Available online at www.sciencedirect.com
ScienceDirect
ELSEVIER Journal of Economic Dynamics & Control 32 (2008) 1569–1599
JOURNAL OF Economic Dynamics & Control
www.elsevier.com/locate/jedc

A numerical analysis of the evolutionary stability of learning rules

Jens Josephson*

Department of Economics and Business, Universitat Pompeu Fabra, Ramon Trias Fargas 25-27, 08005 Barcelona, Spain

Received 10 November 2005; accepted 22 June 2007
Available online 4 July 2007

Abstract

In this paper, we define an evolutionary stability criterion for learning rules. Using simulations, we then apply this criterion to three types of symmetric 2×2 games for a class of learning rules that can be represented by the parametric model of Camerer and Ho [1999. Experience-weighted attraction learning in normal form games. Econometrica 67, 827–874]. This class contains stochastic versions of reinforcement and fictitious play as extreme cases. We find that only learning rules with high or intermediate levels of hypothetical reinforcement are evolutionarily stable, but that the stable parameters depend on the game.

JEL classification: C72; C73; C15

Keywords: Monte Carlo simulation; Evolutionary stability; Learning in games; Fictitious play; Reinforcement learning; EWA learning

1. Introduction

The bounded rationality paradigm is based on the assumption that people learn to play games by using simple rules of adaptation, often referred to as *learning rules*. In almost all of the literature, it is assumed that all players of the game employ the same

*Tel.: +34 93 542 22 68; fax: +34 93 542 17 46.
E-mail address: jens.josephson@upf.edu

0165-1889/$ - see front matter
doi:10.1016/j.jedc.2007.06.008

type of learning rule. Moreover, the type of rule is generally motivated by its asymptotic properties in such a homogeneous setting, such as convergence to Nash equilibrium. This seems questionable from an evolutionary perspective. Evolutionary selection should favor the use of learning rules that perform well in a heterogeneous setting with mutant learning rules present in the population. Furthermore, since most individual interactions are of limited length, survival should to a great extent be determined by the payoffs to using the learning rule in the short or intermediate run in such a heterogeneous setting.

In this paper, we implement these ideas by defining an evolutionary stability criterion for learning rules and applying this criterion to a set of well-known learning rules using Monte Carlo simulations. More specifically, we ask if there is a rule such that, if applied by a homogeneous population of individuals, it will survive any sufficiently small invasion of mutants using a different rule. We call such an uninvadable rule an *evolutionarily stable learning rule* (ESLR). This concept is an extension of Taylor and Jonker's (1978) definition of evolutionarily stable strategies (ESSs) (a concept originally due to Maynard Smith and Price, 1973; Maynard Smith, 1974) to learning rules and behavioral strategies.

The setting is a world where the members of a large population, consisting of an even number of individuals, in each of a finite number of periods are all randomly matched in pairs to play a finite two-player game. Each individual uses a learning rule, which is a function of his private history of past play, and fitness is measured in terms of expected average payoff. This framework provides a rationale for the use of learning rules and it is of particular interest since very little analysis of learning in this 'repeated rematching' context has previously been done (see Hopkins, 1999, for an exception).

Technically, learning rules are mappings from the history of past play to the set of pure or mixed strategies. There are many models of learning and we therefore restrict the numerical analysis to a class of learning rules that can be described by the general parametric model of Camerer and Ho (1999), called *experience-weighted attraction (EWA) learning*. The rules in this class have experimental support and perform well in an environment where the game changes from time to time.[1] Moreover, the class contains rules which differ considerably in their use of information. Two of the most well-known learning rules, *reinforcement learning* and *fictitious play*, are special cases of this model for specific parameter values.

Reinforcement learning is an important model in the psychological literature on individual learning. It was introduced by Bush and Mosteller (1951) and further developed by Erev and Roth (1998), although the principle behind the model, that choices which have led to good outcomes in the past are more likely to be repeated in the future, originally is due to Thorndike (1898). Under reinforcement learning in games, players assign probability distributions to their available pure strategies. If a pure strategy is employed in a particular period, the probability of the same pure strategy being used in the subsequent period increases as a function of the realized payoff. The model has very low information and rationality requirements in the

[1] See, for example, Camerer and Ho (1999) and Stahl (2003).