Decentralized Dynamics for Finite Opinion Games*

Diodato Ferraioli[1], Paul W. Goldberg[2], and Carmine Ventre[3]

[1] Dipartimento di Informatica, Università di Salerno, Fisciano, Italy
ferraioli@dia.unisa.it
[2] Department of Computer Science, University of Liverpool, Liverpool, UK
P.W.Goldberg@liverpool.ac.uk
[3] School of Computing, Teesside University, Middlesbrough, UK
C.Ventre@tees.ac.uk

Abstract. Game theory studies situations in which strategic players can modify the state of a given system, due to the absence of a central authority. Solution concepts, such as Nash equilibrium, are defined to predict the outcome of such situations. In the spirit of the field, we study the computation of solution concepts by means of decentralized dynamics. These are algorithms in which players move in turns to improve their own utility and the hope is that the system reaches an "equilibrium" quickly.

We study these dynamics for the class of opinion games, recently introduced by [1]. These are games, important in economics and sociology, that model the formation of an opinion in a social network. We study best-response dynamics and show that the convergence to Nash equilibria is polynomial in the number of players. We also study a noisy version of best-response dynamics, called logit dynamics, and prove a host of results about its convergence rate as the noise in the system varies. To get these results, we use a variety of techniques developed to bound the mixing time of Markov chains, including coupling, spectral characterizations and bottleneck ratio.

1 Introduction

Social networks are widespread in physical and digital worlds. The following scenario therefore becomes of interest. Consider a group of individuals, connected in a social network, who are members of a committee, and suppose that each individual has her own opinion on the matter at hand. How can this group of people reach *consensus*? This is a central question in economic theory, especially for processes in which people repeatedly average their own opinions. This line of work, see e.g. [2–5], is based on a model defined by DeGroot [6]. In this

* Work partially supported by EPSRC grant EP/G069239/1 "Efficient Decentralised Approaches in Algorithmic Game Theory" and by PRIN 2008 research project COGENT (COmputational and GamE-theoretic aspects of uncoordinated NeTworks), funded by the Italian Ministry of University and Research.

M. Serna (Ed.): SAGT 2012, LNCS 7615, pp. 144–155, 2012.

model, each person i holds an opinion given by a real number x_i, which might for example represent a position on a political spectrum. There is an undirected graph $G = (V, E)$ representing a social network, and node i is influenced by the opinions of her neighbors in G. In each time step, node i updates her opinion to be an average of her current opinion with the current opinions of her neighbors. A variation of this model of interest to our study is due to Friedkin and Johnsen [7]. In [7] it is additionally assumed that each node i maintains a persistent *internal belief* b_i, which remains constant even as node i updates her overall opinion x_i through averaging. (See Sect. 2 for the formal framework.)

However, as recently observed by Bindel et al. [1], consensus is hard to reach, the case of political opinions being a prominent example. The authors of [1] justify the absence of consensus by interpreting repeated averaging as a decentralized dynamics for selfish players. Consensus is not reached as players will not compromise further when this diminishes their *utility.* Therefore, these dynamics will converge to an equilibrium in which players might disagree; Bindel et al. study the cost of disagreement by bounding the price of anarchy in this setting.

In this paper, we continue the study of [1] and ask the question of how quickly equilibria are reached by decentralized dynamics in opinion games. We focus on the setting in which players have only a finite number of strategies available. This is motivated by the fact that in many cases although players have personal beliefs which may assume a continuum of values, they only have a limited number of strategies available. For example, in political elections, people have only a limited number of parties they can vote for and usually vote for the party which is *closer* to their own opinions. Motivated by several electoral systems around the world, we concentrate in this study on the case in which players only have two strategies available. This setting already encodes a number of interesting technical challenges as outlined below.

1.1 Our Contribution

For the finite version of the opinion games considered in [1], we firstly note that this is a potential game [8, 9] thus implying that these games admit pure Nash equilibria. The set of pure Nash equilibria is then characterized. We also notice the interesting fact that while the games in [1] have a price of anarchy of 9/8, our games have unbounded price of anarchy, thus implying that for finite games disagreeing has far deeper consequences on the social cost. These basic facts turn out to be useful in the study of decentralized dynamics for finite opinion games.

Given that the potential function is polynomial in the number of players, by proving that the potential decreases by a constant at each step of the best-response dynamics, we can prove that this dynamics quickly converges to pure Nash equilibria. This result is proved by "reducing" an opinion game to a version of it in which the internal beliefs can only take certain values. The reduced version is equivalent to the original one, as long as best-response dynamics is concerned. Note that the convergence rate for the version of the game considered in [1] is unknown.

In real life, however, there is some noise in the decision process of players. Arguably, people are not fully rational. On the other hand, even if they were, they might not exactly know what strategy represents the best response to a given strategy profile due to the incapacity to correctly determine their utility functions. To model this, we study *logit dynamics* [10] for opinion games. Logit dynamics features a *rationality level* $\beta \geq 0$ (equivalently, a noise level $1/\beta$) and each player is assumed to play a strategy with a probability which is proportional to the corresponding utility to the player and β. So the higher β is, the less noise there is and the more the dynamics is similar to best-response dynamics. Logit dynamics for potential games defines a Markov chain that has a nice structure. As in [11, 12] we exploit this structure to prove bounds on the convergence rate of logit dynamics to the so-called *logit equilibrium*. The logit equilibrium corresponds to the stationary distribution of the Markov chain. Intuitively, a logit equilibrium is a probability distribution over strategy profiles of the game; the distribution is concentrated around pure Nash equilibrium profiles.[1] It is observed in [12] how this notion enjoys a number of desiderata one would like solution concepts to have.

We prove a host of results on the convergence rate of logit dynamics that give a pretty much complete picture as β varies. We give an upper bound in terms of the cutwidth of the graph modeling the social network. The bound is exponential in β and the cutwidth of the graph, thus yielding an exponential guarantee for some topology of the social network. We complement this result by proving a polynomial upper bound when β takes a small value, namely, for β at most the inverse of the maximum degree of nodes of the graph. We complete the preceding upper bound in terms of the cutwidth with lower bounds. Firstly, we prove that in order to get an (essentially) matching lower bound it is necessary to evaluate the size of a certain subset of strategy profiles. For large enough β relative to this subset then we can prove that the upper bound is tight for any social network (specifically, we roughly need β bigger than $n \log n$ over the cutwidth of the graph). For smaller values of β, we are unable to prove a lower bound which holds for every graph. However, we prove that the lower bound holds in this case at both ends of the spectrum of possible social networks. In details, we look at two cases of graphs encoding social networks: cliques, which model monolithic, highly interconnected societies, and complete bipartite graphs, which model more sparse "antitransitive" societies. For these graphs, we firstly evaluate the cutwidth and then relate the latter to the size of the aforementioned set of states. This allows to prove a lower bound exponential in β and the cutwidth of the graph for (almost) any value of β. As far as we know, no previous result was known about the cutwidth of a complete bipartite graph; this might be of independent interest. The result on cliques is instead obtained by generalizing arguments in [13].

To prove the convergence rate of logit dynamics to logit equilibrium we adopt a variety of techniques developed to bound the mixing time of Markov chains.

[1] It is worth noting that the focus of best-response dynamics and logit dynamics is on two different solution concepts.

To prove the upper bounds we use some spectral properties of the transition matrix of the Markov chain defined by the logit dynamics, and coupling of Markov chains. To prove the lower bounds, we instead relay on the concept of bottleneck ratio and the relation between the latter and mixing time. (The interested reader might refer to [13] for a discussion of these concepts.)

Due to the lack of space some of the proofs are omitted or sketched.

1.2 Related Work

In addition to the papers mentioned above, our paper is related to the work on logit dynamics. This dynamics is introduced by Blume [10] and it is mainly adopted in the analysis of graphical coordination games [14–16], in which players are placed on vertices of a graph embedding social relations and each player wants to coordinate with neighbors: we highlight that an unique game is played on every edge, whereas, for opinion games, we need different games in order to encode beliefs (see below). Asadpour and Saberi [17] adopt the logit dynamics for analyzing a class of congestion games. However, none of these works evaluates the time the logit dynamics takes in order to reach the stationary distribution: this line of research is conducted in [11, 12].

A number of papers study the efficient computation of (approximate) pure Nash equilibria for 2-strategy games, such as, *party affiliation games* [18, 19] and *cut games* [20]. Similarly to these works, we focus on a class of 2-strategy games and study efficient computation of pure Nash equilibria; additionally we also study the convergence rate to logit equilibria.

Another related work is [21] by Dyer and Mohanaraj. They study graphical games, called *pairwise-interaction games*, and prove among other results, quick convergence of best-response dynamics for these games. However, our games do not fall in their class. The difference is that, in their case, there is a unique game being played on the edges of the graph; as noted above, we instead need a different game to encode the internal beliefs of the players.

2 The Game

Let $G = (V, E)$ be an undirected connected graph[2] with $|V| = n$. Every vertex of the graph represents a player. Each player i has an *internal belief* $b_i \in [0, 1]$ and only two strategies or *opinions* are available, namely 0 and 1. Motivated by the model in [1], we define the utility of player i in a strategy profile $\mathbf{x} \in \{0, 1\}^n$ as

$$u_i(\mathbf{x}) = -\left((x_i - b_i)^2 + \sum_{j\colon (i,j)\in E} (x_i - x_j)^2\right) .$$

[2] A number of papers, including [1], assume that the graph is weighted to model neighbors' different levels of influence. Here we focus on the case in which all neighbors exert the same kind of "political" weight.

We call such a game an n-player opinion game on a graph G. Let $D_i(\mathbf{x}) = \{j : (i,j) \in E \wedge x_i \neq x_j\}$ be the set of neighbors of i that have an opinion different from i. Then $u_i(\mathbf{x}) = -(x_i - b_i)^2 - |D_i(\mathbf{x})|$.

Let $D(\mathbf{x}) = \{(u,v) \in E \colon x_u \neq x_v\}$ be the set of *discording edges* in the strategy profile $\mathbf{x}$, that is the set of all edges in G whose endpoints have different opinions. Then it is not hard to check that the function $\Phi(\mathbf{x}) = \sum_i (x_i - b_i)^2 + |D(\mathbf{x})|$ is an exact potential function for the opinion game described above. Interestingly, the potential function looks very similar to (but not the same as) the social cost $\mathsf{SC}(\mathbf{x}) = -\sum_{i=1}^n u_i(\mathbf{x}) = \sum_i (x_i - b_i)^2 + 2|D(\mathbf{x})|$.

Let B_i be the integer closer to the internal belief of the player i: that is, $B_i = 0$ if $b_i \leq 1/2$, $B_i = 1$ if $b_i > 1/2$. Moreover, let $N_i^s(\mathbf{x}) = |\{j \colon (i,j) \in E \text{ and } x_j = s\}|$ be the number of neighbors of i that play strategy s in the strategy profile $\mathbf{x}$.

It is not hard to verify that in Nash equilibria each player i selects B_i if and only if at least half his neighborhood has selected this opinion. The only special cases occur when players have beliefs in $\{0, 1/2, 1\}$: if $b_i = 1/2$ player i will be additionally indifferent when exactly half (assuming that Δ_i is even) of his neighbors are playing the same strategy and the other half are playing the other strategy; if $b_i = 0$ or $b_i = 1$ player i will also be indifferent when Δ_i is odd and only $\lfloor \Delta_i/2 \rfloor$ neighbors are playing B_i. Roughly speaking, in a Nash equilibrium players tend to form large coalitions, by preferring to play what the majority plays to their own beliefs.

It is easy to check that this game has infinite Price of Anarchy. Consider the opinion game on a clique where each player has internal belief 0: the profile where each player has opinion 0 has social cost 0. The profile where each player has opinion 1 is a Nash equilibrium and its social cost is $n > 0$. This is in sharp contrast with the bound 9/8 proved in [1].

3 Best-Response Dynamics

Given two games we say they are *best-response equivalent* if each player has identical best responses to every combination of opponents' strategies. For the opinion games the following observation is straightforward.

Observation 1. Let $\mathcal{G}$ be an opinion game where the player i has belief $b_i \in (0, 1/2)$: then $\mathcal{G}$ is best-response equivalent to the same game where the belief of i is set to $b_i = 1/4$. Similarly, if the player i has opinion $b_i \in (1/2, 1)$ the game is best-response equivalent to the same game where the belief of i is set to $b_i = 3/4$.

The following theorem shows that, for this class of games, the best-response dynamics quickly converges to a Nash equilibrium.

Theorem 2. *The best-response dynamics for an n-player opinion game $\mathcal{G}$ converges to a Nash equilibrium after a polynomial number of steps.*

Proof (Sketch). From Observation 1 we know that each opinion game is best-response equivalent to an opinion game where each player i has $b_i \in S = \left\{0, \frac{1}{4}, \frac{1}{2}, \frac{3}{4}, 1\right\}$. So, for a given opinion game $\mathcal{G}$ we construct a game $\mathcal{G}'$ with beliefs restricted to belong to S by "rounding" the beliefs of the original game and show that best-response dynamics converges quickly on $\mathcal{G}'$. We begin by observing that for every profile $\mathbf{x}$, we have $0 \leq \Phi(\mathbf{x}) \leq n^2 + n$. Thus, the theorem follows by showing that at each time step the cost of a player decreases by at least a constant value. □

4 Logit Dynamics for Opinion Games

Let $\mathcal{G}$ be an opinion game as from the above; moreover, let $S = \{0,1\}^n$ denote the set of all strategy profiles. For two vectors $\mathbf{x}, \mathbf{y} \in S$, we denote with $H(\mathbf{x}, \mathbf{y}) = |\{i\colon x_i \neq y_i\}|$ the Hamming distance between $\mathbf{x}$ and $\mathbf{y}$. The *Hamming graph* of the game $\mathcal{G}$ is defined as $\mathcal{H} = (S, \mathsf{E})$, where two profiles $\mathbf{x} = (x_1, \ldots, x_n), \mathbf{y} = (y_1, \ldots, y_n) \in S$ are adjacent in $\mathcal{H}$ if and only if $H(\mathbf{x}, \mathbf{y}) = 1$.

The *logit dynamics* for $\mathcal{G}$ runs as follows: at every time step (i) Select one player $i \in [n]$ uniformly at random; (ii) Update the strategy of player i according to the *Boltzmann distribution* with parameter β over the set $S_i = \{0,1\}$ of her strategies. That is, a strategy $s_i \in S_i$ will be selected with probability

$$\sigma_i(s_i \mid \mathbf{x}_{-i}) = \frac{1}{Z_i(\mathbf{x}_{-i})} e^{\beta u_i(\mathbf{x}_{-i}, s_i)}, \tag{1}$$

where $\mathbf{x}_{-i} \in \{0,1\}^{n-1}$ is the profile of strategies played at the current time step by players different from i, $Z_i(\mathbf{x}_{-i}) = \sum_{z_i \in S_i} e^{\beta u_i(\mathbf{x}_{-i}, z_i)}$ is the normalizing factor, and $\beta \geq 0$. As mentioned above, from (1), it is easy to see that for $\beta = 0$ player i selects her strategy uniformly at random, for $\beta > 0$ the probability is biased toward strategies promising higher payoffs, and for β that goes to ∞ player i chooses her best response strategy (if more than one best response is available, she chooses one of them uniformly at random).

The above dynamics defines a *Markov chain* $\{X_t\}_{t \in \mathbb{N}}$ with the set of strategy profiles as state space, and where the probability $P(\mathbf{x}, \mathbf{y})$ of a transition from profile $\mathbf{x} = (x_1, \ldots, x_n)$ to profile $\mathbf{y} = (y_1, \ldots, y_n)$ is zero if $H(\mathbf{x}, \mathbf{y}) \geq 2$ and it is $\frac{1}{n}\sigma_i(y_i \mid \mathbf{x}_{-i})$ if the two profiles differ exactly at player i. More formally, we can define the logit dynamics as follows.

Definition 3 (Logit dynamics [10]). *Let $\mathcal{G}$ be an opinion game as from the above and let $\beta \geq 0$. The* logit dynamics *for $\mathcal{G}$ is the Markov chain $\mathcal{M}_\beta = (\{X_t\}_{t \in \mathbb{N}}, S, P)$ where $S = \{0,1\}^n$ and*

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

where $\sigma_i(y_i \mid \mathbf{x}_{-i})$ is defined in (1).

The Markov chain defined by (2) is ergodic. Hence, from every initial profile $\mathbf{x}$ the distribution $P^t(\mathbf{x}, \cdot)$ of chain X_t starting at $\mathbf{x}$ will eventually converge to a *stationary distribution* π as t tends to infinity.[3] As in [12], we call the stationary distribution π of the Markov chain defined by the logit dynamics on a game $\mathcal{G}$, the *logit equilibrium* of $\mathcal{G}$. In general, a Markov chain with transition matrix P and state space S is said to be *reversible* with respect to the distribution π if, for all $\mathbf{x}, \mathbf{y} \in S$, it holds that $\pi(\mathbf{x})P(\mathbf{x}, \mathbf{y}) = \pi(\mathbf{y})P(\mathbf{y}, \mathbf{x})$. If the chain is reversible with respect to π, then π is its stationary distribution. For the class of potential games the stationary distribution is the well-known *Gibbs measure*.

Theorem 4 ([10]). *If $\mathcal{G} = ([n], \mathcal{S}, \mathcal{U})$ is a potential game with potential function Φ, then the Markov chain given by* (2) *is reversible with respect to the Gibbs measure $\pi(\mathbf{x}) = \frac{1}{Z} e^{-\beta\Phi(\mathbf{x})}$, where $Z = \sum_{\mathbf{y} \in S} e^{-\beta\Phi(\mathbf{y})}$ is the normalizing constant.*

Mixing Time of Markov Chains. The most prominent measures of the rate of convergence of a Markov chain to its stationary distribution is the *mixing time*. For a Markov chain with transition matrix P and state space S, let us set $d(t) = \max_{\mathbf{x} \in S} \|P^t(\mathbf{x}, \cdot) - \pi\|_{\text{TV}}$, where the *total variation distance* $\|\mu - \nu\|_{\text{TV}}$ between two probability distributions μ and ν on the same state space S is defined as $\|\mu - \nu\|_{\text{TV}} = \max_{A \subset S} |\mu(A) - \nu(A)|$. For $0 < \varepsilon < 1/2$, the mixing time is defined as $t_{\text{mix}}(\varepsilon) = \min\{t \in \mathbb{N} \colon d(t) \leq \varepsilon\}$. It is usual to set $\varepsilon = 1/4$ or $\varepsilon = 1/2e$. If not explicitly specified, when we write t_{mix} we mean $t_{\text{mix}}(1/4)$. Observe that $t_{\text{mix}}(\varepsilon) \leq \lceil \log_2 \varepsilon^{-1} \rceil t_{\text{mix}}$.

Bottleneck Ratio. An important concept to establish our lower bounds is represented by the *bottleneck ratio*. Consider an ergodic Markov chain with finite state space S, transition matrix P, and stationary distribution π. The probability distribution $Q(\mathbf{x}, \mathbf{y}) = \pi(\mathbf{x})P(\mathbf{x}, \mathbf{y})$ is of particular interest and is sometimes called the *edge stationary distribution*. Note that if the chain is reversible then $Q(\mathbf{x}, \mathbf{y}) = Q(\mathbf{y}, \mathbf{x})$. For any $L \subseteq S$, we let $Q(L, S \setminus L) = \sum_{\mathbf{x} \in L, \mathbf{y} \in S \setminus L} Q(\mathbf{x}, \mathbf{y})$. The bottleneck ratio of $L \subseteq S$, L non-empty, is $B(L) = \frac{Q(L, S \setminus R)}{\pi(L)}$.

The following theorem relates bottleneck ratio and mixing time.

Theorem 5 (Bottleneck ratio [13]). *Let $\mathcal{M} = \{X_t \colon t \in \mathbb{N}\}$ be an irreducible and aperiodic Markov chain with finite state space S, transition matrix P, and stationary distribution π. Then the mixing time is $t_{\text{mix}} \geq \max_{L \colon \pi(L) \leq 1/2} \frac{1}{4 \cdot B(L)}$.*

4.1 Upper Bounds

For Every $\boldsymbol{\beta}$. Consider the bijective function $\sigma \colon V \to \{1, \ldots, |V|\}$: it represents an ordering of vertices of G. Let $\mathcal{L}$ be the set of all orderings of vertices of G and set $V_i^\sigma = \{v \in V \colon \sigma(v) < i\}$. Then, the *cutwidth* of G is $\mathsf{CW}(G) = \min_{\sigma \in \mathcal{L}} \max_{1 < i \leq |V|} |E(V_i^\sigma, V \setminus V_i^\sigma)|$.

[3] The notation $P^t(\mathbf{x}, \cdot)$, standard in Markov chains literature [13], denotes the probability distribution over states of S after the chain has taken t steps starting from $\mathbf{x}$.

Theorem 6. *Let $\mathcal{G}$ be an n-player opinion game on a graph $G = (V, E)$. The mixing time of the logit dynamics for $\mathcal{G}$ is* $t_{\mathrm{mix}} \leq (1+\beta) \cdot \mathsf{poly}\,(n) \cdot e^{\beta\Theta(\mathsf{CW}(G))}$.

The proof is a generalization of a similar proof given by Berger et al. [22] based on spectral arguments.

For Small β. The following theorem shows that for small values of β the mixing time is polynomial. We remark that there are network topologies for which this theorem gives a bound higher than that guaranteed by Theorem 6 on the values of β for which the mixing time is polynomial.

Theorem 7. *Let $\mathcal{G}$ be an n-player opinion game on a connected graph G, with $n > 2$. Let $\Delta_{\max}$ be the maximum degree in the graph. If $\beta \leq 1/\Delta_{\max}$, then the mixing time of the logit dynamics for $\mathcal{G}$ is $\mathcal{O}(n \log n)$.*

Proof (Sketch). Consider two profiles $\mathbf{x}$ and $\mathbf{y}$ that differ only in the strategy played by player j and consider the coupling described in [11] for two chains X and Y starting respectively from $X_0 = \mathbf{x}$ and $Y_0 = \mathbf{y}$. We show the expected distance between X_1 and Y_1 after one step of the coupling is less then $e^{-1/(3n)}$. The bound on the mixing time follows from the well-known path coupling technique [23]. □

4.2 Lower Bounds

Recall that $\mathcal{H}$ is the Hamming graph on the set of profiles of an opinion games on a graph G. The following observation easily follows from the definition of cutwidth.

Observation 8. For every path on $\mathcal{H}$ between the profile $\mathbf{0} = (0, \ldots, 0)$ and the profile $\mathbf{1} = (1, \ldots, 1)$ there exists a profile for which there are at least $\mathsf{CW}(G)$ discording edges.

From now on, let us write CW as a shorthand for $\mathsf{CW}(G)$, when the reference to the graph is clear from the context. For sake of compactness, we set $\mathbf{b}(\mathbf{x}) = \sum_i (x_i - b_i)^2$. We denote as $\mathbf{b}^\star$ the minimum of $\mathbf{b}(\mathbf{x})$ over all profiles with CW discording edges.

Let $R_{\mathbf{0}}$ ($R_{\mathbf{1}}$) be the set of profiles $\mathbf{x}$ for which a path from $\mathbf{0}$ (resp., $\mathbf{1}$) to $\mathbf{x}$ exists on $\mathcal{H}$ such that every profile along the path has potential value less than $\mathbf{b}^\star + \mathsf{CW}$. To establish the lower bound we use the technical result given by Theorem 5 which requires to compute the bottleneck ratio of a subset of profiles that is weighted at most a half by the stationary distribution. Accordingly, we set $R = R_{\mathbf{0}}$ if $\pi(R_{\mathbf{0}}) \leq 1/2$ and $R = R_{\mathbf{1}}$ if $\pi(R_{\mathbf{1}}) \leq 1/2$. (If both sets have stationary distribution less than one half, the best lower bound is achieved by setting R to $R_{\mathbf{0}}$ if and only if $\Phi(\mathbf{0}) \leq \Phi(\mathbf{1})$.) W.l.o.g., in the remaining of this section we assume $R = R_{\mathbf{0}}$.

For Large β. Let ∂R be the set of profiles in R that have at least a neighbor $\mathbf{y}$ in the Hamming graph $\mathcal{H}$ such that $\mathbf{y} \notin R$. Moreover let $\mathcal{E}(\partial R)$ the set of edges $(\mathbf{x}, \mathbf{y})$ in $\mathcal{H}$ such that $\mathbf{x} \in \partial R$ and $\mathbf{y} \notin R$: note that $|\mathcal{E}(\partial R)| \leq n|\partial R|$. The following lemma bounds the bottleneck ratio of R.

Lemma 9. *For the set of profiles R defined above, we have $B(R) \leq n \cdot |\partial R| \cdot e^{-\beta(\mathsf{CW}+\mathbf{b}^\star-\mathbf{b}(\mathbf{0}))}$.*

Proof. Since $\mathbf{0} \in R$, it holds $\pi(R) \geq \pi(\mathbf{0}) = \frac{e^{-\beta\mathbf{b}(\mathbf{0})}}{Z}$. Moreover, by (1) we have

$$\begin{aligned}
Q(R,\overline{R}) &= \sum_{\substack{(\mathbf{x},\mathbf{y})\in\mathcal{E}(\partial R):\\ \mathbf{y}=(\mathbf{x}_{-i},y_i)}} \frac{e^{-\beta\Phi(\mathbf{x})}}{Z} \frac{e^{\beta u_i(\mathbf{y})}}{e^{\beta u_i(\mathbf{x})}+e^{\beta u_i(\mathbf{y})}}\\
&= \sum_{\substack{(\mathbf{x},\mathbf{y})\in\mathcal{E}(\partial R):\\ \mathbf{y}=(\mathbf{x}_{-i},y_i)}} \frac{e^{-\beta\Phi(\mathbf{x})}}{Z} \frac{e^{-\beta\Phi(\mathbf{y})}e^{\beta(u_i(\mathbf{x})+\Phi(\mathbf{x}))}}{e^{-\beta\Phi(\mathbf{x})}e^{\beta(u_i(\mathbf{x})+\Phi(\mathbf{x}))}+e^{-\beta\Phi(\mathbf{y})}e^{\beta(u_i(\mathbf{x})+\Phi(\mathbf{x}))}}\\
&= \frac{1}{Z}\sum_{(\mathbf{x},\mathbf{y})\in\mathcal{E}(\partial R)} \frac{e^{-\beta\Phi(\mathbf{x})}e^{-\beta\Phi(\mathbf{y})}}{e^{-\beta\Phi(\mathbf{x})}+e^{-\beta\Phi(\mathbf{y})}} = \frac{1}{Z}\sum_{(\mathbf{x},\mathbf{y})\in\mathcal{E}(\partial R)} \frac{e^{-\beta\Phi(\mathbf{y})}}{1+e^{\beta(\Phi(\mathbf{x})-\Phi(\mathbf{y}))}}\\
&\leq \frac{1}{Z}\sum_{(\mathbf{x},\mathbf{y})\in\mathcal{E}(\partial R)} e^{-\beta\Phi(\mathbf{y})} \leq |\mathcal{E}(\partial R)| \cdot \frac{e^{-\beta(\mathbf{b}^\star+\mathsf{CW})}}{Z}.
\end{aligned}$$

The second equality follows from the definition of potential function which implies $\Phi(\mathbf{y}) - \Phi(\mathbf{x}) = -u_i(\mathbf{y}) + u_i(\mathbf{x})$ for $\mathbf{x}$ and $\mathbf{y}$ as above; last inequality holds because if by contradiction $\Phi(\mathbf{y}) < \mathbf{b}^\star + \mathsf{CW}$ then, by definition of R, it would be $\mathbf{y} \in R$, a contradiction. □

From Lemma 9 and Theorem 5 we obtain a lower bound to the mixing time of the opinion games that holds for every value of β, every social network G and every vector $(b_1, \ldots, b_n)$ of internal beliefs. However, it is not clear how close this bound is to the one given in Theorem 6. Nevertheless, by taking $b_i = 1/2$ for each player i and β high enough, we can state the following theorem.

Theorem 10. *Let $\mathcal{G}$ be an n-player opinion game on a graph G. Then, there exist a vector of internal beliefs such that for $\beta = \Omega\left(\frac{n \log n}{\mathsf{CW}}\right)$ it holds $t_{\mathrm{mix}} \geq e^{\beta\Theta(\mathsf{CW})}$.*

Proof. If $b_i = 1/2$ for every player i, from Lemma 9 and Theorem 5, since $|\partial R| \leq 2^n$ then $t_{\mathrm{mix}} \geq \frac{e^{\beta\mathsf{CW}}}{n2^n} = e^{\beta\mathsf{CW}-n\log(2n)} = e^{\beta\Theta(\mathsf{CW})}$. □

For Smaller β. Theorem 10 gives an almost tight lower bound for high values of β for each network topology. It would be interesting to prove a matching bound also for lower values of the rationality parameter: in this section we prove such a bound for specific classes of graphs: complete bipartite graphs and cliques.

We start by considering the class of complete bipartite graphs $K_{m,m}$.

Theorem 11. *Let $\mathcal{G}$ be an n-player opinion game on $K_{m,m}$. Then, there exist a vector of internal beliefs such that, for every $\beta = \Omega\left(\frac{1}{m}\right)$, we have $t_{\mathrm{mix}} \geq \frac{e^{\beta\Theta(\mathsf{CW})}}{n}$.*

To prove the theorem above, we start by evaluating the cutwidth of $K_{m,m}$: in particular, we characterize the best ordering from which the cutwidth is obtained. We will denote with A and B the two sides of the bipartite graph. Then it is not hard to see that the ordering that obtains the cutwidth in $K_{m,m}$ is the one that selects alternatively a vertex from A and a vertex from B. Moreover, it turns out that the cutwidth of $K_{m,m}$ is $\lceil m^2/2 \rceil$. The following lemma gives a bound to the size of ∂R for this graph.

Lemma 12. *For the opinion game on the graph $K_{m,m}$ with $b_i = 1/2$ for every player i, there exists a constant c_1 such that $|\partial R| \leq e^{c_1\sqrt{\mathsf{CW}}}$.*

Proof (Sketch). Since $b_i = 1/2$ for every player i, we have that $\mathbf{b}(\mathbf{x}) = n/4$ for every profile $\mathbf{x}$. Therefore, by definition of R, all profiles in R (and therefore ∂R) have less then CW discording edges. Indeed, for $\mathbf{x} \in R$ we have $\mathbf{b}(\mathbf{x}) + |D(\mathbf{x})| = \Phi(\mathbf{x}) < \mathbf{b}^\star + \mathsf{CW}$. Moreover, if a profile $\mathbf{y}$ has less then $\mathsf{CW} - m$ discording edges, then $\mathbf{y}$ is not in ∂R as a state neighbor of $\mathbf{y}$ has at most $m-1$ additional discording edges.

Consequently, to bound the size of ∂R, we need to count the number of profiles in R that have potential between $\mathbf{b}^\star + \mathsf{CW} - m$ and $\mathbf{b}^\star + \mathsf{CW} - 1$ (i.e., the number of profiles with at least $\mathsf{CW} - m$ and at most $\mathsf{CW} - 1$ discording edges). By using the facts about the cut-width of bipartite graphs stated above, we have $|\partial R| \leq (5e)^m \leq e^{3m}$. The lemma follows since $m \leq \sqrt{2}\sqrt{\mathsf{CW}}$. □

Proof (of Theorem 11). If $b_i = 1/2$ for every player i, from Lemmata 9 and 12, we have $B(R) \leq n \cdot e^{c_1\sqrt{\mathsf{CW}}} \cdot e^{-\beta\mathsf{CW}} \leq n \cdot e^{-\beta\mathsf{CW}(1-c_2)}$, where $c_2 = \frac{c_1\sqrt{\mathsf{CW}}}{\beta\mathsf{CW}} < 1$ since by hypothesis $\beta > \frac{c_1}{\sqrt{\mathsf{CW}}} = \Omega(1/m)$; we also notice that c_2 goes to 0 as β increases. The theorem follows from Theorem 5. □

We remark that it is possible to prove a result similar to Theorem 11 also for the clique K_n: the proof follows from a simple generalization of Theorem 15.3 in [13] and by observing that the cutwidth of a clique is $\lfloor n^2/4 \rfloor$.

5 Conclusions and Open Problems

In this work we analyze two decentralized dynamics for binary opinion games: the best-response dynamics and the logit dynamics. For the best-response dynamics we show that it takes time polynomial in the number of players to reach a Nash equilibrium, the latter being characterized by the existence of clusters in which players have a common opinion. On the other hand, for the logit dynamics we show polynomial convergence when the level of noise is high enough and that it increases as β grows.

It is important to highlight, as noted above, that the convergence time of the two dynamics are computed with respect to two different equilibrium concepts, namely Nash equilibrium for the best-response dynamics and logit equilibrium for the logit dynamics. This explains why the convergence times of these two dynamics asymptotically diverge even though the logit dynamics becomes similar to the best response dynamics as β goes to infinity.

Theorem 6 and 10 which prove bounds to the convergence of logit dynamics can also be read in a positive fashion. Indeed, for social networks that have a bounded cutwidth, the convergence rate of the dynamics depends only on the value of β. (We highlight that checking if a graph has bounded cutwidth can be done in polynomial time [24].) In general, we have the following picture: as long as β is less than the maximum of (roughly) $\frac{\log n}{\mathsf{CW}}$ and $\frac{1}{\Delta}$ the convergence time to the logit equilibrium is polynomial. Moreover, Theorem 10 shows that for β lower bounded by (roughly) $\frac{n \log n}{\mathsf{CW}}$ the convergence time to the logit equilibrium is super-polynomial. Then for some network topology, there is a gap in our knowledge which is naturally interesting to close.

In [25] the concept of metastable distributions has been introduced in order to predict the outcome of games for which the logit dynamics takes too much time to reach the stationary distribution for some value of β. It would be interesting to investigate existence and structure of such distributions for our opinion games.

We also note that our proofs for logit dynamics can be extended to the case in which the social graph is weighted. In such a setting, however, we obtain non-matching bounds: it would be interesting to develop more sophisticated techniques in order to get tight bounds.

References

1. Bindel, D., Kleinberg, J., Oren, S.: How bad is forming your own opinion? In: 2011 IEEE 52nd Annual Symposium on Foundations of Computer Science (FOCS), pp. 57–66 (October 2011)
2. Acemoglu, D., Ozdaglar, A.: Opinion dynamics and learning in social networks. Dynamic Games and Applications 1, 3–49 (2011)
3. DeMarzo, P.M., Vayanos, D., Zwiebel, J.: Persuasion bias, social influence, and unidimensional opinions. The Quarterly Journal of Economics 118(3), 909–968 (2003)
4. Golub, B., Jackson, M.O.: Naïve learning in social networks and the wisdom of crowds. American Economic Journal: Microeconomics 2(1), 112–149 (2010)
5. Jackson, M.O.: Social and Economic Networks. Princeton University Press, Princeton (2008)
6. DeGroot, M.H.: Reaching a consensus. Journal of the American Statistical Association 69(345), 118–121 (1974)
7. Friedkin, N.E., Johnsen, E.C.: Social influence and opinions. The Journal of Mathematical Sociology 15(3-4), 193–206 (1990)
8. Rosenthal, R.: A class of games possessing pure-strategy nash equilibria. International Journal of Game Theory 2, 65–67 (1973)
9. Monderer, D., Shapley, L.S.: Potential games. Games and Economic Behavior 14(1), 124–143 (1996)
10. Blume, L.E.: The statistical mechanics of strategic interaction. Games and Economic Behavior 5(3), 387–424 (1993)

11. Auletta, V., Ferraioli, D., Pasquale, F., Persiano, G.: Mixing Time and Stationary Expected Social Welfare of Logit Dynamics. In: Kontogiannis, S., Koutsoupias, E., Spirakis, P.G. (eds.) SAGT 2010. LNCS, vol. 6386, pp. 54–65. Springer, Heidelberg (2010)
12. 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 2011, pp. 197–206. ACM (2011)
13. Levin, D.A., Peres, Y., Wilmer, E.L.: Markov chains and mixing times. American Mathematical Society (2006)
14. Ellison, G.: Learning, local interaction, and coordination. Econometrica 61(5), 1047–1071 (1993)
15. Young, H.P.: The diffusion of innovations in social networks. In: Blume, B.L., Durlauf, S.N. (eds.) Economy as an Evolving Complex System. Proceedings volume in the Santa Fe Institute studies in the sciences of complexity, vol. 3, pp. 267–282. Oxford University Press, US (2006)
16. Montanari, A., Saberi, A.: Convergence to equilibrium in local interaction games. In: 50th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2009, pp. 303–312 (October 2009)
17. Asadpour, A., Saberi, A.: On the Inefficiency Ratio of Stable Equilibria in Congestion Games. In: Leonardi, S. (ed.) WINE 2009. LNCS, vol. 5929, pp. 545–552. Springer, Heidelberg (2009)
18. Fabrikant, A., Papadimitriou, C., Talwar, K.: The complexity of pure nash equilibria. In: Proceedings of the Thirty-Sixth Annual ACM Symposium on Theory of Computing, STOC 2004, pp. 604–612. ACM, New York (2004)
19. Balcan, M.F., Blum, A., Mansour, Y.: Improved equilibria via public service advertising. In: Proceedings of the Twentieth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2009, pp. 728–737. Society for Industrial and Applied Mathematics (2009)
20. Bhalgat, A., Chakraborty, T., Khanna, S.: Approximating pure nash equilibrium in cut, party affiliation, and satisfiability games. In: Proceedings of the 11th ACM Conference on Electronic Commerce, EC 2010, pp. 73–82. ACM (2010)
21. Dyer, M., Mohanaraj, V.: Pairwise-Interaction Games. In: Aceto, L., Henzinger, M., Sgall, J. (eds.) ICALP 2011, Part I. LNCS, vol. 6755, pp. 159–170. Springer, Heidelberg (2011)
22. Berger, N., Kenyon, C., Mossel, E., Peres, Y.: Glauber dynamics on trees and hyperbolic graphs. Probability Theory and Related Fields 131, 311–340 (2005)
23. Bubley, R., Dyer, M.: Path coupling: A technique for proving rapid mixing in markov chains, p. 223. IEEE Computer Society, Los Alamitos (1997)
24. Thilikos, D., Serna, M., Bodlaender, H.: Constructive Linear Time Algorithms for Small Cutwidth and Carving-Width. In: Lee, D.T., Teng, S.-H. (eds.) ISAAC 2000. LNCS, vol. 1969, pp. 192–203. Springer, Heidelberg (2000)
25. Auletta, V., Ferraioli, D., Pasquale, F., Persiano, G.: Metastability of logit dynamics for coordination games. In: Proceedings of the Twenty-Third Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2012, pp. 1006–1024. SIAM (2012)