Metastability of Asymptotically Well-Behaved Potential Games (Extended Abstract)

Diodato Ferraioli[1] and Carmine Ventre[2]

[1] Dip. di Informatica, Università di Salerno, Italy dferraioli@unisa.it
[2] School of Computing, Teesside University, UK C.Ventre@tees.ac.uk

Abstract. One of the main criticisms to game theory concerns the assumption of full rationality. Logit dynamics is a decentralized algorithm in which a level of irrationality (a.k.a. "noise") is introduced in players' behavior. In this context, the solution concept of interest becomes the logit equilibrium, as opposed to Nash equilibria. Logit equilibria are distributions over strategy profiles that possess several nice properties, including existence and uniqueness. However, there are games in which their computation may take exponential time. We therefore look at an approximate version of logit equilibria, called *metastable distributions*, introduced by Auletta et al. [4]. These are distributions which remain stable (i.e., players do not go too far from it) for a large number of steps (rather than forever, as for logit equilibria). The hope is that these distributions exist and can be reached quickly by logit dynamics.
We identify a class of potential games, that we name *asymptotically well-behaved*, for which the behavior of the logit dynamics is not chaotic as the number of players increases, so to guarantee meaningful asymptotic results. We prove that any such game admits distributions which are metastable no matter the level of noise present in the system, and the starting profile of the dynamics. These distributions can be quickly reached if the rationality level is not too big when compared to the inverse of the maximum difference in potential. Our proofs build on results which may be of independent interest, including some spectral characterizations of the transition matrix defined by logit dynamics for generic games and the relationship among convergence measures for Markov chains.

1 Introduction

One of the most prominent assumptions in game theory dictates that people are rational. This is contrasted by many concrete instances of people making irrational choices in certain strategic situations, such as stock markets [20]. This might be due to the incapacity of exactly determining one's own utilities: the strategic game is played with utilities perturbed by some noise. Logit dynamics [5] incorporates this noise in players' actions and then is advocated to be a good model for people behavior. In more detail, logit dynamics features a rationality level $\beta \geq 0$ (equivalently, a noise level $1/\beta$) and each player is assumed to play a strategy with a probability which is proportional to her corresponding utility and β. So the higher β is, the less noise there is and the more rational players are. Logit dynamics can then be seen as a noisy best-response dynamics.

The natural equilibrium concept for logit dynamics is defined by a probability distribution over the pure strategy profiles of the game. Whilst for best-response dynamics Pure Nash equilibria (PNE) are stable states, in logit dynamics there is a chance, which is inversely proportional to β, that players deviate from such strategy profiles. PNE are then not an adequate solution concept for this dynamics. However, the random process defined by the logit dynamics can be modeled via an ergodic Markov chain. Stability in Markov chains is represented by the concept of stationary distributions. These distributions, dubbed logit equilibria, are suggested as a suitable solution concept in this context due to their properties [3]. For example, from the results known in Markov chain literature, we know that any game possesses a logit equilibrium and that this equilibrium is unique. The absence of either of these guarantees is often considered a weakness of PNE. Nevertheless, as for (P)NE, the computation of logit equilibria may be computationally hard depending on whether the chain mixes rapidly or not [2].

As the hardness of computing (P)NE justifies approximate notions of the concept [15], so Auletta et al. [4] look at an approximation of logit equilibria that they call *metastable distributions.* These distributions aim to *describe* regularities arising during the transient phase of the dynamics before stationarity has been reached. Indeed, they are distributions that remain stable for a time which is long enough for the observer (in computer science terms, this time is assumed to be super-polynomial) rather than forever. Roughly speaking, the stability of the distributions in this concept is measured in terms of the generations living some historical era, while stationary distributions remain stable throughout all the generations. When the convergence to logit equilibria is too slow, then there are generations which are outlived by the computation of the stationary distribution. For these generations, metastable distributions grant an otherwise impossible descriptive power. (We refer the interested reader to [4] for a complete overview of the rationale of metastability.) It is unclear whether and which strategic games possess these distributions and if logit dynamics quickly reaches them.

The focus of this paper is the study of metastable distributions for the class of potential games [17]. Potential games are an important and widely studied class of games modeling many strategic settings. These games satisfy several appealing properties, the existence of PNE being one of them. A general study of metastability of potential games was left open by [4] and assumes particular interest due to the known hardness results (e.g., [9]) which suggest that computing a PNE for them is an intractable problem even for centralized algorithms.

Our Contribution. We aim to prove asymptotic results, in terms of the number of players n of potential games, concerning the super-polynomially long stability of metastable distributions, and the polynomial convergence time to them. This desiderata imposes some requirement on the potential games of interest, for otherwise a chaotic behavior (w.r.t. n) of the logit dynamics run on a game would not allow any meaningful asymptotic guarantee for it. We therefore identify a simple-to-describe class of potential games, termed *asymptotically well-behaved*, for which the behavior of the dynamics is "almost" the same for any number of players. Intuitively, the potential function of a game in this class has a shape which is, in a sense, immaterial from the actual value of n. As an example, consider the potential minimized when all the players agree on a strategy x, maximized when no player plays x, and that increases as the number of players playing x decreases. The technical definition of this notion is given in Section 4. We stress that similar assumptions are made in related literature on logit dynamics either implicitly (as in [18, 2], where it is assumed that certain properties of the potential function do not change as n changes), or explicitly, by considering specific games that clearly enjoy this property [4]. Moreover, asymptotic results on the mixing time of Markov chains require an assumption on the behavior of the chain (i.e., the minimum bottleneck ratio must either be a polynomial or a super-polynomial) usually implicitly guaranteed. Given that our objective is much more complex than bounding the mixing time (i.e., measuring asymptotically the transient phase of the chain and ascertain stability of *and* convergence time to metastable distributions) we need a similar, yet stronger, requirement.

Together with the formalization of the class of games of interest, we formalize, building upon [4], the concept of asymptotic convergence/closeness to a metastable distribution, as a function of the number of players of a game. We then note, via the construction of an ad-hoc n-player potential game that not all potential games admit metastable distributions (cf. Section 2), thus showing formally that some restriction on games under consideration is necessary.

Our main result proves that any asymptotically well-behaved n-player potential game has a metastable distribution for each starting profile of the logit dynamics. These distributions remain stable for a time which is super-polynomial in n, if one is content with being within distance $\varepsilon > 0$ from the distributions. (The distance is defined in this context as the total variation distance, see below.) We also prove that the convergence rate to these distributions, called *pseudo-mixing time*, is polynomial in n for values of β not too big when compared to the (inverse of the) maximum difference in potential of neighboring profiles. Note that when β is very high then logit dynamics is "close" to the best-response dynamics and therefore it is impossible to prove quick convergence results due to the aforementioned hardness results. We then give a picture which is as complete as possible relatively to the class of well-behaved potential games.

The proof of the above results consists of two main steps. We first devise a sufficient property for any n-player (not necessarily potential) game to have, for any starting profile, a distribution that is metastable for a super-polynomial number of steps and reached in polynomial time. The main idea behind this sufficient condition is that when the dynamics starts from a subset from which it is "hard to leave" and in which it is "easy to mix", then the dynamics will stay for a long time close to the stationary distribution restricted to that subset. Moreover, if a subset is "easy-to-leave," then the dynamics will quickly reach a "hard-to-leave" subset. The sufficient property consists of a rather technical definition that is intuitively a partition of the profiles into subsets that are asymptotically "hard-to-leave & easy-to-mix" or "easy-to-leave".

The second step amounts to showing that any asymptotically well-behaved potential game admits such a partition. The proof of this result builds on a number of involved technical contributions, some of which might be of independent interest. They mainly concern Markov chains. The concepts of interest are mixing time (how long the chain takes to mix), bottleneck ratio (intuitively, how hard it is for the stationary distribution to leave a subset of states), hitting time (how long the chain takes to hit a certain subset of states) and spectral properties of the transition matrix of Markov chains. We define a procedure which computes the required partition for these games. It iteratively identifies in the set of pure strategy profiles the "hard-to-leave" subsets. To prove that these subsets are "easy-to-mix", we firstly relate the pseudo-mixing time to the mixing time of a certain family of restricted Markov chains. We then prove that the mixing time of these chains is polynomial by using a spectral characterization of these restricted chains. The proof that the remaining profiles are "easy-to-leave" mainly relies on a connection between bottleneck ratio and hitting time.

We remark that, as a byproduct of our result, we essentially close an open problem of [4] about metastability of the Curie-Weiss game.

Related Works. Logit dynamics is defined in [5]. Early works about this dynamics have focused on its long-term behavior: [5] considered 2×2 coordination games and potential games, whereas a general characterization for wider classes of games is given in [1]. Several works gave bounds on the time the dynamics takes to reach specific Nash equilibria of a game for graphical coordination games on cliques and rings [8] and more general families of graphs [19, 18]. The stationary distribution of the logit dynamics Markov chain is proposed as a new equilibrium concept in game theory in [3] and the convergence time to it studied in [2]. The logit response function has also been used for defining another equilibrium concept, known as *quantal response equilibrium* [16]. This differs from the logit equilibrium since it is a product distribution (like Nash equilibrium).

In physics, chemistry, and biology, metastability is a phenomenon related to the evolution of systems under noisy dynamics. In particular, metastability concerns moves between regions of the state spaces and the existence of multiple, well separated time scales: at short time scales, the system appears to be in a quasi-equilibrium, but really explores only a confined region of the available space state, while, at larger time scales, it undergoes transitions between such different regions. Research in physics about metastability aims at expressing typical features of a metastable state and to evaluate the transition time between metastable states. Several monographs on the subject are available in physics literature (see, e.g., [12]). The idea of metastability of probability distributions is introduced in [4] together with the concepts of metastable distribution and pseudo-mixing time for some specific potential games.

2 Preliminary Definitions

A *strategic game* $\mathcal{G}$ is a triple $([n], S_1, \dots, S_n, \mathcal{U})$, where $[n] = \{1, \dots, n\}$ is a finite set of players, $(S_1, \dots, S_n)$ is a family of non-empty finite sets (S_i is the set of strategies available to player i), and $\mathcal{U} = (u_1, \dots, u_n)$ is a family of utility functions (or payoffs), where $u_i \colon S \to \mathbb{R}$, $S = S_1 \times \dots \times S_n$ being the set of all strategy profiles, is the utility function of player i. We focus on (exact) *potential games*, i.e., games for which there exists a function $\Phi \colon S \to \mathbb{R}$ such that for any pair of $\mathbf{x}, \mathbf{y} \in S$, $\mathbf{y} = (\mathbf{x}_{-i}, y_i)$, we have $\Phi(\mathbf{x}) - \Phi(\mathbf{y}) = u_i(\mathbf{y}) - u_i(\mathbf{x})$. Note that we use the standard game theoretic notation $(\mathbf{x}_{-i}, s)$ to mean the vector obtained from $\mathbf{x}$ by replacing the i-th entry with s; i.e. $(\mathbf{x}_{-i}, s) = (x_1, \dots, x_{i-1}, s, x_{i+1}, \dots, x_n)$. For two vectors $\mathbf{x}, \mathbf{y}$, we denote with $H(\mathbf{x}, \mathbf{y}) = |\{i \colon x_i \neq y_i\}|$ their Hamming distance. For two probability distributions μ and ν on the same state space S, the *total variation distance* $\|\mu - \nu\|_{\mathrm{TV}}$ is defined as $\|\mu - \nu\|_{\mathrm{TV}} = \max_{A \subset S} |\mu(A) - \nu(A)| = \frac{1}{2} \sum_{\mathbf{x} \in S} |\mu(\mathbf{x}) - \nu(\mathbf{x})|$, where $\mu(A) = \sum_{\mathbf{x} \in A} \mu(\mathbf{x})$ and $\nu(A) = \sum_{\mathbf{x} \in A} \nu(\mathbf{x})$.

Logit Dynamics. The logit dynamics has been introduced in [5] and runs as follows: at every time step (i) Select one player $i \in [n]$ uniformly at random; (ii) Update the strategy of player i according to the *Boltzmann distribution* with parameter β over the set S_i of her strategies. That is, a strategy $s_i \in S_i$ will be selected with probability $\sigma_i(s_i \mid \mathbf{x}_{-i}) = \frac{1}{Z_i(\mathbf{x}_{-i})} e^{\beta u_i(\mathbf{x}_{-i}, s_i)}$, where $\mathbf{x}_{-i}$ is the profile of strategies played at the current time step by players different from i, $Z_i(\mathbf{x}_{-i}) = \sum_{z_i \in S_i} e^{\beta u_i(\mathbf{x}_{-i}, z_i)}$ is the normalizing factor, and $\beta \geq 0$. One can see parameter β as the inverse of the noise or, equivalently, 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 β that goes to infinity player i chooses her best response strategy (if more than one best response is available, she chooses one of them uniformly at random). The above dynamics defines a *Markov chain* $\{X_t\}_{t \in \mathbb{N}}$ with the set of strategy profiles as state space, and where the transition probability from profile $\mathbf{x} = (x_1, \dots, x_n)$ to profile $\mathbf{y} = (y_1, \dots, y_n)$, denoted $P(\mathbf{x}, \mathbf{y})$, is zero if $H(\mathbf{x}, \mathbf{y}) \geq 2$ and it is $\frac{1}{n} \sigma_i(y_i \mid \mathbf{x}_{-i})$ if the two profiles differ exactly at player i. The Markov chain defined by the logit dynamics is ergodic [5]. Hence, from every initial profile $\mathbf{x}$ the distribution $P^t(\mathbf{x}, \cdot)$ over states of S of the chain X_t starting at $\mathbf{x}$ will eventually converge to a *stationary distribution* π as t tends to infinity. We denote with $\mathbf{P}_{\mathbf{x}}(\cdot)$ the probability on an event given that the logit dynamics starts from profile $\mathbf{x}$. The principal notion to measure the rate of convergence of a Markov chain to its stationary distribution is the *mixing time*. For $0 < \varepsilon < 1/2$, the mixing time of the logit dynamics is defined as $t_{\mathrm{mix}}(\varepsilon) = \min\{t \in \mathbb{N} \colon d(t) \leq \varepsilon\}$, where usually $\varepsilon = 1/4$ or $\varepsilon = 1/(2e)$ and $d(t) = \max_{\mathbf{x} \in S} \|P^t(\mathbf{x}, \cdot) - \pi\|_{\mathrm{TV}}$. Other important concepts related to Markov chains important for our study are the *relaxation time*, which is related to the spectra of the transition matrix P, *hitting time* and *bottleneck ratio*. The *hitting time* τ_L of $L \subseteq S$ is the first time a Markov chain is in a profile of L. For an ergodic Markov chain with finite state space S, transition matrix P, and stationary distribution π, we define the *bottleneck ratio* of a non-empty $L \subseteq S$ as $B(L) = \sum_{\mathbf{x} \in L, \mathbf{y} \in S \setminus L} \pi(\mathbf{x}) P(\mathbf{x}, \mathbf{y}) \cdot (\pi(L))^{-1}$. As in [3], we call the stationary distribution π of the Markov chain defined by the logit dynamics on a game $\mathcal{G}$, the *logit equilibrium* of $\mathcal{G}$. In general, a Markov chain with transition matrix P

and state space S is said to be *reversible* with respect to a distribution π if, for all $\mathbf{x}, \mathbf{y} \in S$, it holds that $\pi(\mathbf{x})P(\mathbf{x},\mathbf{y}) = \pi(\mathbf{y})P(\mathbf{y},\mathbf{x})$. If an ergodic chain is reversible with respect to π, then π is its stationary distribution. Therefore when this happens, to simplify our exposition we simply say that the matrix P is reversible. It is known [5] that the logit dynamics for the class of potential games is reversible and the stationary distribution is the *Gibbs measure* $\pi(\mathbf{x}) = \frac{1}{Z}e^{-\beta\Phi(\mathbf{x})}$, where $Z = \sum_{\mathbf{y}\in S} e^{-\beta\Phi(\mathbf{y})}$.

Metastability. We now give formal definitions of *metastable distributions* and *pseudo-mixing time*. For a more detailed description we refer the reader to [4].

Definition 1. *Let P be the transition matrix of a Markov chain with state space S. A probability distribution μ over S is $(\varepsilon, \mathcal{T})$-*metastable *for P, with $\varepsilon > 0$ and $\mathcal{T} \in \mathbb{N}$, if for every $0 \le t \le \mathcal{T}$ it holds that $\|\mu P^t - \mu\|_{\mathrm{TV}} \le \varepsilon$. Moreover, let $L \subseteq S$ be a non-empty set of states. We define the* pseudo-mixing time *$t_\mu^L(\varepsilon)$ as $t_\mu^L(\varepsilon) = \inf\{t \in \mathbb{N}\colon\ \|P^t(\mathbf{x},\cdot) - \mu\|_{\mathrm{TV}} \le \varepsilon$ for all $\mathbf{x} \in L\}$.*

The definition of metastable distribution captures the idea of a distribution that behaves approximately like the stationary distribution: if we start from such a distribution and run the chain we stay close to it for a "long" time.

These notions apply to a single Markov chain. Anyway, [4] adopted them to evaluate the asymptotic behavior of the logit dynamics for parametrized classes of potential game, where the parameter is the number n of players. They consider a *sequence of n-player games* $\mathbf{G}$, one for each number n of players, and analyze the asymptotic properties that the logit dynamics enjoys when run on each $\mathcal{G}_k$, where $\mathcal{G}_k$ is the game in the sequence $\mathbf{G}$ with exactly k players. Thus, we do not have a single Markov chain but a sequence of them, one for each number n of players, and need to consider an asymptotic counterpart of the notions above. [4], in fact, showed that the logit dynamics for specific classes of n-player potential games enjoys the following property, that we name *asymptotic metastability.*

Definition 2. *Let $\mathbf{G}$ be a sequence of n-player strategic games. We say that the logit dynamics for $\mathbf{G}$ is* asymptotically metastable *for the rationality level β if for any $\varepsilon > 0$ there is a polynomial $p = p_\varepsilon$ and a super-polynomial $q = q_\varepsilon$ such that for each n, the logit dynamics with rationality level β for the game $\mathcal{G}_n$ converges in time at most $p(n)$ from each profile of $\mathcal{G}_n$ to a $(\varepsilon, q(n))$-metastable distribution.*

When the logit dynamics for a game is (not) asymptotically metastable, we say for brevity that the game itself is (not) asymptotically metastable. Unfortunately, asymptotic metastability cannot be proved for every (potential) game.

Lemma 1. *There is a sequence of n-player (potential) games $\mathbf{G}$ which is not asymptotically metastable for any β sufficiently high and any $\varepsilon < \frac{1}{4}$.*

Proof (idea). Consider pairs (p_j, q_j), where p_j is a polynomial asymptotically greater than p_{j-1} and q_j is a super-polynomial asymptotically smaller than q_{j-1}. Let $\mathcal{T}$ be a function sandwiched for any j between p_j and q_j. Let $\mathbf{G}$ be a sequence of n-player games such that each strategic game $\mathcal{G}_n$ in the sequence assigns to each player exactly two strategies, 0 and 1, and has potential function Φ defined as follows: For $t \leq n-1$ and profile $\mathbf{x}$ wherein exactly t players play strategy 1 we have $\Phi(\mathbf{x}) = n - t$, whereas $\Phi(\mathbf{1}) = 1 + \beta^{-1} \log(\mathcal{T}(n)/\varepsilon - 1)$, with $\mathbf{1} = (1, \ldots, 1)$. Intuitively, (p_j, q_j) will describe the behavior of the game with n_j players and will not be precise enough for the game with more than n_j players. Indeed, we can show that for any $\varepsilon < 1/4$, infinitely many values of n, each polynomial p and each super-polynomial q, the logit dynamics for $\mathcal{G}_n$ does not converge in time $O(p(n))$ from $\mathbf{1}$ to any $(\varepsilon, q(n))$-metastable distribution. $\square$

3 Asymptotic Metastability and Partitioned Games

Motivated by the result above, we next give a sufficient property for *any* (not necessarily potential) game to be asymptotically metastable. We will introduce the concept of game *partitioned* by the logit dynamics and give examples of games satisfying this notion. Then, we prove that games partitioned by the logit dynamics are asymptotically metastable. Note that we focus only on games and values of β such that the mixing time of the logit dynamics is at least super-polynomial in n, otherwise the stationary distribution enjoys the desired properties of stability and convergence. Throughout the rest of the paper we denote with β_0 the smallest value of β such that the mixing time is not polynomial.

Let $\mathbf{G}$ be a sequence of n-player games. Let P be the transition matrix of the logit dynamics on $\mathcal{G}_n$, for some $n > 0$, and let π be the corresponding stationary distribution. For $L \subseteq S$ non-empty, we define a Markov chain with state space L and transition matrix $\mathring{P}_L$ defined as follows: $\mathring{P}_L(\mathbf{x}, \mathbf{y}) = P(\mathbf{x}, \mathbf{y})$ if $\mathbf{x} \neq \mathbf{y}$, and $\mathring{P}_L(\mathbf{x}, \mathbf{y}) = 1 - \sum_{\substack{\mathbf{z} \in L, \\ \mathbf{z} \neq \mathbf{x}}} P(\mathbf{x}, \mathbf{z}) = P(\mathbf{x}, \mathbf{x}) + \sum_{\mathbf{z} \in S \setminus L} P(\mathbf{x}, \mathbf{z})$, otherwise. It easy to check that the stationary distribution of this Markov chain is given by the distribution $\pi_L(\mathbf{x}) = \frac{\pi(\mathbf{x})}{\pi(L)}$, for every $\mathbf{x} \in L$. Note also that the Markov chain defined upon $\mathring{P}_L$ is aperiodic, since the Markov chain defined upon P is, and it will be irreducible if L is a connected set. For a fixed $\varepsilon > 0$, we denote with $t_{\text{mix}}^L(\varepsilon)$ the mixing time of the chain described above. We also denote with $B_L(A)$ the bottleneck ratio of $A \subset L$ in the Markov chain with state space L and transition matrix $\mathring{P}_L$. We now introduce the definition of partitioned games.

Definition 3. *Let* $\mathbf{G}$ *be a sequence of* n*-player strategic games.* $\mathbf{G}$ *is* partitioned *by the logit dynamics for the rationality level* β *if for any* $\varepsilon > 0$ *there is a polynomial* $p = p_\varepsilon$ *and a super-polynomial* $q = q_\varepsilon$ *s. t. for any* n *there is a family of connected subsets* $R_1, \ldots, R_k$ *of the set* S *of profiles of* $\mathcal{G}_n$*, with* $k \geq 1$*, and a partition* $T_1, \ldots, T_k, N$ *of* S*, with* $T_i \subseteq R_i$ *for any* $i = 1, \ldots, k$*, such that*
1. the bottleneck ratio of R_i *is at most* $1/q(n)$*, for any* $i = 1, \ldots, k$*;*
2. the mixing time $t_{\text{mix}}^{R_i}(\varepsilon)$ *is at most* $p(n)$*, for any* $i = 1, \ldots, k$*;*
3. for any $i = 1, \ldots, k$ *and for any* $\mathbf{x} \in T_i$*, it holds that* $\mathbf{P}_{\mathbf{x}}\left(\tau_{S \setminus R_i} \leq t_{\text{mix}}^{R_i}(\varepsilon)\right) \leq \varepsilon$*;*
4. for any $\mathbf{x} \in N$*, it holds that* $\mathbf{P}_{\mathbf{x}}\left(\tau_{\bigcup_i T_i} \leq p(n)\right) \geq 1 - \varepsilon$*.*

Note that we allow in the above definition that T_i, for some $i = 1, \ldots, k$, or N are empty. Linking back to the intuition discussed in the introduction, $R_1, \ldots, R_k$

represent the "easy-to-mix" subsets of states (condition (2)); these sets play a crucial role in defining distributions that are metastable for very long time (condition (1)). However, when the logit dynamics starts close to the boundary of some R_i, it is likely to leave R_i quickly. Since we are interested in "easy-to-mix & hard-to-leave" subsets of profiles, for each R_i we identify its *core* T_i as the set of profiles from which the logit dynamics takes long time to leave R_i (condition (3)). The distinction between core and non-core profiles will help in proving that metastable distributions are quickly reached from any starting profile.

The main result of this section proves that a game partitioned by the logit dynamics is asymptotically metastable.

Theorem 1. *If a sequence of games* $\mathbf{G}$ *is partitioned by the logit dynamics for* β*, then the logit dynamics for* $\mathbf{G}$ *is asymptotically metastable for* β*.*

We remark that Definition 3 and Theorem 1 do not use any specific property of the game, and then can be extended to consider Markov chains in general.

The details of the proof are deferred to the full version [11]. Here we give examples of an actual game satisfying it (more examples can be found in [11]) as well as comments on a game not enjoying it. Consider, indeed, the following game-theoretic formulation of the *Curie-Weiss model* (the *Ising model* on the complete graph), that we call CW-*game*: each one of n players has two strategies, -1 and $+1$, and the utility of player i for profile $\mathbf{x} = (x_1, \ldots, x_n) \in \{-1, +1\}^n$ is $u_i(\mathbf{x}) = x_i \sum_{j \neq i} x_j$. Observe that for every player i it holds that $u_i(\mathbf{x}_{-i}, +1) - u_i(\mathbf{x}_{-i}, -1) = \mathcal{H}(\mathbf{x}_{-i}, -1) - \mathcal{H}(\mathbf{x}_{-i}, +1)$, where $\mathcal{H}(\mathbf{x}) = -\sum_{j \neq k} x_j x_k$, hence the CW-game is a potential game with potential function $\mathcal{H}$. It is known (see Chapter 15 in [14]) that the logit dynamics for this game (or equivalently the *Glauber dynamics* for the Curie-Weiss model) has mixing time polynomial in n for $\beta < 1/n$ and super-polynomial as long as $\beta > 1/n$. Moreover, [4] describes metastable distributions for $\beta > c \log n/n$ and shows that such distributions are quickly reached from profiles where the number of $+1$ (respectively -1) is a sufficiently large majority, namely if the magnetization k is such that $k^2 > c \log n/\beta$, where the *magnetization* of a profile $\mathbf{x}$ is defined as $M(\mathbf{x}) = \sum_i x_i$.

It has been left open what happens when β lies in the interval $(1/n, c \log n/n)$ and if a metastable distribution is quickly reached when in the starting point the number of $+1$ is close to the number of -1. We observe that next lemma, along with Theorem 1 essentially closes this problem by showing that CW-games are asymptotic metastable for $\beta \geq c/n$ for some constant $c > 1$.

Lemma 2. *Let* $\mathbf{G}$ *be a sequence of* n*-player* CW*-games. Then* $\mathbf{G}$ *is partitioned by the logit dynamics for any* $\beta > c/n$*, for constant* $c > 1$*.*

Proof (idea). Fix n and let S_+ (resp., S_-) be the set of profiles with positive (resp., negative) magnetization in $\mathcal{G}_n$. Let us set $R_1 = S_+$ and $R_2 = S_-$. It is known that the bottleneck ratio of these subset is super-polynomial for any $\beta > c/n$, for constant $c > 1$, (see, e.g., Chapter 15 in [14]). Moreover, in [13] it has been proved that the mixing time of the chain restricted to S_+ (resp. S_-) is actually $c_1 n \log n$ for some constant $c_1 > 0$. Moreover, from [6, 7] it follows that there are subsets of R_1 (resp., R_2) from which the dynamics hits a profile $\mathbf{y} \in S_-$ (resp., S_+) in a time equivalent to the mixing time of the chain restricted to S_+ with probability at most ε. Thus, these subsets correspond to T_1 and T_2. Finally, observe that from [6], we have that for each profile $\mathbf{x} \in N$, the hitting time of $T_1 \cup T_2$ is polynomial with high probability. □

Consider now the game of Lemma 1. We can then prove the following lemma.

Lemma 3. *Let* $\mathbf{G}$ *be the sequence of games defined in Lemma 1. Then, for any* β *sufficiently large,* $\mathbf{G}$ *is not partitioned by the logit dynamics.*

Proof. Fix n. The bottleneck of $\mathbf{1} = (1, \ldots, 1)$ is asymptotically larger than any polynomial and smaller than any super-polynomial. Hence, this profile cannot be contained in N and, for all i, it must be that $R_i = L$ for some $L \subseteq S$, with $\mathbf{1} \in L$, and $L \neq \{\mathbf{1}\}$. Then, for β sufficiently large $\pi_L(\mathbf{1}) \leq 1/2$. But, since the bottleneck of $\mathbf{1}$ is asymptotically larger than any polynomial, the mixing time of the chain restricted to L is not polynomial. □

4 Asymptotically Well-Behaved Potential Games

We now ask what class of potential games are partitioned by the logit dynamics. We know already that the answer must differ from the whole class of potential games, due to Lemma 1. However, it is important to understand to what extent it is possible to prove asymptotic metastability for potential games.

Our main aim is to give results, asymptotic in the number n of players, about the behavior of logit dynamics run on potential games. Clearly, it makes sense to give asymptotic results about the property of an object, only if this property is asymptotically well-defined, that is, the object is uniquely defined for infinitely many values of the parameter according to which we compute the asymptotic and the property of this object does not depend “chaotically” on this parameter. For example when we say that a graph has large expansion, we actually mean that there is a sequence of graphs indexed by the number of vertices, such that the expansion of each graph can be bounded by a single function of this number. Similarly, when we say that a Markov Chain has large mixing time, we actually mean that there is a sequence of Markov chains indexed by the number of states, such that the mixing time of each Markov chain can be bounded by a single function of this number. Yet another example arises in algorithm game theory: when we say that the Price of Anarchy of a game is large, we actually mean that there is a sequence of games indexed, for example, by the number of players such that the Price of Anarchy of each game can be bounded by a single function of this number. In this work the object of interest is a potential game and the property of interest is the behavior of the logit dynamics for this game. And thus, in our setting, it makes sense to give asymptotic results only when a potential game is uniquely defined for infinitely many n and the behavior of the logit dynamics for the potential game is not chaotic as n increases. However, giving a formal definition of what this means is not as immediate as in the case of the expansion of a graph or of the Price of Anarchy of a game. Thus, in order to gain insight on how to formalize this concept, let us look at an example of a game for which the behavior of the logit dynamics is evidently "almost the same" as n increases (these include the examples of partitioned games analyzed above) and examples in which this behavior instead changes infinitely often.

The behavior of logit dynamics for the CW-game can be described in a way that is immaterial from the actual value of n. Indeed, the potential function has two equivalent minima when either all players adopt strategy -1 or all players adopt strategy $+1$ and it increases as the number of players adopting a strategy different from the one played by the majority of agents increases. The potential reaches its maximum when each strategy is adopted by the same number of players. Moreover, regardless of the actual value of n it is easy to see that if the number of $+1$ strategies is sufficiently larger than the number of -1 strategies[1], then it must be hard for the logit dynamics to reach a profile with more -1's than $+1$'s, whereas it must be easy for the dynamics to converge to the potential minimizer in which all players are playing $+1$. Thus the logit dynamics for the CW-game is asymptotically well-behaved for our purposes. Indeed, the evolution of the dynamics when the number of players is n can be mapped into the evolution of the dynamics with more or less players, so that the time necessary to some events to happen (e.g., for reaching or leaving certain sets of profiles) can be bounded by the same function of the number of players.

A similar argument holds even for a number of other games considered in literature (i.e., pure coordination games, graphical coordination game on the ring, Pigou's congestion game) [11]. Moreover, it is not too hard to see that the finite opinion games on well-defined classes of graphs (cliques, rings, complete bipartite graphs, etc.) studied in [10] are also asymptotically well behaved. Thus it can be seen how our set of examples covers much of the spectrum of games considered in the logit dynamics literature.

A Game for Which the Logit Dynamics Chaotically Depends on n. Observe that, even though the potential function of the game in Lemma 1 can be easily described solely as a function of n, just as done above, we cannot describe how the logit dynamics for this game and a given rationality parameter β behaves as n changes. In particular, we are forced to describe the time that is necessary for the logit dynamics to leave the profile in which all players are adopting strategy 1 by enumerating infinitely many cases and not through a single function of n (thus preventing us to give asymptotic results in n). Indeed, this time, by construction, changes infinitely often and, for any tentative bound, there will always be a value of n from which that bound will turn out to be incorrect.

Asymptotically Well-Behaved Games: the Definition. From the analysis of these games, it is evident that the behavior of the logit dynamics for potential games is asymptotically well-defined when profiles of the n-player game can be associated to profiles of the n'-player game such that the probability of leaving associated

[1] The extent to which the number of $+1$ must be larger than the number of -1 can depend on n, but it can be bounded by a single function F on the number of players.

profiles can be always bounded by the same function of the number of players. Formally, we have the following definition.

Definition 4. *The logit dynamics for a sequence of n-player (potential) games is asymptotically well-behaved if there is n_0 and a small constant $0 < \lambda < 1$ such that for every $n \geq n_0$ and for every $L' \subseteq \mathcal{S}_n$, where $\mathcal{S}_n$ is the set of profiles in $\mathcal{G}_n$, there is a subset $L \subseteq \mathcal{S}_{n_0}$ and a function F_L such that $B(L) = F_L(n_0)$ and $B(L') \in [F_L(n)(1-\lambda), F_L(n)(1+\lambda)]$.*

For sake of compactness, we will simply say that the potential game is asymptotically well-behaved whenever the logit dynamics (run on it) is.

The main result of our paper follows.

Theorem 2. *Let $\mathbf{G}$ be an asymptotically well-behaved sequence of n-player potential games. Fix $\Delta(n) := \max\left\{\Phi_n(\mathbf{x}) - \Phi_n(\mathbf{y})\colon H(\mathbf{x},\mathbf{y}) = 1\right\}$, where Φ_n is the potential function of $\mathcal{G}_n$. Then, for any function ρ at most polynomial, $\mathbf{G}$ is asymptotically metastable for $\beta_0 \leq \beta \leq \frac{\rho(n)}{\Delta(n)}$.*

The dependence on $\Delta(n)$ is a by-product of the fact that the logit dynamics is not invariant to scaling of the utility function, i.e., scaling the utility function of a certain factor requires to inversely scale β to get the same logit dynamics (see [3] for a discussion). In a sense, $\beta\Delta(n)$ is the natural parameter that describes the logit dynamics. Then, according to this point of view, the requirement on β in the above theorem becomes almost natural: we, indeed, require that $\beta\Delta(n)$ is sufficiently large in order for the mixing time being not polynomial, but we also require that $\beta\Delta(n)$ is a polynomial. This assumption on $\beta\Delta(n)$ is in general necessary because when it is high enough logit dynamics roughly behaves as best-response dynamics. Moreover, in this case, the only metastable distributions have to be concentrated around the set of Nash equilibria. This is because for $\beta\Delta(n)$ very high, it is extremely unlikely that a player leaves a Nash equilibrium. Then, the hardness results about the convergence of best-response dynamics for potential games, cf. e.g. [9], imply that the convergence to metastable distributions for high $\beta\Delta(n)$ is similarly computationally hard.

The proof builds upon Theorem 1 and proves that any asymptotic well-behaved potential game is partitioned by the logit dynamics for β not too large.

5 Conclusions and Open Problems

In this work we prove that for any asymptotically well-behaved potential game and any starting point of this game there is a distribution that is metastable for super-polynomial time and it is quickly reached. In the proof we also give a sufficient condition for a game to enjoy this metastable behavior. It is a very interesting open problem to prove that this property is also *necessary.* The main obstacle for this direction consists of the fact that we do not know any tool for proving or disproving metastability of distributions that are largely different from the ones considered in this work. Also, given that our arguments are game-independent, it would be interesting to see whether sufficient and necessary conditions can be refined for specific subclass of games.

Acknowledgments. This work is supported by PRIN 2010-2011 project ARS TechnoMedia and EPSRC grant EP/M018113/1. Authors wish to thank Paul W. Goldberg for many invaluable discussions related to a number of results discussed in this paper, and an anonymous reviewer for the enlightening comments on an earlier version of this work.

References

1. C. Alós-Ferrer and N. Netzer. The logit-response dynamics. *Games and Economic Behavior*, 68(2):413 – 427, 2010.
2. V. Auletta, D. Ferraioli, F. Pasquale, P. Penna, and G. Persiano. Convergence to equilibrium of logit dynamics for strategic games. In *SPAA*, pages 197–206, 2011.
3. V. Auletta, D. Ferraioli, F. Pasquale, and G. Persiano. Mixing time and stationary expected social welfare of logit dynamics. In *SAGT*, pages 54–65, 2010.
4. V. Auletta, D. Ferraioli, F. Pasquale, and G. Persiano. Metastability of logit dynamics for coordination games. In *SODA*, pages 1006–1024, 2012.
5. L. E. Blume. The statistical mechanics of strategic interaction. *Games and Economic Behavior*, 5:387–424, 1993.
6. J. Ding, E. Lubetzky, and Y. Peres. Censored glauber dynamics for the mean field ising model. *Statistical Physics*, 137(3):407–458, 2009.
7. J. Ding, E. Lubetzky, and Y. Peres. The mixing time evolution of glauber dynamics for the mean-field ising model. *Comm. in Mathematical Physics*, 289(2):725–764, 2009.
8. G. Ellison. Learning, local interaction, and coordination. *Econometrica*, 61(5): 1047-1071, 1993.
9. A. Fabrikant, C. H. Papadimitriou, and K. Talwar. The complexity of pure nash equilibria. In *STOC*, pages 604–612, 2004.
10. D. Ferraioli, P. W. Goldberg, and C. Ventre. Decentralized dynamics for finite opinion games. In *SAGT*, pages 144–155, 2012.
11. D. Ferraioli and C. Ventre. Metastability of potential games. *CoRR*, abs/1211.2696, 2012.
12. F. Hollander. Metastability under stochastic dynamics. *Stochastic Processes and their Applications*, 114(1):1–26, 2004.
13. D. Levin, M. Luczak, and Y. 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.
14. D. Levin, Yuval P., and E. L. Wilmer. *Markov Chains and Mixing Times*. 2008.
15. R. J. Lipton, E. Markakis, and A. Mehta. Playing large games using simple strategies. In *ACM EC*, pages 36–41, 2003.
16. R. D. McKelvey and T. R. Palfrey. Quantal response equilibria for normal form games. *Games and economic behavior*, 10(1):6–38, 1995.
17. D. Monderer and L. S. Shapley. Potential games. *Games and Economic Behavior*, 14(1):124 – 143, 1996.
18. A. Montanari and A. Saberi. Convergence to equilibrium in local interaction games. In *FOCS*, 2009.
19. H. Peyton Young. *The diffusion of innovations in social networks*, chapter in “The Economy as a Complex Evolving System”, vol. III. 2003.
20. R. J. Schiller. *Irrational Exuberance.* Wiley, 2000.