Games and Economic Behavior 87 (2014) 178–203

Contents lists available at ScienceDirect

Games and Economic Behavior

www.elsevier.com/locate/geb

A generalized approach to belief learning in repeated games

Christos A. Ioannou [a], Julian Romero [b,*]

[a] Department of Economics, University of Southampton, Southampton, SO17 1BJ, United Kingdom
[b] Department of Economics, Krannert School of Management, Purdue University, Lafayette, IN 47907, United States

ARTICLE INFO

Article history:
Received 17 January 2013
Available online 24 May 2014

JEL classification:
C51
C72
C92
D03

Keywords:
Adaptive models
Belief learning
Repeated-game strategies
Finite automata
Prisoner's Dilemma
Battle of the Sexes
Stag-Hunt
Chicken

ABSTRACT

We propose a methodology that is generalizable to a broad class of repeated games in order to facilitate operability of belief-learning models with repeated-game strategies. The methodology consists of (1) a generalized repeated-game strategy space, (2) a mapping between histories and repeated-game beliefs, and (3) asynchronous updating of repeated-game strategies. We implement the proposed methodology by building on three proven action-learning models. Their predictions with repeated-game strategies are then validated with data from experiments with human subjects in four, symmetric 2×2 games: Prisoner's Dilemma, Battle of the Sexes, Stag-Hunt, and Chicken. The models with repeated-game strategies approximate subjects' behavior substantially better than their respective models with action learning. Additionally, inferred rules of behavior in the experimental data overlap with the predicted rules of behavior.

1. Introduction

The limitations of adaptive models with actions have been well recognized in the literature. For instance, Erev and Roth (1998) note that learning behavior generally cannot be analyzed in terms of actions alone (p. 872). Along similar lines, Camerer and Ho (1999) point out that actions "are not always the most natural candidates for the strategies that players learn about" (p. 871). Yet developing models with repeated-game strategies has been inhibited by several obstacles. First, since the set of possible strategies in repeated games is infinite (uncountable), expecting a player to fully explore such an infinite set is unrealistic and impractical. Second, as McKelvey and Palfrey (2001) note, "players face an inference problem going from histories to beliefs about opponents' strategies" in repeated games (p. 25). A player's beliefs about the opponent's repeated-game strategy become more complex because several different strategies can lead to the same history. Consequently, even though the history of play is publicly observed, a player may not know the opponent's precise strategy. Third, repeated-game strategies need several periods to be evaluated. Traditionally, action-learning models required that the updating of each player's action set occurs synchronously at the end of each period. This is sensible if a player uses actions, but if a player uses repeated-game strategies, then the player ought to play the stage game a number of times before assessing the payoff consequences of the repeated-game strategy chosen.

* Corresponding author.
E-mail addresses: c.ioannou@soton.ac.uk (C.A. Ioannou), jnromero@purdue.edu (J. Romero).

http://dx.doi.org/10.1016/j.geb.2014.05.007

Table 1
Proposed methodology for strategy learning.

Generalized repeated-game strategy space We propose a generalized repeated-game strategy space where players' strategies are implemented by a type of finite automaton, called a Moore machine (Moore, 1956).
Mapping between histories and repeated-game beliefs We propose a fitness function that counts the number of consecutive fits of each candidate strategy of the opponent with the observed action profile sequence, starting from the most recent action profile and going backwards. Beliefs for each candidate strategy of the opponent are derived by normalizing each strategy's respective fitness by the total number of fitness of all candidate strategies of the opponent.
Asynchronous updating of repeated-game strategies We propose that the updating of repeated-game strategies is endogenous and based on the "surprise-triggers-change" regularity identified by Erev and Haruvy (2013). Thus, a player's strategy set is updated asynchronously with the completion of a block of periods, and not, necessarily, at the end of each period.

In this study, we propose a methodology that is generalizable to a broad class of repeated games in order to facilitate operability of belief-learning models with repeated-game strategies. The methodology consists of (1) a generalized repeated-game strategy space, (2) a mapping between histories and repeated-game beliefs, and (3) asynchronous updating of repeated-game strategies. The first step in operationalizing the proposed framework is to use generalizable rules, which require a relatively small repeated-game strategy set but may implicitly encompass a much larger space (see, for instance, Stahl's rule learning in Stahl and Dale (1999) and Stahl et al. (2012)). A large number of repeated-game strategies is impractical for updating under most existing learning models because the probability of observing any particular strategy in the space is near zero. We propose, instead, a generalized repeated-game strategy space where players' strategies are implemented by a type of finite automaton, called a *Moore machine* (Moore, 1956); thus, the strategy space includes only a subset of the theoretically large set of possible strategies in repeated games. The second step establishes a mapping between histories and repeated-game beliefs. In particular, we propose a fitness function that counts the number of consecutive fits of each candidate strategy of the opponent with the observed action profile sequence, starting from the most recent action profile and going backwards. Beliefs for each candidate strategy of the opponent are derived by normalizing each strategy's respective fitness by the total fitness of all candidate strategies of the opponent. This novel approach solves the inference problem of going from histories to beliefs about opponents' strategies in a manner consistent with belief learning.[1] The third step accommodates asynchronous updating of repeated-game strategies. A player's strategy set is updated with the completion of a *block of periods*, and not necessarily at the end of each period, as traditional action-learning models require. Furthermore, the probability of updating the strategy set is endogenous and based on the "surprise-triggers-change" regularity identified by Erev and Haruvy (2013).[2] Surprise is defined as the difference between actual and expected (anticipated) payoff. Thus, if a player receives a payoff similar to what is expected, then surprise is low; hence, the probability of updating the strategy set increases by a relatively small amount. However, if a player receives a payoff that is drastically different from the one anticipated, then surprise is high; hence, the probability of updating the strategy set increases by a relatively large amount. Henceforth, for brevity, we refer to learning of repeated-game strategies as *strategy learning*. Table 1 shows the proposed methodology.

We assess the impact of the proposed methodology by building on three leading action-learning models: a self-tuning Experience Weighted Attraction model (Ho et al., 2007), a γ-Weighted Beliefs model (Cheung and Friedman, 1997), and an Inertia, Sampling and Weighting model (Erev et al., 2010). The predictions of the three models with strategy learning are validated with data from experiments with human subjects by Mathevet and Romero (2012) in four symmetric 2×2 games: Prisoner's Dilemma, Battle of the Sexes, Stag-Hunt, and Chicken. We use the experimental dataset to also validate the predictions of their respective models with action learning, which enables us to determine the improvement in fit in moving from action-learning models to strategy-learning ones. Finally, we infer rules of behavior in the experimental dataset and compare them to those predicted by the strategy-learning models.

We find that the strategy-learning models approximate subjects' behavior substantially better than their respective models with action learning. Furthermore, inferred rules of behavior in the experimental data overlap with the predicted rules of behavior. More specifically, the most prevalent rules of behavior in the experimental dataset in the Prisoner's Dilemma, Stag-Hunt, and Chicken games are the cooperative rules of behavior "Grim-Trigger" and "Tit-For-Tat." The same two rules emerge as the most prevalent in the simulations. Likewise, in the Battle of the Sexes, the same cooperative rules of behavior implementing alternations that prevail in the experimental dataset, also prevail in the simulations.

[1] Alternatively, Hanaki et al. (2005) develop a model of learning of repeated-game strategies with standard reinforcement. Reinforcement learning responds only to payoffs obtained by strategies chosen by the player and, thus, evades the inference problem highlighted above. Yet reinforcement models are most sensible when players do not know the foregone payoffs of unchosen strategies. Several studies show that providing foregone payoff information affects learning, which suggests that players do not simply reinforce chosen strategies (see Mookherjee and Sopher, 1994; Rapoport and Erev, 1998; Camerer and Ho, 1999; Costa-Gomes et al., 2001; Nyarko and Schotter, 2002 and Van Huyck et al., 2007).

[2] Erev and Haruvy (2013) observe that subjects exhibit a positive relationship (inertia) between recent and current action choices (see, also, Cooper and Kagel, 2003 and Erev and Haruvy, 2005). Yet the probability of terminating the inertia mode increases with surprise; that is, surprise triggers change.