Logit Dynamics with Concurrent Updates for Local Interaction Games*

Vincenzo Auletta[1], Diodato Ferraioli[2], Francesco Pasquale[3], Paolo Penna[4], and Giuseppe Persiano[1]

[1] Università di Salerno, Italy [auletta,giuper]@dia.unisa.it
[2] Université Paris-Dauphine, France diodato.ferraioli@dauphine.fr
[3] "Sapienza" Università di Roma, Italy pasquale@di.unroma1.it
[4] ETH Zurich, Switzerland paolo.penna@inf.ethz.ch

Abstract. *Logit dynamics* are a family of randomized best response dynamics based on the *logit choice function* [21] that is used to model players with limited rationality and knowledge. In this paper we study the *all-logit dynamics*, where at each time step *all* players concurrently update their strategies according to the logit choice function. In the well studied *one-logit dynamics* [7] 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 games*, a class of games that has been used to model complex social phenomena [7, 23, 26] and physical systems [19]. In a local interaction 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 (weighted) sum of the payoffs from each of the games.

We prove that local interaction games characterize the class of games for which the all-logit dynamics are 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.

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

* This work was partially supported by PRIN 2008 research project COGENT (COmputational and GamE-theoretic aspects of uncoordinated NeTworks), funded by the Italian Ministry of University and Research. Diodato Ferraioli is supported by ANR, project COCA, ANR-09-JCJC-0066.

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 neighbors. Game theoretic approaches have been often proposed for modeling phenomena in a complex social network, such as the formation of the social network itself [16, 5, 3, 8], the formation of opinions [6, 11] and the spread of innovation [26, 23] in the social network. Many of these models are based on *local interaction games*, 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 [13], the logit dynamics [7], fictitious play [12] and no-regret dynamics [15]. Any such dynamics can be seen as made of two components: (i) *Selection rule:* by which the set of players that update their state (strategy) is determined; (ii) *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 [9] 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 [12] and the no-regret dynamics [15].

In this paper, we study a specific class of randomized update rules called the *logit choice function* [21, 7] 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* [7] 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.

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?

For example, it is easy to construct games for which the best response converges to a Nash equilibrium when only one player is selected at each step and does not converge to any state when more players are chosen to concurrently update their strategies.

In this paper we study how the logit choice function behaves in an extremal case of concurrency. Specifically, we couple this update rule 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.

Our Contribution. We study the all-logit dynamics for local interaction games [10, 23]. 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 [10] used to model the spread of innovation in social networks [26], and the Ising model [20] for magnetism. In particular, we study the all-logit dynamics for local interaction games at every possible value of the inverse noise β and we are interested in properties of the original one-logit dynamics that are preserved by the all-logit.

We first consider *reversibility*, 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 (specifically, the Markov chains resulting from the all-logit dynamics) are reversible and it turns out that this class coincides with the class of local interaction games. This is to be compared with the well-known result saying that the one-logit dynamics are reversible for every potential game [7]. We find remarkable that a non-trivial property, as reversibility is, of Markov chains modeling the one-logit for potential games holds even for Markov chains modeling all-logit for a large and widely-used subclass of potential games.

Then, we focus on the *observables* of local interaction games. An observable is a function of the strategy profile (that is, the set 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 dynamics. A prominent example of observable is the difference Diff between the number of players adopting two given strategies in a game. In a local interaction 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 whereas in the Ising model it corresponds to the magnetic field of a magnet.

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 game is bipartite. Note that in many of these cases the stationary distributions of one- and all-logit dynamics are completely different. We highlight that the class of observables for which our result holds includes the Diff observable. 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). This implies that, for the Ising model, the all-logit are dynamics that are compatible with the observations and it are arguably more natural than the one-logit dynamics (that postulate that at any given time step only one particle updates its status and then the updated strategy is instantaneously propagated). 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.

In the full version of the paper [4] we also give preliminary bounds on the convergence time of the all-logit dynamics to their stationary distribution.

Related Works. There is a substantial body of work on the logit dynamics: the interested reader can refer to [24] and references therein.

Specifically, the all-logit dynamics for strategic games have been studied in [1], where the authors consider the logit-choice function combined with general selection rules (including the selection rule of the all-logit) and investigate conditions for which a state is *stochastically stable*. A stochastically stable state is a state that has non-zero probability as β goes to infinity. We focus instead on a specific selection rule that is used by several remarkable dynamics (Cournot, fictitious play, and no-regret) and consider the whole range of values of β.

The one-logit dynamics have been actively studied starting from the work of Blume [7] that showed that for 2×2 coordination games, the risk dominant equilibria (see [14]) are stochastically stable. The one-logit for local interaction games have been analyzed in several papers with the aim of modeling and understanding the spread of innovations in a social network, see e.g. [10, 26].

Remark. For readability sake, in Sections 3 and 4 most of the lemmas and theorems have "proof ideas" instead of full proofs. For full proofs and more detailed descriptions we refer the reader to the full version of the paper [4].

2 Definitions

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

Strategic Games. Let $\mathcal{G} = ([n], S_1, \ldots, S_n, u_1, \ldots, u_n)$ be a finite normal-form strategic game. The set $[n] = \{1, \ldots, n\}$ is the player set, S_i is the set of *strategies* for player $i \in [n]$, $S = S_1 \times S_2 \times \cdots \times S_n$ is the set of *strategy profiles* and $u_i \colon S \to \mathbb{R}$ is the *utility* function of player $i \in [n]$. We adopt the standard game-theoretic notation and for $\mathbf{x} = (x_1, \ldots, x_n) \in S$ and $s \in S_i$, we denote by $(\mathbf{x}_{-i}, s)$ the strategy profile $(x_1, \ldots, x_{i-1}, s, x_{i+1}, \ldots, x_n) \in S$.

Potential games [22] are an important class of games. We say that function $\Phi \colon S \to \mathbb{R}$ is an *exact potential* (or simply a *potential*) for game $\mathcal{G}$ if for every $i \in [n]$ and every $\mathbf{x} \in S$ it holds that $u_i(\mathbf{x}_{-i}, s) - u_i(\mathbf{x}_{-i}, z) = \Phi(\mathbf{x}_{-i}, z) - \Phi(\mathbf{x}_{-i}, s)$ for all $s, z \in S_i$. A game $\mathcal{G}$ that admits a potential is called a *potential game*.

Local Interaction Games. In a *local interaction 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-player 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 game $\mathcal{G}$ is $u_i(\mathbf{x}) = \sum_{e=(i,j)} u_i^e(x_i, x_j)$. It is easy to check that the function $\Phi = \sum_e \Phi_e$ is a potential function for the local interaction game $\mathcal{G}$.

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* [21, 7] described as follows: from profile $\mathbf{x} \in S$ player $i \in [n]$ updates her strategy to $s \in S_i$ with probability $\sigma_i(s \mid \mathbf{x}) = \frac{e^{\beta u_i(\mathbf{x}_{-i}, s)}}{\sum_{z \in S_i} e^{\beta u_i(\mathbf{x}_{-i}, z)}}$. In other words, the logit choice function leans towards strategies promising higher utility. The parameter $\beta \geqslant 0$ 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, where *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 these 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 $\mathbf{x} = (x_1, \ldots, x_n)$ to profile $\mathbf{y} = (y_1, \ldots, y_n)$ is

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

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 $U(\mathbf{x}, \mathbf{y}) = \sum_i u_i(\mathbf{x}_{-i}, y_i)$ [1]. Indeed, by observing that $\prod_{i=1}^{n} \sum_{z \in S_i} e^{\beta u_i(\mathbf{x}_{-i}, z)} = \sum_{\mathbf{z} \in S} \prod_{i=1}^{n} e^{\beta u_i(\mathbf{x}_{-i}, z_i)}$, we can rewrite (1) as

$$P(\mathbf{x}, \mathbf{y}) = \frac{e^{\beta U(\mathbf{x}, \mathbf{y})}}{D(\mathbf{x})} \,, \tag{2}$$

where $D(\mathbf{x}) = \sum_{\mathbf{z} \in S} e^{\beta U(\mathbf{x}, \mathbf{z})}$. For a potential game $\mathcal{G}$ with potential Φ, we can define the *cumulative potential* of $\mathbf{x}$ with respect to $\mathbf{y}$ as $\Psi(\mathbf{x}, \mathbf{y}) = \sum_i \Phi(\mathbf{x}_{-i}, y_i)$. Simple algebraic manipulations show that, for a potential game, we can rewrite the transition probabilities in (2) as $P(\mathbf{x}, \mathbf{y}) = \frac{e^{-\beta \Psi(\mathbf{x}, \mathbf{y})}}{T(\mathbf{x})}$, where $T(\mathbf{x})$ is a shorthand for $\sum_{\mathbf{z} \in S} e^{-\beta \Psi(\mathbf{x}, \mathbf{z})}$.

It is easy to see that a Markov chain with transition matrix (1) is ergodic. Indeed, for example, ergodicity follows from the fact that all entries of the transition matrix are strictly positive.

Reversibility & Observables. In this work we focus on two 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 S. $\mathcal{M}$ is *reversible* with respect to a distribution π if, for every pair of states $x, y \in S$, the following *detailed balance condition* holds $\pi(x)P(x, y) = \pi(y)P(y, x)$.

An *observable* O is a function $O\colon S \to \mathbb{R}$, i.e. it is a function that assigns a value to each strategy profile of the game.

3 Reversibility and Stationary Distribution

It is easy to see that the one-logit dynamics 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 dynamics. However, we will prove that the class of games for which the all-logit dynamics are reversible is exactly the class of local interaction games.

Reversibility Criteria. As previously stated, a Markov chain $\mathcal{M}$ is reversible if there exists a distribution π such that the detailed balance condition is satisfied. The following 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 $x \in S$ to state $y \in S$ is a sequence of states $\langle x_0, x_1, \dots, x_\ell \rangle$ such that $x_0 = x$ and $x_\ell = y$. The probability $\mathbf{P}(\Gamma)$ of path Γ is defined as $\mathbf{P}(\Gamma) = \prod_{j=1}^{\ell} P(x_{j-1}, x_j)$. The *inverse of path* $\Gamma = \langle x_0, x_1, \dots, x_\ell \rangle$ is the path $\Gamma^{-1} = \langle x_\ell, x_{\ell-1}, \dots, x_0 \rangle$. Finally, a cycle C is simply a path from a state x to itself. We are now ready to state the Kolmogorov reversibility criterion (see, for example, [17]).

Theorem 1. *An irreducible Markov chain $\mathcal{M}$ with state space S 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 useful for proving reversibility conditions for the all-logit dynamics and for stating a closed expression for its stationary distribution.

Lemma 1. *Let $\mathcal{M}$ be an irreducible Markov chain with transition probability P and state space S. $\mathcal{M}$ is reversible if and only if for every pair of states $x, y \in S$, there exists a constant $c_{x,y}$ such that for all paths Γ from x to y, it holds that $\frac{\mathbf{P}(\Gamma)}{\mathbf{P}(\Gamma^{-1})} = c_{x,y}$.*

Proof (idea). One direction follows directly from the Kolmogorov reversibility criterion, since each cycle can be seen as a concatenation of two paths from x to y (actually, a path and the inverse of another path). As for the other direction, fix z and check that the distribution $\tilde{\pi}(x) = c_{z,x}/Z$, where Z is the normalizing constant, satisfies the detailed balance equation. $\square$

All-Logit Reversibility Implies Potential Games. Now we prove that if the all-logit dynamics 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 for $\mathcal{G}$.

Lemma 2. *The all-logit dynamics for game $\mathcal{G}$ are reversible if and only if the following property holds for every $\mathbf{x}, \mathbf{y}, \mathbf{z} \in S$: $U(\mathbf{x}, \mathbf{y}) - U(\mathbf{y}, \mathbf{x}) = \Big(U(\mathbf{x}, \mathbf{z}) + U(\mathbf{z}, \mathbf{y})\Big) - \Big(U(\mathbf{y}, \mathbf{z}) + U(\mathbf{z}, \mathbf{x})\Big)$.*

Proof (idea). One direction follows from Lemma 1. As for the other direction, the hypothesis implies that, for any fixed $\mathbf{z}$, $\tilde{\pi}(\mathbf{x}) = \frac{P(\mathbf{z},\mathbf{x})}{Z \cdot P(\mathbf{x},\mathbf{z})}$ satisfies the detailed balance equation, where Z is the normalizing constant. □

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

Proposition 1. *If the all-logit dynamics for game $\mathcal{G}$ are reversible then $\mathcal{G}$ is a potential game.*

Proof (idea). We show that if the all-logit dynamics are reversible then the utility improvement over any cycle of length 4 is 0. The thesis then follows by a known characterization of potential games (Theorem 2.8 of [22]). □

A Necessary and Sufficient Condition for All-Logit Reversibility. Previously we established that the all-logit dynamics are reversible only for potential games and therefore, from now on, we only consider potential games $\mathcal{G}$ with potential function Φ. Now we present in Proposition 2 a necessary and sufficient condition for reversibility that involves the potential and the cumulative potential. The condition will then be used to prove that local interaction games are exactly the games whose all-logit dynamics are reversible.

Proposition 2. *The all-logit dynamics for a game $\mathcal{G}$ with potential Φ and cumulative potential Ψ are reversible if and only if, for all strategy profiles $\mathbf{x}, \mathbf{y} \in S$,*

$$\Psi(\mathbf{x},\mathbf{y}) - \Psi(\mathbf{y},\mathbf{x}) = (n-2)\left(\Phi(\mathbf{x}) - \Phi(\mathbf{y})\right). \tag{3}$$

Proof (idea). We rewrite Lemma 2 in terms of cumulative potential as $\Psi(\mathbf{x},\mathbf{y}) - \Psi(\mathbf{y},\mathbf{x}) = \Big(\Psi(\mathbf{x},\mathbf{z}) + \Psi(\mathbf{z},\mathbf{y})\Big) - \Big(\Psi(\mathbf{y},\mathbf{z}) + \Psi(\mathbf{z},\mathbf{x})\Big)$. Simple algebraic manipulations shows that (3) implies the above equation. As for the other direction, we proceed by induction on the Hamming distance between $\mathbf{x}$ and $\mathbf{y}$. □

Reversibility and Local Interaction Games. Here we prove that the games for which all-logit dynamics are reversible are exactly the local interaction games.

A potential $\Phi : S_1 \times \cdots \times S_n \to \mathbb{R}$ 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. It is easy to see that any two-player potential satisfies (3).

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

Observe that if two potentials satisfy (3), then such is also their sum. Hence we have the following proposition.

Proposition 3. *The all-logit dynamics for local interaction games are reversible.*

Next we prove that also the reverse implication holds.

Proposition 4. *If an n-player potential Φ satisfies* (3) *then it can be written as the sum of at most $N = \binom{n}{2}$ two-player potentials, $\Phi_1, \ldots, \Phi_N$ and thus it represents a local interaction game.*

Proof (idea). Let $z_i^\star$ denote the first strategy in each player's strategy set and let $\mathbf{z}^\star$ be the strategy profile $(z_1^\star, \ldots, z_n^\star)$. Moreover, we fix an arbitrary ordering $(u_1, v_1), \ldots, (u_N, v_N)$ of the N unordered pairs of players. For a potential Φ we define the sequence $\vartheta_0, \ldots, \vartheta_N$ of potentials as follows: $\vartheta_0 = \Phi$ and, for $i = 1, \ldots, N$, set $\vartheta_i = \vartheta_{i-1} - \Phi_i$ where, for $\mathbf{x} \in S$, $\Phi_i(\mathbf{x})$ is defined as $\Phi_i(\mathbf{x}) = \vartheta_{i-1}(x_{u_i}, x_{v_i}, \mathbf{z}^\star_{-u_i v_i})$. Observe that, for $i = 1, \ldots, N$, Φ_i is a two-player potential and its players are u_i and v_i. Moreover, $\sum_{i=1}^N \vartheta_i = \sum_{i=0}^{N-1} \vartheta_i - \sum_{i=1}^N \Phi_i$. Thus $\Phi - \vartheta_N = \sum_{i=1}^N \Phi_i$. We show that, if Φ satisfies (3), 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 game. $\square$

We can thus conclude that if the all-logit dynamics for a potential game $\mathcal{G}$ are reversible then $\mathcal{G}$ is a local interaction game. By combining this result with Proposition 1 and Proposition 3, we obtain

Theorem 2. *The all-logit dynamics for game $\mathcal{G}$ are reversible if and only if $\mathcal{G}$ is a local interaction game.*

Stationary Distribution of the All-Logit for Local Interaction Games.

Theorem 3 (Stationary Distribution). *Let $\mathcal{G}$ be a local interaction game with potential function Φ. Then the stationary distribution of the all-logit for $\mathcal{G}$ is $\pi(\mathbf{x}) \propto e^{(n-2)\beta\Phi(\mathbf{x})} \cdot T(\mathbf{x})$, where $T(\mathbf{x}) = \sum_{\mathbf{z} \in S} e^{-\beta\Psi(\mathbf{x},\mathbf{z})}$.*

Proof (idea). Fix any profile $\mathbf{y}$. The detailed balance equation and Proposition 2 give $\pi(\mathbf{x}) = e^{(n-2)\beta\Phi(\mathbf{x})} \cdot T(\mathbf{x}) \left(\frac{\pi(\mathbf{y})}{e^{(n-2)\beta\Phi(\mathbf{y})} \cdot T(\mathbf{y})}\right)$, for every $\mathbf{x} \in S$. Since the term in parenthesis does not depend on $\mathbf{x}$ the theorem follows. $\square$

For a local interaction game $\mathcal{G}$ with potential function Φ we write $\pi_1(\mathbf{x})$, the stationary distribution of the one-logit for $\mathcal{G}$, as $\pi_1(\mathbf{x}) = \gamma_1(\mathbf{x})/Z_1$ where $\gamma_1(\mathbf{x}) = e^{-\beta\Phi(\mathbf{x})}$ (also termed *Boltzmann factor*) and $Z_1 = \sum_{\mathbf{x}} \gamma_1(\mathbf{x})$ is the partition function. From Theorem 3, we derive that $\pi_A(\mathbf{x})$, the stationary distribution of the all-logit for $\mathcal{G}$, can be written in similar way, i.e., $\pi_A(\mathbf{x}) = \frac{\gamma_A(\mathbf{x})}{Z_A}$, where $\gamma_A(\mathbf{x}) = \sum_{\mathbf{y} \in S} e^{-\beta[\Psi(\mathbf{x},\mathbf{y}) - (n-2)\Phi(\mathbf{x})]}$ and $Z_A = \sum_{\mathbf{x} \in S} \gamma_A(\mathbf{x})$ is the partition function of the all-logit. Simple algebraic manipulations show that, by setting $K(\mathbf{x}, \mathbf{y}) = 2 \cdot \Phi(\mathbf{x}) + \sum_{i \in [n]} \mathbf{d}_{\mathbf{x},\mathbf{y}}(i) \cdot (\Phi(\mathbf{x}_{-i}, y_i) - \Phi(\mathbf{x}))$ where $\mathbf{d}_{\mathbf{x},\mathbf{y}}$ is the characteristic vector of positions i in which $\mathbf{x}$ and $\mathbf{y}$ differ (i.e., $\mathbf{d}_{\mathbf{x},\mathbf{y}}(i) = 1$ if $x_i \neq y_i$ and 0 otherwise), we can write $\gamma_A(\mathbf{x})$ and Z_A as

$$\gamma_A(\mathbf{x}) = \sum_{\mathbf{y} \in S} e^{-\beta K(\mathbf{x},\mathbf{y})} \quad \text{and} \quad Z_A = \sum_{\mathbf{x},\mathbf{y}} e^{-\beta K(\mathbf{x},\mathbf{y})}. \tag{4}$$

4 Observables of Local Information Games

In this section we study observables of local interaction 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. In Theorem 5, we give a sufficient condition for an observable to be invariant, that is for having the two expected values to coincide. The sufficient condition is related to the existence of a *decomposition* of the set $S \times S$ that decomposes the quantity K appearing in the expression for the stationary distribution of the all-logit for the local interaction game $\mathcal{G}$ (see Eq. 4) into a sum of two potentials. In Theorem 5 we show that if $\mathcal{G}$ admits such a decomposition μ and in addition observable O is also decomposed by μ (see Definition 2) then O has the same expected value at the stationarity of the one-logit and of the all-logit dynamics. We then show that all local interaction games on *bipartite* social graphs admit a decomposition permutation (see Theorem 4) and give an example of invariant observable.

The above finding follows from a relation between the partition functions of the one-logit and of the all-logit dynamics that might be of independent interest. More precisely, in Theorem 4 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 dynamics. 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 [18]). 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 4 can be seen as a further evidence that the all-logit can be decomposed into two independent one-logit dynamics.

Throughout this section we assume, for the sake of ease of presentation, that each player has just two strategies available. Extending our results to any number of strategies is straightforward.

We start by introducing the concept of a *decomposition* and then we define the concept of a *decomposable* observable.

Definition 1. *A permutation* $\mu\colon (\mathbf{x}, \mathbf{y}) \mapsto (\mu_1(\mathbf{x}, \mathbf{y}), \mu_2(\mathbf{x}, \mathbf{y}))$ *of* $S \times S$ *is a* decomposition *for a local interaction game* $\mathcal{G}$ *with potential* Φ *if, for all* $(\mathbf{x}, \mathbf{y})$*, we have that* $K(\mathbf{x}, \mathbf{y}) = \Phi(\mu_1(\mathbf{x}, \mathbf{y})) + \Phi(\mu_2(\mathbf{x}, \mathbf{y}))$*,* $\mu_1(\mathbf{x}, \mathbf{y}) = \mu_2(\mathbf{y}, \mathbf{x})$ *and* $\mu_2(\mathbf{x}, \mathbf{y}) = \mu_1(\mathbf{y}, \mathbf{x})$.

Theorem 4. *If a decomposition* μ *for a local interaction game* $\mathcal{G}$ *exists, then* $Z_A = Z_1^2$.

Proof. From (4) we have $Z_A = \sum_{\mathbf{x},\mathbf{y}} e^{-\beta K(\mathbf{x},\mathbf{y})} = \sum_{\mathbf{x},\mathbf{y}} e^{-\beta[\Phi(\mu_1(\mathbf{x},\mathbf{y}))+\Phi(\mu_2(\mathbf{x},\mathbf{y}))]}$. Since μ is a permutation of $S \times S$, we have $Z_A = \sum_{\mathbf{x},\mathbf{y}} e^{-\beta[\Phi(\mathbf{x})+\Phi(\mathbf{y})]} = Z_1^2$. □

Definition 2. *An observable* O *is* decomposable *if there exists a decomposition* μ *such that, for all* $(\mathbf{x}, \mathbf{y})$*, it holds 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 dynamics.

Theorem 5. *If observable O is decomposable then $\langle O, \pi_1\rangle = \langle O, \pi_A\rangle$.*

Proof (idea). Suppose that O is decomposed by μ. Then we have that, for all $\mathbf{x} \in S$, $\gamma_A(\mathbf{x}) = \sum_{\mathbf{y}} \gamma_1(\mu_1(\mathbf{x},\mathbf{y})) \cdot \gamma_1(\mu_2(\mathbf{x},\mathbf{y}))$ and thus

$$\langle O, \pi_A\rangle = \frac{1}{2}\frac{1}{Z_A}\sum_{\mathbf{x},\mathbf{y}} [O(\mathbf{x}) + O(\mathbf{y})]\, \gamma_1(\mu_1(\mathbf{x},\mathbf{y}))\gamma_1(\mu_2(\mathbf{x},\mathbf{y})),$$

where we used the property that $\mu_1(\mathbf{x},\mathbf{y}) = \mu_2(\mathbf{y},\mathbf{x})$ and $\mu_2(\mathbf{x},\mathbf{y}) = \mu_1(\mathbf{y},\mathbf{x})$. The theorem follows since O is decomposable. □

We next prove that for all local interaction games on a bipartite social graph there exists a decomposition. We start with the following sufficient condition for a permutation to be a decomposition.

Lemma 3. *Let $\mathcal{G}$ be a social interaction game on social graph G with potential Φ and let μ be a permutation of $S \times S$ such that, for all $\mathbf{x},\mathbf{y} \in S$, we have $\mu_1(\mathbf{x},\mathbf{y}) = \mu_2(\mathbf{y},\mathbf{x})$, $\mu_2(\mathbf{x},\mathbf{y}) = \mu_1(\mathbf{y},\mathbf{x})$ and for all edges $e = (u,v)$ of G and for all $\mathbf{x},\mathbf{y} \in S$ either $(\tilde{x}_u, \tilde{x}_v, \tilde{y}_u, \tilde{y}_v) = (x_u, y_v, y_u, x_v)$ or $(\tilde{x}_u, \tilde{x}_v, \tilde{y}_u, \tilde{y}_v) = (y_u, x_v, x_u, y_v)$, where $\tilde{\mathbf{x}} = \mu_1(\mathbf{x},\mathbf{y})$ and $\tilde{\mathbf{y}} = \mu_2(\mathbf{x},\mathbf{y})$. Then μ is a decomposition of $\mathcal{G}$.*

Proof (idea). We prove by simple case analysis that the contribution of each edge $e = (u,v)$ to $K(\mathbf{x},\mathbf{y})$ is $\Phi_e(\tilde{x}_u, \tilde{x}_v) + \Phi_e(\tilde{y}_u, \tilde{y}_v)$. The lemma is then obtained by summing over all edges e. □

Theorem 6. *Let $\mathcal{G}$ be a social interaction game on a bipartite graph G. Then $\mathcal{G}$ admits a decomposition.*

Proof (idea). Let (L,R) be the set of vertices in which G is bipartite. For each $(\mathbf{x},\mathbf{y}) \in S \times S$ we define $\tilde{\mathbf{x}} = \mu_1(\mathbf{x},\mathbf{y})$ and $\tilde{\mathbf{y}} = \mu_2(\mathbf{x},\mathbf{y})$ as follows: for every vertex u of G, (i) if $u \in L$ then we set $\tilde{x}_u = x_u$ and $\tilde{y}_u = y_u$; (ii) if $u \in R$ then we set $\tilde{x}_u = y_u$ and $\tilde{y}_u = x_u$.

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 $\mu_2(\mathbf{x},\mathbf{y}) = \mu_1(\mathbf{y},\mathbf{x})$. From the bipartiteness of G it follows that for each edge one of the conditions of Lemma 3 is satisfied. Then we can conclude that the mapping is a decomposition. □

We now give an example of decomposable observable. Consider the observable Diff that returns the (signed) difference between the number of vertices adopting strategy 0 and the number of vertices adopting strategy 1. That is, $\mathsf{Diff}(\mathbf{x}) = n - 2\sum_u x_u$. In local interaction games used to model the diffusion of innovations in social networks and the spread of new technology (see, for example, [26]), this observable is a measure of how wide is the adoption of the innovation. The Diff observable is also meaningful in the Ising model for ferromagnetism (see, for example, [20]) as it is the measured magnetism.

To prove that Diff is decomposable we consider the mapping used in the proof of Theorem 6 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 $O(\mathbf{x}) + O(\mathbf{y}) = O(\tilde{\mathbf{x}}) + O(\tilde{\mathbf{y}})$.

Decomposable Observables for General Graphs. We can show that for local interaction games $\mathcal{G}$ on general social graphs G the expected values of a decomposable observable O with respect to the stationary distributions of the one-logit and of the all-logit dynamics 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 [25]). Due to lack of space we omit the details of that result and refer the interested reader to the full version of this paper [4].

5 Future Directions

In this paper we considered the selection rule where all players play concurrently. A natural extension of this selection rule assigns a different probability to each subset of players. What is the impact of such a probabilistic selection rule on reversibility and on observables? Some interesting results along that direction have been obtained in [1, 2]. Notice that if we consider the selection rule that selects player i with probability $p_i > 0$ (the one-logit dynamics 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 Gibbs distribution, that is the stationary distribution of the one-logit dynamics (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 the results in Section 4: at least in some cases the stationary distribution of the all-logit dynamics can be seen as a composition of simpler distributions.

References

1. Carlos Alós-Ferrer and Nick Netzer. The logit-response dynamics. *Games and Economic Behavior*, 68(2):413–427, 2010.
2. Carlos Alós-Ferrer and Nick Netzer. Robust stochastic stability. ECON - Working Papers 063, Department of Economics - University of Zurich, Feb 2012.
3. Elliot Anshelevich, Anirban Dasgupta, Éva Tardos, and Tom Wexler. Near-optimal network design with selfish agents. *Theory of Computing*, 4(1):77–109, 2008.
4. Vincenzo Auletta, Diodato Ferraioli, Francesco Pasquale, Paolo Penna, and Giuseppe Persiano. Logit dynamics with concurrent updates for local-interaction games. http://arxiv.org/abs/1207.2908, 2012.
5. Venkatesh Bala and Sanjeev Goyal. A noncooperative model of network formation. *Econometrica*, 68(5):1181–1229, 2000.

6. David Bindel, Jon M. Kleinberg, and Sigal Oren. How bad is forming your own opinion? In *Foundations of Computer Science (FOCS), 2011 IEEE 52nd Annual Symposium on*, pages 57–66, 2011.
7. Lawrence E. Blume. The statistical mechanics of strategic interaction. *Games and Economic Behavior*, 5(3):387–424, 1993.
8. Christian Borgs, Jennifer T. Chayes, Jian Ding, and Brendan Lucier. The hitch-hiker's guide to affiliation networks: A game-theoretic approach. In *ICS*, pages 389–400. Tsinghua University Press, 2011.
9. Antoine A. Cournot. *Recherches sur le Principes mathematiques de la Theorie des Richesses*. L. Hachette, 1838.
10. Glenn Ellison. Learning, local interaction, and coordination. *Econometrica*, 61(5):1047–1071, 1993.
11. Diodato Ferraioli, Paul Goldberg, and Carmine Ventre. Decentralized dynamics for finite opinion games. In *Algorithmic Game Theory (SAGT '12)*, pages 144–155. Springer Berlin Heidelberg, 2012.
12. Drew Fudenberg and David K. Levine. *The Theory of Learning in Games*. MIT Press, 1998.
13. Drew Fudenberg and Jean Tirole. *Game Theory*. MIT Press, 1992.
14. John C. Harsanyi and Reinhard Selten. *A General Theory of Equilibrium Selection in Games*. MIT Press, 1988.
15. Sergiu Hart and Andreu Mas-Colell. A general class of adaptive procedures. *Journal of Economic Theory*, 98(1):26 – 54, 2001.
16. Matthew O. Jackson and Asher Wolinsky. A strategic model of social and economic networks. *Journal of Economic Theory*, 71(1):44–74, 1996.
17. Frank Kelly. *Reversibility and Stochastic Networks*. Cambridge University Press, 2011.
18. Lev D. Landau and Evgenij M. Lifshitz. *Statistical Physics*, volume 5. Elsevier Science, 1996.
19. David A. Levin, Malwina Luczak, and Yuval Peres. Glauber dynamics for the mean-field Ising model: cut-off, critical power law, and metastability. *Probability Theory and Related Fields*, 146(1-2):223–265, 2010.
20. Fabio Martinelli. Lectures on Glauber dynamics for discrete spin models. In *Lectures on Probability Theory and Statistics*, volume 1717 of *Lecture Notes in Mathematics*, pages 93–191. Springer Berlin Heidelberg, 1999.
21. Daniel L. McFadden. Conditional logit analysis of qualitative choice behavior. In *Frontiers in Econometrics*, pages 105–142. Academic Press, 1974.
22. Dov Monderer and Lloyd S. Shapley. Potential games. *Games and Economic Behavior*, 14:124–143, 1996.
23. Andrea Montanari and Amin Saberi. Convergence to equilibrium in local interaction games. In *Foundations of Computer Science, 2009. FOCS '09. 50th Annual IEEE Symposium on*, pages 303–312, 2009.
24. William H. Sandholm. *Population Games and Evolutionary Dynamics*. MIT Press, 2011.
25. Luca Trevisan. Max cut and the smallest eigenvalue. In *Proceedings of the 41st annual ACM symposium on Theory of computing*, STOC '09, pages 263–272. ACM, 2009.
26. Peyton H. Young. The diffusion of innovations in social networks. Technical report, 2002.