UNIVERSITÀ DEGLI STUDI DI SALERNO
DIPARTIMENTO DI INFORMATICA "RENATO M. CAPOCELLI"

CORSO DI DOTTORATO IN INFORMATICA
X CICLO – NUOVA SERIE

ANNO ACCADEMICO 2010-2011

TESI DI DOTTORATO IN INFORMATICA

Logit Dynamics for Strategic Games

Mixing time and Metastability

Tutor
prof. **Vincenzo Auletta**

Candidato
Diodato Ferraioli

Coordinatore
prof. **Giuseppe Persiano**

Contents

Introduction

A *complex system* is generally defined as a system emerging from the interaction of several and different components, each one with their properties and their goals, usually subject to external influences. Nowadays, complex systems are ubiquitous and they are found in many research areas: examples can be found in Economy (e.g., markets), Physics (e.g., ideal gases, spin systems), Biology (e.g., evolution of life) and Computer Science (e.g., Internet and social networks). Modeling complex systems, understanding how they evolve and predicting the future status of a complex system are major research endeavors.

Historically, physicists, economists, sociologists and biologists have separately studied complex systems, developing their own tools that, however, often are not suitable for being adopted in different areas. Recently, the close relation between phenomena in different research areas has been highlighted. Hence, the aim is to have a powerful tool that is able to give us insight both about Nature and about Society, an universal language spoken both in natural and in social sciences, a modern code of nature. In a recent book [119], Tom Siegfried pointed out *game theory* as such a powerful tool, able to embrace complex systems in Economics [7, 8, 9], Biology [84], Physics [53], Computer Science [70, 76], Sociology [80] and many other disciplines.

Game theory deals with selfish agents or *players*, each with a set of possible actions or *strategies*. An agent chooses a strategy evaluating her utility or *payoff* that does not depend only on agent's own strategy, but also on the strategies played by the other players. The way players update their strategies in response to changes generated by other players defines the *dynamics* of the game and describes how the game evolves. If the game eventually reaches a fixed point, i.e., a state stable under the dynamics considered, then it is said that the game is in an *equilibrium*, through which we can make predictions about the future status of a game.

The classical game theory approach assumes that players have *complete knowledge* about the game and they are always able to select the strategy that maximizes their utility: in this *rational* setting, the evolution of a system is modeled by *best response dynamics* and predictions can be done by looking at well-known *Nash equilibrium*. Another approach is followed by *learning dynamics*: here, players are supposed to "learn" how to play in the next rounds by analyzing the history of previous plays.

By examining the features and the drawbacks of these dynamics, we can detect the basic requirements to model the evolution of complex systems and to predict their future status. Usually, in these systems, environmental factors can influence the way each agent selects her own strategy: for example, the temperature and the pressure play a fundamental role in the dynamics of particle systems, whereas the limited computational power is the main influence in computer and social settings. Moreover, as already pointed by Harsanyi and Selten [57], the complete knowledge assumption can fail due to limited information about external factors that could influence the game (e.g., if it will rain tomorrow), or about the attitude of other players (if they are risk taking), or about the amount of knowledge available to other players.

Equilibria are usually used to make predictions about the future status of a game: for this reason, we like that an equilibrium always exists and that the game converges to it. Moreover, in case that multiple equilibria exist, we like to know which equilibrium will be selected (about this important issue see [57]), otherwise we could make wrong predictions. Finally, if the dynamics takes too long time to reach its fixed point, then this equilibrium cannot be taken to describe the state of the players, unless we are willing to wait super-polynomially long transient time.

Thus we would like to have dynamics that models bounded rationality and induces an equilibrium that always exists, it is unique and is quickly reached. *Logit dynamics*, introduced by Blume [18], models a noisy-rational behavior in a clean and tractable way. In the logit dynamics for a game, at each time step, a player is randomly selected for strategy update and the update is performed with respect to an *inverse noise* parameter β (that represents the degree of rationality or knowledge) and of the state of the system, that is the strategies currently played by the players. Intuitively, a low value of β represents the situation where players choose their strategies "nearly at random" because they are subject to strong noise or they have very limited knowledge of the game; instead, an high value of β represents the situation where players "almost surely" play the best response, that is, they pick the strategies yielding high payoff with higher probability. This model is similar to the one used by physicists to describe particle systems, where the behavior of each particle is influenced by temperature: here, low temperature means high rationality and high temperature means low rationality. It is well known [18] that this dynamics defines an ergodic finite *Markov chain* over the set of strategy profiles of the game, and thus it is known that a *stationary distribution* always exists, it is unique and the chain converges to such distribution, independently of the starting profile.

Since the logit dynamics models bounded rationality in a clean and tractable way, several works have been devoted to this subject. Early works about this dynamics have focused about *long-term behavior* of the dynamics: Blume [18] showed that, for 2×2 *coordination games* and *potential games*, the long-term behavior of the system is concentrated in a specific Nash equilibrium; Alós-Ferrer and Netzer [1] gave a general characterization of long term behavior of logit dynamics for wider classes of games. A lot of works have been devoted to evaluating the time that the dynamics takes to reach specific Nash equilibria of a game, called *hitting time*: Ellison [40] considered logit dynamics for *graphical coordination games* on cliques and rings; Peyton Young [109] extended this work for more general families of graphs; Montanari and Saberi [94] gave the exact graph theoretic property of the underlying interaction network that characterizes the hitting time in graphical coordination games; Asadpour and Saberi [2] studied the hitting time for a class of *congestion games*.

Our approach is different: indeed, our first contribution is to propose the stationary distribution of the logit dynamics Markov chain as a new equilibrium concept in game theory. Our new solution concept, sometimes called *logit equilibrium*, always exists, it is unique and the game converges to it from any starting point. Instead, previous works only take in account the classical equilibrium concept of Nash equilibrium, that it is known to not satisfying all the requested properties. Moreover, the approach of previous works forces to consider only specific values of the rationality parameter, whereas we are interested to analyze the behavior of the system for each value of β.

In order to validate the logit equilibrium concept we follow two different lines of research: from one hand we evaluate the performance of a system when it reaches this equilibrium; on the other hand we look for bounds to the time that the dynamics takes to reach this equilibrium, namely the *mixing time*. This approach is trained on some simple but interesting games, such as 2×2 coordination games, congestion games and two team games (i.e., games where every player has the same utility).

Then, we give bounds to the convergence time of the logit dynamics for very interesting classes of games, such as potential games, *games with dominant strategies* and graphical coordination games. Specifically, we prove a twofold behavior of the mixing time: there are games for which it exponentially depends on β, whereas for other games there exists a function independent of β such that the mixing time is always bounded by this function. Unfortunately, we show also that there are games where the mixing time can be exponential in the number of players.

When the mixing is slow, in order to describe the future status of the system through the logit equilibrium, we need to wait a long transient phase. But in this case, it is natural to ask if we can make predictions about the future status of the game even if the equilibrium has not been reached yet. In order to answer this question we introduce the concept of *metastable distribution*, a probability distribution such that the dynamics quickly reaches it and spends a lot of time therein: we show that there are graphical coordination games where there are some distributions such that for almost every starting profile the logit dynamics rapidly converges to one of these distributions and remains close to it for an huge number of steps. In this way, even if the logit equilibrium is no longer a meaningful description of the future status of a game, the metastable distributions resort the predictive power of the logit dynamics.

Organization. In chapter 1 we introduce complex systems and game theory concepts. In particular, we survey some of the dynamics and related equilibrium concepts that were presented in literature: the analysis of these dynamics highlights the requirements desired in a dynamics. In this chapter, we also summarize the Markov chain theory and highlight the main tools that will be used in our analysis. In chapter 2 we introduce the logit dynamics, by highlighting its main properties and the reasons that make this dynamics and its stationary distribution so attractive. The chapter 3 reviews previous literature about logit dynamics. Previous works are classified in three main categories: works about the long term behavior of the logit dynamics; works about the hitting time of specific Nash equilibria in the logit dynamics; works about the mixing time of the Glauber dynamics for the Ising model. Next chapters describe our technical contribution: in chapter 4 we analyze the logit dynamics for some introductory games, by evaluating the expected social welfare at equilibrium and the mixing time of the dynamics; in chapter 5 we give mixing time bounds that hold for very large classes of games: the results about mixing time given in this two chapters differ also for the adopted techniques, *coupling* in chapter 4 and *spectral techniques* in chapter 5. Finally, in chapter 6 we arise the quest for metastability and we show our result about metastability of the logit dynamics for some introductory games. The conclusions of this work and the future directions of this line of research are presented in chapter 7.

Notations. We write $|S|$ for the size of a set S and $\overline{S}$ for its complementary set. We use bold symbols for vectors; when $\mathbf{x} = (x_1, \ldots, x_n) \in \Sigma^n$, for some alphabet Σ, we write $|\mathbf{x}|_a$ for the number of occurrences of $a \in \Sigma$ in $\mathbf{x}$; i.e., $|\mathbf{x}|_a = |\{i \in [n] : x_i = a\}|$. We use the standard game theoretic notation $(\mathbf{x}_{-i}, y)$ to mean the vector obtained from $\mathbf{x}$ by replacing the i-th entry with y; i.e., $(\mathbf{x}_{-i}, y) = (x_1, \ldots, x_{i-1}, y, x_{i+1}, \ldots, x_n)$. For two vectors $\mathbf{x}, \mathbf{y}$, we denote as $d(\mathbf{x}, \mathbf{y}) = |\{i\colon x_i \neq y_i\}|$ the Hamming distance between $\mathbf{x}$ and $\mathbf{y}$: we write $\mathbf{x} \sim \mathbf{y}$ if $d(\mathbf{x}, \mathbf{y}) = 1$. We denote by $\mathsf{poly}\,(n)$ a polynomial function in n and by $\mathsf{negl}\,(n)$ a function in n is smaller than the inverse of any polynomial in n.

Chapter 1

Preliminaries

1.1 Complex Systems

Even if nowadays the term *complex system* is broadly used in many different scientific disciplines such as anthropology, biology, chemistry, computer science, ecology, economics, meteorology, neuro-science, physics, psychology and sociology, a consensus does not exist yet about what it means.

Anyway, all the different definitions proposed agree on that an intrinsic property of a complex system is that it is a system constituted by several and different subunits or components, whose interaction gives rise to complex collective behavior. This complexity is not the effect of incomplete information about the "causes" or "inputs" by an external observer, but it is inherently rooted in the nature of the system, i.e., it is given by a non obvious relation between causes and effects, inputs and outputs [59, 100].

Two main features of a system contribute to make it complex: the emergence of traits regarding the system as a whole, that cannot be explained by properties of single components, but can evolve only from the interaction of these subunits; the presence in the story of the system, alternated with long periods of stable behavior and regular trends and periods of unpredictable and apparently random changes. This last feature is particularly annoying, because it prevents observers from predicting the future status of the system. Thus, it is of particular interest to find techniques that make possible to give meaningful previsions: in chapter 6, we will introduce a tool to face this problem.

Historically, the theory of complex systems stems from three main branches: *theoretical ecology*, that models ecology as population dynamics, whose evolution is subject to random motion [87]; *systems theory*, proposed by Von Bertalanffy [126], that refused to break a system in components and proposed to analyze systems as the whole of components and interactions; *cybernetics* [129], whose goal is to understand and define functions and processes of regulatory systems.

Complex system theory has been applied to understand natural systems, such as the atmosphere or climate, or human systems, as the immune system or the brain, and in developing "intelligent" devices in contexts like biotechnology and robotics, pattern recognition and optimization. Tools from complex system theory have been also adopted by social and economic sciences, especially in the analysis of market dynamics, management, transport or decision making.

Researchers that work on complex systems have several different goals: first of all, they would like to understand how to describe a complex system and how to model the interactions between the components; then, they would like to explain how a complex system evolves; finally, they would like to make reliable predictions about the future status of the system and about the occurrence of specific events during the evolution.

The tools adopted to achieve these goals came from many different areas: nonlinear dynamics, statistical physics, information theory, data analysis, numerical simulation are only some examples. Two of the more successful tools are *probability theory* and *game theory*. Indeed, many complex systems are characterized by high sensibility to environmental changes and by a multiplicity of aspects that make the systems intrinsically random: thus, trying to describe a system and its evolution by means of deterministic rules looks too hard, whereas the probability theory offers simpler alternatives.

On the other side, the actions of components in complex systems are not completely random, but led by some objective. Hence, modeling systems as games can be useful for understanding generic complex systems, since these models can capture otherwise intractable nonlinear effects and thereby reveal global patterns that would have been previously out of reach.

For these reasons, in this work we approach complex systems through game theory: we describe systems through games and the way components evolves through a probabilistic model. First to do it, we introduce in next sections the main concepts of game theory and of Markov chain theory.

1.2 Game theory

A game is a formal model that tries to represent how agents interact in a setting of strategic inter-dependence, i.e., when the welfare of an agent depends not only on his/her actions, but also on the actions taken by other agents. Von Neumann introduced this tool in a German paper published in 1928 [98], whose English Translation, titled "On the Theory of Games of Strategy", appear in [123]. However, only after the 1944, when Von Neumann published the book "Theory of Games and Economic Behavior" with Oskar Morgenstern [99], game theory has become a field able to attract the interest of mathematicians, economists, sociologists, psychologists, biologists, philosophers and recently also physicists and computer scientists.

Formally, a *strategic game* is a triple $([n], \mathcal{S}, \mathcal{U})$, where $[n] = \{1, \ldots, n\}$ is a finite set of *players*, $\mathcal{S} = \{S_1, \ldots, S_n\}$ is a family of non-empty finite sets (S_i is the set of *strategies* for player i), and $\mathcal{U} = \{u_1, \ldots, u_n\}$ is a family of utility functions (or *payoffs*), where $u_i \colon S_1 \times \ldots \times S_n \to \mathbb{R}$ is the utility function of player i. Players have to choose a strategy from their own strategy set in order to maximize their utility function. In the classical game theory setting, every player is supposed to have *complete knowledge* of the game and to be *selfish* and *rational*, i.e., players perfectly know the strategies and the utility functions of all players, they aim to maximize only their own utility function and they are able to make arbitrarily complex computations towards this end. In several settings complete knowledge and rationality are unrealistic since the strategy choice can be influenced by scarce knowledge of the game and by environmental limitations and a lot of works have considered models where these assumptions are weakened [115].

The computation of the strategy played by the players in a game can be seen as an evolving process in which each player reacts to the actions of other players: this process is named *dynamics* and the eventual outcome of this computation is named *equilibrium*. The dynamics of a game describes how agents change their strategy over time and how the game evolves, and it is possible to predict the future status of the game by looking at the equilibrium states induced by the dynamics. Obviously, predictions are possible only if such an induced equilibrium exists and the dynamics will converge to it.

Many dynamics and related equilibria have been proposed: in the remaining of this section we briefly describe some of these, highlighting the main features and their drawbacks. The rest of this work will focus on one of these dynamics, the logit dynamics.

1.2.1 Nash equilibrium and price of anarchy

Nash equilibrium concept is one of the most influential concepts in game theory. A (pure) *Nash equilibrium* is a strategy profile (a vector of strategies, one for every player) such that every player has no alternative strategy that increases own utility. Formally, a profile $\mathbf{x} \in S = S_1 \times \ldots \times S_n$ is a Nash equilibrium if for every player i, and every strategy $s \in S_i$ holds that

$$u_i(\mathbf{x}) \geqslant u_i(s, \mathbf{x}_{-i}) .$$

Observe that, given a profile $\mathbf{x}$ that is a Nash equilibrium, the strategy $x_i \in S_i$ maximizes $u_i(s, \mathbf{x}_{-i})$ and thus it is the *best response* that player i can give in a setting where other players are playing according to $\mathbf{x}_{-i}$. The *best response dynamics* is a dynamics that, at every time step, selects a player whose currently played strategy is not a best response and updates her strategy to a best response. It is easy to see that the Nash equilibrium is the equilibrium induced by the best response dynamics.

However, this equilibrium concept has a big issue: indeed, in many games a Nash equilibrium does not exist, thwarting the predictive power of game theory. If we focus on games where a Nash equilibrium always exists, then two classes of games are noteworthy: potential games and games with dominant strategies.

A strategic game is an exact *potential game* [93] if there exists a function $\Phi \colon S \to \mathbb{R}$ such that for every player i, every profile $\mathbf{x} \in S$, and every pair of strategies $s, z \in S_i$, it holds that

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

i.e., the increase in the utility of player i when i switches from a strategy s to a strategy z is equivalent to the decrease in the potential function between the two corresponding profiles. Since the range of Φ is finite, this function has a minimum and the profile at which the minimum is obtained is, necessarily, a Nash equilibrium. Monderer and Shapley [93] showed that the class of potential game is equivalent to the class of *congestion games* [114], where the strategies are seen as resources and the utility functions only depend on the number of players using resources. Congestion games well model a lot of problems arising in real applications, like routing problems.

A strategy $s \in S_i$ is *dominant* for player i if it yields the maximum payoff regardless of the strategies of the other players; that is, for every $z \in S_i$ and every $\mathbf{x}_{-i} \in S_1 \times \ldots \times S_{i-1} \times S_{i+1} \times \ldots \times S_n$ holds that

$$u_i(s, \mathbf{x}_{-i}) \geqslant u_i(z, \mathbf{x}_{-i}) .$$

In a *game with dominant strategies* every player has a dominant strategy. Obviously, the profile where every player plays a dominant strategy is a Nash equilibrium. This class of games has a lot of applications: specifically, it catches the interest of *Auction Theory* [69, 91, 72] and *Mechanism Design Theory* [52, 83, 101].

Even if for these classes a Nash equilibrium always exists, it might be not unique: in this case, it is not clear which equilibrium will be effectively played by players and thus which equilibrium we can predict as future status of the game. This issue is even more serious since there are a lot of games for which different equilibrium profiles have very different properties.

For class of games where the Nash equilibrium is not guaranteed to exist, we can randomize over the strategies. Specifically, a *mixed strategy* for player i is a distribution of probability μ_i over S_i; a profile $(\mu_1, \ldots, \mu_n)$ of mixed strategies is a *mixed Nash equilibrium* if for every player i and every mixed strategy ν_i over S_i holds that

$$\sum_{\mathbf{x}} u_i(\mathbf{x})\mu(\mathbf{x}) \geqslant \sum_{\mathbf{x}} u_i(\mathbf{x})\nu_i(x_i)\mu_{-i}(\mathbf{x}) ,$$

where $\mu(\mathbf{x}) = \prod_j \mu_j(x_j)$ and $\mu_{-i}(\mathbf{x}) = \prod_{j \neq i} \mu_j(x_j)$.

A celebrated result by John Nash [97] establishes that for every finite strategic game a mixed Nash equilibrium always exists. Again, the uniqueness of such an equilibrium concept is not guaranteed and thus it is unsuitable to prediction.

Moreover, recent results of Daskalakis, Goldberg and Papadimitriou [37] and Chen, Deng and Teng [32] showed that computing a Nash equilibrium is hard also for 2-player games. Specifically, the problem is complete for the complexity class **PPAD** (Polynomial Parity Argument, Directed version): this class, introduced by Papadimitriou in 1991 [106], contains problems whose decisional version is easy (indeed, it is immediate to answer the question "does a Nash equilibrium exist?"), but the functional version (to find a Nash Equilibrium) is supposed to be hard. These results cast further doubts about the usage of Nash equilibria in the analysis of complex systems.

Despite that, Nash and mixed Nash equilibria have been often adopted to describe systems and to evaluate their performance. One of the prominent measure of performance is the *Price of Anarchy* (*PoA*), introduced by Koutsoupias and Papadimitriou in 1999 [70], that assesses how much the lack of central coordination influences the performance: given an objective function W that we are interested in optimizing, the PoA is the ratio between the value of W when evaluated in the worst Nash equilibrium and the optimal value that W can assume. The Price of Anarchy has been extended and generalized in order to overcame the shortcomings of Nash equilibrium and to take in account different equilibrium concepts [55, 16, 34]: nowadays, PoA and its variants are the main tools for evaluating of the performance of any equilibrium.

1.2.2 Best response dynamics and sink equilibrium

Besides the shortcomings pointed out in the previous section, further issues are raised about Nash equilibria. The first one concerns the usage of randomization in strategies: as established above, only if players are supposed to randomize between different strategies, then the existence of an equilibrium is assured; but, the assumption that a player "sees" the mixed strategies of other players and deals with them looks very unrealistic. The second issue is raised by the relation between Nash equilibrium and best response dynamics: indeed, it is known that, even if a pure Nash equilibrium surely exists, the best response dynamics converges to it only if a specific condition is satisfied, namely the game is *weakly acyclic* [107, 90]. Moreover, computing a pure Nash equilibrium is a **PLS**-complete[1] problem [43] even for games where every player has the same strategy set: hence, there are games where the best response dynamics, in order to reach a Nash equilibrium, requires a number of iterations exponential in the number of players.

To overcome these issues, a new equilibrium concept, namely *sink equilibrium*, was proposed in 2005 by Goemans, Mirrokni and Vetta [55]. They consider the profiles of a game as vertices of a graph and they put a directed edge from $\mathbf{x}$ to $\mathbf{y}$ if $\mathbf{y}$ can be reached from $\mathbf{x}$ in one step of the best response dynamics: in this graph, any strongly connected component with no outgoing edges is a sink equilibrium. Roughly speaking, in games where the best response dynamics does not converge, this happens because there is a subset of profiles in which the dynamics sticks: by defining this subset as a sink, we assure both the existence of the equilibrium and that best response dynamics converges to it. It is obvious that any Nash equilibrium is also a sink equilibrium.

Unfortunately, sink equilibria do not solve other issues raised from Nash equilibria, as uniqueness and computational tractability. In particular, finding a sink equilibrium is **PLS**-complete [55] (and finding the best or the worst sink equilibrium is even **PSPACE**-complete [42]), whereas deciding if a profile belongs to a sink is **PSPACE**-complete [42].

[1]**PLS** (Polynomial Local Search) is the complexity class of functional problems solvable by local search. This class is strictly contained in **NP** unless "$\mathbf{NP} \neq \mathbf{co} - \mathbf{NP}$" is false [65, 89], but it is conjectured it is not contained in **P**.

1.2.3 Learning dynamics and fictious play

The best response dynamics heavily rests on the assumption of complete knowledge and perfect rationality, making the dynamics and related equilibria inopportune for modeling many settings arising from real world. Differently, *learning dynamics* assume that players "learn" how to respond to other players' strategies from what occurred in previous stages of the game. Several and different models of learning have been proposed and we survey some of those here and in the following sections (for a broad and detailed analysis about learning dynamics see the excellent books of Fudenberg and Levine [50] and Peyton Young [108]).

Fictitious play is a learning dynamics proposed by Brown in the 1951 [25][2]: the frequency of a strategy in the past iterations of the dynamics is used as a probability that this strategy will be played in the next round, and players play the best response to these probability distributions.

Specifically, let $\mathbf{x}^t$ be the profile of strategies played at time step t and let $h^t = (\mathbf{x}^1, \ldots, \mathbf{x}^t)$ be the *history* of previously played profiles. The probability $p_i^t(s)$ that player i plays strategy $s \in S_i$ at next round is the fraction of the times that player i has played this strategy in the history h^t. Fictitious play assumes that, at the next round, player i will play the strategy $s \in S_i$ that maximizes the expected utility according this empirical distribution, i.e.

$$\arg\max_{s \in S_i} \left\{ \sum_{\mathbf{x}_{-i}} u_i(s, \mathbf{x}_{-i}) p_{-i}^t(\mathbf{x}_{-i}) \right\} ,$$

where $p_{-i}^t(\mathbf{x}_{-i}) = \prod_{j \neq i} p_j^t(x_j)$.

We say that fictitious play converges if the probability distributions p_i^t converge as t increase: unfortunately, it is known that this dynamics does not converge for generic games [118] and sufficient conditions for the convergence are given only for restricted classes of 2-player games [113, 96, 92, 13]. Interestingly, Fudenberg and Kreps [47] show that, in 2-player games, if fictitious play converges, then the product distribution $p^t = \prod_j p_j^t$ is a mixed Nash equilibrium.

Recently, Brandt, Fischer and Harrenstein [24] showed another important drawback with fictitious play: this dynamics takes exponentially many rounds (in the size of the representation of the game) to converge to the equilibrium in almost every class of games where the convergence is proved.

1.2.4 No regret dynamics and correlated equilibrium

In a repeated game, at every round, every player selects a strategy: the *regret* of the sequence of strategies selected by the player i is the difference between the utility of the best fixed solution in hindsight and the average utility gained during the game. Formally, if $\mathbf{x}^t$ is the profile of strategies selected at round t, then the regret of the player i after T rounds is

$$\frac{1}{T} \left(\max_{s \in S_i} \sum_{t=1}^{T} u_i(s, \mathbf{x}_{-i}^t) - \sum_{t=1}^{T} u_i(\mathbf{x}^t) \right) .$$

We highlight that the average utility of player i is not compared with the maximum average utility obtained by the best possible sequence of strategies, but it is compared with the best average utility that i could obtain by choosing a strategy at beginning of the game and never changing her mind. In this way we are mimicking the behavior of a selfish agent in a repeated game: she can select a strategy in advance and thereafter she always plays the same strategy, or she can select an opportune and different strategy at every round. The regret is a measure

[2] Brown proposed fictitious play as an heuristic for computing Nash equilibria in some class of games and not as a learning dynamics.

of which of these two choices is better: if the regret is high then the fixed choice was preferable, but if the regret is low (goes to 0 as the time increase) or is negative, then the second choice was right.

A *regret-minimizing algorithm* is an algorithm such that the regret approaches to 0 as the time increases. By the discussion above, using such an algorithm is a reasonable behavior for selfish agents. A *no regret dynamics* is simply a dynamics where every player is assumed to run a regret-minimizing algorithm.

The first of these algorithm was proposed in 1957 by Hannan [56]. After that, a lot of really interesting regret-minimizing algorithms have been proposed, like weighted majority algorithm [79], multiplicative weights or Hedge algorithm [46], regret matching [58], polynomial weights algorithm [31].

In 1974 Aumann [6] defined the concept of correlated equilibrium as a generalization of the Nash equilibrium. A *correlated equilibrium* is a probability distribution μ over the profile space S such that for every player i and every two different strategies $s', s'' \in S_i$ holds that

$$\sum_{\mathbf{x}_{-i}} \mu(s', \mathbf{x}_{-i}) u_i(s', \mathbf{x}_{-i}) \geqslant \sum_{\mathbf{x}_{-i}} \mu(s', \mathbf{x}_{-i}) u_i(s'', \mathbf{x}_{-i}) \,.$$

In order to understand the above condition, assume there exists an external correlation device that announces to all players a probability distribution μ over the profiles of the game, then it selects a strategy profile according to μ and finally it suggests to every player the corresponding strategy. The probability distribution μ will be a correlated equilibrium if for every player, assuming that other players follow the recommendation, the expected utility gained playing the suggested strategy is higher than the expected utility obtained by playing a different strategy. Notice that correlated equilibria generalize the mixed Nash equilibrium concept: indeed, the last one is a correlated equilibrium such that the distribution μ is a product distribution of mixed strategies.

No regret dynamics are strictly related to correlated equilibria. Indeed, It is known that any no regret dynamics converges to the class of correlated equilibria [44, 58, 17]. We highlight that we are dealing with a weak notion of convergence: we only know that the dynamics converges to that class, but we do not know neither which correlated equilibrium is effectively reached nor the probability that a specific correlated equilibrium is reached. We also note that no regret dynamics converges to Nash equilibria for some remarkable class of games, namely non atomic routing games [15] and congestion games [68][3].

1.2.5 Evolutionary game theory and the replicator dynamics

Since 1973, after the seminal work of Maynard Smith and Price [85], Biology has applied game theory to explain complex systems arising in nature. This area merges game theory with the evolutionary concepts of natural selection and mutation, giving rise to a new discipline called *evolutionary game theory*. After the publication of two seminal books, "Evolution and the Theory of Games" by Maynard Smith in 1982 [84] and "The Evolution of Cooperation" by Robert Axelrod in 1984 [10], this area attracted the interest of economists and social scientists, too.

Evolutionary game theory deals with populations of agents subject to mutations and selection, as established by the Darwinian approach to evolution [36]. *Selection mechanism* can be stated in game theory terms: every individual (player) adopts an action (strategy) and has a

[3] Actually, in [68] is proved that the dynamics converges to Nash equilibria only if players run a specific (continuous time) regret-minimizing algorithm, the aggregate monotonic selection (AMS) algorithm introduced in [116, 112].

fitness landscape (the utility function); individuals with higher fitness have more possibility to reproduce and thus, in the next generation, the fraction of population that plays the action that best performed in current generation will increase. This dynamics, known as *replicator dynamics* and introduced in 1978 by Taylor and Jonker [120], is one of the main concepts in evolutionary game theory (for more details about replicator dynamics and evolutionary game theory see the excellent books of Weibull [128] and of Hofbauer and Sigmund [61]).

Formally, the replicator dynamics considers a large but finite population of homogeneous individuals (any individual has the same strategy set $S_\star$ and the same utility function u); for every strategy $s \in S_\star$ let x_s be the fraction of population using the strategy s at time t and let $\mathbf{x} = (x_s)_{s \in S_\star}$ be the profile of the population; the average payoff when the population is in the profile $\mathbf{x}$ is

$$u^\star(\mathbf{x}) = \sum_{s \in S_\star} x_s u(s, \mathbf{x}) \,.$$

Then, given the initial population profile $\mathbf{x}$, the dynamics is controlled by the continuous-time dynamic system

$$\dot{x}_s = [u(s, \mathbf{x}) - u^\star(\mathbf{x})] x_s \,.$$

Notice that, as expected, the fraction of population playing a strategy that performs worse than the average payoff will decrease, whereas the number of individuals playing an action that returns an utility higher than average payoff will increase.

The evolution eventually will reach an *attractor*, i.e., a profile $\mathbf{x}$ such that for every strategy $s \in S_\star$ such that $x_s > 0$ we have that $u(s, \mathbf{x}) = u^\star(\mathbf{x})$: thus, an attractor is stationary for the replicator dynamics and it represents the main equilibrium concept related to this dynamics. There exists a strong relation between attractors and Nash equilibria: indeed, any Nash equilibrium is an attractor and, moreover, if any strategy appears in the population, then an attractor is a Nash equilibrium. This means that the only attractors that are not Nash equilibria are the ones where there is at least a strategy never played. Unfortunately, it is also known that replicator dynamics does not converge for every game.

Notice that in the replicator dynamics if a strategy is not present in the population at the beginning, then it never appears and, similarly, any strategy present in the population at the beginning never disappears. To resolve this issue, it is possible to introduce *mutations*: that is, we assume that it is possible to randomly change the strategy played by a small fraction of population. Unfortunately, it is known that there are games where no profile is resistant against mutations, i.e., there exists no stable state such that the dynamics always come back to that state after a mutation occurs.

Recently, a dynamics similar in spirit to the replicator dynamics has been proposed by Lieberman, Hauert and Nowak [78]. This dynamics introduce social relationship in the evolutionary game theory: indeed, they consider a structured population where individuals are placed at vertices of a graph and every individual only interacts with neighbors. Moreover, a recent result by Kleinberg at al. [67] shows that there are games where the replicator dynamics, even if it does not converge, outperforms the performance of a Nash equilibrium.

However, the biggest issue with replicator dynamics is that this dynamics could be unsuitable to model complex systems that do not concern with populations of agents.

1.2.6 Behavioral game theory and EWA learning

Economic theory has often used psychology in order to explain complex and apparently unforeseeable economic and financial events: this approach dates back to famous economists like Irving Fisher and John Maynard Keynes. Thereafter, the developing of behavioral models that integrate insight from psychology and economic theory increased and led to the creation of the new areas of *behavioral economics theory* and *behavioral game theory*. The methodology adopted by these theories is experiment-driven: any proposed theory is developed from or evaluated by means of experimental observations (often involving the adoptions of the magnetic resonance imaging in order to determine which area of the brain responds to economic stimuli) and survey responses, that try to shine a light on the psychological processes beyond any strategic choice. Readers interested in behavioral game theory can refer to the book "Behavioral Game Theory" of Colin Camerer [27].

In particular, Camerer and Ho [29] proposed in 1999 a new dynamics, called *"experience-weighted attraction" (EWA) learning dynamics*, whose predictions well resemble the real behavior adopted by players in experiments (see Section 5 in [29] and Section 3 in [28]). This dynamics combines and generalizes two different learning approaches: the first one is the *fictitious play* approach, where agents form beliefs about other players and then best-respond to such beliefs; the second one is the *reinforcement learning* approach, where strategies that well behaved in previous iterations are more attractive and thus they have a greater probability to be selected in the next rounds. In order to combine these orthogonal approaches, the EWA learning dynamics introduces numerous parameters: a pair of parameters, namely $N(t)$ and ρ, take in account the experience of a player, where $N(t)$ can be seen as the number of iterations remembered by the player at time t and ρ is a factor that depreciates the impact of previous experience in favor of the most recent ones; the parameter δ measures if we are only interested in reinforce strategies that we actually played in previous iteration or to reinforce also other strategies; the parameter $A_i^s(t)$ describes how much attractive is the strategy s for the player i at time t; the parameter ϕ weights how much the attraction in previous iteration influences the attraction in the next round; finally, a parameter λ measures how much players are sensible to changes in attractions.

Obviously, the EWA learning dynamics has too many parameters and thus it is intractable for many practical purposes. For this reason, Camerer et al. [28] proposed a simpler version, called *self-tuning EWA learning dynamics*, where some parameter are replaced with fixed numerical values and some other parameters are replaced by specific functions on the history of the game. Despite this, also this version of the dynamics looks too hard to analyze in order to predict the behavior of components in a complex system.

Recently, another issue with EWA learning dynamics has been pointed out: there are circumstances where the dynamics, even for two-player games, does not converge and cycles in a subset of the profile space, whose dimension can be very high [54] (see also [117] for similar results about reinforcement learning). The high dimension of the attractor cycle implies that the behavior of the dynamics is effectively indistinguishable from random behavior: thus, adopting this dynamics to predict the evolution of a complex system can be useless.

1.2.7 Which properties we look for in a dynamics?

Our discussion about different proposed dynamics highlights the main features that we look for in a dynamics. First of all, we want a dynamics that *take into account scarce knowledge of players or environmental limitations*, since our aim is to model complex systems that, as established above, are characterized by inherent randomness. We notice that this property also explains why in real world the same person can take extremely different actions in very similar situations.

Since the main goal in the analysis of complex systems is to predict the future status, the dynamics has to enable such predictions. This obviously means that *an equilibrium has to exist*; anyway the *uniqueness of equilibrium* and the *quick convergence of the dynamics to the equilibrium* are basic to make any prediction meaningful.

A dynamics that could be adopted to analyze different complex systems arising in the reality need to be *universal*, i.e. not linked with a particular structure of the game. Moreover, our dynamics has to enable further analysis about the system: for this reason we like to deal with *analytically tractable* dynamics and not with too complex dynamics, even if the latter can be more realistic. Finally, a dynamics gets more attractive if it resembles experimental results about the behavior of components in a system.

In Section 2.3 we will show that the logit dynamics satisfies almost any of these requirements, becoming a very appealing choice for modeling the evolution of complex systems that can be described as games.

1.3 Markov chain and mixing time

In this section we survey some basic facts about Markov chains. Fore a more detailed description and for notation conventions we refer the reader to [75].

A sequence of random variables $(X_0, X_1, \ldots)$ is a *Markov chain* with *state space* Ω and *transition matrix* P if for all $x, y \in \Omega$, all $t \geqslant 1$, and all events $H_{t-1} = \bigcap_{s=0}^{t-1}\{X_s = x_s\}$ satisfying $\mathbf{P}\left(H_{t-1} \cap \{X_t = x\}\right) > 0$, we have

$$\mathbf{P}\left(X_{t+1} = y \mid H_{t-1} \cup \{X_t = x\}\right) = \mathbf{P}\left(X_{t+1} = y \mid X_t = x\right) = P(x, y)\,.$$

This condition, called *Markov property*, means that the probability to reach the state y from x is independent from the states visited before x. Notice that P is a *stochastic matrix*. By denoting with $\mathbf{P}_x(\cdot)$ and $\mathbf{E}_x[\cdot]$ the probability and the expectation conditioned on the starting state of the Markov chain being x, i.e., on the event $\{X_0 = x\}$, we have that

$$\mathbf{P}_x\left(X^t = y\right) = P^t(x, y)\,,$$

where the matrices P^t is the *t-step transition matrix*.

A Markov chain $\mathcal{M} = (\{X_i\}_{i=0}^{\infty}, \Omega, P)$ is *irreducible* if for any two states $x, y \in \Omega$ there exists an integer $t = t(x, y)$ such that $P^t(x, y) > 0$, i.e., it is possible to reach any state from any other one. The *period* of an irreducible Markov chain is the greatest common divisor of $\{t \geqslant 1 \colon P^t(x, x) > 0\}$ for some $x \in \Omega$. If the period of a Markov chain is greater than 1, then the chain is called *periodic*, otherwise *aperiodic*. If a Markov chain is finite (i.e., the space state Ω is a finite set), irreducible and aperiodic then the chain is *ergodic*: for an ergodic chain there is an integer r such that, for all $x, y \in \Omega$, $P^r(x, y) > 0$.

It is a classical result that if $\mathcal{M}$ is ergodic there exists an unique *stationary distribution*; that is, a distribution π on Ω such that

$$\pi P = \pi\,.$$

A Markov chain $\mathcal{M}$ is said *reversible* if for all $x, y \in \Omega$, it holds that

$$\pi(x)P(x, y) = \pi(y)P(y, x)\,.$$

The probability distribution $Q(x, y) = \pi(x)P(x, y)$ over $\Omega \times \Omega$ is sometimes called *edge stationary distribution*.

The *total variation distance* $\|\mu - \nu\|_{\mathrm{TV}}$ between two probability distributions μ and ν on the same state space Ω is defined as

$$\|\mu - \nu\|_{\mathrm{TV}} = \max_{A \subset \Omega} |\mu(A) - \nu(A)| = \frac{1}{2}\sum_{x \in \Omega} |\mu(x) - \nu(x)| = \sum_{\substack{x \in \Omega \colon \\ \mu(x) > \nu(x)}} \left(\mu(x) - \nu(x)\right)\,.$$

An ergodic Markov chain $\mathcal{M}$ *converges* to its stationary distribution π; specifically, there exists $0 < \alpha < 1$ such that

$$d(t) \leqslant \alpha^t ,$$

where

$$d(t) = \max_{x \in \Omega} \left\| P^t(x, \cdot) - \pi \right\|_{\mathrm{TV}} .$$

Observe that $d(t) \leqslant \bar{d}(t) \leqslant 2d(t)$, where

$$\bar{d}(t) = \max_{x,y \in \Omega} \left\| P^t(x, \cdot) - P^t(x, \cdot) \right\|_{\mathrm{TV}} .$$

For $0 < \varepsilon < 1/2$, the *mixing time* is defined as

$$t_{\mathrm{mix}}(\varepsilon) = \min\{t \in \mathbb{N} : d(t) \leqslant \varepsilon\}.$$

It is usual to set $\varepsilon = 1/4$ or $\varepsilon = 1/2e$. If not explicitly specified, when we write t_{mix} we mean $t_{\mathrm{mix}}(1/4)$. Observe that $t_{\mathrm{mix}}(\varepsilon) \leqslant \lceil \log_2 \varepsilon^{-1} \rceil t_{\mathrm{mix}}$. Next sections show the techniques for bounding mixing time that will be used in this work.

1.3.1 Coupling

A *coupling* of two probability distributions μ and ν on Ω is a pair of random variables (X, Y) defined on $\Omega \times \Omega$ such that the marginal distribution of X is μ and the marginal distribution of Y is ν. A *coupling of a Markov chain* $\mathcal{M}$ with transition matrix P is a process $(X_t, Y_t)_{t=0}^{\infty}$ with the property that both X_t and Y_t are Markov chains with transition matrix P. When the two coupled chains start at $(X_0, Y_0) = (x, y)$, we write $\mathbf{P}_{x,y}(\cdot)$ and $\mathbf{E}_{x,y}[\cdot]$ for the probability and the expectation on the space where the two coupled chains are both defined.

We denote by τ_{couple} the first time the two chains meet; that is,

$$\tau_{\mathrm{couple}} = \min\{t : X_t = Y_t\} .$$

We will consider only couplings of Markov chains with the property that for $s \geqslant \tau_{\mathrm{couple}}$, it holds $X_s = Y_s$. The following theorem establishes the importance of this tool (see, for example, Theorem 5.2 in [75]).

Theorem 1.3.1 (Coupling). *Let $\mathcal{M}$ be a Markov chain with finite state space Ω and transition matrix P. For each pair of states $x, y \in \Omega$ consider a coupling (X_t, Y_t) of $\mathcal{M}$ with starting states $X_0 = x$ and $Y_0 = y$. Then*

$$\left\| P^t(x, \cdot) - P^t(y, \cdot) \right\|_{\mathrm{TV}} \leqslant \mathbf{P}_{x,y}\left(\tau_{\mathrm{couple}} > t\right) .$$

Consider a partial order $\preceq$ over the states in Ω. A coupling of a Markov chain is said to be *monotone* w.r.t. $(\Omega, \preceq)$ if, for every $t \geqslant 0$, $X_t \preceq Y_t \Rightarrow X_{t+1} \preceq Y_{t+1}$. For a state $z \in \Omega$, the *hitting time* τ_z of z is the first time the chain is in state z, $\tau_z = \inf\{t \geqslant 0 : X_t = z\}$. Then, the following lemma relates coupling time and hitting time.

Lemma 1.3.2. *Let $\mathcal{M}$ be a Markov chain with finite state space Ω and transition matrix P. Let $\preceq$ be a partial order over Ω. For each pair of states $x, y \in \Omega$ consider a coupling (X_t, Y_t) of $\mathcal{M}$ with starting states $X_0 = x$ and $Y_0 = y$ that is monotone w.r.t. $(\Omega, \preceq)$. Moreover, suppose the ordered set $(\Omega, \preceq)$ has an unique maximum at z. Then*

$$\mathbf{P}_{x,y}\left(\tau_{\mathrm{couple}} > t\right) \leqslant 2 \cdot \max\{\mathbf{P}_x\left(\tau_z > t\right), \mathbf{P}_y\left(\tau_z > t\right)\} .$$

Sometimes it is difficult to specify a coupling and to analyze the coupling time τ_{couple} for each pair of starting states x and y. The *path coupling* theorem says that it is sufficient to define a coupling only for pairs of Markov chains starting from *adjacent* states and an upper bound on the mixing time can be obtained if each of these couplings contracts their distance on average. More precisely, consider a Markov chain $\mathcal{M}$ with state space Ω and transition matrix P; let $G = (\Omega, E)$ be a connected graph and let $w : E \to \mathbb{R}$ be a function assigning weights to the edges such that $w(e) \geqslant 1$ for every edge $e \in E$; for $x, y \in \Omega$, we denote by $\rho(x, y)$ the weight of the shortest path in G between x and y. The following theorem holds.

Theorem 1.3.3 (Path Coupling [26])**.** *Suppose that for every edge $\{x, y\} \in E$ a coupling (X_t, Y_t) of $\mathcal{M}$ with $X_0 = x$ and $Y_0 = y$ exists such that $\mathbf{E}_{x,y}[\rho(X_1, Y_1)] \leqslant e^{-\alpha} \cdot w(\{x, y\})$ for some $\alpha > 0$. Then*

$$t_{\text{mix}}(\varepsilon) \leqslant \frac{\log(\text{diam}(G)) + \log(1/\varepsilon)}{\alpha}$$

where $\text{diam}(G)$ *is the (weighted) diameter of* G*.*

1.3.2 Spectral techniques

Let P be the transition matrix of a Markov chain with finite state space Ω and let us label the eigenvalues of P in non-increasing order

$$\lambda_1 \geqslant \lambda_2 \geqslant \cdots \geqslant \lambda_{|\Omega|}.$$

It is well-known (see, for example, Lemma 12.1 in [75]) that $\lambda_1 = 1$ and, if P is irreducible and aperiodic, then $\lambda_2 < 1$ and $\lambda_{|\Omega|} > -1$. We set $\lambda^\star$ as the largest absolute value among eigenvalues other than λ_1,

$$\lambda^\star = \max_{i=2,\ldots,|\Omega|} \{|\lambda_i|\} .$$

The *relaxation time* t_{rel} of a Markov chain $\mathcal{M}$ is defined as

$$t_{\text{rel}} = \frac{1}{1-\lambda^\star} = \max\left\{\frac{1}{1-\lambda_2}, \frac{1}{1+\lambda_{|\Omega|}}\right\} .$$

The relaxation time is related to the mixing time by the following theorem (see, for example, Theorems 12.3 and 12.4 in [75]).

Theorem 1.3.4 (Relaxation time)**.** *Let P be the transition matrix of a reversible, irreducible, and aperiodic Markov chain with state space Ω and stationary distribution π. Then*

$$(t_{\text{rel}} - 1) \log\left(\frac{1}{2\varepsilon}\right) \leqslant t_{\text{mix}}(\varepsilon) \leqslant \log\left(\frac{1}{\varepsilon \pi_{\min}}\right) t_{\text{rel}} ,$$

where $\pi_{\min} = \min_{x \in \Omega} \pi(x)$*.*

The following theorem allows to relate two chains by comparing their stationary and edge stationary distributions (it is derived from Lemma 13.11 and Lemma 13.22 in [75]).

Theorem 1.3.5 (Comparison Theorem)**.** *Let P and $\hat{P}$ be the transition matrices of two reversible, irreducible, and aperiodic Markov chains with the same state space Ω, stationary distributions π and $\hat{\pi}$ respectively, and edge stationary distributions Q and $\hat{Q}$ respectively. Suppose that two constants α, γ exist such that, for all $x, y \in \Omega$,*

$$\begin{aligned} \hat{Q}(x,y) &\leqslant \alpha \cdot Q(x,y) \\ \pi(x) &\leqslant \gamma \cdot \hat{\pi}(x). \end{aligned}$$

Then

$$\frac{1}{1-\lambda_2} \leqslant \alpha\gamma \cdot \frac{1}{1-\hat{\lambda}_2},$$

where λ_2 is the second eigenvalue of P and $\hat{\lambda}_2$ is the second eigenvalue of $\hat{P}$.

Sometimes, better bounds on relaxation time can be obtained by using the following lemma (see Corollary 13.24 in [75]).

Lemma 1.3.6. *Let P the transition matrix of an irreducible, aperiodic and reversible Markov chain with state space Ω and stationary distribution π. Consider the graph $G = (\Omega, E)$, where $E = \{(x, y) \colon P(x, y) > 0\}$, and to every pair of states $x, y \in \Omega$ we assign a path $\Gamma_{x,y}$ from x to y on G. We define*

$$\rho = \max_{e=(z,w)\in E} \frac{1}{Q(e)} \sum_{\substack{x,y\colon \\ e\in\Gamma_{x,y}}} \pi(x)\pi(y)|\Gamma_{x,y}|.$$

Then $\frac{1}{1-\lambda_2} \leqslant \rho$.

Lemma 1.3.6 can be explained by means of *multi-commodity network flows*: consider the graph G as a network where every edge e has capacity $Q(e)$; every vertex x in the network (a state of the Markov chain) want to exchange a flow of size $\pi(x)\pi(y)$ with every other vertex y; a pair of vertices x and y agree on a path $\Gamma_{x,y}$ and they exchange their flow over this path; for every edge $e \in E$

$$\sum_{\substack{x,y\colon \\ e\in\Gamma_{x,y}}} \frac{\pi(x)\pi(y)}{Q(e)}$$

is now the cost over e, i.e. how much the load on e exceeds its capacity; thus, $1/(1-\lambda_2)$ is bounded by the product of the maximal length of a path and the maximum cost over an edge.

1.3.3 Bottleneck Ratio

Let $\mathcal{M} = \{X_t \colon t \in \mathbb{N}\}$ be an irreducible and aperiodic Markov chain with finite state space Ω, transition matrix P, and stationary distribution π. Let $R \subseteq \Omega$, the *bottleneck ratio* of R is

$$B(R) = \frac{Q(R, \overline{R})}{\pi(R)},$$

where $Q(R, \overline{R}) = \sum_{x\in R,\, y\in\overline{R}} Q(x, y)$. We will use the following theorem to derive lower bounds to the mixing time (see, for example, Theorem 7.3 in [75]).

Theorem 1.3.7 (Bottleneck ratio)**.** *Let $\mathcal{M} = \{X_t : t \in \mathbb{N}\}$ be an irreducible and aperiodic Markov chain with finite state space Ω, transition matrix P, and stationary distribution π. Let $R \subseteq \Omega$ be any set with $\pi(R) \leqslant 1/2$. Then the mixing time is*

$$t_{\mathrm{mix}}(\varepsilon) \geqslant \frac{1-2\varepsilon}{2B(R)}.$$

Chapter 2

Logit Dynamics

In this chapter we introduce the *logit dynamics*, a *randomized* dynamics introduced in [18], and we highlight its main properties and motivations.

The logit dynamics runs as follows: At every time step

1. Select one player $i \in [n]$ uniformly at random;
2. Update the strategy of player i according to the *Boltzmann distribution* with parameter β over the set S_i of her strategies. That is, a strategy $y \in S_i$ will be selected with probability

$$\sigma_i(y \mid \mathbf{x}) = \frac{1}{T_i(\mathbf{x})} \, e^{\beta u_i(\mathbf{x}_{-i}, y)} \,, \tag{2.1}$$

where $\mathbf{x} \in S_1 \times \cdots \times S_n$ is the strategy profile played at the current time step, $T_i(\mathbf{x}) = \sum_{z \in S_i} e^{\beta u_i(\mathbf{x}_{-i}, z)}$ is the normalizing factor, and $\beta \geqslant 0$.

We can see parameter β as the *rationality level* of the system: indeed, it is easy to see that for $\beta = 0$ player i selects her strategy uniformly at random, for $\beta > 0$ the probability is biased toward strategies promising higher payoffs, and for $\beta \to \infty$ player i chooses her best response strategy (if more than one best response is available, she chooses uniformly at random one of them). Moreover observe that probability $\sigma_i(y \mid \mathbf{x})$ does not depend on the strategy x_i currently adopted by player i.

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

Definition 2.0.8 (Logit dynamics [18])**.** *Let $\mathcal{G} = ([n], \mathcal{S}, \mathcal{U})$ be a strategic game and let $\beta \geqslant 0$. The* logit dynamics *for $\mathcal{G}$ is the Markov chain $\mathcal{M}_\beta = (\{X_t\}_{t \in \mathbb{N}}, S, P)$ where $S = S_1 \times \cdots \times S_n$ and*

$$P(\mathbf{x}, \mathbf{y}) = \frac{1}{n} \cdot \begin{cases} \sigma_i(y_i \mid \mathbf{x}), & \text{if } \mathbf{y}_{-i} = \mathbf{x}_{-i} \text{ and } y_i \neq x_i; \\ \sum_{i=1}^{n} \sigma_i(y_i \mid \mathbf{x}), & \text{if } \mathbf{y} = \mathbf{x}; \\ 0, & \text{otherwise;} \end{cases} \tag{2.2}$$

where $\sigma_i(y_i \mid \mathbf{x})$ is defined in (2.1).

As will be showed in Section 2.1.1, the Markov chain defined in (2.2) 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.

For the class of *potential games* the stationary distribution is the well known *Gibbs measure*.

Theorem 2.0.9 ([18]). *If $\mathcal{G} = ([n], \mathcal{S}, \mathcal{U})$ is a potential game with potential function Φ, then the Markov chain given by (2.2) is reversible with respect to its stationary distribution, that is the Gibbs measure*

$$\pi(\mathbf{x}) = \frac{1}{Z} e^{-\beta\Phi(\mathbf{x})}, \tag{2.3}$$

where $Z = \sum_{\mathbf{y}\in S} e^{-\beta\Phi(\mathbf{y})}$ is the normalizing constant.

2.1 Properties

In this section we give some useful properties of the logit dynamics.

2.1.1 Ergodicity

It is easy to see that the Markov chain defined by the logit dynamics is ergodic. Indeed, let $\mathbf{x} = (x_1, \dots, x_n)$ and $\mathbf{y} = (y_1, \dots, y_n)$ be two profiles and let $(\mathbf{z}^0, \dots, \mathbf{z}^n)$ be a *path* of profiles where $\mathbf{z}^0 = \mathbf{x}$, $\mathbf{z}^n = \mathbf{y}$ and $\mathbf{z}^i = (y_1, \dots, y_i, x_{i+1}, \dots x_n)$ for $i = 1, \dots, n-1$. The probability that the chain starting at $\mathbf{x}$ is in $\mathbf{y}$ after n steps is

$$P^n(\mathbf{x}, \mathbf{y}) = P^n(\mathbf{z}^0, \mathbf{z}^n) \geqslant P^{n-1}(\mathbf{z}^0, \mathbf{z}^{n-1}) P(\mathbf{z}^{n-1}, \mathbf{z}^n),$$

and recursively

$$P^n(\mathbf{x}, \mathbf{y}) \geqslant \prod_{i=1}^{n} P(\mathbf{z}^{i-1}, \mathbf{z}^i) > 0,$$

where the last inequality follows from (2.2), since for all $i = 1, \dots, n$, the Hamming distance between $\mathbf{z}^{i-1}$ and $\mathbf{z}^i$ is at most 1.

2.1.2 Invariance under utility translation

Let $\mathcal{G} = ([n], \mathcal{S}, \mathcal{U})$ be a game and let $\tilde{\mathcal{G}} = ([n], \mathcal{S}, \tilde{\mathcal{U}})$ be a new game obtained from $\mathcal{G}$ by substituting the utility functions with a new family $\tilde{\mathcal{U}} = \{\tilde{u}_i : i \in [n]\}$ of utility functions as follows

$$\tilde{u}_i(\mathbf{x}) := u_i(\mathbf{x}) + c_i \qquad \text{for all } \mathbf{x} \text{ and for all } i.$$

We observe that $\mathcal{G}$ and $\tilde{\mathcal{G}}$ have the same logit dynamics. Indeed, according to (2.1), when the game is at profile $\mathbf{x}$ player i chooses strategy y with probability

$$\begin{aligned} \tilde{\sigma}_i(y \mid \mathbf{x}) &= \frac{e^{\beta\tilde{u}_i(\mathbf{x}_{-i}, y)}}{\sum_{z\in S_i} e^{\beta\tilde{u}_i(\mathbf{x}_{-i}, z)}} = \frac{1}{\sum_{z\in S_i} e^{\beta[\tilde{u}_i(\mathbf{x}_{-i}, z) - \tilde{u}_i(\mathbf{x}_{-i}, y)]}} \\ &= \frac{1}{\sum_{z\in S_i} e^{\beta[u_i(\mathbf{x}_{-i}, z) - u_i(\mathbf{x}_{-i}, y)]}} = \sigma_i(y \mid \mathbf{x}). \end{aligned}$$

2.1.3 Noise changes under utility rescaling

While translations of utilities do not affect logit dynamics, a *rescaling* of the utility functions for a constant $\alpha > 0$ changes the inverse noise from β to $\alpha \cdot \beta$. Indeed, if for every player i and every profile $\mathbf{x}$ we set

$$\tilde{u}_i(\mathbf{x}) := \alpha \cdot u_i(\mathbf{x}),$$

then, from (2.1), we have

$$\tilde{\sigma}_i(y \mid \mathbf{x}) = \frac{e^{\beta\tilde{u}_i(\mathbf{x}_{-i}, y)}}{\sum_{z\in S_i} e^{\beta\tilde{u}_i(\mathbf{x}_{-i}, z)}} = \frac{e^{\alpha\beta u_i(\mathbf{x}_{-i}, y)}}{\sum_{z\in S_i} e^{\alpha\beta u_i(\mathbf{x}_{-i}, z)}}.$$

Notice that, unlike the previous property that holds even if for each player i we add a different constant c_i to the utility functions, here we must have the same rescaling constant α for all utility functions.

2.1.4 Logit dynamics vs. Glauber dynamics

In a potential game, the logit dynamics is equivalent to the well-studied *Glauber dynamics*. In fact, let $S = S_1 \times \cdots \times S_n$ be a state space and μ be a probability distribution over S, then the Glauber dynamics for μ proceeds as follows: from profile $\mathbf{x} \in S$, pick a player $i \in [n]$ uniformly at random and update her strategy at $y \in S_i$ with probability μ conditioned on the other players being at $\mathbf{x}_{-i}$, i.e.,

$$\mu(y \mid \mathbf{x}_{-i}) = \frac{\mu(\mathbf{x}_{-i}, y)}{\sum_{z \in S_i} \mu(\mathbf{x}_{-i}, z)} .$$

It is easy to see that the Markov chain defined by the Glauber dynamics is irreducible, aperiodic, and reversible with stationary distribution μ. When $\mathcal{G} = ([n], \mathcal{S}, \mathcal{U})$ is a potential game with potential function Φ, the logit dynamics defines the same Markov chain as the Glauber dynamics for the Gibbs distribution π in (2.3). Indeed, in this case we have

$$\begin{aligned} \sigma_i(y \mid \mathbf{x}) &= \frac{e^{\beta u_i(\mathbf{x}_{-i}, y)}}{\sum_{z \in S_i} e^{\beta u_i(\mathbf{x}_{-i}, z)}} = \frac{1}{\sum_{z \in S_i} e^{\beta(u_i(\mathbf{x}_{-i}, z) - u_i(\mathbf{x}_{-i}, y))}} \\ &= \frac{1}{\sum_{z \in S_i} e^{\beta(\Phi(\mathbf{x}_{-i}, y) - \Phi(\mathbf{x}_{-i}, z))}} = \frac{e^{-\beta \Phi(\mathbf{x}_{-i}, y)}}{\sum_{z \in S_i} e^{-\beta \Phi(\mathbf{x}_{-i}, z)}} = \frac{\pi(\mathbf{x}_{-i}, y)}{\sum_{z \in S_i} \pi(\mathbf{x}_{-i}, z)} . \end{aligned}$$

Hence, logit dynamics for potential games and Glauber dynamics for Gibbs distributions are two ways of looking at the same Markov chain: in the former case the dynamics is derived from the potential function, in the latter case from the stationary distribution. However, observe that in general the Glauber dynamics for the stationary distribution of the logit dynamics is different from the logit dynamics (see, for example, the Matching Pennies case in Subsection 2.2.1).

Due to the analogies between logit and Glauber dynamics, we will sometimes adopt the terminology used by physicists to indicate the quantities involved (see for example Section 3.2); in particular we will denote parameter β as inverse noise or *inverse temperature* and the normalizing constant Z of the Gibbs distribution (2.3) as *partition function*.

2.2 Some simple examples

Let us consider now the logit dynamics to a pair of very simple games. Our examples will highlight some interesting features of the dynamics and its related equilibrium concept.

2.2.1 Matching Pennies

Consider the classical *Matching Pennies* game:

	H	T
H	$+1, -1$	$-1, +1$
T	$-1, +1$	$+1, -1$

(2.4)

According to (2.1), the update probabilities for the logit dynamics are, for every $x \in \{H, T\}$

$$\begin{aligned} \sigma_1(H \mid (x, H)) = \sigma_1(T \mid (x, T)) = \frac{1}{1+e^{-2\beta}} = \sigma_2(T \mid (H, x)) = \sigma_2(H \mid (T, x)) , \\ \sigma_1(T \mid (x, H)) = \sigma_1(H \mid (x, T)) = \frac{1}{1+e^{2\beta}} = \sigma_2(H \mid (H, x)) = \sigma_2(T \mid (T, x)) . \end{aligned}$$

Hence the transition matrix (see (2.2)) is

$$
P = \left(\begin{array}{c|cccc} & HH & HT & TH & TT \\ \hline HH & 1/2 & b/2 & (1-b)/2 & 0 \\ HT & (1-b)/2 & 1/2 & 0 & b/2 \\ TH & b/2 & 0 & 1/2 & (1-b)/2 \\ TT & 0 & (1-b)/2 & b/2 & 1/2 \end{array} \right) \tag{2.5}
$$

where, for readability sake, we named $b = \frac{1}{1+e^{-2\beta}}$.

You can see that each player has a positive probability to play every strategy in her strategy set: however, the probability that the selected player plays the best response increases, as the rationality level β increases.

Since every column of the matrix adds up to 1, the matrix is doubly stochastic and the uniform distribution π over the set of strategy profiles is the stationary distribution for the logit dynamics.

We observe that the Glauber dynamics for π generates a transition matrix different from P: this proves that the perfect match between logit dynamics and Glauber dynamics holds only for potential games and not for every game.

2.2.2 A *stairs* game

Let $\mathcal{G}$ be a potential game where every player has two strategies, say *upstairs* (or 1) and *downstairs* (or 0). Define the potential of a profile $\mathbf{x} \in \{0,1\}^n$ as the number of players that are upstairs, i.e. $\Phi(\mathbf{x}) = -|\mathbf{x}|_1$.

Then, the partition function is

$$
Z(\beta) = \sum_{\mathbf{x} \in \{0,1\}^n} e^{-\beta|\mathbf{x}|_1} = \sum_{k=0}^{n} \binom{n}{k} e^{-\beta k} = \left(1 + e^{-\beta}\right)^n ,
$$

and the stationary distribution is

$$
\pi(\mathbf{x}) = \frac{e^{-\beta|\mathbf{x}|_1}}{(1+e^{-\beta})^n} .
$$

Observe that the probability that the player selected for the update plays strategy 1 (or equivalently strategy 0) is independent of the strategies played by other players. Indeed, according to (2.1), for every $\mathbf{x}$ it holds that

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

This example shows that for potential games it is very easy to describe the stationary distribution and the update probabilities even for n-player games: this is one of the main reasons for which almost every game that we consider in this work will be a potential game.

2.3 Motivations

In the original work by Blume [18] that introduced logit dynamics, it is highlighted that this dynamics handles two key features of strategic behavior: *lock-in* and *bounded rationality.*

The lock-in property establishes that, once a player makes a choice, she is committed to it for some while: assuming this property holds in reality (and this is admissible if we consider decision costs or costs for strategy update) justifies a strategy selection rule that takes into account only the strategies actually played by other players and not the strategies they previously played. This is exactly what the logit dynamics update rule (2.1) does.

The bounded rationality property appears in the logit dynamics in two different aspects: the *myopic* behavior, i.e., players contemplate only the present reward and not the expected stream of rewards, and the *limited information* available to players. As suggested by Harsanyi and Selten [57], the last aspect can be handled by considering a probabilistic model that involves suitably chosen random moves. It is evident that the logit dynamics update rule (2.1) embodies both these aspects of bounded rationality.

The probabilistic approach of the logit dynamics is also motivated by the goal to model the evolution of complex systems, that are intrinsically random and thus they can be described only by means of probabilistic models.

Bounded rationality is just one of the properties that, as pointed out in Section 1.2.7, make a dynamics attractive for analyzing complex systems. We will show that logit dynamics enjoys many other ones.

Since the logit dynamics defines a Markov chain, the evolution of a system is modeled in a clean and *tractable* way, by allowing deep analysis by means of tools from Markov chain theory. Moreover, no restriction is given about the structure of the game or the utility functions: thus, this dynamics can be *universally* applied to every system we are interested.

The natural equilibrium concept of the logit dynamics is the stationary distribution of the Markov chain, already known as *logit equilibrium* [88]. Since the Markov chain described by the logit dynamics is ergodic, we have that this equilibrium *always exists* and is *unique*.

Fudenberg and Levine [49] suggested *universal consistency* as a desiderata for any learning algorithm. Specifically, a learning rule ρ_i for player i is ε-universally consistent if, regardless of the play of the other players, almost surely holds that

$$\lim_{T\to\infty} \sup \max_{x_i \in S_i} u_i(x_i, \gamma_i^T) - \frac{1}{T}\sum_t u_i(\rho_i(h_{t-1})) \leqslant \varepsilon\,, \tag{2.6}$$

where h_{t-1} is the history of play in the previous round of the game and γ_i^T is the empirical frequency of strategies by other players. Roughly speaking, a learning rule is universally consistent if a player, using the learning algorithm, takes at least as much utility as she could have gained had she known the frequency but not the order of observations in advance. The same authors in [51] showed that the logit dynamics update rule in 2.1 is an universally consistent learning rule.

Thus, logit dynamics satisfies a lot of interesting properties, that motivates its adoption as a model of evolution of games. However, these properties do not explain the reasons beyond the introduction of this dynamics. Therefore, it is natural to raise the questions: Why does the parameter β in (2.1) represent the rationality level? And why does a player with rationality β play according to (2.1)? In the remaining of this section we will give two different answers to these questions: the first one based on the random utility model in Economics, the second one on information theory concepts.

2.3.1 An economic interpretation of the update rule

The *random utility model* [14] has been broadly adopted in economic theory, and in particular in the area of rational stochastic choice theory. This model claims that the difference between observed utilities and real utilities is given by stochastic error terms concentrated around 0. The random utility model strongly derives from the pioneering work of Thurstone in *psychometrics* and especially in the fields of measurement [121] and factor analysis [122].

One of the prominent random utility models is the *log-linear model*, that considers the error terms as independent random variables identically distributed according to a standard *log-Weibull distribution* (known also as *Gumbel distribution*): this distribution is essentially the same as the normal distribution (except in the tail), but it will result analytically more tractable.

Based on the log-linear utility model, we can set for every player i, for every strategy $s \in S_i$ and for every strategy profile $\mathbf{x}_{-i} \in S_1 \times \ldots \times S_{i-1} \times S_{i+1} \times \ldots \times S_n$,

$$U_i(s, \mathbf{x}_{-i}) = \beta u_i(s, \mathbf{x}_{-i}) + \varepsilon_i(s, \mathbf{x}_{-i}) \,, \tag{2.7}$$

where $U_i(s, \mathbf{x}_{-i})$ is the real utility of player i, $u_i(s, \mathbf{x}_{-i})$ is the guessed utility and $\varepsilon_i(s, \mathbf{x}_{-i})$ the error term. Here, β represents how much we are confident that the guess is right (except for a small error term). Alternatively, we can think $u_i(s, \mathbf{x}_{-i})$ as the utility computed by the player i and β as a measure of "how much the computation is right": thus, a large value of β, by neutralizing any error term, means we are assuming player i is able to do any computation in order to maximize utility, so she is *fully rational*; conversely, small β means we are assuming player i has *bounded rationality*, since the utility that she is able to compute is only a far approximation of the real utility and, thus, computed preferences could not match real preferences.

Assuming the utilities are as described in (2.7), we can motivate the adoption of the update rule (2.1) by following the approach of McFadden [86] about random utility maximization. Since a selfish player wants to maximize the real utility, the probability that player i plays the strategy $s \in S_i$, given that other players are in the profile $\mathbf{x}_{-i}$, is

$$\begin{aligned}
\sigma_i(s \mid \mathbf{x}_{-i}) &= \mathbf{P}\left(U_i(s, \mathbf{x}_{-i}) \geqslant U_i(z, \mathbf{x}_{-i}),\ \text{for every } z \neq s\right) \\
&= \mathbf{P}\left(\varepsilon_i(z, \mathbf{x}_{-i}) \leqslant \beta(u_i(s, \mathbf{x}_{-i}) - u_i(z, \mathbf{x}_{-i})) + \varepsilon_i(s, \mathbf{x}_{-i}),\ \text{for every } z \neq s\right) .
\end{aligned}$$

For sake of readability, we will use Δ_{sz} as a shorthand for $\beta(u_i(s, \mathbf{x}_{-i}) - u_i(z, \mathbf{x}_{-i}))$. Since error terms are independent and they take values in $(-\infty, +\infty)$, we have

$$\begin{aligned}
\sigma_i(s \mid \mathbf{x}_{-i}) &= \int_{-\infty}^{+\infty} \mathbf{P}\left(\varepsilon_i(z, \mathbf{x}_{-i}) \leqslant \Delta_{sz} + \varepsilon,\ \text{for every } z \neq s\right) \mathbf{P}\left(\varepsilon_i(s, \mathbf{x}_{-i}) = \varepsilon\right) d\varepsilon \\
&= \int_{-\infty}^{+\infty} \mathbf{P}\left(\varepsilon_i(s, \mathbf{x}_{-i}) = \varepsilon\right) \prod_{z \neq s} \mathbf{P}\left(\varepsilon_i(z, \mathbf{x}_{-i}) \leqslant \Delta_{sz} + \varepsilon\right) d\varepsilon \,.
\end{aligned}$$

Since error terms are log-Weibull distributed, we have that

$$\mathbf{P}\left(\varepsilon_i(s, \mathbf{x}_{-i}) = \varepsilon\right) = e^{-\varepsilon} e^{-e^{-\varepsilon}} = e^{-\varepsilon} e^{-e^{-(\Delta_{ss} + \varepsilon)}} ,$$

and for every $z \neq s$

$$\mathbf{P}\left(\varepsilon_i(z, \mathbf{x}_{-i}) \leqslant \Delta_{sz} + \varepsilon\right) = e^{-e^{-(\Delta_{sz} + \varepsilon)}} .$$

Hence,

$$\sigma_i(s \mid \mathbf{x}_{-i}) = \int_{-\infty}^{+\infty} e^{-\varepsilon} \prod_{z \in S_i} e^{-e^{-(\Delta_{sz} + \varepsilon)}} d\varepsilon = \int_{-\infty}^{+\infty} e^{-\varepsilon} e^{-e^{-\varepsilon} \sum_{z \in S_i} e^{-\Delta_{sz}}} d\varepsilon \,.$$

By setting $t = e^{-\varepsilon}$, we will have

$$\sigma_i(s \mid \mathbf{x}_{-i}) = \int_{0}^{+\infty} e^{-t \sum_{z \in S_i} e^{-\Delta_{sz}}} dt = \frac{1}{\sum_{z \in S_i} e^{-\Delta_{sz}}} = \frac{e^{\beta u_i(s, \mathbf{x}_{-i})}}{\sum_{z \in S_i} e^{\beta u_i(z, \mathbf{x}_{-i})}} .$$

Thus, the update rule (2.1) is exactly the probability that player i adopts the strategy s when her computed utility is $u_i(\cdot)$ and the level of exactness of the computation, i.e., the *rationality level*, is β.

2.3.2 An information theoretical interpretation of the update rule

Wolpert defined in [130] a measure of rationality based on *information theory* concepts and then he showed that, given that a player has a level of rationality β, the best prediction about the behavior of this player is given by the Boltzmann distribution in (2.1).

Before to show these results, we have to review some useful concepts from information theory and optimization theory.

Entropy, Maxent Principle, and Kullback-Leibler distance

In the following we summarize some basic concepts of information theory useful for understanding the remaining of the section: for a more detailed description we refer the reader to [35].

Entropy. In information theory, the (Shannon) *entropy* H_b of a probability distribution σ over a space S is

$$H_b(\sigma) = \sum_{s \in S} \sigma(s) \log_b \frac{1}{\sigma(s)}\,.$$

Usually $b = 2$, but we will use $b = e$.

The entropy $H_b(\sigma)$ is a measure of the *uncertainty* associated with a random variable X that takes values in S with probability distribution σ. Or, alternatively, the entropy is a measure of the *amount of information* in a probability distribution σ: indeed, the amount of information is equivalent to "how much uncertainty the distribution clears up" with respect to some prior knowledge μ (if we have no prior knowledge then μ is the distribution with maximal uncertainty, i.e., the uniform distribution) and thus it can be measured as $H_b(\mu) - H_b(\sigma)$.

More precisely, the entropy $H_b(\sigma)$ is the expected number of b-ary *information units* that we need to represent a random variable X that takes values in S with probability distribution σ. If $b = 2$, we are considering as information unit the *binary unit* (bit), whereas if $b = e$, we are considering *natural units* (nats).

The *Shannon source coding theorem* establishes that, if we want to encode a random variable X that takes values in S with probability distribution σ with an alphabet of size b, the expected word-length is at least $H_b(\sigma)$.

Maxent Principle. How we can estimate a distribution σ, given some incomplete prior information about it?

The *maximum entropy (maxent) principle* [64] suggests that the best estimation we can do is the one that contains the minimal amount of extra information beyond the prior knowledge. Or equivalently, the best approximation for σ is the distribution that maximizes the entropy between all distributions that agree with the prior knowledge.

Kullback-Leibler distance. Given two probability distributions σ_1 and σ_2, how they are distant with respect to the information provided?

Let us consider a random variable X_2 that assumes values in S with probability distribution σ_2: above discussion about entropy and Shannon source coding theorem shows that the optimal encoding of X_2 assigns to every symbol $s \in S$ a string of at least $\log_b \sigma_2(s)^{-1}$ b-ary information units and the expected word-length of X is at least $H_b(\sigma_2)$.

Now, consider the random variable X_1 with distribution σ_1 on the same space S: it is evident that the expected word-length of X_1 when we encode it by using the optimal encoding for X_2

is at least

$$H_b(\sigma_1, \sigma_2) = \sum_{s \in S} \sigma_1(s) \log_b \frac{1}{\sigma_2(s)} ,$$

that is called *cross entropy*.

Intuitively, if σ_1 and σ_2 convey similar information we expect that the optimal code for X_2 is good enough also for encoding X_1. For this reason, the *Kullback-Leibler distance*,

$$KL_b(\sigma_1 \| \sigma_2) = H_b(\sigma_1, \sigma_2) - H_b(\sigma_1) = \sum_{s \in S} \sigma_1(s) \log_b \frac{\sigma_1(s)}{\sigma_2(s)} , \tag{2.8}$$

is used to measure the gap between the information provided by σ_1 and by σ_2. We notice that $KL_b(\sigma_1 \| \sigma_2)$ is not a metric and in particular it is not symmetric. However, it is known that $KL_b(\sigma_1 \| \sigma_2)$ is always non-negative and it is zero if and only if the distributions are identical (this result is known as *Gibbs' inequality* [81]).

The method of Lagrange multipliers

Consider a combinatorial optimization problem

$$\begin{aligned} &\text{maximize} \quad f(x_1, \dots, x_n) \\ &\text{subject to} \quad g(x_1, \dots, x_n) = c\,. \end{aligned} \tag{2.9}$$

The *method of Lagrange multipliers* is a technique to solve such problems. It introduces a new variable $\lambda \geqslant 0$ (the *Lagrange multiplier*) and considers the following function, called *Lagrangian*,

$$L(x_1, \dots x_n, \lambda) = f(x_1, \dots, x_n) - \lambda \left[c - g(x_1, \dots, x_n) \right] .$$

We can solve Problem (2.9), by maximizing the Lagrangian $L(x_1, \dots x_n, \lambda)$. In particular, a solution to Problem (2.9) as function of the Lagrange multiplier λ, is found by solving the following system of differential equations:

$$\left\{ \frac{\partial L(x_1, \dots x_n, \lambda)}{\partial x_i} = 0 \qquad \text{for } 1 \leqslant i \leqslant n\,. \right.$$

An exact solution to Problem (2.9) follows if we add the equation $c - g(x_1, \dots, x_n) = 0$ to the system.

For a more detailed description of optimization theory concepts we refer the reader to [23].

A measure of rationality

We look for an operator R that, given an objective function to maximize, like $u_i(\cdot \mid \mathbf{x}_{-i})$, and a probability distribution σ_i on S_i, returns a measure of how much rational σ_i is with respect to the objective. In this paragraph, we motivate the adoption of the measure

$$R_{u_i(\cdot|\mathbf{x}_{-i})}(\sigma_i) = \arg\min_{\beta} KL_e \left(\sigma_i \| \mu^{\beta}_{u_i(\cdot|\mathbf{x}_{-i})} \right) , \tag{2.10}$$

where $\mu^{\beta}_{u_i(\cdot|\mathbf{x}_{-i})}$ is the Boltzmann distribution given in (2.1). In the following we will use $R(\sigma_i)$ and μ^{β} when the reference to $u_i(\cdot \mid \mathbf{x}_{-i})$ is clear from the context. Moreover, from now on, whenever we refer to entropy and Kullback-Leibler distance, we assume $b = e$, even if the subscript is omitted.

A selfish and rational player i plays the (mixed) strategy that obtains the maximum utility $u_i^\star$ given $\mathbf{x}_{-i}$: this is our prior knowledge about the behavior of player i. The maxent principle suggests that the best estimation about the strategy played by i will be the probability distribution σ_i that

$$\begin{aligned} &\text{maximize} \quad H(\sigma_i) \\ &\text{subject to} \quad \mathbf{E}_{\sigma_i}\left[u_i(\cdot, \mathbf{x}_{-i})\right] = \sum_{s \in S_i} \sigma_i(s) u_i(s, \mathbf{x}_{-i}) = u_i^\star \,. \end{aligned} \tag{2.11}$$

This combinatorial optimization problem can be solved via the method of Lagrange multipliers: the *maxent Lagrangian* is

$$M(\sigma_i, \beta) = H(\sigma_i) - \beta\left[u_i^\star - \mathbf{E}_{\sigma_i}\left[u_i(\cdot, \mathbf{x}_{-i})\right]\right] \,. \tag{2.12}$$

Before solving the optimization problem, let us discuss the role of β (the Lagrange multiplier) in (2.12). If β is small, then $M(\sigma_i, \beta)$ can be very large even if the expected utility with respect to σ_i is far from the maximum utility: thus, small β means we are assuming player i could have a *bounded rational* behavior. On the other hand, if β is large, then $M(\sigma_i, \beta)$ can be large only if the expected utility in according to σ_i is close to the maximum utility: thus, large β means we are considering a *fully rational* player. Thus, the Lagrange multiplier β in the maxent Lagrangian perfectly specifies the balance between the rational and irrational behavior of a player.

Now let look to the solution of the maxent Lagrangian.

Lemma 2.3.1. *The solution to the Problem (2.11) as function of the Lagrange multiplier is given by the Boltzmann distribution in (2.1).*

Proof. Assume $S_i = \{s_1, \ldots, s_{|S_i|}\}$. For $j = 1, \ldots, |S_i|$, we have to satisfy the equations

$$\frac{\partial M(\sigma_i)}{\partial \sigma_i(s_j)} = \beta u_i(s_j, \mathbf{x}_{-i}) - \ln \sigma_i(s_j) - 1 = 0 \,.$$

Hence follows that $\sigma_i(s_j) = \frac{1}{e} e^{\beta u_i(s_i, \mathbf{x}_{-i})}$ and, by normalizing, we complete the proof. □

Thus, the Boltzmann distribution with parameter β is the best approximation of the strategy that will be played by a player i, and β measures how much the player is rational.

However, we do not have yet a rationality operator. Obviously, if the distribution σ_i that we want to test is distributed as a Boltzmann distribution with parameter λ, then the above discussion suggests that we can consider such λ as the measure of the rationality of σ_i. What about σ_i that is non-Boltzmann? Notice that if $KL\left(\sigma_i \| \mu^\lambda\right)$ is small, then the information provided by σ_i is more on less the same as the information convoyed by the Boltzmann distribution of parameter λ, or, equivalently, the latter is a good approximation of σ_i: thus λ will be a good approximation of the rationality of σ_i.

In this way we have motivated the *rationality operator* R given in (2.10). We observe that R is always non negative. However, if $u_i(\cdot, \mathbf{x}_{-i})$ is constant, then R is not really meaningful (this depend on the fact that every σ_i obtains the same expected utility, thus every probability distribution is equivalently rational). For this reason in the next paragraph, we will consider only non-constant utility functions.

Approximating the behaviour of a bounded-rational player

Once we have an operator that measures the rationality of a choice, we can reverse the question: if we know that the player i has rationality $\beta^\star$, which is the best estimation that we can do about the behavior of player i?

In order to answer this question, we use again the maxent principle: given our prior knowledge about rationality of player i, the best estimation is the probability distribution σ_i that

$$\begin{array}{ll} \text{maximize} & H(\sigma_i) \\ \text{subject to} & R(\sigma_i) = \beta^\star \, . \end{array} \tag{2.13}$$

The following lemma shows an alternative way to express the constraint of Problem (2.13).

Lemma 2.3.2. *$R(\sigma_i) = \beta^\star$ if and only if $\mathbf{E}_{\sigma_i}[u_i(\cdot, \mathbf{x}_{-i})] = \mathbf{E}_{\mu^{\beta^\star}}[u_i(\cdot, \mathbf{x}_{-i})]$.*

Proof. The proof proceeds in three steps.
Step 1: We observe that $\mathbf{E}_{\mu^\lambda}[u_i(\cdot, \mathbf{x}_{-i})]$ is a strictly increasing function in the parameter λ. Indeed,

$$\begin{aligned} \frac{\partial \mathbf{E}_{\mu^\lambda}[u_i(\cdot, \mathbf{x}_{-i})]}{\partial \lambda} &= \sum_{s \in S_i} \mu^\lambda(s) u_i(s, \mathbf{x}_{-i})^2 - \left(\sum_{s \in S_i} \mu^\lambda(s) u_i(s, \mathbf{x}_{-i}) \right)^2 \\ &= Var_{\mu^\lambda}[u_i(\cdot, \mathbf{x}_{-i})] > 0 \, , \end{aligned}$$

where the last inequality is strict since we are considering non-constant utility functions. As a consequence, we obtain that $R(\sigma_i) = \beta^\star$ if and only if $\mathbf{E}_{\mu^{R(\sigma_i)}}[u_i(\cdot, \mathbf{x}_{-i})] = \mathbf{E}_{\mu^{\beta^\star}}[u_i(\cdot, \mathbf{x}_{-i})]$.
Step 2: We observe that, for every $\lambda \geqslant 0$,

$$\frac{\partial \ln \sum_{s \in S_i} e^{\lambda u_i(s, \mathbf{x}_{-i})}}{\partial \lambda} = \sum_{s \in S_i} u_i(s, \mathbf{x}_{-i}) \frac{e^{\lambda u_i(s, \mathbf{x}_{-i})}}{\sum_{s \in S_i} e^{\lambda u_i(s, \mathbf{x}_{-i})}} = \mathbf{E}_{\mu^\lambda}[u_i(\cdot, \mathbf{x}_{-i})] \, .$$

Step 3: From the definition of rationality measure in (2.10) and the definition of Kullback-Leibler distance in (2.8), we have

$$R(\sigma_i) = \arg\min_\lambda \left[\ln \sum_{s \in S_i} e^{\lambda u_i(s, \mathbf{x}_{-i})} - \lambda \mathbf{E}_{\sigma_i}[u_i(\cdot, \mathbf{x}_{-i})] \right] .$$

This is equivalent to say that

$$\begin{aligned} 0 &= \frac{\partial \left(\ln \sum_{s \in S_i} e^{R(\sigma_i) u_i(s, \mathbf{x}_{-i})} - R(\sigma_i) \mathbf{E}_{\sigma_i}[u_i(\cdot, \mathbf{x}_{-i})] \right)}{\partial R(\sigma_i)} \\ &= \frac{\partial \ln \sum_{s \in S_i} e^{R(\sigma_i) u_i(s, \mathbf{x}_{-i})}}{\partial R(\sigma_i)} - \mathbf{E}_{\sigma_i}[u_i(\cdot, \mathbf{x}_{-i})] \, . \end{aligned}$$

□

Now, setting $u^\star = \mathbf{E}_{\mu^{\beta^\star}}[u_i(\cdot, \mathbf{x}_{-i})]$, the combinatorial problem in (2.13) is exactly the same as Problem (2.11). Thus, by Lemma 2.3.1, follows that the solution for Problem (2.13) is given by the Boltzmann distribution in (2.1). Finally, since we are solving the optimization problem with the method of Lagrange multipliers, the exact solution is found by imposing the condition

$$\mathbf{E}_{\mu^\beta}[u_i(\cdot, \mathbf{x}_{-i})] = \mathbf{E}_{\mu^{\beta^\star}}[u_i(\cdot, \mathbf{x}_{-i})] \, ,$$

that is, $\beta = \beta^\star$. Thus we have obtained that, given that player i has rationality level $\beta^\star$, the best estimation about the behavior of i is that she plays according to the Boltzmann distribution in (2.1) with parameter $\beta^\star$.

2.4 Relations with other dynamics

In this section we show relations between the logit dynamics and some of the dynamics discussed in Section 1.2.

2.4.1 Best response dynamics

The logit dynamics can be viewed as a *randomized* version of the best response dynamics. Actually, as already underlined in [18] and in [88], our dynamics is only an instance of the class of *noisy best response dynamics* [48, 66]: in these dynamics, player i, if selected for the update when other players are playing $\mathbf{x}_{-i}$, will play a best response with a probability $\alpha_i(\mathbf{x}_{-i})$ and a non-best response with probability $1 - \alpha_i(\mathbf{x}_{-i})$.

It is easy to see that, as the rationality level β in (2.1) goes to infinity, thus the logit dynamics approaches to the best response dynamics and, if the game is a potential game, the logit equilibrium approaches to a Nash equilibrium (or a distribution over Nash equilibria). Unfortunately, this means that the logit dynamics inherits some of the drawbacks of the best response dynamics: in particular, as established in Section 1.2.2, there are games where, when β goes to infinity, the logit dynamics takes an exponential (in the number of players) number of steps to reach the logit equilibrium.

One of the main goals of this work, that will be carried out in next chapters, is to evaluate the convergence time of the logit dynamics to the stationary distribution for each value of β.

2.4.2 No regret dynamics

We notice that the definition of universal consistency in (2.6) is very similar in spirit to (but not the same as) the definition of regret minimizing algorithm. Moreover, the logit equilibrium, the equilibrium concept related to the logit dynamics, is a probability distribution over the profile space, like the correlated equilibrium, that, as pointed in Section 1.2.4, is the equilibrium concept related to the no regret dynamics. Given such similarities, it is natural to ask if logit dynamics is a no regret dynamics and in particular if the stationary distribution is a correlated equilibrium. We can give a negative answer to these questions. Consider, indeed, the following game:

	0	1
0	$1,2$	$1,2$
1	$0,1$	$0,1$

It is easy to see that the strategy 0 is dominant for the row player: thus any correlated equilibrium has to assign zero probability to profiles 10 and 11. It is also easy to see that the game is a potential game with potential function

$$\Phi(00) = \Phi(01) = 1 \qquad \text{and} \qquad \Phi(10) = \Phi(11) = 0\,.$$

Thus the logit equilibrium is given by the Gibbs distribution in (2.3) and we can check that, for every finite β, the probability assigned to profiles 10 and 11 is greater than zero.

2.5 Experimental results

Several experiments that compare predictions of logit dynamics and real data have been presented in literature. In particular, in this section we will report the ones presented by McKelvey and Palfrey [88] and by Camerer et al. [28]: here, we are only concerned in evaluating and discussing those results. For more details we refer to the original papers.

McKelvey and Palfrey [88] focus on experiments involving two-person games with a unique Nash equilibrium where there are not outcomes Pareto preferred to the Nash equilibrium. Real data are collected from different experiments run in more than 30 years: a 3-strategy zero sum game repeated 200 times [77], a 4-strategy zero sum game repeated 150 times [105], a 5-strategy zero sum game repeated 120 times [111], a 2-strategy non zero sum game repeated 640 times [103]. In [88], it is calculated, for each experiment, a maximum likelihood estimate of the parameter β and it is analyzed how well the model fits the data. Results show that the logit dynamics predicts systematic deviations from Nash equilibrium. Nevertheless, authors also notice that there are aspects of real data that remain unexplained by the logit dynamics, such as if there exists some consistency in the rationality parameter β across experiments.

Camerer et al. [28] consider seven games: two matrix games with unique mixed Nash equilibrium [95]; a patent race game [110]; a median-action order statistic coordination game [125]; a continental-divide coordination game [124]; a dominance solvable beauty contest [60]; a price matching game [30]. In [28], parameter β is estimated by using the 70% of the subjects and then those estimations are used to predict choices by the remaining 30%. In another test, parameters are estimated on six of the seven games and then such parameters are used to predict choices in the remaining seventh game. From these tests, it is possible to see that there are games where the predictions of the logit dynamics are comparable to the ones of more powerful dynamics like EWA learning dynamics and its self tuning version, introduced in Section 1.2.6: we remember that these dynamics are analytically intractable and they are not supposed to converge in any game. However, there are class of games where logit dynamics fits poorly real data: it will be interesting to understand why this happens and if a characterization exists for games that are experimentally well described by the logit dynamics.

Chapter 3

State of Art

In this chapter we survey previous results about logit dynamics. Moreover, we introduce the Ising model, used by physicists to modeling spin systems, and we show some results about the mixing time of Glauber dynamics for the Ising model. We start our discussion by giving evidence of the relations between our research and the literature on this subject.

3.1 Logit dynamics

Research on the logit dynamics mainly focused on two lines: results that try to characterize the dynamics as β goes to infinity, that we summarize in Section 3.1.1, and results that estimate the hitting time of specific profiles when β is large, that we discuss in Section 3.1.2.

3.1.1 Equilibrium selection

The *equilibrium selection* problem was introduced by Harsanyi and Selten in 1988: in their excellent book [57] they pointed out that the Nash equilibrium concept has several weaknesses, in particular the presence of multiple Nash equilibria. Specifically, they noticed that almost every nontrivial game has many (and sometimes infinitely many) different equilibrium points, and this is a limit to the predictive power of game theory, since it is impossible to establish which equilibrium will be selected by players. To overcame these weaknesses, they proposed several new solution concepts and provided a mathematical criterion in order to select one equilibrium point as the solution of the game.

The criterion adopted by Harsanyi and Selten involves the introduction of vanishing noise in the game that allow to go out from "weak" equilibria, but not from the "strongest" equilibrium. Since logit dynamics introduces noise too, a lot of interest has been devoted to understand which equilibrium will be selected by the dynamics when the perturbation vanishes.

The first class of games that has been considered is the class of 2×2 *coordination games*. These are games in which players have an advantage in selecting the same strategy. These games are often used to model the spread of a new technology [109]: two players have to decide whether to adopt or not a new technology and it is assumed that the players would prefer choosing the same technology as the other one.

The game is formally described by the following payoff matrix

$$\begin{array}{c|c|c|}
 \multicolumn{1}{c}{} & \multicolumn{1}{c}{0} & \multicolumn{1}{c}{1} \\ \cline{2-3}
0 & a, a & c, d \\ \cline{2-3}
1 & d, c & b, b \\ \cline{2-3}
\end{array} \tag{3.1}$$

We assume that $a > d$ and $b > c$ which implies that players have an advantage in selecting the same strategy of their opponents. This game has two pure Nash equilibria: $(0, 0)$, where each player has utility a, and $(1,1)$, where each player has utility b. If $a - d > b - c$ we say that $(0,0)$ is the *risk dominant equilibrium*; if $a - d < b - c$ we say that $(1,1)$ is the risk dominant equilibrium; otherwise we say that no risk dominant equilibrium exists.

The risk dominant equilibrium concept is one of the refinements of Nash equilibrium proposed by Harsanyi and Selten [57], and it is the equilibrium that is more likely to be played by players that do not know the actions taken by other players. In the spread of technology example, we can assume that every player, if she does not known anything about the other player action, prefers to choose the new technology: with this assumption the outcome where both select the new technology will be the risk dominant equilibrium.

Blume [18] proved the following result about logit dynamics for 2×2 coordination games.

Theorem 3.1.1 ([18])**.** *Let $\mathcal{G}$ be a 2×2 coordination game such that a risk dominant equilibrium $\mathbf{x}$ exists and let P be the transition matrix of the logit dynamics for $\mathcal{G}$. Then, for every starting profile $\mathbf{y}$, we have that*

$$\lim_{\beta\to\infty} \lim_{t\to\infty} P^t(\mathbf{y},\mathbf{x}) = 1\,.$$

That is, the logit dynamics will select the risk dominant equilibrium as the noise vanishes.

A similar result holds even if both players have more than 2 strategies.

A *graphical coordination games* is a game in which n players are connected by a network $G = (V,E)$ (encoding, for example, social relationships) and every player plays the basic coordination game (3.1) with each of the adjacent players. Specifically, when a player selects her strategy, such a strategy is played against each one of her adjacent players. The payoff of a player is given by the sum of the payoffs gained from each instance of the basic coordination game. It is easy to see that the profiles where all players play the same strategy are Nash equilibria: moreover, if $(0,0)$ is risk dominant for the basic coordination game, then profile $\mathbf{0} = (0,\ldots,0)$ will be the risk dominant profile for the graphical coordination game. Ellison [40] gave a result similar to Theorem 3.1.1 for graphical coordination games on two network topologies: the *clique*, where every player is connected to every other players, and the *ring*, where every player has exactly two neighbors.

Theorem 3.1.2 ([40])**.** *Let $\mathcal{G}$ be a graphical coordination game on a graph G such that a risk dominant equilibrium $\mathbf{x}$ exists and let P be the transition matrix of the logit dynamics for $\mathcal{G}$. If G is a clique or a ring, then, for every starting profile $\mathbf{y}$, we have that*

$$\lim_{\beta\to\infty} \lim_{t\to\infty} P^t(\mathbf{y},\mathbf{x}) = 1\,.$$

That is, the logit dynamics will select the risk dominant equilibrium as the noise vanishes.

A strategy $s \in S_i$ is said *dominated* if, for every profile $\mathbf{x}_{-i}$, there exists a strategy $z \in S_i$, with $z \neq s$, such that

$$u_i(z,\mathbf{x}_{-i}) > u_i(s,\mathbf{x}_{-i})\,.$$

Consider a game $\mathcal{G}_0$: starting from $i = 1$, we obtain the game $\mathcal{G}_i$ by deleting all dominated strategies in $\mathcal{G}_{i-1}$, until it is possible. The set of *iteratively dominated strategies* of $\mathcal{G}_0$ contains all the strategies eliminated during this process.

Blume [18] showed that profiles involving iteratively dominated strategies cannot be selected by the logit dynamics as the noise vanishes.

Theorem 3.1.3 ([18])**.** *Let $\mathcal{G}$ be a game, S be the set of strategy profiles of $\mathcal{G}$ and $V \subseteq S$ be the set of profiles that involves iteratively dominated strategies. Let P be the transition matrix of the logit dynamics for $\mathcal{G}$. Then, for every starting profile $\mathbf{x}$, we have that*

$$\lim_{\beta\to\infty} \lim_{t\to\infty} P^t(\mathbf{x}, S\setminus V) = 1\,.$$

Logit dynamics for potential games attracted a lot of attention, because of peculiarity of this class of games (see Section 2.1.4). In particular, as a corollary of Theorem 2.0.9, we have the following result.

Corollary 3.1.4 ([18])**.** *Let $\mathcal{G}$ be a potential game such that an unique potential minimizer $\mathbf{x}$ exists and let P be the transition matrix of the logit dynamics for $\mathcal{G}$. Then, for every starting profile $\mathbf{y}$, we have that*

$$\lim_{\beta \to \infty} \lim_{t \to \infty} P^t(\mathbf{y}, \mathbf{x}) = 1\,.$$

That is, the logit dynamics will select the potential minimizer as the noise vanishes.

Several generalizations of the class of potential games have been presented in literature, such as weighted potential games, ordinal potential games, best response potential games and generalized potential games. A game is a *weighted potential games* [93] if there exists a function $\Phi\colon S \to \mathbb{R}$ and weights $w_1, \ldots, w_n$ such that for every player i, every profile $\mathbf{x} \in S$, and every pair of strategies $s, z \in S_i$, it holds that

$$u_i(s, \mathbf{x}_{-i}) - u_i(z, \mathbf{x}_{-i}) = w_i \left(\Phi(z, \mathbf{x}_{-i}) - \Phi(s, \mathbf{x}_{-i}) \right)\,.$$

A game is an *ordinal potential game* [93] if there exists a function $\Phi\colon S \to \mathbb{R}$ such that for every player i, every profile $\mathbf{x} \in S$, and every pair of strategies $s, z \in S_i$, it holds that $u_i(s, \mathbf{x}_{-i}) - u_i(z, \mathbf{x}_{-i})$ and $\Phi(z, \mathbf{x}_{-i}) - \Phi(s, \mathbf{x}_{-i})$ have the same sign. A game is a *best-response potential game* [127] if there exists a function $\Phi\colon S \to \mathbb{R}$ such that for every player i, every profile $\mathbf{x} \in S$, it holds that

$$\arg\max_{s \in S_i} u_i(s, \mathbf{x}_{-i}) = \arg\min_{s \in S_i} \Phi(s, \mathbf{x}_{-i})\,.$$

Last, a game is a *generalized ordinal potential game* [93] if there exists a function $\Phi\colon S \to \mathbb{R}$ such that for every player i, every profile $\mathbf{x} \in S$, and every pair of strategies $s, z \in S_i$, we have that $u_i(s, \mathbf{x}_{-i}) - u_i(z, \mathbf{x}_{-i}) > 0$ implies $\Phi(z, \mathbf{x}_{-i}) - \Phi(s, \mathbf{x}_{-i}) > 0$. It is easy to see that every potential game is a weighted potential game and every weighted potential game is an ordinal potential game. Moreover, the classes of best-response potential games and generalized ordinal potential games are both generalizations of the class of ordinal potential game: however, Voorneveld [127] showed that these two classes are distinct.

Unfortunately, Alós-Ferrer and Netzer [1] showed recently that Corollary 3.1.4 has not an equivalent for generalizations of potential games. Specifically, they prove the following result.

Theorem 3.1.5 ([1])**.** *Let $\mathcal{G}$ be a best-response potential game and let P be the transition matrix of the logit dynamics for $\mathcal{G}$. For every starting profile $\mathbf{y}$, let $V_{\mathbf{y}}$ be the set of profiles $\mathbf{x}$ such that*

$$\lim_{\beta \to \infty} \lim_{t \to \infty} P^t(\mathbf{y}, \mathbf{x}) > 0\,.$$

Then $V_{\mathbf{y}}$ is a subset of the set of Nash equilibria of $\mathcal{G}$.

Alternatively, we can say that, if the game is a best-response potential game, then the logit dynamics selects a local minimum of the potential function as the noise vanishes. However, there are games where, differently from exact potential games, the global minimum is not selected and this is showed by the following example.

$$\begin{array}{c|c|c|} & 0 & 1 \\ \hline 0 & 6,2 & 0,0 \\ \hline 1 & 0,0 & 3,3 \\ \hline \end{array} \tag{3.2}$$

This game is a coordination game between two different players: indeed, $(0,0)$ and $(1,1)$ are the Nash equilibria. We can check that the game is a weighted potential game with potential function

$$\Phi(0,0) = -6 \qquad \Phi(0,1) = -\frac{12}{5} \qquad \Phi(1,0) = 0 \qquad \Phi(1,1) = -\frac{27}{5}$$

and weights $w_1 = 1$ and $w_2 = 5/9$. Thus the profile $(0,0)$ is the potential minimizer, whereas Alós-Ferrer and Netzer [1] showed that the logit dynamics will select $(1,1)$ as noise vanishes (intuitively, their result follows from the fact that it is easier that a player selects the wrong strategy when the loss in utility is only 2, than when the loss is 3).

Alós-Ferrer and Netzer [1] also prove that Theorem 3.1.5 is tight, that is they prove that for the class of generalized ordinal potential games the logit dynamics does not necessarily select a Nash equilibrium. This is showed by the following game:

	0	1
0	$0,0$	$0,1$
1	$0,1$	$1,0$

(3.3)

The game is a generalized ordinal potential game with potential function

$$\Phi(0,0) = 0 \qquad \Phi(0,1) = -1 \qquad \Phi(1,0) = -3 \qquad \Phi(1,1) = -2\,.$$

The unique Nash equilibrium is $(1,0)$. Anyway, it is easy to see that, even when the noise vanishes, any profile has a positive probability to be reached. In particular, there is a non-zero probability that the game will be in a profile where the column player selects the strategy 1, whereas in every mixed Nash equilibria this strategy has zero probability.

3.1.2 The hitting time of the Nash equilibrium

Since we know that there are games such that for high values of β the logit dynamics converges to a specific profile, a natural question is: how much time does the system take to reach such a profile? This question was raised for the first time by Ellison [40] about graphical coordination games. In particular, he proved that when players are on a clique, the hitting time of the risk dominant equilibrium is exponential in the number of players[1]. Ellison [40] also gave evidence that the topology of the graph influences the time the dynamics takes to reach specific profiles, and proved that the rate of convergence in the clique is slower than the rate of convergence in the ring.

Peyton Young [109] has continued the work of Ellison, by proving a sufficient condition for fast convergence to the risk dominant strategy in graphical coordination game when β is large. Consider a undirected connected graph $G = (V,E)$ and two non-empty subsets S', S'' of V. Let $|E(S', S'')|$ be the number of edges with an endpoint in S' and the other one in S'' and $|E(S')|$ be the number of edges with at least one endpoint in S'. A subset of vertices S' is said *r-close knit*, for $0 \leqslant r \leqslant 1$, if

$$\min_{S'' \subseteq S'} \frac{|E(S', S'')|}{|E(S')|} = r\,.$$

That is, every member of S' has at least a fraction r of its neighbors in S''. A graph $G = (V,E)$ is *(r,k)-close knit* if for every vertices $v \in V$, there exists $S' \subset V$ such that $v \in S'$, $|S'| \leqslant k$ and S' is r-close knit. A family of graphs is *close knit* if for every $0 \leqslant r \leqslant 1/2$ there exists an integer $k(r)$ such that every graph in the family is $(r, k(r))$-close knit. Examples of close-knit families of graphs are the class of all polygons and the family of square lattices. The following theorem holds.

[1] In [40] a different dynamics is considered, but it is easy to see that the proof holds also for the logit dynamics with slightly modifications.

Theorem 3.1.6 ([109]). *Let $\mathcal{G}$ be an n-player graphical coordination game on a graph G from a close knit family of graphs. Assume that a risk dominant equilibrium $\mathbf{x}$ exists. Then, given a small $\delta > 0$, there exist β_δ and $\tau = \tau(\beta, \delta, k(\cdot))$, such that for every $\beta > \beta_\delta$ and every starting profile $\mathbf{y}$, the probability that after τ steps of the logit dynamics for $\mathcal{G}$ a fraction $1-\delta$ of players are playing the risk dominant strategy is at least $1-\delta$.*

Roughly speaking, the above theorem says that for every starting profile, the logit dynamics gets close to the risk dominant equilibrium before time τ (large but independent of n) with high probability if β is sufficiently large.

Montanari and Saberi [94] extended results about hitting time of risk dominant equilibrium in graphical coordination games, by deriving the graph theoretical quantities that characterize it. Consider a graph $G = (V, E)$ such that $|V| = n$. Let $\mathcal{L}$ be the set of all *linear orderings* $L = \{i(1), i(2), \dots, i(n)\}$ of vertices: for a linear ordering $L \in \mathcal{L}$ and an integer $1 \leqslant t \leqslant n$, we define $L_t = \{i(1), \dots, i(t)\}$. Given the basic coordination game (3.1), we set $h = \frac{a-d-b+c}{a-d+b-c}$ and, for every player i, $h_i = h|N(i)|$, where $N(i)$ is the number of neighbors of player i in G. The *tilted cutwidth* of a graphical coordination game on a graph G is

$$\Gamma(G, \mathbf{h}) = \min_{L \in \mathcal{L}} \max_{1 \leqslant t \leqslant n} \left(|E(L_t, V \setminus L_t)| - \sum_{i \in L_t} h_i \right) ,$$

where $\mathbf{h} = (h_1, \dots, h_n)$. Roughly speaking, this quantity corresponds to the maximum increase of the potential function along the *lowest* sequence of profiles from an equilibrium profile to the risk dominant profile, that is the sequence that minimizes this increase.

Let R be a collection of subsets of V such that $\emptyset \in R$ and $V \notin R$: such collection is said *monotone* if whenever $A \in R$, then $A' \in R$, for every $A' \subseteq A$. We denote with $\mathcal{R}$ the set of all monotone collections of subsets of V. Moreover, let ∂R be the set of pairs $(A, A \cup \{i\})$ such that $A \in R$ and $A \cup \{i\} \notin R$. The *tilted cut* of a graphical coordination game on a graph G is

$$\Delta(G, \mathbf{h}) = \max_{R \in \mathcal{R}} \min_{(A_1, A_2) \in \partial R} \max_{i=1,2} \left(|E(A_i, V \setminus A_i)| - \sum_{i \in A_i} h_i \right) .$$

The tilted cut can be seen as a dual quantity of the tilted cutwidth: indeed, it corresponds to the lowest value of the potential function between the highest pair of neighboring profiles.

For an induced subgraph $F \subseteq G$, we set $h_i^F = h_i + |N(i)|_{G \setminus F}$, where $|N(i)|_{G \setminus F}$ is the degree of i in $G \setminus F$, and $\mathbf{h}^F = \left(h_1^F, \dots, h_n^F\right)$.

Last, for a profile $\mathbf{x} \in S$, we define the *typical hitting time*

$$T_{\mathbf{x}} = \sup_{\mathbf{y} \in S} \inf \left\{ t \geqslant 0 \colon \mathbf{P}_{\mathbf{y}} \left(\tau_{\mathbf{x}} \geqslant t \right) \leqslant e^{-1} \right\} ,$$

where $\tau_{\mathbf{x}}$ is the hitting time as defined in Section 1.3. Then, the following theorem holds.

Theorem 3.1.7 ([94]). *Let $\mathcal{G}$ be an n-player graphical coordination game on a graph G such that a risk dominant equilibrium $\mathbf{x}$ exists. For the logit dynamics, we have that the typical hitting time of $\mathbf{x}$ is*

$$T_{\mathbf{x}} = \exp\{2\beta\Gamma_\star(G, \mathbf{h}) + o(\beta)\} ,$$

where $\Gamma_\star(G, \mathbf{h}) = \max_{F \subseteq G} \Gamma(F, \mathbf{h}^F) = \max_{F \subseteq G} \Delta(F, \mathbf{h}^F)$.

In their work, Montanari and Saberi [94] bound $\Gamma_\star$ for several families of graphs, like random graphs and small-world networks. It is important to note that the term $o(\beta)$ in the above theorem can hide n-dependent factors.

The hitting time of specific profiles has been analyzed also for games other than graphical coordination games. In particular, Asadpour and Saberi [2] consider two subclasses of potential games, namely unweighted routing games and load balancing games, and they evaluate the typical hitting time of the set A_ε of profiles $\mathbf{x}$ such that $\Phi(\mathbf{x}) \leqslant (1+\varepsilon)\Phi_{\min}$, where $\Phi_{\min}$ is the minimum value that the potential function assumes.

In an *unweighted routing game*, we have a directed graph $G = (V, E)$ and n pairs of vertices $(s_1, t_1), \ldots, (s_n, t_n)$, where every s_i is called *source* and every t_i is called *sink*. Each player corresponds to a different pair. Every edge has a nonnegative continuous nondecreasing *cost function* $c_e \colon \mathbb{R}^+ \to \mathbb{R}^+$. Player i has to select a path from s_i to t_i. When every player has selected a path, we denote as f_e the number of players that go through the edge e: the cost that player i pays is the sum of the costs $c_e(f_e)$ of the edges in the path that i selected. Then, the following theorem holds.

Theorem 3.1.8 ([2]). *Let $\mathcal{G}$ an n-players unweighted routing game such that the graph G has at most K vertices and M edges and the cost functions are polynomials of degree at most d. Then, for every constant $\varepsilon > 0$ there exists a value $\beta_0 = \beta_0(M, K, d, \varepsilon)$ such that, for any $\beta \geqslant \beta_0$, the typical hitting time of A_ε is at most Kn^3. Moreover, the Markov chain will almost always be in A_ε after hitting it.*

In a *load balancing game*, there are n jobs of integer positive weights $w_1, \ldots, w_n$, each one controlled by a different player, and m identical machines. Every job has to select a machine, where it will be run. The load l_j of the machine j is the sum of the weights of the jobs assigned to j. Player i incurs in a cost equivalent to the load of the machine where job i has been assigned. We denote with l_{avg} the average load of the machines, i.e. $l_{\text{avg}} = \sum_j l_j/m$, and with $w_{\max}$ the maximum weights among jobs. Then, the following theorem holds.

Theorem 3.1.9 ([2]). *Let $\mathcal{G}$ a load balancing game with n jobs and m machines. Then, for every constant $\varepsilon > 0$ there exists a value $\beta_0 = \beta_0(w_{\max}, \varepsilon)$ such that, for any $\beta \geqslant \beta_0$, the typical hitting time of A_ε is $\mathcal{O}\left(l_{\text{avg}}^2 nm^3\right)$. Moreover, the Markov chain will almost always be in A_ε after hitting it.*

We emphasize that both in Theorem 3.1.8 and in Theorem 3.1.9 the total size of the resources, namely $K, M, l_{\text{avg}}, w_{\max}$ and m, is constant in the number of players. Indeed, Asadpour and Saberi [2] show that there are games where a convergence polynomial in the size of the resources is possible only for small values of β, but for such values of β the dynamics does not remain close to the potential minimizer for long time.

3.2 Glauber dynamics for Ising model

The *Ising model* is a mathematical model used in Statistical Physics for ferro-magnetism: it represent a set of magnets, each having one of the two possible orientations, or *spins*, positive or negative. Magnets can influence each other: in order to represent such interactions we can assume that n magnets are placed on the vertices of a graph $G = (V, E)$ with $|V| = n$ and an edge between two magnets means that they influence each other. We denote with $x_v \in \{+1, -1\}$ the spin of the magnet at vertex $v \in V$, with $\mathbf{x} = (x_v)_{v \in V}$ a configuration of magnets and with S the set of all possible configurations. Moreover, we define the *energy* of a configuration

$$\Phi(\mathbf{x}) = - \sum_{\substack{u,v \in V\colon \\ (u,v) \in E}} x_u x_v \,.$$

Thus, the energy decreases as the number of pairs of magnets whose spins agree decreases. The strength of the interaction between magnets depends on the *temperature*: if the temperature is low, then interaction is strong and magnets' behavior tends to minimize the energy of the system; on the contrary, when the temperature is low, the interaction between magnets is weak and the system could not minimize the energy. This behavior allows to describe the system through the *Gibbs measure*, defined by

$$\pi(\mathbf{x}) = \frac{e^{-\beta\Phi(\mathbf{x})}}{Z(\beta)} \,,$$

where β can be interpreted as the inverse of the temperature and $Z(\beta)$, called *partition function*, is the normalizing constant required to make π a probability distribution, i.e.,

$$Z(\beta) = \sum_{\mathbf{x}\in S} e^{-\beta\Phi(\mathbf{x})} \,.$$

The evolution of a system towards this probability distribution is usually modeled by the *Glauber dynamics* for the Gibbs measure: a configuration $\mathbf{x}$ is updated by taking a vertex $v \in V$ uniformly at random and updating the spin of the magnet placed at v according the Gibbs distribution with the condition that the resulting profile accords with $\mathbf{x}$ everywhere except at vertex v. That is, the transition matrix on S is given by

$$P(\mathbf{x},\mathbf{y}) = \frac{1}{n} \cdot \begin{cases} \sigma_v(y_v \,|\, \mathbf{x}), & \text{if } \mathbf{y}_{-v} = \mathbf{x}_{-v} \text{ and } y_v \neq x_v; \\ \sum_{v\in V} \sigma_v(y_v \,|\, \mathbf{x}), & \text{if } \mathbf{y} = \mathbf{x}; \\ 0, & \text{otherwise;} \end{cases}$$

where

$$\sigma_v(y_v \,|\, \mathbf{x}) = \frac{e^{\beta u_v(y_v,\mathbf{x}_{-v})}}{e^{\beta u_v(y_v,\mathbf{x}_{-v})} + e^{\beta u_v(\mathbf{x})}} \,,$$

and, for every profile $\mathbf{x}$, we have

$$u_v(\mathbf{x}) = x_v \sum_{\substack{u\in V:\\ (v,u)\in E}} x_u \,.$$

For a more detailed description of the Ising model and its dynamics we refer the reader to [82]. Here, we highlight that the Glauber dynamics for the Gibbs measure in the Ising model corresponds to the logit dynamics for a specific graphical coordination games, that we call *Ising game*. Here, the magnets are the agents that play with each neighbor the following basic coordination game.[2]

	$+1$	-1
$+1$	$+1,+1$	$-1,-1$
-1	$-1,-1$	$+1,+1$

Moreover, we can check that $u_v(\cdot)$ is exactly the utility function of the player placed on the vertex v and $\Phi(\cdot)$ is the potential function of the game: then, the correspondence between the two dynamics follows from the discussion in Section 2.1.4.

We introduced this game because, differently from other graphical coordination games, the *mixing time* of the logit dynamics for this game has been bounded for several of different topologies of the underlying graph. A lot of work has been done about the Ising game on the *square lattice* and we refer the interested reader to [82] for a survey of the major results for this setting.

Recently, several and different results have been found for general graphs. Specifically, it has been proved that the mixing time is fast when β is small (see Theorem 15.1 in [75]).

[2] Note that in this coordination game no risk dominant strategy exists.

Theorem 3.2.1. *Consider the Ising game on a graph with n vertices and maximal degree d. Let $c(\beta) = 1 - d\frac{e^{\beta}-e^{-\beta}}{e^{\beta}+e^{-\beta}}$. If $c(\beta) > 0$, then the mixing time of the logit dynamics for this game is*

$$t_{\text{mix}}(\varepsilon) \leqslant \left\lceil \frac{n(\log n - \log \varepsilon)}{c(\beta)} \right\rceil . \tag{3.4}$$

In particular, (3.4) holds whenever $\beta < \frac{1}{d}$.

Berger et al. [12] gave a bound to the mixing time that depends on a structural property of the graph G. Specifically, consider the bijective function $L \colon V \to \{1, \ldots, |V|\}$: it represents an ordering of vertices of G. Let $\mathcal{L}$ be the set of all orderings of vertices of G and set $V_i^L = \{v \in V \colon L(v) < i\}$. Then, the *cutwidth* of G is

$$\chi(G) = \min_{L \in \mathcal{L}} \max_{1 < i \leqslant |V|} |E(V_i^L, V \setminus V_i^L)| . \tag{3.5}$$

The following theorem holds.

Theorem 3.2.2 ([12])**.** *Consider the Ising game on a graph G with n vertices and maximal degree d. The relaxation time of the logit dynamics for this game is at most*

$$n \cdot e^{(4\chi(G)+2d)\beta} .$$

Then, the bound to the mixing time follows from Theorem 1.3.4.

Other results are known for specific graph structures: in particular, for the *clique* and the *ring*. We start with the clique (see Theorem 15.3 in [75]).

Theorem 3.2.3. *Consider the Ising game on a clique with n vertices. If $\beta < \frac{1}{n}$, the mixing time of the logit dynamics for this game is*

$$t_{\text{mix}}(\varepsilon) \leqslant \frac{n(\log n - \log \varepsilon)}{1 - n\beta} .$$

If $\beta > \frac{1}{n}$, then there exists $r = r(n \cdot \beta)$ positive such that

$$t_{\text{mix}} = \mathcal{O}\left(e^{rn}\right) .$$

Above theorem does not say what happens when $\beta = 1/n$. This case is considered by the following theorem due to Levin et al. [74].

Theorem 3.2.4 ([74])**.** *Consider the Ising game on a clique with n vertices. If $\beta = \frac{1}{n}$, then the mixing time of the logit dynamics for this game is*

$$t_{\text{mix}} = \Theta\left(n^{3/2}\right) .$$

Lastly, Ding et al. [39] gave a full characterization of the mixing time evolution as β increase, showing how the mixing time changes in function of the distance between β and $1/n$.

The mixing time for the ring topology is faster as showed by the following theorem (see Theorem 15.4 in [75]).

Theorem 3.2.5. *Consider the Ising game on a ring with n vertices. For any $\beta > 0$ and fixed $\varepsilon > 0$, the mixing time of the logit dynamics for this game is*

$$t_{\text{mix}}(\varepsilon) = \Theta\left(\left(1 + e^{4\beta}\right) n \log n\right) .$$

3.3 How our work relates to the previous literature

In next chapters we present our results about the mixing time of the logit dynamics for different games or classes of games. Specifically, we first show our approach with some introductory games in Chapter 4 and then we show our main results in Chapter 5.

Our main contribution is to propose logit equilibrium as a new equilibrium concept for dealing with the evolution of complex systems. Differently, works summarized in Section 3.1 focus on Nash equilibrium and its refinements as equilibrium concepts.

This difference between our work and previous literature about logit dynamics has another meaningful effect: our results holds for every β, whereas results given in Section 3.1 assume β very large. This is because, for small values of β the probability that the dynamics will not be in a Nash equilibrium is high, even at stationarity: analyzing the convergence to this equilibrium is not meaningful, since the chain can leave this profile quickly. Nevertheless, many ideas arising from the previous results about logit dynamics guided us in our research work.

Our approach is instead similar to the one pursued by the works about Glauber dynamics for the Ising model cited in Section 3.2: in particular, some of our findings extend results given in that section to more generic graphical coordination games. However, we also consider wider classes of games, such as *potential games*. In general, techniques and ideas from results in Section 3.2 have been really useful to our research.

Lastly, in Chapter 6, we introduce *metastability* as a way to deal with logit dynamics for games when the mixing time is exponential. Even if metastability was known, especially in Physics literature (that we will survey in Chapter 6), our approach is completely original, in which we move our focus from states to distributions.

Chapter 4

Mixing time and stationary expected social welfare

One of the main goal of this work is to introduce the *logit equilibrium* as equilibrium concept for games evolving according the logit dynamics. In this chapter, in order to justify this adoption, we give bounds to the time that the dynamics takes to converge to the equilibrium and we estimate the quality of the equilibrium. Specifically, we study the mixing time of the logit dynamics and the stationary expected social welfare for some introductory games[1]:

- We start by analyzing in Section 4.2 the logit dynamics for a simple 3-player linear *congestion game* (the CK game [33]) which exhibits the worst Price of Anarchy among linear congestion games. We show that the mixing time of the logit dynamics is upper bounded by a constant independent of β. Moreover, we show that the stationary expected social welfare is larger than the social welfare of the worst Nash equilibrium for all β;

- Then, in Section 4.3, we analyze the basic *coordination games* given in (3.1). Here we show that, under some conditions, the stationary expected social welfare is larger than the social welfare of the worst Nash equilibrium. We give upper and lower bounds on the mixing time exponential in β. We also observe that the same bounds apply to anti-coordination games;

- Finally, in Sections 4.4 and 4.5, we apply our analysis to two simple n-player games: the OR game and the XOR game. We give upper and lower bounds on the mixing time: we show that the mixing time of the OR game can be upper bounded by a function independent of β, while the mixing time of the XOR game increases exponentially in β. We also prove that for $\beta = \mathcal{O}(\log n)$ the mixing time is polynomial in n for both games. Despite their *game-theoretic* simplicity, the analytical study of the mixing time of the logit dynamics for the two n-player games is far from trivial.

Before showing these results, we give some preliminary definitions in Section 4.1.

4.1 Preliminary definitions

Let $\mathcal{G}$ be a game with profile space S. Let $W \colon S \longrightarrow \mathbb{R}$ be a *social welfare function* (in this chapter we assume that W is simply the sum of all the utility functions $W(\mathbf{x}) = \sum_{i=1}^{n} u_i(\mathbf{x})$, but clearly any other function of interest can be analyzed). We define the *stationary expected*

[1]Most of the results in this chapter already appeared in [5].

social welfare as the expectation of W when the strategy profiles are random according to the stationary distribution π of the Markov chain, i.e.

$$\mathbf{E}_{\pi}[W] = \sum_{\mathbf{x} \in S} W(\mathbf{x})\pi(\mathbf{x}) .$$

In this chapter we will bound the mixing time and we evaluate the stationary expected social welfare of the logit dynamics for some simple but interesting games. We illustrate the approach of this chapter on the two simple examples given in Section 2.2.

Matching Pennies. Since the uniform distribution is stationary for this game, the stationary expected social welfare is 0 for every inverse noise β.

As for the mixing time, it is easy to see that it is upper bounded by a constant independent of β. Indeed, a direct calculation shows that, for every $\mathbf{x} \in \{HH, HT, TH, TT\}$ and for every $\beta \geqslant 0$ it holds that

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

A *stairs* game. One of the main techniques used to give upper bounds on the mixing time of Markov chains is the coupling technique (see Theorem 1.3.1). In this example we use it to upper bound the mixing time of the logit dynamics for a simple game.

We can define a coupling of two Markov chains starting at two different profiles as follows: choose $i \in [n]$ uniformly at random and perform the same update at player i in both chains[2]. When every player has been chosen at least once the two chains have coalesced. From the coupon collector's argument, it takes $\mathcal{O}(n \log n)$ to have that, with probability at least 3/4, all players have been chosen at least once. By applying Theorem 1.3.1 we have that the mixing time is $\mathcal{O}(n \log n)$.

In the above examples, it turned out that the mixing time of the logit dynamics can be upper bounded by functions that do not depend on the inverse noise β. As we shall see in the next sections, this is not always the case. Moreover, the analysis of the mixing time is usually far from trivial.

4.1.1 Description of the Coupling

Throughout the chapter we will use the *coupling* and *path coupling* techniques (see Theorem 1.3.1 and Theorem 1.3.3) to give upper bounds on mixing times. Since we will use the same coupling idea in several proofs, we describe it here and we will refer to this description when we will need it.

Consider an n-player 2-strategy game $\mathcal{G}$ and let us rename 0 and 1 the strategies of every player. For every pair of strategy profiles $\mathbf{x} = (x_1, \ldots, x_n), \mathbf{y} = (y_1, \ldots, y_n) \in \{0,1\}^n$ we define a coupling (X_1, Y_1) of two copies of the Markov chain with transition matrix P defined in (2.2) for which $X_0 = \mathbf{x}$ and $Y_0 = \mathbf{y}$.

The coupling proceeds as follows: first, pick a player i uniformly at random; then, update the strategies x_i and y_i of player i in the two chains, by setting

$$(x_i, y_i) = \begin{cases} (0,0), & \text{with probability } \min\{\sigma_i(0 \mid \mathbf{x}), \sigma_i(0 \mid \mathbf{y})\} ; \\ (1,1), & \text{with probability } \min\{\sigma_i(1 \mid \mathbf{x}), \sigma_i(1 \mid \mathbf{y})\} ; \\ (0,1), & \text{with probability } \sigma_i(0 \mid \mathbf{x}) - \min\{\sigma_i(0 \mid \mathbf{x}), \sigma_i(0 \mid \mathbf{y})\} ; \\ (1,0), & \text{with probability } \sigma_i(1 \mid \mathbf{x}) - \min\{\sigma_i(1 \mid \mathbf{x}), \sigma_i(1 \mid \mathbf{y})\} . \end{cases}$$

[2] This is the same coupling used in the analysis of the lazy random walk on the hypercube (e.g. see Section 5.3.3 in [75]), the only difference being that the probability of choosing 0 or 1 is not $1/2, 1/2$ but $1/(1+e^{\beta})$, $1/(1+e^{-\beta})$.

Three easy observations are in order: if $\sigma_i(0 \mid \mathbf{x}) = \sigma_i(0 \mid \mathbf{y})$ and player i is chosen, then, after the update, we have $x_i = y_i$; for every player i, at most one of the updates $(x_i, y_i) = (0, 1)$ and $(x_i, y_i) = (1, 0)$ has positive probability; if i is chosen for update, then the marginal distributions of x_i and y_i agree with $\sigma_i(\cdot \mid \mathbf{x})$ and $\sigma_i(\cdot \mid \mathbf{y})$ respectively, indeed, for $b \in \{0, 1\}$, the probability that $x_i = b$ is

$$\min\{\sigma_i(b \mid \mathbf{x}), \sigma_i(b \mid \mathbf{y})\} + \sigma_i(b \mid \mathbf{x}) - \min\{\sigma_i(b \mid \mathbf{x}), \sigma_i(b \mid \mathbf{y})\} = \sigma_i(b \mid \mathbf{x}),$$

and the probability that $y_i = b$ is

$$\begin{aligned} &\min\{\sigma_i(b \mid \mathbf{x}), \sigma_i(b \mid \mathbf{y})\} + \sigma_i(1 - b \mid \mathbf{x}) - \min\{\sigma_i(1 - b \mid \mathbf{x}), \sigma_i(1 - b \mid \mathbf{y})\} = \\ &= \min\{\sigma_i(b \mid \mathbf{x}), \sigma_i(b \mid \mathbf{y}) \ + (1 - \sigma_i(b \mid \mathbf{x})) - (1 - \max\{\sigma_i(b \mid \mathbf{x}), \sigma_i(b \mid \mathbf{y})\}) = \sigma_i(b \mid \mathbf{y}). \end{aligned}$$

We define $G = (\Omega, E)$ as the *Hamming graph* of the game, where $\Omega = \{0, 1\}^n$ is the set of strategy profiles, and two profiles $\mathbf{x} = (x_1, \ldots, x_n), \mathbf{y} = (y_1, \ldots, y_n) \in \Omega$ are adjacent if they differ only for the strategy of one player, i.e.

$$\{\mathbf{x}, \mathbf{y}\} \in E \iff \mathbf{x} \sim \mathbf{y}. \tag{4.1}$$

For the path coupling technique (see Theorem 1.3.3), the coupling described above is applied only to pairs of adjacent starting profiles.

4.2 A 3-player congestion game

In this section we analyze the CK game, a simple 3-player linear *congestion game* introduced in [33]. This game is interesting because it highlights the weakness of the Price of Anarchy notion for the logit dynamics. Indeed, the CK game exhibits the worst Price of Anarchy with respect to the average social welfare among all linear congestion games with 3 or more players. But, as we shall see soon, the stationary expected social welfare of the logit dynamics is always larger than the social welfare of the worst Nash equilibrium and, for large enough β, players spend most of the time in the best Nash equilibrium. Moreover, we will show that the mixing time of the logit dynamics can be bounded independently from β: that is, the stationary distribution guarantees a good social welfare and it is quickly reached by the system.

Let us now describe the CK game. We have 3 players and 6 facilities divided into two sets: $G = \{g_1, g_2, g_3\}$ and $H = \{h_1, h_2, h_3\}$. Player $i \in \{0, 1, 2\}$ has two strategies: Strategy "0" consists in selecting facilities (g_i, h_i); Strategy "1" consists in selecting facilities $(g_{i+1}, h_{i-1}, h_{i+1})$ (index arithmetic is modulo 3). The cost of a facility is the number of players choosing such facility, and the utility of a player is minus the sum of the costs of the facilities she selected. It easy to see that this game has two pure Nash equilibria: the solution where every player plays strategy 0 (each player pays 2, which is optimal), and the solution where every player plays strategy 1 (each player pays 5). The game is a congestion game, and thus, by [114], it is also a potential game and its potential function is:

$$\Phi(\mathbf{x}) = \sum_{j \in G \cup H} \sum_{i=1}^{L_{\mathbf{x}}(j)} i,$$

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

Stationary expected social welfare. It is easy to see that the update probabilities given by the logit dynamics for this game (see Equation (2.1)) only depend on the number of players playing strategy 1 and not on which player is actually playing that strategy. In particular, we have that, from a profile $\mathbf{x}$, the player i, if selected for update, plays strategy 0 with the following probabilities:

$$\sigma_i(0 \mid |\mathbf{x}_{-i}| = 0) = \frac{1}{1+e^{-4\beta}}\,, \quad \sigma_i(0 \mid |\mathbf{x}_{-i}| = 1) = \frac{1}{1+e^{-2\beta}}\,, \quad \sigma_i(0 \mid |\mathbf{x}_{-i}| = 2) = \frac{1}{2}\,, \tag{4.2}$$

and strategy 1 with the remaining probabilities. Next theorem evaluates the stationary expected social welfare for this game.

Theorem 4.2.1 (Expected social welfare)**.** *The stationary expected social welfare* $\mathbf{E}_\pi[W]$ *of the logit dynamics for the* CK *game is*

$$\mathbf{E}_\pi[W] = -\frac{6+39e^{-4\beta}+63e^{-6\beta}}{1+3e^{-4\beta}+4e^{-6\beta}}\,.$$

Proof. We notice that two profiles with the same number of players playing strategy 1 have both the same potential (and, by Equation (2.3), the same stationary distribution) and the same social welfare. Thus, $\pi(\mathbf{x}) = \pi[k]$ and $W(\mathbf{x}) = W[k]$ for a profile $\mathbf{x}$ such that $|\mathbf{x}|_1 = k$, with

$$\pi[0] = \frac{e^{-6\beta}}{Z(\beta)}\,, \qquad \pi[1] = \frac{e^{-10\beta}}{Z(\beta)}\,, \qquad \pi[2] = \pi[3] = \frac{e^{-12\beta}}{Z(\beta)}\,,$$

where $Z(\beta) = e^{-6\beta} + 3e^{-10\beta} + 4e^{-12\beta}$, and

$$W[0] = -6\,, \qquad W[1] = -13\,, \qquad W[2] = -16\,, \qquad W[3] = -15\,.$$

Hence, the stationary expected social welfare is

$$\mathbf{E}_\pi[W] = -\frac{6\cdot e^{-6\beta} + 3\cdot 13\cdot e^{-10\beta} + (3\cdot 16+15)\cdot e^{-12\beta}}{e^{-6\beta}+3e^{-10\beta}+4e^{-12\beta}} = -\frac{6+39e^{-4\beta}+63e^{-6\beta}}{1+3e^{-4\beta}+4e^{-6\beta}}\,.$$

□

Notice that for $\beta = 0$ we have $\mathbf{E}_\pi[W] = -27/2$, which is better than the social welfare of the worst Nash equilibrium. This means that, even if each player selects her strategy at random, the logit dynamics drives the system to a random profile whose expectation is better than the worst Nash equilibrium. We also observe that $\mathbf{E}_\pi[W]$ increases with β and thus the long-term behavior of the logit dynamics gives a better social welfare than the worst Nash equilibrium for any $\beta \geqslant 0$. Moreover, the stationary expected social welfare approaches the optimal social welfare as β tends to ∞.

Mixing time. Now we study the mixing time of the logit dynamics for the CK game and we show that it is bounded by a constant for any $\beta \geqslant 0$. The proof will use the coupling theorem (see Theorem 1.3.1).

Theorem 4.2.2 (Mixing time)**.** *There exists a constant* τ *such that the mixing time* t_{mix} *of the logit dynamics of the* CK *game is upper bounded by* τ *for every* $\beta \geqslant 0$.

Proof. First, we notice that the update probabilities given in Equation (4.2) imply that

$$\forall\, i, \forall\, \mathbf{x}, \forall\, \beta, \qquad \sigma_i(0 \mid \mathbf{x}) \geqslant 1/2\,. \tag{4.3}$$

Let X_t and Y_t be two copies of the logit dynamics for the CK game, starting in $\mathbf{x}$ and $\mathbf{y}$ respectively, coupled as described in Section 4.1.1. It is easy to check that, by Equation (4.3), the player selected for update, chooses strategy 0 in both chain with probability at least 1/2.

Finally, we bound the probability that after three steps the two coupled chains coalesce: it is at least as large as the probability that we choose three different players and all of them play strategy 0 at their turn, i.e.

$$\mathbf{P}_{\mathbf{x},\mathbf{y}}\left(X_3 = Y_3\right) \geqslant \frac{1}{2} \cdot \frac{1}{3} \cdot \frac{1}{6} = \frac{1}{36}\,.$$

Since this bound holds for every starting pair $(\mathbf{x}, \mathbf{y})$, we have that the probability the two chains have not yet coalesced after $3t$ steps is

$$\mathbf{P}_{\mathbf{x},\mathbf{y}}\left(X_{3t} \neq Y_{3t}\right) \leqslant \left(1 - \frac{1}{36}\right)^t \leqslant e^{-t/36}\,.$$

The thesis follows from Theorem 1.3.1. □

4.3 Two player games

In this section we analyze the performance of the logit dynamics for 2×2 coordination games given in (3.1) and 2×2 anti-coordination games.

Coordination games. For convenience sake we name

$$\Delta := a - d \quad \text{and} \quad \delta := b - c\,. \tag{4.4}$$

It is easy to see that the coordination game is a potential game and the following function is an exact potential for it:

$$\Phi(0,0) = -\Delta \qquad \Phi(0,1) = \Phi(1,0) = 0 \qquad \Phi(1,1) = -\delta.$$

This game has two pure Nash equilibria: $(0,0)$, where each player has utility a, and $(1,1)$, where each player has utility b. As $d + c < a + b$, the social welfare is maximized at one of the two equilibria.

We analyze the mixing time of the logit dynamics for 2×2 coordination games and compute its stationary expected social welfare as a function of β.

Stationary expected social welfare. The logit dynamics for the coordination game defined by the payoffs in Table 3.1 establishes that, from a profile $\mathbf{x}$, player i selected for update plays according to the following probability distribution (see Equation (2.1)):

$$\begin{array}{lclclcl} \sigma_i(0 \mid \mathbf{x}_{-i} = 0) & = & \frac{1}{1+e^{-\Delta\beta}}\,, & \qquad & \sigma_i(1 \mid \mathbf{x}_{-i} = 0) & = & \frac{1}{1+e^{\Delta\beta}}\,, \\ \sigma_i(0 \mid \mathbf{x}_{-i} = 1) & = & \frac{1}{1+e^{\delta\beta}}\,, & \qquad & \sigma_i(1 \mid \mathbf{x}_{-i} = 1) & = & \frac{1}{1+e^{-\delta\beta}}\,. \end{array}$$

Next theorem bounds the stationary expected social welfare $\mathbf{E}_\pi\left[W\right]$ obtained by the logit dynamics and gives conditions for which $\mathbf{E}_\pi\left[W\right]$ is better than the social welfare SW_N of the worst Nash equilibrium.

Theorem 4.3.1 (Expected social welfare)**.** *The stationary expected social welfare $\mathbf{E}_\pi\left[W\right]$ of the logit dynamics for the coordination game in Table 3.1 is*

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

Moreover, if $a \neq b$ then $\mathbf{E}_\pi\left[W\right] \geqslant \mathsf{SW}_N$ for β sufficiently large.

Proof. The stationary distribution π of the logit dynamics is

$$\pi(0,0) = \frac{e^{\Delta\beta}}{Z(\beta)} \qquad \pi(1,1) = \frac{e^{\delta\beta}}{Z(\beta)} \qquad \pi(0,1) = \pi(1,0) = \frac{1}{Z(\beta)}$$

where $Z(\beta) = e^{\Delta\beta} + e^{\delta\beta} + 2$.

Since $\mathbf{E}_\pi[W] = 2\cdot\mathbf{E}_\pi[u_i]$, we compute the expected utility $\mathbf{E}_\pi[u_i]$ of player i at the stationary distribution,

$$\begin{aligned}
\mathbf{E}_\pi[u_i] &= \sum_{\mathbf{x}\in\{0,1\}^2} u_i(\mathbf{x})\pi(\mathbf{x}) \\
&= \frac{ae^{\Delta\beta} + be^{\delta\beta} + c + d}{e^{\Delta\beta} + e^{\delta\beta} + 2} \\
&= \frac{a + be^{-(\Delta-\delta)\beta} + (c+d)e^{-\Delta\beta}}{1 + e^{-(\Delta-\delta)\beta} + 2e^{-\Delta\beta}}\,.
\end{aligned}$$

Thus, if $a > b$ and $\beta \geqslant \max\left\{0, \frac{1}{\Delta}\log\frac{2b-c-d}{a-b}\right\}$, we have

$$\mathbf{E}_\pi[W] - \mathsf{SW}_N = 2\cdot\frac{a + be^{-(\Delta-\delta)\beta} + (c+d)e^{-\Delta\beta}}{1 + e^{-(\Delta-\delta)\beta} + 2e^{-\Delta\beta}} - 2b = 2\cdot\frac{(a-b) - (2b-c-d)e^{-\Delta\beta}}{1 + e^{-(\Delta-\delta)\beta} + 2e^{-\Delta\beta}} \geqslant 0\,.$$

Similarly, we obtain $\mathbf{E}_\pi[W] - \mathsf{SW}_N \geqslant 0$ if $b > a$ and $\beta \geqslant \max\left\{0, \frac{1}{\delta}\log\frac{2a-c-d}{b-a}\right\}$. □

Mixing time. Now we study the mixing time of the logit dynamics for coordination games and we show that it is exponential in β and in the minimum potential difference between adjacent profiles.

Theorem 4.3.2 (Mixing Time)**.** *The mixing time of the logit dynamics for the coordination game 3.1 is $\Theta\left(e^{\delta\beta}\right)$ for every $\beta \geqslant 0$.*

Proof. Upper bound: We apply the Path Coupling technique (see Theorem 1.3.3) with the Hamming graph defined in (4.1) and all the edge-weights set to 1. Let $\mathbf{x}$ and $\mathbf{y}$ be two profiles differing only for the player j and consider the coupling defined in Section 4.1.1 for this pair of profiles. Now we bound the expected distance of the two coupled chains after one step.

We denote by $b_i(\mathbf{x},\mathbf{y})$ the probability that both chains perform the same update given that player i has been selected for strategy update. Clearly, $b_i(\mathbf{x},\mathbf{y}) = 1$ for $i = j$, while for $i \neq j$, we have

$$\begin{aligned}
b_i(\mathbf{x},\mathbf{y}) &= \min\{\sigma_i(0 \mid \mathbf{x}),\, \sigma_i(0 \mid \mathbf{y})\} + \min\{\sigma_i(1 \mid \mathbf{x}),\, \sigma_i(1 \mid \mathbf{y})\} \\
&= \frac{1}{1+e^{\Delta\beta}} + \frac{1}{1+e^{\delta\beta}}\,.
\end{aligned}$$

For sake of readability we set

$$p = \frac{1}{1+e^{\Delta\beta}} \qquad \text{and} \qquad q = \frac{1}{1+e^{\delta\beta}}$$

and thus $b_i(\mathbf{x},\mathbf{y}) = p + q$. To compute $\mathbf{E}_{\mathbf{x},\mathbf{y}}[\rho(X_1,Y_1)]$, we observe that the logit dynamics chooses player j with probability $1/2$. In this case, as $b_j(\mathbf{x},\mathbf{y}) = 1$, the coupling updates both chains in the same way, resulting in $X_1 = Y_1$. Similarly, player $i \neq j$ is chosen for strategy update with probability $1/2$. In this case, with probability $b_i(\mathbf{x},\mathbf{y})$ the coupling performs the same update in both chains resulting in $\rho(X_1,Y_1) = 1$. Instead with probability $1 - b_i(\mathbf{x},\mathbf{y})$, the coupling performs different updates on the chains resulting in $\rho(X_1, Y_1) = 2$. Therefore we have,

$$\begin{aligned}
\mathbf{E}_{\mathbf{x},\mathbf{y}}\left[\rho(X_1, Y_1)\right] &= \frac{1}{2} b_i(\mathbf{x},\mathbf{y}) + 2 \cdot \frac{1}{2}(1 - b_i(\mathbf{x},\mathbf{y})) \\
&= 1 - \frac{1}{2} b_i(\mathbf{x},\mathbf{y}) = 1 - \frac{1}{2}(p+q) \leqslant e^{-\frac{1}{2}(p+q)} .
\end{aligned}$$

Since the diameter of the Hamming graph is 2, from Theorem 1.3.3, with $\alpha = \frac{1}{2}(p+q)$, follows that

$$t_{\mathrm{mix}}(\varepsilon) \leqslant \frac{2\left(\log 2 + \log(1/\varepsilon)\right)}{p+q} = \frac{1}{p+q} \log \frac{4}{\varepsilon^2} .$$

Lower bound: We use the *relaxation time* bound (see Theorem 1.3.4). The transition matrix of the logit dynamics is

$$P = \left(\begin{array}{c|cccc}
 & 00 & 01 & 10 & 11 \\
\hline
00 & 1-p & p/2 & p/2 & 0 \\
01 & \frac{1-p}{2} & \frac{p+q}{2} & 0 & \frac{1-q}{2} \\
10 & \frac{1-p}{2} & 0 & \frac{p+q}{2} & \frac{1-q}{2} \\
11 & 0 & q/2 & q/2 & 1-q
\end{array}\right)$$

It is easy to see that the second largest eigenvalue of P is $\lambda_\star = \frac{(1-p)+(1-q)}{2}$, hence the relaxation time is $t_{\mathrm{rel}} = 1/(1-\lambda_\star) = \frac{2}{p+q}$, and for the mixing time we have

$$\begin{aligned}
t_{\mathrm{mix}}(\varepsilon) &\geqslant (t_{\mathrm{rel}} - 1) \log \frac{1}{2\varepsilon} = \frac{2-(p+q)}{p+q} \log \frac{1}{2\varepsilon} \\
&\geqslant \frac{1}{p+q} \log \frac{1}{2\varepsilon} .
\end{aligned}$$

In the last inequality we used that p and q are both smaller than $1/2$.

Finally, the theorem follows by observing that

$$\frac{1}{p+q} = \frac{1}{\frac{1}{1+e^{\Delta\beta}} + \frac{1}{1+e^{\delta\beta}}} = \Theta\left(e^{\delta\beta}\right) .$$

□

Notice that, if we used the relaxation time to upper bound the mixing time (see Theorem 1.3.4) we would get a non-tight bound, hence in the above proof we had to resort to the path coupling for the upper bound.

Anti-coordination games. Very similar results can be obtained for *anti-coordination games*. These are two-player games in which the players have an advantage in selecting different strategies. They model many settings where there is a common and exclusive resource: two players have to decide whether to use the resource or to drop it. If they both try to use it, then a deadlock occurs and this is bad for both players. Usually, these games are described by a payoff matrix like (3.1), where we assume that $d > a$ and $c > b$ and that $d - a \geqslant c - b$. Notice that Nash equilibria of this game are *unfair*, as one player has utility $\max\{c, d\}$ and the other $\min\{c, d\}$.

For the logit dynamics, we have that, for all β, the stationary expected social welfare is worse than the one guaranteed by a Nash equilibrium. On the other hand, for sufficiently large β we have that the expected utility of a player is always better than $\min\{c, d\}$: that is, in the logit dynamics each player expects to gain more than in the worst Nash equilibrium. Moreover, the stationary distribution is a fair equilibrium, since every player has the same expected utility. As for the coordination games, the mixing time is exponential in β and in the minimum potential difference between adjacent profiles.

4.4 The OR game

In this section we consider the following simple n-player potential game that we here call *OR game*. Every player has two strategies, say $\{0, 1\}$, and each player pays the OR of the strategies of all players (including herself). More formally, the utility function of player $i \in [n]$ is

$$u_i(\mathbf{x}) = \begin{cases} 0, & \text{if } \mathbf{x} = \mathbf{0}\,; \\ -1, & \text{otherwise.} \end{cases}$$

Notice that the OR game has $2^n - n$ Nash equilibria. The only profiles that are not Nash equilibria are the n profiles with exactly one player playing 1. Nash equilibrium $\mathbf{0}$ has social welfare 0, while all the others have social welfare $-n$.

In Theorem 4.4.1 we show that the stationary expected social welfare is always better than the social welfare of the worst Nash equilibrium, and it is *significantly* better for large β. Unfortunately, in Theorem 4.4.2 we show that if β is large enough to guarantee a *good* stationary expected social welfare, then the time needed to get close to the stationary distribution is exponential in n. Finally, in Theorem 4.4.3 we give upper bounds on the mixing time showing that if β is relatively small then the mixing time is polynomial in n, while for large β the upper bound is exponential in n and it is almost-tight with the lower bound. Despite the simplicity of the game, the analysis of the mixing time is far from trivial.

Theorem 4.4.1 (Expected social welfare)**.** *The stationary expected social welfare of the logit dynamics for the OR game is* $\mathbf{E}_\pi[W] = -\alpha n$ *where* $\alpha = \alpha(n, \beta) = \frac{(2^n-1)e^{-\beta}}{1+(2^n-1)e^{-\beta}}$.

Proof. Observe that the OR game is a potential game with exact potential Φ where $\Phi(\mathbf{0}) = 0$ and $\Phi(\mathbf{x}) = 1$ for every $\mathbf{x} \neq \mathbf{0}$. Hence the stationary distribution is

$$\pi(\mathbf{x}) = \begin{cases} 1/Z, & \text{if } \mathbf{x} = \mathbf{0}\,; \\ e^{-\beta}/Z, & \text{if } \mathbf{x} \neq \mathbf{0}\,; \end{cases}$$

where the normalizing factor is $Z = 1 + (2^n - 1)e^{-\beta}$. The expected social welfare is thus

$$\mathbf{E}_\pi[W] = \sum_{\mathbf{x} \in \{0,1\}^n} W(\mathbf{x})\pi(\mathbf{x}) = -n \cdot \frac{(2^n - 1)e^{-\beta}}{1 + (2^n - 1)e^{-\beta}}\,.$$

□

In the next theorem we show that the mixing time can be polynomial in n only if $\beta \leqslant c \log n$ for some constant c.

Theorem 4.4.2 (Lower bound on mixing time)**.** *The mixing time of the logit dynamics for the OR game is*

1. $\Omega(e^\beta)$ *if* $\beta < \log(2^n - 1)$*;*
2. $\Omega(2^n)$ *if* $\beta > \log(2^n - 1)$.

Proof. Consider the set $R \subseteq \{0, 1\}^n$ containing only the state $\mathbf{0} = (0, \ldots, 0)$ and observe that $\pi(\mathbf{0}) \leqslant 1/2$ for $\beta \leqslant \log(2^n - 1)$. The bottleneck ratio is

$$B(\mathbf{0}) = \frac{1}{\pi(\mathbf{0})} \sum_{\mathbf{y} \in \{0,1\}^n} \pi(\mathbf{0})P(\mathbf{0}, \mathbf{y}) = \sum_{\mathbf{y} \in \{0,1\}^n : |\mathbf{y}|_1 = 1} P(\mathbf{0}, \mathbf{y}) = n \cdot \frac{1}{n} \frac{1}{1 + e^\beta}\,.$$

Hence, by applying Theorem 1.3.7, the mixing time is

$$t_{\text{mix}} \geqslant \frac{1}{B(\mathbf{0})} = 1 + e^{\beta}\,.$$

If $\beta > \log(2^n - 1)$ instead we consider the set $R \subseteq \{0,1\}^n$ containing all states except state $\mathbf{0}$, and observe that

$$\pi(R) = \frac{1}{Z}(2^n - 1)e^{-\beta} = \frac{(2^n - 1)e^{-\beta}}{1 + (2^n - 1)e^{-\beta}}\,.$$

and $\pi(R) \leqslant 1/2$ for $\beta > \log(2^n - 1)$. It holds that

$$Q(R, \overline{R}) = \sum_{\mathbf{x} \in R} \pi(\mathbf{x})P(\mathbf{x}, \mathbf{0}) = \sum_{\mathbf{x} \in \{0,1\}^n \,:\, |\mathbf{x}|_1 = 1} \pi(\mathbf{x})P(\mathbf{x}, \mathbf{0}) = n\frac{e^{-\beta}}{Z}\frac{1}{n}\frac{1}{1 + e^{-\beta}}\,.$$

The bottleneck ratio is

$$B(R) = \frac{Q(R, \overline{R})}{\pi(R)} = \frac{Z}{(2^n - 1)e^{-\beta}}\frac{e^{-\beta}}{Z}\frac{1}{1 + e^{-\beta}} = \frac{1}{(2^n - 1)(1 + e^{-\beta})} < \frac{1}{2^n - 1}\,.$$

Hence, by applying Theorem 1.3.7, the mixing time is

$$t_{\text{mix}} \geqslant \frac{1}{B(R)} > 2^n - 1\,.$$

□

In the next theorem we give upper bounds on the mixing time depending on the value of β. The theorem shows that, if $\beta \leqslant c \log n$ for some constant c, the mixing time is effectively polynomial in n with degree depending on c. The use of the path coupling technique in the proof of the theorem requires a careful choice of the edge-weights.

Theorem 4.4.3 (Upper bound on mixing time)**.** *The mixing time of the logit dynamics for the OR game is $\mathcal{O}(n^{5/2}2^n)$ for every β. Moreover, for small values of β the mixing time is*

1. *$\mathcal{O}(n \log n)$ if $\beta < (1 - \varepsilon)\log n$, for an arbitrary small constant $\varepsilon > 0$;*
2. *$\mathcal{O}(n^{c+3} \log n)$ if $\beta \leqslant c \log n$, where $c \geqslant 1$ is an arbitrary constant.*

Proof. We apply the path coupling technique (see Theorem 1.3.3) with the Hamming graph defined in (4.1). Let $\mathbf{x}, \mathbf{y} \in \{0,1\}^n$ be two profiles differing only at player $j \in [n]$ and, without loss of generality, let us assume $|\mathbf{x}| = k - 1$ and $|\mathbf{y}| = k$ for some $k = 1, \ldots, n$. We set the weight of edge $\{\mathbf{x}, \mathbf{y}\}$ depending only on k, i.e. $\ell(\mathbf{x}, \mathbf{y}) = \delta_k$ where $\delta_k \geqslant 1$ will be chosen later. Consider the coupling defined in Section 4.1.1.

Now we evaluate the expected distance after one step $\mathbf{E}_{\mathbf{x},\mathbf{y}}\left[\rho(X_1, Y_1)\right]$ of the two coupled chains (X_t, Y_t) starting at $(\mathbf{x}, \mathbf{y})$. Let i be the player chosen for the update. Observe that if $i = j$, i.e. if we update the player where $\mathbf{x}$ and $\mathbf{y}$ are different (this holds with probability $1/n$), then the distance after one step is zero, otherwise we distinguish four cases depending on the value of k.

<u>Case $k = 1$</u>: In this case profile $\mathbf{x}$ is all zeros and profile $\mathbf{y}$ has only one 1 and the length of edge $\{\mathbf{x}, \mathbf{y}\}$ is $\ell(\mathbf{x}, \mathbf{y}) = \delta_1$. When choosing a player $i \neq j$ (this happens with probability $(n-1)/n$), at the next step the two chains will be at distance δ_1 (if in both chains player i chooses strategy 0, and this holds with probability $\min\{\sigma_i(0 \mid \mathbf{x}), \sigma_i(0 \mid \mathbf{y})\}$), or at distance δ_2 (if in both chains player i chooses strategy 1, and this holds with probability $\min\{\sigma_i(1 \mid \mathbf{x}), \sigma_i(1 \mid \mathbf{y})\}$), or at

distance $\delta_1 + \delta_2$ (if player i chooses strategy 0 in chain X_1 and strategy 1 in chain Y_1, and this holds with the remaining probability). Notice that, from the definition of the coupling, it will never happen that player i chooses strategy 1 in chain X_1 and strategy 0 in chain Y_1, indeed we have that

$$\begin{aligned} \min\{\sigma_i(0 \mid \mathbf{x}),\, \sigma_i(0 \mid \mathbf{y})\} &= \sigma_i(0 \mid \mathbf{y}) = \frac{1}{2} \qquad \text{and} \\ \min\{\sigma_i(1 \mid \mathbf{x}),\, \sigma_i(1 \mid \mathbf{y})\} &= \sigma_i(1 \mid \mathbf{x}) = \frac{1}{1+e^{\beta}}\,. \end{aligned} \tag{4.5}$$

Hence the expected distance after one step is

$$\begin{aligned} \mathbf{E}_{\mathbf{x},\mathbf{y}}\left[\rho(X_1, Y_1)\right] &= \frac{n-1}{n}\left(\frac{1}{2}\delta_1 + \frac{1}{1+e^{\beta}}\delta_2 + \left(1 - \frac{1}{2} - \frac{1}{1+e^{\beta}}\right)(\delta_1 + \delta_2)\right) \\ &= \frac{n-1}{n}\left(\frac{\delta_1}{1+e^{-\beta}} + \frac{\delta_2}{2}\right). \end{aligned}$$

<u>Case $k = 2$</u>: In this case we have $x_j = 0$ and $y_j = 1$, there is another player $h \in [n] \setminus \{j\}$ where $x_h = y_h = 1$, and for all the other players $i \in [n] \setminus \{j, h\}$ it holds $x_i = y_i = 0$. Hence the length of edge $\{\mathbf{x}, \mathbf{y}\}$ is $\ell(\mathbf{x}, \mathbf{y}) = \delta_2$.

When player h is chosen (this holds with probability $1/n$) we have that $\sigma_h(s \mid \mathbf{x})$ and $\sigma_h(s \mid \mathbf{y})$ for $s \in \{0, 1\}$ are the same as in (4.5). At the next step the two chains will be at distance δ_2 (if player h stays at strategy 1 in both chains), or at distance δ_1 (if player h chooses strategy 0 in both chains), or at distance $\delta_1 + \delta_2$ (if player h stays at strategy 0 in chain X_1 and chooses strategy 1 in chain Y_1).

When a player $i \notin \{h, j\}$ is chosen (this holds with probability $(n-2)/n$) we have that $\sigma_i(0, \mathbf{x}) = \sigma_i(1, \mathbf{x}) = \sigma_i(0, \mathbf{y}) = \sigma_i(1, \mathbf{y}) = 1/2$. Thus in this case the two coupled chains always perform the same choice at player i, and at the next step they will be at distance δ_2 (if player i stays at strategy 0 in both chains) or at distance δ_3 (if player i chooses strategy 1 in both chains).

Hence the expected distance after one step is

$$\begin{aligned} \mathbf{E}_{\mathbf{x},\mathbf{y}}\left[\rho(X_1, Y_1)\right] &= \frac{1}{n}\left(\frac{\delta_1}{2} + \frac{\delta_2}{1+e^{\beta}} + \left(1 - \frac{1}{2} - \frac{1}{1+e^{\beta}}\right)(\delta_1 + \delta_2)\right) + \frac{n-2}{n}\left(\frac{1}{2}\delta_2 + \frac{1}{2}\delta_3\right) \\ &= \frac{1}{2n}\left(\frac{2}{1+e^{-\beta}}\delta_1 + (n-1)\delta_2 + (n-2)\delta_3\right). \end{aligned}$$

<u>Case $3 \leqslant k \leqslant n-1$</u>: When a player $i \neq j$ is chosen such that $x_i = y_i = 1$ (this holds with probability $(k-1)/n$) then at the next step the two chains will be at distance δ_k (if i stays at strategy 1) or at distance δ_{k-1} (if i moves to strategy 0). When a player $i \neq j$ is chosen such that $x_i = y_i = 0$ (this holds with probability $(n-k)/n$) then at the next step the two chains will be at distance δ_k (if i chooses to stay at strategy 0) or at distance δ_{k+1} (if i chooses to move to strategy 0). Hence the expected distance after one step is

$$\begin{aligned} \mathbf{E}_{\mathbf{x},\mathbf{y}}\left[\rho(X_1, Y_1)\right] &= \frac{k-1}{n}\left(\frac{1}{2}\delta_k + \frac{1}{2}\delta_{k-1}\right) + \frac{n-k}{n}\left(\frac{1}{2}\delta_k + \frac{1}{2}\delta_{k+1}\right) \\ &= \frac{1}{2n}\left((n-1)\delta_k + (k-1)\delta_{k-1} + (n-k)\delta_{k+1}\right). \end{aligned}$$

<u>Case $k = n$</u>: When a player $i \neq j$ is chosen, then at the next step the two chains will be at distance δ_n or at distance δ_{n-1}. Hence the expected distance after one step is

$$\mathbf{E}_{\mathbf{x},\mathbf{y}}\left[\rho(X_1, Y_1)\right] = \frac{n-1}{n}\left(\frac{1}{2}\delta_n + \frac{1}{2}\delta_{n-1}\right) = \frac{n-1}{2n}(\delta_n + \delta_{n-1})\,. \tag{4.6}$$

In order to apply Theorem 1.3.3 we now have to show that it is possible to choose the edge weights $\delta_1, \dots, \delta_n$ and a parameter $\alpha > 0$ such that

$$\begin{aligned}
\frac{n-1}{n}\left(\frac{\delta_1}{1+e^{-\beta}} + \frac{\delta_2}{2}\right) &\leqslant \delta_1 e^{-\alpha}\,, \\
\frac{1}{2n}\left(\frac{2}{1+e^{-\beta}}\delta_1 + (n-1)\delta_2 + (n-2)\delta_3\right) &\leqslant \delta_2 e^{-\alpha}\,, \\
\frac{1}{2n}\left((n-1)\delta_k + (k-1)\delta_{k-1} + (n-k)\delta_{k+1}\right) &\leqslant \delta_k e^{-\alpha}\,, \qquad \text{for } k = 3, \dots, n-1\,, \\
\frac{n-1}{2n}(\delta_n + \delta_{n-1}) &\leqslant \delta_n e^{-\alpha}\,.
\end{aligned} \tag{4.7}$$

For different values of β, we make different choices for α and for the weights δ_k. For clarity's sake we split the proof in three different lemmas. We denote by $\delta^{\max}$ the largest δ_k.

In Lemma 4.4.4 we show that Inequalities (4.7) are satisfied for every value of β by choosing the weights as follows

$$\delta_k = \begin{cases} \frac{1}{2}[(n-1)\delta_2 + 1], & \text{if } k = 1; \\ \frac{n-k}{k}\delta_{k+1} + 1, & \text{if } 2 \leqslant k \leqslant n-1; \\ 1, & \text{if } k = n; \end{cases}$$

and by setting $\alpha = 1/(2n\delta^{\max})$. From Corollary 4.4.9, we have $\delta^{\max} = \mathcal{O}(\sqrt{n}2^n)$. Observe that the diameter of the Hamming graph is $\sum_{i=1}^{n} \delta_i \leqslant n\delta^{\max}$, hence from Theorem 1.3.3 we obtain $t_{\text{mix}} = \mathcal{O}\left(n^{5/2}2^n\right)$.

In Lemma 4.4.5 we show that, if $\beta < (1-\varepsilon)\log n$ for an arbitrarily small constant $\varepsilon > 0$, Inequalities (4.7) are satisfied, for sufficiently large n, by choosing weights $\delta_1 = n^{1-\varepsilon}$, $\delta_2 = 4/3, \delta_3 = \ldots = \delta_n = 1$, and $\alpha = 1/n$. In this case the diameter is $\mathcal{O}(n)$ and, by Theorem 1.3.3, $t_{\text{mix}} = \mathcal{O}(n \log n)$.

In Lemma 4.4.6 we show that Inequalities (4.7) are satisfied by choosing weights as follows

$$\delta_k = \begin{cases} \frac{1+e^{-\beta}}{2}\left[\frac{a_1}{b_1}\delta_2 + 1\right], & \text{if } k = 1; \\ \frac{a_k}{b_k}\delta_{k+1} + 1, & \text{if } 2 \leqslant k \leqslant n-1; \\ 1, & \text{if } k = n; \end{cases}$$

where $a_1 = n-1$ and $b_1 = ne^{-\beta} + 1$ and, for every $k = 2, \dots, n-1$

$$a_k = (n-k)b_{k-1} \qquad \text{and} \qquad b_k = (n+1)b_{k-1} - (k-1)a_{k-1}\,;$$

and by setting $\alpha = 1/(2n\delta^{\max})$. From Corollary 4.4.12 it follows that, if $\beta \leqslant c \log n$ for a constant $c \in \mathbb{N}$, we have that $\delta_{\max} = \mathcal{O}(n^{c+2})$ and the diameter of the Hamming graph is $\mathcal{O}(n^{c+3})$. Thus, by Theorem 1.3.3 it follows that $t_{\text{mix}} = \mathcal{O}(n^{c+3}\log n)$. □

4.4.1 Technical lemmas

In this section we prove the technical lemmas needed for completing the proof of Theorem 4.4.3.

Lemma 4.4.4. *Let $\delta_1, \dots, \delta_n$ be as follows*

$$\delta_k = \begin{cases} \frac{1}{2}[(n-1)\delta_2 + 1], & \text{if } k = 1; \\ \frac{n-k}{k}\delta_{k+1} + 1, & \text{if } 2 \leqslant k \leqslant n-1; \\ 1, & \text{if } k = n; \end{cases} \tag{4.8}$$

and let $\alpha = 1/(2n\delta^{\max})$ where $\delta^{\max} = \max\{\delta_k : k = 1, \dots, n\}$. Then Inequalities (4.7) are satisfied for every $\beta \geqslant 0$.

Proof. Observe that, for every $k = 1, \dots, n$, the right-hand side of the k-th inequality in (4.7) is

$$\delta_k e^{-\alpha} = \delta_k e^{-1/(2n\delta^{\max})} \geqslant \delta_k \left(1 - \frac{1}{2n\delta^{\max}}\right) = \delta_k - \frac{\delta_k}{2n\delta^{\max}} \geqslant \delta_k - \frac{1}{2n}. \tag{4.9}$$

Now we check that the left-hand side is at most $\delta_k - 1/(2n)$.

<u>First inequality $(k = 1)$:</u> $\frac{n-1}{n}\left(\frac{\delta_1}{1+e^{-\beta}} + \frac{\delta_2}{2}\right) \leqslant \delta_1 e^{-\alpha}$.

From the definition of δ_1 in (4.8) we have that

$$\delta_2 = \frac{2\delta_1 - 1}{n-1}.$$

Hence the left-hand side is

$$\begin{aligned}\frac{n-1}{n}\left(\frac{\delta_1}{1+e^{-\beta}} + \frac{\delta_2}{2}\right) \leqslant \frac{n-1}{n}\left(\delta_1 + \frac{\delta_2}{2}\right) &= \\ \frac{n-1}{n}\left(\delta_1 + \frac{2\delta_1 - 1}{2(n-1)}\right) = \frac{1}{2n}(2n\delta_1 - 1) &= \delta_1 - \frac{1}{2n}.\end{aligned}$$

<u>Second inequality $(k = 2)$:</u> $\frac{1}{2n}\left(\frac{2}{1+e^{-\beta}}\delta_1 + (n-1)\delta_2 + (n-2)\delta_3\right) \leqslant \delta_2 e^{-\alpha}$.

From the definition of δ_2 in (4.8) we have that

$$\delta_3 = \frac{2}{n-2}(\delta_2 - 1).$$

Hence the left-hand side of the second inequality is

$$\begin{aligned}\frac{1}{2n}\left(\frac{2}{1+e^{-\beta}}\delta_1 + (n-1)\delta_2 + (n-2)\delta_3\right) &\leqslant \frac{1}{2n}\left(2\delta_1 + (n-1)\delta_2 + (n-2)\delta_3\right) \\ &= \frac{1}{2n}\left((n-1)\delta_2 + 1 + (n-1)\delta_2 + 2(\delta_2 - 1)\right) \\ &= \frac{1}{2n}(2n\delta_2 - 1) = \delta_2 - \frac{1}{2n}.\end{aligned}$$

<u>Other inequalities $(k = 3, \dots, n-1)$:</u> $\frac{1}{2n}\left((n-1)\delta_k + (k-1)\delta_{k-1} + (n-k)\delta_{k+1}\right) \leqslant \delta_k e^{-\alpha}$.

From the definition of δ_k in (4.8) we have that

$$\delta_{k+1} = \frac{k}{n-k}(\delta_k - 1).$$

Hence the left-hand side is

$$\begin{aligned}\frac{((n-1)\delta_k + (k-1)\delta_{k-1} + (n-k)\delta_{k+1})}{2n} &= \frac{((n-1)\delta_k + (n-k+1)\delta_k + (k-1) + k\delta_k - k)}{2n} \\ &= \frac{1}{2n}(2n\delta_k - 1) = \delta_k - \frac{1}{2n}.\end{aligned}$$

<u>Last inequality $(k = n)$:</u> $\frac{n-1}{2n}(\delta_n + \delta_{n-1}) \leqslant \delta_n e^{-\alpha}$.

Since $\delta_n = 1$ and $\delta_{n-1} = \frac{1}{n-1}\delta_n + 1 = \frac{n}{n-1}$, the left-hand side of the last inequality is

$$\frac{n-1}{2n}(\delta_n + \delta_{n-1}) = \frac{n-1}{2n}(1 + \frac{n}{n-1}) = 1 - \frac{1}{2n}.$$

□

Lemma 4.4.5. *Let $\delta_1, \dots, \delta_n$ be as follows*

$$\delta_1 = n^{1-\varepsilon},\ \delta_2 = 4/3,\ \delta_3 = \cdots = \delta_n = 1$$

where $\varepsilon > 0$ is an arbitrary small constant and let $\alpha = 1/n$. Then Inequalities (4.7) are satisfied for every $\beta \leqslant (1-\varepsilon)\log n$ and n sufficiently large.

Proof. We check that all the inequalities in (4.7) are satisfied.
First inequality ($k = 1$): $\frac{n-1}{n}\left(\frac{\delta_1}{1+e^{-\beta}} + \frac{\delta_2}{2}\right) \leqslant \delta_1 e^{-\alpha}$.
For the left-hand side we have

$$\begin{aligned}
\frac{n-1}{n}\left(\frac{\delta_1}{1+e^{-\beta}} + \frac{\delta_2}{2}\right) &= \left(1-\frac{1}{n}\right)\left(\frac{n^{1-\varepsilon}}{1+e^{-\beta}} + \frac{2}{3}\right) \\
&\leqslant \left(1-\frac{1}{n}\right)\left(\frac{n^{1-\varepsilon}}{1+\frac{1}{n^{1-\varepsilon}}} + \frac{2}{3}\right) = \left(1-\frac{1}{n}\right)\left(\frac{n^{2(1-\varepsilon)}}{n^{1-\varepsilon}+1} + \frac{2}{3}\right) \\
&= \left(1-\frac{1}{n}\right)\left(\frac{(n^{1-\varepsilon}+1)(n^{1-\varepsilon}-1)+1}{n^{1-\varepsilon}+1} + \frac{2}{3}\right) \\
&= \left(1-\frac{1}{n}\right)\left(n^{1-\varepsilon} + \frac{1}{n^{1-\varepsilon}+1} - \frac{1}{3}\right).
\end{aligned}$$

For the right-hand side we have

$$\delta_1 e^{-\alpha} = n^{1-\varepsilon} e^{-1/n} \geqslant n^{1-\varepsilon}\left(1-\frac{1}{n}\right).$$

Hence the left-hand side is smaller than the right-hand one (for n sufficiently large).
Second inequality ($k = 2$): $\frac{1}{2n}\left(\frac{2}{1+e^{-\beta}}\delta_1 + (n-1)\delta_2 + (n-2)\delta_3\right) \leqslant \delta_2 e^{-\alpha}$.
For the left-hand side we have

$$\begin{aligned}
\frac{1}{2n}\left(\frac{2}{1+e^{-\beta}}\delta_1 + (n-1)\delta_2 + (n-2)\delta_3\right) &= \frac{1}{2n}\left(\frac{2}{1+e^{-\beta}}n^{1-\varepsilon} + (n-1)\frac{4}{3} + (n-2)\right) \\
&\leqslant \frac{1}{2n}\left(2n^{1-\varepsilon} + \frac{7}{3}n\right) = \frac{7}{6} + \frac{1}{n^{\varepsilon}}.
\end{aligned}$$

And for the right-hand side we have

$$\delta_2 e^{-\alpha} = \frac{4}{3}e^{-1/n} \geqslant \frac{4}{3}\left(1-\frac{1}{n}\right) \geqslant \frac{4}{3} - \frac{1}{n}.$$

Hence the left-hand side is smaller than the right-hand one (for n sufficiently large).
Third inequality ($k = 3$): $\frac{1}{2n}\left((n-1)\delta_3 + 2\delta_2 + (n-3)\delta_4\right) \leqslant \delta_3 e^{-\alpha}$.
For the left-hand side we have

$$\begin{aligned}
\frac{1}{2n}\left((n-1)\delta_3 + 2\delta_2 + (n-3)\delta_4\right) &= \frac{1}{2n}\left((n-1) + 2\frac{4}{3} + (n-3)\right) \\
&= \frac{1}{2n}(2n-3) \leqslant \left(1-\frac{1}{n}\right).
\end{aligned}$$

And for the right-hand side we have

$$\delta_3 e^{-\alpha} = e^{-1/n} \geqslant \left(1-\frac{1}{n}\right).$$

Hence the left-hand side is smaller than the right-hand one.

Other inequalities ($k \geqslant 4$): $\frac{1}{2n}\left((n-1)\delta_k + (k-1)\delta_{k-1} + (n-k)\delta_{k+1}\right) \leqslant \delta_k e^{-\alpha}$.

Since $\delta_k = \delta_{k-1} = \delta_{k+1} = 1$ the left-hand side is equal to $\frac{n-1}{n}$ and the right-hand side is $e^{-1/n} \geqslant \frac{n-1}{n}$. □

Lemma 4.4.6. *Let $\delta_1, \dots, \delta_n$ be as follows*

$$\delta_k = \begin{cases} \frac{1+e^{-\beta}}{2}\left[\frac{a_1}{b_1}\delta_2 + 1\right], & \text{if } k = 1; \\ \frac{a_k}{b_k}\delta_{k+1} + 1, & \text{if } 2 \leqslant k \leqslant n-1; \\ 1, & \text{if } k = n; \end{cases} \tag{4.10}$$

where $a_1 = n-1$ and $b_1 = ne^{-\beta}+1$ and for every $k = 2, \dots, n-1$

$$a_k = (n-k)b_{k-1} \qquad \text{and} \qquad b_k = (n+1)b_{k-1} - (k-1)a_{k-1}\,,$$

and let $\alpha = 1/(2n\delta^{\max})$ where $\delta^{\max} = \max\{\delta_k \;:\; k = 1, \dots, n\}$. Then Inequalities (4.7) are satisfied for every $\beta \geqslant 0$.

Before to prove the Lemma 4.4.6 we consider the following lemma.

Lemma 4.4.7. *Let b_k defined as in the Lemma 4.4.6. Then, for every $k \geqslant 2$, it holds that $b_k \geqslant k b_{k-1}$.*

Proof. We proceed by induction on k. The base case $k = 2$ follows from

$$b_2 = (n+1)(ne^{-\beta}+1) - (n-1) = (n+1)ne^{-\beta} + 2 > 2(ne^{-\beta}+1) = 2b_1\,.$$

Now suppose the claim holds for $k-1$, that is $b_{k-1} \geqslant (k-1)b_{k-2}$. Then

$$\begin{aligned} b_k &= (n+1)b_{k-1} - (k-1)a_{k-1} \\ &= (n+1)b_{k-1} - (k-1)(n-k+1)b_{k-2} \\ &\geqslant \left[(n+1) - (n-k+1)\right] b_{k-1} = k b_{k-1}\,. \end{aligned}$$

□

Proof (Lemma 4.4.6). Observe that, as in Equation (4.9), for every $k = 1, \dots, n$, the right-hand side of the k-th inequality in (4.7) is

$$\delta_k e^{-\alpha} \geqslant \delta_k - \frac{1}{2n}\,.$$

Now we check that the left-hand side is at most $\delta_k - 1/(2n)$.

First inequality ($k = 1$): $\frac{n-1}{n}\left(\frac{\delta_1}{1+e^{-\beta}} + \frac{\delta_2}{2}\right) \leqslant \delta_1 e^{-\alpha}$.

From the definition of δ_1 in (4.10) we have that

$$\delta_2 = \frac{ne^{-\beta}+1}{n-1}\left(\frac{2\delta_1}{1+e^{-\beta}} - 1\right)\,.$$

Hence the left-hand side is

$$\begin{aligned} \frac{n-1}{n}\left(\frac{\delta_1}{1+e^{-\beta}} + \frac{\delta_2}{2}\right) &= \frac{n-1}{n}\left[\frac{\delta_1}{1+e^{-\beta}} + \frac{ne^{-\beta}+1}{n-1}\left(\frac{\delta_1}{1+e^{-\beta}} - \frac{1}{2}\right)\right] \\ &= \frac{n-1}{n}\frac{\delta_1}{1+e^{-\beta}}\left(1 + \frac{ne^{-\beta}+1}{n-1}\right) - \frac{ne^{-\beta}+1}{2n} \\ &\leqslant \delta_1 - \frac{1}{2n}\,. \end{aligned}$$

Second inequality ($k=2$): $\frac{1}{2n}\left(\frac{2}{1+e^{-\beta}}\delta_1 + (n-1)\delta_2 + (n-2)\delta_3\right) \leqslant \delta_2 e^{-\alpha}$.
From the definition of δ_2 in (4.10) we have that

$$\delta_3 = \frac{b_2}{a_2}(\delta_2 - 1) = \frac{(n+1)b_1 - a_1}{(n-2)b_1}(\delta_2 - 1)\,.$$

Hence the left-hand side is

$$\begin{aligned}
\frac{1}{2n}\left(\frac{2}{1+e^{-\beta}}\delta_1 + (n-1)\delta_2 + (n-2)\delta_3\right) &= \\
\frac{1}{2n}\left[\left(\frac{a_1}{b_1}\delta_2 + 1\right) + (n-1)\delta_2 + \frac{(n+1)b_1 - a_1}{b_1}(\delta_2 - 1)\right] &= \\
\delta_2 - \frac{1}{2n}\frac{nb_1 - a_1}{b_1} = \delta_2 - \frac{1}{2n}\left(n - \frac{n-1}{ne^{-\beta}+1}\right) &\leqslant \delta_2 - \frac{1}{2n}\,.
\end{aligned}$$

Other inequalities ($k = 3, \dots, n-1$): $\frac{1}{2n}\left((n-1)\delta_k + (k-1)\delta_{k-1} + (n-k)\delta_{k+1}\right) \leqslant \delta_k e^{-\alpha}$.
From the definition of δ_k in (4.10) we have that

$$\delta_{k+1} = \frac{b_k}{a_k}(\delta_k - 1) = \frac{(n+1)b_{k-1} - (k-1)a_{k-1}}{(n-k)b_{k-1}}(\delta_k - 1)\,.$$

Hence the left-hand side is

$$\begin{aligned}
\frac{1}{2n}\left((n-1)\delta_k + (k-1)\delta_{k-1} + (n-k)\delta_{k+1}\right) &= \\
\frac{1}{2n}\left[(n-1)\delta_k + (k-1)\left(\frac{a_{k-1}}{b_{k-1}}\delta_k + 1\right) + \frac{(n+1)b_{k-1} - (k-1)a_{k-1}}{b_{k-1}}(\delta_k - 1)\right] &= \\
\delta_k - \frac{1}{2n}\frac{(n-k+2)b_{k-1} - (k-1)a_{k-1}}{b_{k-1}} &= \\
\delta_k - \frac{1}{2n}\left((n-k+2) - (k-1)(n-k+1)\frac{b_{k-2}}{b_{k-1}}\right) &\leqslant \delta_k - \frac{1}{2n}\,.
\end{aligned}$$

where the inequality follows from the Lemma 4.4.7.
Last inequality ($k=n$): $\frac{n-1}{2n}(\delta_n + \delta_{n-1}) \leqslant \delta_n e^{-\alpha}$.
Since $\delta_n = 1$ and $\delta_{n-1} = \frac{a_{n-1}}{b_{n-1}}\delta_n + 1 = \frac{a_{n-1}}{b_{n-1}} + 1$, the left-hand side of the last inequality is

$$\begin{aligned}
\frac{n-1}{2n}(\delta_n + \delta_{n-1}) &= \frac{n-1}{2n}\left(2 + \frac{a_{n-1}}{b_{n-1}}\right) = \frac{n-1}{2n}\left(2 + \frac{b_{n-2}}{b_{n-1}}\right) \\
&\leqslant \frac{n-1}{2n}\left(2 + \frac{1}{n-1}\right) = 1 - \frac{1}{2n}\,.
\end{aligned}$$

where the inequality follows from the Lemma 4.4.7. □

In order to apply the path coupling theorem, we need to bound $\delta_{\max}$: the next lemma will represent the main tool to achieve this goal.

Lemma 4.4.8. *Let $\delta_1, \dots, \delta_n$ be defined recursively as follows: $\delta_n = 1$ and*

$$\delta_k = \gamma_k \delta_{k+1} + 1\,,$$

where $\gamma_k > 0$ for every $k = 1, \dots, n-1$. Let $\delta^{\max} = \max\{\delta_k \,:\, k = 1, \dots, n\}$. Then

$$\delta^{\max} \leqslant n \max\left\{\prod_{i=h}^{j} \gamma_i \;:\; 1 \leqslant h \leqslant j \leqslant n-1\right\}\,.$$

Proof. The lemma follows from the fact that, for $k = 1, \dots, n-1$, we have

$$\delta_k = 1 + \sum_{j=k}^{n-1} \prod_{i=k}^{j} \gamma_i \,.$$

$\square$

Corollary 4.4.9. *Let $\delta_1, \dots, \delta_n$ be defined as in Lemma 4.4.4. Then $\delta^{\max} \leqslant c\sqrt{n}2^n$ for a suitable constant c.*

Proof. From Lemma 4.4.8 and the definition of $\delta_1, \dots, \delta_n$, it holds that

$$\begin{aligned} \delta^{\max} &\leqslant n \max \left\{ \prod_{i=h}^{j} \frac{n-i}{i} \; : \; 1 \leqslant h \leqslant j \leqslant n \right\} \\ &\leqslant n \prod_{i=1}^{\lfloor n/2 \rfloor} \frac{n-i}{i} \leqslant n \binom{n}{\lfloor n/2 \rfloor} \leqslant c\sqrt{n}2^n \,. \end{aligned}$$

for a suitable constant c. $\square$

In order to bound $\delta_{\max}$ when $\delta_1, \dots, \delta_n$ are defined as in Lemma 4.4.6 and $\beta \leqslant c \log n$ for a constant $c \in \mathbb{N}$, we define

$$\gamma_k = \frac{a_k}{b_k} = \frac{p_k e^{-\beta} + l_k}{q_k e^{-\beta} + r_k} \,. \tag{4.11}$$

We can check that $p_1 = 0$, $q_1 = n$ and for every $k > 1$,

$$p_k = (n-k)q_{k-1} \qquad \text{and} \qquad q_k = (n+1)q_{k-1} - (k-1)p_{k-1} \,;$$

we note that $p_k = (n+1)q_{k-1} - (k+1)q_{k-1} \leqslant q_k$ for every k. We can also prove the following simple lemma about q_k.

Lemma 4.4.10. *For every $k \geqslant 1$ constant, we have $q_k \geqslant 2^{-k}n^k$.*

Proof. We proceed by induction on k, with the base $k = 1$ being obvious. Suppose the claim holds for $k-1$, that is $q_{k-1} \geqslant 2^{-(k-1)}n^{k-1}$, then

$$q_k = (n+1)q_{k-1} - (k-1)p_{k-1} \geqslant \frac{n}{2} q_{k-1} \geqslant 2^{-k}n^k \,.$$

$\square$

Moreover, we can check that $l_1 = n-1$, $r_1 = 1$ and for every $k > 1$,

$$l_k = (n-k)r_{k-1} \qquad \text{and} \qquad r_k = (n+1)r_{k-1} - (k-1)l_{k-1} \,;$$

we notice that above recursion gives $l_k = (n-k)(k-1)!$ and $r_k = k!$. Next lemma bounds γ_k defined in Equation 4.11.

Lemma 4.4.11. *Let $\delta_1, \dots, \delta_n$ be defined as in Lemma 4.4.6, γ_k defined as in Equation (4.11) and $\beta \leqslant c \log n$ for a constant $c \in \mathbb{N}$. Then, for sufficiently large n, it holds that*

$$\begin{cases} \gamma_k < n & \forall\, k; \\ \gamma_k < 1 & \text{if } k > c+2; \\ \gamma_{c+2} = \mathcal{O}(1). \end{cases}$$

Proof. Since $p_k \leqslant q_k$, then $(nq_k - p_k)e^{-\beta} > 0$; instead, $l_k - nr_k = (k-1)!(n-k-nk) < 0$. Hence we have for every k

$$\gamma_k - n = \frac{p_k e^{-\beta} + l_k}{q_k e^{-\beta} + r_k} - n = \frac{(l_k - nr_k) - (nq_k - p_k)e^{-\beta}}{q_k e^{-\beta} + r_k} < 0\,.$$

Inductively, we show that for every $k \geqslant c+3$, we have $\gamma_k < 1$. Set $k = c+3$: c is a constant, thus Lemma 4.4.10 holds for $k-1$; hence and since $e^{-\beta} \geqslant n^{-c}$, we have that

$$(q_{c+3} - p_{c+3})e^{-\beta} = [(n+1)q_{c+2} - (c+2)p_{c+2} - (n-c-3)q_{c+2}]e^{-\beta} \geqslant 2q_{c+2}e^{-\beta} \geqslant 2^{-(c+1)}n^2.$$

Instead, $l_{c+3} - r_{c+3} = (c+2)!(n-2c-6) \leqslant (c+2)! \cdot n$. Thus,

$$\gamma_{c+3} - 1 = \frac{(l_{c+3} - r_{c+3}) - (q_{c+3} - p_{c+3})e^{-\beta}}{q_{c+3}e^{-\beta} + r_{c+3}} \leqslant \frac{(c+2)! \cdot n - 2^{-(c+1)}n^2}{q_{c+3}e^{-\beta} + r_{c+3}} < 0\,,$$

for n sufficiently large. Now, suppose that $\gamma_{k-1} < 1$; then, we have

$$\gamma_k - 1 = \frac{a_k - b_k}{b_k} = \frac{(k-1)a_{k-1} - (k+1)b_{k-1}}{b_k} < 0\,,$$

where $a_{k-1} < b_{k-1}$ is implied by the inductive hypothesis.

In order to complete the proof, we need to show that $\gamma_{c+2} = \mathcal{O}(1)$. Similarly to the case $k = c+3$, we obtain $(q_{c+2} - p_{c+2})e^{-\beta} \geqslant 2^{-c}n$ and $l_{c+2} - r_{c+2} \leqslant (c+1)! \cdot n$. Hence,

$$\gamma_{c+2} \leqslant \frac{p_{c+2} + r_{c+2} + (c+1)! \cdot n}{p_{c+2} + r_{c+2} + 2^{-c}n} \leqslant (c+1)! \cdot 2^c = \mathcal{O}(1)\,.$$

□

Corollary 4.4.12. *Let $\delta_1, \ldots, \delta_n$ and c be defined as in Lemma 4.4.6. Then $\delta^{\max} = \mathcal{O}(n^{c+2})$.*

Proof. From Lemma 4.4.8, Lemma 4.4.11 and the definition of $\delta_1, \ldots, \delta_n$ it follows that

$$\begin{aligned} \delta^{\max} &\leqslant n \max\left\{ \prod_{i=h}^{j} \frac{a_i}{b_i} \ : \ 1 \leqslant h \leqslant j \leqslant n \right\} \\ &\leqslant n \prod_{i=1}^{c+2} \frac{a_i}{b_i} = \mathcal{O}(n^{c+2}). \end{aligned}$$

□

4.5 The XOR game

In this section we analyze the logit dynamics for another simple n-player game, the *XOR game*. The XOR game is a symmetric n-player game in which each player has two strategies, denoted by 0 and 1, and each player pays the XOR of the strategies of all players (including herself). More formally, for each $i \in [n]$, the utility function $u_i(\cdot)$ is defined as follows

$$u_i(\mathbf{x}) = \begin{cases} -1, & \text{if } \mathbf{x} \text{ has an odd number of 1's;} \\ 0, & \text{if } \mathbf{x} \text{ has an even number of 1's.} \end{cases}$$

Notice that the XOR game has 2^{n-1} Nash equilibria, namely all profiles with an even number of players playing strategy 1. Nash equilibria have social welfare 0 and profiles not in equilibria have social welfare $-n$. Observe that the XOR game is a potential game with exact potential Φ where $\Phi(\mathbf{x}) = -u_i(\mathbf{x})$ for every $\mathbf{x}$ and every $i \in [n]$. Hence, the stationary distribution is

$$\pi(\mathbf{x}) = \begin{cases} e^{-\beta}/Z, & \text{if } \mathbf{x} \text{ has an odd number of 1's;} \\ 1/Z, & \text{if } \mathbf{x} \text{ has an even number of 1's;} \end{cases}$$

where the normalizing factor is $Z = 2^{n-1}(1 + e^{-\beta})$.

Even if this game looks similar to the OR game, it exhibits a different behavior. Theorem 4.5.1 gives the stationary expected social welfare of the XOR game and we can see that, as β increases, the expected social welfare tends from below to the social welfare at the Nash equilibria. In contrast the expected social welfare of the OR game is better than the worst Nash equilibrium for all values of β. Moreover, in Theorem 4.5.2 and Theorem 4.5.3 we show that the mixing time for the XOR game is polynomial in n and exponential in β, whereas the mixing time for the OR game can be bounded independently from β.

Theorem 4.5.1 (Expected social welfare)**.** *The stationary expected social welfare of the logit dynamics for the XOR game is* $\mathbf{E}_\pi[W] = -\frac{n}{1+e^\beta}$.

Proof. The expected social welfare is

$$\mathbf{E}_\pi[W] = \sum_{\mathbf{x} \in \{0,1\}^n} W(\mathbf{x})\pi(\mathbf{x}) = -n \cdot \frac{2^{n-1}e^{-\beta}}{2^{n-1}(1+e^{-\beta})} = -\frac{n}{1+e^\beta}.$$

□

The next theorem shows that the mixing time is exponential in β for every $\beta > 0$.

Theorem 4.5.2 (Lower bound on mixing time)**.** *The mixing time of the logit dynamics for the XOR game is* $\Omega(e^\beta)$.

Proof. Consider the set $S \subseteq \{0,1\}^n$ containing only the state $\mathbf{0} = (0, \dots, 0)$. Observe that $\pi(\mathbf{0}) \leqslant 1/2$. The bottleneck ratio is

$$B(\mathbf{0}) = \frac{1}{\pi(\mathbf{0})} \sum_{\mathbf{y} \in \{0,1\}^n} \pi(\mathbf{0})P(\mathbf{0}, \mathbf{y}) = \sum_{\mathbf{y} \in \{0,1\}^n \,:\, |\mathbf{y}|_1 = 1} P(\mathbf{0}, \mathbf{y}) = n \cdot \frac{1}{n} \cdot \frac{1}{1+e^\beta}\,.$$

Hence, by applying Theorem 1.3.7, the mixing time is

$$t_{\text{mix}} \geqslant \frac{1}{B(\mathbf{0})} = 1 + e^\beta\,.$$

□

Finally, in the next theorem we give an almost matching upper bound to the mixing time.

Theorem 4.5.3 (Upper bound on mixing time)**.** *The mixing time of the logit dynamics for the XOR game is* $\mathcal{O}(n^3 e^\beta)$.

The theorem is proved using coupling (see Theorem 1.3.1) and proof is presented in the next sections. Specifically, we use the coupling described in Section 4.1.1; in Section 4.5.1 we show that if the coupled chains are at even distance then distance does not increase after one step of the coupling; in Section 4.5.2 we show that if the coupled chains are at odd distance then they get closer distance with probability independent from β; finally, in Section 4.5.3 we bound the expected time needed by the two chains to coalesce and use Theorem 1.3.1 to derive an upper bound for the mixing time.

4.5.1 Even Hamming distance

Let X_t and Y_t be two chains coupled as described in Section 4.1.1. Suppose that $X_t = \mathbf{x}$, $Y_t = \mathbf{y}$, and $d(\mathbf{x}, \mathbf{y}) = 2\ell$, for $\ell > 0$. In this case, $u_i(\mathbf{x}) = u_i(\mathbf{y}) = b$ for all $i \in [n]$ and some $b \in \{-1, 0\}$.

Let i be the index selected for update and let us distinguish two cases. In the first case $x_i = y_i$ and we have

$$u_i(\mathbf{x}_{-i}, 0) = u_i(\mathbf{y}_{-i}, 0) \quad \text{and} \quad u_i(\mathbf{x}_{-i}, 1) = u_i(\mathbf{y}_{-i}, 1)$$

and thus

$$\sigma_i(0 \mid \mathbf{x}) = \sigma_i(0 \mid \mathbf{y}) \quad \text{and} \quad \sigma_i(1 \mid \mathbf{x}) = \sigma_i(1 \mid \mathbf{y}).$$

Therefore the coupling always update the strategy of player i in the same way in the two chains and thus $d(X_{t+1}, Y_{t+1}) = 2\ell$.

In the second case we have $x_i \neq y_i$ and we assume, without loss of generality, that $x_i = 0$ and $y_i = 1$. We observe that, for $b \in \{-1, 0\}$,

$$u_i(\mathbf{x}_{-i}, 0) = u_i(\mathbf{y}_{-i}, 1) = b \quad \text{and} \quad u_i(\mathbf{y}_{-i}, 0) = u_i(\mathbf{x}_{-i}, 1) = -(1 + b).$$

Therefore we have

$$\sigma_i(0 \mid \mathbf{x}) = \sigma_i(1 \mid \mathbf{y}) = \frac{1}{1 + e^{-(1+2b)\beta}} \quad \text{and} \quad \sigma_i(1 \mid \mathbf{x}) = \sigma_i(0 \mid \mathbf{y}) = \frac{1}{1 + e^{(1+2b)\beta}}$$

and thus we have three possible updates for the strategy of player i:

1. both chains update to 0 (and thus $d(X_{t+1}, Y_{t+1}) = 2\ell - 1$) with probability

$$\min\left\{\frac{1}{1 + e^{(1+2b)\beta}}, \frac{1}{1 + e^{-(1+2b)\beta}}\right\} = \frac{1}{1 + e^{\beta}};$$

2. both chains update to 1 (and thus $d(X_{t+1}, Y_{t+1}) = 2\ell - 1$) with probability

$$\min\left\{\frac{1}{1 + e^{(1+2b)\beta}}, \frac{1}{1 + e^{-(1+2b)\beta}}\right\} = \frac{1}{1 + e^{\beta}};$$

3. chain X and Y choose two different strategies for updating the strategy of player i (and thus $d(X_{t+1}, Y_{t+1}) = 2\ell$) with probability

$$1 - \frac{2}{1 + e^{\beta}}.$$

The following lemma summarizes the above observations.

Lemma 4.5.4. *Suppose that $d(X_t, Y_t) = 2\ell$, for $\ell > 0$. Then*

$$d(X_{t+1}, Y_{t+1}) = \begin{cases} 2\ell - 1, & \text{with probability } \frac{2\ell}{n} \cdot \frac{2}{1+e^{\beta}}; \\ 2\ell, & \text{with probability } 1 - \frac{2\ell}{n} \cdot \frac{2}{1+e^{\beta}}. \end{cases}$$

4.5.2 Odd Hamming distance

Let X_t and Y_t be two chains coupled as described in Section 4.1.1. Suppose that $X_t = \mathbf{x}$, $Y_t = \mathbf{y}$, and $d(\mathbf{x}, \mathbf{y}) = 2\ell - 1$, for $\ell > 0$. In this case we have $u_i(\mathbf{x}) = b$ and $u_i(\mathbf{y}) = -(1+b)$ for some $b \in \{-1, 0\}$. Let i be the index selected for update and let us distinguish two cases.

In the case in which $x_i = y_i = c$ for some $c \in \{0, 1\}$, we have

$$u_i(\mathbf{x}_{-i}, c) = u_i(\mathbf{y}_{-i}, 1-c) = b \quad \text{and} \quad u_i(\mathbf{x}_{-i}, 1-c) = u_i(\mathbf{y}_{-i}, c) = -(1+b).$$

Therefore

$$\sigma_i(c \mid \mathbf{x}) = \sigma_i(1-c \mid y) = \frac{1}{1+e^{-(1+2b)\beta}} \quad \text{and} \quad \sigma_i(1-c \mid x) = \sigma_i(c \mid y) = \frac{1}{1+e^{(1+2b)\beta}}$$

and thus we have three possible updates:

1. both chains update to c (and thus $d(X_{t+1}, Y_{t+1}) = 2\ell - 1$) with probability

$$\min\left\{\frac{1}{1+e^{-(1+2b)\beta}}, \frac{1}{1+e^{(1+2b)\beta}}\right\} = \frac{1}{1+e^{\beta}};$$

2. both chains update to $1-c$ (and thus $d(X_{t+1}, Y_{t+1}) = 2\ell - 1$) with probability

$$\min\left\{\frac{1}{1+e^{-(1+2b)\beta}}, \frac{1}{1+e^{(1+2b)\beta}}\right\} = \frac{1}{1+e^{\beta}};$$

3. chains X and Y choose two different strategies for updating the strategy player i (and thus $d(X_{t+1}, Y_{t+1}) = 2\ell$) with probability $1 - \frac{2}{1+e^{\beta}}$.

In the second case we have $x_i \neq y_i$ and we assume, without loss of generality, that $x_i = 0$ and $y_i = 1$. We observe that

$$u_i(\mathbf{x}_{-i}, 0) = u_i(\mathbf{y}_{-i}, 0) = b \quad \text{and} \quad u_i(\mathbf{x}_{-i}, 1) = u_i(\mathbf{y}_{-i}, 1) = -(1+b).$$

Therefore we have

$$\sigma_i(0 \mid \mathbf{x}) = \sigma_i(0 \mid \mathbf{y}) \quad \text{and} \quad \sigma_i(1 \mid \mathbf{x}) = \sigma_i(1 \mid \mathbf{y})$$

and thus in this case $d(X_{t+1}, Y_{t+1}) = 2\ell - 2$.

The following lemma summarizes the above observations.

Lemma 4.5.5. *Suppose that $d(X_t, Y_t) = 2\ell - 1$, for $\ell > 0$. Then*

$$d(X_{t+1}, Y_{t+1}) = \begin{cases} 2\ell - 2, & \textit{with probability } \frac{2\ell-1}{n}; \\ 2\ell - 1, & \textit{with probability } \frac{n-2\ell+1}{n} \frac{2}{1+e^{\beta}}; \\ 2\ell, & \textit{with probability } \frac{n-2\ell+1}{n}\left(1 - \frac{2}{1+e^{\beta}}\right). \end{cases}$$

4.5.3 Time to coalesce

We denote with τ_k the random variable indicating the first time at which the two coupled chains have distance k. More precisely,

$$\tau_k = \min\{t \colon d(X_t, Y_t) = k\}\,.$$

Therefore, $\tau_{\text{couple}} = \tau_0$ is the time needed for the two chains to coalesce. We next give a bound on the expected time $\mathbf{E}_{\mathbf{x},\mathbf{y}}[\tau_{\text{couple}}]$ for the two chains to coalesce starting from $\mathbf{x}$ and $\mathbf{y}$. If $\mathbf{x}$ and $\mathbf{y}$ have distance 2ℓ, we denote by μ_ℓ the expected time to reach distance $2\ell - 2$. That is,

$$\mu_\ell = \mathbf{E}_{\mathbf{x},\mathbf{y}}[\tau_{2\ell-2}] \, .$$

Similarly, if $\mathbf{x}$ and $\mathbf{y}$ have distance $2\ell - 1$, we denote by ν_ℓ the expected time to reach distance $2\ell - 2$. That is,

$$\nu_\ell = \mathbf{E}_{\mathbf{x},\mathbf{y}}[\tau_{2\ell-2}] \, .$$

Notice that, if $d(\mathbf{x},\mathbf{y}) = d(\mathbf{x}',\mathbf{y}')$ then

$$\mathbf{E}_{\mathbf{x},\mathbf{y}}[\tau_k] = \mathbf{E}_{\mathbf{x}',\mathbf{y}'}[\tau_k]$$

for all k, and thus the μ_ℓ and ν_ℓ are well defined.

From Lemma 4.5.4 and Lemma 4.5.5, we have the following relations

$$\begin{aligned}
\mu_\ell &= 1 + \mu_\ell \cdot \left(1 - \frac{2\ell}{n} \cdot \frac{2}{1+e^\beta}\right) + \nu_\ell \cdot \frac{2\ell}{n} \cdot \frac{2}{1+e^\beta} \\
\nu_\ell &= 1 + \nu_\ell \cdot \frac{n-2\ell+1}{n} \cdot \frac{2}{1+e^\beta} + \mu_\ell \cdot \frac{n-2\ell+1}{n} \cdot \left(1 - \frac{2}{1+e^\beta}\right) .
\end{aligned}$$

Simple algebraic manipulations give

$$\nu_\ell = \frac{n}{2\ell-1}\left(1 + \frac{n-2\ell+1}{2\ell} \cdot \frac{e^\beta - 1}{2}\right)$$

and

$$\begin{aligned}
\mu_\ell &= \nu_\ell + \frac{n}{2\ell} \cdot \frac{1+e^\beta}{2} \\
&= \frac{n}{2\ell-1} + \frac{n}{2\ell}\left(\frac{n}{2\ell-1} \cdot \frac{e^\beta-1}{2} + 1\right) \\
&\leqslant \frac{n}{2\ell-1}\left(\frac{n}{2\ell-1} \cdot \frac{e^\beta-1}{2} + 2\right) \\
&\leqslant n\left(n \cdot \frac{e^\beta-1}{2} + 2\right) .
\end{aligned}$$

Hence,

$$\mathbf{E}_{\mathbf{x},\mathbf{y}}[\tau_{\text{couple}}] \leqslant 1 + \sum_{\substack{2 \leqslant \ell \leqslant n \\ \ell \text{ even}}} \mu_\ell \leqslant \frac{n^2}{2}\left(n \cdot \frac{e^\beta-1}{2} + 2\right) + 1 = \mathcal{O}\left(n^3 e^\beta\right) .$$

From Markov inequality we have that

$$\mathbf{P}_{\mathbf{x},\mathbf{y}}(\tau_{\text{couple}} > t) \leqslant \frac{\mathbf{E}_{\mathbf{x},\mathbf{y}}[\tau_{\text{couple}}]}{t}$$

and thus, by taking $t_0 = 4\mathbf{E}_{\mathbf{x},\mathbf{y}}[\tau_{\text{couple}}]$, we have $d(t_0) \leqslant 1/4$. Therefore, by using Theorem 1.3.1, we have that

$$t_{\text{mix}} = \mathcal{O}\left(n^3 e^\beta\right) .$$

4.6 Conclusions and open problems

The games analyzed in this chapter highlighted some interesting features of the logit dynamics and of the logit equilibrium.

The most important evidence that rises up from these results is that there exists a separation between games that can be upper bounded by a function independent of β, such as CK game and OR game, and games where the mixing time is necessarily exponential in the noise parameter β, such as coordination games or the XOR game. In the next chapter we will focus on this aspect, trying to characterize the class of games whose mixing time is independent of β and looking for the parameters that influence the mixing time.

Since our aim is to consider and characterize wide classes of games, in the next chapter we have to abandon the analysis of the performance, because the relation between stationary expected social welfare and the social welfare at Nash equilibrium strongly depends on the structure of the underlying game.

Moreover, we noted that the analysis of the mixing time is often far from trivial, even for very simple games: this calls for the use of new tools in order to be able to deal with more complex games.

Another relevant aspect emerged in this chapter is that games that obtains the worst Price of Anarchy bound, as the CK game behave very well under the logit dynamics: this aspect also requires further investigations, not pursued in this work.

Chapter 5

Convergence to equilibrium of logit dynamics

To validate the proposal of the logit dynamics as a model for the evolution of games in which agents have limited knowledge, it is necessary to bound how long the logit dynamics takes to converge to the stationary distribution. In previous chapter we showed that the *mixing time* of the logit dynamics can vary a lot (from linear to exponential): thus, it is natural to ask the following questions: (1) How do the *rationality level* and the *structure* of the game affect the mixing time? (2) Can the mixing time grow *arbitrarily*?

In order to answer above questions, we give general bounds on the mixing time for wide classes of games[1]. Specifically, we prove in Section 5.2 that for all *potential games* the mixing time of the logit dynamics is upper-bounded by a *polynomial* in the number of players and by an *exponential* in the rationality level and in some structural properties of the game. However, for very small values of β the mixing time is always polynomial in the number of players.

We complement the upper bound with a lower bound showing that there exist potential games with mixing time exponential in the rationality level. Thus the mixing time can grow indefinitely in potential games as β increases. In Section 5.4 we also study a special class of potential games, the *graphical coordination games*: we extend the result given in Theorem 3.2.2; then, we give a more careful look at two extreme and well-studied cases, the clique and the ring.

Going to the second question, in Section 5.3 we show that for games with dominant strategies (not necessarily potential games) the mixing time cannot exceed some *absolute bound* T which depends uniquely on the number of players n and on the number of strategies m. Though $T = T(n, m)$ is of the form $\mathcal{O}(m^n)$, it is independent of the rationality level and we show that, in general, such an exponential growth is the best possible.

Our results suggest that the structural properties of the game are important for the mixing time. For high β, players tend to play best response and for those games that have more than one pure Nash equilibrium (PNE) with similar potential the system is likely to remain in a PNE for a long time, whereas the stationary distribution gives each PNE approximately the same weight. This happens for (certain) potential games, whence the exponential growth of mixing time with respect to the rationality level. On the contrary, for games with dominant strategies there is a PNE (a dominant profile) with high stationary probability *and* players are guaranteed to play that profile with non-vanishing probability (regardless of the rationality level).

Before showing our results we describe in Section 5.1 the techniques that will be used for obtaining our bounds.

[1]Most of the results in this chapter already appeared in [3].

5.1 Proof techniques

5.1.1 For upper bounds

To derive our upper bounds, we employ two techniques: *Markov chain coupling* and *spectral techniques* Both are well-established techniques for bounding the mixing time and they are summarized in Section 1.3. Here, we give further details about their application in our results.

The coupling. Let $\mathcal{G}$ be a n-player game with profile space S and P be the transition matrix of the logit dynamics for $\mathcal{G}$. We describe, for each $\mathbf{x}, \mathbf{y} \in S$, a coupling of $P(\mathbf{x}, \cdot)$ and $P(\mathbf{y}, \cdot)$.

For each player i, we partition two copies of the interval $[0, 1]$, called $I_{X,i}$ and $I_{Y,i}$, in sub-intervals each labeled with a strategy from the set $S_i = \{z_1, \ldots, z_{|S_i|}\}$ of strategies of player i. The sub-intervals are constructed as follows. For $k = 1, \ldots, |S_i|$, we take the leftmost not yet labeled interval of length $l_k = \min\{\sigma_i(z_k \mid \mathbf{x}), \sigma_i(z_k \mid \mathbf{y})\}$ of both $I_{X,i}$ and $I_{Y,i}$ and label it with strategy z_k. In addition, we take the rightmost non yet labeled interval of length $\sigma_i(z_k \mid \mathbf{x}) - l_k$ of I_X and the rightmost non yet labeled interval of length $\sigma_i(z_k \mid \mathbf{y}) - l_k$ of I_Y and label both with z_k. Notice that at least one of these two intervals has length 0. Define functions $h_{X,i} \colon I_{X,i} \to S_i$ and $h_{Y,i} \colon I_{Y,i} \to S_i$ that for $s \in [0, 1]$ return the labels $h_{X,i}(s)$ and $h_{Y,i}(s)$ of the sub-intervals containing s. Observe that there is a point $\ell \in (0, 1)$ such that $h_X(s) = h_Y(s)$ for every $s \leqslant \ell$ and $h_X(s) \neq h_Y(s)$ for every $s \leqslant \ell$.

Given the above partitions of $I_{X,i}$ and $I_{Y,i}$ for each i, the coupling can be described as follows: pick $i \in [n]$ and $U \in [0, 1]$ uniformly at random and update X and Y by setting $X_i = h_{X,i}(U)$ and $Y_i = h_{Y,i}(U))$. By construction we have that (X, Y) is a coupling of $P(\mathbf{x}, \cdot)$ and $P(\mathbf{y}, \cdot)$.

We define $\mathcal{H} = (S, E)$ as the *Hamming graph* of the game, whose vertex set is the set of strategy profiles, and two profiles are adjacent if they differ only for the strategy of one player. For the path coupling technique (see Theorem 1.3.3), the coupling described above is applied only to pairs of adjacent starting profiles.

Spectral properties. In order to use Theorems 1.3.5 and 1.3.6 for our bounds we need to show that for the logit dynamics Markov chain the second eigenvalue is larger in absolute value than the last eigenvalue.

Theorem 5.1.1. *Let $\mathcal{G}$ be an n-player potential game with profile space S and let P and π be the transition matrix and the stationary distribution of the logit dynamics for $\mathcal{G}$. Let $1 = \lambda_1 \geqslant \lambda_2 \geqslant \ldots \geqslant \lambda_{|S|}$ be the eigenvalues of P. Then $\lambda_2 \geqslant \left|\lambda_{|S|}\right|$.*

Proof. Assume for sake of contradiction that there exists an eigenvalue $\lambda < 0$ of P. Let f be an eigenfunction of λ. By definition, $f \neq \mathbf{0}$; hence, since $\lambda < 0$, then for every profile $\mathbf{x} \in S$ such that $f(\mathbf{x}) \neq 0$ we have $\mathsf{sign}\left((Pf)(\mathbf{x})\right) \neq \mathsf{sign}\left(f(\mathbf{x})\right)$ and thus

$$\langle Pf, f\rangle_\pi := \sum_{\mathbf{x} \in S} (Pf)(\mathbf{x}) f(\mathbf{x}) < 0\,.$$

For every player i and every strategy sub-profile $\mathbf{z}_{-i}$, we consider the corresponding "single-player" matrix $P^{(i,\mathbf{z}_{-i})}$ defined by

$$P^{(i,\mathbf{z}_{-i})}(\mathbf{x}, \mathbf{y}) := \frac{1}{T_i(\mathbf{z}_{-i})} \begin{cases} e^{\beta u_i(\mathbf{y})}\,, & \text{if } \mathbf{x}_{-i} = \mathbf{y}_{-i} = \mathbf{z}_{-i}\,; \\ 0\,, & \text{otherwise}\,. \end{cases}$$

The transition matrix P is the sum of all such "single-player" matrices:

$$P = \frac{1}{n} \sum_i \sum_{\mathbf{z}_{-i}} P^{(i,\mathbf{z}_{-i})}$$

Let us define
$$S_{i,\mathbf{z}_{-i}} := \{\mathbf{x} \mid \mathbf{x} = (s_i, \mathbf{z}_{-i}) \text{ and } s_i \in S_i\} \,.$$
For any $\mathbf{x}, \mathbf{y} \in S_{i,\mathbf{z}_{-i}}$ we have that
$$\frac{e^{\beta u_i(\mathbf{y})}}{e^{\beta u_i(\mathbf{x})}} = e^{-\beta(\Phi(\mathbf{y})-\Phi(\mathbf{x}))} \quad \text{implies} \quad \frac{e^{\beta u_i(\mathbf{x})}}{e^{-\beta\Phi(\mathbf{x})}} = \frac{e^{\beta u_i(\mathbf{y})}}{e^{-\beta\Phi(\mathbf{y})}}$$
That is, the ratio $r_{i,\mathbf{z}_{-i}} := \frac{e^{\beta u_i(\mathbf{z})}}{e^{-\beta\Phi(\mathbf{z})}}$ is constant over all $\mathbf{z} \in S_{i,\mathbf{z}_{-i}}$. Hence, whenever $P^{(i,\mathbf{z}_{-i})}(\mathbf{x}, \mathbf{y})$ is not zero, it does not depend on $\mathbf{x}$: indeed,
$$P^{(i,\mathbf{z}_{-i})}(\mathbf{x}, \mathbf{y}) = \frac{e^{\beta u_i(\mathbf{y})}}{T_i(\mathbf{z}_{-i})} = \frac{Z \cdot r_{i,\mathbf{z}_{-i}}}{T_i(\mathbf{z}_{-i})}\pi(\mathbf{y}) \,.$$
Let $C_{i,\mathbf{z}_i} = \frac{Z \cdot r_{i,\mathbf{z}_{-i}}}{T_i(\mathbf{z}_{-i})}$, we obtain
$$\langle P^{(i,\mathbf{z}_{-i})} f, f\rangle_\pi = C_{i,\mathbf{z}_{-i}} \sum_{\mathbf{x}\in S_{i,\mathbf{z}_{-i}}} \sum_{\mathbf{y}\in S_{i,\mathbf{z}_{-i}}} \pi(\mathbf{x})\pi(\mathbf{y}) f(\mathbf{x}) f(\mathbf{y}) = C_{i,\mathbf{z}_{-i}} \left(\sum_{\mathbf{x}\in S_{i,\mathbf{z}_{-i}}} \pi(\mathbf{x}) f(\mathbf{x}) \right)^2 \geqslant 0 \,.$$
From the linearity of the inner product, it follows that
$$\langle Pf, f\rangle_\pi = \frac{1}{n} \sum_i \sum_{\mathbf{z}_{-i}} \langle P^{(i,\mathbf{z}_{-i})} f, f\rangle_\pi \geqslant 0 \,,$$
contradicting the hypothesis. □

For a function $\Phi\colon S \to \mathbb{R}$ over a finite set S, let us name $\Delta\Phi$ the difference between the maximum and minimum values of Φ and L its Lipschitz constant, i.e.
$$\begin{aligned} \Delta\Phi = \Phi_{\max} - \Phi_{\min} &= \max\{\Phi(\mathbf{x}) - \Phi(\mathbf{y})\colon \mathbf{x}, \mathbf{y} \in S\} \\ L &= \max\{\Phi(\mathbf{x}) - \Phi(\mathbf{y})\colon d(\mathbf{x}, \mathbf{y}) = 1\} \,. \end{aligned} \tag{5.1}$$
Moreover, let $\mathcal{G}$ be a n-player potential game with profile space S and let P and π be the transition matrix and the stationary distribution of the logit dynamics for $\mathcal{G}$. Then, for every pair of profiles $\mathbf{x}, \mathbf{y} \in S$ we set $\perp_{\mathbf{x},\mathbf{y}} = \arg\min\{\pi(\mathbf{x}), \pi(\mathbf{y})\}$ and $\top_{\mathbf{x},\mathbf{y}} = \arg\max\{\pi(\mathbf{x}), \pi(\mathbf{y})\}$. The following is an easy corollary of Lemma 1.3.6.

Corollary 5.1.2. *Let $\mathcal{G}$ be a n-player potential game with profile space S and let P and π be the transition matrix and the stationary distribution of the logit dynamics for $\mathcal{G}$. For every pair of profiles $\mathbf{x}, \mathbf{y}$ we assign a path $\Gamma_{\mathbf{x},\mathbf{y}}$ on the Hamming graph with vertex set S. Then*
$$t_{\text{rel}} \leqslant mne^{\beta L} \max_{\substack{\mathbf{z},\mathbf{w}\colon \\ \mathbf{z}\sim\mathbf{w}}} \frac{1}{\pi(\perp_{\mathbf{z},\mathbf{w}})} \sum_{\substack{\mathbf{x},\mathbf{y}\colon \\ (\mathbf{z},\mathbf{w})\in\Gamma_{\mathbf{x},\mathbf{y}}}} \pi(\mathbf{x})\pi(\mathbf{y})|\Gamma_{\mathbf{x},\mathbf{y}}| \,.$$
where $m = \max_i\{|S_i|\}$, i.e., the maximum number of strategies available for a player. Moreover, if every player has at most two strategies we have
$$t_{\text{rel}} \leqslant 2n \max_{\substack{\mathbf{z},\mathbf{w}\colon \\ \mathbf{z}\sim\mathbf{w}}} \frac{1}{\pi(\perp_{\mathbf{z},\mathbf{w}})} \sum_{\substack{\mathbf{x},\mathbf{y}\colon \\ (\mathbf{z},\mathbf{w})\in\Gamma_{\mathbf{x},\mathbf{y}}}} \pi(\mathbf{x})\pi(\mathbf{y})|\Gamma_{\mathbf{x},\mathbf{y}}| \,.$$

Proof. From Theorem 5.1.1, follows that $t_{\text{rel}} = \frac{1}{1-\lambda_2}$. Moreover, by reversibility of P, we have

$$Q(\mathbf{z},\mathbf{w}) = \pi(\bot_{\mathbf{z},\mathbf{w}})P(\bot_{\mathbf{z},\mathbf{w}},\top_{\mathbf{z},\mathbf{w}}) \geqslant \pi(\bot_{\mathbf{z},\mathbf{w}})\frac{e^{-\beta\Phi(\top_{\mathbf{z},\mathbf{w}})}}{mne^{-\beta\min_{\mathbf{y}\sim\top_{\mathbf{z},\mathbf{w}}}\Phi(\mathbf{y})}} \geqslant \frac{\pi(\bot_{\mathbf{z},\mathbf{w}})}{mne^{\beta L}}\,.$$

Similarly, if every player has at most two strategies we have

$$Q(\mathbf{z},\mathbf{w}) = \pi(\bot_{\mathbf{z},\mathbf{w}})P(\bot_{\mathbf{z},\mathbf{w}},\top_{\mathbf{z},\mathbf{w}}) \geqslant \frac{\pi(\bot_{\mathbf{z},\mathbf{w}})}{2n}\,.$$

Thus, the corollary follows from Lemma 1.3.6. □

5.1.2 For lower bounds

To derive our lower bounds we will use the the *Bottleneck Ratio Theorem* (see Theorem 1.3.7) and a refinement of it for the logit dynamics of potential games (see Theorem 5.1.3 below).

Let $\mathbf{x} \in S$ be a profile of a potential game and let $M \subseteq S \setminus \{\mathbf{x}\}$ be a set of profiles different from $\mathbf{x}$. We define $R_{\mathbf{x},M}$ as the set of profiles in the connected component of the *Hamming graph* with vertex set $S \setminus M$ that contains $\mathbf{x}$ and define

$$\partial R_{\mathbf{x},M} := \{\mathbf{y} \in R_{\mathbf{x},M} : \mathbf{y} \sim \mathbf{z} \text{ for some } \mathbf{z} \in M\}\,.$$

In other words, $\partial R_{\mathbf{x},M}$ consists exactly of those profiles in $R_{\mathbf{x},M}$ that have a neighbor in M. We have the following theorem.

Theorem 5.1.3. *For any potential game $\mathcal{G}$, for any profile $\mathbf{x} \in S$ and for any $M \subset S \setminus \{\mathbf{x}\}$, if $R = R_{\mathbf{x},M}$ satisfies $\pi(R) \leqslant 1/2$ then the mixing time of the logit dynamics with rationality parameter β for $\mathcal{G}$ satisfies*

$$t_{\text{mix}} = \Omega\left(\frac{e^{\beta(\Phi^M - \Phi^R)}}{(m-1)\,|\partial R|}\right),$$

where Φ^R and Φ^M are the minimum potential among profiles in R and M, respectively.

Proof. Observe that for every pair $\mathbf{y},\mathbf{z}$ of profiles that differs only in the strategy of player j, it holds that

$$\pi(\mathbf{y})P(\mathbf{y},\mathbf{z}) = \frac{e^{-\beta\Phi(\mathbf{y})}}{Z}\cdot\frac{1}{n}\cdot\frac{e^{-\beta\Phi(\mathbf{z})}}{e^{-\beta\Phi(\mathbf{y})}+\sum_{s\in S_j}e^{-\beta\Phi(\mathbf{y}_{-j},s)}} \leqslant \frac{e^{-\beta\Phi(\mathbf{z})}}{nZ}\,.$$

Note that for every $\mathbf{y} \in \partial R$ there are at most $(m-1)n$ neighbors outside R and all of them belong to M by definition, thus

$$\begin{aligned} Q\left(R,\overline{R}\right) &= \sum_{\substack{\mathbf{y}\in R\\ \mathbf{z}\in\overline{R}}}\pi(\mathbf{y})P(\mathbf{y},\mathbf{z}) = \sum_{\substack{\mathbf{y}\in\partial R\\ \mathbf{z}\in M}}\pi(\mathbf{y})P(\mathbf{y},\mathbf{z}) \\ &\leqslant \sum_{\substack{\mathbf{y}\in\partial R\\ \mathbf{z}\in M}}\frac{e^{-\beta\Phi(\mathbf{z})}}{nZ} \leqslant (m-1)\,|\partial R|\,\frac{e^{-\beta\Phi^M}}{Z}\,. \end{aligned}$$

Let $\mathbf{x}^+ \in R$ be a profile with the highest potential in R; that is, $\Phi(\mathbf{x}^+) = \Phi^R$. Obviously

$$\pi(R) \geqslant \pi(\mathbf{x}^+) = \frac{e^{-\beta\Phi^R}}{Z}.$$

These two inequalities yield

$$\frac{Q\left(R,\overline{R}\right)}{\pi(R)} \leqslant \frac{(m-1)\,|\partial R|}{e^{\beta(\Phi^M-\Phi^R)}}$$

and since $\pi(R) \leqslant 1/2$ the thesis follows from the Bottleneck Ratio Theorem (Theorem 1.3.7). □

The above theorem gives good lower bounds when we choose $\mathbf{x}$ and M such that all profiles in M have low potential, the resulting set $R = R_{\mathbf{x},M}$ contains at least one profile of high potential (and thus $\Phi^R - \Phi^M$ is large) and the boundary ∂R is small.

5.2 Potential games

We will start by giving a bound that holds for each value of β. Then we show a slightly better bound for low values of β. Finally, we give the exact characterization of the mixing time for high value of β.

5.2.1 For every β

In this section we shall see that it is possible to give upper bounds on the mixing time of the logit dynamics for potential games depending only on the two quantities L and $\Delta\Phi$ defined in 5.1 that holds for every $\beta > 0$. Moreover we will show that such bounds are nearly tight by providing examples of games whose logit dynamics mixing time is close to the given upper bound.

Upper bound. In order to give the upper bound on the mixing time, we first give an upper bound on the relaxation time and then use Theorem 1.3.4.

In the proof of Theorem 5.2.2 we obtain the upper bound on the relaxation time by comparing the logit dynamics with inverse noise β for a potential game $\mathcal{G}$ and the logit dynamics with inverse noise 0 for the same game. When the inverse noise is zero, the logit dynamics is a random walk on a generalized hypercube. Next lemma evaluates the relaxation time of such a chain. The proof is a simple generalization of the proof for the relaxation time of the lazy random walk on the hypercube and is omitted.

Lemma 5.2.1. *Let $\mathcal{G}$ be an n-player game. The relaxation time of the logit dynamics with rationality level $\beta = 0$ for $\mathcal{G}$ is $t_{\mathrm{rel}} = n$.*

The following theorem is the main result of this section.

Theorem 5.2.2. *Let $\mathcal{G}$ be a n-player potential game with potential function Φ and profile space S. The relaxation time of the logit dynamics for $\mathcal{G}$ with rationality level β is $t_{\mathrm{rel}} = \mathcal{O}\left(n \cdot e^{\beta(\Delta\Phi + L)}\right)$.*

Proof. Remember that for all profiles $\mathbf{x} \in S$, the stationary distribution is

$$\pi_\beta(\mathbf{x}) = \frac{e^{-\beta\Phi(\mathbf{x})}}{Z_\beta} \leqslant \frac{e^{-\beta\Phi_{\min}}}{Z_\beta},$$

where $Z_\beta = \sum_{\mathbf{y}\in S} e^{-\beta\Phi(\mathbf{y})}$ is the partition function. As for the edge-stationary distribution, for two adjacent profiles $\mathbf{x} \sim \mathbf{y}$ that differ at player $i \in [n]$ we have

$$Q_\beta(\mathbf{x},\mathbf{y}) = \frac{e^{-\beta\Phi(\mathbf{x})}}{Z_\beta}\frac{1}{n}\frac{e^{-\beta\Phi(\mathbf{y})}}{\sum_{z\in S_i} e^{-\beta\Phi(\mathbf{x}_{-i},z)}} \geqslant \frac{e^{-\beta\Phi_{\max}}}{Z_\beta}\frac{1}{n}\frac{1}{|S_i|\cdot e^{\beta L}}, \tag{5.2}$$

where we used that

$$\frac{e^{-\beta\Phi(\mathbf{y})}}{\sum_{z\in S_i} e^{-\beta\Phi(\mathbf{x}_{-i},z)}} = \frac{1}{\sum_{z\in S_i} e^{\beta[\Phi(\mathbf{y})-\Phi(\mathbf{x}_{-i},z)]}} \geqslant \frac{1}{|S_i|\cdot e^{\beta L}}.$$

Hence, for all $\mathbf{x}, \mathbf{y} \in S$ it holds that

$$\begin{aligned} \pi_\beta(\mathbf{x}) &\leqslant \frac{Z_0}{Z_\beta} e^{-\beta\Phi_{\min}} \pi_0(\mathbf{x}) \,, \\ Q_\beta(\mathbf{x}, \mathbf{y}) &\geqslant \frac{Z_0}{Z_\beta} \frac{e^{-\beta\Phi_{\max}}}{e^{\beta L}} Q_0(\mathbf{x}, \mathbf{y}) \,. \end{aligned}$$

Since from Lemma 5.2.1 it holds that for $\beta = 0$ the relaxation time is $\mathcal{O}(n)$, the thesis follows by applying the comparison theorem (Theorem 1.3.5) and Theorem 5.1.1 with

$$\alpha = \frac{Z_\beta}{Z_0} \frac{e^{\beta L}}{e^{-\beta\Phi_{\max}}} \qquad \text{and} \qquad \gamma = \frac{Z_0}{Z_\beta} e^{-\beta\Phi_{\min}}.$$

□

A slightly better upper bound holds when the players have two strategies.

Corollary 5.2.3. *Let $\mathcal{G}$ be a n-player 2-strategy potential game with potential function Φ and profile space S. The relaxation time of the logit dynamics for $\mathcal{G}$ with rationality level β is* $t_{\mathrm{rel}} = \mathcal{O}\left(n \cdot e^{\beta\Delta\Phi}\right)$.

Proof. Observe that, when every player has two strategies, in Equation (5.2) we have that

$$\frac{e^{-\beta\Phi(\mathbf{x})} e^{-\beta\Phi(\mathbf{y})}}{\sum_{z \in S_i} e^{-\beta\Phi(\mathbf{x}_{-i}, z)}} = \frac{e^{-\beta\Phi(\mathbf{x})} e^{-\beta\Phi(\mathbf{y})}}{e^{-\beta\Phi(\mathbf{x})} + e^{-\beta\Phi(\mathbf{y})}} \geqslant \frac{e^{-\beta \max\{\Phi(\mathbf{x}), \Phi(\mathbf{y})\}}}{2} \,.$$

Hence, we obtain

$$Q_\beta(\mathbf{x}, \mathbf{y}) \geqslant \frac{e^{-\beta\Phi_{\max}}}{Z_\beta} \frac{1}{2n} \,,$$

and we can apply the comparison theorem with

$$\alpha = \frac{Z_\beta}{Z_0} \frac{1}{e^{-\beta\Phi_{\max}}} \qquad \text{and} \qquad \gamma = \frac{Z_0}{Z_\beta} e^{-\beta\Phi_{\min}} \,.$$

□

Finally, we can obtain the bounds on the mixing time by using Theorem 1.3.4 and the fact that $\pi_{\min} \geqslant 1/\left(e^{\beta\Delta\Phi}|S|\right)$.

Corollary 5.2.4. *Let $\mathcal{G}$ be a n-player potential game with potential function Φ and profile space S. The mixing time of the logit dynamics for $\mathcal{G}$ is*

$$t_{\mathrm{mix}} = \mathcal{O}\left(n \cdot e^{\beta(\Delta\Phi + L)} \left(\beta\Delta\Phi + \log|S|\right)\right) \,.$$

Moreover, if every player has at most two strategies, the mixing time is

$$t_{\mathrm{mix}} = \mathcal{O}\left(n \cdot e^{\beta\Delta\Phi}(\beta\Delta\Phi + n)\right) \,.$$

Lower bound. It is easy to find potential games whose logit dynamics mixing time is $\Omega(e^{\beta\Delta\Phi})$ when $\Delta\Phi = L$; e.g., games whose potential function Φ has only two values and at least two non-adjacent maxima (see, for example, XOR game studied in Section 4.5). One naturally wonders whether a similar lower bound can be achieved for games where the Lipschitz constant L is small compared to $\Delta\Phi$. The following theorem shows that the term $e^{\beta\Delta\Phi}$ in the upper bound in Corollary 5.2.4 cannot be essentially improved for L smaller than $\Delta\Phi$.

Theorem 5.2.5. *For every $0 < \delta < 1$ and for every $L = \omega(\log n)$ a family of potential games with two strategies per player exists such that the potential function Φ has Lipschitz constant L, it satisfies $\Delta\Phi/L > n^\delta$ and the mixing time of the logit dynamics is $\Omega\left(e^{(\beta - o(1))\Delta\Phi}\right)$.*

Proof. Consider the game with n players in which every player has strategies 0 and 1, and whose potential function is

$$\Phi(\mathbf{x}) = \Phi(|\mathbf{x}|_1) = -\min\{c; |c - |\mathbf{x}|_1|\} \cdot L\,,$$

where $c = \lceil n^\delta \rceil$. Note that the minimum of the potential is $\Phi(\mathbf{0}) = -\Delta\Phi = -cL$, while the maximum is zero and is attained at all states in the set $M = \{\mathbf{x} \in S \,:\, |\mathbf{x}|_1 = c\}$.

Consider the set $R_{\mathbf{0},M}$ as defined in Section 5.1 and observe that

$$\begin{aligned} R_{\mathbf{0},M} &= \{\mathbf{x} \in S \,:\, |\mathbf{x}|_1 < c\}\,, \\ \partial R_{\mathbf{0},M} &= \{\mathbf{x} \in S \,:\, |\mathbf{x}|_1 = c - 1\}\,. \end{aligned}$$

By the symmetry of the potential function, the stationary probability of $R_{\mathbf{0},M}$ is $\pi(R_{\mathbf{0},M}) \leqslant \frac{1}{2}$ and the size of its boundary is

$$|\partial R_{\mathbf{0},M}| \leqslant \binom{n}{c} \leqslant e^{c \log n} = e^{(\Delta\Phi/L)\log n}.$$

Thus, from Theorem 5.1.3 we have that the mixing time of the logit dynamics is

$$t_{\text{mix}} = \Omega\left(e^{\beta\Delta\Phi - (\Delta\Phi/L)\log n}\right)$$

and the theorem follows. □

5.2.2 For small β

Corollary 5.2.4 shows that, even for small β the mixing time is at most n^2. In this section we will show that for small values of the rationality parameter this bound can be improved to $n \log n$.

Theorem 5.2.6. *Let $\mathcal{G}$ be an n-player potential game with profile space S. If $\beta \leqslant c/(Ln)$, with $c < 1$ constant and L defined in (5.1), then the mixing time of the logit dynamics for $\mathcal{G}$ is $\mathcal{O}(n \log n)$.*

Proof. We will apply the path-coupling technique. (see Theorem 1.3.3): for every $\mathbf{x}, \mathbf{y}$ such that $d(\mathbf{x}, \mathbf{y}) = 1$ we have to find a coupling (X, Y) of the two distributions $P(\mathbf{x}, \cdot)$ and $P(\mathbf{y}, \cdot)$ such that $\mathbf{E}_{\mathbf{x},\mathbf{y}}\left[d(X, Y)\right]$ is as small as possible.

Specifically, if we consider the coupling described in Section 5.1.1, then $d(X, Y)$ is a random variable taking values $0, 1$, and 2: it is 0 if we are updating the player j on which $\mathbf{x}$ and $\mathbf{y}$ differ, it is 1 if we are updating a player on which the two profiles coincide and $U \leqslant \ell$, and it is 2 if we are updating a player on which the two profiles coincide and $U > \ell$. Hence the expected distance between X and Y is

$$\begin{aligned}
\mathbf{E}_{\mathbf{x},\mathbf{y}}\left[d(X,Y)\right] &= \mathbf{P}_{\mathbf{x},\mathbf{y}}\left(d(X,Y)=1\right) + 2\mathbf{P}_{\mathbf{x},\mathbf{y}}\left(d(X,Y)=2\right) \\
&= \left(1-\frac{1}{n}\right)\ell + 2\left(1-\frac{1}{n}\right)(1-\ell) \\
&= \left(1-\frac{1}{n}\right)(2-\ell) \leqslant e^{-1/n}(2-\ell)\,.
\end{aligned}$$

From the coupling construction we have $\ell = \sum_{z \in S_i} \min\{\sigma_i(z \mid \mathbf{x}), \sigma_i(z \mid \mathbf{y})\}$. Observe that for any profile $\mathbf{x}$, any player i and any strategy $z \in S_i$ it holds that

$$\sigma_i(z \mid \mathbf{x}) = \frac{e^{-\beta\Phi(\mathbf{x}_{-i},z)}}{\sum_{k \in S_i} e^{-\beta\Phi(\mathbf{x}_{-i},k)}} = \frac{1}{\sum_{k \in S_i} e^{\beta[\Phi(\mathbf{x}_{-i},z)-\Phi(\mathbf{x}_{-i},k)]}} \geqslant \frac{1}{|S_i| e^{\beta L}}\,.$$

Hence

$$\ell = \sum_{z \in S_i} \min\{\sigma_i(z \mid \mathbf{x}),\, \sigma_i(z \mid \mathbf{y})\} \geqslant \sum_{z \in S_i} \frac{1}{|S_i| e^{\beta L}} = e^{-\beta L} \geqslant e^{-c/n}\,,$$

where in the last inequality we used the hypothesis on β. Thus, the expected distance between X and Y is upper bounded by

$$\begin{aligned}
\mathbf{E}_{\mathbf{x},\mathbf{y}}\left[d(X,Y)\right] &\leqslant e^{-1/n}(2-\ell) \leqslant e^{-1/n}(2-e^{-c/n}) \\
&= e^{-1/n}(1+(1-e^{-c/n})) \leqslant e^{-1/n}(1+c/n) \leqslant e^{-1/n}e^{c/n} = e^{-\frac{1-c}{n}}\,,
\end{aligned}$$

where in the second line we repeatedly used the well-known inequality $1 + x \leqslant e^x$ for every $x > -1$.

The thesis then follows by applying Theorem 1.3.3 with $\alpha = \frac{1-c}{n}$. □

5.2.3 For high β

In this section we give a bound for the mixing time of the logit dynamics for potential games and we show that for high values of β this bound is tight: the bound depends on a structural property of the potential function that measure how difficult is to visit a profile by starting from a profile with higher potential value.

Specifically, consider a n-player game $\mathcal{G}$ with exact potential Φ and profile space S. For every pair of profiles $(\mathbf{x}, \mathbf{y})$, we assume w.l.o.g. that $\Phi(\mathbf{x}) \geqslant \Phi(\mathbf{y})$ and define $\mathcal{P}_{\mathbf{x},\mathbf{y}}$ as the set of paths from $\mathbf{x}$ to $\mathbf{y}$ in the Hamming graph; for every $\Gamma_{\mathbf{x},\mathbf{y}} \in \mathcal{P}_{\mathbf{x},\mathbf{y}}$ let $\Gamma^i_{\mathbf{x},\mathbf{y}}$, for $0 \leqslant i \leqslant |\Gamma_{\mathbf{x},\mathbf{y}}|$, be the i-th profiles on the path $\Gamma_{\mathbf{x},\mathbf{y}}$ between $\mathbf{x}$ and $\mathbf{y}$ (notice that $\Gamma^0_{\mathbf{x},\mathbf{y}} = \mathbf{x}$ and $\Gamma^{|\Gamma_{\mathbf{x},\mathbf{y}}|}_{\mathbf{x},\mathbf{y}} = \mathbf{y}$). For every pair $(\mathbf{x}, \mathbf{y})$ and every path $\Gamma_{\mathbf{x},\mathbf{y}} \in \mathcal{P}_{\mathbf{x},\mathbf{y}}$, we define

$$\zeta(\Gamma_{\mathbf{x},\mathbf{y}}) = \max_{0 \leqslant i \leqslant |\Gamma_{\mathbf{x},\mathbf{y}}|} \left(\Phi(\Gamma^i_{\mathbf{x},\mathbf{y}}) - \Phi(\mathbf{x})\right)\,,$$

i.e., the maximum increase of the potential function along the path $\Gamma_{\mathbf{x},\mathbf{y}}$. Similarly, we have

$$\zeta(\mathbf{x},\mathbf{y}) = \min_{\Gamma_{\mathbf{x},\mathbf{y}} \in \mathcal{P}_{\mathbf{x},\mathbf{y}}} \zeta(\Gamma_{\mathbf{x},\mathbf{y}})\,,$$

i.e., the maximum increase of the potential function that is necessary for going from $\mathbf{x}$ to $\mathbf{y}$. Finally, we set $\zeta^\star = \max_{\mathbf{x},\mathbf{y}} \zeta(\mathbf{x},\mathbf{y})$: notice that $\zeta^\star \geqslant 0$. Then, we will show that the mixing time exponentially depends on β and $\zeta^\star$. We start by bounding the relaxation time of the logit dynamics for a potential game $\mathcal{G}$.

Lemma 5.2.7. *Let $\mathcal{G}$ be an n-player potential game with profile space S and potential function Φ and let $\zeta_\star$ as defined above. The relaxation time of the logit dynamics for $\mathcal{G}$ is*

$$t_{\text{rel}} \leqslant nm^{2n+1} e^{(\zeta^\star + L)\beta} \,.$$

Moreover, if every player has at most two strategies, we have

$$t_{\text{rel}} \leqslant n2^{2n+1} e^{\zeta^\star \beta} \,.$$

Proof. For every pair of profiles $(\mathbf{x}, \mathbf{y})$ we associate a path $\Gamma^\star_{\mathbf{x},\mathbf{y}}$ such that $\zeta(\Gamma^\star_{\mathbf{x},\mathbf{y}}) = \zeta(\mathbf{x}, \mathbf{y})$: note that $|\Gamma^\star_{\mathbf{x},\mathbf{y}}| \leqslant m^n$. For every pair of profiles $\mathbf{z}, \mathbf{w}$ such that $\mathbf{z} \sim \mathbf{w}$, we have

$$\sum_{\substack{(\mathbf{x},\mathbf{y}):\\ (\mathbf{z},\mathbf{w}) \in \Gamma^\star_{\mathbf{x},\mathbf{y}}}} \frac{\pi(\mathbf{x})\pi(\mathbf{y})}{\pi(\perp_{\mathbf{z},\mathbf{w}})} = \sum_{\mathbf{x}} e^{\beta(\Phi(\perp_{\mathbf{z},\mathbf{w}}) - \Phi(\mathbf{x}))} \sum_{\substack{\mathbf{y}:\\ (\mathbf{z},\mathbf{w}) \in \Gamma^\star_{\mathbf{x},\mathbf{y}}}} \pi(\mathbf{y}) \leqslant m^n e^{\zeta^\star \beta} \,. \tag{5.3}$$

Applying Corollary 5.1.2 the lemma follows. $\square$

Hence, for $\beta = \omega\left(\frac{n \log m}{\zeta^\star}\right)$ and $\zeta^\star > 0$ we have that $t_{\text{rel}} \leqslant e^{(\zeta^\star + L)\beta(1+o(1))}$ ($t_{\text{rel}} \leqslant e^{\zeta^\star \beta(1+o(1))}$ if any player has at most two strategies). Then, the bound to the mixing time given in the next theorem directly follows from Theorem 1.3.4 and the fact that $\pi_{\min} \geqslant 1/\left(e^{\beta \Delta\Phi}|S|\right)$.

Theorem 5.2.8. *Let $\mathcal{G}$ be an n-player potential game and at most m strategies for player and let $\zeta^\star$ as defined above. If $\beta = \omega\left(\frac{n \log m}{\zeta^\star}\right)$, the mixing time of the logit dynamics for $\mathcal{G}$ is*

$$t_{\text{mix}} \leqslant e^{(\zeta^\star + L)\beta(1+o(1))} \left(\beta \Delta\Phi + \log|S|\right) \,.$$

Moreover, if any player has at most two strategies, the mixing time is

$$t_{\text{mix}} \leqslant e^{\zeta^\star \beta(1+o(1))} \left(\beta \Delta\Phi + n\right) \,.$$

The next theorem shows that for such high values of β this bound is almost tight.

Theorem 5.2.9. *Let $\mathcal{G}$ be an n-player potential game such that $\zeta^\star > 0$. For β high enough, the mixing time of the logit dynamics for $\mathcal{G}$ is*

$$t_{\text{mix}} \geqslant e^{\zeta^\star \beta(1-o(1))} \,.$$

Proof. Fix a pair $(\mathbf{x}, \mathbf{y})$ and a path $\Gamma^\star_{\mathbf{x},\mathbf{y}}$ such that $\zeta(\Gamma^\star_{\mathbf{x},\mathbf{y}}) = \zeta(\mathbf{x}, \mathbf{y}) = \zeta^\star$ and let $\mathbf{x}_\perp$ be the profile of maximum potential along $\Gamma^\star_{\mathbf{x},\mathbf{y}}$. Let

$$M = \{\mathbf{z} \mid \exists\, \Gamma \in \mathcal{P}_{\mathbf{x},\mathbf{y}} \colon \mathbf{z} = \arg\max_{\mathbf{w} \in \Gamma} \Phi(\mathbf{w})\} \,,$$

i.e., the set of profiles of maximum potential along some path from $\mathbf{x}$ to $\mathbf{y}$: since $\zeta^\star > 0$, then $M \subset S$. Moreover, notice that the Hamming graph over $S \setminus M$ has two disjoint components: in particular, $\mathbf{x}$ and $\mathbf{y}$ are not in the same component and we let $R_{\mathbf{x},M}$ and $R_{\mathbf{y},M}$ be the component that contains $\mathbf{x}$ and $\mathbf{y}$, respectively.

We first consider the case $\Phi(\mathbf{x}) = \Phi(\mathbf{y})$. Then, either $\pi(R_{\mathbf{x},M}) \leqslant 1/2$ or $\pi(R_{\mathbf{y},M}) \leqslant 1/2$: suppose w.l.o.g. that the first is true, then we can apply Theorem 5.1.3 and, since $\Phi^R \leqslant \Phi(\mathbf{x})$, $\Phi^M = \Phi(\mathbf{x}_\perp)$ and $\Phi(\mathbf{x}_\perp) - \Phi(\mathbf{x}) = \zeta^\star$, we obtain

$$t_{\text{mix}} \geqslant e^{\zeta^\star \beta(1-o(1))} \,,$$

for β sufficiently high.

If $\Phi(\mathbf{x}) > \Phi(\mathbf{y})$, then for every profile $\mathbf{z} \in R_{\mathbf{x},M}$, $\Phi(\mathbf{x}) \leqslant \Phi(\mathbf{z})$: indeed, if we suppose this is not true, then we can consider a path $\Gamma^{\star}_{\mathbf{z},\mathbf{y}}$ from $\mathbf{z}$ to $\mathbf{y}$ where both endpoints have potential higher than $\mathbf{x}$; moreover, since M disconnect $\mathbf{z}$ and $\mathbf{y}$, there exists a profile $\mathbf{w} \in M$ such that $\mathbf{w} \in \Gamma^{\star}_{\mathbf{z},\mathbf{y}}$; but, since $\Phi(\mathbf{w}) \leqslant \Phi(\mathbf{x}_{\perp})$, then $\zeta(\Gamma^{\star}_{\mathbf{z},\mathbf{y}}) > \zeta(\Gamma^{\star}_{\mathbf{x},\mathbf{y}}) = \zeta^{\star}$ and this is a contradiction. Hence $\Phi(\mathbf{y}) > \Phi(\mathbf{z})$ for every $\mathbf{z} \in R_{\mathbf{x},M}$ and thus, for β high enough, $\pi(R_{\mathbf{x},M}) \leqslant 1/2$: as before, applying Theorem 5.1.3, we obtain

$$t_{\text{mix}} \geqslant e^{\zeta^{\star}\beta(1-o(1))} .$$

□

5.3 Mixing time independent of β

Theorems 5.2.9 and 5.2.5 show that there are games where the mixing time is necessarily exponential in β and in structural properties of the game. This fact naturally raises the question about the existence and the features of games where the mixing time can be bounded by a function independent of the rationality level β.

The following lemma suggests a possible characterization for these games.

Lemma 5.3.1. *Let $\mathcal{G}$ be an n-player potential games with state space S such that $\zeta^{\star} = 0$. The relaxation time of the logit dynamics for $\mathcal{G}$ is independent from β. Moreover, there are no potential games with $\zeta^{\star} > 0$ such that the relaxation time is independent from β.*

Proof. For every pair of profiles $(\mathbf{x}, \mathbf{y})$ we associate the path $\Gamma_{\mathbf{x},\mathbf{y}}$ such that $\zeta(\Gamma_{\mathbf{x},\mathbf{y}}) = \zeta(\mathbf{x}, \mathbf{y})$ and, for every $i = 0, \ldots, |\Gamma_{\mathbf{x},\mathbf{y}}| - 1$, the profile $\Gamma^{i+1}_{\mathbf{x},\mathbf{y}}$ is obtained from $\Gamma^{i}_{\mathbf{x},\mathbf{y}}$ by updating the strategy of a player j with one of the best responses of j: it is clear that such a path exists whenever $\zeta^{\star} = 0$. We also notice that $|\Gamma_{\mathbf{x},\mathbf{y}}| \leqslant m^n$.

Let consider an edge $(\mathbf{z}, \mathbf{w})$ of one of these paths. Then, we have

$$Q(\mathbf{z}, \mathbf{w}) = \pi(\mathbf{z})P(\mathbf{z}, \mathbf{w}) \geqslant \frac{\pi(\mathbf{z})}{mn} . \tag{5.4}$$

From Theorem 5.1.1 and Lemma 1.3.6, we obtain $t_{\text{rel}} = \frac{1}{1-\lambda_2} \leqslant 2nm^{2n+1}$, where the inequality follows from (5.3) and (5.4).

The second part of the lemma follows directly from Theorem 5.2.9, by observing that the relaxation time is upper bounded by the mixing time, as established in Theorem 1.3.4. □

We wonder if the same characterization given in Lemma 5.3.1 for the relaxation time holds also for the mixing time. Indeed, by simply invoking the Theorem 1.3.4 we do not obtain the desired result, since $\log \pi_{\min}^{-1}$ depends on the rationality parameter[2].

However, we are able to prove that there are classes of games for which the mixing time is independent of β. It is remarkable that the set of games for which we prove this result is not necessarily a subset of the class of potential games.

[2]However, we can conclude that the class of potential games such that $\zeta^{\star} = 0$ does not depend on β exponentially, but at most linearly.

Games with dominant strategies. We prove that, for the class of games with *dominant strategies*, it is possible to give an upper bound to the mixing time that is independent of β. In other words, the mixing time of the logit dynamics for games with dominant strategies does not grow arbitrarily as β tends to infinity.

Let us name 0 a dominant strategy for all players and consider the profile $\mathbf{0} = (0, \ldots, 0)$. It is easy to see that the following observation holds for the logit dynamics of a game with dominant strategies.

Observation 5.3.2. *In every profile and for every β, if player i is selected for update then the logit dynamics updates her strategy to the dominant strategy with probability at least $1/|S_i|$. That is, for all $\mathbf{x}$, β and i, $\sigma_i(0 \mid \mathbf{x}) \geqslant 1/|S_i|$.*

We are now ready to derive an upper bound on the mixing time of the logit dynamics for dominant strategy games.

Theorem 5.3.3. *Let $\mathcal{G}$ be an n-player games with dominant strategies where each player has at most m strategies. The mixing time of the logit dynamics for $\mathcal{G}$ is $t_{\text{mix}} = \mathcal{O}\left(m^n n \log n\right)$.*

Proof. We apply the coupling technique (see Theorem 1.3.1). Let $\{X_t\}$ and $\{Y_t\}$ be two instances of the logit dynamics starting at $\mathbf{x}$ and $\mathbf{y}$ respectively, and consider a coupling with the following properties: at every step the same player in both chains is chosen for the update, the probability that the strategy of the chosen player is updated to the dominant strategy 0 in both chains is at least $1/|S_i| \geqslant 1/m$ (notice that this is possible because of Observation 5.3.2), and once the two chains coalesce they stay coupled for all the following time steps.

We can check that the coupling described in Section 5.1.1 has the properties required above. Indeed, this coupling always select the same player in both chains. Moreover, if player i is selected for update, the probability that both chains choose strategy s for player i is exactly $\min\{\sigma_i(s \mid \mathbf{x}), \sigma_i(s \mid \mathbf{y})\}$: if s is dominant for player i, we have that $\sigma_i(s \mid \mathbf{x}), \sigma_i(s \mid \mathbf{y}) \geqslant 1/|S_i|$ and thus the probability that the coupling updates to s is at least $1/|S_i|$.

Let τ be the first time such that all the players have been selected at least once and let $t^\star = 2n \log n$. Observe that for all starting profiles $\mathbf{z}$ and $\mathbf{w}$, it holds that

$$\mathbf{P}_{\mathbf{z},\mathbf{w}}\left(X_{t^\star} = \mathbf{0} \text{ and } Y_{t^\star} = \mathbf{0} \mid \tau \leqslant t^\star\right) \geqslant \frac{1}{m^n}\,. \tag{5.5}$$

Indeed, given that all players have been selected at least once within time $t^\star$, both chains are in profile $\mathbf{0}$ at time $t^\star$ if and only if every player chose strategy 0 in both chains the last time she played before time $t^\star$. From the construction of the coupling it follows that such event holds with probability at least $1/m^n$.

Hence, for all starting profiles $\mathbf{z}$ and $\mathbf{w}$, we have that

$$\begin{aligned}\mathbf{P}_{\mathbf{z},\mathbf{w}}\left(X_{t^\star} = Y_{t^\star}\right) &\geqslant \mathbf{P}_{\mathbf{z},\mathbf{w}}\left(X_{t^\star} = \mathbf{0} \text{ and } Y_{t^\star} = \mathbf{0}\right) \\ &\geqslant \mathbf{P}_{\mathbf{z},\mathbf{w}}\left(X_{t^\star} = \mathbf{0} \text{ and } Y_{t^\star} = \mathbf{0} \mid \tau \leqslant t^\star\right)\mathbf{P}_{\mathbf{z},\mathbf{w}}\left(\tau \leqslant t^\star\right) \\ &\geqslant \frac{1}{m^n} \cdot \frac{1}{2}\,,\end{aligned} \tag{5.6}$$

where in the last inequality we used (5.5) and the Coupon Collector's argument.

Therefore, by considering k phases each one lasting $t^\star$ time steps, since the bound in (5.6) holds for every starting states of the Markov chain, we have that the probability that the two chains have not yet coupled after $kt^\star$ time steps is

$$\mathbf{P}_{\mathbf{x},\mathbf{y}}\left(X_{kt^\star} \neq Y_{kt^\star}\right) \leqslant \left(1 - \frac{1}{2m^n}\right)^k \leqslant e^{-k/2m^n}\,,$$

which is less than 1/4, for $k = \mathcal{O}(m^n)$. By applying the Coupling Theorem (see Theorem 1.3.1) we have that $t_{\text{mix}} = \mathcal{O}\left(m^n n \log n\right)$. □

In the previous chapter a n-player game with two strategies per player, namely the OR game is shown whose logit dynamics mixing time is $\Omega(2^n)$ for large values of β. We next prove that, for every $m \geqslant 2$, there are n-player games with m strategies per player whose logit dynamics mixing time is $\Omega\left(m^{n-1}\right)$. Thus the m^n factor in the upper bound given by Theorem 5.3.3 cannot be essentially improved.

Theorem 5.3.4. *For every $m \geqslant 2$ and $n \geqslant 2$, there exists a n-player potential game with dominant strategies where each player has m strategies and such that, for sufficiently large β, $t_{\text{mix}} = \Omega\left(m^{n-1}\right)$.*

Proof. Consider the game with n players, each of them having strategies $\{0, \ldots, m-1\}$, such that for every player i:

$$u_i(\mathbf{x}) = \begin{cases} 0, & \text{if } \mathbf{x} = \mathbf{0}; \\ -1, & \text{otherwise.} \end{cases}$$

Note that 0 is a dominant strategy. This is a potential game with potential $\Phi(\mathbf{x}) = -u_i(\mathbf{x})$ and thus the stationary distribution is given by the Gibbs measure in 2.3. We apply the bottleneck ratio technique (see Theorem 1.3.7) with $R = \{0 \ldots, m-1\}^n \setminus \{\mathbf{0}\}$, for which we have

$$\pi(R) = \frac{e^{-\beta}}{Z}(m^n - 1)$$

with $Z = 1 + e^{-\beta}(m^n - 1)$. It is easy to see that $\pi(R) < 1/2$ for $\beta > \log(m^n - 1)$ and furthermore

$$\begin{aligned} Q(R, \overline{R}) &= \sum_{\mathbf{x} \in R} \pi(\mathbf{x}) P(\mathbf{x}, \mathbf{0}) \\ &= \frac{e^{-\beta}}{Z} \sum_{\mathbf{x} \in R} P(\mathbf{x}, \mathbf{0}) = \frac{e^{-\beta}}{Z} \sum_{\mathbf{x} \in R_1} P(\mathbf{x}, \mathbf{0}), \end{aligned}$$

where R_1 is the subset of R containing all states with exactly one non-zero entry. Notice that, for every $\mathbf{x} \in R_1$, we have

$$P(\mathbf{x}, \mathbf{0}) = \frac{1}{n} \cdot \frac{1}{1 + (m-1)e^{-\beta}}.$$

As $|R_1| = n(m-1)$, we have

$$\begin{aligned} Q(R, \overline{R}) &= \frac{e^{-\beta}}{Z} |R_1| \frac{1}{n} \cdot \frac{1}{1 + (m-1)e^{-\beta}} \\ &= \frac{e^{-\beta}}{Z} \cdot \frac{m-1}{1 + (m-1)e^{-\beta}} \end{aligned}$$

whence

$$\begin{aligned} t_{\text{mix}} &\geqslant \frac{1}{4} \cdot \frac{\pi(R)}{Q(R, \overline{R})} \\ &\geqslant \frac{1}{4} \cdot (m^n - 1) \cdot \frac{1 + (m-1)e^{-\beta}}{m-1} > \frac{1}{4} \cdot \frac{m^n - 1}{m-1}. \end{aligned}$$

□

Max-solvable games. Observe that, by using the same techniques exploited in this section, it is possible to prove an upper bound independent of β for *max-solvable* games [102], a class which contains games with dominant strategies as a special case, albeit with an upper bound that is much larger than $\mathcal{O}(m^n n \log n)$.

5.4 Graphical coordination games

In previous chapter we give tight bounds for the mixing time of 2×2 coordination games. In this section we will consider *graphical coordination games*. We notice that this class of games is a subset of the class of potential games: indeed, it is easy to check that, given a graph $G = (V, E)$ the function $\Phi(\mathbf{x}) = \sum_{e\in E} \Phi_e(\mathbf{x})$ is a potential function for the graphical coordination game, where, for every $e = (u, v) \in E$

$$\Phi_e(\mathbf{x}) = \begin{cases} -\Delta & \text{if } x_u = x_v = 0\,; \\ -\delta & \text{if } x_u = x_v = 1\,; \\ 0 & \text{otherwise}\,; \end{cases} \tag{5.7}$$

with Δ and δ as defined in (4.4).

In Section 5.4.1 we will show a bound, based on the work of Berger et al. [12], that holds for every class of games. Then, we focus on the two more studied network topologies: the clique (Section 5.4.2), where the mixing time dependence on $e^{\beta\Delta\Phi}$ showed in Corollary 5.2.4 cannot be improved, and the ring (Section 5.4.3), where a more local interaction implies a faster convergence to the stationary distribution.

In the rest of this section we will assume w.l.o.g. that $\Delta \geqslant \delta$.

5.4.1 For every graph

Theorem 5.4.1. *Let $\mathcal{G}$ be an n-player graphical coordination game on a graph G. The mixing time of the logit dynamics for $\mathcal{G}$ is*

$$t_{\text{mix}} = \mathcal{O}\left(2n^3 e^{\chi(G)(\Delta+\delta)\beta}(n\Delta\beta + 1)\right)\,.$$

where $\chi(G)$ is the cutwidth *of G defined in (3.5).*

Proof. Consider the ordering of vertices of G that obtains the cutwidth. For every $\mathbf{x}, \mathbf{y} \in S$ that differ in the strategies played at vertices $v_1, v_2, \ldots, v_d$, we consider the path $\Gamma_{\mathbf{x},\mathbf{y}} = (\mathbf{x}^0, \mathbf{x}^1, \ldots, \mathbf{x}^d)$, where

$$\mathbf{x}^i = \left(y_1, \ldots, y_{v_{i+1}-1}, x_{v_{i+1}}, \ldots, x_n\right)\,.$$

(Above we assume $v_{d+1} = n + 1$). Notice that $\mathbf{x}^0 = \mathbf{x}$ and $\mathbf{x}^d = \mathbf{y}$ and $|\Gamma_{\mathbf{x},\mathbf{y}}| \leqslant n$. To every edge $\xi = (\mathbf{x}^i, \mathbf{x}^{i+1})$, we consider the function γ_ξ that assigns to every pair of profiles $\mathbf{x}, \mathbf{y}$ such that $\xi \in \Gamma_{\mathbf{x},\mathbf{y}}$, the following new profile

$$\gamma_\xi(\mathbf{x}, \mathbf{y}) = \begin{cases} \left(x_1, \ldots, x_{v_{i+1}-1}, y_{v_{i+1}}, \ldots, y_n\right) & \text{if } \pi(\mathbf{x}^i) \leqslant \pi(\mathbf{x}^{i+1})\,; \\ \left(x_1, \ldots, x_{v_{i+1}}, y_{v_{i+1}+1}, \ldots, y_n\right) & \text{otherwise.} \end{cases}$$

It is easy to see that γ_ξ is an injective function: indeed, since ξ is known, if $\pi(\mathbf{x}^i) \leqslant \pi(\mathbf{x}^{i+1})$, then we can retrieve v_{i+1}, that is the first vertex where $\mathbf{x}^i$ and $\mathbf{x}^{i+1}$ differ and thus, selecting the first $v_{i+1} - 1$ vertices from $\gamma_\xi(\mathbf{x}, \mathbf{y})$ and the remaining ones from $\mathbf{x}^i$ we are able to reconstruct $\mathbf{x}$ and, similarly, selecting the first $v_{i+1} - 1$ vertices from $\mathbf{x}^i$ and the remaining ones from $\gamma_\xi(\mathbf{x}, \mathbf{y})$

we are able to reconstruct $\mathbf{y}$. Similarly, if $\pi(\mathbf{x}^i) > \pi(\mathbf{x}^{i+1})$, we can retrieve v_{i+2} and we can reconstruct $\mathbf{x}$ and $\mathbf{y}$ from $\gamma_\xi(\mathbf{x},\mathbf{y})$ and $\mathbf{x}^{i+1}$.

Let $E^\star = \{(j,k) \in E \colon j < v_{i+1} \text{ and } k \geqslant v_{i+1}\}$: observe that $|E^\star| \leqslant \chi(G)$. For any edge $e = (j,k) \in E^\star$, for every $\mathbf{x},\mathbf{y} \in S$ and for every $\xi = (\mathbf{x}^i, \mathbf{x}^{i+1}) \in \Gamma_{\mathbf{x},\mathbf{y}}$, we distinguish two cases: If $x_j = y_j$ or $x_k = y_k$, it is easy to see that for all available values of x_j, y_j, x_k and y_k

$$\Phi_e(\mathbf{x}) + \Phi_e(\mathbf{y}) - \Phi_e(\perp_{\mathbf{x}^i,\mathbf{x}^{i+1}}) - \Phi_e(\gamma_\xi(\mathbf{x},\mathbf{y})) = 0\,.$$

If $x_j \neq y_j$ and $x_k \neq y_k$, it is easy to see that for all available values of x_j, y_j, x_k and y_k

$$\Phi_e(\mathbf{x}) + \Phi_e(\mathbf{y}) - \Phi_e(\perp_{\mathbf{x}^i,\mathbf{x}^{i+1}}) - \Phi_e(\gamma_\xi(\mathbf{x},\mathbf{y})) = \pm(\Delta + \delta)\,.$$

Thus, we have that for every $\mathbf{x},\mathbf{y} \in S$ and for every $\xi = (\mathbf{x}^i, \mathbf{x}^{i+1}) \in \Gamma_{\mathbf{x},\mathbf{y}}$,

$$\begin{aligned}\Phi(\mathbf{x}) + \Phi(\mathbf{y}) - \Phi(\perp_{\mathbf{x}^i,\mathbf{x}^{i+1}}) - \Phi(\gamma_\xi(\mathbf{x},\mathbf{y})) &= \sum_{e \in E^*} \left(\Phi_e(\mathbf{x}) + \Phi_e(\mathbf{y}) - \Phi_e(\perp_{\mathbf{x}^i,\mathbf{x}^{i+1}}) - \Phi_e(\gamma_\xi(\mathbf{x},\mathbf{y}))\right) \\ &\geqslant -\chi(G)(\Delta+\delta)\,.\end{aligned} \tag{5.8}$$

Applying Corollary 5.1.2, we obtain

$$t_{\text{rel}} \leqslant 2n^2 e^{\chi(G)(\Delta+\delta)\beta} \sum_{\mathbf{x},\mathbf{y}} \pi(\gamma_\xi(\mathbf{x},\mathbf{y})) \leqslant 2n^2 e^{\chi(G)(\Delta+\delta)\beta}\,.$$

where we first inequality follows from (5.8) and the second one from the fact that γ_ξ is injective.

The theorem follows from Theorem 1.3.4 and by observing that

$$\log \pi_{\min}^{-1} \leqslant \beta\Delta\Phi + \log|S| \leqslant |E|\Delta\beta + \log 2^n \leqslant n(n\Delta\beta + 1)\,.$$

□

5.4.2 On the clique

We now focus on graphical coordination games on one of the most studied network topologies: the *clique*, where every player plays the basic coordination game in (3.1) with every other player. Since the cutwidth of a clique is $\Theta(n^2)$, the Theorem 5.4.1 gives an upper bound to the mixing time of the logit dynamics for this class of games that is exponential in $n^2(\Delta+\delta)\beta$. However, since the the game is a potential game we can obtain a slightly better upper bound by using Corollary 5.2.4. Moreover, we will show that the mixing time for graphical coordination games on the clique turns out to be necessarily exponential in n, even for $\beta = \Theta(1)$.

It is not difficult to see that for this class of games, we can rewrite, for every profile $\mathbf{x}$, the potential value of $\mathbf{x}$ as $\Phi(\mathbf{x}) = -\phi(|\mathbf{x}|_0)$, where

$$\phi(k) = (k^\star - k)\left(\frac{2n - k^\star - k - 1}{2}\delta - \frac{k^\star + k - 1}{2}\Delta\right),$$

and $k^\star = \left\lceil (n-1)\frac{\delta}{\Delta+\delta} \right\rceil$. Notice that the maximum of the potential is attained when $k^\star$ players are playing 0 and, since $\phi(k^\star) = 0$, we have that $\Delta\Phi = \max_k \phi(k)$. Moreover, it is easy to check that $\phi(k)$ monotonically decreases as k goes from 0 to $k^\star$ and then monotonically increases as k goes from $k^\star$ to n. Therefore, $\Delta\Phi = \max\{\phi(0), \phi(n)\}$.

Notice that, since $\Delta \geqslant \delta$, then $\phi(k) \leqslant \phi(n-k)$ for $k < k^\star$, $\Delta\Phi = \phi(n)$ and

$$\sum_{\mathbf{x}\colon |\mathbf{x}|_0 \leqslant k^\star} \Phi(\mathbf{x}) \geqslant \sum_{\mathbf{x}\colon |\mathbf{x}|_0 \geqslant n-k^\star} \Phi(\mathbf{x})\,. \tag{5.9}$$

Since $\Delta\Phi = \phi(n)$, by applying the result on the mixing time of the logit dynamics of potential games (see Corollary 5.2.4) we get $t_{\text{mix}} = \mathcal{O}\left(n \cdot e^{\beta\phi(n)} \cdot (\beta\phi(n) + n)\right)$. By observing that $\phi(n) = \Theta(n^2\Delta)$, we have that the term δ at exponent in the bound of Theorem 5.4.1 can be dropped.

We next state a lower bound on the mixing time for coordination games on a clique.

Lemma 5.4.2. *Let $\mathcal{G}$ be an n-player graphical coordination game on the clique. The mixing time of the logit dynamics for $\mathcal{G}$ is $t_{\text{mix}} = \Omega\left(e^{(\beta - o(1))\phi(0)}\right)$.*

Proof. We obtain our lower bound by applying Theorem 5.1.3 with configuration $\mathbf{x}^\star = (1, \ldots, 1)$ and set $M = \{x \in S \,:\, |\mathbf{x}|_0 = k^\star\}$.

The connected component R of $S \setminus M$ that contains $\mathbf{x}^\star$ is

$$R = \{\mathbf{x} \in S\colon\ |\mathbf{x}|_0 < k^\star\}\ .$$

From (5.9) it follows that $\pi(R) \leqslant \frac{1}{2}$. Finally, notice that

$$\begin{aligned} |\partial R| &\leqslant |\{\mathbf{x} \in \Omega\colon |\mathbf{x}|_0 = k^\star - 1\}| \\ &= \binom{n}{k^\star - 1} \leqslant n^{k^\star} \leqslant n^{\frac{2}{b-c}\frac{\phi(0)}{n-1}}\ . \end{aligned}$$

The lemma follows by applying Theorem 5.1.3 and by observing that the minimum potential among profiles in R and M are $\Phi^R = -\phi(0)$ and $\Phi^M = 0$, respectively. □

We stress that when the basic coordination game has no risk dominant strategy (that is the case $\Delta = \delta$), $\phi(0) = \phi(n)$ and thus the exponents of the upper and lower bound coincide up to a $o(1)$ term. In general, by observing that $\phi(0) = \Theta(n^2\delta)$ and $\phi(n) = \Theta(n^2\Delta)$, we obtain the following theorem.

Theorem 5.4.3. *For every graphical coordination game on a clique there exist two constants C and D such that $C^{\beta n^2\delta} \leqslant t_{\text{mix}} \leqslant D^{\beta n^2\Delta}$.*

5.4.3 On the ring

In this section we give upper and lower bounds on the mixing time for graphical coordination games on the *ring* when there is no risk dominant strategy. Unlike the clique, the ring encodes a very local type of interaction between the players which is more likely to occur in a social context. Our results show that the mixing time is polynomial in the number of players n and e^β.

From the potential function given in (5.7), we can observe that $\Phi(\mathbf{1}) = \Phi(\mathbf{0}) = -n\delta$. Moreover, if n is even, the configuration $\mathbf{x}$ where every player selects a strategy different from the one selected by her neighbors has potential $\Phi(\mathbf{x}) = 0$: thus, there are graphical coordination games on the ring where $\Delta\Phi = n\delta$. If we used Corollary 5.2.4, we would get an upper bound exponential in n, whereas if we used Theorem 5.4.1, since the cutwidth of the ring is 4, we would get an upper bound exponential in $4\delta\beta$ and polynomial in $n^4\delta\beta$. Instead we here show a upper bound that is exponential only in $2\delta\beta$ and polynomial only in $n \log n$.

The proof of the upper bound uses the path coupling technique (see Theorem 1.3.3) and can be seen as a generalization of the upper bound on the mixing time for the Ising model on the ring (see Theorem 3.2.5).

Theorem 5.4.4. *Let $\mathcal{G}$ be an n-player graphical coordination games on the ring with no risk-dominant strategy ($a - d = b - c = \delta$). The mixing time of the logit dynamics for $\mathcal{G}$ is*

$$t_{\text{mix}} = \mathcal{O}\left((e^{2\beta\delta} + 1) n \log n\right)\ .$$

Proof. We identify the n players with the integers in $\{0, \ldots, n-1\}$ and assume that every player i plays the basic coordination game with her two adjacent players, $(i-1) \bmod n$ and $(i+1) \bmod n$. Let $S = \{0,1\}^n$ be the set of profiles of the game and consider the Hamming graph G over S where profiles $\mathbf{x}$ and $\mathbf{y}$ are adjacent if and only if they differ in exactly one position.

Let us consider two adjacent configurations $\mathbf{x}$ and $\mathbf{y}$. Denote by j the position in which they differ and assume, without loss of generality, that $x_j = 1$ and $y_j = 0$. We consider the following coupling for two chains X and Y starting respectively from $X_0 = \mathbf{x}$ and $Y_0 = \mathbf{y}$: pick $i \in \{0, \ldots, n-1\}$ and $U \in [0,1]$ independently and uniformly at random and update position i of $\mathbf{x}$ and $\mathbf{y}$ by setting

$$x_i = \begin{cases} 0, & \text{if } U \leqslant \sigma_i(0 \mid \mathbf{x}); \\ 1, & \text{if } U > \sigma_i(0 \mid \mathbf{x}); \end{cases} \qquad y_i = \begin{cases} 0, & \text{if } U \leqslant \sigma_i(0 \mid \mathbf{y}); \\ 1, & \text{if } U > \sigma_i(0 \mid \mathbf{y}). \end{cases}$$

We next compute the expected distance between X_1 and Y_1 after one step of the coupling. Notice that $\sigma_i(0 \mid \mathbf{x})$ only depends on x_{i-1} and x_{i+1} and $\sigma_i(0 \mid \mathbf{y})$ only on y_{i-1} and y_{i+1}. Therefore, since $\mathbf{x}$ and $\mathbf{y}$ only differ at position j, $\sigma_i(0 \mid \mathbf{x}) = \sigma_i(0 \mid \mathbf{y})$ for $i \neq j-1, j+1$.

We start by observing that if position j is chosen for update (this happens with probability $1/n$) then, by the observation above, both chains perform the same update. Since $\mathbf{x}$ and $\mathbf{y}$ differ only for player j, we have that the two chains are coupled (and thus at distance 0). Similarly, if $i \neq j-1, j, j+1$ (which happens with probability $(n-3)/n$) we have that both chains perform the same update and thus remain at distance 1. Finally, let us consider the case in which $i \in \{j-1, j+1\}$. In this case, since $x_j = 1$ and $y_j = 0$, we have that $\sigma_i(0|\mathbf{x}) \leqslant \sigma_i(0|\mathbf{y})$. Therefore, with probability $\sigma_i(0 \mid \mathbf{x})$ both chains update position i to 0 and thus remain at distance 1; with probability $1 - \sigma_i(0 \mid \mathbf{y})$ both chains update position i to 1 and thus remain at distance 1; and with probability $\sigma_i(0 \mid \mathbf{y}) - \sigma_i(0 \mid \mathbf{x})$ chain X updates position i to 1 and chain Y updates position i to 0 and thus the two chains go to distance 2. By summing up, we have that the expected distance $E[\rho(X_1, Y_1)]$ after one step of coupling of the two chains is

$$\begin{aligned} E[\rho(X_1, Y_1)] = &= \frac{n-3}{n} + \frac{1}{n} \sum_{i \in \{j-1, j+1\}} [\sigma_i(0 \mid \mathbf{x}) + 1 - \sigma_i(0 \mid \mathbf{y}) + 2 \cdot (\sigma_i(0 \mid \mathbf{y}) - \sigma_i(0 \mid \mathbf{x}))] \\ &= \frac{n-3}{n} + \frac{1}{n} \cdot \sum_{i \in \{j-1, j+1\}} (1 + \sigma_i(0 \mid \mathbf{y}) - \sigma_i(0 \mid \mathbf{x})) \\ &= \frac{n-1}{n} + \frac{1}{n} \cdot \sum_{i \in \{j-1, j+1\}} (\sigma_i(0 \mid \mathbf{y}) - \sigma_i(0 \mid \mathbf{x})) \end{aligned}$$

Let us now evaluate the difference $\sigma_i(0 \mid \mathbf{y}) - \sigma_i(0 \mid \mathbf{x})$ for $i = j-1$ (the same computation holds for $i = j+1$). We distinguish two cases depending on the strategies of player $j-2$ and start with the case $x_{j-2} = y_{j-2} = 1$. In this case we have that

$$\sigma_{j-1}(0 \mid \mathbf{x}) = \frac{1}{1 + e^{2\beta\delta}} \quad \text{and} \quad \sigma_{j-1}(0 \mid \mathbf{y}) = \frac{1}{2}\,.$$

Thus,

$$\sigma_{j-1}(0 \mid \mathbf{y}) - \sigma_{j-1}(0 \mid \mathbf{x}) = \frac{1}{2} - \frac{1}{1 + e^{2\beta\delta}}.$$

If instead $x_{j-2} = y_{j-2} = 0$, we have

$$\sigma_{j-1}(0 \mid \mathbf{x}) = \frac{1}{2} \quad \text{and} \quad \sigma_{j-1}(0 \mid \mathbf{y}) = \frac{1}{1 + e^{-2\beta\delta}}\,.$$

Thus

$$\begin{aligned}\sigma_{j-1}(0 \mid \mathbf{y}) - \sigma_{j-1}(0|\mathbf{x}) &= \frac{1}{1+e^{-2\beta\delta}} - \frac{1}{2} \\ &= 1 - \frac{1}{1+e^{2\beta\delta}} - \frac{1}{2} = \frac{1}{2} - \frac{1}{1+e^{2\beta\delta}}\,.\end{aligned}$$

We can conclude that the expected distance after one step of the chain is

$$\begin{aligned}E[\rho(X_1,Y_1)] &= \frac{n-1}{n} + \frac{1}{n}\left(1 - \frac{2}{1+e^{2\beta\delta}}\right) \\ &= 1 - \frac{2}{n(1+e^{2\beta\delta})} \leqslant e^{-\frac{2}{n(1+e^{2\beta\delta})}}\,.\end{aligned}$$

Since the diameter of G is $\mathrm{diam}(G) = n$, by applying Theorem 1.3.3 with $\alpha = \frac{2}{n(1+e^{2\beta\delta})}$, we obtain the theorem. □

The upper bound in Theorem 5.4.4 is nearly tight (up to the $n \log n$ factor). Indeed, a lower bound can be obtained by applying the Bottleneck Ratio technique (see Theorem 1.3.7) to the set $R = \{\mathbf{1}\}$. Notice that $\pi(R) \leqslant \frac{1}{2}$ since profile $\mathbf{0}$ has the same potential as $\mathbf{1}$. Thus set R satisfies the hypothesis of Theorem 1.3.7. Simple computations show that the bottleneck ratio is

$$B(R) = \sum_{\mathbf{y}\neq\mathbf{1}} P(\mathbf{1},\mathbf{y}) = \frac{1}{1+e^{2\beta\delta}}\,.$$

Thus, by applying Theorem 1.3.7, we obtain the following bound.

Theorem 5.4.5. *Let $\mathcal{G}$ be a n-player graphical coordination game on a ring with no risk-dominant strategy. The mixing time of the logit dynamics for $\mathcal{G}$ is $t_{\mathrm{mix}} = \Omega\left(1+e^{2\beta\delta}\right)$.*

5.5 Conclusions and open problems

In this chapter we give different bounds on the mixing time of the logit dynamics for the class of potential games: we showed that the mixing time is fast when β is small enough and we found the structural property of the game that characterizes the mixing time for large β; finally we showed a bound that holds for every value of β. Unfortunately, the last bound does not match the previous ones in every game: this fact raises the quest for an unifying bound that holds for every value of β and matches with bounds given in Theorems 5.2.6 and 5.2.8. We suppose that this unifying bound has to depends on a structural property that, as β increase, evolves from $\Delta\Phi$ to $\zeta^\star$.

A term $e^{\beta L}$ appears in the bound to the mixing time of the logit dynamics for potential game both when β is high (Theorem 5.2.8) and for every β (Corollary 5.2.4). Nevertheless, in all lower bounds we presented this term never appears: it is natural to ask if this term is necessary (by presenting a matching lower bound) or it can be eliminated by a more careful analysis. We suppose that the second hypothesis is right, but we are still unable to prove it.

Another open problem is about potential games where $\zeta^\star = 0$. We showed that this class of games exactly characterizes the potential games for which the relaxation time of the logit dynamics is independent of β: our aim is to prove the same result holds for the mixing time too.

Previous works gave a lot of attention to logit dynamics for graphical coordination games: for this reason, finding tight bounds on the mixing time of our dynamics for this class of games will be very interesting. Unfortunately, our results about the clique and the ring show that the bound of Theorem 5.4.1 is not tight.

If the mixing time of the logit dynamics is fast, then the logit equilibrium gives good predictions about the state of a complex system after a small number of time steps. Anyway, we showed that there are games where the mixing time can be exponential in the number of players. For these games the logit equilibrium does not represent a meaningful description of the game and it becomes interesting to analyze the *transient* phase of the logit dynamics, in order to investigate what kind of predictions can be made about the state of the system in such a phase. This approach will be pursued in the next chapter.

Chapter 6

Metastability

The drawback of using the *logit equilibrium* to describe the behavior of a complex system is that the system may take too long to reach it, unless the chain is rapidly mixing. Previous chapters showed the mixing of the logit dynamics for strategic games can be not rapid depending on the features of the underlying game and on the rationality level.

For this reason, in this chapter we focus on the *transient phase* of the logit dynamics and, in particular, we try to answer the following questions: when the mixing time is exponential in the number of players, is the transient phase completely chaotic, or can we still spot some regularities? Are we able to say something about the behavior of the chain before it reaches the stationary distribution?

Obviously, such a chain is perfectly described by the collections of probability distributions consisting of one distribution for each time step and each starting profile. This should be contrasted with the rapidly mixing case (i.e., a Markov chain with polynomial mixing time) in which one can approximately describe the state of the system (after the mixing time) using an unique distribution (that is, the stationary distribution).

Our results show games for which regularities can be observed even in the transient phase. In particular, we will show that, depending on the starting profile, the dynamics rapidly reaches a distribution and remains close to this distribution for a sufficiently long time (we call such a distribution *metastable*).

We can describe our results also in terms of the quantity of information needed to predict the status of a system that evolves according to the logit dynamics. We know that the long-term behavior of the system can be compactly described in terms of a unique distribution but we have to wait a transient phase of length equal to the mixing time. Thus, if the system is rapidly mixing this description is significant after a short transient phase. However, when the mixing time is super-polynomial this description becomes significant only after a long time. Our results show that, for a large class of n-player games, logit dynamics is not rapidly mixing but the profile (the strategies played by the n players) can still be described with good approximation and for a super-polynomial number of steps by means of a small number of probability distributions. This comes at the price of sacrificing a short polynomial initial transient phase (so far we are on a par with the rapidly mixing case) and requires a few bits of information about the starting profile (this is not needed in the rapidly mixing case).

In this chapter, we obtain results[1] about the metastability of the logit dynamics for different classes of coordination games.

- We start in Section 6.1 by introducing the notion of an (ε, T)-metastable distribution μ and of its *pseudo-mixing time*. Roughly speaking, μ is (ε, T)-metastable for a Markov

[1] Most of the results in this chapter already appeared in [4].

chain if, starting from μ, the Markov chain stays at distance at most ε from μ for at least T steps. The pseudo-mixing time of μ starting from a state x, $t_\mu^x(\varepsilon)$, is the number of steps needed by the Markov chain to get ε-close to μ when started from x.

In a rapidly-mixing Markov chain, after a "short time" and regardless of the starting state the chain converges rapidly to the stationary distribution and remains there. For the case of non-rapidly mixing Markov chains, we replace the notions of "mixing time" and "stationary distribution" by that of "pseudo-mixing time" and that of "metastable distribution". Intuitively speaking, we would like to say that, even when the mixing time is (prohibitively) high, there are "few" distributions which give us an accurate description of the chain over a "reasonable amount of time". Roughly speaking, the state space Ω can be partitioned into a small number of subsets $\Omega_1, \Omega_2, \dots$ of "equivalent" states; that is, if the chain starts in any of the states in Ω_i, then it will rapidly converge to a "metastable" distribution μ_i, where metastable denotes the fact that the chain remains there for "sufficiently" long.

- In Section 6.2, we analyze the metastable distributions of the Ising model on the clique, also known as the Curie-Weiss model.

  As showed in Section 3.2, the mixing time of this dynamics is known to be exponential for every $\beta > 1/n$. For this model, we show that distributions where all magnets have the same magnetization are $(1/n, t)$-metastable for any $t = \mathsf{poly}\,(n)$ when $\beta = \Omega(\log n/n)$. Moreover we show that the pseudo-mixing time of these distributions is polynomial when the dynamics starts from a profile where the difference in the number of positive and negative magnets is large.

- In Section 6.3 we study graphical coordination games on the ring topology. We show that for every starting profile there is a metastable distribution and the dynamics approaches it in a polynomial number of steps.

- Finally, we consider the OR-game, defined in Section 4.4, that highlights the distinctive features of our metastability notion based on distributions.

Previous works about metastability

In Physics, Chemistry or Biology, metastability is a phenomenon related to the evolution of systems under noisy dynamics. In particular, metastability concerns the transition between different regions of the state space and the existence of multiple, well-separated time scales: at short time scales the system appears to be in a *quasi-equilibrium*, and it explores only a confined region of the available space state, while, at larger time scales, it undergoes transitions between such different regions. Examples of metastability can be found in Biology, Climatology, Economics, Materials Science and Physics.

Metastability appears for the first time around 1935 with the works of Eyring [41] and Kramers [71] on diffusion in potential wells, but the mathematically rigorous analysis of metastability phenomena in the context of randomly perturbed dynamical systems start in the early 1970's with the work of Freidlin and Wentzell [45]. Since then, metastability is a very well studied topic in Physics and several monographs on this subject are available (see, for example [62, 104, 19, 63]). The goal of metastability is to model processes showing the following typical behavior: starting from a given profile, the system will rather quickly visit the nearby maximum of the potential function (a metastable state); the dynamics stays very close to such a state for a very long time, avoiding visits to other local maxima; at some point, the system leaves the metastable state (and its neighborhood) and moves to some other local maximum, usually better than the previous one; the process then is repeated. Research in Physics about metastability aims at expressing typical features of a metastable state and to evaluate the transition time between metastable states; the main approaches used to this analysis are based on large deviation theory [45] or on potential theory [20]. Our approach is closest to the one of Bovier et al. [21]. They define the notion of a *metastable point* as a state that is quickly reached and difficult to leave. For every metastable point $\mathbf{x}$ they define the *local valley* of $\mathbf{x}$ as the set of states for which $\mathbf{x}$ is the metastable point with the smallest hitting time and the associated metastable distribution associated with $\mathbf{x}$ is the stationary distribution restricted to the local valley. In [22], Bovier and Manzo apply the approach of [21] in the context of zero temperature limit of Glauber dynamics of spin systems in finite volume and show that the transition times can be expressed in terms of properties of the potential function.

Metastability was analyzed not only for discrete dynamics, but also for continuous Markov processes. In [73] Larralde et al. define a metastable state by two components: spectral feature of a state (namely, isolated eigenstate of the master operator of the Markov Process having an exceptionally low eigenvalue) and the technical condition meaning that the probability of being in a metastable state at equilibrium is vanishingly small. These conditions partition the state space in two disjoint set: the metastable states and the equilibrium states. They show that for any starting profile $\mathbf{x}$, the dynamics quickly reach, with a probability $p_{\mathbf{x}}$, a state which is fully in the metastable region and, with probability $1-p_{\mathbf{x}}$, the equilibrium. Further, if the dynamics start from the metastable region, then the probability of leaving it in short time is very low. Moreover, they consider a *restricted dynamics* in which the process is reflected each time it attempts to leave the metastable region, whose equilibrium is described by the restriction of the stationary distribution to the metastable region: they show that these restricted dynamics well mimic the process when the starting point is in the metastable region.

Very recently, Beltran and Landim [11] describe for the continuous time Markov process of the Ising model all metastable behaviors, defining time scales at which they occur, the metastable set associated to each time scale, and the asymptotic dynamics which specifies at which rate the process jumps from one metastable state to another.

The work on censored Glauber dynamics [74, 38, 39] is also related to ours: the mixing time in a censored dynamics resemble the pseudo mixing time for the metastable distribution on a subset of states. However, we stress that the censored dynamics alters the original evolution of the Markov chain and the techniques developed do not seem useful to answer questions about the pseudo-mixing time.

6.1 Metastability

In this section we give formal definitions of *metastable distributions* and *pseudo-mixing* time. As a simple example we analyze metastability for the logit dynamics for 2-player coordination games and we also highlight some connections between metastability and the bottleneck ratio.

Definition 6.1.1 (Metastable distribution)**.** *Let P be a Markov chain with finite state space Ω. A probability distribution μ over Ω is* (ε, T)-metastable *if for every* $0 \leqslant t \leqslant T$ *it holds that*

$$\left\|\mu P^t - \mu\right\|_{\mathrm{TV}} \leqslant \varepsilon\,.$$

Here are two obvious property of metastable distributions.

1. *Monotonicity*: If μ is (ε, T)-metastable then it is (ε', T')-metastable for every $\varepsilon' \geqslant \varepsilon$ and $T' \leqslant T$;

2. *Stationarity and Metastability*: if μ is $(0, 1)$-metastable, then it is $(0, T)$-metastable for every T; μ is stationary if and only if it is $(0, 1)$-metastable.

A third property is given by the following easy and useful lemma.

Lemma 6.1.2. *If μ is $(\varepsilon, 1)$-metastable for P then μ is $(\varepsilon T, T)$-metastable for P.*

Proof. By using the triangle inequality, we have

$$\begin{aligned}\left\|\mu P^T - \mu\right\|_{\mathrm{TV}} &\leqslant \left\|\mu P^t - \mu P\right\|_{\mathrm{TV}} + \|\mu P - \mu\|_{\mathrm{TV}} \\ &\leqslant \left\|\mu P^{t-1} - \mu\right\|_{\mathrm{TV}} + \varepsilon\,,\end{aligned}$$

where the last inequality follows from the $(\varepsilon, 1)$-metastability of μ and from the fact that if μ and ν are two probability distributions and P is a stochastic matrix then $\|\mu P - \nu P\|_{\mathrm{TV}} \leqslant \|\mu - \nu\|_{\mathrm{TV}}$. □

The definition of metastable distribution captures the idea of a distribution that behaves approximately like the stationary distribution, meaning that if we start from such distribution and run the chain we stay close to it for a long time.

Among all metastable distributions, we are interested in the ones that are quickly reached from a, possibly large, set of states. This motivates the following definition.

Definition 6.1.3 (Pseudo-mixing time)**.** *Let P be a Markov chain with state space Ω, let $R \subseteq \Omega$ be a set of states and let μ be a probability distribution over Ω. We define the pseudo-mixing time $t_\mu^R(\varepsilon)$ as*

$$t_\mu^R(\varepsilon) = \inf\{t \in \mathbb{N} \colon \left\|P^t(x, \cdot) - \mu\right\|_{\mathrm{TV}} \leqslant \varepsilon \text{ for all } x \in R\}\,.$$

Since the stationary distribution π of an ergodic Markov chain is reached within ε in time $t_{\mathrm{mix}}(\varepsilon)$ from every state, according to Definition 6.1.3, we have that $t_\pi^\Omega(\varepsilon) = t_{\mathrm{mix}}(\varepsilon)$.

The following simple lemma connects metastability and pseudo-mixing.

Lemma 6.1.4. *Let μ be a (ε, T)-metastable distribution and let $R \subseteq \Omega$ be a set of states such that $t_\mu^R(\varepsilon) < +\infty$. Then for every $x \in R$ it holds that*

$$\left\|P^t(x, \cdot) - \mu\right\|_{\mathrm{TV}} \leqslant 2\varepsilon \quad \text{for every } t_\mu^R(\varepsilon) \leqslant t \leqslant t_\mu^R(\varepsilon) + T\,.$$

Proof. Let us name $\bar{t} = t - t_\mu^R(\varepsilon)$ for convenience sake. By using the triangle inequality for the total variation distance, the fact that $P^{\bar{t}}$ is a stochastic matrix, and the definitions of metastable distribution and pseudo-mixing, we have that

$$\begin{aligned}\left\|P^t(x, \cdot) - \mu\right\|_{\mathrm{TV}} &= \left\|P^{t_\mu^R(\varepsilon)}(x, \cdot)P^{\bar{t}} - \mu\right\|_{\mathrm{TV}} \\ &\leqslant \left\|P^{t_\mu^R(\varepsilon)}(x, \cdot)P^{\bar{t}} - \mu P^{\bar{t}}\right\|_{\mathrm{TV}} + \left\|\mu P^{\bar{t}} - \mu\right\|_{\mathrm{TV}} \\ &\leqslant \left\|P^{t_\mu^R(\varepsilon)}(x, \cdot) - \mu\right\|_{\mathrm{TV}} + \left\|\mu P^{\bar{t}} - \mu\right\|_{\mathrm{TV}} \leqslant 2\varepsilon\,.\end{aligned}$$

□

Example: A simple three-state Markov chain

As a first example, let us consider the simplest Markov chain that may highlight the concepts of *metastability* and *pseudo-mixing*,

$$P = \begin{pmatrix} \varepsilon & \frac{1-\varepsilon}{2} & \frac{1-\varepsilon}{2} \\ \varepsilon & 1-\varepsilon & 0 \\ \varepsilon & 0 & 1-\varepsilon \end{pmatrix}$$

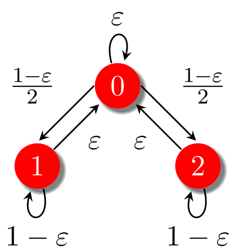

The chain is ergodic with stationary distribution $\pi = (\varepsilon, (1-\varepsilon)/2, (1-\varepsilon)/2)$, and its mixing time is $t_{\text{mix}} = \Theta(1/\varepsilon)$. Hence the mixing time increases as ε tends to zero.

Now observe that, for every $\delta > \varepsilon$, degenerate[2] distributions $\mu_1 = (0,1,0)$ and $\mu_2 = (0,0,1)$ are $(\delta, \Theta(\delta/\varepsilon))$-metastable according to Definition 6.1.1. If we start from the first state (i.e. from degenerate distribution $\nu = (1,0,0)$), after one step we are in the stationary distribution.

Hence, even if the mixing time can be arbitrary large, for every starting state x there is a $(1/4, \Theta(t_{\text{mix}}))$-metastable distribution μ that is quickly (in constant time, independent of ε) reached from x.

Example: Two-player coordination games

Coordination games given in (3.1) are examples of games where the mixing time is a function increasing exponentially in β (see Theorem 4.3.2).

We define

$$\varepsilon = \frac{1}{1+e^{(a-d)\beta}}\,; \qquad \delta = \frac{1}{1+e^{(b-c)\beta}}\,.$$

We can rewrite the stationary distribution of the logit dynamics for coordination games as

$$\pi = \frac{1}{\varepsilon+\delta}\left[\delta(1-\varepsilon),\, \varepsilon\delta,\, \varepsilon\delta,\, \varepsilon(1-\delta)\right]\,.$$

Thus, Theorem 4.3.2 states that the mixing time for such games is $t_{\text{mix}} = \Theta(1/\delta)$.

Let consider the special case when $\varepsilon = \delta$, hence the stationary distribution is

$$\pi = ((1-\delta)/2,\, \delta/2,\, \delta/2,\, (1-\delta)/2)\,.$$

Let $\mu_{(0,0)}$ and $\mu_{(1,1)}$ be the two distributions concentrated in states $(0,0)$ and $(1,1)$ respectively, i.e.,

$$\mu_{(0,0)} = [1,0,0,0] \quad \text{and} \quad \mu_{(1,1)} = [0,0,0,1]\,.$$

Observe that, if we start from $\mu_{(0,0)}$ or $\mu_{(1,1)}$, after one step of the chain we are respectively in distributions

$$\begin{aligned} \mu_{(0,0)}P &= [1-\delta, \delta/2, \delta/2, 0]\,, \\ \mu_{(1,1)}P &= [0, \delta/2, \delta/2, 1-\delta]\,. \end{aligned}$$

Hence

$$\left\|\mu_{(0,0)}P - \mu_{(0,0)}\right\|_{\text{TV}} = \left\|\mu_{(1,1)}P - \mu_{(1,1)}\right\|_{\text{TV}} = \delta\,.$$

By using Lemma 6.1.2, we have that, for every constant $c \leqslant 1/2$, $\mu_{(0,0)}$ and $\mu_{(1,1)}$ are $(c, \Theta(1/\delta))$-metastable according to Definition 6.1.1. Moreover, if the chain starts from state $(0,1)$ or from state $(1,0)$, after 1 step of the chain we are δ-close to the stationary distribution π, indeed

$$\begin{aligned} (0,1,0,0)P &= \left[\frac{1-\delta}{2}, \delta, 0, \frac{1-\delta}{2}\right]\,, \\ (0,0,1,0)P &= \left[\frac{1-\delta}{2}, 0, \delta, \frac{1-\delta}{2}\right]\,, \end{aligned}$$

and

$$\|(0,1,0,0)P - \pi\|_{\text{TV}} = \|(0,1,0,0)P - \pi\|_{\text{TV}} = \delta/2\,.$$

We can summarize what we have just shown in the following theorem.

Theorem 6.1.5. *Let $\mathcal{G}$ be a 2-player coordination game with profile space S and let P be the transition matrix of the logit dynamics for $\mathcal{G}$. For every starting profile $\mathbf{x} \in S$ and every constant $c \leqslant 1/2$ there is a $(c,\, \Theta(t_{\text{mix}}))$-metastable distribution $\mu_{\mathbf{x}}$ such that $t_{\mu_{\mathbf{x}}}^{\{\mathbf{x}\}} = \Theta(1)$.*

[2] A probability distribution is *degenerate* if it is concentrated in one single element.

6.1.1 Metastability and the bottleneck ratio

Consider an ergodic Markov chain P with state space Ω and stationary distribution π. For a subset R of states, let π_R be the stationary distribution conditioned on R, i.e.

$$\pi_R(x) = \begin{cases} \pi(x)/\pi(R), & \text{if } x \in R\,; \\ 0, & \text{otherwise.} \end{cases} \tag{6.1}$$

It is well-known (see e.g. Theorem 7.3 in [75]) that the *bottleneck ratio* at set R equals the total variation distance between π_R and $\pi_R P$, i.e $\|\pi_R P - \pi_R\|_{\mathrm{TV}} = B(R)$. Hence, the following lemma about the metastability of π_R holds.

Lemma 6.1.6. *Let P be a Markov chain with finite state space Ω and let $R \subseteq \Omega$ be a subset of states. Then, π_R is $(B(R), 1)$-metastable.*

6.1.2 Pseudo-mixing time tools

In order to upper bound the mixing time of an ergodic chain, it is often used the fact that, for every starting state $x \in \Omega$ the total variation distance between the distribution of the chain at time t and the stationary distribution π is upper bounded by the maximum, over all states $y \in \Omega$, of the total variation between the chain starting at x and the chain starting at y (see Lemma 4.11 in [75]), i.e.

$$\left\|P^t(x,\cdot) - \pi\right\|_{\mathrm{TV}} \leqslant \max_{y \in \Omega} \left\|P^t(x,\cdot) - P^t(y,\cdot)\right\|_{\mathrm{TV}} .$$

In the following lemma we formalize and prove an analogous statement for metastable distributions.

Lemma 6.1.7. *Let P be a Markov chain with finite state space Ω and let μ be an (ε, T)-metastable distribution supported over a subset $R \subseteq \Omega$ of the state space. Then for every $x \in R$ and every $1 \leqslant t \leqslant T$, it holds that*

$$\left\|P^t(x,\cdot) - \mu\right\|_{\mathrm{TV}} \leqslant \varepsilon + \max_{y \in R} \left\|P^t(x,\cdot) - P^t(y,\cdot)\right\|_{\mathrm{TV}} .$$

Proof. From triangle inequality we have

$$\left\|P^t(x,\cdot) - \mu\right\|_{\mathrm{TV}} \leqslant \left\|P^t(x,\cdot) - \mu P^t\right\|_{\mathrm{TV}} + \left\|\mu P^t - \mu\right\|_{\mathrm{TV}} .$$

Since μ is (ε, t)-metastable for every $t \leqslant T$, we have $\left\|\mu P^t - \mu\right\|_{\mathrm{TV}} \leqslant \varepsilon$. Observe that, since $\mu(y) = 0$ for $y \notin R$, then for every set of states $A \subseteq \Omega$ and for every t it holds that

$$\begin{aligned} |P^t(x,A) - \mu P^t(A)| &= \left| P^t(x,A) - \sum_{y \in R} \mu(y) P^t(y,A) \right| \\ &= \left| \sum_{y \in R} \mu(y) \left(P^t(x,A) - P^t(y,A)\right) \right| \\ &\leqslant \sum_{y \in R} \mu(y) \left|P^t(x,A) - P^t(y,A)\right| \\ &\leqslant \max_{y \in R} \left|P^t(x,A) - P^t(y,A)\right| . \end{aligned}$$

Thus, the total variation between $P^t(x,\cdot)$ and μP^t is

$$\begin{aligned}
\left\|P^t(x,\cdot) - \mu P^t\right\|_{\mathrm{TV}} &= \max_{A\subseteq\Omega} \left|P^t(x,A) - \mu P^t(A)\right| \\
&\leqslant \max_{A\subseteq\Omega}\max_{y\in R} \left|P^t(x,A) - P^t(y,A)\right| \\
&= \max_{y\in R} \left\|P^t(x,\cdot) - P^t(y,\cdot)\right\|_{\mathrm{TV}} .
\end{aligned}$$

□

In some cases metastable distributions are concentrated in one single state. The following lemma shows that, in those cases, the hitting time of such a state can be used to establish the pseudo-mixing time of the metastable distribution.

Lemma 6.1.8. *Let P be a Markov chain with finite state space Ω and let μ be an (ε, T)-metastable distribution concentrated on a single state y. Let τ_y be the hitting time of this state. Then for all $x \in \Omega$ and $1 \leqslant t \leqslant T$, we have*

$$\left\|P^t(x,\cdot) - \mu\right\|_{\mathrm{TV}} \leqslant \varepsilon + (1-\varepsilon)\mathbf{P}_x\left(\tau_y > t\right) .$$

Proof. Since μ is concentrated in y, we have that

$$\begin{aligned}
\left\|P^t(x,\cdot) - \mu\right\|_{\mathrm{TV}} &= \mathbf{P}_x\left(X_t \neq y\right) \\
&= \mathbf{P}_x\left(X_t \neq y, \tau_y \leqslant t\right) + \mathbf{P}_x\left(X_t \neq y, \tau_y > t\right) \\
&= \mathbf{P}_x\left(X_t \neq y \mid \tau_y \leqslant t\right)\mathbf{P}_x\left(\tau_y \leqslant t\right) + \mathbf{P}_x\left(\tau_y > t\right) .
\end{aligned}$$

Moreover, observe that

$$\begin{aligned}
\mathbf{P}_x\left(X_t \neq y \mid \tau_y \leqslant t\right) &= \sum_{k\leqslant t} \mathbf{P}_x\left(X_t \neq y \mid \tau_y = k\right)\mathbf{P}_x\left(\tau_y = k \mid \tau_y \leqslant t\right) \\
&= \sum_{k\leqslant t} \mathbf{P}_y\left(X_{t-k} \neq y\right)\mathbf{P}_x\left(\tau_y = k \mid \tau_y \leqslant t\right) \\
&= \sum_{k\leqslant t} \left\|\mu P^{t-k} - \mu\right\|_{\mathrm{TV}} \mathbf{P}_x\left(\tau_y = k \mid \tau_y \leqslant t\right) \\
&\leqslant \varepsilon \sum_{k\leqslant t} \mathbf{P}_x\left(\tau_y = k \mid \tau_y \leqslant t\right) = \varepsilon ,
\end{aligned}$$

where in the inequality we used the metastability of μ. Hence,

$$\begin{aligned}
\left\|P^t(x,\cdot) - \mu\right\|_{\mathrm{TV}} &= \mathbf{P}_x\left(X_t \neq y \,|\, \tau_y \leqslant t\right)\mathbf{P}_x\left(\tau_y \leqslant t\right) + \mathbf{P}_x\left(\tau_y > t\right) \\
&\leqslant \varepsilon\mathbf{P}_x\left(\tau_y \leqslant t\right) + \mathbf{P}_x\left(\tau_y > t\right) \\
&= \varepsilon + (1-\varepsilon)\mathbf{P}_x\left(\tau_y > t\right) .
\end{aligned}$$

□

6.2 Ising model on the complete graph

Consider the n-player *Ising game* $\mathcal{G}$ defined in Section 3.2, that is the game-theoretic formulation of the well-studied Ising model on a clique. Let $S = \{-1, +1\}^n$ be the profile space of $\mathcal{G}$. For every $\mathbf{x} \in S$, the *magnetization* of $\mathbf{x}$ is defined as $\Lambda(\mathbf{x}) = \sum_{i=i}^n x_i$. Observe that the potential of a profile $\mathbf{x}$ depends only on its magnetization, i.e. if $\Lambda(\mathbf{x}) = k$ then $\Phi(\mathbf{x}) = -\phi(k) = -\frac{1}{2}\left(k^2 - n\right)$. To see this, let us name p and m the number of $+1$ and -1 respectively, in profile $\mathbf{x}$, and observe that $p - m = \Lambda(\mathbf{x}) = k$ and $p + m = n$. Each pair of players with the same sign contributes for $+1$ in $\Phi(\mathbf{x})$ and each pair of players with opposite signs contributes for -1; since there are $\binom{p}{2}$ pairs where both players play $+1$, $\binom{m}{2}$ pairs where both play -1 and $p \cdot m$ pairs where players play opposite strategies, we have that

$$\Phi(\mathbf{x}) = -\left(\binom{p}{2} + \binom{m}{2} - p \cdot m\right) = -\frac{1}{2}((p-m)^2 - (p+m)) .$$

In this section we study the metastability properties of the logit dynamics for the Ising game on the clique from our quantitative point of view. Namely, we show that, if we start from a profile where the number of $+1$ (respectively -1) is a sufficiently large majority, and if β is large enough then, after an initial pseudo-mixing phase, the distribution of the chain at time t is close, in total variation distance, to the degenerate distribution concentrated in the profile with all $+1$'s (respectively all -1's) for all $t = \mathsf{poly}\,(n)$.

Let π_+ and π_- be the two degenerate distributions concentrated in the states with all $+1$ and all -1, respectively. The next lemma shows that, for $\beta = \Omega(\log n/n)$, π_+ and π_- are metastable for a polynomially-long time.

Lemma 6.2.1. *If $\beta > c \log n/n$ then π_+ and π_- are $\left(1/n, n^{c-2}\right)$-metastable distributions of the logit dynamics for the Ising game.*

Proof. We prove the result for π_+, exactly the same proof (by swapping minuses and pluses) works for π_-.

Since $\pi_+(\mathbf{x}) = 0$ for all $\mathbf{x} \neq +\mathbf{1}$ and

$$(\pi_+ P)(\mathbf{x}) = \begin{cases} 0\,, & \text{if } \Lambda(\mathbf{x}) < n-2 \ \ (\text{i.e., } |\mathbf{x}|_{-1} > 1)\,; \\ \frac{1}{n} \cdot \frac{1}{1+e^{\beta(n-2)}}\,, & \text{if } \Lambda(\mathbf{x}) = n-2 \ \ (\text{i.e., } |\mathbf{x}|_{-1} = 1)\,; \\ \frac{1}{1+e^{-\beta(n-2)}}\,, & \text{if } \Lambda(\mathbf{x}) = n \ \ (\text{i.e., } \mathbf{x} = +\mathbf{1})\,, \end{cases}$$

the total variation distance between $\pi_+ P$ and π_+ is

$$\begin{aligned} \|\pi_+ P - \pi_+\|_{\mathrm{TV}} &= \frac{1}{2} \sum_{\mathbf{x} \in S} |\pi_+ P(\mathbf{x}) - \pi_+(\mathbf{x})| \\ &= \frac{1}{1 + e^{\beta(n-2)}} \leqslant e^{-\beta(n-2)} \leqslant n^{c-1} . \end{aligned}$$

In the last inequality we used $\beta \geqslant c \log n/n$. Hence π_+ is $(n^{-(c-1)}, 1)$-metastable. The thesis follows from Lemma 6.1.2. □

In order to give an upper bound on the pseudo-mixing time we need some preliminary results about birth-and-death chains, which we show in the next subsection.

6.2.1 Biased birth-and-death chains

In this section we consider *birth-and-death chains* with state space $\Omega = \{0, 1, \ldots, n\}$ (see Chapter 2.5 in [75] for a detailed description of such chains). For $k \in \{1, \ldots, n-1\}$ let $p_k = \mathbf{P}_k\,(X_1 = k+1)$, $q_k = \mathbf{P}_k\,(X_1 = k-1)$, and $r_k = 1 - p_k - q_k = \mathbf{P}_k\,(X_1 = k)$. We will be interested in the probability that the chain starting at some state $h \in \Omega$ hits state 0 before state n, namely $\mathbf{P}_k\left(X_{\tau_{0,n}} = n\right)$ where $\tau_{0,n} = \min\{t \in \mathbb{N} \colon X_t \in \{0, n\}\}$.

We start by giving an exact formula for such probability for the case when p_k and q_k do not depend on k.

Lemma 6.2.2. *Consider a birth-and-death chain $\{X_t\}$ with state space $\Omega = \{0, 1, \ldots, n\}$. Suppose for all $k \in \{1, \ldots, n-1\}$ it holds that $p_k = \varepsilon$ and $q_k = \delta$, for some ε and δ with $\varepsilon + \delta \leqslant 1$. Then the probability the chain hits state n before state 0 starting from state $h \in \Omega$ is*

$$\mathbf{P}_h\left(X_{\tau_{0,n}} = n\right) = \frac{1 - (\delta/\varepsilon)^h}{1 - (\delta/\varepsilon)^n}\,.$$

Proof. Let α_k be the probability to reach state n before state 0 starting from state k, i.e.

$$\alpha_k = \mathbf{P}_k\left(X_{\tau_{0,n}} = n\right)\,.$$

Observe that for $k = 1, \ldots, n-1$ we have

$$\alpha_k = \delta \cdot \alpha_{k-1} + \varepsilon \cdot \alpha_{k+1} + (1 - (\delta + \varepsilon))\,\alpha_k\,. \tag{6.2}$$

Hence

$$\varepsilon \cdot \alpha_k - \delta \cdot \alpha_{k-1} = \varepsilon \cdot \alpha_{k+1} - \delta \cdot \alpha_k\,,$$

with boundary conditions $\alpha_0 = 0$ and $\alpha_n = 1$. If we name $\Delta_k = \varepsilon \cdot \alpha_k - \delta \cdot \alpha_{k-1}$ we have $\Delta_k = \Delta_{k+1}$ for all k. By simple calculation and using that $\alpha_0 = 0$ it follows that

$$\alpha_k = \frac{\Delta}{\varepsilon} \sum_{i=0}^{k-1} \left(\frac{\delta}{\varepsilon}\right)^i = \frac{\Delta}{\varepsilon - \delta}\left(1 - (\delta/\varepsilon)^k\right)\,.$$

From $\alpha_n = 1$ we get

$$\Delta = \frac{\varepsilon - \delta}{(1 - (\delta/\varepsilon)^n)}\,.$$

Hence

$$\alpha_k = \frac{1 - (\delta/\varepsilon)^k}{1 - (\delta/\varepsilon)^n}\,. \tag{6.3}$$

□

Lemma 6.2.3. *Consider a birth-and-death chain $\{X_t\}$ with state space $\Omega = \{0, 1, \ldots, n\}$. Suppose for all $k \in \{1, \ldots, n-1\}$ it holds that $p_k \geqslant \varepsilon$ and $q_k \leqslant \delta$, for some ε and δ with $\varepsilon + \delta \leqslant 1$. Then the probability to hit state n before state 0 starting from state $h \in \Omega$ is*

$$\mathbf{P}_h\left(X_{\tau_{0,n}} = n\right) \geqslant \frac{1 - (\delta/\varepsilon)^h}{1 - (\delta/\varepsilon)^n}\,.$$

Proof. Let $\{Y_t\}$ be a birth-and-death chain with the same state space as $\{X_t\}$ but different transition rates

$$\mathbf{P}_k(Y_1 = k-1) = \delta \qquad \mathbf{P}_k(Y_1 = k+1) = \varepsilon$$

Consider the following coupling of X_t and Y_t: when (X_t, Y_t) is at state (k, h), consider the two $[0, 1]$ intervals, each one partitioned in three subintervals as in Fig. 6.1. Let U be a uniform random variable over the interval $[0, 1]$ and choose the update for the two chains according to position of U in the two intervals.

Observe that, since $p_k \geqslant \varepsilon$ and $q_k \leqslant \delta$, if the two chains start at the same state $h \in \Omega$, i.e. $(X_0, Y_0) = (h, h)$, then at every time t it holds that $X_t \geqslant Y_t$. Hence if chain Y_t hits state n before state 0, then chain X_t hits state n before state 0 as well. More formally, let $\tau_{0,n}$ and $\hat{\tau}_{0,n}$ be the random variables indicating the first time chains X_t and Y_t respectively hit state 0 or n, hence

$$\left\{Y_{\hat{\tau}_{0,n}} = n\right\} \Rightarrow \left\{X_{\tau_{0,n}} = n\right\}\,.$$

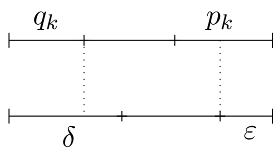

Figure 6.1: Partition for the coupling in Lemma 6.2.3.

Thus

$$\mathbf{P}_h\left(X_{\tau_{0,n}} = n\right) \geqslant \mathbf{P}_h\left(Y_{\hat{\tau}_{0,n}} = n\right) \geqslant \frac{1 - (\delta/\varepsilon)^h}{1 - (\delta/\varepsilon)^n}\,.$$

In the last inequality we used Lemma 6.2.2. □

Lemma 6.2.4. *Consider a birth-and-death chain $\{X_t\}$ with state space $\Omega = \{0, 1, \ldots, n\}$. Suppose for all $k \in \{1, \ldots, n-1\}$ it holds that $q_k/p_k \leqslant \alpha$, for some $\alpha < 1$. Then the probability to hit state 0 before state n starting from state $h \in \Omega$ is*

$$\mathbf{P}_h\left(X_{\tau_{0,n}} = 0\right) \leqslant \alpha^h\,.$$

Proof. Let $\hat{p}_k = \frac{p_k}{p_k+q_k}$ and $\hat{q}_k = \frac{q_k}{p_k+q_k}$ and let $\{Y_t\}$ be the birth-and-death chain with transition rates $\hat{p}_k$ and $\hat{q}_k$.

Let $\{U_t\}$ be an array of random variables such that $U_t = -1$ with probability q_{Y_t}, $U_t = +1$ with probability p_{Y_t} and $U_t = 0$ with remaining probability. We will use U_t to update chains X_t and Y_t at different time steps. Specifically, we denote with u the index of the first variables U_t not used for updating X_t (thus, at the beginning $u = 1$) and : set $Y_{t+1} = Y_t + U_t$; for chain X_t, we toss a coin that gives head with probability $p_{X_t} + q_{X_t}$ and if it gives tail we set $X_{t+1} = X_t$, otherwise we set $X_{t+1} = X_t + U_u$. Roughly speaking, we have that the chain X_t *follows the path* traced by chain Y_t: indeed, it is easy to see that, if they start at the same place, the sequence of states visited by the two chains is the same and in the same order. Hence chain X_t hits state 0 before state n if and only if chain Y_t hits state 0 before state n and thus $\mathbf{P}_k\left(X_{\tau^X_{0,n}} = 0\right) = \mathbf{P}_k\left(Y_{\tau^Y_{0,n}} = 0\right)$.

Finally, observe that $\frac{\hat{q}_k}{\hat{p}_k} = \frac{q_k}{p_k} \leqslant \alpha$ and $\hat{p}_k + \hat{q}_k = 1$. Hence, for every $k \in \{1, \ldots, n-1\}$, we have that $\hat{p}_k \geqslant \frac{1}{1+\alpha}$ and $\hat{q}_k \leqslant \frac{\alpha}{1+\alpha}$. This implies, from Lemma 6.2.3, that, from any state $h \in \Omega$,

$$\mathbf{P}_h\left(Y_{\tau^Y_{0,n}} = n\right) \geqslant \frac{1-\alpha^h}{1-\alpha^n} \geqslant 1 - \alpha^h\,.$$

The lemma follows. □

6.2.2 Convergence time at low temperature

If X_t is the logit dynamics for the Ising game, the magnetization process $Y_t = \Lambda(X_t)$ is itself a Markov chain, with state space $\Omega = \{-n, -n+2, \ldots, n-4, n-2, n\}$. When at state $k \in \Omega$, the probability to go right (to state $k+2$) or left (to state $k-2$) is respectively

$$\begin{aligned}\mathbf{P}_k\left(Y_1 = k+2\right) = p_k &= \frac{n-k}{2n}\,\frac{1}{1+e^{-2(k+1)\beta}}\,;\\ \mathbf{P}_k\left(Y_1 = k-2\right) = q_k &= \frac{n+k}{2n}\,\frac{1}{1+e^{2(k-1)\beta}}\,.\end{aligned} \tag{6.4}$$

Indeed, let us evaluate the probability to jump from a profile $\mathbf{x}$ with magnetization k to a profile with magnetization $k+2$. If $\Lambda(\mathbf{x}) = k$, then there are $(n+k)/2$ players playing $+1$ and $(n-k)/2$ players playing -1. The chain moves to a profile with magnetization $k+2$ if a player playing -1 is selected, this happens with probability $(n-k)/2n$, and she updates her strategy to $+1$, this happens with probability

$$\frac{e^{\beta u_i(\mathbf{x}_{-i},+1)}}{e^{\beta u_i(\mathbf{x}_{-i},+1)} + e^{\beta u_i(\mathbf{x}_{-i},-1)}} = \frac{1}{1 + e^{\beta[u_i(\mathbf{x}_{-i},-1) - u_i(\mathbf{x}_{-i},+1)]}}\,.$$

Finally observe that $u_i(\mathbf{x}_{-i},-1) - u_i(\mathbf{x}_{-i},+1) = -2\sum_{j\neq i} x_j = -2\left(\Lambda(\mathbf{x}) - x_i\right) = -2(k+1)$.

For $a, b \in [-n, n]$, with $a < b$, let $\tau_{a,b}$ be the random variable indicating the first time the chain reaches a state h with $h \leqslant a$ or $h \geqslant b$,

$$\tau_{a,b} = \inf\{t \in \mathbb{N} \colon Y_t \leqslant a \text{ or } Y_t \geqslant b\}\,.$$

At time $\tau_{a,b}$, chain $Y_{\tau_{a,b}}$ can be in one out of two states, namely the larger state smaller than a or the smallest state larger than b. We need to give an upper bound on the probability that when the chain exits from interval (a, b), it happens on the left side of the interval.

In the next lemma we show that, if the chain starts from a sufficiently large positive state k, and if $\beta k^2 \geqslant c \log n$ for a suitable constant c, then when chain Y_t gets out of interval $(0, n/2)$, it happens on the $n/2$ side w.h.p.

Lemma 6.2.5. *Consider the birth and death chains $\{Y_t\}$ with state space Ω. Let $k \in \Omega$ be the starting state with $4 \leqslant k \leqslant n/2$. If $\beta \geqslant 6/n$ and $\beta k^2 \geqslant 16 \log n$, then*

$$\mathbf{P}_k\left(Y_{\tau_{0,n/2}} \leqslant 0\right) \leqslant 1/n\,.$$

Proof. According to (6.4), the ratio of q_h and p_h is

$$\frac{q_h}{p_h} = \frac{n+h}{n-h} \cdot \frac{1 + e^{-2(h+1)\beta}}{1 + e^{2(h-1)\beta}}\,.$$

Now observe that for all $h \geqslant 2$ it holds that

$$\frac{1 + e^{-2(h+1)\beta}}{1 + e^{2(h-1)\beta}} \leqslant e^{-2(h-1)\beta} \leqslant e^{-h\beta}\,, \tag{6.5}$$

and for all $h \leqslant n/2$ it holds that

$$\frac{n+h}{n-h} = \frac{1 + h/n}{1 - h/n} \leqslant e^{3h/n}\,.$$

Hence, for every $2 \leqslant h \leqslant n/2$ we can give the following upper bound

$$\frac{q_h}{p_h} \leqslant e^{3h/n} \cdot e^{-\beta h} = e^{-(\beta - 3/n)h} \leqslant e^{-\frac{1}{2}\beta h}\,, \tag{6.6}$$

where in the last inequality we used $\beta \geqslant 6/n$.

Thus, for each state h of the chain with $k/2 \leqslant h \leqslant n/2$ we have that the ratio q_h/p_h is less than $e^{-\frac{1}{4}\beta k}$. If the chain starts at k, by applying Lemma 6.2.4 it follows that the probability of reaching $k/2$ before reaching $n/2$ is less than $\left(e^{-\frac{1}{4}\beta k}\right)^{\ell}$, where ℓ is the number of states between $k/2$ and k, that is $\ell = k/4$. Hence, for every $4 \leqslant k \leqslant n/2$, if $\beta k^2 \geqslant 16 \log n$, the chain starting at k hits state $k/2$ before state $n/2$ with probability

$$\mathbf{P}_k\left(Y_{\tau_{k/2,n/2}} \leqslant k/2\right) \leqslant e^{\frac{1}{16}\beta k^2} \leqslant \frac{1}{n}\,.$$

The thesis follows by observing that $\mathbf{P}_k\left(Y_{\tau_{0,n/2}} \leqslant 0\right) \leqslant \mathbf{P}_k\left(Y_{\tau_{k/2,n/2}} \leqslant k/2\right)$. $\square$

In the next lemma we show that, if the chain starts from a state $k \geqslant n/2$, and if $\beta \geqslant c \log n/n$ for a suitable constant c, then when chain Y_t reaches one of the endpoints of interval $(0, n)$ it is on the n side with probability exponentially close to 1.

Lemma 6.2.6. *Consider the birth and death chains $\{Y_t\}$ with state space Ω. Let $k \in \Omega$ be the starting state with $n/2 \leqslant k \leqslant n-1$. If $\beta \geqslant 8 \log n/n$, then*

$$\mathbf{P}_k\left(Y_{\tau_{0,n}} \leqslant 0\right) \leqslant (2/n)^{n/8}.$$

Proof. Observe that for every $h \leqslant n-1$ it holds that $\frac{n+h}{n-h} \leqslant 2n$, and by using it together with (6.5) we have that $q_h/p_h \leqslant 2ne^{-\beta h}$ for every $2 \leqslant h \leqslant n-1$. Thus, for every $k/2 \leqslant h \leqslant n-1$ it holds that

$$q_h/p_h \leqslant 2ne^{-\frac{1}{2}\beta k} \leqslant 2/n\,, \tag{6.7}$$

where in the last inequality we used $k \geqslant n/2$ and $\beta \geqslant 8 \log n/n$. Hence, if the chain starts at k, by applying Lemma 6.2.4 it follows that the probability of reaching $k/2$ before reaching n is less than $(2/n)^{\ell}$, where $\ell = k/4 \geqslant n/8$ is the number of states between $k/2$ and k. Hence,

$$\mathbf{P}_k\left(Y_{\tau_{0,n}} \leqslant 0\right) \leqslant \mathbf{P}_k\left(Y_{\tau_{k/2,n}} \leqslant k/2\right) \leqslant (2/n)^{n/8}.$$

□

In the next lemma we show that for every starting state between 0 and n, the expected time the chain reaches 0 or n is at most $\mathcal{O}(n^3)$.

Lemma 6.2.7. *Consider the birth and death chains $\{Y_t\}$ with state space Ω. If $\beta \geqslant 8 \log n/n$, for every $k \in \Omega$ with $k \geqslant 2$ it holds that $\mathbf{E}_k\left[\tau_{0,n}\right] \leqslant n^3$.*

Proof. Consider the birth-and-death chain $\{Y_t^\star\}$ on $\Omega_+ = \{l \in \Omega \mid l \geqslant 0\}$ and let $\tau_n^\star$ the hitting time of n in this chain. It is obvious that $\mathbf{E}_k\left[\tau_{0,n}\right] \leqslant \mathbf{E}_k\left[\tau_n^\star\right]$. It is well-known (see, for example, Section 2.5 in [75]) that

$$\mathbf{E}_k\left[\tau_n^\star\right] = \sum_{l=\frac{k+2}{2}}^{n/2} \frac{1}{q_{2l} w_{2l}} \sum_{j=\frac{n \bmod 2}{2}}^{l-1} w_{2j}\,,$$

where $w_{n \bmod 2} = 1$ and $w_{2j} = \prod_{i=1}^{j} p_{2(j-1)}/q_{2j}$. From simple computations, we obtain

$$\begin{aligned}
\mathbf{E}_k\left[\tau_n^\star\right] &= \sum_{l=\frac{k+2}{2}}^{n/2} \sum_{j=\frac{n \bmod 2}{2}}^{l-1} \frac{1}{p_{2j}} \prod_{i=j+1}^{l-1} \frac{q_{2i}}{p_{2i}} \\
&\leqslant \sum_{l=\frac{k+2}{2}}^{n/2} \sum_{j=\frac{n \bmod 2}{2}}^{l-1} \frac{1}{p_{2j}}\,,
\end{aligned}$$

where the inequality follows from (6.6) and (6.7). Finally, the Lemma follows by observing that $p_{2j} \geqslant \frac{1}{2n}$ for every $j \geqslant 0$. □

Now we can state and prove the main theorem of this section.

Theorem 6.2.8. *Let $\mathcal{G}$ be the n-player Ising game and consider the logit dynamics for $\mathcal{G}$. If $\beta \geqslant c \log n/n$ and $k^2 > c \log n/\beta$ then*

$$t_{\pi_+}^{S_k}(\mathcal{O}(1/n)) \leqslant n^4\,, \quad \textit{and} \quad t_{\pi_-}^{S_{-k}}(\mathcal{O}(1/n)) \leqslant n^4\,,$$

where S_k and S_{-k} are the sets of profiles $\mathbf{x} \in \{-1,+1\}^n$ whose magnetization $\Lambda(\mathbf{x}) \geqslant k$ and $\Lambda(\mathbf{x}) \leqslant -k$, respectively.

Proof. We only consider starting states with positive magnetization: by symmetry, the results holds also for starting states with negative magnetization.

Let τ_n be the first time the chain hits state with all $+1$ and let $\tau_{0,n}$ be the first time the magnetization of the chain is either n or less than or equal to 0,

$$\tau_{0,n} = \min\{t \in \mathbb{N} \,:\, \Lambda(X_t) = n \text{ or } \Lambda(X_t) \leqslant 0\}\,.$$

Since $\{\tau_n > t,\, \tau_{0,n} \leqslant t\}$ implies that the magnetization chain reaches 0 before reaching n we have, that for every starting profile $\mathbf{x}$ with magnetization $l \geqslant k$

$$\begin{aligned}
\mathbf{P}_{\mathbf{x}}\left(\tau_n > t\right) &= \mathbf{P}_{\mathbf{x}}\left(\tau_n > t,\, \tau_{0,n} > t\right) + \mathbf{P}_{\mathbf{x}}\left(\tau_n > t,\, \tau_{0,n} \leqslant t\right) \\
&\leqslant \mathbf{P}_{\mathbf{x}}\left(\tau_{0,n} > t\right) + \mathbf{P}_{\mathbf{x}}\left(\Lambda(X_{\tau_{0,n}}) \leqslant 0\right) \\
&\leqslant \frac{\mathbf{E}_{\mathbf{x}}\left[\tau_{0,n}\right]}{t} + \mathbf{P}_l\left(Y_{\tau_{0,n}} \leqslant 0\right)\,,
\end{aligned}$$

where Y_t is the birth-and-death chain with state space Ω and transition rates as in (6.4). As for the first term of the sum, from Lemma 6.2.7 it follows that $\mathbf{E}_{\mathbf{x}}\left[\tau_{0,n}\right]/t \leqslant 1/n$ for $t \geqslant n^4$. As for the second term, by conditioning on the position of the chain when it gets out of subinterval $(0, n/2)$ we have

$$\begin{aligned}
\mathbf{P}_l\left(Y_{\tau_{0,n}} \leqslant 0\right) &= \mathbf{P}_l\left(Y_{\tau_{0,n}} \leqslant 0 \,|\, Y_{\tau_{0,n/2}} \leqslant 0\right)\mathbf{P}_l\left(Y_{\tau_{0,n/2}} \leqslant 0\right) + \\
&\quad + \mathbf{P}_l\left(Y_{\tau_{0,n}} \leqslant 0 \,|\, Y_{\tau_{0,n/2}} \geqslant n/2\right)\mathbf{P}_l\left(Y_{\tau_{0,n/2}} \geqslant n/2\right) \\
&\leqslant \mathbf{P}_l\left(Y_{\tau_{0,n/2}} \leqslant 0\right) + \mathbf{P}_l\left(Y_{\tau_{0,n}} \leqslant 0 \,|\, Y_{\tau_{0,n/2}} \geqslant n/2\right)\,.
\end{aligned}$$

From Lemma 6.2.5 we have that $\mathbf{P}_l\left(Y_{\tau_{0,n/2}} \leqslant 0\right) \leqslant 1/n$, and observe that

$$\begin{aligned}
\mathbf{P}_l\left(Y_{\tau_{0,n}} \leqslant 0 \,|\, Y_{\tau_{0,n/2}} \geqslant n/2\right) &\leqslant \mathbf{P}_{n/2}\left(Y_{\tau_{0,n}} \leqslant 0\right) \\
&\leqslant (2/n^{n/8})\,,
\end{aligned}$$

where the last inequality follows from Lemma 6.2.6. Hence for every $t \geqslant n^4$ it holds that $\mathbf{P}_{\mathbf{x}}\left(\tau_n > t\right) \leqslant 3/n$. Hence, for every starting profile $\mathbf{x}$ with magnetization $l \geqslant k$ and for $t = n^4$ we have

$$\left\|P^t(\mathbf{x}, \cdot) - \pi_+\right\|_{\mathrm{TV}} \leqslant \left\|\pi_+ P^t - \pi_+\right\|_{\mathrm{TV}} + \frac{1}{2}\mathbf{P}_{\mathbf{x}}\left(\tau_n > t\right) = \mathcal{O}(1/n)\,,$$

where we used that π_+ is $(1/n, n^{c-2})$-metastable when $\beta \geqslant c \log n/n$. $\square$

6.3 Graphical coordination games on rings

In this section we study metastability for graphical coordination games on a ring.

The set of profiles of the game is $S = \{0,1\}^n$. We define $S_d \subseteq S$ as the set of profiles where *exactly* d players are playing 0 and $R \subseteq S$ as the set of profiles in which at least two *adjacent* players are playing 0. We also set $\tilde{S}_d = \bigcup_{i=d}^n S_i$.

In previous chapter it is showed that the mixing time of the logit dynamics for this game is polynomial in n for $\beta = \mathcal{O}(\log n)$ and greater than any polynomial in n, for $\beta = \omega(\log n)$.

6.3.1 Games with risk dominant strategies

In this section we study the case $\Delta > \delta$ and prove that for $\beta = \omega(\log n)$, the logit dynamics reaches in polynomial time a metastable distribution and remains close to it for super-polynomial time. On the other hand, we know that for $\beta = \mathcal{O}(\log n)$ the logit dynamics is rapidly mixing (and thus reaches in polynomial time the stationary distribution and stays close to it forever).

Theorem 6.3.1. *Let $\mathcal{G}$ be the n-player graphical coordination game on the ring. If $\beta = \omega(\log n)$ then for every $\mathbf{x} \in S$ there exists a distribution μ that is (ε, T)-metastable for the logit dynamics on $\mathcal{G}$, where $\varepsilon < 2/5$ and $T = T(n)$ is a super-polynomial function, and the dynamics starting from $\mathbf{x}$ approaches μ in polynomial time.*

Throughout this section, we set $S_d^\star = R \cup \tilde{S}_d$.

Metastable distributions

The next theorem identifies three classes of metastable distributions. We remind the reader that π_S is the stationary distribution (that is the Gibbs measure defined in 2.3) restricted to the set S of profiles (see 6.1).

Theorem 6.3.2. *Let $\mathcal{G}$ be the n-player graphical coordination game on the ring. For every $\varepsilon > 0$, $\pi_{\mathbf{1}}$ is $(\varepsilon, \varepsilon \cdot e^{2\delta\beta})$-metastable for the logit dynamics on $\mathcal{G}$ and $\pi_{\mathbf{0}}$ is $(\varepsilon, \varepsilon \cdot e^{2\Delta\beta})$-metastable. Moreover, for $\beta = \omega(\log n)$ and constant $d > 0$, $\pi_{S_d^\star}$ is $(\varepsilon, e^{\Omega(n \log n)})$-metastable.*

Proof. The bottleneck ratio of $\mathbf{1}$ is

$$B(\mathbf{1}) = \frac{\sum_{\mathbf{x} \neq \mathbf{1}} \pi(\mathbf{1}) P(\mathbf{1}, \mathbf{x})}{\pi(\mathbf{1})} = e^{-2\delta\beta}\,.$$

Thus from Lemma 6.1.6, we have that $\pi_{\mathbf{1}}$ is $(e^{-2\delta\beta}, 1)$-metastable. By applying Lemma 6.1.2, we obtain that $\pi_{\mathbf{1}}$ is $(\varepsilon, \varepsilon \cdot e^{2\delta\beta})$-metastable for every $\varepsilon > 0$.

Similarly, the bottleneck ratio of $\mathbf{0}$ is

$$B(\mathbf{0}) = \frac{\sum_{\mathbf{x} \neq \mathbf{0}} \pi(\mathbf{0}) P(\mathbf{0}, \mathbf{x})}{\pi(\mathbf{0})} = e^{-2\Delta\beta}\,,$$

and thus, by applying Lemma 6.1.6 and Lemma 6.1.2, we have that $\pi_{\mathbf{0}}$ is $(\varepsilon, \varepsilon \cdot e^{2\Delta\beta})$-metastable for every ε.

Finally, the bottleneck ratio of $S_d^\star$ is

$$\begin{aligned} B(S_d^\star) &= \frac{\sum_{\mathbf{x} \in \partial S_d^\star} \pi(\mathbf{x}) \sum_{\mathbf{y} \in \S \setminus S_d^\star} P(\mathbf{x}, \mathbf{y})}{\sum_{\mathbf{x} \in S_d^\star} \pi(\mathbf{x})} \\ &\leqslant \frac{\sum_{\mathbf{x} \in \partial S_d^\star} \pi(\mathbf{x})}{\sum_{\mathbf{x} \in S_d^\star} \pi(\mathbf{x})} \\ &\leqslant \frac{n^{d+1} \max_{\mathbf{x} \in \partial S_d^\star} \pi(\mathbf{x})}{\max_{\mathbf{x} \in S_d^\star} \pi(\mathbf{x})} \\ &\leqslant \frac{n^{d+1} e^{[(n-d-1)\delta + (d-1)\Delta]\beta}}{e^{n\Delta\beta}} \\ &\leqslant n^{d+1} e^{-(n-d-1)(\Delta - \delta)\beta} \\ &\leqslant n^{-n+2(d+1)} = e^{-\Omega(n \log n)}\,. \end{aligned}$$

From the second to the third line we used that $\partial S_d^\star$, the set of all profiles in $S_d^\star$ with at least a neighbor in $S \setminus S_d^\star$, has size at most n^{d+1}; from the third to the fourth line we used that the maximum is attained when d 0's are adjacent; then we used that $\beta = \omega(\log n)$ and $(\Delta - \delta)$ and d are positive constants.

Thus, by applying Lemma 6.1.6 and Lemma 6.1.2, we have that $\pi_{S_d^\star}$ is $(\varepsilon, e^{\Omega(n \log n)})$-metastable for every ε. $\square$

Pseudo-mixing time

In this subsection we look at the the pseudo-mixing time of the metastable distributions described in Theorem 6.3.2 and we show that, for every starting profile, the dynamics rapidly approaches one of them. We remind the reader that the interesting case is $\beta = \omega(\log n)$ as for $\beta = \mathcal{O}(\log n)$ the mixing time of the logit dynamics is polynomial in n.

Not to overburden our notation, we will denote distribution $\pi_{S_d^\star}$ by π_d.

Our proof distinguishes cases depending on the starting profile $\mathbf{x}$ of the chain. We start by considering $\mathbf{x} \in S_d^\star$, for constant d, and show (see Theorem 6.3.7) that the pseudo-mixing time of π_d is polynomial. Finally, in Theorem 6.3.9, we show that if the dynamics starts from one of the remaining profiles, in polynomial time it hits either $S_d^\star$ or $\mathbf{1}$ with high probability. For sake of readability we postpone proofs to Section 6.3.3.

Starting from $S_d^\star$. From the definition of pseudo mixing time and by using Lemma 6.1.7, for any S, we can bound $t_{\pi_S}^S(\gamma)$ in terms of the total variation distance of two copies of the same Markov chain starting in different states. Next lemma relates this quantity, for the logit dynamics we are studying, to the hitting time $\tau_{\mathbf{0}}$ of profile $\mathbf{0}$.

Lemma 6.3.3. *Let $\mathcal{G}$ be the n-player graphical coordination game on the ring and P be the transition matrix of the logit dynamics for $\mathcal{G}$. For every $T \subseteq S$, $t > 0$ and for $\mathbf{x}, \mathbf{y} \in T$ we have*

$$\left\| P^t(\mathbf{x}, \cdot) - P^t(\mathbf{y}, \cdot) \right\|_{\mathrm{TV}} \leqslant 2 \cdot \max_{\mathbf{z} \in T} \mathbf{P_z}\left(\tau_{\mathbf{0}} \geqslant t\right) .$$

Thus, from the previous lemma and Lemma 6.1.7, in order to bound $t_{\pi_d}^{S_d^\star}(\gamma)$, it is sufficient to give an upper bound on $\mathbf{P_x}\left(\tau_{\mathbf{0}} \geqslant t\right)$, for $\mathbf{x} \in S_d^\star$. The next lemma bounds the hitting time of $\mathbf{0}$ when starting from R.

Lemma 6.3.4. *Let $\mathcal{G}$ be the n-player graphical coordination game on the ring and consider the logit dynamics for $\mathcal{G}$. For $\beta = \omega\left(\log n\right)$, for every $\lambda > 0$ and every $\mathbf{x} \in R$, we have*

$$\mathbf{P_x}\left(\tau_{\mathbf{0}} > \frac{(8-\lambda)n^2}{\lambda}\right) \leqslant \frac{\lambda}{4} .$$

Next we show that, when starting from $\mathbf{x} \in S_d$, the dynamics hits R in polynomial time.

Lemma 6.3.5. *Let $\mathcal{G}$ be the n-player graphical coordination game on the ring and consider the logit dynamics for $\mathcal{G}$. For $\beta = \omega\left(\log n\right)$, for every $d > 0$ and for every profile $\mathbf{x} \in S_d$,*

$$\mathbf{P_x}\left(\tau_R \leqslant n^2\right) \geqslant \frac{2d}{2d+1}\left(1 - \mathsf{negl}\left(n\right)\right) .$$

Finally, we have

Lemma 6.3.6. *Let $\mathcal{G}$ be the n-player graphical coordination game on the ring and consider the logit dynamics for $\mathcal{G}$. For $\beta = \omega\left(\log n\right)$, for every $d > 0$, for $\mathbf{x} \in S_d^\star$ and for every $\lambda > 0$,*

$$\mathbf{P_x}\left(\tau_{\mathbf{0}} > \frac{8n^2}{\lambda}\right) \leqslant \frac{1}{2d+1} + \frac{\lambda}{4} .$$

Proof. We need to consider only $\mathbf{x} \in S_d^\star \setminus R$. By Lemma 6.3.4 and Lemma 6.3.5, we have

$$\begin{aligned}
\mathbf{P_x}\left(\tau_\mathbf{0} \leqslant \frac{8n^2}{\lambda}\right) &\geqslant \mathbf{P_x}\left(\tau_\mathbf{0} \leqslant \frac{8n^2}{\lambda} \mid \tau_R \leqslant n^2\right) \mathbf{P_x}\left(\tau_R \leqslant n^2\right) \\
&\geqslant \mathbf{P}_{X_{\tau_R}}\left(\tau_\mathbf{0} \leqslant \frac{(8-\lambda)n^2}{\lambda}\right) \mathbf{P_x}\left(\tau_R \leqslant n^2\right) \\
&\geqslant \left(1 - \frac{\lambda}{4}\right) \frac{2d}{2d+1} \left(1 - \mathsf{negl}\left(n\right)\right) \\
&\geqslant \frac{2d}{2d+1} - \frac{\lambda}{4}\,.
\end{aligned}$$

□

We are now ready to prove an upper bound on the pseudo-mixing time of π_d, when starting from a profile in $S_d^\star$.

Theorem 6.3.7. *Let $\mathcal{G}$ be the n-player graphical coordination game on the ring and consider the logit dynamics for $\mathcal{G}$. For $\beta = \omega\left(\log n\right)$, for constant $d > 1$ and for every $\lambda > 0$,*

$$t_{\pi_d}^{S_d^\star}(\gamma) \leqslant \frac{8n^2}{\lambda}\,,$$

where $\gamma = \frac{2}{2d+1} + \lambda$.

Proof. By Theorem 6.3.2, we have that π_d is $(\lambda/2, 8n^2/\lambda)$-metastable for every $\lambda > 0$ and sufficiently large n. Therefore, for every $\mathbf{x} \in S_d^\star$ and $t \geqslant 8n^2/\lambda$, we have

$$\begin{aligned}
\left\|P^t(\mathbf{x},\cdot) - \pi_d\right\|_{\mathrm{TV}} &\leqslant \frac{\lambda}{2} + \max_{\mathbf{y}\in S_d^\star} \left\|P^t(\mathbf{x},\cdot) - P^t(\mathbf{y},\cdot)\right\|_{\mathrm{TV}} \\
&\leqslant \frac{\lambda}{2} + \max_{\mathbf{z}\in S_d^\star} \mathbf{P_z}\left(\tau_\mathbf{0} \geqslant t\right) \\
&\leqslant \frac{2}{2d+1} + \lambda\,.
\end{aligned}$$

where the first inequality follows from Lemma 6.1.7, the second one from Lemma 6.3.3, and the third one follows from Lemma 6.3.4 and Corollary 6.3.6. □

Starting from outside $S_d^\star$. Observe that when $\mathbf{x} = \mathbf{1}$, metastable distribution $\pi_\mathbf{1}$ is trivially reached immediately. Thus, it only remains to analyze $\mathbf{x} \notin S_d^\star \cup \{\mathbf{1}\}$. For this, it is enough to prove that for such an $\mathbf{x}$ the hitting time of $S_d^\star \cup \{\mathbf{1}\}$ is polynomial.

Lemma 6.3.8. *Let $\mathcal{G}$ be the n-player graphical coordination game on the ring and consider the logit dynamics for $\mathcal{G}$. For $\beta = \omega\left(\log n\right)$, for every $d > 0$ and for every $\mathbf{x} \in S_d$, we have*

$$\mathbf{P_x}\left(\tau_\mathbf{1} \leqslant n^2\right) \geqslant \frac{1}{3^d}\left(1 - \mathsf{negl}\left(n\right)\right)\,.$$

We can now state the following theorem.

Theorem 6.3.9. *Let $\mathcal{G}$ be the n player graphical coordination game on the ring and consider the logit dynamics for $\mathcal{G}$. For $\beta = \omega\left(\log n\right)$, for every $d > z > 0$ and for every profile $\mathbf{x} \in S_z$, we have*

$$\mathbf{P_x}\left(\tau_{S_d^\star \cup \{\mathbf{1}\}} \leqslant n^2\right) \geqslant \left(\frac{2z}{2z+1} + \frac{1}{3^z}\right)\left(1 - \mathsf{negl}\left(n\right)\right)\,.$$

Theorem 6.3.1 follows from Theorems 6.3.7 and 6.3.9.

Staying arbitrarily close. We observe that, in Theorem 6.3.7, the distance between the dynamics and the metastable distribution cannot be made arbitrarily small. We can achieve this at the cost of slightly reducing the set of starting states from which convergence is proved. Specifically, the next theorem shows that, for $d = \omega(1)$ and arbitrarily small $\gamma > 0$, the logit dynamics starting from $S_d^\star$ is within distance γ from $\pi_{\mathbf{0}}$ in a number of steps that is polynomial in n and in $1/\gamma$.

Theorem 6.3.10. *Let $\mathcal{G}$ be the n-player graphical coordination game on the ring and consider the logit dynamics for $\mathcal{G}$. For $\beta = \omega(\log n)$, $d = \omega(1)$ and every $\gamma > 0$,*

$$t_{\pi_{\mathbf{0}}}^{S_d^\star}(\gamma) \leqslant \frac{8n^2}{\gamma} \,.$$

Proof. Since $\beta = \omega(\log n)$, Theorem 6.3.2 implies that $\pi_{\mathbf{0}}$ is $(\gamma/2, 8n^2/\gamma)$-metastable for every $\gamma > 0$ and sufficiently large n. Therefore, for every $\mathbf{x} \in S_d^\star$ and $t \geqslant 8n^2/\gamma$, we have

$$\begin{aligned} \left\| P^t(\mathbf{x}, \cdot) - \pi_{\mathbf{0}} \right\|_{\mathrm{TV}} &\leqslant \frac{\gamma}{2} + \mathbf{P}_{\mathbf{x}}\left(\tau_{\mathbf{0}} > t\right) \\ &\leqslant \frac{1}{2d+1} + \frac{3}{4}\gamma \leqslant \gamma \,. \end{aligned}$$

where the first inequality follows from Lemma 6.1.8, the second one from Lemma 6.3.4 and Corollary 6.3.6, and the third one holds because $d = \omega(1)$. □

6.3.2 Games without risk dominant strategies.

In this section we study the case of graphical coordination games without risk dominant strategies (that is, $\Delta = \delta$) played on a ring by n players. Next theorem identifies a class of metastable distributions.

Theorem 6.3.11. *Let $\mathcal{G}$ be the n-player graphical coordination game on the ring and consider the logit dynamics for $\mathcal{G}$. For every $\varepsilon > 0$ and for every $0 \leqslant d \leqslant n$, distribution $\mu_d = \frac{d}{n}\pi_{\mathbf{0}} + \left(1 - \frac{d}{n}\right)\pi_{\mathbf{1}}$ is $(\varepsilon, \varepsilon e^{2\Delta\beta})$-metastable.*

Proof. We notice that

$$\begin{aligned} \|\mu_d P - \mu_d\|_{\mathrm{TV}} &= \frac{1}{2}\sum_{\mathbf{x}} |(\mu_d P)(\mathbf{x}) - \mu_d(\mathbf{x})| \\ &= \frac{1}{2}\sum_{\mathbf{x}} \left| \sum_{\mathbf{y}} \mu_d(\mathbf{y}) P(\mathbf{y}, \mathbf{x}) - \mu_d(\mathbf{x}) \right| \\ &= d \sum_{\mathbf{x} \in S_{n-1}} P(\mathbf{0}, \mathbf{x}) + (1-d) \sum_{\mathbf{x} \in S_1} P(\mathbf{1}, \mathbf{x}) \\ &= \frac{1}{1 + e^{2\Delta\beta}} \,. \end{aligned}$$

Thus μ is $\left(\frac{1}{1+e^{2\Delta\beta}}, 1\right)$-metastable. The Theorem follows from Lemma 6.1.2. □

The main and quite surprising result in this section is that for every starting profile $\mathbf{x} \in S_d$ the dynamics starting in $\mathbf{x}$ converges in polynomial time to μ_d, for $d = 1, \ldots, n$. In order to prove this result, we define $\tau_{\mathbf{0},\mathbf{1}} = \min\{\tau_{\mathbf{0}}, \tau_{\mathbf{1}}\}$ and prove that this quantity is polynomial in n with very high probability; then we show that the dynamics starting at $\mathbf{x} \in S_d$ after $\tau_{\mathbf{0},\mathbf{1}}$ steps is distributed as a metastable distribution very close to μ_d. We formalize these arguments in two technical lemmas, whose proofs will be given in Section 6.3.4.

Lemma 6.3.12. *Let $\mathcal{G}$ be n-player graphical coordination game on the ring and consider the logit dynamics for $\mathcal{G}$. If $\beta = \omega(\log n)$, then for every $\mathbf{x} \in S$, $\mathbf{P}_{\mathbf{x}}\left(\tau_{\mathbf{0},\mathbf{1}} \leqslant n^5\right) \geqslant 1 - o(1)$.*

Lemma 6.3.13. *Let $\mathcal{G}$ be the n-player graphical coordination game on the ring and consider the logit dynamics for $\mathcal{G}$. For every d, $\mathbf{x} \in S_d$ and $\beta = \omega(\log n)$, the random variable $X_{\tau_{\mathbf{0},\mathbf{1}}}$ given that $X_0 = \mathbf{x}$, has distribution $\nu_x = \left(\frac{d}{n} + \lambda_{\mathbf{x}}\right)\pi_{\mathbf{0}} + \left(1 - \frac{d}{n} - \lambda_{\mathbf{x}}\right)\pi_{\mathbf{1}}$, with $|\lambda_{\mathbf{x}}| = o(1)$.*

The pseudo mixing time of distributions μ_d, for $d = 0, 1, \ldots, n$, is given by the next Theorem.

Theorem 6.3.14. *Let $\mathcal{G}$ be the n-player graphical coordination game on the ring and consider the logit dynamics for $\mathcal{G}$. If $\beta = \omega(\log n)$, for every d and every $\gamma > 0$*

$$t_{\mu_d}^{S_d}(\gamma) \leqslant n^5.$$

Proof. From Lemma 6.3.12, for every $\mathbf{x} \in S$ we have

$$\begin{aligned}
\left\|P^{n^5}(\mathbf{x}, \cdot) - \mu_d\right\|_{\mathrm{TV}} &= \max_{A \subset S} \left|\mathbf{P}_{\mathbf{x}}\left(X_{n^5} \in A\right) - \mu_d(A)\right| \\
&= o(1) + \max_{A \subset S} \left|\mathbf{P}_{\mathbf{x}}\left(X_{n^5} \in A \mid \tau_{\mathbf{0},\mathbf{1}} \leqslant n^5\right) - \mu_d(A)\right| \\
&= o(1) + \left\|\mathbf{P}_{\mathbf{x}}\left(X_{n^5} \mid \tau_{\mathbf{0},\mathbf{1}} \leqslant n^5\right) - \mu_d\right\|_{\mathrm{TV}} \\
&\leqslant o(1) + \left\|\mathbf{P}_{\mathbf{x}}\left(X_{n^5} \mid \tau_{\mathbf{0},\mathbf{1}} \leqslant n^5\right) - \nu_{\mathbf{x}}\right\|_{\mathrm{TV}} + \left\|\nu_{\mathbf{x}} - \mu_d\right\|_{\mathrm{TV}} ,
\end{aligned}$$

where the last inequality follows from the triangle inequality of the total variation distance. Moreover, from Lemma 6.3.13, for every $\mathbf{x} \in S_d$, we have

$$\|\nu_{\mathbf{x}} - \mu_d\|_{\mathrm{TV}} = \|P^{\tau_{\mathbf{0},\mathbf{1}}}(\mathbf{x}, \cdot) - \mu_d\|_{\mathrm{TV}} = o(1)\,.$$

Finally, from Theorem 6.3.11 we have that $\mu_{\mathbf{x}}$ is $(o(1), n^5)$-metastable and thus,

$$\begin{aligned}
\left\|\mathbf{P}_{\mathbf{x}}\left(X_{n^5} \mid \tau_{\mathbf{0},\mathbf{1}} \leqslant n^5\right) - \nu_{\mathbf{x}}\right\|_{\mathrm{TV}} &= \left\|P^{\tau_{\mathbf{0},\mathbf{1}}}(\mathbf{x}, \cdot) P^{n^5 - \tau_{\mathbf{0},\mathbf{1}}} - \nu_{\mathbf{x}}\right\|_{\mathrm{TV}} \\
&= \left\|\nu_{\mathbf{x}} P^{n^5 - \tau_{\mathbf{0},\mathbf{1}}} - \nu_{\mathbf{x}}\right\|_{\mathrm{TV}} = o(1)\,.
\end{aligned}$$

□

6.3.3 Proofs from Section 6.3.1

Proof of Lemma 6.3.3

Proof of Lemma 6.3.3. Consider the following partial order $\preceq$ over S: for profiles $\mathbf{x}, \mathbf{y} \in S$, $\mathbf{x} \preceq \mathbf{y}$ if and only if for every $0 \leqslant i \leqslant n-1$, we have that $x_i \geqslant y_i$. That is, if $\mathbf{x} \preceq \mathbf{y}$ then $\mathbf{x}$ can be obtained from $\mathbf{y}$ by changing 0's into 1's. We note that, in according to this order, $\mathbf{0}$ is the unique maximum.

We next show that for every two profiles $\mathbf{x}, \mathbf{y} \in S$ there exists a monotone (w.r.t. $(S, \preceq)$) coupling (X_1, Y_1) of two copies of the logit dynamics for the graphical coordination game on the ring for which $X_0 = \mathbf{x}$ and $Y_0 = \mathbf{y}$. The Lemma then follows from Theorem 1.3.1 and Lemma 1.3.2.

Consider the coupling described in Section 4.1.1: this coupling is monotone w.r.t. $(S, \preceq)$. Indeed, suppose $\mathbf{x} \preceq \mathbf{y}$ and that the player i was selected for update. Since $\mathbf{x} \preceq \mathbf{y}$, the number of neighbors of i playing 0 in $\mathbf{x}$ is less or equal than in $\mathbf{y}$ and thus $\sigma_i(0 \mid \mathbf{x}) \leqslant \sigma_i(0 \mid \mathbf{y})$. Thus, the coupling either sets $x_i = y_i$ or $x_i = 1$ and $y_i = 0$. In both cases, $X_1 \preceq Y_1$. □

Proof of Lemma 6.3.4

Lemma 6.3.4 gives an upper bound on the hitting time, $\tau_{\mathbf{0}}$, of $\mathbf{0}$, for a dynamics starting from a profile $\mathbf{x} \in R$ (profiles in R are those in which at least two adjacent players play 0). For convenience, we rename players so that $x_0 = x_1 = 0$. Intuitively, for $\beta = \omega(\log n)$, each of player 0 and 1 changes her strategy with very low probability. Moreover, player 2, when selected for update, plays 0 with high probability. Similarly, after player 2 has played 0, we have that each of player $0, 1$ and 2 changes her strategy with very low probability and player 3, when selected for update, plays 0 with high probability. This process repeats until every player is playing 0. In the following, we estimate the number of steps sufficient to have all players playing strategy 0 with high probability.

For sake of compactness, we will denote the strategy of player i at time step t by X_t^i. We start with a simple observation that lower bounds the probability that a player picks strategy 0 when selected for update, given that at least one of their neighbors is playing 0.

Observation 6.3.15. *Let $\mathcal{G}$ be the n-player graphical coordination game on the ring and consider the logit dynamics for $\mathcal{G}$. For every player i, if i is selected for update at time t, then, for $b \in \{-1, 1\}$*

$$\mathbf{P}\left(X_t^i = 0 \mid X_t^{i+b} = 0\right) \geqslant \left(1 - \frac{1}{1 + e^{(\Delta-\delta)\beta}}\right).$$

We start by evaluating the probability that the dynamics selects players $2, \ldots, n-1$ at least once in this order before time t. To this aim, we set $\rho_1 = 0$ and, for $i = 2, \ldots, n-1$, we define ρ_i as the first time player i is selected for update after time step ρ_{i-1}. Thus, at time ρ_i player i is selected for update and players $2, \ldots, i-1$ have been selected at least once in this order. In particular, ρ_{n-1} is the first time step at which every player i, $i \geqslant 3$, has been selected at least once after his left neighbor. Obviously, $\rho_i > \rho_{i-1}$ for $i = 2, \ldots, n-1$. The next lemma lower bounds the probability that $\rho_{n-1} \leqslant t$.

Lemma 6.3.16. *Let $\mathcal{G}$ be the n-player graphical coordination game on the ring and consider the logit dynamics for $\mathcal{G}$. For every $\mathbf{x} \in R$ and every $t > 0$, we have*

$$\mathbf{P}_{\mathbf{x}}\left(\rho_{n-1} \leqslant t\right) \geqslant 1 - \frac{n^2}{t}.$$

Proof. Every player i has probability $\frac{1}{n}$ of being selected at any given time step. Therefore, $\mathbf{E}\left[\rho_2\right] = \mathbf{E}\left[\rho_2 - \rho_1\right] = n$ and $\mathbf{E}\left[\rho_i - \rho_{i-1}\right] = n$, for $i = 3, \ldots, n-1$. Thus, by linearity of expectation,

$$\mathbf{E}\left[\rho_{n-1}\right] = \sum_{i=2}^{n-1} \mathbf{E}\left[\rho_i - \rho_{i-1}\right] \leqslant n^2.$$

The lemma follows from the Markov inequality. □

Suppose now that $t \geqslant \rho_{n-1}$. The next lemma shows that, for all players i, the probability that $X_t^i = 0$ is high.

Lemma 6.3.17. *Let $\mathcal{G}$ be the n-player graphical coordination game on the ring and consider the logit dynamics for $\mathcal{G}$. For every starting profile $\mathbf{x} \in R$, for every player i and for every time step $t > 0$, we have*

$$\mathbf{P}_{\mathbf{x}}\left(X_t^i = 0 \mid \rho_{n-1} \leqslant t\right) \geqslant \left(1 - \frac{1}{1 + e^{(\Delta-\delta)\beta}}\right)^t.$$

We prove the lemma first for $i \geqslant 2$. Then we deal with players 0 and 1.

Fix player $i \geqslant 2$, time step t and set $s_{i+1} = t$. Starting from time step t and going backward to time step 0, we identify, the sequence of time steps $s_i > s_{i-1} > \ldots > s_2 > 0$ such that, for $j = i, i-1, \ldots, 2$, s_j is the last time player j has been selected before time s_{j+1}. We remark that, since $t \geqslant \rho_{n-1} > \rho_i$ we have that players $2, \ldots, i$ are selected at least once in this order and thus all the s_j^i are well defined. Strictly speaking, the sequence $s_i, \ldots, s_2$ depends on i and t and thus a more precise, and more cumbersome, notation would have been $s_{i,j,t}$. Since player i and time step t will be clear from the context, we drop i and t.

In order to lower bound the probability that $X_t^i = 0$ for $i \geqslant 2$, we first bound it in terms of the the probability that player 2 plays 0 at time s_2 and then we evaluate this last quantity. The next lemma is the first step.

Lemma 6.3.18. *Let $\mathcal{G}$ be the n-player graphical coordination game on the ring and consider the logit dynamics for $\mathcal{G}$. For every $\mathbf{x} \in R$, every player $2 \leqslant i \leqslant n-1$ and every time step t, we have*

$$\mathbf{P}_{\mathbf{x}}\left(X_t^i = 0 \mid \rho_{n-1} \leqslant t\right) \geqslant \left(1 - \frac{1}{1+e^{(\Delta-\delta)\beta}}\right)^{i-2} \mathbf{P}_{\mathbf{x}}\left(X_{s_2}^2 = 0 \mid \rho_{n-1} \leqslant t\right) .$$

Proof. For every fixed i, s_i is the last time the player i is selected for update before t and thus $X_t^i = X_{s_i}^i$. Hence, for $i = 2$ the lemma obviously holds. For $i > 2$ and $j = 2, \ldots, i$, we observe that, since $t \geqslant \rho_{n-1} > \rho_i$, player j has been selected for update at time s_j and s_j is the last time that player j is selected for update before time s_{j+1} and thus $X_{s_{j+1}}^j = X_{s_j}^j$.

From Observation 6.3.15, we have

$$\begin{aligned}
\mathbf{P}_{\mathbf{x}}\left(X_{s_j}^j = 0 \mid \rho_{n-1} \leqslant t\right) &\geqslant \mathbf{P}_{\mathbf{x}}\left(X_{s_j}^j = 0 \mid X_{s_j}^{j-1} = 0, \rho_{n-1} \leqslant t\right) \cdot \mathbf{P}_{\mathbf{x}}\left(X_{s_j}^{j-1} = 0 \mid \rho_{n-1} \leqslant t\right) \\
&\geqslant \left(1 - \frac{1}{1+e^{(\Delta-\delta)\beta}}\right) \mathbf{P}_{\mathbf{x}}\left(X_{s_{j-1}}^{j-1} = 0 \mid \rho_{n-1} \leqslant t\right) ,
\end{aligned}$$

and the lemma follows. □

We now bound the probability that player 2 plays 0 at time step s_2. If player 1 has not been selected for update before time s_2, then $X_{s_2}^1 = X_0^1 = 0$, and, from Observation 6.3.15, we have

$$\begin{aligned}
\mathbf{P}_{\mathbf{x}}\left(X_{s_2}^2 = 0 \mid \rho_{n-1} \leqslant t\right) &\geqslant \mathbf{P}_{\mathbf{x}}\left(X_{s_2}^2 = 0 \mid X_{s_2}^1 = 0, \rho_{n-1} \leqslant t\right) \\
&\geqslant \left(1 - \frac{1}{1+e^{(\Delta-\delta)\beta}}\right) .
\end{aligned}$$

It remains to consider the case when player 1 has been selected for update at least once before time s_2. For fixed player i and time step t we define a new sequence of time steps $r_0 > r_1, \ldots > 0$ in the following way. We set $r_0 = s_2$, and, starting from time step s_2 and going backward to time step 0, r_j, for $j > 0$, is the last time player $j \bmod 2$ has been selected before time r_{j-1}. For the last element in the sequence, r_k, it holds that player $(k+1) \bmod 2$ is not selected before time step r_k.

Since r_1 is the last time player 1 has been selected for update before $r_0 = s_2$, we have $X_{s_2}^1 = X_{r_1}^1$ and, by Observation 6.3.15,

$$\begin{aligned}
\mathbf{P}_{\mathbf{x}}\left(X_{s_2}^2 = 0 \mid \rho_{n-1} \leqslant t\right) &\geqslant \mathbf{P}_{\mathbf{x}}\left(X_{s_2}^2 = 0 \mid X_{s_2}^1 = 0, \rho_{n-1} \leqslant t\right) \cdot \mathbf{P}_{\mathbf{x}}\left(X_{s_2}^1 = 0 \mid \rho_{n-1} \leqslant t\right) \\
&\geqslant \left(1 - \frac{1}{1+e^{(\Delta-\delta)\beta}}\right) \cdot \mathbf{P}_{\mathbf{x}}\left(X_{r_1}^1 = 0 \mid \rho_{n-1} \leqslant t\right) .
\end{aligned} \tag{6.8}$$

Finally, we bound $\mathbf{P}_{\mathbf{x}}\left(X_{r_1}^1 = 0 \mid \rho_{n-1} \leqslant t\right)$.

Lemma 6.3.19. *Let $\mathcal{G}$ be the n-player graphical coordination game on the ring and consider the logit dynamics for $\mathcal{G}$. For every starting profile $\mathbf{x} \in R$ and for every time step t, for every fixed player i, let $r_0, \ldots, r_k$ be defined as above. If $k > 0$, we have*

$$\mathbf{P_x}\left(X^1_{r_1} = 0 \mid \rho_{n-1} \leqslant t\right) \geqslant \left(1 - \frac{1}{1 + e^{(\Delta-\delta)\beta}}\right)^k .$$

Proof. For sake of compactness, in this proof we denote the parity of integer a with $\mathcal{P}(a) = a \bmod 2$. Thus, the definition of sequence r_j gives that player $\mathcal{P}(j)$ has been selected for update at time r_j and

$$\mathbf{P_x}\left(X^{\mathcal{P}(j)}_{r_j} = 0 \mid \rho_{n-1} \leqslant t\right) \geqslant \mathbf{P_x}\left(X^{\mathcal{P}(j)}_{r_j} = 0 \mid X^{\mathcal{P}(j+1)}_{r_j} = 0, \rho_{n-1} \leqslant t\right) \cdot \mathbf{P_x}\left(X^{\mathcal{P}(j+1)}_{r_j} = 0 \mid \rho_{n-1} \leqslant t\right) .$$

If $j \neq k$ player $\mathcal{P}(j+1)$ has not been selected for update between time r_{j+1} and time r_j and by Observation 6.3.15

$$\begin{aligned}\mathbf{P_x}\left(X^{\mathcal{P}(j)}_{r_j} = 0 \mid \rho_{n-1} \leqslant t\right) &\geqslant \mathbf{P_x}\left(X^{\mathcal{P}(j)}_{r_j} = 0 \mid X^{\mathcal{P}(j+1)}_{r_j} = 0, \rho_{n-1} \leqslant t\right) \cdot \mathbf{P_x}\left(X^{\mathcal{P}(j+1)}_{r_{j+1}} = 0 \mid \rho_{n-1} \leqslant t\right) \\ &\geqslant \left(1 - \frac{1}{1 + e^{(\Delta-\delta)\beta}}\right) \mathbf{P_x}\left(X^{\mathcal{P}(j+1)}_{r_{j+1}} = 0 \mid \rho_{n-1} \leqslant t\right) .\end{aligned}$$

If $j = k$, instead, player $\mathcal{P}(k+1)$ has not been selected for update before time r_k and thus $X^{\mathcal{P}(k+1)}_{r_k} = X^{\mathcal{P}(k+1)}_0 = 0$. By Observation 6.3.15, we have

$$\begin{aligned}\mathbf{P_x}\left(X^{\mathcal{P}(k)}_{r_k} = 0 \mid \rho_{n-1} \leqslant t\right) &\geqslant \mathbf{P_x}\left(X^{\mathcal{P}(k)}_{r_k} = 0 \mid X^{\mathcal{P}(k+1)}_{r_k} = 0, \rho_{n-1} \leqslant t\right) \\ &\geqslant \left(1 - \frac{1}{1 + e^{(\Delta-\delta)\beta}}\right) .\end{aligned}$$

□

Let k be the index of the last term in the sequence $r_0, r_1, \ldots$ previously defined. Then, from Lemma 6.3.18, Equation (6.8) and Lemma (6.3.19) we have for every player $i \geqslant 2$ and for every time step $t > 0$

$$\begin{aligned}\mathbf{P_x}\left(X^i_t = 0 \mid \rho_{n-1} \leqslant t\right) &\geqslant \left(1 - \frac{1}{1 + e^{(\Delta-\delta)\beta}}\right)^{i-2} \mathbf{P_x}\left(X^2_{s_2} = 0 \mid \rho_{n-1} \leqslant t\right) \\ &\geqslant \left(1 - \frac{1}{1 + e^{(\Delta-\delta)\beta}}\right)^{i-1} \mathbf{P_x}\left(X^1_{r_1} = 0 \mid \rho_{n-1} \leqslant t\right) \\ &\geqslant \left(1 - \frac{1}{1 + e^{(\Delta-\delta)\beta}}\right)^{i-1+k} \\ &\geqslant \left(1 - \frac{1}{1 + e^{(\Delta-\delta)\beta}}\right)^{t} ,\end{aligned}$$

where in the last inequality we used $i - 1 + k \leqslant t$.

This ends the proof of Lemma 6.3.17 for player $i \geqslant 2$. The cases $i = 0, 1$ can be proved in similar way. Clearly, if player i has never been selected for update before time t, we have that $X^i_t = 0$ with probability 1. If player i has been selected at least once we have to distinguish the cases $i = 0$ and $i = 1$. If $i = 1$, we define $r_0 = t + 1$ and we identify a sequence of time step $r_1 > r_2 > \ldots > 0$ as above: we have that $X^1_t = X^1_{r_1}$ and from Lemma 6.3.19 follows that

$$\mathbf{P_x}\left(X^i_t = 0 \mid \rho_{n-1} \leqslant t\right) \geqslant \left(1 - \frac{1}{1 + e^{(\Delta-\delta)\beta}}\right)^k \geqslant \left(1 - \frac{1}{1 + e^{(\Delta-\delta)\beta}}\right)^t ,$$

where k is the last index of the sequence $r_1, r_2, \ldots$. Finally, the probability that player 0 plays the strategy 0 at time t, given that she was selected for update at least once, can be handled similarly to the probability that player 2 plays the strategy 0 at time s_2. This concludes the proof of Lemma 6.3.17.

The following lemma gives the probability that the hitting time of the profile $\mathbf{0}$ is less or equal to t, given that $\rho_{n-1} \leqslant t$.

Lemma 6.3.20. *Let $\mathcal{G}$ be the n-player graphical coordination game on the ring and consider the logit dynamics for $\mathcal{G}$. For every $\mathbf{x} \in R$ and every $t > 0$, we have*

$$\mathbf{P_x}\left(\tau_{\mathbf{0}} \leqslant t \mid \rho_{n-1} \leqslant t\right) \geqslant \left(1 - \frac{1}{1 + e^{(\Delta-\delta)\beta}}\right)^{nt}.$$

Proof. To prove our lemma we will show a bound on the probability that, conditioned on $\rho_{n-1} \leqslant t$, all players are playing 0 at time t.

Let f be the permutation that sort players in order of last selection for update: i.e., $f(0)$ is the last player that is selected for update, $f(1)$ is the next to last one, and so on. We have

$$\begin{aligned}
\mathbf{P_x}\left(\tau_{\mathbf{0}} \leqslant t \mid \rho_{n-1} \leqslant t\right) &\geqslant \mathbf{P_x}\left(\bigcap_{j=0}^{n-1} X_t^{f(j)} = 0 \mid \rho_{n-1} \leqslant t\right) \\
&= \prod_{j=0}^{n-1} \mathbf{P_x}\left(X_t^{f(j)} = 0 \mid \bigcap_{i=j+1}^{n-1} X_t^{f(i)} = 0, \rho_{n-1} \leqslant t\right) \\
&\geqslant \prod_{j=0}^{n-1} \mathbf{P_x}\left(X_t^{f(j)} = 0 \mid \rho_{n-1} \leqslant t\right) \\
&\geqslant \left(1 - \frac{1}{1 + e^{(\Delta-\delta)\beta}}\right)^{nt},
\end{aligned}$$

where the last inequality follows from Lemma 6.3.17. □

Now we are ready to give the actual proof of Lemma 6.3.4.

Proof of Lemma 6.3.4. From Lemma 6.3.16 and Lemma 6.3.20, we have that for every $\mathbf{x} \in R$ and every $t > 0$

$$\begin{aligned}
\mathbf{P_x}\left(\tau_{\mathbf{0}} \leqslant t\right) &\geqslant \mathbf{P_x}\left(\rho_{n-1} \leqslant t\right) \mathbf{P_x}\left(\tau_{\mathbf{0}} \leqslant t \mid \rho_{n-1} \leqslant t\right) \\
&\geqslant \left(1 - \frac{n^2}{t}\right)\left(1 - \frac{1}{1 + e^{(\Delta-\delta)\beta}}\right)^{nt}.
\end{aligned}$$

Thus for every $\lambda > 0$, we have for $t = \frac{(8-\lambda)n^2}{\lambda}$

$$\begin{aligned}
\mathbf{P_x}\left(\tau_{\mathbf{0}} \leqslant t\right) &\geqslant \left(1 - \frac{\lambda}{8-\lambda}\right)\left(1 - \frac{1}{1 + \frac{(8-\lambda)n^3}{\lambda \log\left(\frac{8}{8-\lambda}\right)}}\right)^{\frac{(8-\lambda)n^3}{\lambda}} \\
&\geqslant \frac{8\left(1 - \frac{\lambda}{4}\right)}{8-\lambda}\frac{8-\lambda}{8} = 1 - \frac{\lambda}{4},
\end{aligned}$$

where the first inequality follows from the fact that, since $\beta = \omega(\log n)$, for n large enough, we have $\beta \geqslant \frac{\log\left(\frac{(8-\lambda)n^3}{\lambda}\right) - \log\log\left(\frac{8}{8-\lambda}\right)}{\Delta-\delta}$, whereas the second inequality follows from the well known approximation $1 - a \geqslant e^{-\frac{a}{1-a}}$. □

Proof of Lemma 6.3.5

Let $\theta^\star$ be the first time at which all players have been selected at least once. The following lemma directly follows from coupon collector argument; we include a proof for completeness.

Lemma 6.3.21. *Let $\mathcal{G}$ be the n-player graphical coordination game on the ring and consider the logit dynamics for $\mathcal{G}$. For every $t > 0$,*

$$\mathbf{P}_{\mathbf{x}}\left(\theta^\star \leqslant t\right) \geqslant 1 - ne^{-t/n}\,.$$

Proof. The logit dynamics at each time step selects a player for update uniformly and independently of the previous selections. Thus the probability that i players are never selected for update in t steps is $\left(1-\frac{i}{n}\right)^t$ and

$$\mathbf{P}_{\mathbf{x}}\left(\theta^\star > t\right) \leqslant \sum_{i=1}^{n-1}\left(1-\frac{i}{n}\right)^t \leqslant \sum_{i=1}^{n-1} e^{-\frac{it}{n}} \leqslant ne^{-t/n}\,.$$

□

We define players playing 0 in profile $\mathbf{x}$ as the *zero-players* of $\mathbf{x}$ and their neighbors as *border-players*; we also define $l(\mathbf{x}) \geqslant d+1$ as the number of border-players in $\mathbf{x}$.

Let $\tau^\star$ be the first time step at which a border-player is selected for update before one of its neighboring zero-players; if this event does not occur then $\tau^\star = +\infty$. The next lemma bounds the probability that $\tau^\star$ is finite given that all players have been selected at least once within time t.

Lemma 6.3.22. *Let $\mathcal{G}$ be the n-player graphical coordination game on the ring and consider the logit dynamics for $\mathcal{G}$. For every $\mathbf{x} \in S_d \setminus R$ and every $t > 0$*

$$\mathbf{P}_{\mathbf{x}}\left(\tau^\star \leqslant t \mid \theta^\star \leqslant t\right) \geqslant \frac{2d}{2d+1}\,.$$

Proof. Observe that if $\theta^\star \leqslant t$, then $\tau^\star > t$ is equivalent to say that $\tau^\star$ is infinite: thus we will consider $\mathbf{P}_{\mathbf{x}}\left(\tau^\star \text{ is finite} \mid \theta^\star \leqslant t\right)$.

The proof proceeds by induction on d. Let $d = 1$ and denote by i the one zero-player. Notice that $\tau^\star$ is finite if and only if one of the two neighbors of i is selected for update before i is selected. Since we are conditioning on $\theta^\star \leqslant t$, all players are selected at least once by time t and thus the probability of this event is $\frac{2}{3} = \frac{2d}{2d+1}$.

Suppose now that the claim holds for $d-1$ and let $\mathbf{x} \in S_d \setminus R$. Denote by $T_{\mathbf{x}}$ the set of all the zero-players in $\mathbf{x}$ and their border-players and let i be the first player in $T_{\mathbf{x}}$ to be selected for update (notice that i is well defined since $\theta^\star \leqslant t$). Observe that, if i is a border-player, then $\tau^\star$ is finite and this happens with probability $\frac{l(\mathbf{x})}{l(\mathbf{x})+d}$. If i is a zero-player, we consider the subset $\overline{T}_{\mathbf{x}} \subset T_{\mathbf{x}}$ of the remaining $d-1$ zero-players and their border-players. $\tau^\star$ is finite only if at least one border-player in $\overline{T}_{\mathbf{x}}$ is selected before one of its neighboring zero-players. Notice though that $\overline{T}_{\mathbf{x}} = T_{\mathbf{y}}$, for $\mathbf{y} \in S_{d-1} \setminus R$ such that $y_i = 1$ and $\mathbf{y}_{-i} = \mathbf{x}_{-i}$. Thus, by inductive hypothesis, $\mathbf{P}_{\mathbf{y}}\left(\tau^\star \text{ is finite} \mid \theta^\star \leqslant t\right) \geqslant \frac{2d-2}{2j-1}$. Finally,

$$\begin{aligned}
\mathbf{P}_{\mathbf{x}}\left(\tau^\star \text{ is finite} \mid \theta^\star \leqslant t\right) &= \frac{l(\mathbf{x})}{l(\mathbf{x})+d} + \frac{d}{l(\mathbf{x})+d}\cdot \mathbf{P}_{\mathbf{y}}\left(\tau^\star \text{ is finite} \mid \theta^\star \leqslant t\right)\\
&\geqslant \frac{l(\mathbf{x})}{l(\mathbf{x})+d} + \frac{d}{l(\mathbf{x})+d}\cdot\frac{2d-2}{2d-1}\\
&= 1 - \frac{d}{(l(\mathbf{x})+d)(2d-1)}\\
&\geqslant 1 - \frac{1}{2d+1}\,,
\end{aligned}$$

where the last inequality follows from $l(\mathbf{x}) \geqslant d+1$. □

Now we are ready to give the actual proof of Lemma 6.3.5.

Proof of Lemma 6.3.5. Suppose $\tau^\star$ is finite and let i be the border-player selected for update at time $\tau^\star$. Then, at time $\tau^\star$, i has at least one neighbor playing 0 and thus i plays 0 with probability

$$\mathbf{P}_{\mathbf{x}}\left(X^i_{\tau^\star} = 0 \mid \tau^\star \leqslant t\right) \geqslant \left(1 - \frac{1}{1+e^{(\Delta-\delta)\beta}}\right).$$

Moreover, if i plays strategy 0, then at time $\tau^\star$ the dynamics hits a profile in R. Thus, for a finite $t > 0$, we have

$$\begin{aligned}
\mathbf{P}_{\mathbf{x}}\left(\tau_R \leqslant t\right) &\geqslant \mathbf{P}_{\mathbf{x}}\left(X^i_{\tau^\star} = 0 \wedge \tau^\star \leqslant t\right) \\
&= \mathbf{P}_{\mathbf{x}}\left(X^i_{\tau^\star} = 0 \mid \tau^\star \leqslant t\right) \mathbf{P}_{\mathbf{x}}\left(\tau^\star \leqslant t\right) \\
&\geqslant \mathbf{P}_{\mathbf{x}}\left(X^i_{\tau^\star} = 0 \mid \tau^\star \leqslant t\right) \mathbf{P}_{\mathbf{x}}\left(\tau^\star \leqslant t \mid \theta^\star \leqslant t\right) \mathbf{P}_{\mathbf{x}}\left(\theta^\star \leqslant t\right) \\
&\geqslant \left(1 - \frac{1}{1+e^{(\Delta-\delta)\beta}}\right) \frac{2d}{2d+1} \left(1 - ne^{-t/n}\right),
\end{aligned}$$

where the last inequality follows from Lemma 6.3.21 and Lemma 6.3.22. Finally, the lemma follows since $\beta = \omega(\log n)$ and by taking $t = n^2$. □

Proof of Lemma 6.3.8

This proof is very similar to the proof of Lemma 6.3.5: in particular, we refer to notation defined in that proof. If all zero-players are selected before both neighboring border-players, we set $\tau^\star$ be the time step at which the last zero-players is selected, otherwise we set $\tau^\star$ be infinity.

Lemma 6.3.23. *Let $\mathcal{G}$ be the n-player graphical coordination game on the ring and consider the logit dynamics for $\mathcal{G}$. For every $\mathbf{x} \in S_d \setminus R$*

$$\mathbf{P}_{\mathbf{x}}\left(\tau^\star \leqslant t \mid \theta^\star \leqslant t\right) \geqslant \frac{1}{3^d}.$$

Proof. Observe that if $\theta^\star \leqslant t$, then $\tau^\star > t$ is equivalent to say that $\tau^\star$ is infinite: thus we will consider $\mathbf{P}_{\mathbf{x}}\left(\tau^\star \text{ is finite} \mid \theta^\star \leqslant t\right)$.

The proof proceeds by induction on d. For the base case, we denote by i the only zero-player and $\tau^\star$ is finite if and only if i is selected for update before her neighbors. Since we are conditioning on $\theta^\star \leqslant t$, we know that all players have been selected at least once and thus the probability of this event is $\frac{1}{3}$.

Suppose that the claim holds for $d-1$ and let $\mathbf{x} \in S_d \setminus R$. Denote by $T_{\mathbf{x}}$ the set of all the zero-players in $\mathbf{x}$ along with their border-players and let i be the first player in $T_{\mathbf{x}}$ to be selected for update (notice that i is well defined since $\theta^\star \leqslant t$ and thus all players has been selected at least once). Observe that, if i is a border-player, then $\tau^\star$ is *infinite* and this happens with probability $\frac{l(\mathbf{x})}{l(\mathbf{x})+j}$. Otherwise, we consider the subset $\overline{T}_{\mathbf{x}} \subset T_{\mathbf{x}}$ of the remaining $d-1$ zero-players and their border-players. $\tau^\star$ will be finite only if all zero-players in $\overline{T}_{\mathbf{x}}$ are selected before their border-players. However, $\overline{T}_{\mathbf{x}} = T_{\mathbf{y}}$, where $\mathbf{y} \in S_{d-1} \setminus R$ is the profile obtained from $\mathbf{x}$ by setting $y_i = 1$. By inductive hypothesis the probability $\mathbf{P}_{\mathbf{y}}\left(\tau^\star \text{ is finite} \mid \theta^\star \leqslant t\right) \geqslant \frac{1}{3^{d-1}}$. Thus, we have

$$\begin{aligned}
\mathbf{P}_{\mathbf{x}}\left(\tau^\star \text{ is finite} \mid \theta^\star \leqslant t\right) &= \frac{d}{l(\mathbf{x})+d} \cdot \mathbf{P}_{\mathbf{y}}\left(\tau^\star \text{ is finite} \mid \theta^\star \leqslant t\right) \\
&\geqslant \frac{1}{3} \cdot \frac{1}{3^{d-1}}.
\end{aligned}$$

□

Now we give the actual proof of Lemma 6.3.8.

Proof of Lemma 6.3.8. Notice that if $\tau^\star$ is finite, every time a player is selected for update she have both neighbors that are playing 0 and thus

$$\mathbf{P}_{\mathbf{x}}\left(X_{\tau^\star} = \mathbf{1} \mid \tau^\star \leqslant t\right) \geqslant \left(1 - \frac{1}{1+e^{2\delta\beta}}\right)^t \geqslant \left(1 - \frac{t}{e^{2\delta\beta}}\right),$$

where the last inequality follows from the approximations $1 - a \leqslant e^{-x}$ and $1 - a \geqslant e^{-\frac{x}{1-x}}$.

Obviously, if $X_{\tau^\star} = \mathbf{1}$, then $\tau_{\mathbf{1}} \leqslant \tau^\star$. Thus, for a finite $t > 0$, we have

$$\begin{aligned}
\mathbf{P}_{\mathbf{x}}\left(\tau_{\mathbf{1}} \leqslant t\right) &\geqslant \mathbf{P}_{\mathbf{x}}\left(X_{\tau^\star} = \mathbf{1} \wedge \tau^\star \leqslant t\right) \\
&= \mathbf{P}_{\mathbf{x}}\left(X_{\tau^\star} = \mathbf{1} \mid \tau^\star \leqslant t\right) \mathbf{P}_{\mathbf{x}}\left(\tau^\star \leqslant t\right) \\
&\geqslant \mathbf{P}_{\mathbf{x}}\left(X_{\tau^\star} = \mathbf{1} \mid \tau^\star \leqslant t\right) \mathbf{P}_{\mathbf{x}}\left(\tau^\star \leqslant t \mid \theta^\star \leqslant t\right) \mathbf{P}_{\mathbf{x}}\left(\theta^\star \leqslant t\right) \\
&\geqslant \left(1 - \frac{t}{e^{2\delta\beta}}\right) \frac{1}{3^z} \left(1 - ne^{-t/n}\right),
\end{aligned}$$

where the last inequality follows from Lemma 6.3.21 and Lemma 6.3.23. Finally, the lemma follows since $\beta = \omega(\log n)$ and by taking $t = n^2$. □

6.3.4 Proofs from Section 6.3.2

We say that a profile $\mathbf{x}$ has a *zero-block* of size l starting at player i if $x_i = x_{i+1} = \ldots = x_{i+l-1} = 0$ and $x_{i-1} = x_{i+l} = 1$. Players i and $i + h - 1$ are the *border* players of the block. A similar definition is given for *one-blocks*. Notice that every profile $\mathbf{x} \neq \mathbf{0}, \mathbf{1}$ has the same number of zero-blocks and one-blocks and this number is called the *level* of $\mathbf{x}$ and is denoted by $\ell(\mathbf{x})$. We set $\ell(\mathbf{0}) = \ell(\mathbf{1}) = 0$.

The following observation gives the *level structure* of the potential function (note that we are studying the case $\Delta = \delta$).

Observation 6.3.24. *Let $\mathcal{G}$ be the n-player graphical coordination game with profile space S. For every profile $\mathbf{x} \in S$, the potential of $\mathbf{x}$ is $\Phi(\mathbf{x}) = (n - 2\ell(\mathbf{x}))\Delta$, regardless of the sizes of the zero-blocks and one-blocks.*

Moreover, for a profile $\mathbf{x}$, we defines $s_0(\mathbf{x})$ as the number of zero-blocks of size 1, $s_1(\mathbf{x})$ as the number of one-blocks of size 1 and set $s(\mathbf{x}) = s_0(\mathbf{x}) + s_1(\mathbf{x})$.

Proof of Lemma 6.3.12

We would like to study how long it takes for the logit dynamics to reach $\mathbf{0}$ or $\mathbf{1}$. Starting from profile $\mathbf{x}$ at level $i \geqslant 1$, the logit dynamics needs to go down i levels to hit a profile at level 0; and to go down one level, it is necessary for one monochromatic block (that is, either a zero-block or a one-block) to disappear. We next show that we do not have to wait too long for this to happen.

Our first step bounds the time τ_i needed to go from level $i + 1$ to level i. Consider a profile $\mathbf{x}$ at level $i + 1$ and number arbitrarily the $2(i + 1)$ monochromatic blocks of $\mathbf{x}$ and denote by $k_j(\mathbf{x})$ the size of the j-th monochromatic block. We define $\tau_{i,j}$ in the following way. Suppose that the dynamics reaches level i for the first time after t steps and suppose that this happens because the j-th monochromatic block disappears. Then we set $\tau_{i,j} = t$ and $\tau_{i,j'} = +\infty$ for all $j' \neq j$. Obviously, for any starting profile $\mathbf{x}$ we have that

$$\mathbf{E_x}[\tau_i] = \mathbf{E_x}\left[\min_j \tau_{i,j}\right] \leqslant \max_j \mathbf{E_x}\left[\tau_{i,j} \mid \tau_{i,j} < \tau_{i,j'} \text{ for all } j' \neq j\right].$$

For sake of compactness of notation we define

$$\gamma_{i,l} = \max_{1 \leqslant j \leqslant 2(i+1)} \max_{\substack{\mathbf{x}\colon \ell(\mathbf{x})=i+1 \\ k_j(\mathbf{x})=l}} \mathbf{E_x}\left[\tau_{i,j} \mid \tau_{i,j} < \tau_{i,j'} \text{ for all } j' \neq j\right],$$

set $\gamma_i = \max_l \gamma_{i,l}$ and observe that $\mathbf{E_x}[\tau_i] \leqslant \gamma_i$. It is also easy to see that $\gamma_{i,l}$ is non-decreasing with l. Next we bound γ_i in terms of γ_{i+1}.

Lemma 6.3.25. *Let $\mathcal{G}$ be the n-player graphical coordination game on the ring and consider the logit dynamics for $\mathcal{G}$. For every $i \geqslant 0$*

$$\gamma_i \leqslant n^2 b_i$$

where $b_i = n + \frac{n-4(i+1)+s(\mathbf{x}^\star)}{1+e^{2\Delta\beta}}\gamma_{i+1}$.

Proof. We next bound $\gamma_{i,l}$ by distinguishing cases depending on the size l. For each case, we let $\mathbf{x}$ and j be the profile and the monochromatic block that attains the maximum $\gamma_{i,l}$.

<u>$l = 1$:</u>

- if the unique player of the j-th monochromatic block is selected for update and she changes her strategy then $\tau_{i,j} = 1$. This happens with probability $\frac{1}{n} \cdot \left(1 - \frac{1}{1+e^{2\Delta\beta}}\right)$.
- if a neighbor v of the unique player of the j-th monochromatic block is selected for update and she changes her strategy then the dynamics reaches a profile $\mathbf{y}$ at same level $i+1$ and the size of the j-th block increases to 2. If v belongs to a monochromatic block of size 1, this has probability 0 (we are conditioning on $\tau_{i,j} < \tau_{i,j'}$ for all $j' \neq j$); otherwise, the probability is at most $1/2 \cdot 2/n = 1/n$.
- if we select for update a player that is not at the borders of a monochromatic block and she changes her strategy, then the dynamics reaches a profile $\mathbf{y}$ at level $i+2$. This has probability $\frac{n-4(i+1)+s(\mathbf{x})}{1+e^{2\Delta\beta}}$ of occurring.
- in the remaining cases neither the level nor the length of the j-th monochromatic block changes.

Hence, by observing that $\gamma_{i,2} \geqslant \gamma_{i,1}$ we have

$$\begin{aligned}
\gamma_{i,1} \quad &\leqslant \quad \frac{1}{n}\left(1 - \frac{1}{1+e^{2\Delta\beta}}\right) + \frac{1}{n}(1+\gamma_{i,2}) + \frac{n-4(i+1)+s(\mathbf{x})}{n}\frac{1}{1+e^{2\Delta\beta}}(1+\gamma_{i+1}) \\
&+ \quad \left(\frac{n-2}{n} - \frac{n-4(i+1)+s(\mathbf{x})-1}{n}\frac{1}{1+e^{2\Delta\beta}}\right)(1+\gamma_{i,1}).
\end{aligned}$$

By simple calculations and using that $n - 4(i+1) + s(\mathbf{x}) \geqslant 0$, we obtain

$$\gamma_{i,1} \leqslant \left(\frac{1}{2} + \frac{1}{4e^{2\Delta\beta}+2}\right)\left(n + \gamma_{i,2} + \frac{n-4(i+1)+s(\mathbf{x})}{1+e^{2\Delta\beta}}\gamma_{i+1}\right).$$

Since $\left(\frac{1}{2} + \frac{1}{4e^{2\Delta\beta}+2}\right) \leqslant \frac{2}{3}$ for $\beta = \omega(\log n)$, we have

$$\gamma_{i,1} \leqslant \frac{2}{3}(\gamma_{i,2} + b_i). \tag{6.9}$$

<u>$1 < l < n - 2i - 1$:</u>

- if a player at the borders of the j-th monochromatic block is selected for update (there are two of these players) and she changes her strategy (this happens with probability $1/2$), then the dynamics reaches a profile $\mathbf{y}$ at same level $i+1$ and the length of the j-th monochromatic block decreases to $l-1$;
- if a neighbor v of the border players of the j-th monochromatic block is selected for update and she changes her strategy, then the number of monochromatic blocks does not change (and thus we are at still at level $i+1$) but the j-th monochromatic block increases in size.

  Notice that, in this case, player v does not belong to a monochromatic block of size 1, since we are conditioning on the fact that the j-th monochromatic block is the first to disappear ($\tau_{i,j} < \tau_{i,j'}\ \forall j' \neq j$). Therefore the two neighbors of v are playing two different strategies and thus v adopts any of the two with probability $1/2$. Since there are two players adjacent to the border players of block j, this case happens with probability at most $1/n$.
- if a player v that is not at the borders of a monochromatic block is selected for update and she changes her strategy then the two new adjacent monochromatic blocks are created and the level increases 1. Notice that there are $n-4(i+1)+s(\mathbf{x}^\star)$ such player v and each has probability $\frac{1}{1+e^{2\Delta\beta}}$ of changing her strategy.
- in the remaining cases neither the level nor the length of the j-th monochromatic block changes.

Hence,

$$\begin{aligned}
\gamma_{i,l} &\leqslant \frac{1}{n}(1+\gamma_{i,l-1}) + \frac{1}{n}(1+\gamma_{i,l+1}) + \frac{n-4(i+1)+s(\mathbf{x})}{n}\frac{1}{1+e^{2\Delta\beta}}(1+\gamma_{i+1}) \\
&+ \left(\frac{n-2}{n} - \frac{n-4(i+1)+s(\mathbf{x})}{n}\frac{1}{1+e^{2\Delta\beta}}\right)(1+\gamma_{i,l})\,.
\end{aligned}$$

By simple calculations, similar to the ones for the case $l=1$, we obtain

$$\gamma_{i,l} \leqslant \frac{1}{2}(\gamma_{i,l-1} + \gamma_{i,l+1} + b_i)\,.$$

From the previous inequality and Equation 6.9, a simple induction on l shows that, for every $1 \leqslant l < n-2i-1$, we have

$$\gamma_{i,l} \leqslant \frac{1}{l+2}\left((l+1)\gamma_{i,l+1} + \frac{l(l+3)}{2}b_i\right)\,. \tag{6.10}$$

Moreover, from Equation 6.10, a simple inductive argument shows that, for every $h \geqslant 1$,

$$\begin{aligned}
\gamma_{i,l} &\leqslant \frac{l+1}{l+h+1}\gamma_{i,l+h} + \frac{l+1}{2}b_i\sum_{j=l}^{l+h-1}\frac{j(j+3)}{(j+1)(j+2)} \\
&\leqslant \frac{l+1}{l+h+1}\gamma_{i,l+h} + \frac{l+1}{2}hb_i\,. \qquad (6.11)
\end{aligned}$$

$\underline{l = n-2i-1:}$ in this case all blocks other than the j-th have size 1 and thus every time we select one of these players, she changes her strategy with probability 0 (we are conditioning on j being the first monochromatic block to disappear). This means that that the size of the j-th monochromatic block cannot increase. Reasoning similar to the ones used in the previous cases, we obtain that

$$\gamma_{i,n-2i-1} \leqslant \gamma_{i,n-2i-2} + b_i\,.$$

By using Equation 6.10, we have

$$\gamma_{i,n-2i-1} \leqslant \frac{(n-2i-2)(n-2i+1)+2(n-2i)}{2} b_i \leqslant \frac{n^2}{2} b_i \,.$$

Finally, for every $l \geqslant 1$, by using Equation 6.11 with $h = n-2i-1-l$, we have

$$\gamma_{i,l} \leqslant \frac{l+1}{n-2i}\gamma_{i,n-2i-1} + \frac{(l+1)(n-2i-1-l)}{2} b_i \leqslant n^2 b_i \,.$$

□

Corollary 6.3.26. *Let $\mathcal{G}$ be the n-player graphical coordination game on the ring and consider the logit dynamics for $\mathcal{G}$. If $\beta = \omega(\log n)$, then for every $i \geqslant 0$, $\gamma_i = \mathcal{O}(n^3)$.*

Proof. Note that $b_{\lfloor n/2 \rfloor - 1} = n$ and thus, by using Lemma 6.3.25, $\gamma_{\lfloor n/2 \rfloor - 1} \leqslant n^3$. Moreover, for $0 \leqslant i < \lfloor n/2 \rfloor - 1$, since $n - 4(i+1) + s(\mathbf{x}) \leqslant n$, we have

$$\gamma_i \leqslant n^3 \left(1 + \frac{1}{1+e^{2\Delta\beta}}\gamma_{i+1}\right) \leqslant n^3 \left(1 + \sum_{j=1}^{\lfloor n/2 \rfloor - i - 1} \left(\frac{n^3}{1+e^{2\Delta\beta}}\right)^j\right) . \tag{6.12}$$

The corollary follows by observing that, if $\beta = \omega\left(\frac{\log n}{\Delta}\right)$, then the summation in Equation 6.12 is $o(1)$. □

The above corollary gives a polynomial bound to the time that the dynamics take to go from a profile at level $i+1$ to a profile at level i. Lemma 6.3.12 easily follows.

Proof of Lemma 6.3.12. Obviously, for every $\mathbf{x}$ at level $1 \leqslant k \leqslant n/2$,

$$\mathbf{E}_{\mathbf{x}}\left[\tau_{\mathbf{0},\mathbf{1}}\right] \leqslant \sum_{i=0}^{k-1} \max_{\mathbf{y}\colon \ell(\mathbf{y})=i+1} \mathbf{E}_{\mathbf{y}}\left[\tau_i\right] \leqslant \sum_{i=0}^{k-1} \gamma_i = \mathcal{O}(n^4) \,.$$

The lemma follows from the Markov inequality. □

Proof of Lemma 6.3.13

Proof of Lemma 6.3.13. For a profile $\mathbf{x}$, we denote by $p_{\mathbf{x}}$ the probability that the logit dynamics starting from $\mathbf{x}$ at step $\tau_{\mathbf{0},\mathbf{1}}$ is in profile $\mathbf{0}$; in other words, $p_{\mathbf{x}} = \mathbf{P}_{\mathbf{x}}\left(\tau_{\mathbf{0}} < \tau_{\mathbf{1}}\right)$. Trivially, $p_{\mathbf{0}} = 1$ and $p_{\mathbf{1}} = 0$.

Clearly, at time $\tau_{\mathbf{0},\mathbf{1}}$ the dynamics is either in the state $\mathbf{0}$ (this happens with probability $p_{\mathbf{x}}$) or in the state $\mathbf{1}$ (this happens with probability 1 - $p_{\mathbf{x}}$). Thus, the state of the dynamics at time $\tau_{\mathbf{0},\mathbf{1}}$ when starting from profile x is distributed according to the probability distribution

$$\nu_{\mathbf{x}} = p_{\mathbf{x}}\pi_{\mathbf{0}} + (1-p_{\mathbf{x}})\pi_{\mathbf{1}}.$$

We next show that for $\beta = \omega(\log n)$ and $\mathbf{x} \in S_d$, $p_{\mathbf{x}} = \frac{d}{n} + \lambda_{\mathbf{x}}$, for some $\lambda_{\mathbf{x}} = o(1)$. By the definition of Markov chains we know that

$$p_{\mathbf{x}} = P(\mathbf{x},\mathbf{x}) \cdot p_{\mathbf{x}} + \sum_{y \in N(\mathbf{x})} P(\mathbf{x},\mathbf{y}) \cdot p_{\mathbf{y}} \,.$$

We then partition the neighborhood $N(\mathbf{x})$ of profile $\mathbf{x}$ of level i in 5 subsets, $N_1(\mathbf{x})$, $N_2(\mathbf{x})$, $N_3(\mathbf{x})$, $N_4(\mathbf{x})$, $N_5(\mathbf{x})$, such that, for two profiles $\mathbf{y}_1, \mathbf{y}_2$ in the same subsets it holds that $P(\mathbf{x},\mathbf{y}_1) = P(\mathbf{x},\mathbf{y}_2)$.

- $N_1(\mathbf{x})$ is the set of profiles $\mathbf{y}$ obtained from $\mathbf{x}$ by changing the strategy of a player of a zero-block of size 1. Observe that $|N_1(\mathbf{x})| = s_0(\mathbf{x})$. Moreover, for every $\mathbf{y} \in N_1(\mathbf{x})$, $\mathbf{y}$ is at level $i-1$, has $|\mathbf{x}|_0 - 1$ players playing 0 and $P(\mathbf{x},\mathbf{y}) = \frac{1}{n} \cdot (1 - \frac{1}{1+e^{2\Delta\beta}})$.
- $N_2(\mathbf{x})$ is the set of profiles $\mathbf{y}$ obtained from $\mathbf{x}$ by changing the strategy of a player of a one-block of size 1. Observe that $|N_2(\mathbf{x})| = s_1(\mathbf{x})$. Moreover, for every $\mathbf{y} \in N_2(\mathbf{x})$, $\mathbf{y}$ is at level $i-1$, has $|\mathbf{x}|_0 + 1$ players playing 0 and $P(\mathbf{x},\mathbf{y}) = \frac{1}{n} \cdot (1 - \frac{1}{1+e^{2\Delta\beta}})$.
- $N_3(\mathbf{x})$ is the set of profiles $\mathbf{y}$ obtained from $\mathbf{x}$ by changing the strategy of a border player of a zero-block of size greater than 1. Observe that $|N_3(\mathbf{x})| = 2(i - s_0(\mathbf{x}))$. Moreover, for every $\mathbf{y} \in N_3(\mathbf{x})$, $\mathbf{y}$ is at level i, and has $|\mathbf{x}|_0 - 1$ players playing 0 and $P(\mathbf{x},\mathbf{y}) = 1/2n$.
- $N_4(\mathbf{x})$ is the set of profiles $\mathbf{y}$ obtained from $\mathbf{x}$ by changing the strategy of a border player of a one-block of size greater than 1. Observe that $|N_4(\mathbf{x})| = 2(i - s_1(\mathbf{x}))$. Moreover, for every $\mathbf{y} \in N_4(\mathbf{x})$, $\mathbf{y}$ is at level i, and has $|\mathbf{x}|_0 + 1$ players playing 0 and $P(\mathbf{x},\mathbf{y}) = 1/2n$.
- $N_5(\mathbf{x})$ is the set of all the profiles $\mathbf{y} \in N(\mathbf{x})$ that do not belong to any of the previous 4 subsets. Observe that $|N_5(\mathbf{x})| = n - 4i + s(\mathbf{x})$. Moreover, for every $\mathbf{y} \in N_5(\mathbf{x})$, $\mathbf{y}$ is at level $i+1$, and $P(\mathbf{x},\mathbf{y}) = \frac{1}{n} \cdot \frac{1}{1+e^{2\Delta\beta}}$.

Moreover, we have that

$$P(\mathbf{x},\mathbf{x}) = \frac{s(\mathbf{x})}{n}\frac{1}{1+e^{2\Delta\beta}} + \frac{2i - s(\mathbf{x})}{n} + \frac{n-4i+s(\mathbf{x})}{n}\left(1 - \frac{1}{1+e^{2\Delta\beta}}\right).$$

Then, we have

$$\begin{aligned}
p_{\mathbf{x}} \quad &= \quad \frac{1}{n}\left(1 - \frac{1}{1+e^{2\Delta\beta}}\right)\left(\sum_{\mathbf{y}\in N_1(\mathbf{x})} p_{\mathbf{y}} + \sum_{\mathbf{y}\in N_2(\mathbf{x})} p_{\mathbf{y}}\right) + \frac{1}{2n}\left(\sum_{\mathbf{y}\in N_3(\mathbf{x})} p_{\mathbf{y}} + \sum_{\mathbf{y}\in N_4(\mathbf{x})} p_{\mathbf{y}}\right) \\
&\quad + \frac{1}{n}\frac{1}{1+e^{2\Delta\beta}} \sum_{\mathbf{y}\in N_5(\mathbf{x})} p_{\mathbf{y}} \\
&\quad + \left(\frac{s(\mathbf{x})}{n}\frac{1}{1+e^{2\Delta\beta}} + \frac{2i - s(\mathbf{x})}{n} + \frac{n-4i+s(\mathbf{x})}{n}\left(1 - \frac{1}{1+e^{2\Delta\beta}}\right)\right) p_{\mathbf{x}} \\
&= \quad \frac{1}{n}\left(\sum_{\mathbf{y}\in N_1(\mathbf{x})} p_{\mathbf{y}} + \sum_{\mathbf{y}\in N_2(\mathbf{x})} p_{\mathbf{y}}\right) + \frac{1}{2n}\left(\sum_{\mathbf{y}\in N_3(\mathbf{x})} p_{\mathbf{y}} + \sum_{\mathbf{y}\in N_4(\mathbf{x})} p_{\mathbf{y}}\right) + \frac{n-2i}{n} \cdot p_{\mathbf{x}} + \frac{c}{1+e^{2\Delta\beta}}, \qquad (6.13)
\end{aligned}$$

where

$$c = \frac{1}{n}\left(\sum_{\mathbf{y}\in N_5(\mathbf{x})} p_{\mathbf{y}} - \sum_{\mathbf{y}\in N_1(\mathbf{x})\cup N_2(\mathbf{x})} p_{\mathbf{y}} - (n-4i)p_{\mathbf{x}}\right).$$

We notice that, since $1 \leqslant i \leqslant n/2$ and $|N_1(\mathbf{x})| + |N_2(\mathbf{x})|, |N_5(\mathbf{x})| \leqslant n$, we have $|c| \leqslant 2$ and thus the last term in Equation 6.13 is negligible in n (since $\beta = \omega(\log n)$). Hence we have that the following condition holds for every level $i \geqslant 1$ and every profile $\mathbf{x}$ at level i:

$$p_{\mathbf{x}} = \frac{1}{2i}\left(\sum_{\mathbf{y}\in N_1(\mathbf{x})} p_{\mathbf{y}} + \sum_{\mathbf{y}\in N_2(\mathbf{x})} p_{\mathbf{y}}\right) + \frac{1}{4i}\left(\sum_{\mathbf{y}\in N_3(\mathbf{x})} p_{\mathbf{y}} + \sum_{\mathbf{y}\in N_4(\mathbf{x})} p_{\mathbf{y}}\right) + \eta_{\mathbf{x}},$$

where $\eta_{\mathbf{x}}$ is negligible in n. This gives us a linear system of equations in which the number of equations is the same as the number of variables.

Next we find a solution to a "modified" version of above system, where we omit the negligible part in every equation, and then we show that this solution cannot be very different from the solution of the "original" system.

We build the solution for the "modified" system inductively on the level i: for every profile $\mathbf{x} \in S_d$ at level 0 (this is only possible for $d = 0$ or $d = n$), we have, as discussed above, $p_{\mathbf{x}} = \frac{d}{n}$. Now, we assume that for every profile $\mathbf{x} \in S_d$ at level $i-1$, $p_{\mathbf{x}} = \frac{d}{n}$ is a solution for the system. For $\mathbf{x} \in S_d$ at level i, we can rewrite the "modified" condition as follows:

$$p_{\mathbf{x}} = \frac{s_0(\mathbf{x})}{2i} \cdot \frac{d-1}{n} + \frac{s_1(\mathbf{x})}{2i} \cdot \frac{d+1}{n} + \frac{1}{4i} \left(\sum_{\mathbf{y} \in N_3(\mathbf{x})} p_{\mathbf{y}} + \sum_{\mathbf{y} \in N_4(\mathbf{x})} p_{\mathbf{y}} \right). \tag{6.14}$$

Equation 6.14 gives another linear system of equations. This system has a unique solution: indeed, it has the same number of equations and variables and the matrix of coefficients is a diagonally dominant matrix (since $|N_3(\mathbf{x}) \cup N_4(\mathbf{x})| \leqslant 4i$) and thus it is nonsingular. Moreover, if we set, for every profile $\mathbf{x}$ at level i, $p_{\mathbf{x}} = \frac{|\mathbf{x}|_0}{n}$, then the right hand side of the Equation 6.14 becomes

$$\frac{s_0(\mathbf{x})}{2i}\frac{d-1}{n} + \frac{s_1(\mathbf{x})}{2i}\frac{d+1}{n} + \frac{i - s_0(\mathbf{x})}{2i}\frac{d-1}{n} + \frac{i - s_1(\mathbf{x})}{2i}\frac{d+1}{n} = \frac{d}{n},$$

and hence the system is satisfied by this assignment. Summarizing, we have found that the "modified" system has a unique solution $p_{\mathbf{x}} = \frac{|\mathbf{x}|_0}{n}$ for every profile $\mathbf{x}$. Now, let $p^\star_{\mathbf{x}} = p_{\mathbf{x}} + \lambda_{\mathbf{x}}$ be the assignment that satisfies all "original" conditions: since, as n grows unbounded, these conditions approach the "modified" ones, we have that $p^\star_{\mathbf{x}}$ has to approach to $p_{\mathbf{x}}$ and thus we have that $|\lambda_{\mathbf{x}}| = o(1)$ for every profile $\mathbf{x}$. □

6.4 The OR game

In this section we consider the *OR game* defined in Section 4.4. There we showed that the mixing time of the logit dynamics for the OR game is roughly e^β for $\beta = \mathcal{O}(\log n)$ and it is roughly 2^n for larger β. Here we study the metastability properties of the OR game to highlight the distinguishing features of our quantitative notion of metastability based on distributions. Namely, we show that if we start the logit dynamics at a profile where at least one player is playing 1, then after $\mathcal{O}(\log n)$ time steps the distribution of the chain is close to uniform, and it stays close to uniform for exponentially long time. Hence, even if there is no small set of the state space where the chain stays close for a long time, we can still say that the chain is "metastable" meaning that the "distribution" of the chain stays close to some well-defined distribution for a long time.

6.4.1 Ehrenfest urns

We first need two simple lemmas that will be used in the proof of Theorem 6.4.4.

The *Ehrenfest urn* is the Markov chain with state space $\Omega = \{0, 1, \ldots, n\}$ that, when at state k, moves to state $k-1$ or $k+1$ with probability k/n and $(n-k)/n$ respectively (see, for example, Section 2.3 in [75] for a detailed description). The next lemma gives an upper bound on the probability that the Ehrenfest urn starting at state k hits state 0 within time step t.

Lemma 6.4.1. *Let $\{Z_t\}$ be the Ehrenfest urn over $\{0, 1, \ldots, n\}$ and let τ_0 be the first time the chain hits state 0. Then for every $k \geqslant 1$ it holds that*

$$\mathbf{P}_k\left(\tau_0 < n \log n + cn\right) \leqslant c'/n$$

for suitable positive constants c and c'.

Proof. First observe that for any $t \geqslant 3$ the probability of hitting 0 before time t for the chain starting at 1 is only $\mathcal{O}(1/n)$ larger than for the chain starting at 2, which in turn is only $\mathcal{O}(1/n)$ larger than for the chain starting at 3. Indeed, by conditioning on the first step of the chain, we have

$$\begin{aligned}
\mathbf{P}_1(\tau_0 < t) &= \mathbf{P}_1(\tau_0 < t \,|\, Z_1 = 0)\,\mathbf{P}_1(Z_1 = 0) \\
&\quad + \mathbf{P}_1(\tau_0 < t \,|\, Z_1 = 2)\,\mathbf{P}_1(Z_1 = 2) \\
&= \frac{1}{n} + \frac{n-1}{n}\mathbf{P}_2(\tau_0 < t-1) \\
&\leqslant \frac{1}{n} + \mathbf{P}_2(\tau_0 < t)\ ,
\end{aligned}$$

and

$$\begin{aligned}
\mathbf{P}_2(\tau_0 < t) &= \mathbf{P}_2(\tau_0 < t \,|\, Z_1 = 1)\,\mathbf{P}_2(Z_1 = 1) \\
&\quad + \mathbf{P}_2(\tau_0 < t \,|\, Z_1 = 3)\,\mathbf{P}_2(Z_1 = 3) \\
&= \frac{2}{n}\mathbf{P}_1(\tau_0 < t-1) + \frac{n-2}{n}\mathbf{P}_3(\tau_0 < t-1) \\
&\leqslant \frac{2}{n}\left(\frac{1}{n} + \mathbf{P}_2(\tau_0 < t)\right) + \frac{n-2}{n}\mathbf{P}_3(\tau_0 < t)\ .
\end{aligned}$$

Hence,

$$\begin{aligned}
\mathbf{P}_2(\tau_0 < t) &\leqslant \frac{2}{n-2} + \mathbf{P}_3(\tau_0 < t) \leqslant \frac{3}{n} + \mathbf{P}_3(\tau_0 < t)\ , \\
\mathbf{P}_1(\tau_0 < t) &\leqslant \frac{4}{n} + \mathbf{P}_3(\tau_0 < t)\ .
\end{aligned}$$

Moreover observe that the probability that the chain starting at k hits state 0 before time t is decreasing in k, in particular, for every $k \geqslant 3$ it holds that $\mathbf{P}_k(\tau_0 < t) \leqslant \mathbf{P}_3(\tau_0 < t)$. Now we show that $\mathbf{P}_3(\tau_0 < n\log n + cn) = \mathcal{O}(1/n)$ and this will complete the proof.

Let us consider a path $\mathcal{P}$ of length t starting at state 3 and ending at state 0. Observe that any such path must contain the sub-path going from state 3 to state 0 whose probability is $6/n^3$. Moreover, for all the other $t-3$ moves we have that if the chain crosses an edge $(i, i+1)$ from left to right then it must cross the same edge from right to left (and vice versa). The probability for any such pair of moves is

$$\frac{n-i}{n} \cdot \frac{i+1}{n} \leqslant \frac{e^{1/n}}{4}\,,$$

for every i. Hence, for any path $\mathcal{P}$ of length t going from 3 to 0, the probability that the chain follows exactly path $\mathcal{P}$ is[3]

$$\begin{aligned}
\mathbf{P}_3((X_1, \ldots, X_t) = \mathcal{P}) &\leqslant \frac{6}{n^3} \cdot \left(\frac{e^{2/n}}{4}\right)^{(t-3)/2} \\
&= \frac{6}{n^3} \cdot \frac{2^3}{e^{3/n}} \cdot \left(\frac{e^{1/n}}{2}\right)^t \\
&\leqslant \frac{48}{n^3} \cdot \left(\frac{e^{1/n}}{2}\right)^t .
\end{aligned}$$

[3]Notice that such probability is zero if $t-3$ is odd.

Let ℓ and r be the number of left and right moves respectively in path $\mathcal{P}$ then $\ell + r = t$ and $\ell - r = 3$. Hence the total number of paths of length t going from 3 to 0 is less than

$$\binom{t}{\ell} = \binom{t}{\frac{t-3}{2}} \leqslant 2^t .$$

Thus, the probability that starting from 3 the chain hits 0 for the first time exactly at time t is

$$\mathbf{P}_3\left(\tau_0 = t\right) \leqslant \binom{t}{\frac{t-3}{2}} \frac{48}{n^3} \cdot \left(\frac{e^{1/n}}{2}\right)^t \leqslant \frac{48}{n^3} e^{t/n} .$$

Finally, the probability that the hitting time of 0 is less than t is

$$\begin{aligned} \mathbf{P}_3\left(\tau_0 < t\right) &\leqslant \sum_{i=3}^{t-1} \mathbf{P}_3\left(\tau_0 = i\right) \\ &\leqslant \frac{48}{n^3} \sum_{i=3}^{t-1} e^{i/n} = \frac{48}{n^3} \cdot \frac{e^{t/n} - 1}{e^{1/n} - 1} \leqslant \frac{48 e^c}{n} . \end{aligned}$$

In the last inequality we used that $e^{1/n} - 1 \geqslant 1/n$ and $t = n \log n + cn$. □

In the proof of Theorem 6.4.4 we will be dealing with the lazy version of the Ehrenfest urn. The next lemma, which is folklore, allows us to use the bound we achieved in Lemma 6.4.1 for the non-lazy chain.

Lemma 6.4.2. *Let $\{X_t\}$ be an irreducible Markov chain with finite state space Ω and transition matrix P and let $\{\hat{X}_t\}$ be its lazy version, i.e. the Markov chain with the same state space and transition matrix $\hat{P} = \frac{P+I}{2}$ where I is the $\Omega \times \Omega$ identity matrix. Let τ_a and $\hat{\tau}_a$ be the hitting time of state $a \in \Omega$ in chains $\{X_t\}$ and $\{\hat{X}_t\}$ respectively. Then, for every starting state $b \in \Omega$ and for every time $t \in \mathbb{N}$ it holds that*

$$\mathbf{P}_b\left(\hat{\tau}_a \leqslant t\right) \leqslant \mathbf{P}_b\left(\tau_a \leqslant t\right) .$$

6.4.2 OR game metastability.

The next lemma shows that, if we start from the uniform distribution, the distribution of the logit dynamics stays ε-close to uniform for $\varepsilon 2^n$ time steps.

Lemma 6.4.3. *Let P be the transition matrix of the logit dynamics for the n-player OR-game, let U be the uniform distribution over $\{0,1\}^n$. Then U is $(\varepsilon, \varepsilon 2^n)$-metastable.*

Proof. Observe that, by starting from the stationary distribution, the probability of being in $\mathbf{y} \in \{0,1\}^n$ after one step of the chain is

$$\begin{aligned} UP(\mathbf{y}) &= \sum_{\mathbf{x} \in \{0,1\}^n} U(\mathbf{x}) P(\mathbf{x}, \mathbf{y}) = \frac{1}{2^n} \sum_{\mathbf{x} \in \{0,1\}^n} P(\mathbf{x}, \mathbf{y}) \\ &= \begin{cases} 2^{-n} & \text{if } |\mathbf{y}|_1 \geqslant 2 \\ 2^{-n}\left(\frac{n-1}{n} + \frac{1}{n}\frac{2}{1+e^{\beta}}\right) & \text{if } |\mathbf{y}|_1 = 1 \\ 2^{-n}\left(\frac{2}{1+e^{-\beta}}\right) & \text{if } \mathbf{y} = \mathbf{0} . \end{cases} \end{aligned}$$

Hence, the total variation distance between the uniform distribution and the distribution of the chain after one step is

$$\begin{aligned}\|UP - U\|_{\mathrm{TV}} &= \frac{1}{2} \sum_{\mathbf{y} \in \{0,1\}^n} |UP(\mathbf{y}) - U(\mathbf{y})| \\ &= 2^{-n} \frac{e^\beta - 1}{e^\beta + 1} \leqslant 2^{-n} .\end{aligned}$$

Thus, the uniform distribution is $(2^{-n}, 1)$-metastable and the thesis follows from Lemma 6.1.2. □

In the next theorem we show that, if the chain starts from a state containing at least one 1, then after $\mathcal{O}(\log n)$ time steps the distribution of the chain is ε-close to the uniform distribution, and it stays ε-close to uniform for exponential time.

Theorem 6.4.4. *Let P be the transition matrix of the logit dynamics for the n-player OR game, let U be the uniform distribution over $S = \{0,1\}^n$, let S_1 be the set of profiles $\mathbf{x} \in S$ with $|\mathbf{x}|_1 = k \geqslant 1$, and let $\varepsilon > 0$, then it holds that*

$$t_U^{S_1}(\varepsilon) \leqslant n \log(2n/\varepsilon) .$$

Proof. Let $\{X_t\}$ be the Markov chain starting at $\mathbf{x}$ and let $\{Y_t\}$ be a lazy random walk on the n-cube starting at the uniform distribution, so that X_t is distributed according to $P^t(\mathbf{x}, \cdot)$ and Y_t is uniformly distributed over $\{0,1\}^n$. Consider the following coupling (X_t, Y_t): when chain $\{X_t\}$ is at state $\mathbf{y} \in \{0,1\}^n$ then choose a position $i \in [n]$ u.a.r. and

- If $|\mathbf{y}|_1 \geqslant 2$ then, choose an action $a \in \{0,1\}$ u.a.r. and update both chains X_t and Y_t in position i with action a;
- If $|\mathbf{y}|_1 = 1$ then
  - if X_t has 0 in position i than proceed as in the previous case;
  - if X_t has 1 in position i then
    - update both chains at 0 in position i with probability $1/2$;
    - update both chains at 1 in position i with probability $1/(1 + e^\beta)$;
    - update chain X_t at 0 and chain Y_t at 1 in position i with probability $1/(1 + e^{-\beta}) - 1/2$.
- If $|\mathbf{y}|_1 = 0$ then
  - update both chains at 0 in position i with probability $1/2$;
  - update both chains at 1 in position i with probability $1/(1 + e^\beta)$;
  - update chain X_t at 0 and chain Y_t at 1 in position i with probability $1/(1 + e^{-\beta}) - 1/2$.

By construction we have that (X_t, Y_t) is a coupling of $P^t(\mathbf{x}, \cdot)$ and U, hence $\|P^t(\mathbf{x}, \cdot) - U\| \leqslant \mathbf{P}_{\mathbf{x},U}(X_t \neq Y_t)$. Moreover observe that, if at time t all players have been selected at least once and chain X_t has not yet hit profile $\mathbf{0} = (0, \cdots, 0) \in \{0,1\}^n$, then the two random variables X_t and Y_t have the same value. Hence

$$\begin{aligned}\|P^t(\mathbf{x}, \cdot) - U\| &\leqslant \mathbf{P}_{\mathbf{x},U}(X_t \neq Y_t) \\ &\leqslant \mathbf{P}_{\mathbf{x},U}(\tau_{\mathbf{0}} \leqslant t \cup \eta < t) \\ &\leqslant \mathbf{P}_{\mathbf{x},U}(\tau_{\mathbf{0}} \leqslant t) + \mathbf{P}_{\mathbf{x},U}(\eta < t) ,\end{aligned}$$

where $\tau_{\mathbf{0}}$ is the hitting time of $\mathbf{0}$ for chain X_t, and η is the first time all players have been selected at least once.

From the coupon collector's argument it follows that for every $t \geqslant n \log(2n/\varepsilon)$

$$\mathbf{P}_{\mathbf{x},U}\left(\eta < t\right) \leqslant \varepsilon/2 \,. \tag{6.15}$$

As for the second term observe that $\mathbf{P}_{\mathbf{x},U}\left(\tau_{\mathbf{0}} \leqslant t\right) \leqslant \mathbf{P}_k\left(\rho_0 \leqslant t\right)$ where ρ_0 is the hitting time of state 0 for the lazy Ehrenfest urn. More formally, consider the equivalence relation over $\Omega = \{0,1\}^n$ such that two profiles $\mathbf{x}$ and $\mathbf{y}$ are equivalent if they have the same number of 1's and let $\{Z_t\}$ be the *projection* of chain $\{X_t\}$ over the quotient space $\Omega_{\#} = \{0, 1, \dots, n\}$ of such equivalence relation. Then $\{Z_t\}$ is a Markov chain with state space $\Omega_{\#}$ and transition matrix

$$\begin{array}{rcl} P_{\#}(i, i-1) & = & \frac{i}{2n} \\ P_{\#}(i, i) & = & \frac{1}{2} \\ P_{\#}(i, i+1) & = & \frac{n-i}{2n} \,, \end{array} \tag{6.16}$$

for $i = 2, \dots, n$, and

$$\begin{array}{rcl} P_{\#}(1, 0) & = & \dfrac{1}{n(1+e^{-\beta})} \leqslant \dfrac{1}{n} \\ P_{\#}(1, 1) & = & \dfrac{n-1}{2n} + \dfrac{1}{n(1+e^{\beta})} \\ P_{\#}(1, 2) & = & \dfrac{n-1}{2n} \,. \end{array}$$

The hitting time $\tau_{\mathbf{0}}$ of state $\mathbf{0} \in S$ for chain $\{X_t\}$ coincide with the hitting time $\hat{\rho}_0$ of state $0 \in \Omega_{\#}$ for the projection Z_t.

Observe that, from the transition probabilities in (6.16), chain $\{Z_t\}$ is *almost* the lazy Ehrenfest urn, the only difference being at states 1 and 0. Moreover, the transition from state 1 to state 0 in the Z_t holds with probability smaller than the probability of the same transition in the Ehrenfest urn. From Lemmas 6.4.1 and 6.4.2 it follows that

$$\mathbf{P}_{\mathbf{x},U}\left(\tau_{\mathbf{0}} \leqslant n \log n + n \log(2/\varepsilon)\right) \leqslant c/n \,, \tag{6.17}$$

for a suitable constant $c = c(\varepsilon)$. Hence, for $t = n \log n + n \log(2/\varepsilon)$, by combining (6.15) and (6.17) it holds that

$$\|P^t(\mathbf{x}, \cdot) - U\| \leqslant \frac{\varepsilon}{2} + \frac{c}{n} \leqslant \varepsilon \,,$$

for n sufficiently large. □

6.5 Conclusions and open problems

In this chapter we introduced the concepts of *metastable distributions* and *pseudo-mixing time* and analyzed the metastability properties of the logit dynamics for some coordination games. We showed that, even when the mixing time is exponential, it is possible to find some distributions that well-approximate the behavior of the system for a very large time window. Such metastable distributions can be found even in the case of the OR-game, where no partition of the state space in metastable states exists.

In such way the logit dynamics gives us a way to describe the system even if the equilibrium has not been reached yet: this increases the power of this dynamics as a very useful tool to represent the evolution of complex systems that can be modeled as games and to predict their future status. Anyway, it is interesting to evaluate if there are other dynamics for which it is possible to prove a similar metastable behavior.

In this chapter, we have only considered some introductory games: a natural open question is whether the metastability properties observed here hold in general for any potential game.

In the case of the Ising model on the complete graph, we showed that when $\beta > c \log n/n$ the two degenerate distributions are metastable for $\mathsf{poly}(n)$ time and they are quickly reached from a large fraction of the state space. It would be interesting to investigate the metastability properties when $1/n < \beta < \log n/n$. Indeed, in that range the mixing time is exponential but the distributions concentrated in the two extremal states are not metastable.

Chapter 7

Conclusion and future directions

In this work we proposed logit dynamics as a tool to describe the evolution and to predict the future status of complex systems that can be modeled as games. This proposal has been motivated by two main properties of this dynamics: first of all, it introduces randomness in the strategy selection, allowing to represent in this way bounded rationality of agents and the intrinsic noise of complex systems; on the other hand, it induces an equilibrium that always exists and it is unique.

In order to validate our new equilibrium concept, we evaluated the performance of some simple games at this equilibrium. Our analysis concentrated in particular on the time that the dynamics requires to reach its equilibrium. To this aim, we studied the mixing time of the logit dynamics for wide classes of games.

Our results on mixing time showed that, even if there are many cases for which the convergence of the logit dynamics to the equilibrium is fast, there are games where the dynamics takes an exponential number of steps. This means that the description of the system status given by the equilibrium induced by this dynamics becomes meaningful only after a very long time. However we showed that there are games for which the logit dynamics allow to make meaningful predictions about future status even if the equilibrium has not been reached yet.

A lot of open problems still remains about the expected social welfare, the mixing time and the metastability of the logit dynamics, as discussed in Sections 4.6, 5.5, 6.5. Solving these problems will allow to gain a better understanding of the evolution process defined by the logit dynamics.

Almost every result presented in this work is relative to potential games, since for these games the Markov chain induced by the dynamics is reversible and the stationary distribution is given by the Gibbs measure. It would be of great interest to analyze the logit dynamics for games that have not an exact potential function.

The logit dynamics, as used in this work, has two main limitations: only one player updates her strategy at any time and every player has the same rationality level. The second assumption is based on an analogy with similar models in Physics, like the Ising model, where, usually, the influence on components comes from external factors, like temperature or electro-magnetic fields. Anyway, there are settings where it is possible that different players have simultaneous updates or that the updates are not immediately revealed to other players. There are also settings where the rationality level depends on the personal attitude of a player and thus may be different from a player to another. For this reason, it would be interesting to extend the logit dynamics to include simultaneous updates or different rationality levels between players. The extensions proposed above have the nice property that the resulting dynamics is a Markov chain: thus, we know that an equilibrium always exists and it is unique and a wide set of tools exists in order to handle these dynamics. However, nothing is known about which is the stationary distribution, how large is the mixing time and if metastable distributions exist.

Another interesting point is that the Markov property of the logit dynamics is a consequence of the fact that the behavior of the players disregards the past experience. It would be interesting to understand what happens if the update probability depends also on past experience. For example, is it possible to prove the convergence of randomized dynamics as stochastic fictitious play or reinforcement learning? These dynamics assume that the update probability is proportional to how much attractive a strategy is and to the rationality level: whereas the attractiveness of a strategy changes at each time step according to specific rules, the rationality level is assumed to be constant. However, it would be interesting also to know what happens if the rationality level changes over the time, for example by effect of learning.

Acknowledgements

I thank Vincenzo Auletta, Francesco Pasquale, Paolo Penna and Giuseppe Persiano for valuable discussion and contribution to my work.

A special thank goes to my beloved Principessa, who supported me during PhD, pushed me to go head even when I faced issues and difficulties, was by my side every time I need. Her presence has been really important for me during the PhD and I enjoy in thinking that she will be with me all my life.

Bibliography

[1] Carlos Alós-Ferrer and Nick Netzer. The logit-response dynamics. *Games and Economic Behavior*, 68(2):413 – 427, 2010.

[2] Arash Asadpour and Amin Saberi. On the inefficiency ratio of stable equilibria in congestion games. In *Proc. of the 5th International Workshop on Internet and Network Economics (WINE'09)*, volume 5929 of *Lecture Notes in Computer Science*, pages 545–552. Springer, 2009.

[3] Vincenzo Auletta, Diodato Ferraioli, Francesco Pasquale, Paolo Penna, and Giuseppe Persiano. Convergence to equilibrium of logit dynamics for strategic games. In *Proceedings of the 23rd ACM symposium on Parallelism in algorithms and architectures (SPAA'11)*, pages 197–206. ACM, 2011.

[4] Vincenzo Auletta, Diodato Ferraioli, Francesco Pasquale, and Giuseppe Persiano. Metastability of logit dynamics for coordination games. `http://www.dia.unisa.it/dottorandi/ferraioli/papers/meta_paper.pdf`.

[5] Vincenzo Auletta, Diodato Ferraioli, Francesco Pasquale, and Giuseppe Persiano. Mixing time and stationary expected social welfare of logit dynamics. In *3rd Int. Symp. on Algorithmic Game Theory (SAGT'10)*, volume 6386 of *LNCS*, pages 54–65. Springer, 2010.

[6] Robert J. Aumann. Subjectivity and correlation in randomized strategies. *Journal of Mathematical Economics*, 1(1):67–96, 1974.

[7] Robert J. Aumann and S. Hart, editors. *Handbook of Game Theory with Economic Applications*, volume 1. Elsevier, 1992.

[8] Robert J. Aumann and S. Hart, editors. *Handbook of Game Theory with Economic Applications*, volume 2. Elsevier, 1994.

[9] Robert J. Aumann and S. Hart, editors. *Handbook of Game Theory with Economic Applications*, volume 3. Elsevier, 2002.

[10] Robert M. Axelrod. *The evolution of cooperation*. Political Science / Science. Basic Books, 1984.

[11] Joel Beltran and Claudio Landim. Metastability of reversible finite state Markov processes. *Stochastic Processes and their Applications*, 121:1633–1677, 2011.

[12] Noam Berger, Claire Kenyon, Elchanan Mossel, and Yuval Peres. Glauber dynamics on trees and hyperbolic graphs. *Probability Theory and Related Fields*, 131:311–340, 2005.

[13] Ulrich Berger. Fictitious play in 2 x n games. *Journal of Economic Theory*, 120(2):139–154, 2005.

[14] H. D. Block and Jacob Marschak. Random Orderings and Stochastic Theories of Responses. In Olkin, Ghurye, Hoeffding, Madow, and Mann, editors, *Contributions to Probability and Statistics*, pages 97–132. Stanford University Press, 1960.

[15] Avrim Blum, Eyal Even-Dar, and Katrina Ligett. Routing without regret: on convergence to Nash equilibria of regret-minimizing algorithms in routing games. In Eric Ruppert and Dahlia Malkhi, editors, *Proceedings of the Twenty-Fifth Annual ACM Symposium on Principles of Distributed Computing (PODC)*, pages 45–52. ACM, 2006.

[16] Avrim Blum, MohammadTaghi Hajiaghayi, Katrina Ligett, and Aaron Roth. Regret minimization and the price of total anarchy. In *Proc. of 40th ACM Symposium on Theory of Computing (STOC'08)*, pages 373–382. ACM, 2008.

[17] Avrim Blum and Yishay Mansour. From external to internal regret. In *COLT*, pages 621–636, 2005.

[18] Lawrence E. Blume. The statistical mechanics of strategic interaction. *Games and Economic Behavior*, 5:387–424, 1993.

[19] Anton Bovier. Metastability: a potential theoretic approach. In *Proc. of the International Congress of Mathematicians*, volume III, pages 499–518. European Mathematical Society, 2006.

[20] Anton Bovier, Michael Eckhoff, Véronique Gayrard, and Markus Klein. Metastability in stochastic dynamics of disordered mean-field models. *Probability Theory and Related Fields*, 119:99–161, 2001.

[21] Anton Bovier, Michael Eckhoff, Véronique Gayrard, and Markus Klein. Metastability and low lying spectra in reversible Markov chains. *Communications in Mathematical Physics*, 228:219–255, 2002.

[22] Anton Bovier and Francesco Manzo. Metastability in Glauber dynamics in the low-temperature limit: Beyond exponential asymptotics. *Journal of Statistical Physics*, 107:757–779, 2002.

[23] Stephen P. Boyd and Lieven Vandenberghe. *Convex optimization*. Cambridge University Press, 2004.

[24] Felix Brandt, Felix A. Fischer, and Paul Harrenstein. On the rate of convergence of fictitious play. In *SAGT*, pages 102–113, 2010.

[25] George W. Brown. Iterative solution of games by fictitious play. *Activity Analysis of Production and Allocation*, pages 374–376, 1951.

[26] Ross Bubley and Martin E. Dyer. Path Coupling: A technique for proving rapid mixing in Markov chains. In *Proceedings of the 38th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 223–. IEEE Computer Society, 1997.

[27] Colin F. Camerer. *Behavioral game theory: experiments in strategic interaction*. The Roundtable series in behavioral economics. Russell Sage Foundation, 2003.

[28] Colin F. Camerer, Juin-Kuan Chong, and Teck-Hua Ho. Self-tuning experience weighted attraction learning in games. *Journal of Economic Theory*, 133(1):177 – 198, 2007.

[29] Colin F. Camerer and Teck-Hua Ho. Experienced-weighted attraction learning in normal form games. *Econometrica*, 67(4):827–874, 1999.

[30] C. Monica Capra, Jacob K. Goeree, Rosario Gomez, and Charles A. Holt. Anomalous behavior in a traveler's dilemma? *American Economic Review*, 89(3):678–690, 1999.

[31] Nicolò Cesa-Bianchi, Yishay Mansour, and Gilles Stoltz. Improved second-order bounds for prediction with expert advice. *Machine Learning*, 66(2-3):321–352, 2007.

[32] Xi Chen, Xiaotie Deng, and Shang-Hua Teng. Settling the complexity of computing two-player Nash equilibria. *Journal of the ACM*, 56(3), 2009.

[33] George Christodoulou and Elias Koutsoupias. The price of anarchy of finite congestion games. In *Proc. of the 37th Annual ACM Symposium on Theory of Computing (STOC'05)*, pages 67–73. ACM, 2005.

[34] Christine Chung, Katrina Ligett, Kirk Pruhs, and Aaron Roth. The price of stochastic anarchy. In *Proc. of 1st International Symposium on Algorithmic Game Theory (SAGT'08)*, volume 4997 of *LNCS*, pages 303–314. Springer, 2008.

[35] Thomas M. Cover and Joy A. Thomas. *Elements of information theory.* Wiley series in telecommunications. Wiley, 1991.

[36] Charles Darwin. *On the Origin of Species by Means of Natural Selection.* Murray, 1859.

[37] Constantinos Daskalakis, Paul W. Goldberg, and Christos H. Papadimitriou. The complexity of computing a Nash equilibrium. *SIAM Journal on Computing*, 39(1):195–259, 2009.

[38] Jian Ding, Eyal Lubetzky, and Yuval Peres. Censored Glauber dynamics for the mean field Ising model. *Journal of Statistical Physics*, 137:407–458, 2009.

[39] Jian Ding, Eyal Lubetzky, and Yuval Peres. The mixing time evolution of glauber dynamics for the mean-field ising model. *Communications in Mathematical Physics*, 289:725–764, 2009.

[40] Glenn Ellison. Learning, local interaction, and coordination. *Econometrica*, 61(5):1047–1071, 1993.

[41] Henry Eyring. The activated complex in chemical reactions. *Journal of Chemical Physics*, 3:107–115, 1935.

[42] Alex Fabrikant and Christos H. Papadimitriou. The complexity of game dynamics: Bgp oscillations, sink equilibria, and beyond. In *Proc. of the 19th annual ACM-SIAM symposium on Discrete algorithms (SODA'08)*, pages 844–853. ACM, 2008.

[43] Alex Fabrikant, Christos H. Papadimitriou, and Kunal Talwar. The complexity of pure nash equilibria. In *STOC*, pages 604–612, 2004.

[44] Dean P. Foster and Rakesh V. Vohra. Asymptotic calibration. *Biometrika*, 85(2):379–390, 1998.

[45] Mark Iosifovich Freidlin and Alexander D. Wentzell. *Random Perturbations of Dynamical Systems.* Springer, 1984.

[46] Yoav Freund and Robert E. Schapire. Adaptive game playing using multiplicative weights. *Games and Economic Behavior*, 29(1-2):79–103, 1999.

[47] Drew Fudenberg and David Marc Kreps. *Lectures on learning and equilibrium in strategic form games.* Core lecture series. Core Foundation, 1992.

[48] Drew Fudenberg and David Marc Kreps. Learning mixed equilibria. *Games and Economic Behavior*, 5(3):320–367, 1993.

[49] Drew Fudenberg and David K. Levine. Consistency and cautious fictitious play. *Journal of Economic Dynamics and Control*, 19(5-7):1065–1089, 1995.

[50] Drew Fudenberg and David K. Levine. *The Theory of Learning in Games.* MIT, 1998.

[51] Drew Fudenberg and David K. Levine. Conditional universal consistency. *Games and Economic Behavior*, 29(1-2):104 – 130, 1999.

[52] Drew Fudenberg and Jean Tirole. *Game theory.* MIT Press, 1991.

[53] Serge Galam and Bernard Walliser. Ising model versus normal form game. *Physica A: Statistical Mechanics and its Applications*, 389(3):481 – 489, 2010.

[54] Tobias Galla and J. Doyne Farmer. Complex dynamics in learning complicated games. Personal communication.

[55] Michel Goemans, Vahab Mirrokni, and Adrian Vetta. Sink equilibria and convergence. In *Proc. of the 46th Annual IEEE Symposium on Foundations of Computer Science (FOCS'05)*, pages 142–154. IEEE, 2005.

[56] James Hannan. Approximation to Bayes risk in repeated plays. *Contributions to the Theory of Games*, 3:97–139, 1957.

[57] John C. Harsanyi and Reinhard Selten. *A General Theory of Equilibrium Selection in Games.* MIT Press, 1988.

[58] Sergiu Hart and Andreu Mas-Colell. A simple adaptive procedure leading to correlated equilibrium. *Econometrica*, 68(5):1127–1150, 2000.

[59] Philip F. Henshaw. Complex systems. In Encyclopedia of Earth. `http://www.eoearth.org/article/Complex_systems?topic=49492`.

[60] Teck-Hua Ho, Colin Camerer, and Keith Weigelt. Iterated dominance and iterated best response in experimental p-beauty contests. *American Economic Review*, 88(4):947–69, 1998.

[61] Josef Hofbauer and Karl Sigmund. *Evolutionary games and population dynamics.* Cambridge University Press, 1998.

[62] Frank den Hollander. Metastability under stochastic dynamics. *Stochastic Processes and their Applications*, 114(1):1–26, 2004.

[63] Frank den Hollander. Three lectures on metastability under stochastic dynamics. In *Methods of Contemporary Mathematical Statistical Physics*, volume 1970 of *Lecture Notes in Mathematics*, pages 1–24. Springer Berlin / Heidelberg, 2009.

[64] Edwin T. Jaynes. Information theory and statistical mechanics. *Phys. Rev.*, 106:620–630, 1957.

[65] David S. Johnson, Christos H. Papadimitriou, and Mihalis Yannakakis. How easy is local search? *J. Comput. Syst. Sci.*, 37(1):79–100, 1988.

[66] Michihiro Kandori, George J. Mailath, and Rafael Rob. Learning, mutation, and long run equilibria in games. *Econometrica*, 61(1):pp. 29–56, 1993.

[67] Robert Kleinberg, Katrina Ligett, Georgios Piliouras, and Eva Tardos. Beyond the nash equilibrium barrier. In *Symposium on Innovations in Computer Science (ICS)*, 2011.

[68] Robert Kleinberg, Georgios Piliouras, and Éva Tardos. Multiplicative updates outperform generic no-regret learning in congestion games: extended abstract. In Michael Mitzenmacher, editor, *Proceedings of the 41st Annual ACM Symposium on Theory of Computing (STOC 2009)*, pages 533–542. ACM, 2009.

[69] Paul Klemperer. *Auctions: theory and practice.* Toulouse lectures in economics. Princeton University Press, 2004.

[70] Elias Koutsoupias and Christos H. Papadimitriou. Worst-case equilibria. *Computer Science Review*, 3(2):65–69, 2009. Preliminary version in STACS 1999.

[71] Hendrik Anthony Kramers. Brownian motion in a field of force and the diffusion model of chemical reactions. *Physica*, 7(4):284–304, 1940.

[72] Vijay Krishna. *Auction Theory.* Academic Press/Elsevier, 2009.

[73] Hernán Larralde, Francois Leyvraz, and David P. Sanders. Metastability in Markov processes. *Journal of Statistical Mechanics: Theory and Experiment*, 2006(8):1–29, 2006.

[74] David Levin, Malwina Luczak, and Yuval Peres. Glauber dynamics for the mean-field ising model: cut-off, critical power law, and metastability. *Probability Theory and Related Fields*, 146:223–265, 2010.

[75] David Levin, Yuval Peres, and Elizabeth L. Wilmer. *Markov Chains and Mixing Times.* American Mathematical Society, 2008.

[76] Hagay Levin, Michael Schapira, and Aviv Zohar. Interdomain routing and games. In *STOC*, pages 57–66, 2008.

[77] Bernhardt Lieberman. Human behavior in a strictly determined 3x3 matrix game. *Behavioral Science*, 5(4):317–322, 1960.

[78] Erez Lieberman, Christoph Hauert, and Martin A. Nowak. Evolutionary dynamics on graphs. *Nature*, 433(7023):312–316, 2005.

[79] Nick Littlestone and Manfred K. Warmuth. The weighted majority algorithm. *Inf. Comput.*, 108:212–261, 1994.

[80] Jan Lorenz, Heiko Rauhut, Frank Schweitzer, and Dirk Helbing. How social influence can undermine the wisdom of crowd effect. *Proceedings of the National Academy of Sciences*, 108(22):9020–9025, 2011.

[81] David J.C. MacKay. *Information theory, inference, and learning algorithms.* Cambridge University Press, 2003.

[82] Fabio Martinelli. Lectures on Glauber dynamics for discrete spin models. In *Lectures on Probability Theory and Statistics (Saint-Flour, 1997)*, volume 1717 of *Lecture Notes in Math.*, pages 93–191. Springer, 1999.

[83] Andreu Mas-Colell, Michael Dennis Whinston, and Jerry R. Green. *Microeconomic theory.* Oxford University Press, 1995.

[84] John Maynard Smith. *Evolution and the theory of games.* Cambridge University Press, 1982.

[85] John Maynard Smith and George R. Price. The logic of animal conflict. *Nature*, 246(5427):15–18, 1973.

[86] Daniel McFadden. *Conditional logit analysis of qualitative choice behavior*, volume 1, pages 105–142. Academic Press, 1974.

[87] Robert P. McIntosh. *The background of ecology: concept and theory.* Cambridge studies in ecology. Cambridge University Press, 1986.

[88] Richard D. McKelvey and Thomas R. Palfrey. Quantal response equilibria for normal form games. *Games and Economic Behavior*, 10(1):6–38, 1995.

[89] Nimrod Megiddo and Christos H. Papadimitriou. On total functions, existence theorems and computational complexity. *Theor. Comput. Sci.*, 81(2):317–324, 1991.

[90] Igal Milchtaich. Congestion games with player-specific payoff functions. *Games and Economic Behavior*, 13(1):111–124, 1996.

[91] Paul Robert Milgrom. *Putting auction theory to work.* Churchill lectures in economics. Cambridge University Press, 2004.

[92] Dov Monderer and Lloyd S. Shapley. Fictitious play property for games with identical interests. *Journal of Economic Theory*, 68(1):258 – 265, 1996.

[93] Dov Monderer and Lloyd S. Shapley. Potential games. *Games and Economic Behavior*, 14(1):124 – 143, 1996.

[94] Andrea Montanari and Amin Saberi. Convergence to equilibrium in local interaction games. In *Proc. of the 50th Annual Symposium on Foundations of Computer Science (FOCS'09)*. IEEE, 2009.

[95] Dilip Mookherjee and Barry Sopher. Learning and decision costs in experimental constant sum games. *Games and Economic Behavior*, 19(1):97–132, 1997.

[96] John H. Nachbar. "Evolutionary" selection dynamics in games: Convergence and limit properties. *International Journal of Game Theory*, 19(1):59–89, 1990.

[97] John Nash. Non-Cooperative Games. *The Annals of Mathematics*, 54(2):286–295, 1951.

[98] John Von Neumann. Zur theorie der gesellschaftsspiele. *Mathematische Annalen*, 100(1):295–320, 1928.

[99] John Von Neumann and Oskar Morgenstern. *Theory of Games and Economic Behavior.* Princeton University Press, 1944.

[100] Gregoire Nicolis and Catherine Rouvas-Nicolis. *Foundations of complex systems: nonlinear dynamics, statistical physics, information and prediction*, volume 10. World Scientific, 2007.

[101] Noam Nisan, Tim Roughgarden, Eva Tardos, and Vijay V. Vazirani. *Algorithmic Game Theory*. Cambridge University Press, 2007.

[102] Noam Nisan, Michael Schapira, and Aviv Zohar. Asynchronous best-reply dynamics. In *Proceedings of the 4th International Workshop on Internet and Network Economics*, WINE '08, pages 531–538. Springer-Verlag, 2008.

[103] Jack Ochs. Games with unique, mixed strategy equilibria: An experimental study. *Games and Economic Behavior*, 10(1):202 – 217, 1995.

[104] Enzo Olivieri and Maria Eulália Vares. *Large deviation and metastability*. Cambridge University Press, 2005.

[105] Barry O'Neill. Nonmetric test of the minimax theory of two-person zerosum games. *Proceedings of the National Academy of Science*, 84:2106–2109, 1987.

[106] Christos H. Papadimitriou. On the complexity of the parity argument and other inefficient proofs of existence. *J. Comput. Syst. Sci.*, 48(3):498–532, 1994.

[107] Hobart Peyton Young. The evolution of conventions. *Econometrica*, 61(1):57–84, 1993.

[108] Hobart Peyton Young. *Individual Strategy and Social Structure: An Evolutionary Theory of Institutions*. Princeton University Press, 1998.

[109] Hobart Peyton Young. *The diffusion of innovations in social networks*, chapter in "The Economy as a Complex Evolving System", vol. III, Lawrence E. Blume and Steven N. Durlauf, eds. Oxford University Press, 2003.

[110] Amnon Rapoport and Wilfred Amaldoss. Mixed strategies and iterative elimination of strongly dominated strategies: an experimental investigation of states of knowledge. *Journal of Economic Behavior & Organization*, 42(4):483–521, 2000.

[111] Amnon Rapoport and Richard B. Boebel. Mixed strategies in strictly competitive games: A further test of the minimax hypothesis. *Games and Economic Behavior*, 4(2):261 – 283, 1992.

[112] Klaus Ritzberger and Jorgen W Weibull. Evolutionary selection in normal-form games. *Econometrica*, 63(6):1371–99, 1995.

[113] Julia Robinson. An iterative method of solving a game. *The Annals of Mathematics*, 54(2):pp. 296–301, 1951.

[114] Robert W. Rosenthal. A class of games possessing pure-strategy nash equilibria. *International Journal of Game Theory*, 2(1):6567, 1973.

[115] Ariel Rubinstein. *Modeling Bounded Rationality*. The MIT Press, 1997.

[116] Larry Samuelson and Jianbo Zhang. Evolutionary stability in asymmetric games. *Journal of Economic Theory*, 57(2):363–391, 1992.

[117] Yuzuru Sato, Eizo Akiyama, and J. Doyne Farmer. Chaos in learning a simple two-person game. *Proceedings of the National Academy of Sciences of the United States of America*, 99(7):4748–4751, 2002.

[118] Lloyd S. Shapley. Some topics in two-person games. In Melvin Dresher, Lloyd S. Shapley, and Albert William Tucker, editors, *Advances in Game Theory*, volume 52 of *Annals of Mathematics Studies*, pages 1–29. Princeton University Press, 1964.

[119] Tom Siegfried. *A Beautiful Math: John Nash, Game Theory, and the Modern Quest for a Code of Nature*. Joseph Henry Press, 1st ed edition, 2006.

[120] Peter D. Taylor and Leo B. Jonker. Evolutionary stable strategies and game dynamics. *Mathematical Biosciences*, 40(1-2):145 – 156, 1978.

[121] Louis Leon Thurstone. A law of comparative judgment. *Psychological Review*, 34(4):273–286, 1927.

[122] Louis Leon Thurstone. Multiple factor analysis. *Psychological Review*, 38(5):406 – 427, 1931.

[123] Albert William Tucker and Robert Duncan Luce. *Contributions to the theory of games*. Number v. 4 in Annals of Mathematics Studies. Princeton University Press, 1959.

[124] John B. Van Huyck, Joseph P. Cook, and Raymond C. Battalio. Adaptive behavior and coordination failure. *Journal of Economic Behavior & Organization*, 32(4):483–503, 1997.

[125] John B. Van Huyck, John M. Wildenthal, and Raymond C. Battalio. Tacit cooperation, strategic uncertainty, and coordination failure: Evidence from repeated dominance solvable games. *Games and Economic Behavior*, 38(1):156–175, 2002.

[126] Ludwig von Bertalanffy. *General System Theory: Foundations, Development, Applications*. George Brazilier, 1969.

[127] Mark Voorneveld. Best-response potential games. *Economics Letters*, 66(3):289 – 295, 2000.

[128] Jorgen W. Weibull. *Evolutionary game theory*. MIT Press, 1997.

[129] Norbert Wiener. *Cybernetics: or, Control and communications in the animal and the machine*. Actualités scientifiques et industrielles. Hermann, 1948.

[130] David H. Wolpert. Information theory – the bridge connecting bounded rational game theory and statistical physics. In Dan Braha, Ali A. Minai, and Yaneer Bar-Yam, editors, *Complex Engineered Systems*, volume 14 of *Understanding Complex Systems*, pages 262–290. Springer Berlin / Heidelberg, 2006.

Index