

Dynamics in Near-Potential Games

Ozan Candogan^{a,*}, Asuman Ozdaglar^a, Pablo Parrilo^a

^a*Laboratory of Information and Decision Systems, Massachusetts Institute of Technology,
77 Massachusetts Avenue, Room 32-D608 Cambridge, MA 02139.*

Abstract

Except for special classes of games, there is no systematic framework for analyzing the dynamical properties of multi-agent strategic interactions. Potential games are one such special but restrictive class of games that allow for tractable dynamic analysis. Intuitively, games that are “close” to a potential game should share similar properties. In this paper, we formalize and develop this idea by quantifying to what extent the dynamic features of potential games extend to “near-potential” games.

We study convergence of three commonly studied classes of adaptive dynamics: discrete-time better/best response, logit response, and discrete-time fictitious play dynamics. For better/best response dynamics, we focus on the evolution of the sequence of pure strategy profiles and show that this sequence converges to a (pure) approximate equilibrium set, whose size is a function of the “distance” from a close potential game. We then study logit response dynamics parametrized by a smoothing parameter that determines the frequency with which the best response strategy is played. Our analysis uses a Markov chain representation for the evolution of pure strategy profiles. We provide a characterization of the stationary distribution of this Markov chain in terms of the distance of the game from a close potential game and the corresponding potential function. We further show that the stochastically stable strategy profiles (defined as those that have positive probability under the stationary distribution in the limit as the smoothing parameter goes to 0) are pure approximate equilibria. Finally, we turn attention to fictitious play, and establish that in near-potential games, the sequence of empirical frequencies of player actions converges to a neighborhood of (mixed) equilibria of the game, where the size of the neighborhood increases with distance of the game to a potential game. Thus, our results suggest that games that are close to a potential game inherit the dynamical properties of potential games. Since a close potential game to a given game can be found by solving a convex optimization problem, our approach also provides a systematic framework for studying convergence behavior of adaptive learning dynamics in arbitrary finite strategic form games.

Keywords: Dynamics in games, near-potential games, best response dynamics, logit response dynamics, fictitious play.

JEL: C61, C72, D83,

*Corresponding author. Phone: (+1)(617) 253-9842, Fax: 617-253-3578.

Email addresses: candogan@mit.edu (Ozan Candogan), asuman@mit.edu (Asuman Ozdaglar), parrilo@mit.edu (Pablo Parrilo)

1. Introduction

The study of multi-agent strategic interactions both in economics and engineering mainly relies on the concept of Nash equilibrium. This raises the question whether Nash equilibrium makes approximately accurate predictions of the user behavior. One possible justification for Nash equilibrium is that it arises as the long run outcome of dynamical processes, in which less than fully rational players search for optimality over time. However, unless the game belongs to special (but restrictive) classes of games, such dynamics do not converge to a Nash equilibrium, and there is no systematic analysis of their limiting behavior (Jordan 1993, Fudenberg and Levine 1998, Shapley 1964).

Potential games is a class of games for which many of the simple user dynamics, such as best response dynamics and fictitious play, converge to a Nash equilibrium (Monderer and Shapley 1996b;a, Fudenberg and Levine 1998, Sandholm 2010, Young 2004). Intuitively, dynamics in potential games and dynamics in games that are “close” (in terms of the payoffs of the players) to potential games should be related. Our goal in this paper is to make this intuition precise and provide a systematic framework for studying dynamics in finite strategic form games by exploiting their relation to close potential games.

We start by illustrating via examples that this “continuity” property of limiting dynamics need not hold for arbitrary games, i.e., games that are close in terms of payoffs may have significantly different limiting behavior under simple user dynamics. Our first example focuses on better response dynamics in which at each step or strategy profile, a player (chosen consecutively or at random) updates its strategy unilaterally to one that yields a better payoff.¹

Example 1.1. *Consider two games with two players and payoffs given in Figure 1. The entries of these tables indexed by row X and column Y show payoffs of the players when the first player uses strategy X and the second player uses strategy Y . Let $0 < \theta \ll 1$. Both games have a unique Nash equilibrium: (B, B) for \mathcal{G}_1 , and the mixed strategy profile $(\frac{2}{3}A + \frac{1}{3}B, \frac{\theta}{1+\theta}A + \frac{1}{1+\theta}B)$ for \mathcal{G}_2 .*

We consider convergence of the sequence of pure strategy profiles generated by the better response dynamics. In \mathcal{G}_1 , the sequence converges to strategy profile (B, B) . In \mathcal{G}_2 , the sequence does not converge (it can be shown that the sequence follows the better response cycle (A, A) , (B, A) , (B, B) and (A, B)). Thus, trajectories are not contained in any ϵ -equilibrium set for $\epsilon < 2$.

The second example considers fictitious play dynamics, where at each step, each player maintains an (independent) empirical frequency distribution of other player’s strategies and plays a best response against it.

Example 1.2. *Consider two games with two players and payoffs given in Figure 2. Let $0 < \theta \ll 1$. It can be seen that \mathcal{G}_1 has multiple equilibria (including pure equilibria (A, A) ,*

¹Consider a game where players are not indifferent between their strategies at any strategy profile. Arbitrarily small payoff perturbations of this game lead to games which have the same better response structure as the original game. Hence, for a given game there may exist a close enough game such that the outcome of the better response dynamics in two games are identical. However, for payoff differences of given size it is always possible to find games with different better response properties as illustrated in Example 1.1.

	A	B
A	0, 1	0, 0
B	1, 0	$\theta, 2$

\mathcal{G}_1

	A	B
A	0, 1	0, 0
B	1, 0	$-\theta, 2$

\mathcal{G}_2

Figure 1: A small change in payoffs results in significantly different behavior for the pure strategy profiles generated by the better response dynamics.

(B, B) and (C, C)), whereas \mathcal{G}_2 has a unique equilibrium given by the mixed strategy profile where both players assign $1/3$ probability to each of its strategies.

	A	B	C
A	1, 1	1, 0	0, 1
B	0, 1	1, 1	1, 0
C	1, 0	0, 1	1, 1

\mathcal{G}_1

	A	B	C
A	$1 - \theta, 1 - \theta$	1, 0	0, 1
B	0, 1	$1 - \theta, 1 - \theta$	1, 0
C	1, 0	0, 1	$1 - \theta, 1 - \theta$

\mathcal{G}_2

Figure 2: A small change in payoffs results in significantly different behavior for the empirical frequencies generated by the fictitious play dynamics.

We focus on the convergence of the sequence of empirical frequencies generated by the fictitious play dynamics (under the assumption that initial empirical frequency distribution assigns probability 1 to a pure strategy profile, and whenever players are indifferent between different strategies, they choose the lexicographically smaller one). In \mathcal{G}_1 , this sequence converges to a pure equilibrium starting from any pure strategy profile. In \mathcal{G}_2 , the sequence displays oscillations similar to those seen in the Shapley game (see Shapley (1964), Fudenberg and Levine (1998)). To see this, assume that the initial empirical frequency distribution assigns probability 1 to the strategy profile (A, A) . Observe that since the underlying game is a symmetric game, empirical frequency distribution of each player will be identical at all steps. Starting from (A, A) , both players update their strategy to C . After sufficiently many updates, the empirical frequency of A falls below $\theta/(1 + \theta)$, and that of C exceeds $1/(1 + \theta)$. Thus, the payoff specifications suggest that both players start using strategy B . Similarly, after empirical frequency of B exceeds $1/(1 + \theta)$, and that of C falls below $\theta/(1 + \theta)$, then both players start playing A . Observe that update to a new strategy takes place only when one of the strategies is being used with very high probability (recall that $\theta \ll 1$) and this feature of empirical frequencies is preserved throughout. For this reason the sequence of empirical frequencies does not converge to $(1/3, 1/3, 1/3)$, the unique Nash equilibrium of \mathcal{G}_2 .

In this paper, in contrast with the preceding examples, we will show that games that are close (in terms of payoffs of players) to potential games have similar limiting dynamics to those in potential games. In particular, many reasonable adaptive dynamics “converge” to an approximate equilibrium set, whose size is a function of the distance of the game to a close potential game. Our approach relies on using the potential function of a close

potential game for the analysis of commonly studied update rules.² We note that our results hold for arbitrary strategic form games, however our characterization of limiting behavior of dynamics is more informative for games that are close to potential games. We therefore focus our investigation to such games in this paper and refer to them as *near-potential games*.

We start our analysis by introducing *maximum pairwise difference*, a measure of “closeness” of games. Let \mathbf{p} and \mathbf{q} be two strategy profiles, which differ in the strategy of a single player, say player m . We refer to the change in the payoff of player m between these two strategy profiles, as the pairwise comparison of \mathbf{p} and \mathbf{q} . Intuitively, this quantity captures how much player m can improve its utility by unilaterally deviating from strategy profile \mathbf{p} to strategy profile \mathbf{q} . For given games, the maximum pairwise difference is defined as the maximum difference between the pairwise comparisons of these games. Thus, the maximum pairwise difference captures how different two games are in terms of the utility improvements due to unilateral deviations. Since equilibria of games, and strategy updates in various update rules (such as better/best response dynamics) can be expressed in terms of unilateral deviations, maximum pairwise difference provides a measure of strategic similarities of games. We show that the closest potential game to a given game, in the sense of maximum pairwise difference, can be obtained by solving a convex optimization problem. This provides a systematic way of approximating a given game with a potential game that has a similar equilibrium set and dynamic properties, as illustrated in Example 1.3.

Example 1.3. Consider a two-player game \mathcal{G} , which is not a potential game, and the closest potential game to this game (in terms of maximum pairwise difference), $\hat{\mathcal{G}}$, given in Figure 3. The maximum pairwise difference of these games is 2, since the utility improvements in these games due to unilateral deviations differ by at most 2 (For instance consider the deviation of the column player from (A, A) to (A, B) . In \mathcal{G} this leads to a utility improvement of 6, whereas, in $\hat{\mathcal{G}}$ the improvement amount is 4). It can be seen that for both games (B, B) is the unique equilibrium. Moreover, trajectories of better response dynamics and empirical frequencies of fictitious play dynamics converge to this equilibrium in both games.

	A	B
A	8, 2	8, 8
B	2, 2	12, 10

\mathcal{G}

	A	B
A	7, 3	9, 7
B	3, 1	11, 11

$\hat{\mathcal{G}}$

Figure 3: A game (\mathcal{G}) and a nearby potential game ($\hat{\mathcal{G}}$) share similar equilibrium set and dynamic properties.

We focus on three commonly studied user dynamics: discrete-time better/best response, logit response, and discrete-time fictitious play dynamics, and establish different notions of convergence for each. We first study *better/best response dynamics*. It is known that the sequence of pure strategy profiles, which we refer to as *trajectories*, generated by these

²Throughout the paper, we use the terms *learning dynamics* and *update rules* interchangeably.

update rules converge to pure Nash equilibria in potential games (Monderer and Shapley 1996b, Young 2004). In near-potential games, a pure Nash equilibrium need not even exist. For this reason we focus on the notion of *pure approximate equilibria* or ϵ -*equilibria*, and show that in near-potential games trajectories of these update rules converge to a pure approximate equilibrium set. The size of this set only depends on the distance of the original game from a potential game, and is independent of the payoffs in the original game.

We then focus on *logit response* dynamics. With this update rule, agents, when updating their strategies, choose their best responses with high probability, but also explore other strategies with a nonzero probability. Logit response induces a Markov chain on the set of pure strategy profiles. The stationary distribution of this Markov chain is used to explain the limiting behavior of this update rule (Young 1993, Blume 1997; 1993, Alós-Ferrer and Netzer 2010, Marden and Shamma 2008). In potential games, the stationary distribution can be expressed in closed form in terms of the potential function of the game. Additionally, the *stochastically stable strategy profiles*, i.e., the strategy profiles which have nonzero stationary distribution as the exploration probability goes to zero, are those that maximize the potential function (Alós-Ferrer and Netzer 2010, Blume 1997, Marden and Shamma 2008). Exploiting their relation to close potential games, we obtain similar results for near-potential games: (i) we obtain an explicit characterization of the stationary distribution in terms of the distance of the game from a close potential game and the corresponding potential function, and (ii) we show that the stochastically stable strategy profiles are the strategy profiles that approximately maximize the potential of a close potential game, implying that they are pure approximate equilibria of the game. Our analysis relies on a novel perturbation result for Markov chains (see Theorem 5.1) which provides bounds on deviations from a stationary distribution when transition probabilities of a Markov chain are *multiplicatively* perturbed, and therefore may be of independent interest.

A summary of our convergence results on better/best response and logit response dynamics can be found in Table 1.

Update Rule	Convergence Result
Better/Best Response Dynamics	(Theorem 4.1) Trajectories of dynamics converge to $\mathcal{X}_{\delta h}$, i.e., the δh -equilibrium set of \mathcal{G} .
Logit Response Dynamics (with parameter τ)	(Corollary 5.2) Stationary distribution μ_τ of logit response dynamics is such that $\left \mu_\tau(\mathbf{p}) - \frac{e^{\frac{1}{\tau}\phi(\mathbf{p})}}{\sum_{\mathbf{q} \in E} e^{\frac{1}{\tau}\phi(\mathbf{q})}} \right \leq \frac{e^{\frac{2\delta(h-1)}{\tau}} - 1}{e^{\frac{2\delta(h-1)}{\tau}} + 1}$, for all \mathbf{p} .
Logit Response Dynamics	(Corollary 5.3) Stochastically stable strategy profiles of \mathcal{G} are (i) contained in $S = \{\mathbf{p} \phi(\mathbf{p}) \geq \max_{\mathbf{q}} \phi(\mathbf{q}) - 4\delta(h-1)\}$, (ii) $4\delta h$ -equilibria of \mathcal{G} .

Table 1: Convergence properties of better/best response and logit response dynamics in near-potential games. Given a game \mathcal{G} , we use $\hat{\mathcal{G}}$ to denote a nearby potential game with potential function ϕ such that the distance (in terms of the maximum pairwise difference, defined in Section 2) between the two games is δ . We use the notation \mathcal{X}_ϵ to denote the ϵ -equilibrium set of the original game, h to denote the number of strategy profiles, μ_τ and $\hat{\mu}_\tau$ to denote the stationary distributions of logit response dynamics in \mathcal{G} and $\hat{\mathcal{G}}$, respectively.

We finally analyze *fictitious play dynamics* in near-potential games. In potential games trajectories of fictitious play need not converge to a Nash equilibrium, but the empirical frequencies of the played strategies converge to a (mixed) Nash equilibrium (Monderer and Shapley 1996a, Shamma and Arslan 2004). In our analysis of fictitious play dynamics, we first show that in near-potential games if the empirical frequencies are outside some ϵ -equilibrium set, then the potential of the close potential game (evaluated at the empirical frequency distribution) increases with each strategy update. Using this result we establish convergence of fictitious play dynamics to a set which can be characterized in terms of the ϵ -equilibrium set of the game and the level sets of the potential function of a close potential game. This result suggests that in near-potential games, the empirical frequencies of fictitious play converge to a set of mixed strategies that (in the close potential game) have potential almost as large as the potential of Nash equilibria. Moreover, exploring the property that for small ϵ , ϵ -equilibria are contained in disjoint neighborhoods of equilibria, we strengthen our result and establish that if a game is sufficiently close to a potential game, then empirical frequencies of fictitious play dynamics converge to a small neighborhood of equilibria. This result recovers as a special case convergence of empirical frequencies to Nash equilibria in potential games.

A summary of our results on convergence of fictitious play dynamics is given in Table 2.

Update Rule	Convergence Result
Fictitious Play	(Corollary 6.1) Empirical frequencies of dynamics converge to the set of mixed strategies with large enough potential: $\{\mathbf{x} \in \prod_m \Delta E^m \phi(\mathbf{x}) \geq \min_{\mathbf{y} \in \mathcal{X}_{M\delta}} \phi(\mathbf{y})\}$
Fictitious Play	(Theorem 6.2) Assume that \mathcal{G} has finitely many equilibria. There exists some $\bar{\delta} > 0$, and $\bar{\epsilon} > 0$ (which are functions of utilities of \mathcal{G} but not δ) such that if $\delta < \bar{\delta}$, then the empirical frequencies of fictitious play converge to $\left\{ \mathbf{x} \mid \ \mathbf{x} - \mathbf{x}_k\ \leq \frac{4f(M\delta)ML}{\epsilon} + f(M\delta + \epsilon), \text{ for some equilibrium } \mathbf{x}_k \right\},$ for any ϵ such that $\bar{\epsilon} \geq \epsilon > 0$, where $f : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ is an upper semicontinuous function such that $f(x) \rightarrow 0$ as $x \rightarrow 0$.

Table 2: Convergence properties of fictitious play dynamics in near-potential games. We denote the number of players in the game by M , set of mixed strategies of player m by ΔE^m , and the Lipschitz constant of the mixed extension of ϕ by L . Rest of the notation is the same as in Table 1.

The framework provided in this paper enables us to study the limiting behavior of adaptive user dynamics in arbitrary finite strategic form games. In particular, for a given game we can use the proposed convex optimization formulation to find a nearby potential game and use the distance between these games to obtain a quantitative characterization of the limiting approximate equilibrium set. The characterization this approach provides will be tighter if the original game is closer to a potential game.

Related Literature: Potential games play an important role in game-theoretic analysis because of existence of pure strategy Nash equilibrium, and the stability (under various learning dynamics such as better/best response dynamics) of pure Nash equilibria in these games (Monderer and Shapley 1996b, Fudenberg and Levine 1998, Young 2004). Because of these properties, potential games found applications in various control and resource allocation problems (Monderer and Shapley 1996b, Marden et al. 2009a, Candogan et al. 2010a, Arslan et al. 2007).

There is no systematic framework for analyzing the limiting behavior of many of the adaptive update rules in general games (Jordan 1993, Fudenberg and Levine 1998, Shapley 1964). However, for potential games there is a long line of literature establishing convergence of natural adaptive dynamics such as better/best response dynamics (Monderer and Shapley 1996b, Young 2004), fictitious play (Monderer and Shapley 1996a, Shamma and Arslan 2004, Marden et al. 2009b, Hofbauer and Sandholm 2002) and logit response dynamics (Blume 1993; 1997, Alós-Ferrer and Netzer 2010, Marden and Shamma 2008).

It was shown in recent work that a close potential game to a given game can be obtained by solving a convex optimization problem (see Candogan et al. (2011; 2010b)). It was also proved that equilibria of a given game can be characterized by first approximating this game with a potential game, and then using the equilibrium properties of close potential games (Candogan et al. 2011; 2010b). This paper builds on this line of work to study dynamics in games by exploiting their relation to a close potential game.

Paper Organization: The rest of the paper is organized as follows: We present the game theoretic preliminaries for our work in Section 2. In Section 3, we explain how a close potential game to a given game can be found, and discuss possible extensions of this approach. We present an analysis of better and best response dynamics in near-potential games in Section 4. In Section 5, we extend our analysis to logit response, and focus on the stationary distribution and stochastically stable states of logit response. We present the results on fictitious play, and its extensions in Section 6. We close in Section 7 with concluding remarks and future work.

2. Preliminaries

In this section, we present the game-theoretic background that is relevant to our work. Additionally, we introduce the closeness measure for games, which is used in the rest of the paper.

2.1. Finite Strategic Form Games

Our focus in this paper is on finite strategic form games. A (noncooperative) finite game in strategic form consists of:

- A finite set of players, denoted by $\mathcal{M} = \{1, \dots, M\}$.
- Strategy spaces: A finite set of strategies (or actions) E^m , for every $m \in \mathcal{M}$.
- Utility functions: $u^m : \prod_{k \in \mathcal{M}} E^k \rightarrow \mathbb{R}$, for every $m \in \mathcal{M}$.

We denote a (strategic form) game instance by the tuple $\langle \mathcal{M}, \{E^m\}_{m \in \mathcal{M}}, \{u^m\}_{m \in \mathcal{M}} \rangle$, and the joint strategy space of this game instance by $E = \prod_{m \in \mathcal{M}} E^m$. We refer to a collection of strategies of all players as a *strategy profile* and denote it by $\mathbf{p} = (p^1, \dots, p^M) \in E$. The collection of strategies of all players but the m th one is denoted by \mathbf{p}^{-m} .

The basic solution concept in a noncooperative game is that of a Nash Equilibrium (NE). A (pure) Nash equilibrium is a strategy profile from which no player can unilaterally deviate and improve its payoff. Formally, \mathbf{p} is a Nash equilibrium if

$$u^m(q^m, \mathbf{p}^{-m}) - u^m(p^m, \mathbf{p}^{-m}) \leq 0,$$

for every $q^m \in E^m$ and $m \in \mathcal{M}$.

To address strategy profiles that are approximately a Nash equilibrium, we use the concept of ϵ -equilibrium. A strategy profile $\mathbf{p} \triangleq (p^1, \dots, p^M)$ is an ϵ -equilibrium ($\epsilon \geq 0$) if

$$u^m(q^m, \mathbf{p}^{-m}) - u^m(p^m, \mathbf{p}^{-m}) \leq \epsilon$$

for every $q^m \in E^m$ and $m \in \mathcal{M}$. We denote the set of ϵ -equilibria in a game \mathcal{G} by \mathcal{X}_ϵ . Note that a Nash equilibrium is an ϵ -equilibrium with $\epsilon = 0$.

2.2. Potential Games

We next describe a particular class of games that is central in this paper, the class of potential games (Monderer and Shapley 1996b).

Definition 2.1 (Potential Game). *A potential game is a noncooperative game for which there exists a function $\phi : E \rightarrow \mathbb{R}$ satisfying*

$$u^m(p^m, \mathbf{p}^{-m}) - u^m(q^m, \mathbf{p}^{-m}) = \phi(p^m, \mathbf{p}^{-m}) - \phi(q^m, \mathbf{p}^{-m}), \quad (1)$$

for every $m \in \mathcal{M}$, $p^m, q^m \in E^m$, $\mathbf{p}^{-m} \in E^{-m}$. The function ϕ is referred to as a potential function of the game.

This definition ensures that the change in the utility of a player who unilaterally deviates to a new strategy, coincides exactly with the corresponding change in the potential function. Extensions of this definition in which equation (1) holds when each utility function is multiplied with a (possibly different) positive weight, or changes in utility and potential only agree in sign, give rise to weighted and ordinal potential games that share similar properties to potential games. We briefly discuss some of these extensions in Section 3. However, our main focus in this paper is on potential games in Definition 2.1.

Some properties that are specific to potential games are evident from the definition. For instance, it can be seen that unilateral deviations from a strategy profile that maximizes the potential function (weakly) decrease the utility of the deviating player. Hence, this strategy profile corresponds to a Nash equilibrium, and it follows that every potential game has a pure Nash equilibrium.

Another important property of potential games, which will be used for characterizing the limiting behavior of dynamics in near-potential games, is that the total unilateral utility improvement around a ‘‘closed path’’ is equal to zero. Before we formally state this result, we first provide some necessary definitions, which are also used in Section 4 when we analyze better/best response dynamics in near-potential games.

Definition 2.2 (Path – Closed Path – Improvement Path). A path is a collection of strategy profiles $\gamma = (\mathbf{p}_0, \dots, \mathbf{p}_N)$ such that \mathbf{p}_i and \mathbf{p}_{i+1} differ in the strategy of exactly one player. A path is a closed path (or a cycle) if $\mathbf{p}_0 = \mathbf{p}_N$. A path is an improvement path if $u^{m_i}(\mathbf{p}_i) \geq u^{m_i}(\mathbf{p}_{i-1})$ where m_i is the player who modifies its strategy when the strategy profile is updated from \mathbf{p}_{i-1} to \mathbf{p}_i .

The transition from strategy profile \mathbf{p}_{i-1} to \mathbf{p}_i is referred to as *step i of the path*. The length of a path is equal to its number of steps, i.e., the length of the path $\gamma = (\mathbf{p}_0, \dots, \mathbf{p}_N)$ is N . We say that a closed path is *simple* if no strategy profile other than the first and the last strategy profiles is repeated along the path. For any path $\gamma = (\mathbf{p}_0, \dots, \mathbf{p}_N)$ let $I(\gamma)$ represent the total utility improvement along the path, i.e.,

$$I(\gamma) = \sum_{i=1}^N u^{m_i}(\mathbf{p}_i) - u^{m_i}(\mathbf{p}_{i-1}),$$

where m_i is the index of the player that modifies its strategy in the i th step of the path. The following proposition provides a necessary and sufficient condition under which a given game is a potential game.

Proposition 2.1 (Monderer and Shapley (1996b)). A game is a potential game if and only if $I(\gamma) = 0$ for all simple closed paths γ .

We conclude this section by formally defining the measure of “closeness” of games, used in the subsequent sections.

Definition 2.3 (Maximum Pairwise Difference). Let \mathcal{G} and $\hat{\mathcal{G}}$ be two games with set of players \mathcal{M} , set of strategy profiles E , and collections of utility functions $\{u^m\}_{m \in \mathcal{M}}$ and $\{\hat{u}^m\}_{m \in \mathcal{M}}$ respectively. The maximum pairwise difference (MPD) between these games is defined as

$$d(\mathcal{G}, \hat{\mathcal{G}}) \triangleq \max_{\mathbf{p} \in E, m \in \mathcal{M}, q^m \in E^m} \left| (u^m(q^m, \mathbf{p}^{-m}) - u^m(p^m, \mathbf{p}^{-m})) - (\hat{u}^m(q^m, \mathbf{p}^{-m}) - \hat{u}^m(p^m, \mathbf{p}^{-m})) \right|.$$

Note that the pairwise difference $u^m(q^m, \mathbf{p}^{-m}) - u^m(p^m, \mathbf{p}^{-m})$ quantifies how much player m can improve its utility by unilaterally deviating from strategy profile (p^m, \mathbf{p}^{-m}) to strategy profile (q^m, \mathbf{p}^{-m}) . Thus, the MPD captures how different two games are in terms of the utility improvements due to unilateral deviations.³ We refer to pairs of games with small MPD as *close games*, and games that have a small MPD to a potential game as *near-potential games*.

³An alternative distance measure can be given by

$$d_2(\mathcal{G}, \hat{\mathcal{G}}) \triangleq \left(\sum_{\mathbf{p} \in E} \sum_{m \in \mathcal{M}, q^m \in E^m} ((u^m(q^m, \mathbf{p}^{-m}) - u^m(p^m, \mathbf{p}^{-m})) - (\hat{u}^m(q^m, \mathbf{p}^{-m}) - \hat{u}^m(p^m, \mathbf{p}^{-m})))^2 \right)^{\frac{1}{2}},$$

and this quantity corresponds to the 2-norm of the difference of \mathcal{G} and $\hat{\mathcal{G}}$ in terms of the utility improvements due to unilateral deviations. Our analysis of the limiting behavior of dynamics relies on the maximum of such utility improvement differences between a game and a near-potential game. Thus, the measure in Definition 2.3 provides tighter bounds for our dynamics results, and hence is preferred in this paper.

The MPD measures the closeness of games in terms of the difference of unilateral deviations, rather than the difference of their utility functions, i.e., quantities of the form

$$|(u^m(q^m, \mathbf{p}^{-m}) - u^m(p^m, \mathbf{p}^{-m})) - (\hat{u}^m(q^m, \mathbf{p}^{-m}) - \hat{u}^m(p^m, \mathbf{p}^{-m}))|$$

are used to identify close games, rather than quantities of the form $|u^m(p^m, \mathbf{p}^{-m}) - \hat{u}^m(p^m, \mathbf{p}^{-m})|$. This is because the difference in unilateral deviations provides a better characterization of the strategic similarities (equilibrium and dynamic properties) between two games than the difference in utility functions. This can be seen from the following example: Consider two games with utility functions $\{u^m\}$ and $\{u^m + 1\}$, i.e., in the second game players receive an additional payoff of 1 at all strategy profiles. It can be seen from the definition of Nash equilibrium that despite the difference of their utility functions, these two games share the same equilibrium set. Intuitively, since the additional payoff is obtained at all strategy profiles, it does not affect any of the strategic considerations in the game. While the utility differences between these games is nonzero, it can be seen that the MPD is equal to zero. Hence MPD identifies a strategic equivalence between these games. The recent work Candogan et al. (2011) contains a formal treatment of strategic equivalence and its implications for strategic form games.

3. Finding Near-Potential Games

In this section, we present a framework for finding the closest potential game to a given game, where the distance between the games is measured in terms of MPD. We formulate the problem of identifying such a game as a convex optimization problem, and discuss the extensions of this approach. We note that a procedure for finding near-potential games can be found in Candogan et al. (2011) and Candogan et al. (2010b). In these works the distance between games is measured in terms of a 2-norm. In this section we illustrate how similar ideas can be used when the distance is measured in terms of MPD.

It can be seen from Proposition 2.1 that a game is a potential game if and only if it satisfies linear equalities. This suggests that the set of potential games is convex⁴, i.e., if $\mathcal{G} = \langle \mathcal{M}, E, \{u^m\}_m \rangle$ and $\hat{\mathcal{G}} = \langle \mathcal{M}, E, \{\hat{u}^m\}_m \rangle$ are potential games, then $\mathcal{G}_\alpha = \langle \mathcal{M}, E, \{\alpha u^m + (1 - \alpha)\hat{u}^m\}_m \rangle$, is also a potential game provided that $\alpha \in [0, 1]$.

Assume that a game with utility functions $\{u^m\}_m$ is given. The closest potential game (in terms of MPD) to this game, with payoff functions $\{\hat{u}^m\}_m$, and potential function ϕ can

⁴A game is a weighted potential game if (1) in Definition 2.1 is replaced by

$$\phi(p^m, \mathbf{p}^{-m}) - \phi(q^m, \mathbf{p}^{-m}) = w^m (u^m(p^m, \mathbf{p}^{-m}) - u^m(q^m, \mathbf{p}^{-m})),$$

for some positive player-specific weights $\{w^m\}$. If instead of holding with equality, the left and right hand sides of (1) only agree in sign, then the game is referred to as an ordinal potential game (Monderer and Shapley 1996b). Despite the fact that weighted and ordinal potential games have similar desirable properties to potential games, their sets are nonconvex, and finding the closest weighted/ordinal potential game to a given game requires solving a nonconvex optimization problem (Candogan et al. 2010b).

be obtained by solving the following optimization problem:

$$\begin{aligned}
& \min_{\phi, \{\hat{u}^m\}_m} \max_{\mathbf{p} \in E, m \in \mathcal{M}, q^m \in E^m} \left| (u^m(q^m, \mathbf{p}^{-m}) - u^m(p^m, \mathbf{p}^{-m})) \right. \\
& \quad \left. - (\hat{u}^m(q^m, \mathbf{p}^{-m}) - \hat{u}^m(p^m, \mathbf{p}^{-m})) \right| \\
\text{(P :)} \quad & \text{s.t. } \phi(\bar{q}^m, \bar{\mathbf{p}}^{-m}) - \phi(\bar{p}^m, \bar{\mathbf{p}}^{-m}) = \hat{u}^m(\bar{q}^m, \bar{\mathbf{p}}^{-m}) - \hat{u}^m(\bar{p}^m, \bar{\mathbf{p}}^{-m}), \\
& \quad \text{for all } m \in \mathcal{M}, \bar{\mathbf{p}} \in E, \bar{q}^m \in E^m.
\end{aligned}$$

Note that the difference $(u^m(q^m, \mathbf{p}^{-m}) - u^m(p^m, \mathbf{p}^{-m})) - (\hat{u}^m(q^m, \mathbf{p}^{-m}) - \hat{u}^m(p^m, \mathbf{p}^{-m}))$ is linear in $\{\hat{u}^m\}_m$. Thus, the objective function is the maximum of such linear functions, and hence is convex in $\{\hat{u}^m\}_m$. The constraints of this optimization problem guarantee that the game with payoff functions $\{\hat{u}^m\}_m$ is a potential game with potential ϕ . Note that these constraints are linear. Therefore, it follows that (P) is a convex optimization problem that gives the closest potential game to a given game.

Let \mathcal{G}_1 and \mathcal{G}_2 be games with utility functions $\{u^m\}_m$ and $\{w^m u^m\}_m$ respectively, where for all $m \in \mathcal{M}$, $w^m \geq 1$ is a fixed weight. It can be seen that preferences of players are identical in these two games, i.e., $u^m(x^m, x^{-m}) - u^m(y^m, x^{-m}) > 0$ if and only if $w^m u^m(x^m, x^{-m}) - w^m u^m(y^m, x^{-m}) > 0$ for any $m \in \mathcal{M}$, $y^m \in \Delta E^m$ and $\mathbf{x} \in \prod_{k \in \mathcal{M}} \Delta E^k$. Thus, it follows that the equilibrium sets of these game are the same, and for many of the update rules in games⁵, such as better/best response dynamics, and fictitious play (but not logit response), the trajectories of dynamics in \mathcal{G}_1 and \mathcal{G}_2 are identical.

This observation suggests that it may also be of interest to find a close potential game to a ‘‘scaled version’’ of a given game. The following optimization formulation obtains such a potential game:

$$\begin{aligned}
& \min_{\phi, \{\hat{u}^m\}_m, \{w^m\}_m} \max_{\mathbf{p} \in E, m \in \mathcal{M}, q^m \in E^m} \left| w^m (u^m(q^m, \mathbf{p}^{-m}) - u^m(p^m, \mathbf{p}^{-m})) \right. \\
& \quad \left. - (\hat{u}^m(q^m, \mathbf{p}^{-m}) - \hat{u}^m(p^m, \mathbf{p}^{-m})) \right| \\
\text{(P2 :)} \quad & \text{s.t. } w^m \geq 1 \quad \text{for all } m \in \mathcal{M}, \\
& \quad \phi(\bar{q}^m, \bar{\mathbf{p}}^{-m}) - \phi(\bar{p}^m, \bar{\mathbf{p}}^{-m}) = \hat{u}^m(\bar{q}^m, \bar{\mathbf{p}}^{-m}) - \hat{u}^m(\bar{p}^m, \bar{\mathbf{p}}^{-m}), \\
& \quad \text{for all } m \in \mathcal{M}, \bar{\mathbf{p}} \in E, \bar{q}^m \in E^m.
\end{aligned}$$

The solution of (P2) is a potential game with utility functions $\{\hat{u}^m\}_m$. Comparing (P) and (P2) it can be seen that (P2) obtains the closest potential game⁶ (in terms of MPD) to the game with utility functions $\{w^m u^m\}_m$. Since (P2) also minimizes the objective function over $\{w^m\}$, the solution also reveals the ‘‘scaling’’ of the original game, which makes it as close as possible to a potential game.

⁵For formal definitions of these update rules see Sections 4, 5, 6.

⁶Note that the solution of (P2) is not the closest weighted potential game to the original game. Such a game can be obtained by replacing the objective function by

$$\max_{\mathbf{p} \in E, m \in \mathcal{M}, q^m \in E^m} \left| (u^m(q^m, \mathbf{p}^{-m}) - u^m(p^m, \mathbf{p}^{-m})) - w^m (\hat{u}^m(q^m, \mathbf{p}^{-m}) - \hat{u}^m(p^m, \mathbf{p}^{-m})) \right|.$$

However, this objective function leads to a nonconvex optimization formulation due to the multiplication of the terms w^m and \hat{u}^m , and the solution of this problem is different than that of (P2). See Candogan et al. (2010b) for details.

In the rest of the paper, we do not discuss how a close potential game to a given game is obtained, but we just assume that a close potential game with potential ϕ is known and the MPD between this game and the original game is δ . We provide characterization results on limiting dynamics for a given game in terms of ϕ and δ .

4. Better Response and Best Response Dynamics

In this section, we consider better and best response dynamics, and study convergence properties of these update rules in near-potential games. All of the update rules considered in this section are discrete-time update rules, i.e., players are allowed to update their strategies at time instants $t \in \mathbb{Z}_+ = \{1, 2, \dots\}$.

Best response dynamics is an update rule where at each time instant a player chooses its best response to other players' current strategy profile. In better response dynamics, on the other hand, players choose strategies that improve their payoffs, but these strategies need not be their best responses. Formal descriptions of these update rules are given below.

Definition 4.1 (Better and Best Response Dynamics). *At each time instant $t \in \{1, 2, \dots\}$, a single player is chosen at random for updating its strategy, using a probability distribution with full support over the set of players. Let m be the player chosen at some time t , and let $\mathbf{r} \in E$ denote the strategy profile that is used at time $t - 1$.*

1. *Better response dynamics is the update process where player m does not modify its strategy if $u^m(\mathbf{r}) = \max_{q^m} u^m(q^m, \mathbf{r}^{-m})$, and otherwise it updates its strategy to a strategy in $\{q^m | u^m(q^m, \mathbf{r}^{-m}) > u^m(\mathbf{r})\}$, chosen uniformly at random.*
2. *Best response dynamics is the update process where player m does not modify its strategy if $u^m(\mathbf{r}) = \max_{q^m} u^m(q^m, \mathbf{r}^{-m})$, and otherwise it updates its strategy to a strategy in $\arg \max_{q^m} u^m(q^m, \mathbf{r}^{-m})$, chosen uniformly at random.*

For simplicity of the analysis, we assume here that users are chosen randomly to update their strategy. However, this assumption is not crucial for our results, and can be relaxed.

We refer to strategies in $\arg \max_{q^m} u^m(q^m, \mathbf{r}^{-m})$ as *best responses of player m to \mathbf{r}^{-m}* . We denote the strategy profile used at time t by \mathbf{p}_t , and we define the *trajectory of the dynamics* as the sequence of strategy profiles $\{\mathbf{p}_t\}_{t=0}^\infty$. In our analysis, we assume that the trajectory is initialized at a strategy profile $\mathbf{p}_0 \in E$ at time 0 and it evolves according to one of the update rules described above.

The following theorem establishes that in finite games, better and best response dynamics converge to a set of ϵ -equilibria, where the size of this set is characterized by the MPD to a close potential game.

Theorem 4.1. *Consider a game \mathcal{G} and let $\hat{\mathcal{G}}$ be a nearby potential game such that $d(\mathcal{G}, \hat{\mathcal{G}}) \leq \delta$. Assume that best response or better response dynamics are used in \mathcal{G} , and denote the number of strategy profiles in these games by $|E| = h$.*

For both update processes, the trajectories are contained in the δh -equilibrium set of \mathcal{G} after finite time with probability 1, i.e., let T be a random variable such that $\mathbf{p}_t \in \mathcal{X}_{\delta h}$, for all $t > T$, then $P(T < \infty) = 1$.

Proof. We prove the claim by modeling the update process using a Markov chain, and employing the improvement path condition for potential games (cf. Proposition 2.1).

Using Definition 4.1, we can represent the strategy updates in best response dynamics as the state transitions in the following Markov chain: (i) Each state corresponds to a strategy profile and, (ii) there is a nonzero transition probability from state \mathbf{r} to state $\mathbf{q} \neq \mathbf{r}$, if \mathbf{r} and \mathbf{q} differ in the strategy of a single player, say m , and q^m is a (strict) best response of player m to \mathbf{r}^{-m} . The probability of transition from state \mathbf{r} to state \mathbf{q} is equal to the probability that at strategy profile \mathbf{r} , player m is chosen for update and it chooses q^m as its new strategy. In the case of better response dynamics we allow q^m to be any strategy strictly improving payoff of player m , and a similar Markov chain representation still holds.

Since there are finitely many states, one of the recurrent classes of the Markov chain is reached in finite time (with probability 1). Thus, to prove the claim, it is sufficient to show that any state which belongs to some recurrent class of this Markov chain is contained in the ϵ -equilibrium set of \mathcal{G} .

It follows from Definition 4.1 that a recurrence class is a singleton, only if none of the players can strictly improve its payoff by unilaterally deviating from the corresponding strategy profile. Thus, such a strategy profile is a Nash equilibrium of \mathcal{G} and is contained in the ϵ -equilibrium set.

Consider a recurrence class that is not a singleton. Let \mathbf{r} be a strategy profile in this recurrence class. Since the recurrence class is not a singleton, there exists some player m , who can unilaterally deviate from \mathbf{r} by following its best response to another strategy profile \mathbf{q} , and increase its payoff by some $\alpha > 0$. Since such a transition occurs with nonzero probability, \mathbf{r} and \mathbf{q} are in the same recurrence class, and the process when started from \mathbf{r} visits \mathbf{q} and returns to \mathbf{r} in finitely many updates. Since each transition corresponds to a unilateral deviation that strictly improves the payoff of the deviating player, this constitutes a simple closed improvement path containing \mathbf{r} and \mathbf{q} . Let $\gamma = (\mathbf{p}_0, \dots, \mathbf{p}_N)$ be such an improvement path and $\mathbf{p}_0 = \mathbf{p}_N = \mathbf{r}$, $\mathbf{p}_1 = \mathbf{q}$ and $N \leq |E| = h$. Since $u^m(\mathbf{q}) - u^m(\mathbf{r}) = \alpha$, and $u^{m_i}(\mathbf{p}_i) - u^{m_i}(\mathbf{p}_{i-1}) \geq 0$ at every step i of the path, this closed improvement path satisfies

$$\sum_{i=1}^N (u^{m_i}(\mathbf{p}_i) - u^{m_i}(\mathbf{p}_{i-1})) \geq \alpha. \quad (2)$$

On the other hand it follows by Proposition 2.1 that the close potential game satisfies

$$\sum_{i=1}^N (\hat{u}^{m_i}(\mathbf{p}_i) - \hat{u}^{m_i}(\mathbf{p}_{i-1})) = 0. \quad (3)$$

Combining (2) and (3) we conclude that

$$\begin{aligned} \alpha &\leq \sum_{i=1}^N (u^{m_i}(\mathbf{p}_i) - u^{m_i}(\mathbf{p}_{i-1})) - (\hat{u}^{m_i}(\mathbf{p}_i) - \hat{u}^{m_i}(\mathbf{p}_{i-1})) \\ &\leq N\delta. \end{aligned}$$

Since $N \leq |E| = h$, it follows that $\alpha \leq \delta h$. The claim then immediately follows since \mathbf{r} and the recurrence class were chosen arbitrarily, and our analysis shows that the payoff

improvement of player m (chosen for strategy update using a probability distribution with full support as described in Definition 4.1), due to its best response is bounded by δh . \square

As can be seen from the proof of this theorem, extending dynamical properties of potential games to nearby games relies on special structural properties of potential games. As a corollary of the above theorem, we obtain that trajectories generated by better and best response dynamics converge to a Nash equilibrium in potential games, since if \mathcal{G} is a potential game, the close potential game $\hat{\mathcal{G}}$ can be chosen such that $d(\mathcal{G}, \hat{\mathcal{G}}) = 0$.

5. Logit Response Dynamics

In this section we focus on logit response dynamics. Logit response dynamics can be viewed as a smoothed version of the best response dynamics, in which a smoothing parameter determines the frequency with which the best response strategy is picked. The evolution of the pure strategy profiles can be represented in terms of a Markov chain (with state space given by the set of pure strategy profiles). We characterize the stationary distribution and stochastically stable states of this Markov chain (or of the update rule) in near-potential games. Our approach involves identifying a close potential game to a given game, and exploiting features of the corresponding potential function to characterize the limiting behavior of logit response dynamics in the original game.

In Section 5.1, we provide a formal definition of logit response dynamics and review some of its properties. We also present some of the mathematical tools used in the literature to study this update rule. In Section 5.2, we show that the stationary distribution of logit response dynamics in a near-potential game can be approximately characterized using the potential function of a nearby potential game. We also use this result to show that the stochastically stable strategy profiles are contained in approximate equilibrium sets in near-potential games.

5.1. Properties of Logit Response

We start by providing a formal definition of logit response dynamics:

Definition 5.1. *At each time instant $t \in \{1, 2, \dots\}$, a single player is chosen at random for updating its strategy, using a probability distribution with full support over the set of players. Let m be the player chosen at some time t , and let $\mathbf{r} \in E$ denote the strategy profile that is used at time $t - 1$.*

Logit response dynamics with parameter τ is the update process, where player m chooses a strategy $q^m \in E^m$ with probability

$$P_{\tau}^m(q^m | \mathbf{r}) = \frac{e^{\frac{1}{\tau} u^m(q^m, \mathbf{r}^{-m})}}{\sum_{p^m \in E^m} e^{\frac{1}{\tau} u^m(p^m, \mathbf{r}^{-m})}}.$$

In this definition, $\tau > 0$ is a fixed parameter that determines how often players choose their best responses. The probability of not choosing a best response decreases as τ decreases, and as $\tau \rightarrow 0$, players choose their best responses with probability 1. This feature suggests that logit response dynamics can be viewed as a generalization of best response dynamics, where with small but nonzero probability players use a strategy that is not a best response.

For a given $\tau > 0$, this update process can be represented by a finite aperiodic and irreducible Markov chain (Marden and Shamma 2008, Alós-Ferrer and Netzer 2010). The states of the Markov chain correspond to the strategy profiles in the game. Denoting the probability that player m is chosen for a strategy update by α_m , transition probability from strategy profile \mathbf{p} to \mathbf{q} can be given by (assuming $\mathbf{p} \neq \mathbf{q}$, and denoting the transition from \mathbf{p} to \mathbf{q} by $\mathbf{p} \rightarrow \mathbf{q}$):

$$P_\tau(\mathbf{p} \rightarrow \mathbf{q}) = \begin{cases} \alpha_m P_\tau^m(q^m | \mathbf{p}) & \text{if } \mathbf{q}^{-m} = \mathbf{p}^{-m} \text{ for some } m \in \mathcal{M} \\ 0 & \text{otherwise.} \end{cases} \quad (4)$$

The chain is aperiodic and irreducible since a player updating its strategy can choose any strategy (including the current one) with positive probability. Consequently, it has a unique stationary distribution.

We denote the stationary distribution of this Markov chain by μ_τ and refer to it as the stationary distribution of the logit response dynamics. A strategy profile \mathbf{q} such that $\lim_{\tau \rightarrow 0} \mu_\tau(\mathbf{q}) > 0$ is referred to as a *stochastically stable strategy profile* of the logit response dynamics. Intuitively, these strategy profiles are the ones that are used with nonzero probability, as players adopt their best responses more and more frequently in their strategy updates.

In potential games, the stationary distribution of the logit response dynamics can be written as an explicit function of the potential. If \mathcal{G} is a potential game with potential function ϕ , the stationary distribution of the logit response dynamics is given by the distribution (Marden and Shamma 2008, Alós-Ferrer and Netzer 2010, Blume 1997):⁷

$$\mu_\tau(\mathbf{q}) = \frac{e^{\frac{1}{\tau}\phi(\mathbf{q})}}{\sum_{\mathbf{p} \in E} e^{\frac{1}{\tau}\phi(\mathbf{p})}}. \quad (5)$$

It can be seen from (5) that $\lim_{\tau \rightarrow 0} \mu_\tau(\mathbf{q}) > 0$ if and only if $\mathbf{q} \in \arg \max_{\mathbf{p} \in E} \phi(\mathbf{p})$. Thus, in potential games the stochastically stable strategy profiles are those that maximize the potential function.

We next describe a method for obtaining the stationary distribution of Markov chains. This method will be used in the next subsection in characterizing the stationary distribution of logit response. Assume that an irreducible Markov chain over a finite set of states S , with transition probability matrix P is given. Consider a directed tree, T , with nodes given by the states of the Markov chain, and assume that an edge from node \mathbf{q} to node \mathbf{p} can exist only if there is a nonzero transition probability from \mathbf{q} to \mathbf{p} in the Markov chain. We say that the tree is rooted at state \mathbf{p} , if from every state $\mathbf{q} \neq \mathbf{p}$ there exists a unique directed path along the tree to \mathbf{p} . For each state $\mathbf{p} \in S$, denote by $\mathcal{T}(\mathbf{p})$ the set of all trees rooted at \mathbf{p} , and define a weight $w_{\mathbf{p}} \geq 0$ such that

$$w_{\mathbf{p}} = \sum_{T \in \mathcal{T}(\mathbf{p})} \prod_{(\mathbf{q} \rightarrow \mathbf{r}) \in T} P(\mathbf{q} \rightarrow \mathbf{r}). \quad (6)$$

⁷Note that this expression is independent of $\{\alpha_m\}$, i.e., the probability distribution that is used to choose which player updates its strategy has no effect on the stationary distribution of logit response.

The following proposition from the Markov Chain literature (Leighton and Rivest (1983), Anantharam and Tsoucas (1989), Freidlin and Wentzell (1998)), known as the Markov chain tree theorem, expresses the stationary distribution of Markov chains in terms of these weights.

Proposition 5.1. *The stationary distribution of the Markov chain defined over set S is given by $\mu(\mathbf{p}) = \frac{w_{\mathbf{p}}}{\sum_{\mathbf{q} \in S} w_{\mathbf{q}}}$.*

For any $T \in \mathcal{T}(\mathbf{p})$, intuitively, the quantity $\prod_{(\mathbf{q} \rightarrow \mathbf{r}) \in T} P(\mathbf{q} \rightarrow \mathbf{r})$ gives a measure of likelihood of the event that node \mathbf{p} is reached when the chain is initiated from the leaves (i.e., nodes with indegree equal to 0) of T . Thus, $w_{\mathbf{p}}$ captures how likely it is that node \mathbf{p} is visited in this chain, and the normalization in Proposition 5.1 gives the stationary distribution. Since for finite games logit response dynamics can be modeled as an irreducible Markov chain, this result can be used to characterize its stationary distribution.

5.2. Stationary Distribution of Logit Response Dynamics

In this section we show that the stationary distribution of logit response dynamics in near-potential games can be approximated by exploiting the potential function of a close potential game. We start by showing that in games with small MPD logit response dynamics have similar transition probabilities.

Lemma 5.1. *Consider a game \mathcal{G} and let $\hat{\mathcal{G}}$ be a nearby potential game such that $d(\mathcal{G}, \hat{\mathcal{G}}) \leq \delta$. Denote the transition probability matrices of logit response dynamics in \mathcal{G} and $\hat{\mathcal{G}}$ by P_{τ} and \hat{P}_{τ} respectively. For all strategy profiles \mathbf{p} and \mathbf{q} that differ in the strategy of at most one player, we have*

$$e^{-\frac{2\delta}{\tau}} \leq \hat{P}_{\tau}(\mathbf{p} \rightarrow \mathbf{q}) / P_{\tau}(\mathbf{p} \rightarrow \mathbf{q}) \leq e^{\frac{2\delta}{\tau}}.$$

Proof. Assume that $\mathbf{p}^{-m} = \mathbf{q}^{-m}$. In \mathcal{G} the transition probability $P_{\tau}(\mathbf{p} \rightarrow \mathbf{q})$ can be expressed by (see (4)):

$$P_{\tau}(\mathbf{p} \rightarrow \mathbf{q}) = \begin{cases} \alpha_m P_{\tau}^m(q^m | \mathbf{p}) & \text{if } q^m \neq p^m \\ \sum_{k \in \mathcal{M}} \alpha_k P_{\tau}^k(p^k | \mathbf{p}) & \text{otherwise.} \end{cases}$$

A similar expression holds for the transition probability $\hat{P}_{\tau}(\mathbf{p} \rightarrow \mathbf{q})$ in $\hat{\mathcal{G}}$, replacing P_{τ}^m by \hat{P}_{τ}^m . Thus, it is sufficient to prove $e^{-\frac{2\delta}{\tau}} \leq \hat{P}_{\tau}^m(q^m | \mathbf{p}) / P_{\tau}^m(q^m | \mathbf{p}) \leq e^{\frac{2\delta}{\tau}}$ for all \mathbf{p} , m , q^m to prove the claim.

Observe that by the definition of MPD

$$\begin{aligned} u^m(r^m, \mathbf{p}^{-m}) - u^m(p^m, \mathbf{p}^{-m}) - \delta &\leq \hat{u}^m(r^m, \mathbf{p}^{-m}) - \hat{u}^m(p^m, \mathbf{p}^{-m}) \\ &\leq u^m(r^m, \mathbf{p}^{-m}) - u^m(p^m, \mathbf{p}^{-m}) + \delta. \end{aligned} \quad (7)$$

Definition 5.1 suggests that $\hat{P}_{\tau}^m(q^m | \mathbf{p})$ can be written as (by dividing the numerator and the denominator by $e^{\frac{1}{\tau} \hat{u}^m(p^m, \mathbf{p}^{-m})}$):

$$\hat{P}_{\tau}^m(q^m | \mathbf{p}) = \frac{e^{\frac{1}{\tau} (\hat{u}^m(q^m, \mathbf{p}^{-m}) - \hat{u}^m(p^m, \mathbf{p}^{-m}))}}{\sum_{r^m \in E^m} e^{\frac{1}{\tau} (\hat{u}^m(r^m, \mathbf{p}^{-m}) - \hat{u}^m(p^m, \mathbf{p}^{-m}))}}.$$

Therefore, using the bounds in (7) it follows that

$$\hat{P}_\tau^m(q^m|\mathbf{p}) \leq \frac{\kappa(q^m)e^{\frac{\delta}{\tau}}}{\kappa(q^m)e^{\frac{\delta}{\tau}} + \sum_{r^m \neq q^m} \kappa(r^m)e^{-\frac{\delta}{\tau}}}.$$

where, $\kappa(r^m) = e^{\frac{1}{\tau}(u^m(r^m, \mathbf{p}^{-m}) - u^m(p^m, \mathbf{p}^{-m}))}$ for all $r^m \in E^m$. Dividing both the numerator and the denominator of the right hand side by $\sum_{r^m \in E^m} \kappa(r^m)$ and observing that $P_\tau^m(q^m|\mathbf{p}) = \frac{\kappa(q^m)}{\sum_{r^m \in E^m} \kappa(r^m)}$, we obtain

$$\hat{P}_\tau^m(q^m|\mathbf{p}) \leq \frac{e^{\frac{\delta}{\tau}} P_\tau^m(q^m|\mathbf{p})}{e^{\frac{\delta}{\tau}} P_\tau^m(q^m|\mathbf{p}) + e^{-\frac{\delta}{\tau}} (1 - P_\tau^m(q^m|\mathbf{p}))},$$

or equivalently

$$\frac{\hat{P}_\tau^m(q^m|\mathbf{p})}{P_\tau^m(q^m|\mathbf{p})} \leq \frac{e^{\frac{\delta}{\tau}}}{e^{\frac{\delta}{\tau}} P_\tau^m(q^m|\mathbf{p}) + e^{-\frac{\delta}{\tau}} (1 - P_\tau^m(q^m|\mathbf{p}))}.$$

It can be seen that the right hand side is decreasing in $P_\tau^m(q^m|\mathbf{p})$. Thus replacing $P_\tau^m(q^m|\mathbf{p})$ by 0, the right hand side can be upper bounded by $e^{\frac{2\delta}{\tau}}$. Then we obtain $\hat{P}_\tau^m(q^m|\mathbf{p})/P_\tau^m(q^m|\mathbf{p}) \leq e^{\frac{2\delta}{\tau}}$. By symmetry we also conclude that $P_\tau^m(q^m|\mathbf{p})/\hat{P}_\tau^m(q^m|\mathbf{p}) \leq e^{\frac{2\delta}{\tau}}$, and combining these bounds the claim follows. \square

Definition 5.1 suggests that perturbation of utility functions changes the transition probabilities multiplicatively in logit response. The above lemma supports this intuition: if utility gains due to unilateral deviations are modified by δ , the ratio of the transition probabilities can change at most by $e^{\frac{2\delta}{\tau}}$. Thus, if two games are close, then the transition probabilities of logit response in these games should be closely related.

This suggests using results from perturbation theory of Markov chains to characterize the stationary distribution of logit response in a near-potential game (Haviv and Van der Heyden 1984, Cho and Meyer 2001). However, standard perturbation results characterize changes in the stationary distribution of a Markov chain when the transition probabilities are *additively perturbed*. These results, when applied to multiplicative perturbations, yield bounds which are uninformative. We therefore first present a result which characterizes deviations from the stationary distribution of a Markov chain when its transition probabilities are multiplicatively perturbed, and therefore may be of independent interest.⁸

Theorem 5.1. *Let P and \hat{P} denote the probability transition matrices of two finite irreducible Markov chains with the same state space. Denote the stationary distributions of these Markov chains by μ and $\hat{\mu}$ respectively, and let the cardinality of the state space be h . Assume that $\alpha \geq 1$ is a given constant and for any two states \mathbf{p} and \mathbf{q} , the following inequalities hold*

$$\alpha^{-1}P(\mathbf{p} \rightarrow \mathbf{q}) \leq \hat{P}(\mathbf{p} \rightarrow \mathbf{q}) \leq \alpha P(\mathbf{p} \rightarrow \mathbf{q}).$$

⁸A multiplicative perturbation bound similar to ours, can be found in Freidlin and Wentzell (1998). However, this bound is looser than the one we obtain and it does not provide a good characterization of the stationary distribution in our setting. We provide a tighter bound, and obtain stronger predictions on the stationary distribution of logit response.

Then, for any state \mathbf{p} , we have

$$(i) \quad \frac{\alpha^{-(h-1)}\mu(\mathbf{p})}{\alpha^{-(h-1)}\mu(\mathbf{p}) + \alpha^{h-1}(1 - \mu(\mathbf{p}))} \leq \hat{\mu}(\mathbf{p}) \leq \frac{\alpha^{h-1}\mu(\mathbf{p})}{\alpha^{h-1}\mu(\mathbf{p}) + \alpha^{-(h-1)}(1 - \mu(\mathbf{p}))},$$

$$(ii) \quad |\mu(\mathbf{p}) - \hat{\mu}(\mathbf{p})| \leq \frac{\alpha^{h-1} - 1}{\alpha^{h-1} + 1}.$$

Proof. As before, let $\mathcal{T}(\mathbf{p})$ denote the set of directed trees that are rooted at state \mathbf{p} . Using the characterization of the stationary distribution in Proposition 5.1, for the Markov chain with probability transition matrix P , we have $\mu(\mathbf{p}) = \frac{w_{\mathbf{p}}}{\sum_{\mathbf{q}} w_{\mathbf{q}}}$, where for each state \mathbf{p} ,

$$w_{\mathbf{p}} = \sum_{T \in \mathcal{T}(\mathbf{p})} \prod_{(\mathbf{x} \rightarrow \mathbf{y}) \in T} P(\mathbf{x} \rightarrow \mathbf{y}).$$

For the Markov chain with probability transition matrix \hat{P} , we define $\hat{w}_{\mathbf{p}}$, by replacing P in the above equation with \hat{P} and $\hat{\mu}(\mathbf{p})$ similarly satisfies $\hat{\mu}(\mathbf{p}) = \frac{\hat{w}_{\mathbf{p}}}{\sum_{\mathbf{q}} \hat{w}_{\mathbf{q}}}$.

Since the Markov chain has h states, $|T| = h - 1$ for all $T \in \mathcal{T}(\mathbf{p})$. Hence, it follows from the assumption of the theorem and the above definitions of $w_{\mathbf{p}}$ and $\hat{w}_{\mathbf{p}}$ that

$$\begin{aligned} \alpha^{-(h-1)}w_{\mathbf{p}} &= \alpha^{-(h-1)} \sum_{T \in \mathcal{T}(\mathbf{p})} \prod_{(\mathbf{x} \rightarrow \mathbf{y}) \in T} P(\mathbf{x} \rightarrow \mathbf{y}) \\ &\leq \hat{w}_{\mathbf{p}} = \sum_{T \in \mathcal{T}(\mathbf{p})} \prod_{(\mathbf{x} \rightarrow \mathbf{y}) \in T} \hat{P}(\mathbf{x} \rightarrow \mathbf{y}) \\ &\leq \alpha^{h-1} \sum_{T \in \mathcal{T}(\mathbf{p})} \prod_{(\mathbf{x} \rightarrow \mathbf{y}) \in T} P(\mathbf{x} \rightarrow \mathbf{y}) = \alpha^{h-1}w_{\mathbf{p}}. \end{aligned}$$

This inequality implies that for all \mathbf{q} , $\hat{w}_{\mathbf{q}}$ is upper bounded by $\alpha^{h-1}w_{\mathbf{q}}$ and lower bounded by $\alpha^{-(h-1)}w_{\mathbf{q}}$. Using this observation together with the identity $\hat{\mu}(\mathbf{p}) = \frac{\hat{w}_{\mathbf{p}}}{\sum_{\mathbf{q}} \hat{w}_{\mathbf{q}}}$, we obtain

$$\frac{\alpha^{-(h-1)}w_{\mathbf{p}}}{\alpha^{-(h-1)}w_{\mathbf{p}} + \alpha^{h-1} \sum_{\mathbf{q} \neq \mathbf{p}} w_{\mathbf{q}}} \leq \hat{\mu}(\mathbf{p}) = \frac{\hat{w}_{\mathbf{p}}}{\sum_{\mathbf{q}} \hat{w}_{\mathbf{q}}} \leq \frac{\alpha^{h-1}w_{\mathbf{p}}}{\alpha^{h-1}w_{\mathbf{p}} + \alpha^{-(h-1)} \sum_{\mathbf{q} \neq \mathbf{p}} w_{\mathbf{q}}}.$$

Dividing the numerators and denominators of the left and right hand sides of the inequality by $\sum_{\mathbf{q}} w_{\mathbf{q}}$, using Proposition 5.1, and observing that $\sum_{\mathbf{q} \neq \mathbf{p}} \mu(\mathbf{q}) = 1 - \mu(\mathbf{p})$ the first part of the theorem follows.

Consider functions f and g defined on $[0, 1]$ such that $f(x) = \frac{\alpha^{h-1}x}{\alpha^{h-1}x + \alpha^{-(h-1)}(1-x)} - x$ and $g(x) = \frac{\alpha^{-(h-1)}x}{\alpha^{-(h-1)}x + \alpha^{h-1}(1-x)} - x$ for $x \in [0, 1]$. Checking the first order optimality conditions, it can be seen that $f(x)$ is maximized at $x = \frac{\alpha^{-(h-1)}}{1 + \alpha^{-(h-1)}}$, and the maximum equals to $\frac{\alpha^{h-1} - 1}{\alpha^{h-1} + 1}$. Similarly, the minimum of $g(x)$ is achieved at $x = \frac{\alpha^{h-1}}{1 + \alpha^{h-1}}$ and is equal to $\frac{1 - \alpha^{h-1}}{1 + \alpha^{h-1}}$. Combining these observations with part (i), we obtain

$$\begin{aligned} \frac{1 - \alpha^{h-1}}{1 + \alpha^{h-1}} &\leq g(\mu(\mathbf{p})) = \frac{\alpha^{-(h-1)}\mu(\mathbf{p})}{\alpha^{-(h-1)}\mu(\mathbf{p}) + \alpha^{h-1}(1 - \mu(\mathbf{p}))} - \mu(\mathbf{p}) \leq \hat{\mu}(\mathbf{p}) - \mu(\mathbf{p}) \\ &\leq \frac{\alpha^{h-1}\mu(\mathbf{p})}{\alpha^{h-1}\mu(\mathbf{p}) + \alpha^{-(h-1)}(1 - \mu(\mathbf{p}))} - \mu(\mathbf{p}) = f(\mu(\mathbf{p})) \leq \frac{\alpha^{h-1} - 1}{\alpha^{h-1} + 1}, \end{aligned}$$

hence the second part of the claim follows. \square

Next we use the above theorem to relate the stationary distributions of logit response dynamics in nearby games.

Corollary 5.1. *Let \mathcal{G} and $\hat{\mathcal{G}}$ be finite games with number of strategy profiles $|E| = h$, such that $d(\mathcal{G}, \hat{\mathcal{G}}) \leq \delta$. Denote the stationary distributions of logit response dynamics in these games by μ_τ , and $\hat{\mu}_\tau$ respectively. Then, for any strategy profile \mathbf{p} we have*

$$(i) \quad \frac{e^{-\frac{2\delta(h-1)}{\tau}} \mu_\tau(\mathbf{p})}{e^{-\frac{2\delta(h-1)}{\tau}} \mu_\tau(\mathbf{p}) + e^{\frac{2\delta(h-1)}{\tau}} (1 - \mu_\tau(\mathbf{p}))} \leq \hat{\mu}_\tau(\mathbf{p}) \leq \frac{e^{\frac{2\delta(h-1)}{\tau}} \mu_\tau(\mathbf{p})}{e^{\frac{2\delta(h-1)}{\tau}} \mu_\tau(\mathbf{p}) + e^{-\frac{2\delta(h-1)}{\tau}} (1 - \mu_\tau(\mathbf{p}))},$$

$$(ii) \quad |\mu_\tau(\mathbf{p}) - \hat{\mu}_\tau(\mathbf{p})| \leq \frac{e^{\frac{2\delta(h-1)}{\tau}} - 1}{e^{\frac{2\delta(h-1)}{\tau}} + 1}.$$

Proof. Proof follows from Lemma 5.1 and Theorem 5.1 by setting $\alpha = e^{\frac{2\delta}{\tau}}$. \square

The above corollary can be adapted to near-potential games, by exploiting the relation of stationary distribution of logit response and potential function in potential games (see (5)). We conclude this section by providing such a characterization of the stationary distribution of logit response dynamics in near-potential games.

Corollary 5.2. *Consider a game \mathcal{G} and let $\hat{\mathcal{G}}$ be a nearby potential game such that $d(\mathcal{G}, \hat{\mathcal{G}}) \leq \delta$. Denote the potential function of $\hat{\mathcal{G}}$ by ϕ , and the number of strategy profiles in these games by $|E| = h$. Then, the stationary distribution μ_τ of logit response dynamics in \mathcal{G} is such that*

$$(i) \quad \frac{e^{\frac{1}{\tau}(\phi(\mathbf{p}) - 2\delta(h-1))}}{e^{\frac{1}{\tau}(\phi(\mathbf{p}) - 2\delta(h-1))} + \sum_{\mathbf{q} \neq \mathbf{p} \in E} e^{\frac{1}{\tau}(\phi(\mathbf{q}) + 2\delta(h-1))}} \leq \mu_\tau(\mathbf{p})$$

$$\leq \frac{e^{\frac{1}{\tau}(\phi(\mathbf{p}) + 2\delta(h-1))}}{e^{\frac{1}{\tau}(\phi(\mathbf{p}) + 2\delta(h-1))} + \sum_{\mathbf{q} \neq \mathbf{p} \in E} e^{\frac{1}{\tau}(\phi(\mathbf{q}) - 2\delta(h-1))}},$$

$$(ii) \quad \left| \mu_\tau(\mathbf{p}) - \frac{e^{\frac{1}{\tau}\phi(\mathbf{p})}}{\sum_{\mathbf{q} \in E} e^{\frac{1}{\tau}\phi(\mathbf{q})}} \right| \leq \frac{e^{\frac{2\delta(h-1)}{\tau}} - 1}{e^{\frac{2\delta(h-1)}{\tau}} + 1}.$$

Proof. Proof follows from Corollary 5.1 and (5). \square

With simple manipulations, it can be shown that $(e^x - 1)/(e^x + 1) \leq x/2$ for $x \geq 0$. Thus, (ii) in the above corollary implies that $\left| \mu_\tau(\mathbf{p}) - \frac{e^{\frac{1}{\tau}\phi(\mathbf{p})}}{\sum_{\mathbf{q} \in E} e^{\frac{1}{\tau}\phi(\mathbf{q})}} \right| \leq \frac{\delta(h-1)}{\tau}$. Therefore, the stationary distribution of logit response dynamics in a near-potential game can be characterized in terms of the stationary distribution of this update rule in a close potential game. When τ is fixed and $\delta \rightarrow 0$, i.e., when the original game is arbitrarily close to a potential game, the stationary distribution of logit response is arbitrarily close to the stationary distribution in the potential game. On the other hand, for a fixed δ , as $\tau \rightarrow 0$, the upper bound in (ii) becomes uninformative. This is the case since $\tau \rightarrow 0$ implies that players adopt their best responses with probability 1, and thus the stationary distribution of the update rule becomes very sensitive to the difference of the game from a potential game. In this case we can still characterize the stochastically stable states of logit response using the results of Corollary 5.2, as we show in Corollary 5.3.

Corollary 5.3. *Consider a game \mathcal{G} and let $\hat{\mathcal{G}}$ be a nearby potential game with potential function ϕ and $d(\mathcal{G}, \hat{\mathcal{G}}) \leq \delta$. Denote the potential function of $\hat{\mathcal{G}}$ by $\hat{\phi}$, and the number of strategy profiles in these games by $|E| = h$. The stochastically stable strategy profiles of \mathcal{G} are (i) contained in $S = \{\mathbf{p} | \phi(\mathbf{p}) \geq \max_{\mathbf{q}} \phi(\mathbf{q}) - 4\delta(h-1)\}$, (ii) $4\delta h$ -equilibria of \mathcal{G} .*

Proof. (i) The upper bound in the first part of Corollary 5.2 implies that if \mathbf{p} is a strategy profile such that $\phi(\mathbf{p}) < \max_{\mathbf{q} \in E} \phi(\mathbf{q}) - 4\delta(h-1)$, then the stationary distribution of logit response in \mathcal{G} is such that $\mu_\tau(\mathbf{p}) \rightarrow 0$ as $\tau \rightarrow 0$. Thus, it immediately follows that the stochastically stable states in \mathcal{G} are contained in $\{\mathbf{p} \in E | \phi(\mathbf{p}) \geq \max_{\mathbf{q} \in E} \phi(\mathbf{q}) - 4\delta(h-1)\}$.

(ii) From the definition of S it follows that in $\hat{\mathcal{G}}$, none of the players can deviate from a strategy profile in S and improve its utility by more than $4\delta(h-1)$. Since $d(\mathcal{G}, \hat{\mathcal{G}}) \leq \delta$ it follows from part (i) that in \mathcal{G} , none of the players can unilaterally deviate from a stochastically stable strategy profile and improve its utility by more than $4\delta(h-1) + \delta \leq 4\delta h$. Hence stochastically stable strategy profiles of \mathcal{G} are $4\delta h$ -equilibria. \square

We conclude that in near-potential games, the stochastically stable states of logit response are the strategy profiles that approximately maximize the potential function of a close potential game. This result enables us to characterize the stochastically stable states of logit response dynamics in near-potential games, without explicitly computing the stationary distribution.

6. Fictitious Play

In this section, we investigate the convergence behavior of fictitious play in near-potential games. Unlike better/best response dynamics and logit response, in fictitious play agents maintain an empirical frequency distribution of other players' strategies and play a best response against it. Thus, analyzing fictitious play dynamics requires the notion of mixed strategies and some additional definitions that are provided in Section 6.1. In Section 6.2 we show that in finite games the empirical frequencies of fictitious play converge to a set which can be characterized in terms of the approximate equilibrium set of the game and the level sets of the potential function of a close potential game. When the original game is sufficiently close to a potential game, we strengthen this result and establish that the empirical frequencies converge to a small neighborhood of mixed equilibria of the game, and the size of this neighborhood is a function of the distance of the original game from a potential game. As a special case, our result allows us to recover the result of Monderer and Shapley (1996a), which states that in potential games the empirical frequencies of fictitious play converge to the set of mixed Nash equilibria.

6.1. Mixed Strategies and Equilibria

In this section, we introduce some additional notation and definitions, which will be used in Section 6.2 when studying convergence properties of fictitious play in near-potential games.

We start by introducing the concept of mixed strategies in games. For each player $m \in \mathcal{M}$, we denote by ΔE^m the set of probability distributions on E^m . For $x^m \in \Delta E^m$, $x^m(p^m)$ denotes the probability player m assigns to strategy $p^m \in E^m$. We refer to the distribution $x^m \in \Delta E^m$ as a *mixed strategy of player $m \in \mathcal{M}$* and to the collection $\mathbf{x} =$

$\{x^m\}_{m \in \mathcal{M}} \in \prod_m \Delta E^m$ as a *mixed strategy profile*. The mixed strategy profile of all players but the m th one is denoted by \mathbf{x}^{-m} . We use $\|\cdot\|$ to denote the standard 2-norm on $\prod_m \Delta E^m$, i.e., for $\mathbf{x} \in \prod_m \Delta E^m$, we have $\|\mathbf{x}\|^2 = \sum_{m \in \mathcal{M}} \sum_{p^m \in E^m} (x^m(p^m))^2$.

By slight (but standard) abuse of notation, we use the same notation for the mixed extension of utility function u^m of player $m \in \mathcal{M}$, i.e.,

$$u^m(\mathbf{x}) = \sum_{\mathbf{p} \in E} u^m(\mathbf{p}) \prod_{k \in \mathcal{M}} x^k(p^k), \quad (8)$$

for all $\mathbf{x} \in \prod_m \Delta E^m$. In addition, if player m uses some pure strategy q^m and other players use the mixed strategy profile \mathbf{x}^{-m} , the payoff of player m is denoted by

$$u^m(q^m, \mathbf{x}^{-m}) = \sum_{\mathbf{p}^{-m} \in E^{-m}} u^m(q^m, \mathbf{p}^{-m}) \prod_{k \in \mathcal{M}, k \neq m} x^k(p^k).$$

Similarly, we denote the mixed extension of the potential function by $\phi(\mathbf{x})$, and we use the notation $\phi(q^m, \mathbf{x}^{-m})$ to denote the potential when player m uses some pure strategy q^m and other players use the mixed strategy profile \mathbf{x}^{-m} .

A mixed strategy profile $\mathbf{x} = \{x^m\}_{m \in \mathcal{M}} \in \prod_m \Delta E^m$ is a *mixed ϵ -equilibrium* if for all $m \in \mathcal{M}$ and $p^m \in E^m$,

$$u^m(p^m, \mathbf{x}^{-m}) - u^m(x^m, \mathbf{x}^{-m}) \leq \epsilon. \quad (9)$$

Note that if the inequality holds for $\epsilon = 0$, then \mathbf{x} is referred to as a *mixed Nash equilibrium* of the game. In the rest of the paper, we use the notation \mathcal{X}_ϵ to denote the set of mixed ϵ -equilibria.

Our characterization of the limiting mixed strategy set of fictitious play depends on the number of players in the game. We use $M = |\mathcal{M}|$ as a short-hand notation for this number.

We conclude this section with two technical lemmas which summarize some properties of mixed equilibria and mixed extensions of potential and utility functions. Proofs of these lemmas can be found in the Appendix.

The first lemma establishes the Lipschitz continuity of the mixed extensions of the payoff functions and the potential function. It also shows a natural implication of continuity: for any $\epsilon' > \epsilon$, a small enough neighborhood of the ϵ -equilibrium set is contained in the ϵ' -equilibrium set.

Lemma 6.1. (i) Let $\nu : \prod_{m \in \mathcal{M}} E^m \rightarrow \mathbb{R}$ be a mapping from pure strategy profiles to real numbers. Its mixed extension is Lipschitz continuous with a Lipschitz constant of $M \sum_{p \in E} |\nu(\mathbf{p})|$ over the domain $\prod_{m \in \mathcal{M}} \Delta E^m$.

(ii) Let $\alpha \geq 0$ and $\gamma > 0$ be given. There exists a small enough $\theta > 0$ such that for any $\|\mathbf{x} - \mathbf{y}\| < \theta$ if $\mathbf{x} \in \mathcal{X}_\alpha$, then $\mathbf{y} \in \mathcal{X}_{\alpha+\gamma}$.

Lipschitz continuity follows from the fact that mixed extensions are multilinear functions (8), with bounded domains. The proof of the second part immediately follows from the Lipschitz continuity of mixed extensions of payoff functions and the definition of approximate equilibria (9). Note that the second part implies that for any $\epsilon' > 0$, there exists a small enough neighborhood of equilibria that is contained in the ϵ' -equilibrium set of the game.

We next study the continuity properties of the approximate equilibrium mapping. We first provide the relevant definitions (see Berge (1963), Fudenberg and Tirole (1991)).

Definition 6.1 (Upper Semicontinuous Function). *A function $g : X \rightarrow Y \subset \mathbb{R}$ is upper semicontinuous at x_* , if, for each $\epsilon > 0$ there exists a neighborhood U of x_* , such that $g(x) < g(x_*) + \epsilon$ for all $x \in U$. We say g is upper semicontinuous, if it is upper semicontinuous at every point in its domain.*

Alternatively, g is upper semicontinuous if $\limsup_{x_n \rightarrow x_} g(x_n) \leq g(x_*)$ for every x_* in its domain.*

Definition 6.2 (Upper Semicontinuous Correspondence). *A correspondence $g : X \rightrightarrows Y$ is upper semicontinuous at x_* , if for any open neighborhood V of $g(x_*)$ there exists a neighborhood U of x_* such that $g(x) \subset V$ for all $x \in U$. We say g is upper semicontinuous, if it is upper semicontinuous at every point in its domain and $g(x)$ is a compact set for each $x \in X$.*

Alternatively, when Y is compact, g is upper semicontinuous if its graph is closed, i.e., the set $\{(x, y) | x \in X, y \in g(x)\}$ is closed.

We next establish upper semicontinuity of the approximate equilibrium mapping.⁹

Lemma 6.2. *(i) Let $\nu : \prod_{m \in \mathcal{M}} \Delta E^m \rightarrow \mathbb{R}$ be an upper semicontinuous function. The correspondence $g : \mathbb{R} \rightrightarrows \prod_{m \in \mathcal{M}} \Delta E^m$ such that $g(v) = \{\mathbf{x} | \nu(\mathbf{x}) \geq -v\}$ is upper semicontinuous.*

(ii) Let $g : \mathbb{R} \rightrightarrows \prod_{m \in \mathcal{M}} \Delta E^m$ be the correspondence such that $g(\alpha) = \mathcal{X}_\alpha$. This correspondence is upper semicontinuous.

Upper semicontinuity of the approximate equilibrium mapping implies that for any given neighborhood of the ϵ -equilibrium set, there exists an $\epsilon' > \epsilon$ such that ϵ' -equilibrium set is contained in this neighborhood. In particular, this implies that every neighborhood of equilibria of the game contains an ϵ' -equilibrium set for some $\epsilon' > 0$. Hence, if disjoint neighborhoods of equilibria are chosen (assuming there are finitely many equilibria), this implies that there exists some $\epsilon' > 0$, such that the ϵ' -equilibrium set is contained in disjoint neighborhoods of equilibria. In the next section, we use this observation to establish convergence of fictitious play to small neighborhoods of equilibria of near-potential games.

6.2. Discrete-Time Fictitious Play

Fictitious play is a classical update rule studied in the learning in games literature. In this section, we consider the fictitious play dynamics, proposed in Brown (1951), and explain how the limiting behavior of this dynamical process can be characterized in near-potential games. In particular, we show that the empirical frequencies of fictitious play converge to a set which can be characterized in terms of the ϵ -equilibrium set of the game, and the level sets of the potential function of a close potential game. We also establish that for games sufficiently close to a potential game, the empirical frequencies of fictitious play converge to a neighborhood of the (mixed) equilibrium set. Moreover, the size of this neighborhood depends on the distance of the original game from a nearby potential game. This generalizes

⁹Here we fix the game, and discuss upper semicontinuity with respect to the ϵ parameter characterizing the ϵ -equilibrium set. We note that this is different than the common results in the literature which discuss upper semicontinuity of the equilibrium set with respect to changes in the utility functions of the underlying game (see Fudenberg and Tirole (1991)).

the result of Monderer and Shapley (1996a), on convergence of empirical frequencies to mixed Nash equilibria in potential games.

In this paper, we only consider the discrete-time version of fictitious play, i.e., the update process starts at a given strategy profile at time $t = 0$, and players can update their strategies at discrete time instants $t \in \{1, 2, \dots\}$. Throughout this subsection we denote the strategy used by player m at time instant t by p_t^m , and we denote by $\mathbf{1}(p_t^m = p^m)$ the indicator function which equals to 1 if $p_t^m = p^m$, and 0 otherwise. A formal definition of discrete-time fictitious play dynamics is given next.

Definition 6.3 (Discrete-Time Fictitious Play). *Let $\mu_T^m(q^m) = \frac{1}{T} \sum_{t=0}^{T-1} \mathbf{1}(p_t^m = q^m)$ denote the empirical frequency that player m uses strategy q^m from time instant 0 to time instant $T - 1$, and μ_T^{-m} denote the collection of empirical frequencies of all players but m . A game play, where at each time instant t , every player m , chooses a strategy p_t^m such that*

$$p_t^m \in \arg \max_{q^m \in E^m} u^m(q^m, \mu_t^{-m})$$

is referred to as discrete-time fictitious play. That is, fictitious play dynamics is the update process, where each player chooses its best response to the empirical frequencies of the actions of other players.

We refer to μ_t^m as the distribution of empirical frequencies of player m 's strategies at time t . Note that μ_t^m can be thought of as vector with length $|E^m|$, whose entries are indexed by strategies of player m , i.e., $\mu_t^m(p^m)$ denotes the entry of the vector corresponding to the empirical frequency player m uses strategy p^m with. Similarly, we define the joint empirical frequency distribution of all players as $\mu_t = \{\mu_t^m\}_{m \in \mathcal{M}}$. Note that $\mu_t^m \in \Delta E^m$, i.e., empirical frequency distributions are mixed strategies, and similarly $\mu_t \in \prod_{m \in \mathcal{M}} \Delta E^m$.

Observe that the evolution of this empirical frequency distribution can be captured by the following equation:

$$\mu_{t+1} = \frac{t}{t+1} \mu_t + \frac{1}{t+1} I_t, \quad (10)$$

where $I_t = \{I_t^m\}_{m \in \mathcal{M}}$, and I_t^m is a vector which has the same size as μ_t^m and its entry corresponding to strategy p^m is given by $I_t^m(p^m) = \mathbf{1}(p_t^m = p^m)$. Rearranging the terms in (10), and observing that $I_t, \mu_t \in \prod_{m \in \mathcal{M}} \Delta E^m$ are vectors with entries in $[0, 1]$ we conclude

$$\|\mu_{t+1} - \mu_t\| = \frac{1}{t+1} \|I_t - \mu_t\| = O\left(\frac{1}{t}\right), \quad (11)$$

where $O(\cdot)$ stands for the big-O notation, i.e., $f(x) = O(g(x))$, implies that there exists some x_0 and a constant c such that $|f(x)| \leq c|g(x)|$ for all $x \geq x_0$.

We start analyzing discrete-time fictitious play in near-potential games, by first focusing on the change in the value of the potential function along the fictitious play updates in the original game. In particular, we show that in near-potential games if the empirical frequencies are outside some ϵ -equilibrium set, then the potential of the close potential game (evaluated at the empirical frequency distribution) increases by discrete-time fictitious play updates.¹⁰

¹⁰Our approach here is similar to the one used in Monderer and Shapley (1996a) to analyze discrete-time fictitious play in potential games.

Lemma 6.3. Consider a game \mathcal{G} and let $\hat{\mathcal{G}}$ be a close potential game such that $d(\mathcal{G}, \hat{\mathcal{G}}) \leq \delta$. Denote the potential function of $\hat{\mathcal{G}}$ by ϕ . Assume that in \mathcal{G} players update their strategies according to discrete-time fictitious play dynamics, and at some time instant $T > 0$, the empirical frequency distribution μ_T is outside an ϵ -equilibrium set of \mathcal{G} . Then,

$$\phi(\mu_{T+1}) - \phi(\mu_T) \geq \frac{\epsilon - M\delta}{T+1} + O\left(\frac{1}{T^2}\right).$$

Proof. Consider the mixed extension of the potential function $\phi(\mathbf{x}) = \sum_{\mathbf{p} \in E} \phi(\mathbf{p}) \prod_{m \in \mathcal{M}} x^m(p^m)$, where $\mathbf{x} = \{x^m\}_m$ and $x^m(p^m)$ denotes the probability player m plays strategy p^m . The expression for $\phi(\mathbf{x})$ implies that Taylor expansion of ϕ around μ_T satisfies

$$\phi(\mu_{T+1}) = \phi(\mu_T) + \sum_{m \in \mathcal{M}} \sum_{p^m \in E^m} (\mu_{T+1}^m(p^m) - \mu_T^m(p^m)) \phi(p^m, \mu_T^{-m}) + O(\|\mu_{T+1} - \mu_T\|^2).$$

Observing from (10) that $\mu_{t+1} - \mu_t = \frac{1}{t+1}(I_t - \mu_t)$, and noting from (11) that $\|\mu_{t+1} - \mu_t\| = O(\frac{1}{t})$ the above equality can be rewritten as

$$\phi(\mu_{T+1}) = \phi(\mu_T) + \sum_{m \in \mathcal{M}} \sum_{p^m \in E^m} \frac{1}{T+1} (\mathbf{1}(p_T^m = p^m) - \mu_T^m(p^m)) \phi(p^m, \mu_T^{-m}) + O\left(\frac{1}{T^2}\right).$$

Rearranging the terms, and noting that $\sum_{p^m \in E^m} \mu_T^m(p^m) \phi(p^m, \mu_T^{-m}) = \phi(\mu_T^m, \mu_T^{-m})$, it follows that

$$\begin{aligned} \phi(\mu_{T+1}) &= \phi(\mu_T) + \sum_{m \in \mathcal{M}} \frac{1}{T+1} \phi(p_T^m, \mu_T^{-m}) - \sum_{m \in \mathcal{M}} \frac{1}{T+1} \phi(\mu_T^m, \mu_T^{-m}) + O\left(\frac{1}{T^2}\right) \\ &= \phi(\mu_T) + \frac{1}{T+1} \sum_{m \in \mathcal{M}} (\phi(p_T^m, \mu_T^{-m}) - \phi(\mu_T^m, \mu_T^{-m})) + O\left(\frac{1}{T^2}\right). \end{aligned}$$

Since $d(\mathcal{G}, \hat{\mathcal{G}}) \leq \delta$, the above equality and the definition of MPD imply

$$\phi(\mu_{T+1}) \geq \phi(\mu_T) + \frac{1}{T+1} \sum_{m \in \mathcal{M}} (u^m(p_T^m, \mu_T^{-m}) - u^m(\mu_T^m, \mu_T^{-m}) - \delta) + O\left(\frac{1}{T^2}\right). \quad (12)$$

By definition of the fictitious play dynamics, every player m plays its best response to μ_T^{-m} , therefore $u^m(p_T^m, \mu_T^{-m}) - u^m(\mu_T^m, \mu_T^{-m}) \geq 0$ for all m . Additionally, if μ_T is outside the ϵ -equilibrium set, as in the statement of the lemma, then it follows that $u^m(p_T^m, \mu_T^{-m}) - u^m(\mu_T^m, \mu_T^{-m}) \geq \epsilon$ for at least one player. Therefore, (12) implies

$$\phi(\mu_{T+1}) \geq \phi(\mu_T) + \frac{\epsilon - M\delta}{T+1} + O\left(\frac{1}{T^2}\right),$$

hence, the claim follows. \square

The above theorem implies that if μ_T is not in the ϵ -equilibrium set for some $\epsilon > M\delta$, and T sufficiently large, then the potential evaluated at empirical frequencies increases when

players update their strategies. Since the mixed extension of the potential is a bounded function, the potential cannot increase unboundedly, and this observation suggests that the ϵ -equilibrium set is eventually reached by the empirical frequency distribution. On the other hand, at a later time instant μ_T can still leave this equilibrium set, and before it does so the potential cannot be lower than the lowest potential in this set (since μ_T itself belongs to this set). Moreover, after μ_T leaves the ϵ -equilibrium set the potential keeps increasing. Thus, the empirical frequencies are contained in the set of mixed strategy profiles, which have potential at least as large as the minimum potential in this approximate equilibrium set. We next make this intuition precise, and characterize the set of limiting mixed strategies for fictitious play in near-potential games. We adopt the following convergence notion: we say that empirical frequencies of fictitious play converge to a set $S \subset \prod_{m \in \mathcal{M}} \Delta E^m$, if $\inf_{\mathbf{x} \in S} \|\mu_t - \mathbf{x}\| \rightarrow 0$ as $t \rightarrow \infty$.

Theorem 6.1. *Consider a game \mathcal{G} and let $\hat{\mathcal{G}}$ be a close potential game such that $d(\mathcal{G}, \hat{\mathcal{G}}) \leq \delta$. Denote the potential function of $\hat{\mathcal{G}}$ by ϕ . Assume that in \mathcal{G} players update their strategies according to discrete-time fictitious play dynamics, and let \mathcal{X}_α denote the α -equilibrium set of \mathcal{G} . For any $\epsilon > 0$, there exists a time instant $T_\epsilon > 0$ such that for all $t > T_\epsilon$*

$$\mu_t \in C_\epsilon \triangleq \left\{ \mathbf{x} \in \prod_{m \in \mathcal{M}} \Delta E^m \mid \phi(\mathbf{x}) \geq \min_{\mathbf{y} \in \mathcal{X}_{M\delta+\epsilon}} \phi(\mathbf{y}) \right\}.$$

Proof. Let ϵ' be such that $\epsilon > \epsilon' > 0$. It can be seen from the definition of C_ϵ that $\mathcal{X}_{M\delta+\epsilon'} \subset \mathcal{X}_{M\delta+\epsilon} \subset C_\epsilon$. We prove the claim in two steps: (i) We first show that in this update process $\mathcal{X}_{M\delta+\epsilon'}$ is visited infinitely often by μ_t , i.e., for all T' , there exists $t > T'$ such that $\mu_t \in \mathcal{X}_{M\delta+\epsilon'}$, (ii) We prove that there exists a T'' such that if $\mu_t \in C_\epsilon$ for some $t > T''$, then for all $t' > t$ we have $\mu_{t'} \in C_\epsilon$. Thus, the second step guarantees that if C_ϵ is visited at a sufficiently later time instant, then μ_t remains in C_ϵ . Since $\mathcal{X}_{M\delta+\epsilon'} \subset C_\epsilon$ the first step ensures that such a time instant exists, and the claim in the theorem immediately follows from (ii). Moreover, this time instant corresponds to T_ϵ in the theorem statement.

Proof of both steps rely on the following simple observation: Lemma 6.3 implies that there exists a large enough T , such that if the empirical frequencies do not belong to $\mathcal{X}_{M\delta+\epsilon'}$ at a time instant $t > T$, then ϕ increases:

$$\phi(\mu_{t+1}) - \phi(\mu_t) \geq \frac{M\delta + \epsilon' - M\delta}{(t+1)} + O\left(\frac{1}{t^2}\right) > \frac{\epsilon'}{2(t+1)} > 0. \quad (13)$$

We prove (i) by contradiction. Assume that there exists a T' such that $\mu_t \notin \mathcal{X}_{M\delta+\epsilon'}$ for $t > T'$, and let $T_m = \max\{T, T'\}$. Then, (13) holds for all $t = \{T_m + 1, \dots\}$, and summing both sides of this inequality over this set we obtain

$$\limsup_{t \rightarrow \infty} \phi(\mu_{t+1}) - \phi(\mu_{T_m+1}) \geq \sum_{t=T_m+1}^{\infty} \frac{\epsilon'}{2(t+1)}.$$

Since the mixed extension of the potential is a bounded function, it follows that the left hand side of the above inequality is bounded, but the right hand side grows unboundedly. Hence, we reach a contradiction, and (i) follows.

Lemma 6.1 (ii) implies that there exists some $\theta > 0$ such that if a strategy profile \mathbf{x} is an $(M\delta + \epsilon')$ -equilibrium, then any strategy profile \mathbf{y} that satisfies $\|\mathbf{x} - \mathbf{y}\| < \theta$ is an $(M\delta + \epsilon)$ -equilibrium (recall that $\epsilon > \epsilon' > 0$). Since $\|\mu_{t+1} - \mu_t\| = O(1/t)$ by (11), this implies that there exists some $T'' > T$, such that for all $t > T''$ if $\mu_t \in \mathcal{X}_{M\delta + \epsilon'}$, then we have

$$\mu_{t+1} \in \mathcal{X}_{M\delta + \epsilon}. \quad (14)$$

Let $\mu_t \in C_\epsilon$ for some time instant $t > T''$. If $\mu_t \in \mathcal{X}_{M\delta + \epsilon'}$, then by (14) $\mu_{t+1} \in \mathcal{X}_{M\delta + \epsilon} \subset C_\epsilon$. If, on the other hand, $\mu_t \in C_\epsilon - \mathcal{X}_{M\delta + \epsilon'}$, then by (13) and the definition of C_ϵ we have

$$\phi(\mu_{t+1}) > \phi(\mu_t) \geq \min_{\mathbf{y} \in \mathcal{X}_{M\delta + \epsilon}} \phi(\mathbf{y}), \quad (15)$$

and hence $\mu_{t+1} \in C_\epsilon$. Thus, we have established that there exists some T'' such that if $\mu_t \in C_\epsilon$ for some $t > T''$, then $\mu_{t+1} \in C_\epsilon$, and hence (ii) follows. \square

The above theorem establishes that after finite time μ_t is contained in the set C_ϵ for any $\epsilon > 0$. Corollary 6.1, establishes that in the limit this result can be strengthened: as $t \rightarrow \infty$, μ_t converges to a set, which is a subset of C_ϵ for every $\epsilon > 0$. The proof can be found in the Appendix.

Corollary 6.1. *The empirical frequencies of discrete-time fictitious play converge to*

$$C \triangleq \left\{ \mathbf{x} \in \prod_{m \in \mathcal{M}} \Delta E^m \mid \phi(\mathbf{x}) \geq \min_{\mathbf{y} \in \mathcal{X}_{M\delta}} \phi(\mathbf{y}) \right\}.$$

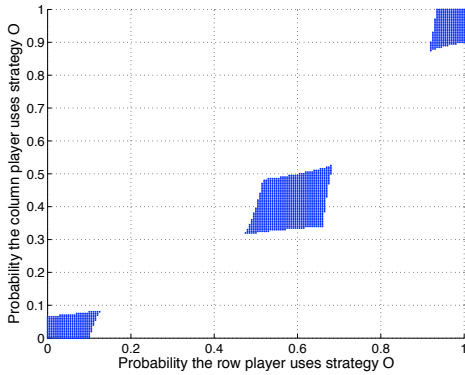
This result suggests that in near-potential games, the empirical frequencies of fictitious play converge to a set where the potential is at least as large as the minimum potential in an approximate equilibrium set. For exact potential games, it is known that the empirical frequencies converge to a Nash equilibrium (Monderer and Shapley 1996a). It can be seen from Definition 2.1 that in potential games, maximizers of the potential function are equilibria of the game. Thus, in potential games with a unique equilibrium the equilibrium is the unique maximizer of the potential function. Hence, for such games, we have $\delta = 0$, $\min_{\mathbf{y} \in \mathcal{X}_{M\delta}} \phi(\mathbf{y}) = \max_{\mathbf{x} \in \prod_{m \in \mathcal{M}} \Delta E^m} \phi(\mathbf{x})$, and Corollary 6.1 implies that empirical frequencies of fictitious play converge to the unique equilibrium of the game, recovering the convergence result of Monderer and Shapley (1996a). However, when there are multiple equilibria Corollary 6.1 suggests that empirical frequencies converge to the set of mixed strategy profiles that have potential weakly larger than the minimum potential attained by the equilibria. While this set contains equilibria, it may contain a continuum of other mixed strategy profiles. This suggests that in games with multiple equilibria our result may provide a loose characterization of the limiting behavior of fictitious play dynamics.

We next show that by exploiting the properties of mixed approximate equilibrium sets, it is possible to obtain a stronger result. Before we present our result, we discuss a feature of mixed equilibrium sets which will be key in our analysis: For small ϵ , the ϵ -equilibrium set is contained in a small neighborhood of equilibria (this statement follows from Lemma 6.2 (ii) by considering the upper semicontinuity of the approximate equilibrium correspondence $g(\alpha)$ at $\alpha = 0$). This property is illustrated in Example 6.1.

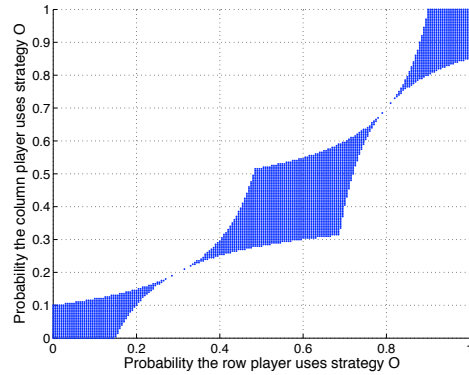
	O	F
O	3, 2	0, 0
F	0, 0	2, 3

Table 3: Payoffs in BoS.

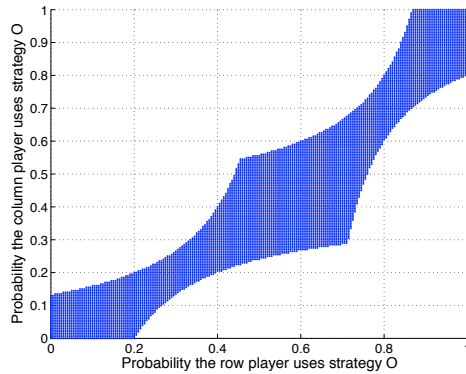
Example 6.1 (Mixed equilibrium set of Battle of the Sexes:). *Consider the two-player battle of the sexes (BoS) game: Each player has two possible actions $\{O, F\}$, and the payoffs of players are as given in Table 3. This game has three equilibria: (i) both players use O, (ii) Both players use F, (iii) Row player uses O with probability 0.6, and column player uses O with probability 0.4. Note that since this is a game where each player has only two strategies, the probability of using strategy O, in the third case uniquely identifies the corresponding mixed equilibrium. For different values of ϵ , the set of ϵ -equilibria of this game is shown in Figure 4. It follows that the set of ϵ -equilibria is contained in disjoint neighborhoods of equilibria for small values of ϵ .*



(a) 0.2-equilibrium set.



(b) 0.3-equilibrium set.



(c) 0.4-equilibrium set.

Figure 4: Approximate equilibrium sets in BoS are contained in disjoint neighborhoods of equilibria for small ϵ .

It was established in Lemma 6.3 that the potential function of a nearby potential game (with MPD δ to the original game), evaluated at the empirical frequency distribution, increases when this distribution is outside the $M\delta$ -equilibrium set of the original game (where M is the number of players). If δ is sufficiently small, then the $M\delta$ -equilibria of the game will be contained in a small neighborhood of the equilibria, as illustrated above and shown in Lemma 6.2 (ii). Thus, for sufficiently small δ , it is possible to establish that the potential of a close potential game increases outside a small neighborhood of the equilibria of the game. In Theorem 6.2, we use this observation to show that for sufficiently small δ the empirical frequencies of fictitious play dynamics converge to a neighborhood of an equilibrium. We state the theorem under the assumption that the original game has finitely many equilibria. This assumption generically holds, i.e., for any game a (nondegenerate) random perturbation of payoffs will lead to such a game with probability one (see Fudenberg and Tirole (1991)). When stating our result, we make use of the Lipschitz continuity of the mixed extension of the potential function, as established in Lemma 6.1.

Theorem 6.2. *Consider a game \mathcal{G} and let $\hat{\mathcal{G}}$ be a close potential game such that $d(\mathcal{G}, \hat{\mathcal{G}}) \leq \delta$. Denote the potential function of $\hat{\mathcal{G}}$ by ϕ , and the Lipschitz constant of the mixed extension of ϕ by L . Assume that \mathcal{G} has finitely many equilibria, and in \mathcal{G} players update their strategies according to discrete-time fictitious play dynamics.*

There exists some $\bar{\delta} > 0$, and $\bar{\epsilon} > 0$ (which are functions of utilities of \mathcal{G} but not δ) such that if $\delta < \bar{\delta}$, then the empirical frequencies of fictitious play converge to

$$\left\{ \mathbf{x} \mid \|\mathbf{x} - \mathbf{x}_k\| \leq \frac{4f(M\delta)ML}{\epsilon} + f(M\delta + \epsilon), \text{ for some equilibrium } \mathbf{x}_k \right\}, \quad (16)$$

for any ϵ such that $\bar{\epsilon} \geq \epsilon > 0$, where $f : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ is an upper semicontinuous function satisfying $f(x) \rightarrow 0$ as $x \rightarrow 0$.

The proof of this theorem can be found in the Appendix, and it has three main steps illustrated in Figures 5 and 6. As explained earlier, for small δ and ϵ , the $M\delta + \epsilon$ -equilibrium set of the game is contained in disjoint neighborhoods of the equilibria of the game. Lemma 6.3 implies that potential evaluated at μ_t increases outside this approximate equilibrium set with strategy updates. In the proof, we first quantify the increase in the potential, when μ_t leaves this approximate equilibrium set and returns back to it at a later time instant (see Figure 5a). Then, using this increase condition we show that for sufficiently large t , μ_t can visit the approximate equilibrium set infinitely often only around one equilibrium, say $\mathbf{x}_{k'}$ (see Figure 5b). This holds since, the increase condition guarantees that the potential increases significantly when μ_t leaves the neighborhood of an equilibrium \mathbf{x}_k , and reaches to that of $\mathbf{x}_{k'}$. Finally, using the increase condition one more time, we establish that if after time T , μ_t visits the approximate equilibrium set only in the neighborhood of $\mathbf{x}_{k'}$, we can construct a neighborhood of $\mathbf{x}_{k'}$, which contains μ_t for all $t > T$ (see Figure 6). This neighborhood is expressed in (16).

Observe that if $\delta = 0$, i.e., the original game is a potential game, then $f(M\delta) = 0$, and Theorem 6.2 implies that empirical frequencies of fictitious play converge to the $f(\epsilon)$ -neighborhood of equilibria for any ϵ such that $\bar{\epsilon} \geq \epsilon > 0$. Thus, choosing ϵ arbitrarily small, and observing that $\lim_{x \rightarrow 0} f(x) = 0$, our result implies that in potential games, empirical

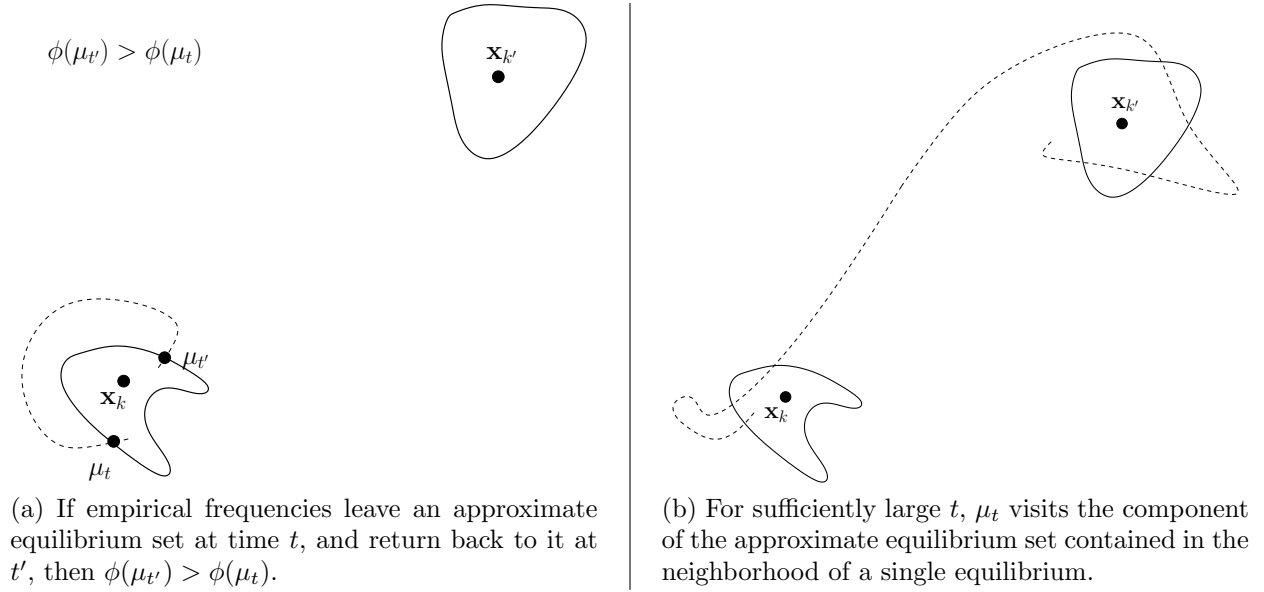


Figure 5: For small δ and ϵ , $M\delta + \epsilon$ -equilibrium set (enclosed by solid lines around equilibria $\mathbf{x}_{k'}$ and \mathbf{x}_k) is contained in disjoint neighborhoods of equilibria. If the empirical frequency distribution, μ_t , is outside this approximate equilibrium set, then the potential increases with each strategy update. Assume that empirical frequency distribution leaves an approximate equilibrium set (at time t) and returns back to it at a later time instant ($t' > t$). We first quantify the resulting increase in the potential (left). If μ_t travels from the component of the approximate equilibrium set in the neighborhood of equilibrium \mathbf{x}_k to that in the neighborhood of equilibrium $\mathbf{x}_{k'}$, then the increase in the potential is significant, and consequently μ_t cannot visit the approximate equilibrium set in the neighborhood of equilibrium \mathbf{x}_k at a later time instant (right).

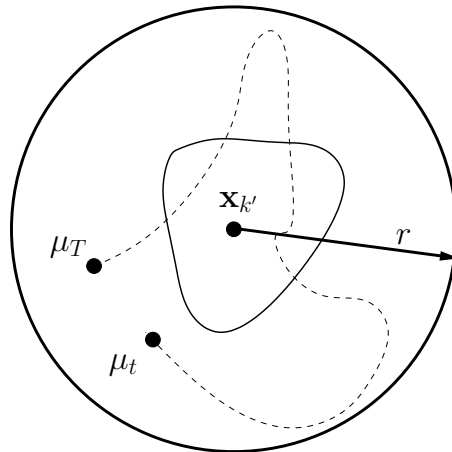


Figure 6: If after time T , μ_t only visits the approximate equilibrium set in the neighborhood of a single equilibrium $\mathbf{x}_{k'}$, then we can establish that μ_t never leaves a neighborhood of this equilibrium for $t > T$. The size of this neighborhood is denoted by r in the figure and can be expressed as in Theorem 6.2.

frequencies converge to the set of Nash equilibria. Hence, as a special case of Theorem 6.2, we obtain the convergence result of Monderer and Shapley (1996a).

Assume that $\delta \neq 0$ and a small $\epsilon < \bar{\epsilon}$ is given. If δ is sufficiently small then $f(M\delta)/\epsilon \approx 0$, since $\lim_{x \rightarrow 0} f(x) = 0$. Consequently, $\frac{4f(M\delta)ML}{\epsilon} + f(M\delta + \epsilon)$ is small, and Theorem 6.2 establishes convergence of empirical frequencies to a small neighborhood of equilibria. Thus, we conclude that for games that are close to potential games, i.e., for $\delta \ll 1$, Theorem 6.2 establishes convergence of empirical frequencies to a small neighborhood of equilibria.

7. Conclusions

In this paper, we present a framework for studying the limiting behavior of adaptive learning dynamics in finite strategic form games by exploiting their relation to nearby potential games. We restrict our attention to better/best response, logit response and fictitious play dynamics. We show that for near-potential games trajectories of better/best response dynamics converge to ϵ -equilibrium sets, where ϵ depends on closeness to a potential game. We study the stochastically stable strategy profiles of logit response dynamics and prove that they are contained in the set of strategy profiles that approximately maximize the potential function of a nearby potential game. In the case of fictitious play we focus on the empirical frequencies of players' actions, and establish that they converge to a small neighborhood of equilibria in near-potential games. Our results suggest that games that are close to a potential game inherit the dynamical properties (such as convergence to approximate equilibrium sets) of potential games. Additionally, since a close potential game to a given game can be found by solving a convex optimization problem, as discussed in Section 3, this enables us to study dynamical properties of strategic form games by first identifying a nearby potential game to this game, and then studying the dynamical properties of the nearby potential game.

The framework presented in this paper opens up a number of interesting research directions. Among them, we mention the following:

Heterogeneous update rules: In this paper we only analyzed the update rules in which players update their strategies using the same mechanism. For instance, we assumed that all players adopt best response, or logit response dynamics with the same parameter. The limiting behavior of dynamic processes, where players adhere to different update rules is still an open question, even for potential games. An interesting future research question is whether the techniques in this paper can be used to understand the limiting behavior of such update rules. For example, consider a potential game where all players update their strategies using logit response with different but “close” τ parameters. Can the outcome of this dynamic process be approximated with the outcome of logit response in a close potential game where all players use the same parameter for their updates?

Guaranteeing desirable limiting behavior: Another promising research direction is to use our understanding of simple update rules, such as better/best response and logit response dynamics to design mechanisms that guarantee desirable limiting behavior, such as low efficiency loss and “fair” outcomes. It is well known that equilibria in games can be very different in terms of such properties (Roughgarden 2005). Hence, it is of interest to develop

update rules that converge to a particular equilibrium, thus providing equilibrium refinement in the limit, or to find mechanisms that modify the underlying game in a way that can induce desirable limiting behavior. It has been shown in some cases that simple pricing mechanisms can ensure convergence to desirable equilibria in near-potential games (Candogan et al. 2010a). It is an interesting research direction to extend such mechanisms to general games.

Dynamics in “near” zero-sum and supermodular games: Dynamical properties of simple update rules in zero-sum games and supermodular games are also well understood (Shamma and Arslan 2004, Milgrom and Roberts 1990). If a game is close to a zero-sum game or a supermodular game, does it still inherit some of the dynamical properties of the original game? If such “continuity” properties do not hold, then the results on dynamical properties of these classes of games may be fragile. Hence, it would be interesting to investigate whether analogous results to the ones in this paper can be established for these classes of games.

References

- Alós-Ferrer, C., Netzer, N., 2010. The logit-response dynamics. *Games and Economic Behavior* 68 (2), 413–427.
- Anantharam, V., Tsoucas, P., 1989. A proof of the markov chain tree theorem. *Statistics & Probability Letters* 8 (2), 189–192.
- Arslan, G., Marden, J., Shamma, J., 2007. Autonomous vehicle-target assignment: A game-theoretical formulation. *Journal of Dynamic Systems, Measurement, and Control* 129 (5), 584–596.
URL <http://link.aip.org/link/?JDS/129/584/1>
- Berge, C., 1963. *Topological spaces*. Oliver & Boyd.
- Blume, L., 1993. The statistical mechanics of strategic interaction. *Games and Economic Behavior* 5 (3), 387 – 424.
URL <http://www.sciencedirect.com/science/article/pii/S0899825683710237>
- Blume, L., 1997. Population games. In: Arthur, W., Durlauf, S., Lane, D. (Eds.), *The economy as an evolving complex system II*. Addison-Wesley, pp. 425–460.
- Brown, G., 1951. Iterative solution of games by fictitious play. *Activity analysis of production and allocation* 13 (1), 374–376.
- Candogan, O., Menache, I., Ozdaglar, A., Parrilo, P., March 2010a. Near-optimal power control in wireless networks: A potential game approach. In: *INFOCOM 2010 Proceedings IEEE*. pp. 1–9.
- Candogan, O., Menache, I., Ozdaglar, A., Parrilo, P. A., 2011. Flows and decompositions of games: Harmonic and potential games. Forthcoming in *Mathematics of Operations Research*.
URL <http://mor.journal.informs.org/cgi/content/abstract/moor.1110.0500v1>

- Candogan, O., Ozdaglar, A., Parrilo, P., December 2010b. A projection framework for near-potential games. In: 49th IEEE Conference on Decision and Control (CDC). pp. 244–249.
- Cho, G., Meyer, C., 2001. Comparison of perturbation bounds for the stationary distribution of a Markov chain. *Linear Algebra and its Applications* 335 (1-3), 137–150.
- Freidlin, M., Wentzell, A., 1998. *Random perturbations of dynamical systems*. Springer Verlag.
- Fudenberg, D., Levine, D., 1998. *The theory of learning in games*. MIT press.
- Fudenberg, D., Tirole, J., 1991. *Game Theory*. MIT Press.
- Haviv, M., Van der Heyden, L., 1984. Perturbation bounds for the stationary probabilities of a finite Markov chain. *Advances in Applied Probability* 16 (4), 804–818.
- Hofbauer, J., Sandholm, W., 2002. On the global convergence of stochastic fictitious play. *Econometrica*, 2265–2294.
- Jordan, J., 1993. Three problems in learning mixed-strategy Nash equilibria. *Games and Economic Behavior* 5 (3), 368–386.
- Leighton, F., Rivest, R., 1983. The markov chain tree theorem. Tech. Rep. MIT/LCS/TM-249, MIT, Laboratory for Computer Science.
- Marden, J., Arslan, G., Shamma, J., 2009a. Cooperative Control and Potential Games. *IEEE Transactions on Systems, Man, and Cybernetics, Part B: Cybernetics* 39 (6), 1393–1407.
- Marden, J., Arslan, G., Shamma, J., 2009b. Joint strategy fictitious play with inertia for potential games. *IEEE Transactions on Automatic Control* 54 (2), 208–220.
- Marden, J., Shamma, J., 2008. Revisiting log-linear learning: Asynchrony, completeness and a payoff-based implementation. Under Submission.
- Milgrom, P., Roberts, J., 1990. Rationalizability, learning, and equilibrium in games with strategic complementarities. *Econometrica: Journal of the Econometric Society* 58 (6), 1255–1277.
- Monderer, D., Shapley, L., 1996a. Fictitious play property for games with identical interests. *Journal of Economic Theory* 68 (1), 258–265.
- Monderer, D., Shapley, L., 1996b. Potential games. *Games and economic behavior* 14 (1), 124–143.
- Roughgarden, T., 2005. *Selfish Routing and the Price of Anarchy*. MIT Press.
- Sandholm, W., 2010. *Population games and evolutionary dynamics*. MIT Press, Cambridge, MA.
- Shamma, J., Arslan, G., 2004. Unified convergence proofs of continuous-time fictitious play. *IEEE Transactions on Automatic Control* 49 (7), 1137–1141.

Shapley, L., 1964. Some topics in two-person games. *Advances in game theory* 52, 1–29.

Young, H., 1993. The evolution of conventions. *Econometrica: Journal of the Econometric Society* 61 (1), 57–84.

Young, H., 2004. *Strategic learning and its limits*. Oxford University Press, USA.

Appendix A. Proofs of Section 6

Proof of Lemma 6.1: (i) The mixed extension of ν can be given as in (8):

$$\nu(\mathbf{x}) = \sum_{\mathbf{p} \in E} \nu(\mathbf{p}) \prod_{k \in \mathcal{M}} x^k(p^k).$$

Hence $\nu(\mathbf{x})$ is equal to the sum of Lipschitz continuous functions $\{\nu(\mathbf{p}) \prod_{k \in \mathcal{M}} x^k(p^k)\}_{\mathbf{p} \in E}$. The claim follows since, as a function of $\mathbf{x} \in \prod_{m \in \mathcal{M}} \Delta E^m$, the Lipschitz constant of $\nu(\mathbf{p}) \prod_{k \in \mathcal{M}} x^k(p^k)$ is bounded by $\nu(\mathbf{p})\sqrt{M} \leq \nu(\mathbf{p})M$.

(ii) Let $\nu : \prod_{m \in \mathcal{M}} \Delta E^m \rightarrow \mathbb{R}$ be a function such that

$$\nu(\mathbf{x}) = - \max_{m \in \mathcal{M}, p^m \in E^m} (u^m(p^m, \mathbf{x}^{-m}) - u^m(x^m, \mathbf{x}^{-m})). \quad (\text{A.1})$$

It follows from the definition of ϵ -equilibrium that a strategy profile \mathbf{x} is an ϵ -equilibrium if and only if $\nu(\mathbf{x}) \geq -\epsilon$.

By (i), it follows that mixed extensions of utility functions are Lipschitz continuous. Thus, the difference $u^m(p^m, \mathbf{x}^{-m}) - u^m(x^m, \mathbf{x}^{-m})$ is Lipschitz continuous in \mathbf{x} . Since ν is obtained from maximum of finitely many such functions, we conclude that it is Lipschitz continuous with some constant L . It follows that if $\|\mathbf{x} - \mathbf{y}\| < \theta$, then $|\nu(\mathbf{x}) - \nu(\mathbf{y})| < L\theta$. Thus, choosing $\theta < \epsilon/L$, and recalling that \mathbf{x} is an ϵ -equilibrium if and only if $\nu(\mathbf{x}) \geq -\epsilon$, the claim follows. \square

Proof of Lemma 6.2: (i) Consider the graph of g , i.e., $S = \{(v, \mathbf{x}) | v \in \mathbb{R}, \mathbf{x} \in g(v)\}$. The definition of g suggests that S can alternatively be written as

$$S = \left\{ (v, \mathbf{x}) \in \mathbb{R} \times \prod_{m \in \mathcal{M}} \Delta E^m \mid \nu(\mathbf{x}) \geq -v \right\} \quad (\text{A.2})$$

Since ν is upper semicontinuous, the function $h : \mathbb{R} \times \prod_{m \in \mathcal{M}} \Delta E^m$, such that $h(v, \mathbf{x}) = \nu(\mathbf{x}) + v$, is also an upper semicontinuous function. Since upper level sets of upper semicontinuous functions are closed (see Berge (1963)), the set $\{(v, \mathbf{x}) \in \mathbb{R} \times \prod_{m \in \mathcal{M}} \Delta E^m | h(v, \mathbf{x}) \geq 0\}$, or equivalently S is a closed subset of $\mathbb{R} \times \prod_{m \in \mathcal{M}} \Delta E^m$. Thus, the graph of g is closed, and since $\prod_{m \in \mathcal{M}} \Delta E^m$ is a compact set the claim follows from Definition 6.2.

(ii) Let $\nu : \prod_{m \in \mathcal{M}} \Delta E^m \rightarrow \mathbb{R}$ be as in (A.1). i.e.,

$$\nu(\mathbf{x}) = - \max_{m \in \mathcal{M}, p^m \in E^m} (u^m(p^m, \mathbf{x}^{-m}) - u^m(x^m, \mathbf{x}^{-m})).$$

As explained in the proof of Lemma 6.1, \mathbf{x} is an ϵ -equilibrium if and only if $\nu(\mathbf{x}) \geq -\epsilon$. The claim follows from part (i) by noting that ν is a continuous function, and $g(\alpha) = \mathcal{X}_\alpha = \{\mathbf{x} | \nu(\mathbf{x}) \geq -\alpha\}$. \square

Proof of Corollary 6.1: Let $\epsilon_n = M\delta + \frac{1}{n}$ for $n \in \mathbb{Z}_+$. Observe that since the mixed extension of the potential function is continuous, C and C_{ϵ_n} are closed sets for any $n \in \mathbb{Z}_+$. Since C is closed $\min_{\mathbf{y} \in C} \|\mathbf{x} - \mathbf{y}\|$ is well-defined for any $\mathbf{x} \in \prod_{m \in \mathcal{M}} \Delta E^m$.

We claim that for any $\theta > 0$ the set

$$S_\theta = \left\{ \mathbf{x} \in \prod_{m \in \mathcal{M}} \Delta E^m \mid \min_{\mathbf{y} \in C} \|\mathbf{x} - \mathbf{y}\| < \theta \right\}, \quad (\text{A.3})$$

is such that $C_{\epsilon_n} \subset S_\theta$ for some n . Note that if this claim holds, then it follows from Theorem 6.1 that there exists some T_θ such that for all $t > T_\theta$ we have $\mu_t \in S_\theta$. Using the definition of S_θ given in (A.3), this implies

$$\limsup_{t \rightarrow \infty} \min_{\mathbf{x} \in C} \|\mathbf{x} - \mu_t\| < \theta. \quad (\text{A.4})$$

Moreover, since $\theta > 0$ is arbitrary, and $\|\mathbf{x} - \mu_t\| \geq 0$, using (A.4) we obtain

$$\lim_{t \rightarrow \infty} \min_{\mathbf{x} \in C} \|\mathbf{x} - \mu_t\| = 0.$$

Thus, if we prove $C_{\epsilon_n} \subset S_\theta$ for some n , it follows that μ_t converges to C .

In order to prove $C_{\epsilon_n} \subset S_\theta$ we first obtain a certificate which can be used to guarantee that a mixed strategy profile belongs to S_θ . Then, we show that for large enough n any $\mathbf{z} \in C_{\epsilon_n}$ satisfies this certificate, and hence belongs to S_θ .

It follows from Lemma 6.2 (i) (by setting $\nu = \phi$ and $v = -\min_{\mathbf{y} \in \mathcal{X}_{M\delta}} \phi(\mathbf{y})$) and definition of upper semicontinuity (Definition 6.2) that there exists $\gamma > 0$ such that θ neighborhood of $\{\mathbf{x} \mid \phi(\mathbf{x}) \geq \min_{\mathbf{y} \in \mathcal{X}_{M\delta}} \phi(\mathbf{y})\}$ contains $\{\mathbf{x} \mid \phi(\mathbf{x}) \geq \min_{\mathbf{y} \in \mathcal{X}_{M\delta}} \phi(\mathbf{y}) - \gamma\}$. Hence, for any \mathbf{z} satisfying $\phi(\mathbf{z}) \geq \min_{\mathbf{y} \in \mathcal{X}_{M\delta}} \phi(\mathbf{y}) - \gamma$ there exists some \mathbf{x} satisfying $\phi(\mathbf{x}) \geq \min_{\mathbf{y} \in \mathcal{X}_{M\delta}} \phi(\mathbf{y})$ and $\|\mathbf{x} - \mathbf{z}\| < \theta$. Note that the definition of S_θ implies that \mathbf{z} for which there exists such \mathbf{x} belongs to S_θ . Thus, if $\phi(\mathbf{z}) \geq \min_{\mathbf{y} \in \mathcal{X}_{M\delta}} \phi(\mathbf{y}) - \gamma$ it follows that $\mathbf{z} \in S_\theta$.

We next show that for large enough n , any \mathbf{z} which belongs to C_{ϵ_n} , satisfies the above certificate and hence belongs to S_θ . Let L denote the Lipschitz constant for the mixed extension of ϕ , as given in Lemma 6.1 (i), and define $\theta' = \gamma/L > 0$. Lemma 6.2 (ii) and Definition 6.2 imply that for large enough n , $\mathcal{X}_{M\delta + \frac{1}{n}}$ is contained in θ' neighborhood of $\mathcal{X}_{M\delta}$, i.e., if $\mathbf{y} \in \mathcal{X}_{M\delta + \frac{1}{n}}$ then there exists $\mathbf{x} \in \mathcal{X}_{M\delta}$ such that $\|\mathbf{x} - \mathbf{y}\| < \theta'$. Moreover, by Lemma 6.1 (i), it follows that $\phi(\mathbf{y}) \geq \phi(\mathbf{x}) - L\theta' = \phi(\mathbf{x}) - \gamma$. Thus, we conclude that there exists large enough n such that

$$\min_{\mathbf{y} \in \mathcal{X}_{M\delta + 1/n}} \phi(\mathbf{y}) \geq \min_{\mathbf{y} \in \mathcal{X}_{M\delta}} \phi(\mathbf{y}) - \gamma. \quad (\text{A.5})$$

Let $\mathbf{z} \in C_{\epsilon_n}$ for some n for which (A.5) holds. By definition of C_ϵ it follows that $\phi(\mathbf{z}) \geq \min_{\mathbf{y} \in \mathcal{X}_{M\delta + 1/n}} \phi(\mathbf{y})$. Thus, (A.5) implies that $\phi(\mathbf{z}) \geq \min_{\mathbf{y} \in \mathcal{X}_{M\delta}} \phi(\mathbf{y}) - \gamma$. However, as argued before such \mathbf{z} belong to S_θ . Hence, we have established that for large enough n , if $\mathbf{z} \in C_{\epsilon_n}$ then $\mathbf{z} \in S_\theta$. Therefore, the claim follows. \square

Proof of Theorem 6.2: Assume that \mathcal{G} has l equilibria, denoted by $\mathbf{x}_1, \dots, \mathbf{x}_l$. Define the minimum pairwise distance between the equilibria as $d \triangleq \min_{i \neq j} \|\mathbf{x}_i - \mathbf{x}_j\|$. Let $f : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ be a function such that

$$f(\alpha) = \max_{\mathbf{x} \in \mathcal{X}_\alpha} \min_{k \in \{1, \dots, l\}} \|\mathbf{x} - \mathbf{x}_k\|, \quad (\text{A.6})$$

for all $\alpha \in \mathbb{R}_+$. Note that $\min_{k \in \{1, \dots, l\}} \|\mathbf{x} - \mathbf{x}_k\|$ is continuous in \mathbf{x} , since it is minimum of finitely many continuous functions. Moreover, \mathcal{X}_α is a compact set, since ϵ -equilibria are defined by finitely many inequality constraints of the form (9). Therefore, in (A.6) maximum is achieved and f is well-defined for all $\alpha \geq 0$. From the definition of f , it follows that the union of closed balls of radius $f(\alpha)$, centered at equilibria, contain α -equilibrium set of the game. Thus, intuitively, $f(\alpha)$ captures the size of a closed neighborhood of equilibria, which contains α -equilibria of the underlying game. This is illustrated in Figure A.7.

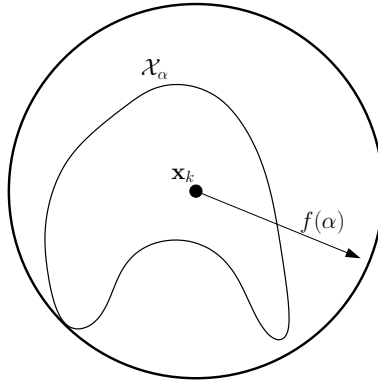


Figure A.7: Consider a game with a unique equilibrium \mathbf{x}_k . The α -equilibrium set of the game (enclosed by a solid line around \mathbf{x}_k) is contained in the $f(\alpha)$ neighborhood of this equilibrium.

Let $a > 0$ be such that $f(a) < d/4$, i.e., every a -equilibrium is at most $d/4$ distant from an equilibrium of a game. Lemma 6.2 (ii) implies (using upper semicontinuity at $\alpha = 0$) that such a exists. Since d is defined as the minimum pairwise distance between the equilibria, it follows that a -equilibria of the game are contained in disjoint $f(a) < d/4$ neighborhoods around equilibria of the game, i.e., if $\mathbf{x} \in \mathcal{X}_a$, then $\|\mathbf{x} - \mathbf{x}_k\| \leq f(a)$ for exactly one equilibrium \mathbf{x}_k . Moreover, for $a_1 \leq a$, since $\mathcal{X}_{a_1} \subset \mathcal{X}_a$, it follows that a_1 -equilibria of the game are contained in disjoint neighborhoods of equilibria.

We prove the theorem in 5 steps summarized below. First two steps explore the properties of function f , and define $\bar{\delta}$ and $\bar{\epsilon}$ presented in the theorem statement. Last three steps are the main steps of the proof, where we establish convergence of fictitious play to a neighborhood of equilibria.

- *Step 1:* We first show that f is (i) weakly increasing, (ii) upper semicontinuous, and it satisfies (iii) $f(0) = 0$, (iv) $f(x) \rightarrow 0$ as $x \rightarrow 0$.
- *Step 2:* We show that there exists some $\bar{\delta} > 0$ and $\bar{\epsilon} > 0$ such that the following inequalities hold:

$$M\bar{\delta} + \bar{\epsilon} < a, \quad (\text{A.7})$$

and

$$f(M\bar{\delta} + \bar{\epsilon}) < \frac{(a - M\bar{\delta})d}{24LM}. \quad (\text{A.8})$$

We will prove the statement of the theorem assuming that $0 \leq \delta < \bar{\delta}$, and establish convergence to the set in (16), for any ϵ such that $0 < \epsilon \leq \bar{\epsilon}$. As can be seen from the definition of a and f (see (A.6)), the first inequality guarantees that $M\bar{\delta} + \bar{\epsilon}$ -equilibrium set is contained in disjoint neighborhoods of equilibria, and the second one guarantees that these neighborhoods are small. In Step 4, we will exploit this observation, and use the inequalities in (A.7) and (A.8) to establish that the empirical frequency distribution μ_t can visit the component of $\mathcal{X}_{M\bar{\delta} + \bar{\epsilon}}$ contained in the neighborhood of only a single equilibrium infinitely often.

- *Step 3:* Let ϵ_1, ϵ_2 be such that $\epsilon_2 > \epsilon_1 > 0$. Assume that (i) at some time instant T , μ_t is contained in $\mathcal{X}_{M\delta + \epsilon_1}$, (ii) at time instants T_1 and T_2 (such that $T_2 > T_1 > T$) μ_t leaves $\mathcal{X}_{M\delta + \epsilon_1}$ and $\mathcal{X}_{M\delta + \epsilon_2}$ respectively and (iii) at time instants T'_2 and T'_1 (such that $T'_1 > T'_2 > T_2$) μ_t returns back to $\mathcal{X}_{M\delta + \epsilon_2}$ and $\mathcal{X}_{M\delta + \epsilon_1}$ respectively. In Figure A.8, the path μ_t follows between T_1 and T'_1 is illustrated.

In this step, we provide a lower bound on $\phi(\mu_{T'_1}) - \phi(\mu_{T_1})$, i.e., the increase in the potential when μ_t follows such a path. This lower bound holds for any ϵ_1 and ϵ_2 provided that $\epsilon_2 > \epsilon_1 > 0$. We use this result by choosing different values for ϵ_1 and ϵ_2 in Steps 4 and 5.

Our lower bound in Step 3 is a function of ϵ_2 . In addition to this lower bound, in Steps 4 and 5, we use the $M\delta + \epsilon_1$ equilibrium set and Lipschitz continuity of the potential to provide an upper bound on $\phi(\mu_{T'_1}) - \phi(\mu_{T_1})$ as a function of ϵ_1 . Thus, properties of $M\delta + \epsilon_1$ and $M\delta + \epsilon_2$ equilibrium sets are exploited for obtaining upper and lower bounds on $\phi(\mu_{T'_1}) - \phi(\mu_{T_1})$ respectively. We establish convergence of fictitious play updates to a neighborhood of an equilibrium by using these bounds together in Steps 4 and 5. We emphasize that allowing for two different approximate equilibrium sets leads to better bounds on $\phi(\mu_{T'_1}) - \phi(\mu_{T_1})$, and a more informative characterization of the limiting behavior of fictitious play, as opposed to using a single approximate equilibrium set, i.e., setting $\epsilon_1 = \epsilon_2$.

- *Step 4:* Our objective in this step is to establish that fictitious play can visit the component of an approximate equilibrium set contained in the neighborhood of only one equilibrium infinitely often.

Let $\epsilon_1 = \bar{\epsilon}$ and $\epsilon_2 = a - M\bar{\delta}$. By (A.7) we have $\epsilon_1 < \epsilon_2$, and using the definition of a we establish that $\mathcal{X}_{M\delta + \epsilon_1}$ and $\mathcal{X}_{M\delta + \epsilon_2}$ are contained in disjoint neighborhoods of equilibria. Assume that μ_t leaves the components of $\mathcal{X}_{M\delta + \epsilon_1}$ and $\mathcal{X}_{M\delta + \epsilon_2}$ in the neighborhood of equilibrium \mathbf{x}_k , and reaches to a similar neighborhood around equilibrium $\mathbf{x}_{k'}$. Using Step 3 we establish a lower bound on the increase in the potential when μ_t follows such a trajectory. We also provide an upper bound, using the Lipschitz continuity of the potential and inequalities (A.7) and (A.8). Comparing these bounds, we establish that the maximum potential in the neighborhood of equilibrium \mathbf{x}_k is lower than the minimum potential in the neighborhood of $\mathbf{x}_{k'}$. Since, \mathbf{x}_k and $\mathbf{x}_{k'}$ are arbitrary, this

observation implies that μ_t cannot visit the component of $\mathcal{X}_{M\delta+\epsilon_1}$ contained in the neighborhood of \mathbf{x}_k at a later time instant. Hence, it follows that μ_t visits only one such component infinitely often.

- *Step 5:* In this step we show that μ_t converges to the approximate equilibrium set given in the theorem statement.

Let ϵ_1, ϵ_2 be such that $0 < \epsilon_1 < \epsilon_2 \leq \bar{\epsilon}$. We consider the equilibrium, whose neighborhood is visited infinitely often (as obtained in Step 4), and a trajectory of μ_t which leaves the components of $\mathcal{X}_{M\delta+\epsilon_1}$ and $\mathcal{X}_{M\delta+\epsilon_2}$ contained in the neighborhood of this equilibrium and returns back to these sets at a later time instant (as illustrated in Figure A.8). As in Step 4, Lipschitz continuity of ϕ is used to obtain an upper bound on the increase in the potential between the end points of this trajectory. Together with the lower bound obtained in Step 3, this provides a bound on how far μ_t can get from the component of $\mathcal{X}_{M\delta+\epsilon_2}$ contained in this neighborhood. Choosing ϵ_1 arbitrarily small (for a fixed ϵ_2) we obtain the tightest such bound. Using this result, we quantify how far μ_t can get from the equilibria of the game (after sufficient time) and the theorem follows.

Next we prove each of these steps.

Step 1: By definition $\mathcal{X}_{\alpha_1} \subset \mathcal{X}_\alpha$ for any $\alpha_1 \leq \alpha$. Since the feasible set of the maximization problem in (A.6) is given by \mathcal{X}_α , this implies that $f(\alpha_1) \leq f(\alpha)$, i.e., f is a weakly increasing function of its argument. Note that the feasible set of the maximization problem in (A.6) can be given by the correspondence $g(\alpha) = \mathcal{X}_\alpha$, which is upper semi continuous in α as shown in Lemma 6.2 (ii). Since as a function of \mathbf{x} , $\min_{k \in \{1, \dots, l\}} \|\mathbf{x} - \mathbf{x}_k\|$ is continuous it follows from Berge's maximum theorem (see Berge (1963)) that for $\alpha \geq 0$, $f(\alpha)$ is an upper semicontinuous function.

The set \mathcal{X}_0 corresponds to the set of equilibria of the game, hence $\mathcal{X}_0 = \{\mathbf{x}_1, \dots, \mathbf{x}_l\}$. Thus, the definition of f implies that $f(0) = 0$. Moreover, upper semicontinuity of f implies that for any $\epsilon > 0$, there exists some neighborhood V of 0, such that $f(x) \leq \epsilon$ for all $x \in V$. Since, $f(x) \geq 0$ by definition, this implies that $\lim_{x \rightarrow 0} f(x)$ exists and equals to 0.

Step 2: Let $\bar{\delta} > 0$ be small enough such that $M\bar{\delta} < a/2$. Since $\lim_{x \rightarrow 0} f(x) = 0$, it follows that for sufficiently small $\bar{\delta}$ and $\bar{\epsilon}$, we obtain $f(M\bar{\delta} + \bar{\epsilon}) < \frac{a\bar{\delta}}{48LM} < \frac{(a-M\bar{\delta})\bar{\delta}}{24LM}$ and $M\bar{\delta} + \bar{\epsilon} < a$.

Step 3: Let ϵ_1, ϵ_2 be such that $0 < \epsilon_1 < \epsilon_2$. Assume $T > 0$ is large enough so that for $t > T$,

$$\begin{aligned} \phi(\mu_{t+1}) - \phi(\mu_t) &\geq \frac{2\epsilon_1}{3(t+1)} \text{ if } \mu_t \notin \mathcal{X}_{M\delta+\epsilon_1}, \text{ and similarly} \\ \phi(\mu_{t+1}) - \phi(\mu_t) &\geq \frac{2\epsilon_2}{3(t+1)} \text{ if } \mu_t \notin \mathcal{X}_{M\delta+\epsilon_2}. \end{aligned} \tag{A.9}$$

Existence of T satisfying these inequalities follows from Lemma 6.3, since for large T and $t > T$, this lemma implies $\phi(\mu_{t+1}) - \phi(\mu_t) \geq \frac{\epsilon_1}{(t+1)} + O\left(\frac{1}{t^2}\right) \geq \frac{2\epsilon_1}{3(t+1)}$ if $\mu_t \notin \mathcal{X}_{M\delta+\epsilon_1}$, and similarly if $\mu_t \notin \mathcal{X}_{M\delta+\epsilon_2}$.

Since $\phi(\mu_t)$ increases outside $M\delta+\epsilon_1$ -equilibrium set for $t > T$, as (A.9) suggests, it follows that μ_t visits $\mathcal{X}_{M\delta+\epsilon_1}$ (and $\mathcal{X}_{M\delta+\epsilon_2}$ since $\mathcal{X}_{M\delta+\epsilon_1} \subset \mathcal{X}_{M\delta+\epsilon_2}$) infinitely often. Otherwise $\phi(\mu_t)$

increases unboundedly, and we reach a contradiction since mixed extension of the potential is a bounded function.

Assume that at some time after T , μ_t leaves $\mathcal{X}_{M\delta+\epsilon_1}$ and $\mathcal{X}_{M\delta+\epsilon_2}$ and returns back to $\mathcal{X}_{M\delta+\epsilon_1}$ at a later time instant. In this step, we quantify how much the potential increases when μ_t follows such a path. We first define time instants T_1 , T_2 , T'_1 , and T'_2 satisfying $T < T_1 \leq T_2 < T'_2 \leq T'_1$, as follows:

- T_1 is a time instant when μ_t leaves $\mathcal{X}_{M\delta+\epsilon_1}$, i.e., $\mu_{T_1-1} \in \mathcal{X}_{M\delta+\epsilon_1}$ and $\mu_t \notin \mathcal{X}_{M\delta+\epsilon_1}$ for $T_1 \leq t < T'_1$.
- T_2 is a time instant when μ_t leaves $\mathcal{X}_{M\delta+\epsilon_2}$, i.e., $\mu_{T_2-1} \in \mathcal{X}_{M\delta+\epsilon_2}$ and $\mu_t \notin \mathcal{X}_{M\delta+\epsilon_2}$ for $T_2 \leq t < T'_2$.
- T'_2 is the first time instant after T_2 when μ_t returns back to $\mathcal{X}_{M\delta+\epsilon_2}$, i.e., $\mu_{T'_2-1} \notin \mathcal{X}_{M\delta+\epsilon_2}$ and $\mu_{T'_2} \in \mathcal{X}_{M\delta+\epsilon_2}$.
- T'_1 is the first time instant after T_1 when μ_t returns back to $\mathcal{X}_{M\delta+\epsilon_1}$, i.e., $\mu_{T'_1-1} \notin \mathcal{X}_{M\delta+\epsilon_1}$ and $\mu_{T'_1} \in \mathcal{X}_{M\delta+\epsilon_1}$.

The definitions are illustrated in Figure A.8. We next provide a lower bound on the quantity $\phi(\mu_{T'_1}) - \phi(\mu_{T_1})$. Note that if there are multiple time instants between T_1 and T'_1 for which μ_t leaves $\mathcal{X}_{M\delta+\epsilon_2}$ (as in the figure), any of these time instants can be chosen as T_2 to obtain a lower bound.

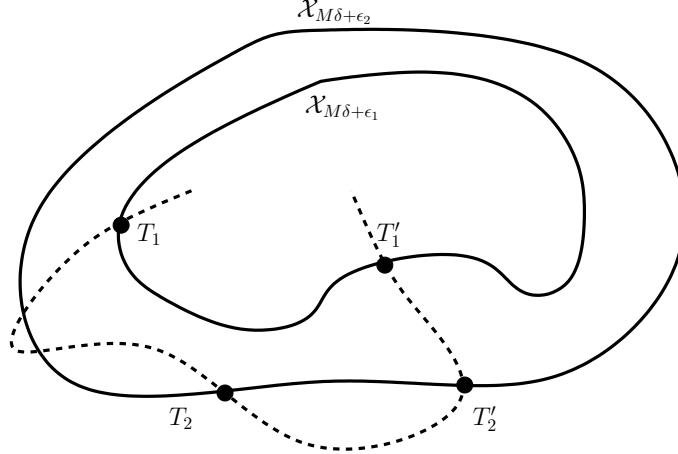


Figure A.8: Trajectory of μ_t (initialized at the left end of the dashed line) is illustrated. T_1 and T_2 correspond to the time instants μ_t leaves $\mathcal{X}_{M\delta+\epsilon_1}$ and $\mathcal{X}_{M\delta+\epsilon_2}$ respectively. T'_1 and T'_2 correspond to the time instants μ_t enters $\mathcal{X}_{M\delta+\epsilon_1}$ and $\mathcal{X}_{M\delta+\epsilon_2}$ respectively.

By definition, for t such that $T_2 \leq t < T'_2$, we have $\mu_t \notin \mathcal{X}_{M\delta+\epsilon_2}$, and for t such that $T_1 \leq t < T_2$ or $T'_2 \leq t < T'_1$, we have $\mu_t \notin \mathcal{X}_{M\delta+\epsilon_1}$. Thus, it follows from (A.9) that

$$\phi(\mu_{t+1}) - \phi(\mu_t) \geq \frac{2\epsilon_2}{3(t+1)} \quad \text{for } T_2 \leq t < T'_2, \quad (\text{A.10})$$

and consequently,

$$\phi(\mu_{T'_2}) - \phi(\mu_{T_2}) = \sum_{t=T_2}^{T'_2-1} \phi(\mu_{t+1}) - \phi(\mu_t) \geq \sum_{t=T_2}^{T'_2-1} \frac{2\epsilon_2}{3(t+1)}. \quad (\text{A.11})$$

Similarly, since $\mu_t \notin \mathcal{X}_{M\delta+\epsilon_1}$ for t such that $T_1 \leq t < T_2$ or $T'_2 \leq t < T'_1$, using (A.9) we establish

$$\phi(\mu_{T'_1}) - \phi(\mu_{T'_2}) = \sum_{t=T'_2}^{T'_1-1} \phi(\mu_{t+1}) - \phi(\mu_t) \geq \sum_{t=T'_2}^{T'_1-1} \frac{2\epsilon_1}{3(t+1)}, \quad (\text{A.12})$$

$$\phi(\mu_{T_2}) - \phi(\mu_{T_1}) = \sum_{t=T_1}^{T_2-1} \phi(\mu_{t+1}) - \phi(\mu_t) \geq \sum_{t=T_1}^{T_2-1} \frac{2\epsilon_1}{3(t+1)}. \quad (\text{A.13})$$

Since $\phi(\mu_{T'_1}) - \phi(\mu_{T_1}) = (\phi(\mu_{T'_1}) - \phi(\mu_{T'_2})) + (\phi(\mu_{T'_2}) - \phi(\mu_{T_2})) + (\phi(\mu_{T_2}) - \phi(\mu_{T_1}))$, it follows from (A.11), (A.12) and (A.13) that

$$\phi(\mu_{T'_1}) - \phi(\mu_{T_1}) \geq \sum_{t=T_2}^{T'_2-1} \frac{2\epsilon_2}{3(t+1)}. \quad (\text{A.14})$$

Step 4: Let $\epsilon_2 = a - M\bar{\delta}$, and $\epsilon_1 = \bar{\epsilon}$. By definition of $\bar{\epsilon}$ and $\bar{\delta}$ (see Step 2), it follows that $\epsilon_2 > \epsilon_1 > 0$. Assume that $\delta < \bar{\delta}$. Since $a = M\bar{\delta} + \epsilon_2 > M\delta + \epsilon_2 > M\delta + \epsilon_1$ we obtain $\mathcal{X}_{M\delta+\epsilon_1} \subset \mathcal{X}_{M\delta+\epsilon_2} \subset \mathcal{X}_a$. By definition of a , \mathcal{X}_a is contained in disjoint neighborhoods of equilibria. Thus, it follows that components of $\mathcal{X}_{M\delta+\epsilon_1}$ and $\mathcal{X}_{M\delta+\epsilon_2}$ are also contained in disjoint neighborhoods of equilibria. Hence, the definition of f suggests that if $\mathbf{x} \in \mathcal{X}_{M\delta+\epsilon_1}$ then $\|\mathbf{x}_k - \mathbf{x}\| \leq f(M\delta + \epsilon_1)$ (similarly if $\mathbf{x} \in \mathcal{X}_{M\delta+\epsilon_2}$, then $\|\mathbf{x}_k - \mathbf{x}\| \leq f(M\delta + \epsilon_2)$) for exactly one equilibrium \mathbf{x}_k .

Let T_1, T_2, T'_1 and T'_2 be defined as in Step 3. In this step, by obtaining an upper bound on $\phi(\mu_{T'_1}) - \phi(\mu_{T_1})$ and refining the lower bound obtained in Step 3 for given values of ϵ_1 and ϵ_2 , we prove that after sufficient time μ_t can visit the component of $\mathcal{X}_{M\delta+\epsilon_1}$ in the neighborhood of a single equilibrium.

Assume that μ_t leaves the component of the $M\delta + \epsilon_1$ -equilibrium set in the neighborhood of equilibrium \mathbf{x}_k , and it reaches to another component in the neighborhood of equilibrium $\mathbf{x}_{k'}$. Since, by definition $\mu_{T_1-1}, \mu_{T'_1} \in \mathcal{X}_{M\delta+\epsilon_1}$, and $\mu_{T_2-1}, \mu_{T'_2} \in \mathcal{X}_{M\delta+\epsilon_2}$, it follows that μ_{T_1-1} and μ_{T_2-1} belong to neighborhoods of equilibrium \mathbf{x}_k , whereas, $\mu_{T'_1}$ and $\mu_{T'_2}$ belong to neighborhoods of $\mathbf{x}_{k'}$, i.e.,

$$\|\mathbf{x}_k - \mu_{T_1-1}\| \leq f(M\delta + \epsilon_1) \quad \text{and} \quad \|\mathbf{x}_k - \mu_{T_2-1}\| \leq f(M\delta + \epsilon_2), \quad \text{whereas,} \quad (\text{A.15})$$

$$\|\mathbf{x}_{k'} - \mu_{T'_1}\| \leq f(M\delta + \epsilon_1) \quad \text{and} \quad \|\mathbf{x}_{k'} - \mu_{T'_2}\| \leq f(M\delta + \epsilon_2). \quad (\text{A.16})$$

By definition of d we have $\|\mathbf{x}_k - \mathbf{x}_{k'}\| \geq d$. Since $a > M\delta + \epsilon_2$, it follows that $f(M\delta + \epsilon_2) \leq f(a) < d/4$, and hence the second inequalities in (A.15) and (A.16) imply

$$\|\mu_{T'_2} - \mu_{T_2-1}\| > \frac{d}{2}. \quad (\text{A.17})$$

Using this inequality, we next refine the lower bound on $\phi(\mu_{T'_1}) - \phi(\mu_{T_1})$ obtained in Step 3. By (10), with an update at time t , the empirical frequency distribution can change by at most

$$\|\mu_{t+1} - \mu_t\| = \frac{1}{t+1} \|\mu_t - I_t\| \leq \frac{1}{t+1} (\|\mu_t\| + \|I_t\|) \leq \frac{2M}{t+1}, \quad (\text{A.18})$$

where the last inequality follows from the fact that $\mu_t = \{\mu_t^m\}_{m \in \mathcal{M}}$, and $I_t = \{I_t^m\}_{m \in \mathcal{M}}$, and $\|\mu_t^m\|, \|I_t^m\| \leq 1$, since $I_t^m, \mu_t^m \in \Delta E^m$. Hence, if T_2 is sufficiently large, then $\|\mu_{T_2} - \mu_{T_2-1}\|$ is small enough so that (A.17) implies $\|\mu_{T'_2} - \mu_{T_2}\| > \frac{d}{2}$. Using this together with (A.18), we conclude

$$\sum_{t=T_2}^{T'_2-1} \frac{2M}{t+1} \geq \sum_{t=T_2}^{T'_2-1} \|\mu_{t+1} - \mu_t\| \geq \left\| \sum_{t=T_2}^{T'_2-1} \mu_{t+1} - \mu_t \right\| = \|\mu_{T'_2} - \mu_{T_2}\| > \frac{d}{2}. \quad (\text{A.19})$$

Thus, the lower bound on $\phi(\mu_{T'_1}) - \phi(\mu_{T_1})$ provided in (A.14) takes the following form:

$$\phi(\mu_{T'_1}) - \phi(\mu_{T_1}) \geq \sum_{t=T_2}^{T'_2-1} \frac{2\epsilon_2}{3(t+1)} \geq \frac{\epsilon_2 d}{6M}. \quad (\text{A.20})$$

Next we provide an upper bound on $\phi(\mu_{T'_1}) - \phi(\mu_{T_1})$, using Lipschitz continuity of the potential and the properties of the $M\delta + \epsilon_1$ equilibrium set. Let $\bar{\phi}_k = \max_{\{\mathbf{x} \mid \|\mathbf{x} - \mathbf{x}_k\| \leq f(M\delta + \epsilon_1)\}} \phi(\mathbf{x})$, and define \mathbf{y}_k as a strategy profile which achieves this maximum. Similarly, let $\underline{\phi}_{k'} = \min_{\{\mathbf{x} \mid \|\mathbf{x} - \mathbf{x}_{k'}\| \leq f(M\delta + \epsilon_1)\}} \phi(\mathbf{x})$ and define $\mathbf{y}_{k'}$ as a strategy profile which achieves this minimum. Observe that

$$\begin{aligned} \underline{\phi}_{k'} - \bar{\phi}_k &= \phi(\mathbf{y}_{k'}) - \phi(\mathbf{y}_k) \\ &= (\phi(\mathbf{y}_{k'}) - \phi(\mu_{T'_1})) + (\phi(\mu_{T'_1}) - \phi(\mu_{T_1})) + (\phi(\mu_{T_1}) - \phi(\mathbf{y}_k)). \end{aligned} \quad (\text{A.21})$$

Note that by (A.15) and (A.16), and the definitions of \mathbf{y}_k and $\mathbf{y}_{k'}$, we have $\mu_{T'_1}, \mathbf{y}_{k'} \in \{\mathbf{x} \mid \|\mathbf{x} - \mathbf{x}_{k'}\| \leq f(M\delta + \epsilon_1)\}$, and $\mu_{T_1-1}, \mathbf{y}_k \in \{\mathbf{x} \mid \|\mathbf{x} - \mathbf{x}_k\| \leq f(M\delta + \epsilon_1)\}$. Hence, using Lipschitz continuity of ϕ (and denoting the Lipschitz constant by L) it follows that $\phi(\mathbf{y}_{k'}) - \phi(\mu_{T'_1}) \geq -2Lf(M\delta + \epsilon_1)$, and $\phi(\mu_{T_1-1}) - \phi(\mathbf{y}_k) \geq -2Lf(M\delta + \epsilon_1)$. Moreover, (A.18) and Lipschitz continuity of ϕ imply that $\phi(\mu_{T_1}) - \phi(\mu_{T_1-1}) = O\left(\frac{1}{T_1}\right)$. Thus, using (A.21) we obtain the following upper bound on $\phi(\mu_{T'_1}) - \phi(\mu_{T_1})$:

$$\underline{\phi}_{k'} - \bar{\phi}_k + 4Lf(M\delta + \epsilon_1) + O\left(\frac{1}{T_1}\right) \geq \phi(\mu_{T'_1}) - \phi(\mu_{T_1}). \quad (\text{A.22})$$

Using the lower and upper bounds we obtained in (A.20) and (A.22), it follows that

$$\underline{\phi}_{k'} - \bar{\phi}_k + 4Lf(M\delta + \epsilon_1) + O\left(\frac{1}{T_1}\right) \geq \frac{\epsilon_2 d}{6M}. \quad (\text{A.23})$$

Since $\epsilon_2 = a - M\bar{\delta}$, and $\epsilon_1 = \bar{\epsilon}$, using the fact that f is an increasing function and $\delta < \bar{\delta}$, it follows from (A.23) that

$$\underline{\phi}_{k'} - \bar{\phi}_k \geq \frac{(a - M\bar{\delta})d}{6M} - 4Lf(M\delta + \bar{\epsilon}) + O\left(\frac{1}{T_1}\right) \geq \frac{(a - M\bar{\delta})d}{6M} - 4Lf(M\bar{\delta} + \bar{\epsilon}) + O\left(\frac{1}{T_1}\right)$$

Note that (A.8) implies $\frac{(a-M\bar{\delta})d}{6M} - 4Lf(M\bar{\delta} + \bar{\epsilon}) > 0$. Thus, for sufficiently large T_1 we obtain $\underline{\phi}_{k'} - \bar{\phi}_k > 0$. Therefore, we conclude when μ_t leaves the component of $\mathcal{X}_{M\delta+\epsilon_1}$ contained in the neighborhood of some equilibrium \mathbf{x}_k , and enters that of another equilibrium $\mathbf{x}_{k'}$, then the minimum potential in the new neighborhood is strictly larger than the maximum potential in the older one (for sufficiently large T_1). Since this is true for arbitrary equilibria \mathbf{x}_k and $\mathbf{x}_{k'}$, it follows that after entering the component of $\mathcal{X}_{M\delta+\epsilon_1}$ in the neighborhood of $\mathbf{x}_{k'}$, μ_t cannot return to the component in the neighborhood of \mathbf{x}_k , as doing so contradicts with the relation between the minimum and maximum potentials in these neighborhoods. Thus, after sufficient time, μ_t can visit the component of $\mathcal{X}_{M\delta+\epsilon_1}$ (or equivalently $\mathcal{X}_{M\delta+\bar{\epsilon}}$) in the neighborhood of a single equilibrium.

Step 5:. Let ϵ_1 , and ϵ_2 be such that $0 < \epsilon_1 < \epsilon_2 \leq \bar{\epsilon}$. As established in Step 4, there exists some T , such that for $t > T$, μ_t visits the component of $\mathcal{X}_{M\delta+\bar{\epsilon}}$, in the neighborhood of a single equilibrium, say \mathbf{x}_k .

Assume that T_1, T_2, T'_1 and T'_2 are defined as in Step 3, and let $T_1 > T + 1$. Since $\epsilon_1 < \epsilon_2 \leq \bar{\epsilon}$, we have $\mathcal{X}_{M\delta+\epsilon_1} \subset \mathcal{X}_{M\delta+\epsilon_2} \subset \mathcal{X}_{M\delta+\bar{\epsilon}}$, and $T_1 > T + 1$ implies that μ_t can only visit the components of $\mathcal{X}_{M\delta+\epsilon_1}$ and $\mathcal{X}_{M\delta+\epsilon_2}$ contained in the neighborhood of \mathbf{x}_k . Following a similar approach to Step 4, we next obtain upper and lower bounds on $\phi(\mu_{T'_1}) - \phi(\mu_{T_1})$, and use these bounds to establish convergence to the mixed equilibrium set given in the theorem statement.

Define d^* as the maximum distance of μ_t from $\mathcal{X}_{M\delta+\epsilon_2}$ for t such that $T + 1 < T_2 \leq t \leq T'_2 - 1$, i.e.,

$$d^* = \max_{\{t|T_2 \leq t \leq T'_2-1\}} \min_{\mathbf{x} \in \mathcal{X}_{M\delta+\epsilon_2}} \|\mu_t - \mathbf{x}\|.$$

Since $\mu_{T_2-1}, \mu_{T'_2} \in \mathcal{X}_{M\delta+\epsilon_2}$ by definition, the total length of the trajectory between $T_2 - 1$ and T'_2 is an upper bound on $2d^*$, i.e.,

$$2d^* \leq \sum_{t=T_2-1}^{T'_2-1} \|\mu_{t+1} - \mu_t\|.$$

As explained in (A.18), $\|\mu_{t+1} - \mu_t\| \leq \frac{2M}{t+1}$, thus the above inequality implies

$$2d^* \leq \sum_{t=T_2-1}^{T'_2-1} \frac{2M}{t+1} = \sum_{t=T_2}^{T'_2-1} \frac{2M}{t+1} + \frac{2M}{T_2}. \quad (\text{A.24})$$

Using this inequality, the lower bound in (A.14) implies

$$\phi(\mu_{T'_1}) - \phi(\mu_{T_1}) \geq \sum_{t=T_2}^{T'_2-1} \frac{2\epsilon_2}{3(t+1)} \geq \left(d^* - \frac{M}{T_2}\right) \frac{2\epsilon_2}{3M} \quad (\text{A.25})$$

We next obtain an upper bound on $\phi(\mu_{T'_1}) - \phi(\mu_{T_1})$. By definition of f , $\mathcal{X}_{M\delta+\epsilon_1}$ is contained in $f(M\delta + \epsilon_1)$ neighborhoods of equilibria. For $T_1 > T + 1$, μ_t can only visit the component of $\mathcal{X}_{M\delta+\epsilon_1}$ in the neighborhood of \mathbf{x}_k , as can be seen from the definition of T . Thus, since $\mu_{T_1-1}, \mu_{T'_1} \in \mathcal{X}_{M\delta+\epsilon_1}$, it follows that $\mu_{T_1-1}, \mu_{T'_1} \in \{\mathbf{x} \mid \|\mathbf{x} - \mathbf{x}_k\| \leq f(M\delta + \epsilon_1)\}$. By Lipschitz

continuity of the potential function it follows that $\phi(\mu_{T_1'}) - \phi(\mu_{T_1-1}) \leq 2f(M\delta + \epsilon_1)L$. Additionally, by (A.18) Lipschitz continuity also implies that $\phi(\mu_{T_1}) - \phi(\mu_{T_1-1}) \leq \frac{2ML}{T_1}$. Combining these we obtain the following upper bound on $\phi(\mu_{T_1'}) - \phi(\mu_{T_1})$:

$$\phi(\mu_{T_1'}) - \phi(\mu_{T_1}) \leq 2f(M\delta + \epsilon_1)L + \frac{2ML}{T_1}. \quad (\text{A.26})$$

It follows from the upper and lower bounds on $\phi(\mu_{T_1'}) - \phi(\mu_{T_1})$ given in (A.25) and (A.26) that

$$\left(d^* - \frac{M}{T_2}\right) \frac{2\epsilon_2}{3M} \leq 2f(M\delta + \epsilon_1)L + \frac{2ML}{T_1}$$

Thus, for sufficiently large T_1 (and hence T_2), we obtain

$$d^* \leq \frac{3f(M\delta + \epsilon_1)ML}{\epsilon_2} + \frac{3M^2L}{\epsilon_2 T_1} + \frac{M}{T_2} \leq \frac{4f(M\delta + \epsilon_1)ML}{\epsilon_2}. \quad (\text{A.27})$$

Note that in the above derivation ϵ_1 is an arbitrary number that satisfies $0 < \epsilon_1 < \epsilon_2$. Thus, (A.27) implies that

$$d^* \leq \limsup_{\epsilon_1 \rightarrow 0} \frac{4f(M\delta + \epsilon_1)ML}{\epsilon_2} \leq \frac{4f(M\delta)ML}{\epsilon_2}, \quad (\text{A.28})$$

where the last inequality follows by upper semicontinuity of f . Thus, by definition of d^* , we conclude that μ_t converges d^* neighborhood of $\mathcal{X}_{M\delta+\epsilon_2}$. Hence, using (A.28), we can establish convergence of μ_t to

$$\left\{ \mathbf{x} \mid \|\mathbf{x} - \mathbf{y}\| \leq \frac{4f(M\delta)ML}{\epsilon_2}, \text{ for some } \mathbf{y} \in \mathcal{X}_{M\delta+\epsilon_2} \right\}. \quad (\text{A.29})$$

Observe that definition of f implies if $\mathbf{y} \in \mathcal{X}_{M\delta+\epsilon_2}$, then for some equilibrium \mathbf{x}_k we have $\|\mathbf{x}_k - \mathbf{y}\| \leq f(M\delta + \epsilon_2)$. Thus, using (A.29) and triangle inequality, we conclude that μ_t converges to

$$\left\{ \mathbf{x} \mid \|\mathbf{x} - \mathbf{x}_k\| \leq \frac{4f(M\delta)ML}{\epsilon_2} + f(M\delta + \epsilon_2), \text{ for some equilibrium } \mathbf{x}_k \right\}. \quad (\text{A.30})$$

Noting that in (A.30) ϵ_2 is an arbitrary number satisfying $0 < \epsilon_2 \leq \bar{\epsilon}$, the theorem follows. \square