



Independent Lazy Better-Response Dynamics on Network Games

Paolo Penna, Laurent Viennot

► **To cite this version:**

Paolo Penna, Laurent Viennot. Independent Lazy Better-Response Dynamics on Network Games. CIAC 2019 2019: 11th International Conference on Algorithms and Complexity, May 2019, Rome, Italy. hal-01373411v3

HAL Id: hal-01373411

<https://hal.inria.fr/hal-01373411v3>

Submitted on 5 Feb 2019

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Independent Lazy Better-Response Dynamics on Network Games

Paolo Penna¹ and Laurent Viennot^{2*}

¹ Department of Computer Science, ETH Zurich, Switzerland,
(paolo.penna@inf.ethz.ch)

² Inria – Université Paris Diderot, France,
(Laurent.Viennot@inria.fr)

Abstract. We study an *independent* best-response dynamics on network games in which the nodes (players) decide to revise their strategies independently with some probability. We provide several bounds on the *convergence time* to an *equilibrium* as a function of this probability, the degree of the network, and the potential of the underlying games. These dynamics are somewhat more suitable for distributed environments than the classical better- and best-response dynamics where players revise their strategies “sequentially”, i.e., no two players revise their strategies simultaneously.

1 Introduction

Complex and distributed systems are often modeled by means of *game dynamics* in which the participants (players) act spontaneously, typically striving to maximize their own payoff. Such selfish behavior often results in a so-called (pure Nash) *equilibrium* which, roughly speaking, corresponds to the situation in which no player has an incentive to change her current strategy.³

Consider the natural scenario in which people interact on a (social) network and take their decisions based on both their personal interests and also on what their friends decided. Situations of this sort are often modeled by means of *games* that are played locally by the nodes of some graph (see, e.g., [14] and [13, Chapter 19]). For example, players may have to choose between two alternatives (strategies), and each strategy becomes more valuable if other friends also choose it (perhaps it is easier to agree than to disagree, or it is better to adopt the same technology for working, rather than different ones).

In many cases, an extremely simple procedure to convergence to an equilibrium is the so-called *best-response* dynamics in which at each step one player

* Supported by IRIF (CNRS UMR 8243) and Inria project-team GANG.

³ In this work we consider only pure Nash equilibria, which are the equilibria that occur in certain games when each player chooses one strategy out of the available ones. Other equilibrium concepts are also studied, most notably the mixed Nash equilibrium, where each player chooses a probability distribution over the available strategies.

revises her strategy so to maximize her own payoff (and the others stay put). These dynamics work in more general settings (not only on network games), where convergence to an equilibrium is proven via a potential argument (every move reduces the value of a global function – called potential). Games of this nature are called potential games and they are used to model a variety of situations. Interestingly, this argument *fails* as soon as two or more players *move at the same time*.

In this work we study a natural variant of best-response dynamics in which we relax the requirement that one player at a time moves. That is, now players become active *independently* with some probability and all active players revise their strategy according to the best-response rule (or more generally any better-response rule). This is similar as before but allowing simultaneous moves. Specifically, we study the *convergence time* of these dynamics when players play on a network a “local” potential game: (1) each player interacts only with her neighbors, meaning that the strategies of the non-neighbors do not affect the payoff of this player, and (2) locally the game is a potential game (see Section 2 for more details and formal definitions).

Simple examples show that convergence is impossible if two players are always active (move all the time), or that the time to converge can be made arbitrarily long if they become active at almost every step. At the other extreme, if the probability of becoming active is too small, then the dynamics will also take a long time to converge since almost all the time nothing happens. The trade-off is between having sufficiently many active players and, at the same time, not too many neighboring players moving simultaneously.

1.1 Our Contribution

We investigate how the convergence time depends on the probabilities of becoming active and on the degree of the network. This is also motivated by the search for simple dynamics that the players can easily implement without global knowledge of the network (namely, they only need to know how many neighbors they have), nor without having complex reasoning (they still myopically better-respond). We first show that for the symmetric coordination game, the convergence time is polynomial whenever the probability of being active is slightly below the inverse of the maximum degree of the network (Theorem 2 and Corollary 1). This generalizes to *arbitrary* potential games on graphs, where every node plays a possibly *different* potential game with each of its neighbors, and the maximum degree is replaced by a *weighted maximum degree* (see Theorem 6). These results indeed hold whenever each active player uses a *better response* (not necessarily the best response). Finally, we prove a lower bound saying that, in general, the probabilities of becoming active *must* depend on the degree for otherwise the convergence time is *exponential* with high probability (Theorem 5 and Corollary 2). Note that this holds also for the simplest scenario of symmetric coordination games.

Our upper bounds can be seen as a probabilistic version of the potential argument (under certain conditions, the potential decreases in expectation *at*

every step by some fixed amount). To the best of our knowledge, this is the first study on the convergence time of these natural variants of best-response dynamics. Prior studies (see next section) either focus on sufficient conditions to guarantee convergence to Nash equilibria, or they consider *noisy* best-response dynamics whose equilibria can be different from best-response.

We note that the general upper bound necessarily depends on the maximum value of the potential, as these games include *max-cut* games which are PLS-complete [29]: for such games, no centralized algorithm for computing a Nash equilibrium in time polynomial in the number of players is known, and these games are hard precisely when the potential can assume arbitrarily large values. Obviously, one cannot hope that simple distributed dynamics do much better than the best centralized procedure.

1.2 Related Work

Several works study convergence to Nash equilibria for simple variants of best-response dynamics. A first line of research concerns the ability to converge to a Nash equilibrium when the strict schedule of the moves of the players (one player at a time) is relaxed [10]; they proved that any “separable” schedule guarantees convergence to a Nash equilibrium. Other works study the convergence time of specific dynamics with limited simultaneous moves: [19] introduce a “local” coordination mechanism for congestion games (which are equivalent to potential games [24]), while [15] shows that with limited simultaneous moves the dynamics reaches quickly a state whose cost is not too far from the worst Nash equilibrium [15]; Note that the reached state need not be a Nash equilibrium, and the required condition is that every T consecutive best responses, each player has moved at least once and at most β times. Fast convergence can be achieved in certain linear congestion games if *approximate* equilibria are considered [9], that is, players keep changing their strategies as long as a significant improvement is possible.

Another well-studied variant of best-response dynamics is that of *noisy* or *logit (response)* dynamics [6,7,1], where players’ responses is *probabilistic* and determined by a noise parameter (as the noise tends to zero, players select almost surely best-responses, while for high noise they respond at random). These dynamics turn out to behave differently from “deterministic” best-response in many aspects. In the original logit dynamics by [6,7], where one randomly chosen player moves at a time, they essentially rest on a subset of potential minimizers. When the players’ schedule is relaxed, this property is lost and additional conditions on the game are required [1,10,2,18,26]. Our *independent* better-response dynamics can be seen as an analog of the independent dynamics of [1] for logit response.

Potential games on graphs (a proper subclass of potential games) are well-studied because of their many applications. In physics, ferro-magnetic systems are modeled as *noisy* best-response dynamics on lattice graphs in which every player (node) plays a coordination game with each neighbor (see, for example,

[23] and Chapter 15 of [22]). The version in which the coordination game is asymmetric (i.e., coordinating on one strategy is more profitable than another) is used to model the diffusion of new technologies [25,21] and opinions [17] in social networks. Finally, potential games on graphs (every node plays some potential game with each neighbor) characterize the class of potential games for which the equilibria of *noisy* best-response dynamics with *all* players updating simultaneously can be “easily” computed [3]. The convergence time of best-response dynamics for games on graphs is studied in [12,17]: Among other results, [12] showed that a polynomial number of steps are sufficient when the same game is played on all edges and the number of strategies is constant. Analogous results are proven for finite opinion games in [17]. Finally, [4] characterize the class of potential games which are also *graphical* games [20], where the potential can be decomposed into the sum of potentials of “maximal” cliques of an underlying graph. Graphical games have been studied in several works (see, e.g., [11,8,27,5]). The class of *local interaction* potential games [3] is the restriction in which the potential can be decomposed into pairwise (edge) potential games. In this work we deal precisely with this class of games. Since this class includes the so-called max-cut games, which are known to be PLS-complete [29], it is considered unlikely that an equilibrium can be computed efficiently, even by a centralized procedure.

Our dynamics are similar to the α -synchronous dynamics in cellular automata [16]. This latter model encompasses ours as long as the response strategy can be implemented through an automata. This is the case for example with symmetric coordination where best-response corresponds to majority rule (where each cellular automaton tries to switch to the majority state of its neighbors). However, most of the literature around cellular automata considers a regular lattice topology and do not cover the case of general graphs we consider here. A notable exception is the study of minority rule on general graphs [28] with similar dynamics as ours. The present work can be seen as a first study of α -synchronous dynamics on general graphs for the rules that follow from best-response to some potential games with neighbors. Our results apply in particular to majority and minority rules (where no action is taken in case of equality, i.e, when the two most frequent states are equally represented among neighbors).

2 Model (Local Interaction Potential Games)

Intuitively speaking we consider a network (graph) where each node is a player who repeatedly plays with her neighbors. We assume that a two-player potential game (defined below) is associated to each edge of the graph. Each player must play the same strategy on all the games associated to its incident edges, and her payoff is the sum of the payoffs obtained in each of these games. We also assume finite strategies, i.e. each player chooses her strategy within a finite set.

Symmetric Coordination Game. One of the simplest (potential) games is the *symmetric coordination game* where each player chooses color B or W (for black or white) and her payoff is 1 if players agree on their strategies, and 0 otherwise

	B	W
B	$1, 1$	$0, 0$
W	$0, 0$	$1, 1$

	B	W
B	$2, 1$	$0, 0$
W	$0, 0$	$1, 2$

	B	W
B	-2	-1
W	0	-2

	B	W
B	$2, 2$	$1, 1$
W	$0, 0$	$2, 2$

(a) Symmetric Coordination Game. (b) Another Coordination Game. (c) Potential for Game (b). (d) Game equivalent to Game (b).

Fig. 1: Examples of (two-players) potential games. The game (b) has potential function (c). Game (d) has same potential function (c) (signs changed for payoff), it is equivalent to Game (b) (both games have same dynamics).

(see Figure 1a where the two numbers are the payoff for the row and the column player, respectively).

General Potential Games. In a general game, we have n players, and each of them can choose one color (strategy) and the combination $c = (c_1, \dots, c_n)$ of all colors gives to each player u some payoff $PAY_u(c)$. In a *potential game*, when the change in the payoff of any player improves by some amount, some global function P called the *potential* will be decreased by the same amount: For any player u and any two configurations c and c' which differ only in u 's strategy, it holds that

$$PAY_u(c') - PAY_u(c) = P(c) - P(c') . \tag{1}$$

A configuration c is a (*pure Nash*) *equilibrium* if no player u can improve her payoff, that is, the quantity above is negative or zero for all $c' = (c_1, \dots, c'_u, \dots, c_n)$. Conversely, c is not an equilibrium if there is a player u who can improve her payoff ($PAY_u(c') - PAY_u(c) > 0$) in which case c'_u is called a *better response* (to strategies c). A *best response* is a better response maximizing this improvement, over the possible strategies of the player. Potential games possess the following nice feature: A configuration c is an equilibrium if and only if no player can improve the potential function by changing her current strategy. In a general (two-player) potential game the payoff of the players is not the same, and the potential function is therefore not symmetric (see the example in Figure 1c).

Local interaction potential games [3]. In a local interaction potential game the potential function can be decomposed into the sum of two-player potential games, one for each edge of the network G :

$$P(c) = \sum_{uv \in E(G)} P_{uv}(c_u, c_v) . \tag{2}$$

No edge exists if the strategies of the two players do not affect each others' payoff (the corresponding potential is constant and can be ignored). This definition captures the following natural class of games on networks: Each edge corresponds to some potential game, and the payoff of a player is the sum of the payoffs of

the games with the neighbors. Note that a player chooses one strategy to be played on all these games.

(Independent) Better-Response Dynamics. A simple procedure for computing an equilibrium consists of repeatedly selecting *one* player who is currently not playing a best response and let her play a better or best response. Every step reduces the potential by a finite amount, and therefore this procedure terminates into an equilibrium in $O(M)$ time steps, where M is the maximum value for the potential (w.l.o.g., we assume that the potential is always non-negative and takes integer values⁴). Here we consider the variant in which, at each time step, each player becomes *independently* active according to some probability, and those who can improve their payoff change strategy accordingly:

Definition 1. *In independent better-response dynamics, at each time step t players do the following:*

- *Each player (node) u becomes active with some probability p_u^t which can change over time (the case in which it is constant over time is a special case of this one).*
- *Every active player (node) revises her strategy according to a better (or best) response rule. If the current strategy is already a best response, then no change is made.*

Note that all players that are active at a certain time step may change their strategies *simultaneously*. So, for example, it may happen that on the symmetric coordination game in Figure 1a the two players move from state BW to state WB and back if they are both active all the time.

Generic upper bound To show that dynamics converge quickly, we show that the potential decreases in expectation at every step. To this end, we consider the probability space of all possible evolutions of the dynamics. A configuration c at a given time t is given by the colors chosen by players at the previous time step (strategy profile) and by the values p_u^t used by users for randomly deciding to be active at time t . The universe Ω is then defined as the set of all infinite sequences c^0, c^1, \dots of configurations.

Definition 2 (δ -improving dynamics). *Dynamics are δ -improving for a given (local interaction) potential game if in expectation the potential decreases by at least δ during each time step, unless the current configuration is an equilibrium. That is, for any configuration c which is not an equilibrium, and any event $F_c^t = \{c^0, c^1, \dots \in \Omega \mid c^t = c\}$ where configuration c is reached at time t , we have*

$$E[P^{t+1} - P^t \mid F_c^t] \leq -\delta$$

⁴ As we assume that strategy sets are finite, the potential function is defined by a finite set of values. Rescaling the potential function so that different values are at least 1 apart, and then truncating the values to integers allows to obtain an equivalent game (with same dynamics). Additionally shifting the values allows to obtain a non-negative potential function for that game.

where P^t denotes the potential at time t .

The proof of the following theorem is based on standard Martingale arguments and it is given in Appendix A.1 for completeness.

Theorem 1. *The expected convergence time of any δ -improving dynamics is $O\left(\frac{M_0}{\delta}\right)$ where M_0 is the expected potential of the game at time 0.*

3 Networks With Symmetric Coordination Games

We first consider the scenario in which every edge of the network is the symmetric coordination game in Figure 1a. The nodes of a graph G (players) can choose between two colors B and W and are rewarded according to the number of neighbors with same color. We are thus considering the dynamics in which nodes attempt to choose the *majority color of their neighbors* and every active node changes its color if more than half of its neighbors has the different color.

In order to analyze the convergence time of these dynamics, we shall relate the probabilities of being active to the number of neighbors having a different color. We say that u is *unstable* at time t if more than half of the neighbors has the other color, that is,

$$dc_u^t > \frac{1}{2}\delta_u$$

where δ_u is the degree of u and dc_u^t is the number of neighbors of u that have a color different from the color of u at time t . By definition, the dynamics converge if no node is unstable. Note that we have $dc_u^t \leq \delta_u \leq \Delta_u \leq \Delta$ where $\Delta = \max_{u \in V(G)} \delta_u$ is the maximum degree of the graph, and $\Delta_u = \max_{uv \in E(G)} \delta_v$ is the local maximum degree in the neighborhood of u .

For the case of symmetric coordination games, the potential function of a configuration is the number of edges whose endpoints have different colors: An edge uv is said to be *conflicting* in configuration c if u and v have different colors. Therefore the potential is at most the number m of edges.

Theorem 2. *Fix some real values $p, q \in (0, 1)$. If we have $p_u^t \in \left[\frac{p}{\Delta}, \frac{q}{\Delta_u}\right]$ for all u, t in a symmetric coordination game, then the expected convergence time is $O\left(\frac{\Delta m_0}{p(1-q)}\right)$ where m_0 is the initial number of conflicting edges, Δ is the maximum degree, and Δ_u is the maximum degree in the neighborhood of u .*

As an immediate corollary, we have the following result for the case in which all nodes are active with the *same* probability p .

Corollary 1. *If all unstable nodes are active with probability $p < \frac{1-\varepsilon}{\Delta}$ for $\varepsilon > 0$, then the dynamics converge to a stable state in $O\left(\frac{m_0}{p\varepsilon}\right)$ expected time.*

Theorem 2 derives from the following lemma and Theorem 1.

Lemma 1. *Any dynamics satisfying the hypothesis of Theorem 2 are δ -improving for $\delta = p(1 - q)/\Delta$.*

Proof. Consider the event F_c^t where a configuration c is reached at time t . Let C^t denote the number of conflicting edges in c , and U^t be the set of unstable nodes at time t respectively. Recall that the number of conflicting edges is equal to the potential, that is, $P^t = C^t$. We now express $E[C^{t+1} - C^t \mid F_c^t]$ as a function of the values $\{p_u^t \mid u \in V(G)\}$ associated to c .

For that purpose, we first analyze the probability that any given edge of c is conflicting after the random choices made at time t . We distinguish the following types of edges. Let S_1 (resp. S_2) denote the set of edges in c with the same color and one unstable extremity (resp. two). Similarly, let C_1 (resp. C_2) denote the set of edges in c with conflicting colors and one unstable extremity (resp. two). Note that $C^t = |C_1| + |C_2|$. A conflicting edge uv will become non-conflicting if only one extremity changes its color. Similarly, a non-conflicting edge uv will become conflicting if only one extremity changes its color. Due to independence of choices, this happens in both cases with probability $p_{uv}^t = p_u^t(1-p_v^t) + (1-p_u^t)p_v^t$ if both u and v are unstable, and with probability p_u^t if u is unstable and v is not. By linearity of expectation, we then obtain:

$$E[C^{t+1} - C^t \mid F_c^t] = \sum_{uv \in S_1} p_u^t + \sum_{uv \in S_2} p_{uv}^t - \sum_{uv \in C_1} p_u^t - \sum_{uv \in C_2} p_{uv}^t . \quad (3)$$

(When we note $uv \in C_1$ (resp. $uv \in S_1$), we assume that u is unstable and v is not.) By definition, each unstable node u sees more conflicting edges than non-conflicting ones, thus implying By multiplying by p_u^t and then summing over all unstable nodes, we obtain:

$$\sum_{u \in U^t} p_u^t + \sum_{uv \in S_1} p_u^t + \sum_{uv \in S_2} (p_u^t + p_v^t) \leq \sum_{uv \in C_1} p_u^t + \sum_{uv \in C_2} (p_u^t + p_v^t) . \quad (4)$$

As $p_{uv}^t = p_u^t + p_v^t - 2p_u^t p_v^t$, we deduce from (3) and (4):

$$E[C^{t+1} - C^t \mid F_c^t] \leq \sum_{uv \in C_2} 2p_u^t p_v^t - \sum_{u \in U^t} p_u^t . \quad (5)$$

Since every edge $uv \in C_2$ has both endpoints in U^t , we can rewrite (5) as

$$E[C^{t+1} - C^t \mid F_c^t] \leq \sum_{u \in U^t} p_u^t \left(-1 + \sum_{v \mid uv \in C_2} p_v^t \right) .$$

Using $p_v^t \leq \frac{q}{\Delta_v} \leq \frac{q}{\delta_u}$ and $p_u^t \geq \frac{p}{\Delta}$, we obtain the following inequality: $E[C^{t+1} - C^t \mid F_c^t] \leq \sum_{u \in U^t} \frac{p}{\Delta} (-1 + q) = -p(1-q) \frac{|U^t|}{\Delta}$. This completes the proof. \square

Adaptive probabilities The upper bound of Theorem 2 can be improved if nodes are aware of the number of neighbors that are willing to change strategy (unstable) and then set accordingly the probability of changing too. More precisely, one can think of active nodes announcing to their neighbors that they are unstable and that they would like to switch to the other color, before actually doing so. Then, each unstable node will switch with a probability inversely

proportional to the number of unstable neighbors. The following theorem shows that this yields an improved upper bound on the convergence time.

Theorem 3. *Fix some real values $p, q \in (0, \frac{1}{2})$. If we have $p_u^t \in [\frac{p}{d_u^t+1}, \frac{q}{d_u^t+1}]$ for all u, t in a symmetric coordination game, where d_u^t is the number of conflicting unstable neighbors of u , then the expected convergence time is $O\left(\frac{m_0}{p(1-2q)}\right)$ where m_0 is the initial number of conflicting edges.*

To prove this theorem we adapt the proof of Lemma 1 and show that these dynamics are δ -improving for $\delta = p(1-2q)$ (see Appendix A.4 for details).

Fully local dynamics Theorem 2 requires that each node is aware of a bound on the maximum degree, or the local maximum degree in her neighborhood for setting p_u^t . Theorem 3 requires knowledge of the number of conflicting unstable neighbors at each time step. We next consider dynamics that are fully local as each node u can set the probabilities p_u^t by only looking at its own degree.

Theorem 4. *Fix some real values $p, q \in (0, \frac{1}{2})$. If we have $p_u^t \in [\frac{p}{\delta_u}, \frac{q}{\delta_u}]$ for all u, t in a symmetric coordination game, where δ_u is the degree of u , then the expected convergence time is $O\left(\frac{\Delta m_0}{p(1-2q)}\right)$ where m_0 is the initial number of conflicting edges.*

The proof of this theorem is similar to that of Theorems 2 and 3 (see Appendix A.5).

Tightness of the results We can use the example in Figure 2 to show that the analysis of Theorems 2, 3, and 4 is tight, and adaptive dynamics are provably faster than non-adaptive ones. Intuitively, the construction is such that the following holds:

- Initially only the extreme nodes of the leftmost path are unstable. After an unstable node become black, the next neighbor in the path (and only that one) becomes unstable.
- The process proceeds from left to right, and nodes in certain path become unstable only after all nodes in the previous path became black (except for the last path with $r/2$ nodes).
- Inside the path containing the only unstable nodes, the process is sequential: the path becomes black from extremities to center. At most two nodes are unstable: the extremities of the sub-path of white nodes. meaning that after one node changes color, its neighbor (and only that) becomes unstable.

These observations imply that any dynamics in which nodes become active with probability $p \simeq \alpha$, require $\Omega(r^2/\alpha) = \Omega(n/\alpha)$ steps.

Since every node has degree $\Theta(r) = \Theta(\sqrt{n}) = \Theta(\Delta)$, and the initial configuration has $m_0 = \Theta(r^2) = \Theta(n)$ conflicting edges (those between the clique and the first path), non-adaptive dynamics take $\Theta(\Delta m_0) = \Theta(n^{3/2})$ time steps.




Fig. 2: Example of network with faster adaptive dynamics than non-adaptive ones. The network is composed of $r/2 + 2$ parts for even r : a clique and $r/2 + 1$ paths. Below each part, we indicate the number of nodes in the part. Each node in a path is connected to all nodes to the right and to the left path (or clique for the first path) as feature by demi-edges with degree indications w.r.t. the previous and the next part of the construction.

On the contrary, adaptive dynamics take $\Theta(m_0) = \Theta(n)$ steps since the number d_u^t of unstable conflicting neighbors of each node u is at most 1. Therefore, the analysis of Theorems 2, 3, and 4 is tight. Moreover, the adaptive dynamics are provably faster than non-adaptive ones.

4 An Exponential Lower Bound when the Degree is Unbounded

In this section we prove a lower bound for the case of symmetric coordination game on each edge and dynamics with constant probabilities, that is, the case in which every node becomes active with some probability p which does not depend on the graph nor on the time, and which is the same over all nodes.

Theorem 5. *For every $p > 0$, there are starting configurations of the complete bipartite graph where the expected number of steps to converge to an equilibrium is exponential in the number of nodes.*

Proof Idea. Consider the *continuous* version of the problem in which, instead of a bipartite graph with n nodes on each side, we imagine L and R being two continuous intervals (see figure below). Start from a symmetric configuration in which a fraction $\alpha > 1/2$ of the players in L has color W and the same fraction in R has the other color B . Suppose that

$$\alpha = \frac{1}{2-p}.$$

Then after one step the system reaches the symmetric configuration, that is, a fraction α of nodes in L has color B and the same fraction in R has color W .

Indeed, the fraction β of players with color B in L after one step is precisely

$$\beta = 1 - \alpha + p \cdot \alpha = \frac{1-p}{2-p} + \frac{p}{2-p} \alpha = \alpha .$$




Fig. 3: The idea of the proof of Theorem 5 (active nodes change color leading to a symmetric configuration for suitable α). All the nodes on the left (L) are adjacent to all nodes on the right (R).

We next prove the theorem via Chernoff bounds. For $\epsilon = p/3$ consider the interval $around(\alpha) := [(1 - \epsilon)\alpha, (1 + \epsilon)\alpha]$, and let $CYCLE(t)$ be the following event:

$$CYCLE(t) := \{\text{At time } t \text{ a fraction } \alpha_L \in around(\alpha) \text{ of the nodes in } L \text{ has some color } c, \text{ and a fraction } \alpha_R \in around(\alpha) \text{ of the nodes in } R \text{ has the other color } \bar{c} \text{ (where } \bar{B} = W \text{ and } \bar{W} = B).\}$$

We say that the configuration is *balanced* at time t when $CYCLE(t)$ holds. Since $\epsilon < p/2$ we have $(1 - \epsilon)\alpha > 1/2$, and thus the best response of every (active) node in a balanced configuration is to switch color (since both $\alpha_L n$ and $\alpha_R n$ are strictly larger than $n/2$). Chernoff bounds guarantee that with high probability enough many nodes will be activated and therefore will switch to obtain a symmetric balanced configuration. Formally, the following lemma holds (see Appendix A.2 for the proof).

Lemma 2. *For any t , it holds that $P[CYCLE(t+1) | CYCLE(t)] \geq 1 - 4 \exp\left(-\frac{\delta^2}{3} \mu\right)$, where $\delta = \frac{\epsilon}{1+\epsilon}$ and $\mu = p(1 + \epsilon)\alpha n$ with $\epsilon = p/3$.*

Proof (of Theorem 5). Consider any starting configuration which is balanced, that is, $CYCLE(0)$ holds. By Lemma 2, the probability that the event $CYCLE(1) \wedge \dots \wedge CYCLE(t)$ holds is at least $(1 - q)^t$ where $q = 4 \exp\left(-\frac{\delta^2}{3} \mu\right)$. Since $CYCLE(t)$ implies that at least one node is unstable, the probability that we reach an equilibrium in t steps is thus at least $(1 - q)^{t-1}$. The expected time for convergence is thus at least $\sum_{t \geq 0} t(1 - q)^{t-1} = 1/q^2$. \square

Simple calculations lead to the following result (see Appendix A.3 for details).

Corollary 2. *Starting from any balanced configuration, the expected number of steps to converge to an equilibrium in the complete bipartite graph is $e^{\Omega(n^{1-3c})}$, as long as $p \geq 1/n^c$ with $0 \leq c < 1/3$.*

5 General Local Interaction Potential Games

In this section we extend the upper bound of Theorem 2 to general local interaction potential games: each edge uv of G is associated with a (two-player) potential game with potential P_{uv} . Without loss of generality, we assume that the potential P_{uv} takes integer non-negative values. The upper bound is given in terms of the following quantity:

$$\Delta_P := \max_u \sum_{v \in N(u)} \Delta_{P_{uv}} , \quad (6)$$

where $\Delta_{P_{uv}}$ denotes the maximum value of P_{uv} . Note that for symmetric coordination games, Δ_P is simply the maximum degree Δ of the graph.

Theorem 6. *For any $p, q \in (0, 1/2)$, if we have $p_u^t \in [\frac{p}{\Delta_P}, \frac{q}{\Delta_P}]$ for all u and t and for Δ_P defined as in (6) in a general local interaction potential game, then the expected convergence time is $O\left(\frac{n\Delta_P^2}{p(1-2q)}\right)$.*

We shall prove later that $n\Delta_P$ is an upper bound on the global potential. Therefore to prove the theorem it suffice to show that in expectation the global potential improves as follows:

Lemma 3. *Any dynamics satisfying the hypothesis of Theorem 6 is δ -improving for $\delta = p(1-q)/\Delta_P$.*

Proof. Consider the event F_c^t where a configuration c is reached at time t . Let $P^t := P(c^t) = \sum_{uv \in E(G)} P_{uv}(c_u^t, c_v^t)$ denote the global potential of the game in configuration c . A node is unstable when it can improve its total payoff by changing c_u^t to some c'_u , which corresponds to an equal improvement on the global potential of the game according to (1). We then let

$$\mu_{uv}^u = -(P_{uv}(c'_u, c_v^t) - P_{uv}(c_u^t, c_v^t))$$

denote the improvement obtained on the game with a neighbor v , and the sum $\sum_{v \in N(u)} \mu_{uv}^u \geq 1$ being the total improvement. Node u then changes for strategy c'_u with probability p_u^t .

Let C_1 denote the edges uv such that u is unstable at time t and v is not. Let C_2 denote the edges uv such that both u and v are unstable. The set of unstable nodes at time t is denoted by U^t . By linearity of expectation, we have

$$E[P^{t+1} - P^t \mid F_c^t] = - \sum_{uv \in C_1} p_u^t \mu_{uv}^u - \sum_{uv \in C_2} (p_u^t(1-p_v^t)\mu_{uv}^u + p_v^t(1-p_u^t)\mu_{uv}^v - p_u^t p_v^t \nu_{uv}^{uv}) \quad (7)$$

where $\nu_{uv}^{uv} = P_{uv}(c'_u, c'_v) - P_{uv}(c_u^t, c_v^t)$ denotes the variation in the potential of edge uv when both u and v change their colors.

As each unstable node expects a gain of at least 1 (since the potential takes integer values), we have

$$\sum_{uv \in C_1} p_u^t \mu_{uv}^u + \sum_{uv \in C_2} (p_u^t \mu_{uv}^u + p_v^t \mu_{uv}^v) \geq \sum_{u \in U^t} p_u^t . \quad (8)$$

By combining (7) and (8) we obtain

$$E[P^{t+1} - P^t \mid F_c^t] \leq - \sum_{u \in U^t} p_u^t + \sum_{uv \in C_2} p_u^t p_v^t (\mu_{uv}^u + \mu_{uv}^v + \nu_{uv}^{uv}) .$$

The last factor in the second sum simplifies to

$$\mu_{uv}^u + \mu_{uv}^v + \nu_{uv}^{uv} = P_{uv}(c_u^t, c_v^t) + P_{uv}(c'_u, c'_v) - P_{uv}(c'_u, c_v^t) - P_{uv}(c_u^t, c'_v) \leq 2\Delta_{P_{uv}} .$$

We thus obtain

$$E[P^{t+1} - P^t \mid F_c^t] \leq - \sum_{u \in U^t} p_u^t + \sum_{uv \in C_2} 2p_u^t p_v^t \Delta_{P_{uv}} = \sum_{u \in U^t} p_u^t (-1 + \sum_{v|uv \in C_2} p_v^t \Delta_{P_{uv}}) .$$

Since $\sum_{v|uv \in C_2} \Delta_{P_{uv}} \leq \Delta_P$, $p_u^t \geq \frac{p}{\Delta_P}$, and $p_v^t \leq \frac{q}{\Delta_P}$, we obtain $E[P^{t+1} - P^t \mid F_c^t] \leq -p(1 - q) \frac{|U^t|}{\Delta_P}$. This completes the proof of the lemma. \square

Theorem 6 then follows from Theorem 1 as we observe that the global potential is bounded from above by $\sum_{uv \in E(G)} \Delta_{P_{uv}} = \sum_u \sum_{v|uv \in E(G)} \Delta_{P_{uv}}/2 \leq n\Delta_P/2$.

5.1 Two Examples of Local Interaction Potential Games

In this section we discuss two examples of local interaction potential games and how the result of Theorem 6 relates to them.

Finite Opinion Games [17] Consider a game where each player must decide between two opinions, 0 or 1, and each player has some internal belief $b_u \in (0, 1)$. Since there are only two strategies, $c_u \in \{0, 1\}$, better-response and best-response coincide. According to [17], the payoff of player u depends on the opinion of her neighbors and on her own belief,

$$PAY_u(c) = - \left((c_u - b_u)^2 + \sum_{v \in N(u)} (c_u - c_v)^2 \right) .$$

and game is a potential game, with the potential function being

$$P(c) = C(c) + \sum_u (c_u - b_u)^2 ,$$

where $C(c)$ is the number of conflicting edges, that is, the edges whose endpoints have different opinions. We can easily see that these are local interaction potential games by viewing the game as follows:

- Between u and every neighbor $v \in N(u)$, we play the symmetric coordination game (payoff is 1 if they have the same opinion, and 0 otherwise).
- Each player u is connected to her own *opinion node*, that is, a dummy node u^* who has only strategy b_u available. Between u and u^* we play a “trivial” coordination game in which u has payoff $P_{uu^*}(c_u) = -(c_u - b_u)^2$.

Finally, in better- or best-response dynamics, one can restrict to $b_u \in \{1/4, 3/4\}$ [17], which makes it possible to have integral potential (simply multiply all payoffs by 16). Overall, we can bound the quantity Δ_P in (6) as

$$16\Delta \leq \Delta_P \leq 16(\Delta + 1)$$

where Δ is the maximum degree in the network (without the dummy nodes u^*). Theorem 6 then yields the following:

Corollary 3. *In finite opinion games on networks of maximum degree Δ , the expected converge time of independent better-response dynamics is $O(n\Delta^2)$ whenever $p_u^t = \frac{\alpha}{\Delta}$ for some $\alpha \in [\frac{p}{16}, \frac{q}{16}]$ with $p, q \in (0, 1/2)$.*

The next example says that one cannot hope to get polynomial convergence time independent of ‘ Δ_P ’ in general.

Max-Cut Games [29] Consider a complete *weighted* graph in which every edge $e = (u, v)$ corresponds to a symmetric coordination game rescaled by the weight w_e of this edge,

	B	W
B	w_e, w_e	$0, 0$
W	$0, 0$	w_e, w_e

The two possible strategies of each player (B or W) determine a partition of the players into two groups, and the potential is simply the value of the cut of the resulting partition (sum of edge weights between the two groups). These problems are PLS-complete [29] meaning that one does not expect to have an algorithm that computes a Nash equilibrium in time polynomial in the number of players, regardless of the weights.

Max-cut games are hard when the weights are arbitrary, that is, when the Δ_P term is not polynomially-bounded. In such instances, even *centralized algorithms* are not expected to find a Nash equilibrium in time polynomial in the number of players.

6 Conclusion

This work provides bounds on the time to converge to a (pure Nash) equilibrium when players are active *independently* with some probability and they better or

best respond to each others current strategies. Our study focuses on a natural (sub)class of potential games, namely, local interaction potential games. The bounds suggest that the time to converge to an equilibrium *must* depend on the *degree* of the nodes in the underlying network:

- These (distributed) dynamics can *converge quickly* if players are *lazy*, that is, if the probability of being active is inversely proportional to the maximum degree (the precise bounds are given by Theorems 2 and 6).
- Conversely, *non-lazy* dynamics can take *exponential* time even in simple cases, where non-lazy means that the probability of being active is too high with respect to the network degree (e.g., when this probability is $1/\Delta^\alpha$ for some $\alpha < 1/3$ and Δ being the maximum degree – see Corollary 2).

This suggests a sort of threshold effect due to the maximum degree of the network. This parameter has a natural interpretation as it corresponds to the number of players (nodes) that can affect the payoff of a single player (node).

Since our bounds hold for local interaction potential games, it would be interesting to investigate whether analogous results hold for *general* potential games (note that one can always construct a graph representing the dependencies between players, by connecting two players whenever payoff of one depends also on the strategy of the other). Here a relevant notion is that of graphical games [20] and the results in [4]. It would also be interesting to sharpen some of our bounds to show that $p \simeq 1/\Delta$ is essentially the threshold between fast and slow convergence: Is it the case that, for every $\alpha < 1$, if all nodes are active with some probability $p \simeq 1/\Delta^\alpha$, then the dynamics does not converge in polynomial time in some graphs of maximum degree Δ ? Finally, it would be interesting to investigate the range $p \in [1/n, 1/n^{1/3}]$ (according to Theorem 2 and Corollary 2, this is where convergence time seems to change from polynomial to exponential).

Acknowledgments We thank Damien Regnault and Nicolas Schabanel for inspiring discussions on closely related problems, and an anonymous reviewer for pointing out the last open question. Part of this work has been done while the first author was at LIAFA, Université Paris Diderot, supported by the French ANR Project DISPLEXITY.

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. On the convergence of logit-response to (strict) nash equilibria. *Economic Theory Bulletin*, 5(1):1–8, 2017.
3. Vincenzo Auletta, Diodato Ferraioli, Francesco Pasquale, Paolo Penna, and Giuseppe Persiano. Logit dynamics with concurrent updates for local interaction potential games. *Algorithmica*, 73(3):511–546, 2015.
4. Yakov Babichenko and Omer Tamuz. Graphical potential games. *Journal of Economic Theory*, 163:889 – 899, 2016.
5. James P. Bailey and Georgios Piliouras. Multiplicative weights update in zero-sum games. In *ACM Conference on Economics and Computation (EC)*, pages 321–338, 2018.

6. Lawrence E. Blume. The statistical mechanics of strategic interaction. *Games and economic behavior*, 5(3):387–424, 1993.
7. Lawrence E. Blume. *Population games*. Addison-Wesley, 1998.
8. Yang Cai and Constantinos Daskalakis. On minmax theorems for multiplayer games. In *Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 217–234, 2011.
9. Steve Chien and Alistair Sinclair. Convergence to approximate nash equilibria in congestion games. *Games and Economic Behavior*, 71(2):315 – 327, 2011.
10. Pierre Coucheney, Stéphane Durand, Bruno Gaujal, and Corinne Touati. General revision protocols in best response algorithms for potential games. In *Network Games, Control and Optimization (NetGCoop)*, 2014.
11. Constantinos Daskalakis and Christos H. Papadimitriou. On a network generalization of the minmax theorem. In *International Colloquium on Automata, Languages, and Programming Automata (ICALP)*, pages 423–434, 2009.
12. Martin Dyer and Velumailum Mohanaraj. Pairwise-interaction games. In *International Colloquium on Automata, Languages and Programming (ICALP)*, volume 6755 of *LNCS*, pages 159–170, 2011.
13. David Easley and Jon Kleinberg. *Networks, crowds, and markets*. Cambridge Univ Press, 2012.
14. Glenn Ellison. Learning, local interaction, and coordination. *Econometrica*, 61(5):1047–1071, 1993.
15. Angelo Fanelli, Luca Moscardelli, and Alexander Skopalik. On the impact of fair best response dynamics. In *Proc. of the 37th international conference on Mathematical Foundations of Computer Science (MFCS)*, volume 7464 of *LNCS*, pages 360–371, 2012.
16. Nazim Fatès, Damien Regnault, Nicolas Schabanel, and Eric Thierry. Asynchronous behavior of double-quiescent elementary cellular automata. In *Latin American Symposium on Theoretical Informatics (LATIN)*, volume 3887 of *LNCS*, pages 455–466, 2006.
17. Diodato Ferraioli, Paul W. Goldberg, and Carmine Ventre. Decentralized dynamics for finite opinion games. *Theoretical Computer Science*, 648:96 – 115, 2016.
18. Diodato Ferraioli and Paolo Penna. Imperfect best-response mechanisms. *Theory of Computing Systems*, 57(3):681–710, Oct 2015.
19. Dimitris Fotakis, Alexis C. Kaporis, and Paul G. Spirakis. Atomic congestion games: Fast, myopic and concurrent. *Theory of Computing Systems*, 47(1):38–59, 2010.
20. Michael Kearns, Michael L. Littman, and Satinder Singh. Graphical models for game theory. In *Conference on Uncertainty in Artificial Intelligence (UAI)*, pages 253–260, 2001.
21. Gabriel E. Kreindler and H. Peyton Young. Rapid innovation diffusion in social networks. *Proc. of the National Academy of Sciences*, 111(Supplement 3):10881–10888, 2014.
22. David Asher Levin, Yuval Peres, and Elizabeth Lee Wilmer. *Markov chains and mixing times*. American Mathematical Soc., 2009.
23. Fabio Martinelli. Lectures on glauber dynamics for discrete spin models. In Pierre Bernard, editor, *Lectures on Probability Theory and Statistics*, volume 1717 of *Lecture Notes in Mathematics*, pages 93–191. Springer Berlin Heidelberg, 1999.
24. Dov Monderer and Lloyd S. Shapley. Potential games. *Games and Economic Behavior*, 14(1):124 – 143, 1996.
25. Andrea Montanari and Amin Saberi. The spread of innovations in social networks. *Proceedings of the National Academy of Sciences*, 107(47):20196–20201, 2010.

26. Paolo Penna. The price of anarchy and stability in general noisy best-response dynamics. *International Journal of Game Theory*, 47(3):839–855, 2018.
27. Georgios Piliouras and Jeff S. Shamma. Optimization despite chaos: Convex relaxations to complex limit sets via poincaré recurrence. In *Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 861–873, 2014.
28. Jean-Baptiste Rouquier, Damien Regnault, and Eric Thierry. Stochastic minority on graphs. *Theoretical Computer Science*, 412(30):3947–3963, 2011.
29. Alejandro A Schäffer and Mihalis Yannakakis. Simple local search problems that are hard to solve. *SIAM journal on Computing*, 20(1):56–87, 1991.

A Postponed proofs

A.1 Proof of Theorem 1

We make use of the following lemma which derives from classical martingale theory (we include a proof for the sake of completeness).

Lemma 4. *Let $(X^t)_{t \in \mathbb{N}}$ a sequence of discrete random variables with values in $\{0, \dots, M\}$, and let $T = \min \{t : X^t = 0\}$ be the stopping time defined as the random variable for the first time t where $X^t = 0$. Suppose that for some $\epsilon > 0$*

$$E[X^{t+1} - X^t \mid X^0 = x^0 \wedge \dots \wedge X^t = x^t] \leq -\epsilon$$

for all tuples of values $x^0, \dots, x^t \in \{1, \dots, M\}$. Then we have $E[T] \leq \frac{E[X^0]}{\epsilon}$.

Proof. We first show $E[T] < \infty$. As the variables X^t are bounded we can obtain a lower bound on the probability that X^t decreases by $\epsilon/2$ at time t : for any tuple of positive values (x^0, \dots, x^t) , we consider the event $F = \{X^0 = x^0 \wedge \dots \wedge X^t = x^t\}$ and prove

$$P[X^{t+1} - X^t < -\epsilon/2 \mid F] \geq \frac{\epsilon}{2M} . \quad (9)$$

From the definition of conditional expectation, we have

$$E[X^{t+1} - X^t \mid F] = \frac{\int_F (X^{t+1} - X^t) dP}{P[F]} .$$

The assumption on the expected variation of X^t at time t then implies

$$-\epsilon P[F] \geq \int_{\{X^{t+1} - X^t < -\epsilon/2\} \cap F} (X^{t+1} - X^t) dP + \int_{\{X^{t+1} - X^t \geq -\epsilon/2\} \cap F} (X^{t+1} - X^t) dP .$$

Using $X^{t+1} - X^t \geq -M$, we get

$$-\epsilon P[F] \geq -MP[\{X^{t+1} - X^t < -\epsilon/2\} \cap F] - \frac{\epsilon}{2} P[\{X^{t+1} - X^t \geq -\epsilon/2\} \cap F] .$$

Equation 9 is then deduced from $P[F] \geq P[\{X^{t+1} - X^t \geq -\epsilon/2\} \cap F]$.

Now, for integral $\tau > 1$, we have $X^{t+\tau} \leq X^t - \epsilon\tau/2$ when $X^{s+1} - X^s < -\epsilon/2$ for $s \in \{t, \dots, t + \tau - 1\}$. We thus have:

$$P[X^{t+\tau} \leq X^t - \epsilon\tau/2 \mid T \geq t + \tau - 1] \geq (\epsilon/2M)^\tau .$$

Fix $\tau > 2M/\epsilon$. As $X^t \leq M$, we have $X^t - \epsilon\tau/2 < 0$ and obtain:

$$P[T \geq t + \tau \mid T \geq t] = 1 - P[X^{t+\tau} = 0 \wedge T \geq t + \tau - 1 \mid T \geq t] < 1 - (\epsilon/2M)^\tau .$$

For integral k , we thus have $P[T \geq k\tau] < \alpha^k$ where $\alpha = 1 - (\epsilon/2M)^\tau$. Using $P[T = t] \leq P[T \geq k\tau]$ for $k\tau \geq t$, we can then write $E[T] \leq \sum_{k \geq 1} k\tau \alpha^k \leq \alpha\tau/(1 - \alpha)^2 < \infty$.

We finally apply Doob's optional stopping theorem to the random variables Y^t defined as follows. Given the outcomes x^0, \dots, x^{t-1} of X^0, \dots, X^{t-1} , we set $Y^t = 0$ if $x^s = 0$ for some $s < t$ and $Y^t = X^t + \epsilon t$ otherwise. First note that Y^t is a supermartingale: Consider values y^0, \dots, y^t . If $y^s = 0$ for some $s \leq t$, we have $y^{s'} = 0$ for $s' \geq s$ and $E[Y^{t+1} \mid Y^0 = y^0 \wedge \dots \wedge Y^t = y^t] = 0 = y^t$. Otherwise, y^0, \dots, y^t are

all positive and $E[Y^{t+1} | Y^0 = y^0 \wedge \dots \wedge Y^t = y^t] = y^t + E[X^{t+1} - y^t + \varepsilon | X^0 = y^0 \wedge \dots \wedge X^t = y^t] \leq y^t$ according to the assumption on the expected variation of X^t at time t . Second, we have $E[T] < \infty$ and $|Y^{t+1} - Y^t| \leq M + \varepsilon$. Doob's theorem thus applies: Y^T is well defined and $E[Y^T] \leq E[Y^0]$. That is $E[X^T + \varepsilon T] \leq E[X^0]$. As $X^T = 0$, linearity of expectation yields $E[T] \leq E[X^0]/\varepsilon$. \square

We are now in a position to prove Theorem 1.

Proof (of Theorem 1). Set $X^t = P^t$ if there are unstable nodes and $X^t = 0$ otherwise where P^t denotes the value of the potential of the game at time t . For a given sequence of positive integers x^0, \dots, x^t , consider the event $F \subset \Omega$ where all configuration sequences in F satisfy $X^0 = x^0, \dots, X^t = x^t$. For a given configuration c that can occur at time t in F , recall that F_c^t denote the set of sequences c^0, c^1, \dots in F where $c^t = c$. Note that

$$E[X^{t+1} - X^t | X^0 = x^0 \wedge \dots \wedge X^t = x^t] = E[P^{t+1} - P^t | F] .$$

Since $F = \cup_c F_c^t$ is a disjoint union we have

$$E[P^{t+1} - P^t | F] = \frac{\sum_c E[P^{t+1} - P^t | F_c^t] P[F_c^t]}{P[F]}$$

since no considered c is an equilibrium and the dynamics is ε -improving

$$E[P^{t+1} - P^t | F] \leq \frac{\sum_c -\varepsilon P[F_c^t]}{P[F]} = -\varepsilon .$$

By Lemma 4 we get $E[T] \leq \frac{M_0}{\varepsilon}$ where T is the expected convergence time of the dynamics (by definition of random variables X^t) and $M_0 = E[P^0]$ is the expected potential at time 0. \square

A.2 Proof of Lemma 2

Proof. By definition of $CYCLE(t)$, at at time t a fraction $\alpha' = \alpha_L$ of nodes in L has color c and a fraction $\alpha_R > 1/2$ of nodes in R has color \bar{c} . The best response for nodes in L is thus \bar{c} . Therefore at time $t + 1$ the fraction α'' of nodes in L having color \bar{c} is

$$\alpha'' = 1 - \alpha' + Y/n$$

where Y is the number of nodes among the $\alpha'n$ of color c that are activated. The expectation of Y is $\mu = p\alpha'n$ and Chernoff bounds imply

$$P[Y/n \geq (1 - \delta)p\alpha'] \geq 1 - \exp\left(-\frac{\delta^2}{2}\mu\right),$$

$$P[Y/n \leq (1 + \delta)p\alpha'] \geq 1 - \exp\left(-\frac{\delta^2}{3}\mu\right).$$

We next show two implications saying that the dynamics keeps oscillating ‘around’ the value α as long as Y is close to its expectation according to $\delta = \frac{\epsilon}{1+\epsilon}$:

$$\begin{aligned}
Y/n \geq (1-\delta)p\alpha' \text{ and } \alpha' \leq (1+\epsilon)\alpha &\Rightarrow \alpha'' \geq 1 - \alpha' + p\alpha' - \delta p\alpha' \\
&= 1 - \alpha'(1-p) - \delta p\alpha' \\
&\geq 1 - (1+\epsilon)\alpha(1-p) - \delta p(1+\epsilon)\alpha \\
&= 1 - \alpha(1-p) - \epsilon\alpha(1-p) - \delta p(1+\epsilon)\alpha \\
&= \alpha(1-\epsilon) + \epsilon\alpha p - \delta p(1+\epsilon)\alpha \\
&\geq \alpha(1-\epsilon) .
\end{aligned}$$

(since $1 - \alpha(1-p) = \alpha$)
(since $\delta = \frac{\epsilon}{1+\epsilon}$)

Similarly,

$$\begin{aligned}
Y/n \leq (1+\delta)p\alpha' \text{ and } \alpha' \geq (1-\epsilon)\alpha &\Rightarrow \alpha'' \leq 1 - \alpha' + p\alpha' + \delta p\alpha' \\
&= 1 - \alpha'(1-p) + \delta p\alpha' \\
&\leq 1 - (1-\epsilon)\alpha(1-p) + \delta p(1+\epsilon)\alpha \\
&= 1 - \alpha(1-p) + \epsilon\alpha(1-p) + \delta p(1+\epsilon)\alpha \\
&= \alpha(1+\epsilon) - \epsilon\alpha p + \delta p(1+\epsilon)\alpha \\
&\leq \alpha(1+\epsilon) .
\end{aligned}$$

(using $\alpha' \leq (1+\epsilon)\alpha$)
(since $1 - \alpha(1-p) = \alpha$)
(since $\delta = \frac{\epsilon}{1+\epsilon}$)

A symmetric argument applies to L and R exchanged. We have thus shown that given $CYCLE(t)$ event $CYCLE(t+1)$ holds unless $Y > (1+\delta)\mu$ or $Y < (1-\delta)\mu$, or a symmetric situation holds with L and R exchanged. Then the union bound implies $P[-CYCLE(t+1) | CYCLE(t)] \leq 4 \exp\left(-\frac{\delta^2}{3}\mu\right)$, which proves the lemma. \square

A.3 Proof of Corollary 2

Proof. In Lemma 2 we have

$$\delta^2 \mu = \left(\frac{p}{3+p}\right)^2 p\left(1 + \frac{p}{3}\right)\alpha n = \frac{p^3 n}{(3+p)3(2-p)} > \frac{p^3 n}{24} \geq \frac{n^{1-3c}}{24}$$

and the expected number of steps to converge to an equilibrium is

$$\Omega(\exp(\delta^2 \mu/3)) = \Omega(\exp(n^{1-3c}/72)) = e^{\Omega(n^{1-3c})}.$$

This completes the proof. \square

A.4 Proof of Theorem 3

Proof. We adapt the proof of Lemma 1 and show that these dynamics are δ -improving for $\delta = p(1-2q)$. Recall that C_2 denotes the conflicting edges between two unstable nodes at time t . We thus have $d_u^t = |\{v \mid uv \in C_2\}|$. From Equation 5 in the proof of Lemma 1, we have

$$E[C^{t+1} - C^t \mid F_c^t] \leq \sum_{uv \in C_2} 2p_u^t p_v^t - \sum_{u \in U^t} p_u^t .$$

In the first sum, we can associate the weight $2p_u^t p_v^t$ to node u if $p_u^t \geq p_v^t$ and node v otherwise. We thus get

$$E[C^{t+1} - C^t \mid F_c^t] \leq \sum_{u \in U^t} p_u^t \left(-1 + \sum_{v \mid uv \in C_2, p_v^t \leq p_u^t} 2p_v^t \right) \leq \sum_{u \in U^t} p_u^t (-1 + 2p_u^t d_u^t) .$$

Using $p_u^t \in [\frac{p}{d_u^t+1}, \frac{q}{d_u^t+1}]$, we obtain

$$E[C^{t+1} - C^t \mid F_c^t] \leq -p(1-2q) \sum_{u \in U^t} \frac{1}{d_u^t+1} .$$

As $|U^t| > \max_u d_u^t$, we finally have

$$E[C^{t+1} - C^t \mid F_c^t] \leq -p(1-2q) .$$

The theorem follows from Theorem 1 and from the fact that the potential is the number of conflicting edges. \square

A.5 Proof of Theorem 4

Proof. Similarly to the proof of Theorem 3, we obtain

$$E[C^{t+1} - C^t \mid F_c^t] \leq -p(1-2q) \sum_{u \in U^t} \frac{1}{\delta_u} \leq -p(1-2q) \frac{|U^t|}{\Delta} .$$

We then conclude similarly as in the proof of Theorem 2. \square