Available online at www.sciencedirect.com

ScienceDirect

Journal of Economic Theory 162 (2016) 21–54

JOURNAL OF Economic Theory

www.elsevier.com/locate/jet

Convergence of best-response dynamics in extensive-form games ☆

Zibo Xu

Pillar of Engineering Systems and Design, Singapore University of Technology and Design, 8 Somapah Road, Singapore

Received 10 March 2015; final version received 22 November 2015; accepted 4 December 2015

Available online 10 December 2015

Abstract

This paper presents a collection of convergence results on best-response dynamics in extensive-form games. We prove that in all finite generic extensive-form games of perfect information, every solution trajectory to the continuous-time best-response dynamic converges to a Nash equilibrium component. We show the robustness of this convergence in the sense that along any interior approximate best-response trajectory, the evolving state is close to the set of Nash equilibria most of the time. We also prove that in any perfect-information game in which every play contains at most one decision node of each player, any interior approximate best-response trajectory converges to the backward-induction strategy profile. Our final result concerns self-confirming equilibria in perfect-information games. If each player always best responds to her conjecture of the current strategy profile, and she updates her conjecture based only on observed moves, then the dynamic will converge to the set of self-confirming equilibria.

JEL classification: C73; D83

Keywords: Convergence; Extensive form; Perfect information; Backward-induction strategy profile; Best-response dynamics; Self-confirming equilibrium

☆ The first version of the paper was completed at Stockholm School of Economics in June 2013. The paper was revised while the author was visiting the Department of Economics and Decision Sciences at HEC Paris in 2014. The current version was completed after the author moved to the Pillar of Engineering Systems and Design at Singapore University of Technology and Design.

E-mail address: zibo_xu@sutd.edu.sg.

http://dx.doi.org/10.1016/j.jet.2015.12.001

0022-0531/

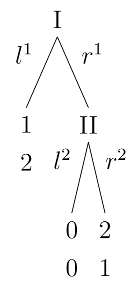

Fig. 1. The best-response dynamic for a two-node game.

1. Introduction

The notion of best response lies at the heart of Nash equilibrium (Nash, 1950): each player selects a best response to the other players' strategies and no player can benefit by unilaterally changing her strategy. If we consider a dynamic setting where each player always chooses a best response to the current strategies of the other players, does this dynamic always converge to a Nash equilibrium?

We adopt the definition of a continuous-time best-response dynamic formulated by a constant revision rate and myopic optimization, as found in Matsui (1989), Gilboa and Matsui (1991), Hofbauer (1995), and Balkenborg et al. (2013). In such a dynamic, a state specifies the strategy profile of all players; the frequency of a strategy increases only if it is a best response to the current state. The continuous-time best-response dynamic has been analyzed in various classes of games in normal form (or strategic form); see Hofbauer and Sigmund (1998) and Sandholm (2010). In particular, the convergence of a continuous-time best-response dynamic to a Nash equilibrium has been shown in Harris (1998), Hofbauer (1995), and Hofbauer and Sorin (2006) for two-player zero-sum games, in Harris (1998) for weighted-potential games, and in Berger (2005) for $2 \times n$ games.

In the present paper, we consider extensive-form games of perfect information. In addition to its explicit representation of players' sequential moves, each perfect-information game admits a backward-induction solution, and backward induction itself is a process to determine a finite sequence of local best-response moves; see Selten (1965, 1975). The present paper studies the predictive power of the best-response dynamic for an extensive-form game of prefect information. In particular, we focus on the convergence results of the continuous-time best-response dynamic concerning the backward-induction strategy profile (or so-called subgame-perfect equilibrium) and the Nash equilibrium components in an extensive-form game.

In a standard evolutionary setting, for each player there is a continuum population of individuals playing the game in that role. The populations are distinct, and each individual plays a pure strategy of the corresponding player. A (population) state specifies the population share of individuals who are programmed to each pure strategy of each player. The payoffs represent fitness of population shares. Even without showing the rigorous model, we can obtain some intuition by analyzing the following classic example of the two-player game of Fig. 1 in the evolutionary setting.

This game has two Nash equilibria in pure strategies (r^1, r^2) and (l^1, l^2), where the first one is the backward-induction equilibrium. Suppose that in the initial state, everyone in population I is playing l^1 and the share of population II playing l^2 is at least $1/2$. Then, on the one hand, pop-