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
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.