Journal of Economic Behavior & Organization
Vol. 45 (2001) 19–35

JOURNAL OF
Economic Behavior
& Organization

www.elsevier.com/locate/econbase

Population rule learning in symmetric normal-form games: theory and evidence

Dale O. Stahl*

Department of Economics, University of Texas, Austin, TX 78712, USA

Received 22 January 1999; received in revised form 31 August 1999; accepted 31 August 1999

Abstract

A model of population rule learning is formulated and estimated using experimental data. When predicting the population distribution of choices and accounting for the number of parameters, the population rule learning model is much better than aggregation of individually estimated rule learning models. Further, rule learning is a statistically significant and important phenomena even when focusing on population statistics, and is much better than one-rule learning dynamics. © 2001 Elsevier Science B.V. All rights reserved.

JEL classification: C15; C52; C72

Keywords: Rules; Learning; Games; Experimental; Testing

1. Introduction

Recent learning research in one-shot games can be divided into two domains: (i) population learning or evolutionary dynamics as typified by replicator dynamics,[1] and (ii) individual learning.[2,3] The first domain focuses on how the population distribution of play changes over time, while the second domain focuses on how an individual's behavior changes over time.

Individualistic models are needed for investigating the nature and characteristics of individual learning patterns and for assessing the amount of diversity in the population. Further,

* Tel.: +1-512-475-8541; fax: +1-512-571-3510.
E-mail address: stahl@eco.utexas.edu (D.O. Stahl).

[1] For example, Hofbauer and Sigmund (1988), Van Huyck et al. (1994), and Cheung and Friedman (1997).

[2] For example, Cheung and Friedman (1997, 1998), Crawford (1994), Mookherjee and Sopher (1994), Cooper and Feltovich (1996), Camerer and Ho (1999a, b), Rapoport et al. (1997), and Stahl (1996, 1997, 1999, 2000).

[3] Studies that examine both include Bush and Mosteller (1955), Friedman et al. (1998), and Roth and Erev (1995).

PII: S0167-2681(00)00169-4

if one wants to construct a cognitive theory of individual behavior in games, then individualistic models are essential because individual details could be masked in population statistics.

From a decision-theoretic framework, however, for one-shot games it is necessary and sufficient for a player to have a belief about the other players' actions, and when the other players are randomly drawn from a population of potential players such a belief is equivalent to a forecast of the population distribution of other players' actions. It is neither necessary nor sufficient to know anything about a single individual's learning dynamics, since one's actual opponents are random draws from a population. For example, to know which side of the road to drive on in the US, I do not need to know any specific history about the driver approaching me on the highway; I only need to know that in the US all sober drivers stay on the right side of the road.

Ideally, as in general equilibrium economics, one would like a theory of individual learning that aggregates up to a theory of population learning. However, we will encounter similar difficulties in finding aggregation theorems with reasonable assumptions. Of course, we can estimate individualistic models and then aggregate. But there is only so much information in any given dataset. If it is used to estimate a multitude of parameters of individualistic models, it does not follow that the prediction following aggregation is better than a prediction from a population (or representative agent) model with far fewer parameters. We will address this pertinent empirical question.

We focus on the class of rule learning models of Stahl (1996, 1997, 1999, 2000) (hereafter S96, S97a,b, and S99). This is a rich class of learning models that encompasses action reinforcement (Roth and Erev, 1995; Erev and Roth, 1998), fictitious play (Brown, 1951), and belief updating (Mookherjee and Sopher, 1994; Camerer and Ho, 1999a, b).[4] Briefly, a "rule" is a mapping from the game and history of play to a mixed strategy. For example, a noisy best response to the recent past is a Cournot-like rule that describes much of the behavior observed in experiments. Iterating once more we have a "level-2" rule that is a noisy best response to the best response to the recent past.

Complicating the econometric estimation of rule learning models is the fact that the rule used by an individual is not directly observable — only the action taken is observable — and in any model with properly specified error structures all rules will have full support on the available undominated actions. In an individualistic model of rule learning (S97b), the posterior probability of the rule conditional on the history was computed, but the computational complexity necessitated the use of precarious approximations. This problem can be potentially avoided by a population learning model because the experience of many individuals using and evaluating different rules gets merged into the population experience, so in essence it is as if the population evaluates all the rules.

In Section 2 we review the individual rule learning model of S97a, spell out aggregation of that model, and develop a population version of rule learning. Section 3 describes the experimental design and data, and Section 4 describes the econometric specification and computational issues. Section 5 presents the results, and Section 6 discusses our findings.

[4] While Camerer and Ho consider both reinforcement learning and belief learning, they formulate a single hybrid rule that combines these two aspects rather than allowing both types of rules to exist simultaneously in the population. In contrast, our rule learning model allows for many rules to exist simultaneously in the population and in the minds of players.