Games and Economic Behavior 82 (2013) 66–90

Contents lists available at SciVerse ScienceDirect

Games and Economic Behavior

www.elsevier.com/locate/geb

Dynamics in near-potential games

Ozan Candogan *, Asuman Ozdaglar, Pablo A. Parrilo

Laboratory of Information and Decision Systems, Massachusetts Institute of Technology, 77 Massachusetts Avenue, Room 32-D608, Cambridge, MA 02139, United States

ARTICLE INFO

Article history:
Received 21 July 2011
Available online 9 July 2013

JEL classification:
C61
C72
D83

Keywords:
Dynamics in games
Near-potential games
Best response dynamics
Logit response dynamics
Fictitious play

ABSTRACT

We consider discrete-time learning dynamics in finite strategic form games, and show that games that are close to a potential game inherit many of the dynamical properties of potential games. We first study the evolution of the sequence of pure strategy profiles under better/best response dynamics. We show that this sequence converges to a (pure) approximate equilibrium set whose size is a function of the "distance" to a given nearby potential game. We then focus on logit response dynamics, and provide a characterization of the limiting outcome in terms of the distance of the game to a given potential game and the corresponding potential function. 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, where the size of the neighborhood increases according to the distance to the set of potential games.

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 (Fudenberg and Levine, 1998; Jordan, 1993; 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 (Fudenberg and Levine, 1998; Monderer and Shapley, 1996a, 1996b; 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 discrete-time dynamics in finite strategic form games by exploiting their relation to close potential games.

We start by illustrating via examples that general games which are close in terms of payoffs may have significantly different limiting behavior under simple user dynamics.[1] Our first example focuses on better response dynamics in which

* Corresponding author. Fax: +1 617 253 3578.
E-mail addresses: candogan@mit.edu (O. Candogan), asuman@mit.edu (A. Ozdaglar), parrilo@mit.edu (P.A. Parrilo).

[1] The games presented in these examples are also close in terms of maximum pairwise difference, discussed in Section 2.

0899-8256/$ – see front matter © 2013 Published by Elsevier Inc.
http://dx.doi.org/10.1016/j.geb.2013.07.001

	A	B
A	0, 1	0, 0
B	1, 0	θ, 2

$\mathcal{G}_1$

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

$\mathcal{G}_2$

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

	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	$1+\theta, 1+\theta$

$\mathcal{G}_1$

	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	$1-\theta, 1-\theta$

$\mathcal{G}_2$

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

at each step or strategy profile, a player (chosen consecutively or at random) updates its strategy unilaterally to one that yields a better payoff.[2]

Example 1.1. Consider two games with two players and payoffs given in Fig. 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 Fig. 2. Let θ be an irrational number such that $0 < \theta \ll 1$. It can be seen that $\mathcal{G}_1$ has multiple equilibria (including pure equilibria (A, A), (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.

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). 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 Fudenberg and Levine, 1998; Shapley, 1964). 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 $\mathcal{G}_2$.

These examples suggest that in general, it may not be possible to characterize the limiting dynamics in a given game, by using knowledge of the limiting behavior in a nearby game. In this paper, in contrast with this observation, 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. Moreover, it is possible to provide a *quantitative measure of the size of the limiting set of dynamics* in terms of the 'distance' of the game from potential games. Our approach relies on using the potential function of a close potential game for the analysis of commonly studied update rules.[3] 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

[2] Consider 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.

[3] Throughout the paper, we use the terms *learning dynamics* and *update rules* interchangeably.

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 $\mathcal{X}_\epsilon$ 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.

Update rule	Convergence result
Better/best response dynamics	(Theorem 3.1) 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 4.2) Stationary distribution μ_τ of logit response dynamics is such that $\left|\mu_\tau(\mathbf{p}) - \frac{e^{\frac{1}{\tau}\phi(\mathbf{p})}}{\sum_{\mathbf{q}\in E} e^{\frac{1}{\tau}\phi(\mathbf{q})}}\right| \leqslant \frac{e^{\frac{2\delta(h-1)}{\tau}}-1}{e^{\frac{2\delta(h-1)}{\tau}}+1}$, for all $\mathbf{p}$.
Logit response dynamics	(Corollary 4.3) Stochastically stable strategy profiles of $\mathcal{G}$ are (i) contained in $S = \{\mathbf{p} \mid \phi(\mathbf{p}) \geqslant \max_{\mathbf{q}} \phi(\mathbf{q}) - 4\delta(h-1)\}$, (ii) $4\delta h$-equilibria of $\mathcal{G}$.

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. 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.

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 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 from the original game to a potential game, and is independent of the payoffs in the original game. In particular, our result for better/best response dynamics establish a 'Lipschitz-type' property, i.e., we can find a constant h (which is equal to the number of strategy profiles in the game as shown in Theorem 3.1) such that in a game that is δ different (in terms of maximum pairwise difference) from a potential game the trajectory converges to the δh-equilibrium set.

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 (Alós-Ferrer and Netzer, 2010; Blume, 1993, 1997; Marden and Shamma, 2008; Young, 1993). 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 4.1) 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.

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, exploiting 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,

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.

Update rule	Convergence result
Fictitious play	(Corollary 5.1) Empirical frequencies of dynamics converge to the set of mixed strategies with large enough potential: $\{\mathbf{x} \in \prod_m \Delta E^m \mid \phi(\mathbf{x}) \geqslant \min_{\mathbf{y}\in\mathcal{X}_{M\delta}} \phi(\mathbf{y})\}$
Fictitious play	(Theorem 5.2) 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 $$\left\{\mathbf{x}\,\middle|\, \|\mathbf{x}-\mathbf{x}_k\| \leqslant \frac{4f(M\delta)ML}{\epsilon} + f(M\delta+\epsilon),\ \text{for some equilibrium } \mathbf{x}_k\right\},$$ for any ϵ such that $\bar{\epsilon} \geqslant \epsilon > 0$, where $f:\mathbb{R}_+ \to \mathbb{R}_+$ is an upper semicontinuous function that quantifies the size of mixed equilibrium sets (defined explicitly in Section 5) such that $f(x) \to 0$ as $x \to 0$.

$\mathcal{G}$

	W	SO	SH
W	90, 90	−12, −12	48, 48
SO	−12, −12	1, −1	24, 24
SH	48, 48	24, 24	1, −1

$\hat{\mathcal{G}}$

	W	SO	SH
W	90, 90	−12, −12	48, 48
SO	−12, −12	0, 0	24, 24
SH	48, 48	24, 24	0, 0

Potential ϕ of $\hat{\mathcal{G}}$

	W	SO	SH
W	90	−12	48
SO	−12	0	24
SH	48	24	0

Fig. 3. A game ($\mathcal{G}$) and a nearby potential game ($\hat{\mathcal{G}}$) with potential ϕ, share similar equilibrium set and dynamic properties.

whose size is explicitly characterized.[4] Our result recovers as a special case convergence of empirical frequencies to Nash equilibria in potential games.[5]

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

The framework provided in this paper enables us to study the limiting behavior of adaptive user dynamics in finite strategic form games that are not potential games. In particular, for a given game we can find a nearby potential game by solving a convex optimization problem, 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.

Example 1.3 illustrates our approach for characterization of the limiting behavior of discrete-time dynamics. For ease of exposition, in this example we focus on a game for which it is straightforward to identify a nearby potential game. For more general settings, a nearby potential game can be identified numerically.

Example 1.3. Consider a two-player game $\mathcal{G}$, with payoffs given in Fig. 3. In this game, each player has three strategies (W: Work, SO: Shirk at Office, SH: Shirk at Home). If both players work, then a project succeeds and they each receive a payoff of 90. If only one of them works, then the project may succeed (and they each receive 48), or it may fail (with a resulting payoff of −12). If none of the players work and they shirk at different places, then each receives a payoff of 24. On the other hand, if both shirk at the same place, then the row player has unit payoff, whereas the column player incurs a disutility of one units.

Game $\mathcal{G}$ is not a potential game, as it involves a utility improvement cycle (SO,SO)–(SO,SH)–(SH,SH)–(SH,SO)–(SO,SO) (see Section 2 for details). A nearby potential game $\hat{\mathcal{G}}$, which has maximum pairwise difference (MPD) $\delta = 1$ to the original game, is provided in Fig. 3. It can be seen that in both games the pure strategy (W, W) is the unique equilibrium. We next show that by exploiting the relation between $\mathcal{G}$ and $\hat{\mathcal{G}}$, the outcome of dynamics in $\mathcal{G}$ can be characterized.

In particular, using the results of Table 1 related to better/best response dynamics, we conclude that in $\mathcal{G}$ these update rules converge to an ϵ-equilibrium set, with $\epsilon = 9$ (note that there are $h = 9$ strategy profiles, and the MPD between $\mathcal{G}$ and $\hat{\mathcal{G}}$ is $\delta = 1$). On the other hand, it can be seen that the only pure strategy profile that belongs to this set is (W, W). Hence we conclude that better/best response dynamics converge to the unique equilibrium in $\mathcal{G}$.

Similarly, the results of Table 1 related to logit response dynamics, suggest that stochastically stable strategy profiles of $\mathcal{G}$, belong to $S = \{p \mid \phi(\mathbf{p}) \geqslant \max_{\mathbf{q}} \phi(\mathbf{q}) - 4\delta(h-1) = 90 - 32 = 58\} = \{(W, W)\}$. That is, (W, W) is the only stochastically stable strategy profile of $\mathcal{G}$.

Finally, the results of Table 2 imply that fictitious play converges to $\{\mathbf{x} \in \prod_m \Delta E^m \mid \phi(\mathbf{x}) \geqslant \min_{\mathbf{y}\in\mathcal{X}_2} \phi(\mathbf{y})\}$, where $\mathcal{X}_2$ denotes the 2-equilibrium set of $\mathcal{G}$. It can be numerically checked that for $\mathbf{x} \in \mathcal{X}_2$, we have $\phi(\mathbf{x}) \geqslant 85$. This implies

[4] The bounds we obtain for the limiting behavior of fictitious play dynamics have a different flavor than those for better/best response dynamics, and logit response. While the bounds we obtain for the latter update rules are independent of the payoffs (and a function of only δ and the number of strategy profiles in the game), for fictitious play they are not. This is because, fictitious play results exploit the structure of mixed (approximate) equilibrium sets, which rely on the actual payoff parameters, whereas other dynamics results do not involve mixed strategies.

[5] This result also implies that in near-potential games fictitious play dynamics are upper semicontinuous with respect to payoff parameters. This upper semicontinuity result could alternatively be proved by considering differential inclusions that represent the limiting behavior of fictitious play (see Benaïm et al., 2005), together with upper semicontinuity results on differential inclusions (Li and Zhang, 2002). Our result, in addition to upper semicontinuity, provides explicit bounds on the size of the limiting set.

that the limiting set is a subset of $\{\mathbf{x} \in \prod_m \Delta E^m | \phi(\mathbf{x}) \geqslant 85\}$. Additionally, it can checked that if $x^m(W) < 0.88$ for some player m, then $\phi(\mathbf{x}) < 85$. Thus, we conclude that the empirical frequencies of fictitious play converge to a set, where each player employs strategy W with probability at least 0.88, i.e., the limiting empirical frequencies are contained in $\{\mathbf{x} \in \prod_m \Delta E^m | x^m(W) \geqslant 0.88\}$.

Observe that in this example $\hat{\mathcal{G}}$ can be used to characterize the outcome of dynamics in $\mathcal{G}$, even if two games may lead to different strategy update trajectories. More precisely, the strategy update probabilities for logit response, and best responses in response to some empirical frequencies may be different for these games. Consequently, logit response and fictitious play may have distinct trajectories in these games. Despite that $\hat{\mathcal{G}}$ reveals useful information about the outcome of strategy updates in $\mathcal{G}$.

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

Recent work (Candogan et al., 2011a) exploits some topological properties of finite strategic form games and potential games to obtain a decomposition of the space of games into orthogonal components. This paper also studies the structure of equilibria in near-potential games defined through these orthogonal components. Another related work (Candogan et al., 2010b) discusses structural properties of sets of exact, weighted and ordinal potential games and provides efficient algorithms for approximating a given game with an exact or weighted potential game. Two other related works are Candogan et al. (2010a) and (2011c). Candogan et al. (2010a) focus on a specific power control game in wireless networks and approximates it with a close potential game to design pricing rules that guarantee that limiting behavior of dynamics remain within a neighborhood of a socially optimal point. Additionally, in Candogan et al. (2011c) we provide an informal discussion of better/best response dynamics and logit response dynamics (but not fictitious play) in near-potential games and partial results on their limiting behavior. The current paper provides a complete picture of applications of near-potential games to the analysis of discrete-time update rules, including the results about discrete-time fictitious play dynamics, and a formal discussion of better/best response dynamics and logit response.

A follow-up paper, Candogan et al. (2013), focuses on geometry of sets of games that are equivalent (with respect to various equivalence relations) to potential games. This paper additionally extends the analysis of the current paper to continuous time dynamics (in particular to generalizations of fictitious play) in near-potential games. Continuous time dynamics provide alternative simplified models of strategy updates. These models often allow for a more tractable characterization of the limiting behavior (see e.g., Krishna and Sjöström, 1998), and hence are commonly studied in the literature. We emphasize that the study of better–best response dynamics and logit response, which we focus on in Sections 3 and 4, does not share a similar methodology to Candogan et al. (2013). Additionally, the study of the limiting behavior of discrete-time fictitious play, given in Section 5, requires a characterization of the change in the potential at each *step* of the update rule (e.g., see proofs of Lemma 5.3, Theorems 5.1, and 5.2). This poses additional challenges in the study of discrete-time update rules, which are not present in the analysis of continuous time dynamics.

Another strand of literature focuses on identifying classes of games with similar properties to potential games. Examples include ordinal potential games (Monderer and Shapley, 1996b), best-response potential games (Voorneveld, 2000), pseudo-potential games (Dubey et al., 2006), and nested potential games (Uno, 2007). Even though these classes of games share similar ordinal properties with potential games, for update processes that involve mixed strategies (such as fictitious play), or that rely on actual payoff values (such as logit response), they do not lead to simple analysis unless further structure is imposed (unlike potential games). For this reason, in this paper we follow a different approach, and characterize dynamic properties of games, by exploiting their closeness to potential games.

There are also papers in the literature, which identify classes of games that are strategically equivalent to potential games (Morris and Ui, 2004). These equivalence notions can be used to extend the dynamical properties of potential games to their equivalence classes. However, we want to emphasize that the framework presented in this paper can be applied for study of dynamics in games that are not strategically equivalent to potential games, thereby providing tools for study of dynamics in arbitrary strategic form games.

Paper organization: The rest of the paper is organized as follows: We present the game theoretic preliminaries for our work in Section 2. We present an analysis of better and best response dynamics in near-potential games in Section 3. In Section 4, 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 in Section 5. We close in Section 6 with concluding remarks and future work.

2. Preliminaries

In this section, we present the game-theoretic background that is relevant to our work and introduce a closeness measure for 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: $u^m : \prod_{k\in\mathcal{M}} E^k \to \mathbb{R}$, 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 mth one is denoted by $\mathbf{p}^{-m}$.

A strategy profile $\mathbf{p} \triangleq (p^1, \dots, p^M)$ is an ϵ-equilibrium ($\epsilon \geqslant 0$) if $u^m(q^m, \mathbf{p}^{-m}) - u^m(p^m, \mathbf{p}^{-m}) \leqslant \epsilon$ 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 $\mathcal{X}_\epsilon$. Nash equilibria of a given game are the ϵ-equilibria for $\epsilon = 0$.

The class of potential games (Monderer and Shapley, 1996b), which we discuss next, is central in this paper.

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

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

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 Eq. (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. Our main focus in this paper is on potential games in Definition 2.1, but we explain that our results generalize to the extensions of potential games at the end of Section 5.

An 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 3 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 $u^{m_i}(\mathbf{p}_i) \geqslant u^{m_i}(\mathbf{p}_{i-1})$ 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 ith 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. *(See Monderer and Shapley, 1996b.) A game is a potential game if and only if $I(\gamma) = 0$ for all simple closed paths γ.*

We next provide a formal definition of 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

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

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.[6] 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*.

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

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

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.[7] 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. For a discussion of different strategic equivalence notions in games, and their applications, see Morris and Ui (2004).

It can be seen from Proposition 2.1 that a game is a potential game if and only if it satisfies certain linear equalities. This suggests that the set of potential games is convex, 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 $\alpha\in[0,1]$. This suggests that the closest potential game (in terms of MPD) to a given game, can be obtained by solving a convex optimization problem.[8]

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.

3. 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. 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 3.1 *(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.

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. For simplicity, we assume here that users are chosen randomly to update their strategy.

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.

---

[6] An alternative distance measure can be given by

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

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.

[7] Note that if a game is close to a potential game in terms of payoffs, it is also close in terms of maximum pairwise difference. Conversely, the definition of potential games suggests that if a game is close to a potential game in maximum pairwise difference, then there exists another potential game that is close to this game in terms of payoffs.

[8] This argument relies on the convexity of MPD as a function of the utility functions (Candogan et al., 2011b). An alternative framework for finding near-potential (and weighted potential) games (using a different norm) can be found in Candogan et al. (2011a and 2010b).

Theorem 3.1. *Consider a game $\mathcal{G}$ and let $\hat{\mathcal{G}}$ be a nearby potential game such that $d(\mathcal{G}, \hat{\mathcal{G}}) \leqslant \delta$. 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. Using Definition 3.1, 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 3.1 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 $N \leqslant |E| = h$. Since $u^m(\mathbf{q}) - u^m(\mathbf{r}) = \alpha$, and $u^{m_i}(\mathbf{p}_i) - u^{m_i}(\mathbf{p}_{i-1}) \geqslant 0$ at every step i of the path, this closed improvement path satisfies $\sum_{i=1}^{N}(u^{m_i}(\mathbf{p}_i) - u^{m_i}(\mathbf{p}_{i-1})) \geqslant \alpha$. On the other hand it follows by Proposition 2.1 that the close potential game satisfies $\sum_{i=1}^{N}(\hat{u}^{m_i}(\mathbf{p}_i) - \hat{u}^{m_i}(\mathbf{p}_{i-1})) = 0$. Combining these inequalities we obtain $\alpha \leqslant \sum_{i=1}^{N}(u^{m_i}(\mathbf{p}_i) - u^{m_i}(\mathbf{p}_{i-1})) - (\hat{u}^{m_i}(\mathbf{p}_i) - \hat{u}^{m_i}(\mathbf{p}_{i-1})) \leqslant N\delta$. Since $N \leqslant |E| = h$, it follows that $\alpha \leqslant \delta h$. 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 3.1), due to its best response is bounded by δh. □

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 $d(\mathcal{G}, \hat{\mathcal{G}}) = 0$. Also, it follows from our proof that our result is applicable in cases where the underlying game has better response cycles. Thus, even when the game does not share similar ordinal properties with potential games, our approach can be used to approximately characterize the limiting behavior of dynamics.

4. Logit response dynamics

In this section we focus on logit response dynamics. Logit response dynamics can be viewed as a smoothened 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 4.1, 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 4.2, 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.

4.1. Properties of logit response

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

Definition 4.1. 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 $P^m_\tau(q^m|\mathbf{r}) = \frac{e^{\frac{1}{\tau}u^m(q^m,\mathbf{r}^{-m})}}{\sum_{p^m\in E^m} e^{\frac{1}{\tau}u^m(p^m,\mathbf{r}^{-m})}}$.

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 $\tau \to 0$, 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 (Alós-Ferrer and Netzer, 2010; Marden and Shamma, 2008). 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 $\mathbf{p} \to \mathbf{q}$):

$$P_\tau(\mathbf{p}\to\mathbf{q}) = \begin{cases} \alpha_m P^m_\tau(q^m|\mathbf{p}) & \text{if } \mathbf{q}^{-m} = \mathbf{p}^{-m} \text{ for some } m \in \mathcal{M}, \\ 0 & \text{otherwise.} \end{cases} \tag{2}$$

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 $\lim_{\tau\to 0}\mu_\tau(\mathbf{q}) > 0$ 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 (Alós-Ferrer and Netzer, 2010; Blume, 1997; Marden and Shamma, 2008)[9]:

$$\mu_\tau(\mathbf{q}) = \frac{e^{\frac{1}{\tau}\phi(\mathbf{q})}}{\sum_{\mathbf{p}\in E} e^{\frac{1}{\tau}\phi(\mathbf{p})}}. \tag{3}$$

It can be seen from (3) that $\lim_{\tau\to 0}\mu_\tau(\mathbf{q}) > 0$ 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 $w_\mathbf{p} \geqslant 0$ such that $w_\mathbf{p} = \sum_{T\in\mathcal{T}(\mathbf{p})}\prod_{(\mathbf{q}\to\mathbf{r})\in T} P(\mathbf{q}\to\mathbf{r})$. The following proposition from the Markov Chain literature (Anantharam and Tsoucas, 1989; Freidlin and Wentzell, 1998; Leighton and Rivest, 1983), known as the Markov chain tree theorem, expresses the stationary distribution of Markov chains in terms of these weights.

Proposition 4.1. *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 $\prod_{(\mathbf{q}\to\mathbf{r})\in T} P(\mathbf{q}\to\mathbf{r})$ 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 4.1 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.

4.2. Stationary distribution of logit response dynamics

We start this section by showing that in games with small MPD logit response dynamics have similar transition probabilities.

Lemma 4.1. *Consider a game $\mathcal{G}$ and let $\hat{\mathcal{G}}$ be a nearby potential game such that $d(\mathcal{G},\hat{\mathcal{G}}) \leqslant \delta$. Denote the transition probability matrices of logit response dynamics in $\mathcal{G}$ and $\hat{\mathcal{G}}$ by P_τ and $\hat{P}_\tau$ respectively. For all strategy profiles $\mathbf{p}$ and $\mathbf{q}$ that differ in the strategy of at most one player, we have*

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

[9] Note 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.

Proof. Assume that $\mathbf{p}^{-m} = \mathbf{q}^{-m}$. In $\mathcal{G}$ the transition probability $P_\tau(\mathbf{p} \to \mathbf{q})$ can be expressed by (see (2)):

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

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

Observe that by the definition of MPD

$$\begin{aligned} u^m\big(r^m, \mathbf{p}^{-m}\big) - u^m\big(p^m, \mathbf{p}^{-m}\big) - \delta &\leqslant \hat{u}^m\big(r^m, \mathbf{p}^{-m}\big) - \hat{u}^m\big(p^m, \mathbf{p}^{-m}\big) \\ &\leqslant u^m\big(r^m, \mathbf{p}^{-m}\big) - u^m\big(p^m, \mathbf{p}^{-m}\big) + \delta. \end{aligned} \tag{4}$$

Definition 4.1 suggests that $\hat{P}_\tau^m(q^m|\mathbf{p})$ can be written as (by dividing the numerator and the denominator by $e^{\frac{1}{\tau}\hat{u}^m(p^m,\mathbf{p}^{-m})}$):

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

Therefore, using the bounds in (4) it follows that $\hat{P}_\tau^m(q^m|\mathbf{p}) \leqslant \frac{\kappa(q^m)e^{\frac{\delta}{\tau}}}{\kappa(q^m)e^{\frac{\delta}{\tau}} + \sum_{r^m \neq q^m} \kappa(r^m)e^{\frac{-\delta}{\tau}}}$, 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 $\hat{P}_\tau^m(q^m|\mathbf{p}) \leqslant \frac{e^{\frac{\delta}{\tau}} P_\tau^m(q^m|\mathbf{p})}{e^{\frac{\delta}{\tau}} P_\tau^m(q^m|\mathbf{p}) + e^{-\frac{\delta}{\tau}}(1 - P_\tau^m(q^m|\mathbf{p}))}$, or equivalently

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

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 $\hat{P}_\tau^m(q^m|\mathbf{p})/P_\tau^m(q^m|\mathbf{p}) \leqslant e^{\frac{2\delta}{\tau}}$. By symmetry we also conclude that $P_\tau^m(q^m|\mathbf{p})/\hat{P}_\tau^m(q^m|\mathbf{p}) \leqslant e^{\frac{2\delta}{\tau}}$, and combining these bounds the claim follows. □

Definition 4.1 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 (Cho and Meyer, 2001; Haviv and Van der Heyden, 1984). 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.[10]

Theorem 4.1. *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 $\alpha \geqslant 1$ is a given constant and for any two states $\mathbf{p}$ and $\mathbf{q}$, the following inequalities hold: $\alpha^{-1} P(\mathbf{p} \to \mathbf{q}) \leqslant \hat{P}(\mathbf{p} \to \mathbf{q}) \leqslant \alpha P(\mathbf{p} \to \mathbf{q})$. Then, for any state $\mathbf{p}$, we have*

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

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

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 4.1, for the Markov chain with probability transition matrix P, we have $\mu(\mathbf{p}) = \frac{w_\mathbf{p}}{\sum_\mathbf{q} w_\mathbf{q}}$,

[10] A 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.

where for each state $\mathbf{p}$, $w_{\mathbf{p}} = \sum_{T\in\mathcal{T}(\mathbf{p})} \prod_{(\mathbf{x}\to\mathbf{y})\in T} P(\mathbf{x}\to\mathbf{y})$. 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

$$\alpha^{-(h-1)} w_{\mathbf{p}} = \alpha^{-(h-1)} \sum_{T\in\mathcal{T}(\mathbf{p})} \prod_{(\mathbf{x}\to\mathbf{y})\in T} P(\mathbf{x}\to\mathbf{y}) \leqslant \hat{w}_{\mathbf{p}} = \sum_{T\in\mathcal{T}(\mathbf{p})} \prod_{(\mathbf{x}\to\mathbf{y})\in T} \hat{P}(\mathbf{x}\to\mathbf{y})$$
$$\leqslant \alpha^{h-1} \sum_{T\in\mathcal{T}(\mathbf{p})} \prod_{(\mathbf{x}\to\mathbf{y})\in T} P(\mathbf{x}\to\mathbf{y}) = \alpha^{h-1} w_{\mathbf{p}}.$$

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

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

Dividing the numerators and denominators of the left- and right-hand sides of the inequality by $\sum_{\mathbf{q}} w_{\mathbf{q}}$, using Proposition 4.1, 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

$$\frac{1-\alpha^{h-1}}{1+\alpha^{h-1}} \leqslant g\big(\mu(\mathbf{p})\big) = \frac{\alpha^{-(h-1)} \mu(\mathbf{p})}{\alpha^{-(h-1)} \mu(\mathbf{p}) + \alpha^{h-1}(1-\mu(\mathbf{p}))} - \mu(\mathbf{p}) \leqslant \hat{\mu}(\mathbf{p}) - \mu(\mathbf{p})$$
$$\leqslant \frac{\alpha^{h-1} \mu(\mathbf{p})}{\alpha^{h-1} \mu(\mathbf{p}) + \alpha^{-(h-1)}(1-\mu(\mathbf{p}))} - \mu(\mathbf{p}) = f\big(\mu(\mathbf{p})\big) \leqslant \frac{\alpha^{h-1}-1}{\alpha^{h-1}+1},$$

hence the second part of the claim follows. □

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

Corollary 4.1. *Let $\mathcal{G}$ and $\hat{\mathcal{G}}$ be finite games with number of strategy profiles $|E| = h$, such that $d(\mathcal{G}, \hat{\mathcal{G}}) \leqslant \delta$. 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*

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

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

Proof. Proof follows from Lemma 4.1 and Theorem 4.1 by setting $\alpha = e^{\frac{2\delta}{\tau}}$. □

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 (3)). We conclude this section by providing such a characterization of the stationary distribution of logit response dynamics in near-potential games.

Corollary 4.2. *Consider a game $\mathcal{G}$ and let $\hat{\mathcal{G}}$ be a nearby potential game such that $d(\mathcal{G}, \hat{\mathcal{G}}) \leqslant \delta$. 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*

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

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

Proof. Proof follows from Corollary 4.1 and (3). □

With simple manipulations, it can be shown that $(e^x - 1)/(e^x + 1) \leqslant x/2$ for $x \geqslant 0$. Thus, (ii) in the above corollary implies that $|\mu_\tau(\mathbf{p}) - \frac{e^{\frac{1}{\tau}\phi(\mathbf{p})}}{\sum_{\mathbf{q}\in E} e^{\frac{1}{\tau}\phi(\mathbf{q})}}| \leqslant \frac{\delta(h-1)}{\tau}$. 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 $\delta \to 0$, 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 $\tau \to 0$, the upper bound in (ii) becomes uninformative. This is the case since $\tau \to 0$ 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 4.2, as we show in Corollary 4.3.

Corollary 4.3. *Consider a game $\mathcal{G}$ and let $\hat{\mathcal{G}}$ be a nearby potential game with potential function ϕ and $d(\mathcal{G}, \hat{\mathcal{G}}) \leqslant \delta$. Denote the potential function of $\hat{\mathcal{G}}$ by ϕ, and the number of strategy profiles in these games by $|E| = h$. The stochastically stable strategy profiles of $\mathcal{G}$ are* (i) *contained in $S = \{\mathbf{p} | \phi(\mathbf{p}) \geqslant \max_{\mathbf{q}} \phi(\mathbf{q}) - 4\delta(h-1)\}$,* (ii) *$4\delta h$-equilibria of $\mathcal{G}$.*

Proof. (i) The upper bound in the first part of Corollary 4.2 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 $\mu_\tau(\mathbf{p}) \to 0$ as $\tau \to 0$. Thus, it immediately follows that the stochastically stable states in $\mathcal{G}$ are contained in $\{\mathbf{p} \in E | \phi(\mathbf{p}) \geqslant \max_{\mathbf{q}\in E} \phi(\mathbf{q}) - 4\delta(h-1)\}$.

(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 $d(\mathcal{G}, \hat{\mathcal{G}}) \leqslant \delta$ 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 $4\delta(h-1) + \delta \leqslant 4\delta h$. Hence stochastically stable strategy profiles of $\mathcal{G}$ are $4\delta h$-equilibria. □

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. Since it is possible to identify a potential game that is close to a given game (as explained in Section 2), Corollaries 4.2 and 4.3 provide a systematic approach for characterizing the stationary distribution and stochastically stable states of logit response, for general games. The characterization is tighter for near-potential games, but it is still informative for general games. Moreover, our results enable robust predictions about stochastically stable strategy profiles in potential games. In particular, we can quantify payoff perturbations that maintain stochastically stable states of a game. For instance, consider potential games where the potential ϕ has a unique maximizer $\mathbf{q}^*$. Corollary 4.3 implies that in such games if the payoffs are perturbed by at most $\frac{1}{8h}(\phi(\mathbf{q}^*) - \max_{\mathbf{q}\neq\mathbf{q}^*} \phi(\mathbf{q}))$ (so that the MPD between the original game and the game obtained after perturbations satisfies $\delta \leqslant \frac{1}{4h}(\phi(\mathbf{q}^*) - \max_{\mathbf{q}\neq\mathbf{q}^*} \phi(\mathbf{q}))$) the stochastically stable strategy profiles do not change. Thus, the corollary allows us to obtain a bound on the payoff perturbations that leave the stochastically stable strategy profiles intact.

5. 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 5.1. In Section 5.2 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.

5.1. Mixed strategies and equilibria

In this section, we introduce some additional notation and definitions, which will be used in Section 5.2 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 mth 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.,

$$u^m(\mathbf{x}) = \sum_{\mathbf{p}\in E} u^m(\mathbf{p}) \prod_{k\in\mathcal{M}} x^k(p^k), \tag{5}$$

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 $p^m \in E^m$,

$$u^m(p^m, \mathbf{x}^{-m}) - u^m(x^m, \mathbf{x}^{-m}) \leqslant \epsilon. \tag{6}$$

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 $\mathcal{X}_\epsilon$ 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 Candogan et al. (2011b).

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 5.1.

(i) *Let $\nu : \prod_{m\in\mathcal{M}} E^m \to \mathbb{R}$ 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 $\alpha \geqslant 0$ 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 (5), 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 (6). 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 5.1 *(Upper semicontinuous function).* A function $g : X \to Y \subset \mathbb{R}$ 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 $\limsup_{x_n \to x_*} g(x_n) \leqslant g(x_*)$ for every x_* in its domain.

Definition 5.2 *(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.[11]

[11] Here 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).

Lemma 5.2.

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

5.2. 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 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 $t\in\{1,2,\ldots\}$. 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 5.3 *(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 $p_t^m\in\arg\max_{q^m\in E^m}u^m(q^m,\mu_t^{-m})$ 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:

$$\mu_{t+1}=\frac{t}{t+1}\mu_t+\frac{1}{t+1}I_t, \tag{7}$$

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 (7), 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

$$\|\mu_{t+1}-\mu_t\|=\frac{1}{t+1}\|I_t-\mu_t\|=O\left(\frac{1}{t}\right), \tag{8}$$

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 $|f(x)|\leqslant c|g(x)|$ for all $x\geqslant x_0$.

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.[12]

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

Lemma 5.3. *Consider a game $\mathcal{G}$ and let $\hat{\mathcal{G}}$ be a close potential game such that $d(\mathcal{G},\hat{\mathcal{G}}) \leqslant \delta$. 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,*

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

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

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

Observing from (7) that $\mu_{t+1} - \mu_t = \frac{1}{t+1}(I_t - \mu_t)$, and noting from (8) that $\|\mu_{t+1} - \mu_t\| = O(\frac{1}{t})$ the above equality can be rewritten as

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

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

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

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

$$\phi(\mu_{T+1}) \geqslant \phi(\mu_T) + \frac{1}{T+1}\sum_{m\in\mathcal{M}} \left(u^m\left(p^m_T, \mu^{-m}_T\right) - u^m\left(\mu^m_T, \mu^{-m}_T\right) - \delta\right) + O\left(\frac{1}{T^2}\right). \tag{9}$$

By definition of the fictitious play dynamics, every player m plays its best response to μ^{-m}_T, therefore $u^m(p^m_T, \mu^{-m}_T) - u^m(\mu^m_T, \mu^{-m}_T) \geqslant 0$ for all m. Additionally, if μ_T is outside the ϵ-equilibrium set, as in the statement of the lemma, then it follows that $u^m(p^m_T, \mu^{-m}_T) - u^m(\mu^m_T, \mu^{-m}_T) \geqslant \epsilon$ for at least one player. Therefore, (9) implies $\phi(\mu_{T+1}) \geqslant \phi(\mu_T) + \frac{\epsilon - M\delta}{T+1} + O(\frac{1}{T^2})$, hence, the claim follows. □

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 $\inf_{\mathbf{x}\in S} \|\mu_t - \mathbf{x}\| \to 0$ as $t \to \infty$.

Theorem 5.1. *Consider a game $\mathcal{G}$ and let $\hat{\mathcal{G}}$ be a close potential game such that $d(\mathcal{G},\hat{\mathcal{G}}) \leqslant \delta$. 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 $\mathcal{X}_\alpha$ denote the α-equilibrium set of $\mathcal{G}$. For any $\epsilon > 0$, there exists a time instant $T_\epsilon > 0$ such that for all $t > T_\epsilon$*

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

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 5.3 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:

$$\phi(\mu_{t+1}) - \phi(\mu_t) \geqslant \frac{M\delta + \epsilon' - M\delta}{(t+1)} + O\left(\frac{1}{t^2}\right) > \frac{\epsilon'}{2(t+1)} > 0. \tag{10}$$

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, (10) holds for all $t = \{T_m + 1, \ldots\}$, and summing both sides of this inequality over this set we obtain $\limsup_{t\to\infty} \phi(\mu_{t+1}) - \phi(\mu_{T_m+1}) \geqslant \sum_{t=T_m+1}^{\infty} \frac{\epsilon'}{2(t+1)}$. 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 5.1(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 (8), 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

$$\mu_{t+1} \in \mathcal{X}_{M\delta+\epsilon}. \tag{11}$$

Let $\mu_t \in C_\epsilon$ for some time instant $t > T''$. If $\mu_t \in \mathcal{X}_{M\delta+\epsilon'}$, then by (11) $\mu_{t+1} \in \mathcal{X}_{M\delta+\epsilon} \subset C_\epsilon$. If, on the other hand, $\mu_t \in C_\epsilon - \mathcal{X}_{M\delta+\epsilon'}$, then by (10) and the definition of C_ϵ we have

$$\phi(\mu_{t+1}) > \phi(\mu_t) \geqslant \min_{\mathbf{y}\in\mathcal{X}_{M\delta+\epsilon}} \phi(\mathbf{y}), \tag{12}$$

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. □

The above theorem establishes that after finite time μ_t is contained in the set C_ϵ for any $\epsilon > 0$. Corollary 5.1, establishes that in the limit this result can be strengthened: as $t \to \infty$, μ_t converges to a set, which is a subset of C_ϵ for every $\epsilon > 0$. The proof can be found in Appendix A.

Corollary 5.1. *The empirical frequencies of discrete-time fictitious play converge to*

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

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 5.1 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 5.1 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. In particular, we make use of the fact that for small ϵ, the ϵ-equilibrium set is contained in a small neighborhood of equilibria (see Candogan et al., 2011b for a detailed discussion).

It was established in Lemma 5.3 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 5.2(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 5.2, 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).

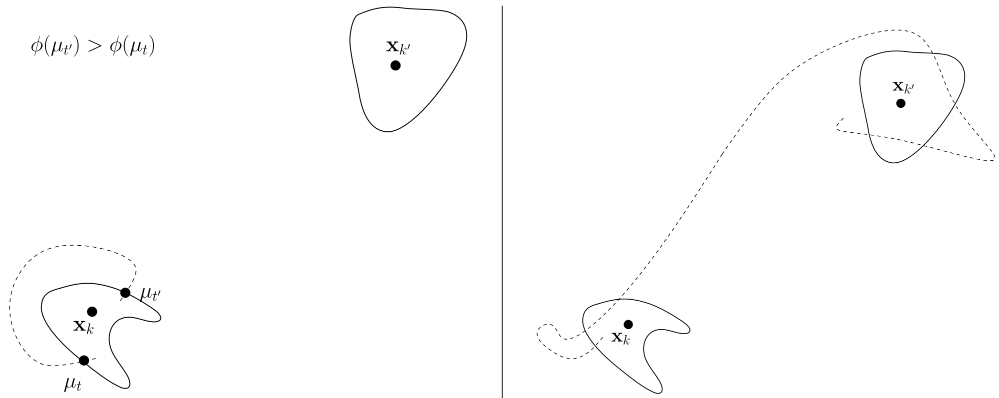

(a) If empirical frequencies leave an approximate equilibrium set at time t, and return back to it at t', then $\phi(\mu_{t'}) > \phi(\mu_t)$.

(b) For sufficiently large t, μ_t visits the component of the approximate equilibrium set contained in the neighborhood of a single equilibrium.

Fig. 4. 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).

When stating our result, we make use of the Lipschitz continuity of the mixed extension of the potential function, as established in Lemma 5.1. We also make use of a function $f:\mathbb{R}_+ \to \mathbb{R}_+$, which quantifies the size of the neighborhood of equilibria which contains the approximate equilibrium sets of games. For a game $\mathcal{G}$ with l equilibria, denoted by $\mathbf{x}_1, \ldots, \mathbf{x}_l$, this function can be formally defined as follows:

$$f(\alpha) = \max_{\mathbf{x}\in\mathcal{X}_\alpha} \min_{k\in\{1,\ldots,l\}} \|\mathbf{x} - \mathbf{x}_k\|, \tag{13}$$

for all $\alpha \in \mathbb{R}_+$. Note that $\min_{k\in\{1,\ldots,l\}} \|\mathbf{x} - \mathbf{x}_k\|$ is continuous in $\mathbf{x}$, since it is minimum of finitely many continuous functions. Moreover, $\mathcal{X}_\alpha$ is a compact set, since ϵ-equilibria are defined by finitely many inequality constraints of the form (6). Therefore, in (13) maximum is achieved and f is well-defined for all $\alpha \geqslant 0$.

Additionally, we define two variables, (a, d), which characterize the approximate equilibrium sets of the underlying game $\mathcal{G}$: (i) the minimum pairwise distance between the equilibria is denoted by $d \triangleq \min_{i\neq j} \|\mathbf{x}_i - \mathbf{x}_j\|$, (ii) $a \triangleq \sup\{\alpha \,|\, f(\alpha) < d/4\} > 0$, i.e., for every $\alpha < a$, the α-equilibrium is at most $d/4$ distant from an equilibrium of $\mathcal{G}$. Next, using these definitions, we state an improved convergence result for fictitious play in near-potential games.

Theorem 5.2. *Consider a game $\mathcal{G}$ and let $\hat{\mathcal{G}}$ be a close potential game such that $d(\mathcal{G}, \hat{\mathcal{G}}) \leqslant \delta$. 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.*

(i) *There exists some $\bar{\delta} > 0$, and $\bar{\epsilon} > 0$ satisfying*

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

(ii) *Consider any $\bar{\delta} > 0$, and $\bar{\epsilon} > 0$ satisfying* (i)*. Provided that $\bar{\delta} \geqslant \delta \geqslant 0$, it can be established that the empirical frequencies of fictitious play converge to*

$$\left\{\mathbf{x} \,\middle|\, \|\mathbf{x} - \mathbf{x}_k\| \leqslant \frac{4f(M\delta)ML}{\epsilon} + f(M\delta + \epsilon),\ \text{for some equilibrium } \mathbf{x}_k\right\}, \tag{14}$$

for any ϵ, such that $\bar{\epsilon} \geqslant \epsilon > 0$.

The proof of this theorem can be found in Appendix A. It has three main steps illustrated in Figs. 4 and 5. 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 5.3 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 Fig. 4a). 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 Fig. 4b). This

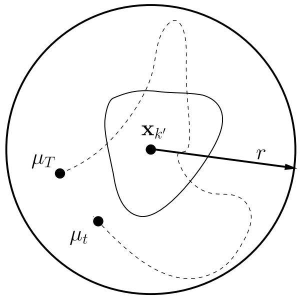

Fig. 5. If after time T, μ_t only visits the approximate equilibrium set in the neighborhood of a single equilibrium $\mathbf{x}_{k'}$, then we can establish that μ_t never leaves a neighborhood of this equilibrium for $t > T$. The size of this neighborhood is denoted by r in the figure and can be expressed as in Theorem 5.2.

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 Fig. 5). In Eq. (14) of Theorem 5.2, we provide bounds on this neighborhood, as a function of δ (that characterizes the "closeness" of the original game to a potential game), and f (that captures how the size of the ϵ-equilibrium sets increase, as a function of ϵ).

Observe that if $\delta = 0$, i.e., the original game is a potential game, then $f(M\delta) = 0$, and Theorem 5.2 implies that empirical frequencies of fictitious play converge to the $f(\epsilon)$-neighborhood of equilibria for any ϵ such that $\bar{\epsilon} \geqslant \epsilon > 0$. Thus, choosing ϵ arbitrarily small, and observing that $\lim_{x\to 0} f(x) = 0$, our result implies that in potential games, empirical frequencies converge to the set of Nash equilibria. Hence, as a special case of Theorem 5.2, 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 $\lim_{x\to 0} f(x) = 0$. Consequently, $\frac{4f(M\delta)ML}{\epsilon} + f(M\delta + \epsilon)$ is small, and Theorem 5.2 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 5.2 establishes convergence of empirical frequencies to a small neighborhood of equilibria.

A strand of the literature characterizes the limiting behavior of discrete-time fictitious play by exploiting its relation to a continuous time update rule (see for instance Benaïm et al., 2005). This framework, can be used to establish convergence of fictitious play to an equilibrium in potential games. Additionally, using this framework, it is also possible to establish upper semicontinuity of the limiting behavior of certain update rules to payoff perturbations. Theorem 5.2 can be seen as a stronger version of these upper semicontinuity results in the special setting of near-potential games. It implies upper semicontinuity of the limiting set (with respect to payoff perturbations), and also provides bounds on the size of this set, as a function of the equilibrium set of the underlying game.

Corollary 5.1 and Theorem 5.2 give a systematic framework for approximately characterizing the limiting behavior of fictitious play in arbitrary games. Moreover, such a characterization can be obtained even in settings where the underlying game does not share similar ordinal properties to potential games. Following a similar argument as in the case of logit response dynamics, our result also allows for characterizing robustness of convergence results for potential games to payoff perturbations.

Remark. So far we have explained how dynamical properties of potential games can be extended to near-potential games. In the cases of better/best response dynamics and fictitious play, it is possible to obtain similar results for games that are close to weighted potential games. To see this, assume that a game with utilities $\{u^m\}$ is given, and this game is close to a weighted potential game with weights $\{w^m \geqslant 1\}$, utilities $\{\hat{u}^m\}$ and potential ϕ. It follows from the definitions of better/best response and fictitious play dynamics that these update rules are invariant under positive scalings of the payoff functions of players. Hence, for these update rules, games with payoffs $\{u^m\}$ and $\{u^m/w^m\}$ have identical limiting sets of (mixed) strategy profiles. On the other hand, from the definition of weighted potential games, it follows that $\{\hat{u}^m/w^m\}$ is a potential game with potential function ϕ. Moreover, since $\{u^m\}$ and $\{\hat{u}^m\}$ are close in terms of their MPD, and $w_m \geqslant 1$, it follows that $\{u^m/w^m\}$ is close to the game with utilities $\{\hat{u}^m/w^m\}$. Thus, the limiting behavior of dynamics in the game with payoffs $\{u^m/w^m\}$ can be characterized using the distance to the nearby potential game $\{\hat{u}^m/w^m\}$, and the results present in Sections 3 and 5. Since for better/best response dynamics and fictitious play, the dynamical properties of $\{u^m/w^m\}$ and $\{u^m\}$ are identical, we conclude that our results immediately provide bounds on the limiting sets of the game with utilities $\{u^m\}$ through a weighted potential game approximation.

We next show that this approximation may give tighter bounds. Assume that the MPD between the game with utilities $\{u^m\}$ and $\{\hat{u}^m\}$ is δ_w. It follows that the MPD between $\{u^m/w^m\}$ and $\{\hat{u}^m/w^m\}$ is bounded by $\max_k \delta_w/w^k = \delta_w$ (without loss of generality, in weighted potential games the smallest weight corresponding to a player can be set to 1). Thus, using the results of Section 3, we conclude that better response dynamics converge to $\delta_w|E|$-equilibrium set of $\{u^m/w^m\}$. Note that this implies that in $\{u^m\}$ better response dynamics converge to $\max_k w_k \delta_w |E|$-equilibrium set. Hence, when all weights are chosen equal to $w_m = 1$ (i.e., when approximation with an exact potential game is considered), this result is identical to Theorem 3.1. Minimizing, $\max_k w_k \delta_w |E|$ over best possible weights, and weighted potential game approximations, a tighter characterization of the limiting set can be obtained. For instance when the original game is a weighted potential game, it follows that for some weights $\{w_m\}$ and weighed potential game $\{\hat{u}^m\}$ we have $\delta_w = 0$, and weighted potential game approximation establishes convergence of dynamics to a Nash equilibrium. Similar bounds can be obtained for the limiting set of fictitious play dynamics.

On the other hand, in the analysis of logit response dynamics, we use approximations with exact potential games rather than weighted potential games since weighted potential games do not have similar appealing convergence properties under this dynamics (e.g., a maximizer of the potential function need not be a stochastically stable strategy profile of logit response dynamics for weighted potential games, Alós-Ferrer and Netzer, 2010).

6. Conclusions

In this paper, we study convergence behavior of discrete-time update processes in near-potential games. introduced by Candogan et al. (2010b, 2011a).[13] 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 2, 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. The first direction that we are exploring involves characterizing the limiting behavior of dynamic processes, where players adhere to different (heterogeneous) update rules (e.g., logit response with different τ parameters), using techniques similar to the ones in this paper. 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. We established in Candogan et al. (2010a) that in some cases simple pricing mechanisms can ensure convergence to desirable equilibria in near-potential games. It is an interesting research direction to extend such mechanisms to more general game-theoretic settings. Finally, other classes of games (such as zero-sum games and supermodular games) admit appealing convergence properties under adaptive dynamics (Milgrom and Roberts, 1990; Shamma and Arslan, 2004). This motivates the question whether a game that is close to a zero-sum game or a supermodular game, still inherits some of the dynamical properties of the original game. Recent related work has considered games with local/monotone potential functions (Morris and Ui, 2005), and additional supermodularity properties, and established that the limiting behavior of dynamics can be characterized in terms of the maximizers of the potential function (Okada and Tercieux, 2008; Oyama et al., 2008). Our goal in future work is to study dynamics in games that are close to supermodular games in the sense defined in this paper.

Acknowledgments

We thank Prof. Sergiu Hart, two anonymous referees and the associate editor for their useful comments and suggestions. This research is partially supported by AFOSR grant #FA9550-09-1-0538 and NSF ECCS-1027922.

Appendix A. Proofs of Section 5

Proof of Corollary 5.1. 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

$$S_\theta = \left\{ \mathbf{x} \in \prod_{m \in \mathcal{M}} \Delta E^m \,\middle|\, \min_{\mathbf{y} \in C} \|\mathbf{x} - \mathbf{y}\| < \theta \right\}, \tag{A.1}$$

is such that $C_{\epsilon_n} \subset S_\theta$ for some n. Note that if this claim holds, then it follows from Theorem 5.1 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.1), this implies

$$\limsup_{t \to \infty} \min_{\mathbf{x} \in C} \|\mathbf{x} - \mu_t\| < \theta. \tag{A.2}$$

[13] See Candogan et al. (2013) for a characterization of the convergence behavior of continuous time update rules in near-potential games.

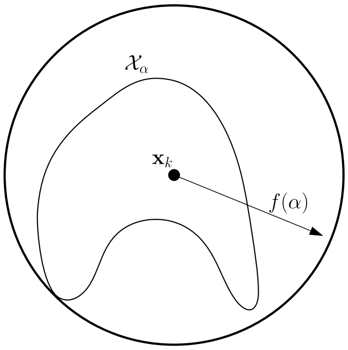

Fig. A.6. Consider a game with a unique equilibrium $\mathbf{x}_k$. The α-equilibrium set of the game (enclosed by a solid line around $\mathbf{x}_k$) is contained in the $f(\alpha)$ neighborhood of this equilibrium.

Moreover, since $\theta > 0$ is arbitrary, and $\|\mathbf{x} - \mu_t\| \geqslant 0$, using (A.2) we obtain $\lim_{t\to\infty} \min_{\mathbf{x}\in C} \|\mathbf{x} - \mu_t\| = 0$. 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 5.2(i) (by setting $\nu = \phi$ and $v = -\min_{\mathbf{y}\in\mathcal{X}_{M\delta}} \phi(\mathbf{y})$) and definition of upper semicontinuity (Definition 5.2) that there exists $\gamma > 0$ such that θ neighborhood of $\{\mathbf{x}|\phi(\mathbf{x}) \geqslant \min_{\mathbf{y}\in\mathcal{X}_{M\delta}} \phi(\mathbf{y})\}$ contains $\{\mathbf{x}|\phi(\mathbf{x}) \geqslant \min_{\mathbf{y}\in\mathcal{X}_{M\delta}} \phi(\mathbf{y}) - \gamma\}$. Hence, for any $\mathbf{z}$ satisfying $\phi(\mathbf{z}) \geqslant \min_{\mathbf{y}\in\mathcal{X}_{M\delta}} \phi(\mathbf{y}) - \gamma$ there exists some $\mathbf{x}$ satisfying $\phi(\mathbf{x}) \geqslant \min_{\mathbf{y}\in\mathcal{X}_{M\delta}} \phi(\mathbf{y})$ 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 $\phi(\mathbf{z}) \geqslant \min_{\mathbf{y}\in\mathcal{X}_{M\delta}} \phi(\mathbf{y}) - \gamma$ 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 5.1(i), and define $\theta' = \gamma/L > 0$. Lemma 5.2(ii) and Definition 5.2 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 5.1(i), it follows that $\phi(\mathbf{y}) \geqslant \phi(\mathbf{x}) - L\theta' = \phi(\mathbf{x}) - \gamma$. Thus, we conclude that there exists large enough n such that

$$\min_{\mathbf{y}\in\mathcal{X}_{M\delta+1/n}} \phi(\mathbf{y}) \geqslant \min_{\mathbf{y}\in\mathcal{X}_{M\delta}} \phi(\mathbf{y}) - \gamma. \tag{A.3}$$

Let $\mathbf{z} \in C_{\epsilon_n}$ for some n for which (A.3) holds. By definition of C_ϵ it follows that $\phi(\mathbf{z}) \geqslant \min_{\mathbf{y}\in\mathcal{X}_{M\delta+1/n}} \phi(\mathbf{y})$. Thus, (A.3) implies that $\phi(\mathbf{z}) \geqslant \min_{\mathbf{y}\in\mathcal{X}_{M\delta}} \phi(\mathbf{y}) - \gamma$. 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. □

Proof of Theorem 5.2. 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 Fig. A.6.

As stated in the theorem statement, we define the minimum pairwise distance between the equilibria as $d \triangleq \min_{i\neq j} \|\mathbf{x}_i - \mathbf{x}_j\|$, and $a = \sup\{\alpha | f(\alpha) < d/4\}$. Lemma 5.2(ii) implies (using upper semicontinuity at 0) that $\alpha > 0$ such that $f(\alpha) < d/4$ exists and hence $a > 0$. Since d is defined as the minimum pairwise distance between the equilibria, it follows that α-equilibria of the game are contained in disjoint $f(\alpha) < d/4$ neighborhoods around equilibria of the game (for $\alpha < a$), i.e., if $\mathbf{x} \in \mathcal{X}_\alpha$, then $\|\mathbf{x} - \mathbf{x}_k\| \leqslant f(\alpha)$ for exactly one equilibrium $\mathbf{x}_k$. Moreover, for $\alpha_1 \leqslant \alpha$, since $\mathcal{X}_{\alpha_1} \subset \mathcal{X}_\alpha$, it follows that α_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 establish existence of $\bar{\delta}$ 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) $f(x) \to 0$ as $x \to 0$.
- *Step 2*: We show that there exists some $\bar{\delta} > 0$ and $\bar{\epsilon} > 0$ such that the following inequalities hold:

$$M\bar{\delta} + \bar{\epsilon} < a, \tag{A.4}$$

and

$$f(M\bar{\delta} + \bar{\epsilon}) < \frac{(a - M\bar{\delta})d}{24LM}. \tag{A.5}$$

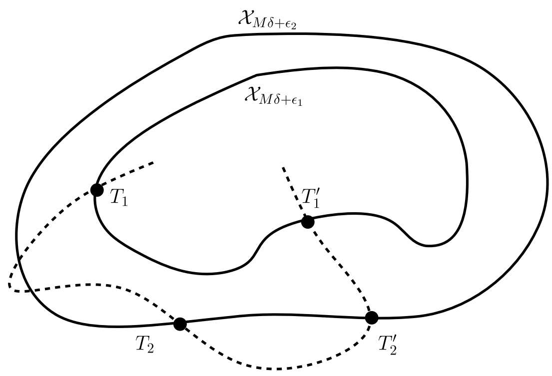

Fig. A.7. Trajectory of μ_t (initialized at the left end of the dashed line) is illustrated. T_1 and T_2 correspond to the time instants μ_t leaves $\mathcal{X}_{M\delta+\epsilon_1}$ and $\mathcal{X}_{M\delta+\epsilon_2}$ respectively. T_1' and T_2' correspond to the time instants μ_t enters $\mathcal{X}_{M\delta+\epsilon_1}$ and $\mathcal{X}_{M\delta+\epsilon_2}$ respectively.

We will prove the statement of the theorem assuming that $0 \leqslant \delta < \bar{\delta}$, and establish convergence to the set in (14), for any ϵ such that $0 < \epsilon \leqslant \bar{\epsilon}$. As can be seen from the definition of a and f (see (13)), 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.4) and (A.5) to establish that the empirical frequency distribution μ_t can visit the component of $\mathcal{X}_{M\delta+\bar{\epsilon}}$ 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 Fig. A.7, 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.4) 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.4) and (A.5). 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 $0 < \epsilon_1 < \epsilon_2 \leqslant \bar{\epsilon}$. 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 Fig. A.7). 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 $\alpha_1 \leqslant \alpha$. Since the feasible set of the maximization problem in (13) is given by $\mathcal{X}_\alpha$, this implies that $f(\alpha_1) \leqslant f(\alpha)$, i.e., f is a weakly increasing function of its argument. Note that the feasible set of the maximization problem in (13) can be given by the correspondence $g(\alpha) = \mathcal{X}_\alpha$, which is upper semicontinuous in α as shown in Lemma 5.2(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 $\alpha \geqslant 0$, $f(\alpha)$ 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 $f(x) \leqslant \epsilon$ for all $x \in V$. Since, $f(x) \geqslant 0$ by definition, this implies that $\lim_{x\to 0} f(x)$ exists and equals to 0.

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

$$\begin{aligned} \phi(\mu_{t+1}) - \phi(\mu_t) &\geqslant \frac{2\epsilon_1}{3(t+1)} \quad \text{if } \mu_t \notin \mathcal{X}_{M\delta+\epsilon_1}, \quad \text{and similarly} \\ \phi(\mu_{t+1}) - \phi(\mu_t) &\geqslant \frac{2\epsilon_2}{3(t+1)} \quad \text{if } \mu_t \notin \mathcal{X}_{M\delta+\epsilon_2}. \end{aligned} \tag{A.6}$$

Existence of T satisfying these inequalities follows from Lemma 5.3, since for large T and $t > T$, this lemma implies $\phi(\mu_{t+1}) - \phi(\mu_t) \geqslant \frac{\epsilon_1}{(t+1)} + O(\frac{1}{t^2}) \geqslant \frac{2\epsilon_1}{3(t+1)}$ 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.6) 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 $T < T_1 \leqslant T_2 < T_2' \leqslant T_1'$, 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 $T_1 \leqslant t < T_1'$.
- 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 $T_2 \leqslant t < T_2'$.
- 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 Fig. A.7. 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.

By definition, for t such that $T_2 \leqslant t < T_2'$, we have $\mu_t \notin \mathcal{X}_{M\delta+\epsilon_2}$, and for t such that $T_1 \leqslant t < T_2$ or $T_2' \leqslant t < T_1'$, we have $\mu_t \notin \mathcal{X}_{M\delta+\epsilon_1}$. Thus, it follows from (A.6) that

$$\phi(\mu_{t+1}) - \phi(\mu_t) \geqslant \frac{2\epsilon_2}{3(t+1)} \quad \text{for } T_2 \leqslant t < T_2', \tag{A.7}$$

and consequently,

$$\phi(\mu_{T_2'}) - \phi(\mu_{T_2}) = \sum_{t=T_2}^{T_2'-1} \phi(\mu_{t+1}) - \phi(\mu_t) \geqslant \sum_{t=T_2}^{T_2'-1} \frac{2\epsilon_2}{3(t+1)}. \tag{A.8}$$

Similarly, since $\mu_t \notin \mathcal{X}_{M\delta+\epsilon_1}$ for t such that $T_1 \leqslant t < T_2$ or $T_2' \leqslant t < T_1'$, using (A.6) we establish

$$\phi(\mu_{T_1'}) - \phi(\mu_{T_2'}) = \sum_{t=T_2'}^{T_1'-1} \phi(\mu_{t+1}) - \phi(\mu_t) \geqslant \sum_{t=T_2'}^{T_1'-1} \frac{2\epsilon_1}{3(t+1)}, \tag{A.9}$$

$$\phi(\mu_{T_2}) - \phi(\mu_{T_1}) = \sum_{t=T_1}^{T_2-1} \phi(\mu_{t+1}) - \phi(\mu_t) \geqslant \sum_{t=T_1}^{T_2-1} \frac{2\epsilon_1}{3(t+1)}. \tag{A.10}$$

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.8), (A.9) and (A.10) that

$$\phi(\mu_{T_1'})-\phi(\mu_{T_1}) \geqslant \sum_{t=T_2}^{T_2'-1} \frac{2\epsilon_2}{3(t+1)}. \tag{A.11}$$

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, 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 $\|\mathbf{x}_k - \mathbf{x}\| \leqslant f(M\delta+\epsilon_1)$ (similarly if $\mathbf{x} \in \mathcal{X}_{M\delta+\epsilon_2}$, then $\|\mathbf{x}_k - \mathbf{x}\| \leqslant f(M\delta+\epsilon_2)$) 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.,

$$\|\mathbf{x}_k - \mu_{T_1-1}\| \leqslant f(M\delta+\epsilon_1) \quad \text{and} \quad \|\mathbf{x}_k - \mu_{T_2-1}\| \leqslant f(M\delta+\epsilon_2), \quad \text{whereas} \tag{A.12}$$

$$\|\mathbf{x}_{k'} - \mu_{T_1'}\| \leqslant f(M\delta+\epsilon_1) \quad \text{and} \quad \|\mathbf{x}_{k'} - \mu_{T_2'}\| \leqslant f(M\delta+\epsilon_2). \tag{A.13}$$

By definition of d we have $\|\mathbf{x}_k - \mathbf{x}_{k'}\| \geqslant d$. Since $a > M\delta+\epsilon_2$, it follows that $f(M\delta+\epsilon_2) < d/4$, and hence the second inequalities in (A.12) and (A.13) imply

$$\|\mu_{T_2'} - \mu_{T_2-1}\| > \frac{d}{2}. \tag{A.14}$$

Using this inequality, we next refine the lower bound on $\phi(\mu_{T_1'})-\phi(\mu_{T_1})$ obtained in Step 3. By (7), with an update at time t, the empirical frequency distribution can change by at most

$$\|\mu_{t+1} - \mu_t\| = \frac{1}{t+1}\|\mu_t - I_t\| \leqslant \frac{1}{t+1}\big(\|\mu_t\| + \|I_t\|\big) \leqslant \frac{2M}{t+1}, \tag{A.15}$$

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 $\|\mu_t^m\|, \|I_t^m\| \leqslant 1$, 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.14) implies $\|\mu_{T_2'} - \mu_{T_2}\| > \frac{d}{2}$. Using this together with (A.15), we conclude

$$\sum_{t=T_2}^{T_2'-1} \frac{2M}{t+1} \geqslant \sum_{t=T_2}^{T_2'-1} \|\mu_{t+1} - \mu_t\| \geqslant \Big\|\sum_{t=T_2}^{T_2'-1} \mu_{t+1} - \mu_t\Big\| = \|\mu_{T_2'} - \mu_{T_2}\| > \frac{d}{2}. \tag{A.16}$$

Thus, the lower bound on $\phi(\mu_{T_1'})-\phi(\mu_{T_1})$ provided in (A.11) takes the following form:

$$\phi(\mu_{T_1'})-\phi(\mu_{T_1}) \geqslant \sum_{t=T_2}^{T_2'-1} \frac{2\epsilon_2}{3(t+1)} \geqslant \frac{\epsilon_2 d}{6M}. \tag{A.17}$$

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 $\overline{\phi}_k = \max_{\{\mathbf{x} \mid \|\mathbf{x}-\mathbf{x}_k\| \leqslant f(M\delta+\epsilon_1)\}} \phi(\mathbf{x})$, and define $\mathbf{y}_k$ as a strategy profile which achieves this maximum. Similarly, let $\underline{\phi}_{k'} = \min_{\{\mathbf{x} \mid \|\mathbf{x}-\mathbf{x}_{k'}\| \leqslant f(M\delta+\epsilon_1)\}} \phi(\mathbf{x})$ and define $\mathbf{y}_{k'}$ as a strategy profile which achieves this minimum. Observe that

$$\underline{\phi}_{k'} - \overline{\phi}_k = \phi(\mathbf{y}_{k'}) - \phi(\mathbf{y}_k) = \big(\phi(\mathbf{y}_{k'}) - \phi(\mu_{T_1'})\big) + \big(\phi(\mu_{T_1'}) - \phi(\mu_{T_1})\big) + \big(\phi(\mu_{T_1}) - \phi(\mathbf{y}_k)\big). \tag{A.18}$$

Note that by (A.12) and (A.13), and the definitions of $\mathbf{y}_k$ and $\mathbf{y}_{k'}$, we have $\mu_{T_1'}, \mathbf{y}_{k'} \in \{\mathbf{x} \mid \|\mathbf{x} - \mathbf{x}_{k'}\| \leqslant f(M\delta+\epsilon_1)\}$, and $\mu_{T_1-1}, \mathbf{y}_k \in \{\mathbf{x} \mid \|\mathbf{x} - \mathbf{x}_k\| \leqslant f(M\delta+\epsilon_1)\}$. Hence, using Lipschitz continuity of ϕ (and denoting the Lipschitz constant by L) it follows that $\phi(\mathbf{y}_{k'}) - \phi(\mu_{T_1'}) \geqslant -2Lf(M\delta+\epsilon_1)$, and $\phi(\mu_{T_1-1}) - \phi(\mathbf{y}_k) \geqslant -2Lf(M\delta+\epsilon_1)$. Moreover, (A.15) and Lipschitz continuity of ϕ imply that $\phi(\mu_{T_1}) - \phi(\mu_{T_1-1}) = O(\frac{1}{T_1})$. Thus, using (A.18) we obtain the following upper bound on $\phi(\mu_{T_1'}) - \phi(\mu_{T_1})$:

$$\underline{\phi}_{k'} - \overline{\phi}_k + 4Lf(M\delta+\epsilon_1) + O\left(\frac{1}{T_1}\right) \geqslant \phi(\mu_{T_1'}) - \phi(\mu_{T_1}). \tag{A.19}$$

Using the lower and upper bounds we obtained in (A.17) and (A.19), it follows that

$$\underline{\phi}_{k'} - \bar{\phi}_k + 4Lf(M\delta + \epsilon_1) + O\left(\frac{1}{T_1}\right) \geqslant \frac{\epsilon_2 d}{6M}. \tag{A.20}$$

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.20) that

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

Note that (A.5) 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 $0 < \epsilon_1 < \epsilon_2 \leqslant \bar{\epsilon}$. 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 $\epsilon_1 < \epsilon_2 \leqslant \bar{\epsilon}$, 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 $T + 1 < T_2 \leqslant t \leqslant T_2' - 1$, i.e.,

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

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., $2d^* \leqslant \sum_{t=T_2-1}^{T_2'-1} \|\mu_{t+1} - \mu_t\|$. As explained in (A.15), $\|\mu_{t+1} - \mu_t\| \leqslant \frac{2M}{t+1}$, thus the above inequality implies

$$2d^* \leqslant \sum_{t=T_2-1}^{T_2'-1} \frac{2M}{t+1} = \sum_{t=T_2}^{T_2'-1} \frac{2M}{t+1} + \frac{2M}{T_2}. \tag{A.21}$$

Using this inequality, the lower bound in (A.11) implies

$$\phi(\mu_{T_1'}) - \phi(\mu_{T_1}) \geqslant \sum_{t=T_2}^{T_2'-1} \frac{2\epsilon_2}{3(t+1)} \geqslant \left(d^* - \frac{M}{T_2}\right)\frac{2\epsilon_2}{3M}. \tag{A.22}$$

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 $\mu_{T_1-1}, \mu_{T_1'} \in \{\mathbf{x} \mid \|\mathbf{x} - \mathbf{x}_k\| \leqslant f(M\delta + \epsilon_1)\}$. By Lipschitz continuity of the potential function it follows that $\phi(\mu_{T_1'}) - \phi(\mu_{T_1-1}) \leqslant 2f(M\delta + \epsilon_1)L$. Additionally, by (A.15) Lipschitz continuity also implies that $\phi(\mu_{T_1}) - \phi(\mu_{T_1-1}) \leqslant \frac{2ML}{T_1}$. Combining these we obtain the following upper bound on $\phi(\mu_{T_1'}) - \phi(\mu_{T_1})$:

$$\phi(\mu_{T_1'}) - \phi(\mu_{T_1}) \leqslant 2f(M\delta + \epsilon_1)L + \frac{2ML}{T_1}. \tag{A.23}$$

It follows from the upper and lower bounds on $\phi(\mu_{T_1'}) - \phi(\mu_{T_1})$ given in (A.22) and (A.23) that $(d^* - \frac{M}{T_2})\frac{2\epsilon_2}{3M} \leqslant 2f(M\delta + \epsilon_1)L + \frac{2ML}{T_1}$. Thus, for sufficiently large T_1 (and hence T_2), we obtain

$$d^* \leqslant \frac{3f(M\delta + \epsilon_1)ML}{\epsilon_2} + \frac{3M^2L}{\epsilon_2 T_1} + \frac{M}{T_2} \leqslant \frac{4f(M\delta + \epsilon_1)ML}{\epsilon_2}. \tag{A.24}$$

Note that in the above derivation ϵ_1 is an arbitrary number that satisfies $0 < \epsilon_1 < \epsilon_2$. Thus, (A.24) implies that

$$d^* \leqslant \limsup_{\epsilon_1 \to 0} \frac{4f(M\delta + \epsilon_1)ML}{\epsilon_2} \leqslant \frac{4f(M\delta)ML}{\epsilon_2}, \tag{A.25}$$

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.25), we can establish convergence of μ_t to $\{\mathbf{x} \mid \|\mathbf{x}-\mathbf{y}\| \leqslant \frac{4f(M\delta)ML}{\epsilon_2}$, for some $\mathbf{y} \in \mathcal{X}_{M\delta+\epsilon_2}\}$. 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 $\|\mathbf{x}_k-\mathbf{y}\| \leqslant f(M\delta+\epsilon_2)$. Thus, using triangle inequality, we conclude that μ_t converges to

$$\left\{\mathbf{x} \,\middle|\, \|\mathbf{x}-\mathbf{x}_k\| \leqslant \frac{4f(M\delta)ML}{\epsilon_2} + f(M\delta+\epsilon_2),\ \text{for some equilibrium } \mathbf{x}_k\right\}. \tag{A.26}$$

Noting that ϵ_2 is an arbitrary number satisfying $0 < \epsilon_2 \leqslant \bar{\epsilon}$, the theorem follows. □

References

Alós-Ferrer, C., Netzer, N., 2010. The logit-response dynamics. Games Econ. Behav. 68 (2), 413–427.
Anantharam, V., Tsoucas, P., 1989. A proof of the Markov chain tree theorem. Stat. Probab. Lett. 8 (2), 189–192.
Benaïm, M., Hofbauer, J., Sorin, S., 2005. Stochastic approximations and differential inclusions. SIAM J. Control Optim. 44 (1), 328–348.
Berge, C., 1963. Topological Spaces. Oliver & Boyd.
Blume, L., 1993. The statistical mechanics of strategic interaction. Games Econ. Behav. 5 (3), 387–424. 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. In: Activity Analysis of Production and Allocation, vol. 13 (1), pp. 374–376.
Candogan, O., Menache, I., Ozdaglar, A., Parrilo, P., 2010a. Near-optimal power control in wireless networks: A potential game approach. In: Proceedings IEEE. INFOCOM, March 2010, pp. 1–9.
Candogan, O., Ozdaglar, A., Parrilo, P., 2010b. A projection framework for near-potential games. In: 49th IEEE Conference on Decision and Control. CDC, pp. 244–249.
Candogan, O., Menache, I., Ozdaglar, A., Parrilo, P.A., 2011a. Flows and decompositions of games: Harmonic and potential games. Math. Operations Res. 36 (3), 474–503.
Candogan, O., Ozdaglar, A., Parrilo, P., 2011b. Dynamics in games and near potential games. Tech. Rep. LIDS, MIT.
Candogan, O., Ozdaglar, A., Parrilo, P.A., 2011c. Learning in near-potential games. In: 50th IEEE Conference on Decision and Control and European Control Conference. CDC–ECC, pp. 2428–2433.
Candogan, O., Ozdaglar, A., Parrilo, P.A., 2013. Near-potential games: Geometry and dynamics. ACM Trans. Econ. Comput. 1 (2), Article No. 11, 32 pp. http://dx.doi.org/10.1145/2465769.2465776.
Cho, G., Meyer, C., 2001. Comparison of perturbation bounds for the stationary distribution of a Markov chain. Linear Algebra Appl. 335 (1–3), 137–150.
Dubey, P., Haimanko, O., Zapechelnyuk, A., 2006. Strategic complements and substitutes, and potential games. Games Econ. Behav. 54 (1), 77–94. http://www.sciencedirect.com/science/article/pii/S089982560400171X.
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. Adv. Appl. Probab. 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 Econ. Behav. 5 (3), 368–386.
Krishna, V., Sjöström, T., 1998. On the convergence of fictitious play. Math. Operations Res. 23 (2), 479–511.
Leighton, F., Rivest, R., 1983. The Markov chain tree theorem. Tech. Rep. MIT/LCS/TM-249. MIT, Laboratory for Computer Science.
Li, D., Zhang, X., 2002. On dynamical properties of general dynamical systems and differential inclusions. J. Math. Anal. Appl. 274 (2), 705–724. http://www.sciencedirect.com/science/article/pii/S0022247X02003529.
Marden, J., Shamma, J., 2008. Revisiting log-linear learning: Asynchrony, completeness and a payoff-based implementation. Under submission.
Marden, J., Arslan, G., Shamma, J., 2009. Joint strategy fictitious play with inertia for potential games. IEEE Trans. Automat. Control 54 (2), 208–220.
Milgrom, P., Roberts, J., 1990. Rationalizability, learning, and equilibrium in games with strategic complementarities. Econometrica 58 (6), 1255–1277.
Monderer, D., Shapley, L., 1996a. Fictitious play property for games with identical interests. J. Econ. Theory 68 (1), 258–265.
Monderer, D., Shapley, L., 1996b. Potential games. Games Econ. Behav. 14 (1), 124–143.
Morris, S., Ui, T., 2004. Best response equivalence. Games Econ. Behav. 49 (2), 260–287. http://www.sciencedirect.com/science/article/pii/S0899825604000132.
Morris, S., Ui, T., 2005. Generalized potentials and robust sets of equilibria. J. Econ. Theory 124 (1), 45–78.
Okada, D., Tercieux, O., 2008. Log-linear dynamics and local potential. Departmental Working Papers.
Oyama, D., Takahashi, S., Hofbauer, J., 2008. Monotone methods for equilibrium selection under perfect foresight dynamics. Theoretical Econ. 3 (2).
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 Trans. Automat. Control 49 (7), 1137–1141.
Shapley, L., 1964. Some topics in two-person games. Adv. Game Theory 52, 1–29.
Uno, H., 2007. Nested potential games. Econ. Bull. 3 (17), 1–8.
Voorneveld, M., 2000. Best-response potential games. Econ. Letters 66 (3), 289–295. http://www.sciencedirect.com/science/article/pii/S0165176599001962.
Young, H., 1993. The evolution of conventions. Econometrica 61 (1), 57–84.
Young, H., 2004. Strategic Learning and Its Limits. Oxford University Press, USA.