Then, consider the profile \mathbf{x}_+ in which each player plays the strategy +1. It is easy to see that $U(\mathbf{x}, \mathbf{y}) \leq U(\mathbf{x}_+, \mathbf{x}_+) = \sum_i a \cdot \text{deg}(i)$, where $\text{deg}(i)$ is the degree of i in G . The case $a < b$ is equivalent except that we now consider the profile \mathbf{x}_- in which each player plays the strategy -1. Similarly, suppose that $c \leq d$. Then $U(\mathbf{x}, \mathbf{y}) \geq U(\mathbf{x}_-, \mathbf{x}_-) = \sum_i c \cdot \text{deg}(i)$. The case $d < c$ is equivalent except we invert the role of \mathbf{x}_- and \mathbf{x}_+ . Hence

$$\begin{aligned} \Delta U &= \sum_i \max\{a, b\} \cdot \text{deg}(i) - \sum_i \min\{c, d\} \cdot \text{deg}(i) \\ &= 2\beta(\max\{a, b\} - \min\{c, d\})|E|. \end{aligned}$$

The thesis then follows from Theorem 6.2. □

This bound shows that the mixing time of the all-logit for graphical coordination games exponentially depends on β , as in the case of the one-logit dynamics. However, the bound given in the previous theorem can be very loose with respect to the known results about the mixing time of the one-logit for graphical coordination games [6]. It would be interesting to understand at which extent the above bounds can be improved (in ‘‘Appendix’’ we slightly improve these bounds for a very special graphical coordination game, namely the Curie–Weiss model for ferromagnetism adopted in Statistical Physics) and, in particular, if it is possible to show that the mixing time of the all-logit is upper bounded by the mixing time of the one-logit.

6.2 Games with Dominant Strategies

Theorems 6.3 shows that for graphical coordination games the mixing time grows with β . In this section we show that for games with a dominant profile, such as the prisoner’s dilemma analyzed in Sect. 3, the time that the all-logit take for converging to the stationary distribution is upper bounded by a function independent of β , as in the case of the one-logit dynamics [6].

Specifically, we say that strategy $s^* \in S_i$ is a *dominant strategy* for player i if for all $s' \in S_i$ and all strategy profiles $\mathbf{x} \in S$,

$$u_i(s^*, \mathbf{x}_{-i}) \geq u_i(s', \mathbf{x}_{-i}).$$

A *dominant profile* $\mathbf{x}^* = (x_1^*, \dots, x_n^*)$ is a profile in which x_i^* is a dominant strategy for player $i = 1, \dots, n$. Then, we can derive the following upper bound on the mixing time of the all-logit dynamics for games with a dominant profile, whose proof resembles the one used for proving a similar result for the one-logit given in [6].

Theorem 6.4 *Let \mathcal{G} be an n -player games with a dominant profile where each player has at most m strategies. The mixing time of the all-logit for \mathcal{G} is*

$$t_{\text{mix}} = \mathcal{O}(m^n).$$

Proof The proof uses the coupling technique (see, for example, Theorem 5.2 in [32]).

Let P be the transition matrix of the all-logit dynamics for \mathcal{G} . For every pair of profiles \mathbf{x} and \mathbf{y} , we consider a coupling (X, Y) of the distributions $P(\mathbf{x}, \cdot)$ and $P(\mathbf{y}, \cdot)$ such that for every player i the probability that both chains choose strategy s for player i is exactly $\min\{\sigma_i(s | \mathbf{x}), \sigma_i(s | \mathbf{y})\}$. Observe that, with such a coupling, once the two chains coalesce, i.e. $X = Y$, they stay together.

We next observe that for all starting profiles \mathbf{x} and \mathbf{y} , it holds that

$$\mathbf{P}_{\mathbf{x},\mathbf{y}}(X_1 = Y_1) \geq \mathbf{P}_{\mathbf{x},\mathbf{y}}(X_1 = \mathbf{x}^* \text{ and } Y_1 = \mathbf{x}^*) \geq \frac{1}{m^n}.$$

Indeed both chains are in profile \mathbf{x}^* after one step if and only if every player chooses strategy x_i^* in both chains. From the properties of the coupling, it follows that this event occurs with probability

$$\prod_i \min\{\sigma_i(x_i^* | \mathbf{x}), \sigma_i(x_i^* | \mathbf{y})\} \geq \prod_i \frac{1}{|S_i|} \geq \frac{1}{m^n},$$

where the first inequality follows from (1) and the fact that x_i^* is a dominant strategy for i .

Therefore we have that the probability that the two chains have not yet coupled after k time steps is

$$\mathbf{P}_{\mathbf{x},\mathbf{y}}(X_k \neq Y_k) \leq \left(1 - \frac{1}{m^n}\right)^k \leq e^{-k/m^n},$$

which is less than $1/4$ for $k = \mathcal{O}(m^n)$. By applying the Coupling Theorem [32, Theorem 5.2] we have that $t_{\text{mix}} = \mathcal{O}(m^n)$. □

7 Conclusions and Open Problems

In this paper, we considered the selection rule that assigns positive probability only to the set of all players. A natural extension of our work would follow the lead of [1,2] and consider general selection rules. What is the impact of these selection rules on reversibility and on observables? Notice that if we consider the selection rule that

selects exactly one player at each round and player i with probability $p_i > 0$ (the one-logit set $p_i = 1/n$ for all i) then the stationary distribution is the same as the stationary distribution of the one-logit. Therefore, all observables have the same expected value and all potential games are reversible.

It is a classical result that the stationary distribution of the one-logit (the micro-canonical ensemble, in Statistical Mechanics parlance) is the distribution that maximizes the entropy among all the distributions with a fixed average potential. Can we say something similar for the stationary distribution of the all-logit? A promising direction along this line of research is suggested by results in Sect. 5: at least in some cases the stationary distribution of the all-logit can be seen as a composition of simpler distributions.

Appendix: Mixing Time of the All-Logit for the Curie–Weiss Model

Here we prove upper and lower bounds on the mixing time of the all-logit dynamics for a special graphical coordination game, the CW-game. In such a game we set $a = b = +1$ and $c = d = -1$. Thus, the utility of player $i \in [n]$ is the sum of the number of players playing the same strategy as i , minus the number of players playing the opposite strategy; that is, the utility of player $i \in [n]$ at profile $\mathbf{x} = (x_1, \dots, x_n) \in \{-1, +1\}^n$ is

$$u_i(\mathbf{x}) = x_i \sum_{j \neq i} x_j.$$

It is easy to see that the potential function for this game is

$$\Phi(\mathbf{x}) = - \sum_{\{i,j\} \in \binom{[n]}{2}} x_i x_j.$$

Due to the high level of symmetry of the game, the potential of a profile \mathbf{x} depends only on the *number* of players playing ± 1 . Indeed, we can rewrite the potential of \mathbf{x} as

$$\Phi(\mathbf{x}) = - \frac{\text{Diff}^2(\mathbf{x}) - n}{2},$$

where Diff is the observable described in Sect. 5.1.

The upper bound Observe that, for the Curie–Weiss model we have $\Delta U = 2n(n - 1)$, hence by using Theorem 6.3 we get directly that

$$t_{\text{mix}} = \mathcal{O}\left(e^{2\beta n(n-1)}\right). \tag{16}$$

Hence it follows that mixing time is $\mathcal{O}(1)$ for $\beta = \mathcal{O}(1/n^2)$ and it is $\mathcal{O}(\text{poly}(n))$ for $\beta = \mathcal{O}(\log n/n^2)$.

In what follows we show that factor “2” at the exponent in (16) can be removed and that a slightly better upper bound can be given for $\beta > \log n/n$.

Lemma 7.1 *For every $\mathbf{x}, \mathbf{y} \in \Omega$ it holds that*

$$P(\mathbf{x}, \mathbf{y}) \geq q^{(n+|\text{Diff}(\mathbf{y})|)/2} (1 - q)^{(n-|\text{Diff}(\mathbf{y})|)/2}$$

where

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

Proof Consider a profile $\mathbf{y} \in \{-1, +1\}^n$. Observe that the number of players playing +1 and -1 in \mathbf{y} can be written as $\frac{n+\text{Diff}(\mathbf{y})}{2}$ and $\frac{n-\text{Diff}(\mathbf{y})}{2}$, respectively. If $\text{Diff}(\mathbf{y}) > 0$, i.e. if the number of players playing +1 is larger than the number of players playing -1, then the profile that minimizes $P(\mathbf{x}, \mathbf{y})$ is profile $\mathbf{x}_- = (-1, \dots, -1)$ where every player plays -1. If we name

$$q = \frac{e^{-\beta(n-1)}}{e^{-\beta(n-1)} + e^{\beta(n-1)}} = \frac{1}{1 + e^{2\beta(n-1)}}$$

the probability that a player in \mathbf{x}_- chooses strategy +1 for the next round, we have that

$$P(\mathbf{x}_-, \mathbf{y}) = q^{\frac{n+\text{Diff}(\mathbf{y})}{2}} (1 - q)^{\frac{n-\text{Diff}(\mathbf{y})}{2}}.$$

On the other hand, if $\text{Diff}(\mathbf{y}) < 0$, then $P(\mathbf{x}, \mathbf{y})$ is minimized when $\mathbf{x} = \mathbf{x}_+ = (+1, \dots, +1)$ and, since q is also the probability that a player in \mathbf{x}_+ chooses strategy -1 for the next round, we have that

$$P(\mathbf{x}_+, \mathbf{y}) = q^{\frac{n-\text{Diff}(\mathbf{y})}{2}} (1 - q)^{\frac{n+\text{Diff}(\mathbf{y})}{2}}$$

and the thesis follows. □

Now we can give an upper bound on the mixing time by using Lemmata 6.1 and 7.1.

Theorem 7.2 (Upper bound) *The mixing time of the all-logit dynamics for the Curie–Weiss model is*

$$t_{\text{mix}} = \mathcal{O}\left(ne^{\beta n^2}\right).$$

If $\beta \geq \log n/n$ the mixing time is

$$t_{\text{mix}} = \mathcal{O}\left(\frac{ne^{\beta n^2}}{2^n}\right).$$

Proof From Lemma 7.1 it follows that for every $\mathbf{y} \in \{-1, +1\}^n$ we have

$$\alpha_{\mathbf{y}} = \min\{P(\mathbf{x}, \mathbf{y}) \mid \mathbf{x} \in \{-1, +1\}^n\} \geq q^{(n+|\text{Diff}(\mathbf{y})|)/2}(1 - q)^{(n-|\text{Diff}(\mathbf{y})|)/2}.$$

Hence

$$\alpha = \sum_{\mathbf{y} \in \{-1, +1\}^n} \alpha_{\mathbf{y}} \geq \sum_{\mathbf{y} \in \{-1, +1\}^n} q^{(n+|\text{Diff}(\mathbf{y})|)/2}(1 - q)^{(n-|\text{Diff}(\mathbf{y})|)/2}. \tag{17}$$

Now observe that there are $\binom{n}{n/2}$ profiles \mathbf{y} such that $\text{Diff}(\mathbf{y}) = k$, and since $q \leq 1/2$, the largest terms in (17) are the ones such that $\text{Diff}(\mathbf{y})$ is as close to zero as possible. In order to give a lower bound to α we will thus consider only profiles \mathbf{y} such that $\text{Diff}(\mathbf{y}) = 0$, when n is even, and profiles \mathbf{y} such that $\text{Diff}(\mathbf{y}) = \pm 1$, when n is odd.

Case n even: If we consider only profiles \mathbf{y} such that $\text{Diff}(\mathbf{y}) = 0$ in (17) we have that

$$\alpha \geq \binom{n}{n/2} [q(1 - q)]^{n/2}.$$

By using a standard lower bound for the binomial coefficient (see e.g. Lemma 9.2 in [36]) we have that

$$\binom{n}{n/2} \geq \frac{2^n}{n + 1}.$$

As for $[q(1 - q)]^{n/2}$ we have that

$$\begin{aligned} q(1 - q) &= \frac{1}{1 + e^{2\beta(n-1)}} \cdot \frac{1}{1 + e^{-2\beta(n-1)}} \\ &= \frac{1}{e^{2\beta(n-1)} + 2 + e^{-2\beta(n-1)}} \\ &= \frac{1}{e^{2\beta(n-1)} (1 + 2e^{-2\beta(n-1)} + e^{-4\beta(n-1)})}. \end{aligned} \tag{18}$$

Now observe that for every $\beta \geq 0$ we can bound $1 + 2e^{-2\beta(n-1)} + e^{-4\beta(n-1)} \leq 4$. Thus we have that

$$[q(1 - q)]^{n/2} \geq \frac{1}{2^n e^{\beta n(n-1)}}. \tag{19}$$

Hence

$$\alpha \geq \binom{n}{n/2} [q(1 - q)]^{n/2} \geq \frac{1}{(n + 1)e^{\beta n(n-1)}}.$$

And by using Lemma 6.1 we have

$$t_{\text{mix}} = \mathcal{O} \left(n e^{\beta n(n-1)} \right).$$

If β is large enough, say $\beta \geq \log n/n$, in (18) we can bound

$$1 + 2e^{-2\beta(n-1)} + e^{-4\beta(n-1)} \leq 1 + \frac{1}{n}.$$

Thus, in this case we have that

$$[q(1 - q)]^{n/2} \geq \frac{1}{e^{\beta n(n-1)} (1 + 1/n)^{(n/2)}} \geq \frac{1}{e^{\beta n(n-1)} \cdot \sqrt{e}}. \tag{20}$$

Hence $\alpha \geq \frac{2^n}{(n+1)e^{1/2+\beta n(n-1)}}$ and

$$t_{\text{mix}} = \mathcal{O}\left(\frac{ne^{\beta n(n-1)}}{2^n}\right).$$

Case n odd: If we consider only profiles \mathbf{y} such that $\text{Diff}(\mathbf{y}) = \pm 1$ in (17) we get

$$\alpha \geq 2 \binom{n}{\frac{n+1}{2}} q^{\frac{n+1}{2}} (1 - q)^{\frac{n-1}{2}} = 2 \binom{n}{\frac{n+1}{2}} (q(1 - q))^{n/2} \sqrt{\frac{q}{1 - q}}.$$

Now observe that

$$\sqrt{\frac{q}{1 - q}} = e^{-\beta(n-1)} \quad \text{and} \quad \binom{n}{\frac{n+1}{2}} \geq \frac{1}{2} \cdot \frac{2^n}{n + 1}.$$

By using bounds (19) and (20) for $[q(1 - q)]^{n/2}$ we get $t_{\text{mix}} = \mathcal{O}(ne^{\beta(n^2-1)})$ for every $\beta \geq 0$ and $t_{\text{mix}} = \mathcal{O}\left(\frac{ne^{\beta(n^2-1)}}{2^n}\right)$ for $\beta \geq \log n/n$. □

The lower bound In order to give a lower bound on the mixing time, we first show that, for the Curie–Weiss model, $K(\mathbf{x}, \mathbf{y})$ can be written as a function of $\text{Diff}(\mathbf{x})$, $\text{Diff}(\mathbf{y})$ and of the Hamming distance between the two profiles.

Lemma 7.3 *Let $\mathbf{x}, \mathbf{y} \in \{-1, +1\}^n$ be two profiles with magnetization $\text{Diff}(\mathbf{x})$ and $\text{Diff}(\mathbf{y})$ respectively and let $h_{\mathbf{x}, \mathbf{y}}$ be their Hamming distance, i.e. the number of players where they differ. Then*

$$K(\mathbf{x}, \mathbf{y}) = n - \text{Diff}(\mathbf{x}) \cdot \text{Diff}(\mathbf{y}) - 2h_{\mathbf{x}, \mathbf{y}}.$$

Proof As stated above, $\Phi(\mathbf{x}) = \frac{n - \text{Diff}^2(\mathbf{x})}{2}$. In order to evaluate $K(\mathbf{x}, \mathbf{y}) = \sum_{i=1}^n \Phi(\mathbf{x}_{-i}, y_i) - (n - 2)\Phi(\mathbf{x})$ let us name n_1, n_2 and n_3 as follows

$$\begin{aligned} n_1 &= \#\{i \in [n] : x_i = y_i\}; \\ n_2 &= \#\{i \in [n] : x_i = +1, y_i = -1\}; \\ n_3 &= \#\{i \in [n] : x_i = -1, y_i = +1\}. \end{aligned}$$

In other words, n_1 is the number of players playing the same strategy in profiles \mathbf{x} and \mathbf{y} , n_2 is the number of players playing $+1$ in \mathbf{x} and -1 in \mathbf{y} , and n_3 the number of players playing -1 in \mathbf{x} and $+1$ in \mathbf{y} . It holds that

$$\begin{aligned} \sum_{i=1}^n \Phi(\mathbf{x}_{-i}, y_i) &= n_1 \frac{n - \text{Diff}^2(\mathbf{x})}{2} + n_2 \frac{n - (\text{Diff}(\mathbf{x}) - 2)^2}{2} + n_3 \frac{n - (\text{Diff}(\mathbf{x}) + 2)^2}{2} \\ &= \frac{1}{2} \left((n_1 + n_2 + n_3)(n - \text{Diff}^2(\mathbf{x})) \right. \\ &\quad \left. + 4(n_2 - n_3)\text{Diff}(\mathbf{x}) - 4(n_2 + n_3) \right). \end{aligned} \tag{21}$$

Now observe that $n_1 + n_2 + n_3 = n$, $2(n_2 - n_3) = \text{Diff}(\mathbf{x}) - \text{Diff}(\mathbf{y})$, and $(n_2 + n_3) = h_{\mathbf{x},\mathbf{y}}$. Hence from (21) we get

$$\begin{aligned} \sum_{i=1}^n \Phi(\mathbf{x}_{-i}, y_i) &= \frac{1}{2} \left(n(n + \text{Diff}^2(\mathbf{x})) + 2(\text{Diff}(\mathbf{x}) - \text{Diff}(\mathbf{y}))\text{Diff}(\mathbf{x}) - 4h_{\mathbf{x},\mathbf{y}} \right) \\ &= \frac{n^2}{2} - \frac{n-2}{2}\text{Diff}^2(\mathbf{x}) - \text{Diff}(\mathbf{x})\text{Diff}(\mathbf{y}) - 2h_{\mathbf{x},\mathbf{y}}. \end{aligned} \tag{22}$$

Thus $K(\mathbf{x}, \mathbf{y}) = n - \text{Diff}(\mathbf{x}) \cdot \text{Diff}(\mathbf{y}) - 2h_{\mathbf{x},\mathbf{y}}$. □

Since the Hamming distance between two profiles is at most n , from the above lemma we get the following observation.

Observation 7.4 Let \mathbf{x}, \mathbf{y} be two profiles with $\text{Diff}(\mathbf{x}) \cdot \text{Diff}(\mathbf{y}) \leq 0$, then $K(\mathbf{x}, \mathbf{y}) \geq -n$.

Now we can give a lower bound on the mixing time by using the bottleneck-ratio technique.

Theorem 7.5 (Lower bound) *The mixing time of the all-logit dynamics for the Curie–Weiss model is*

$$t_{\text{mix}} = \Omega \left(\frac{e^{\beta n(n-2)}}{4^n} \right).$$

Proof Let $S_- \subseteq \{-1, +1\}^n$ be the set of profiles \mathbf{x} such that $\text{Diff}(\mathbf{x}) < 0$, i.e.

$$S_- = \{\mathbf{x} \in \{-1, +1\}^n : \text{Diff}(\mathbf{x}) < 0\}$$

and observe that $\pi(S_-) \leq 1/2$. From Observation 7.4 we have that for every $\mathbf{x} \in S_-$ and $\mathbf{y} \in S_+ = \{-1, +1\}^n \setminus S_-$ it holds that

$$\pi(\mathbf{x})P(\mathbf{x}, \mathbf{y}) = \frac{e^{-\beta K(\mathbf{x},\mathbf{y})}}{Z} \leq \frac{e^{\beta n}}{Z}. \tag{23}$$

Moreover, if we name \mathbf{x}_- the profile where everyone is playing -1 we have that

$$\pi(S_-) \geq \pi(\mathbf{x}_-) \geq \frac{1}{Z} e^{-2\beta\Phi(\mathbf{x}_-)} = \frac{1}{Z} e^{\beta n(n-1)}. \quad (24)$$

Hence, by using bounds (23) and (24), and the fact that the size of S_- is at most 2^{n-1} , we can bound the bottleneck at S_- with

$$B(S_-) = \frac{Q(S_-, S_+)}{\pi(S_-)} = \frac{\sum_{\mathbf{x} \in S_-} \sum_{\mathbf{y} \in S_+} \pi(\mathbf{x}) P(\mathbf{x}, \mathbf{y})}{\pi(S_-)} \leq \frac{2^{2n-2} e^{\beta n}}{e^{\beta n(n-1)}} = \frac{2^{2n-2}}{e^{\beta n(n-2)}}.$$

By using the bottleneck-ratio theorem (see e.g. Theorem 7.3 in [32]) it follows that

$$t_{\text{mix}} = \Omega\left(\frac{e^{\beta n(n-2)}}{2^{2n}}\right).$$

□

Remark In this section we proved upper and lower bounds on the mixing time of the all-logit dynamics for the Curie–Weiss model. In particular, the upper bound shows that for $\beta = \mathcal{O}(1/n^2)$ the mixing time is constant and for $\beta = \mathcal{O}(\log n/n^2)$ it is at most polynomial. The lower bound shows that, for every constant $\varepsilon > 0$, if $\beta > (1 + \varepsilon)(\log 4)/n$ the mixing time is exponential. When β is between $\Theta(\log n/n^2)$ and $\Theta(1/n)$ we still cannot say if mixing is polynomial or exponential. This is to be compared with the mixing time of the one-logit dynamics: in this case, the dynamics are known to quickly converge to the stationary distribution for $\beta = \mathcal{O}(\log n/n^2)$ and to take super-polynomial time for $\beta = \omega(\log n/n^2)$ [6]. Hence, the mixing time of the all-logit asymptotically matches the mixing time of the one-logit.

References

1. Alós-Ferrer, C., Netzer, N.: The logit-response dynamics. *Games Econ. Behav.* **68**(2), 413–427 (2010)
2. Alós-Ferrer, C., Netzer, N.: Robust stochastic stability. *ECON—working papers 063*, Department of Economics, University of Zurich (Feb 2012)
3. Anderson, S.P., Goeree, J.K., Holt, C.A.: Minimum-effort coordination games: stochastic potential and logit equilibrium. *Games Econ. Behav.* **34**(2), 177–199 (2001)
4. Anshelevich, E., Dasgupta, A., Tardos, É., Wexler, T.: Near-optimal network design with selfish agents. *Theory Comput.* **4**(1), 77–109 (2008)
5. Asadpour, A., Saberi, A.: On the inefficiency ratio of stable equilibria in congestion games. In: *Proceedings of the 5th International Workshop on Internet and Network Economics (WINE'09)*, volume 5929 of *Lecture Notes in Computer Science*, pp 545–552. Springer (2009)
6. Auletta, V., Ferraioli, D., Pasquale, F., Penna, P., Persiano, G.: Convergence to equilibrium of logit dynamics for strategic games. In: *Proceedings of the 23rd ACM Symposium on Parallelism in Algorithms and Architectures (SPAA'11)*, pp. 197–206. ACM (2011)
7. Auletta, V., Ferraioli, D., Pasquale, F., Persiano, G.: Metastability of logit dynamics for coordination games. In: *Proceedings of the ACM-SIAM Symposium on Discrete Algorithms (SODA'12)*, pp. 1006–1024. SIAM (2012)
8. Auletta, V., Ferraioli, D., Pasquale, F., Persiano, G.: Mixing time and stationary expected social welfare of logit dynamics. *Theory Comput. Syst.* **53**(1), 3–40 (2013)

9. Babichenko, Y., Tamuz, O.: Graphical potential games. arXiv preprint [arXiv:1405.1481](https://arxiv.org/abs/1405.1481) (2014)
10. Bala, V., Goyal, S.: A noncooperative model of network formation. *Econometrica* **68**(5), 1181–1229 (2000)
11. Baron, R., Durieu, J., Haller, H., Solal, P.: A note on control costs and logit rules for strategic games. *J. Evolut. Econ.* **12**(5), 563–575 (2002)
12. Baron, R., Durieu, J., Haller, H., Solal, P.: Control costs and potential functions for spatial games. *Int. J. Game Theory* **31**(4), 541–561 (2003)
13. Berger, N., Kenyon, C., Mossel, E., Peres, Y.: Glauber dynamics on trees and hyperbolic graphs. *Probab. Theory Relat. Fields* **131**, 311–340 (2005). Preliminary version in FOCS 01
14. Bindel, D., Kleinberg, J.M., Oren, S.: How bad is forming your own opinion? In: Proceedings of the 52nd IEEE Annual Symposium on Foundations of Computer Science (FOCS'11), pp. 57–66 (2011)
15. Blume, L.E.: The statistical mechanics of strategic interaction. *Games Econ. Behav.* **5**(3), 387–424 (1993)
16. Borgs, C., Chayes, J.T., Ding, J., Lucier, B.: The hitchhiker's guide to affiliation networks: a game-theoretic approach. In: Proceedings of the 2nd Symposium on Innovation in Computer Science (ICS'11), pp. 389–400. Tsinghua University Press (2011)
17. Borgs, C., Chayes, J.T., Karrer, B., Meeder, B., Ravi, R., Reagans, R., Sayedi, A.: Game-theoretic models of information overload in social networks. In: Proceedings of the 7th Workshop on Algorithms and Models for the Web Graph (WAW'10), pp. 146–161 (2010)
18. Corbo, J., Parkes, D.C.: The price of selfish behavior in bilateral network formation. In: Proceedings of the 24th Annual ACM Symposium on Principles of Distributed Computing (PODC'05), pp. 99–107 (2005)
19. Cournot, A.A.: *Recherches sur le Principes mathematiques de la Theorie des Richesses*. L. Hachette, Paris (1838)
20. Ellison, G.: Learning, local interaction, and coordination. *Econometrica* **61**(5), 1047–1071 (1993)
21. Fabrikant, A., Luthra, A., Maneva, E.N., Papadimitriou, C.H., Shenker, S.: On a network creation game. In: Proceedings of the 22nd Annual ACM Symposium on Principles of Distributed Computing (PODC'03), pp. 347–351 (2003)
22. Ferraioli, D., Goldberg, P., Ventre, C.: Decentralized dynamics for finite opinion games. In: Proceedings of the 5th International Symposium on Algorithmic Game Theory (SAGT'12), pp. 144–155. Springer, Berlin (2012)
23. Fudenberg, D., Levine, D.K.: *The Theory of Learning in Games*. MIT Press, Cambridge (1998)
24. Fudenberg, D., Tirole, J.: *Game Theory*. MIT Press, Cambridge (1992)
25. Goeree, J.K., Holt, C.A.: Stochastic game theory: for playing games, not just for doing theory. *Proc. Natl. Acad. Sci.* **96**(19), 10564–10567 (1999)
26. John, C.: Harsanyi and Reinhard Selten. *A General Theory of Equilibrium Selection in Games*. MIT Press, Cambridge (1988)
27. Hart, S., Mas-Colell, A.: A general class of adaptive procedures. *J. Econ. Theory* **98**(1), 26–54 (2001)
28. Jackson, M.O., Wolinsky, A.: A strategic model of social and economic networks. *J. Econ. Theory* **71**(1), 44–74 (1996)
29. Kelly, F.: *Reversibility and Stochastic Networks*. Cambridge University Press, Cambridge (2011)
30. Kleinberg, J.M., Oren, S.: Mechanisms for (mis)allocating scientific credit. In: Proceedings of the 43rd ACM Symposium on Theory of Computing (STOC'11), pp. 529–538 (2011)
31. Landau, L.D., Lifshitz, E.M.: *Statistical Physics*, vol. 5. Elsevier Science, Burlington (1996)
32. Levin, D., Peres, Y., Wilmer, E.L.: *Markov Chains and Mixing Times*. American Mathematical Society, Providence (2008)
33. Levin, D.A., Luczak, M., Peres, Y.: Glauber dynamics for the mean-field Ising model: cut-off, critical power law, and metastability. *Probab. Theory Relat. Fields* **146**(1–2), 223–265 (2010)
34. Martinelli, F.: Lectures on Glauber dynamics for discrete spin models. In: *Lectures on Probability Theory and Statistics*, volume 1717 of *Lecture Notes in Mathematics*, pp. 93–191. Springer, Berlin Heidelberg (1999)
35. McFadden, D.L.: Conditional logit analysis of qualitative choice behavior. In: *Frontiers in Econometrics*, pp. 105–142. Academic Press, New York (1974)
36. Mitzenmacher, M., Upfal, E.: *Probability and Computing: Randomized Algorithms and Probabilistic Analysis*. Cambridge University Press, Cambridge (2005)
37. Monderer, D., Shapley, L.S.: Potential games. *Games Econ. Behav.* **14**, 124–143 (1996)

38. Montanari, A., Saberi, A.: Convergence to equilibrium in local interaction games. In: Proceedings of 50th Annual IEEE Symposium on Foundations of Computer Science (FOCS'09), pp. 303–312 (2009)
39. Morris, S.: Contagion. *Rev. Econ. Stud.* **67**(1), 57–78 (2000)
40. Peyton, H.P.: The diffusion of innovations in social networks. In: Blume, L.E., Durlauf, S.N. (eds.) *The Economy as a Complex Evolving System*, vol. III. Oxford University Press, Oxford (2003)
41. Rosenthal, R.W.: A class of games possessing pure-strategy nash equilibria. *Int. J. Game Theory* **2**(1), 65–67 (1973)
42. Sandholm, William H.: *Population Games and Evolutionary Dynamics*. MIT Press, Cambridge (2010)
43. Trevisan, L.: Max cut and the smallest eigenvalue. In: Proceedings of the 41st Annual ACM Symposium on Theory of Computing, (STOC'09), pp 263–272. ACM (2009)
44. Weidenholzer, S.: Coordination games and local interactions: a survey of the game theoretic literature. *Games* **1**(4), 551–585 (2010)
45. Wolpert, D.H.: Information theory—the bridge connecting bounded rational game theory and statistical physics. In: Braha, D., Minai, A.A., Bar-Yam, Y. (eds.) *Complex Engineered Systems. Understanding Complex Systems*, vol. 14, pp. 262–290. Springer Berlin Heidelberg (2006)