Journal of Economic Dynamics & Control 48 (2014) 188–201

Contents lists available at ScienceDirect

Journal of Economic Dynamics & Control

journal homepage: www.elsevier.com/locate/jedc

Learning a population distribution ☆

Seung Han Yoo

Department of Economics, Korea University, Seoul 136-701, Republic of Korea

ARTICLE INFO

Article history:
Received 26 March 2013
Received in revised form
5 September 2014
Accepted 8 September 2014
Available online 21 September 2014

JEL classification:
C72
C73
D83

Keywords:
Bayesian games
Dynamic games
Bayesian learning

ABSTRACT

This paper introduces a dynamic Bayesian game with an unknown population distribution. Players do not know the true population distribution and assess it based on their *private* observations using Bayes' rule. First, we show the existence and characterization of an equilibrium in which each player's strategy is a function not only of the player's type but also of experience. Second, we show that each player's initial belief about the population distribution converges almost surely to a "correct" belief.

1. Introduction

When a (true) population distribution is not known, as opposed to the Bayesian game by Harsanyi (1967), players who are repeatedly matched have incentives to learn from others' actions about the population distribution. For example, in reality, before people decide to take a particular action and anticipate others' strategies, they try to figure out what the relevant population distributions look like. In this paper, we are interested in those game-theoretic situations and their consequences.

We assume that there is a large population, and that players from the population are distributed according to a probability distribution. It is a common knowledge that there is a true population distribution among two distributions in which one "locally" first-order stochastic dominates the other,[1] but players do not know which one is the true population distribution. In each period $t \in \mathbb{N}$, every two players are randomly paired to play a 2×2 game, and after the game, each player's belief about the population distribution is updated by Bayes' rule after observing the other's action. They are paired for just one period, and rematched again after each period. The *types* of each player are relevant for the player's own payoff,

☆ This paper is a revised version of the paper "A Dynamic Bayesian Game with an Unknown Population Distribution." I am grateful to Geir Asheim, James Bullard, Edward Green, Seungjin Han, James Jordan, Atsushi Kajii, Jeong-Yoo Kim, Andrew McLennan, Ehud Lehrer, Yaw Nyarko, Larry Samuelson, seminar participants at the Far Eastern Meeting of the Econometric Society 2008, 5th Pan-Pacific Conference on Game Theory, 4th Asian Decentralization Conference, Midwest Economic Theory Meeting 2009 and International Conference on Game Theory 2013 for helpful comments. I appreciate an anonymous associate editor and two anonymous referees of the journal for valuable comments. I would especially like to thank Fernando Vega-Redondo for helpful suggestions and encouragement in the early stage of this project. I thank Chang-Hoon Lee for research assistance on section 5. Of course, all remaining errors are mine.

E-mail address: shyoo@korea.ac.kr

[1] For a formal definition, see (A5) in Section 3.

http://dx.doi.org/10.1016/j.jedc.2014.09.010
0165-1889/

and observations in each match are *private*. Hence, the matching mechanism of the game is "public," but the information flow of the game is "private."[2]

2×2 games, although limited, provide the birth-places for important equilibrium selection concepts such as risk dominance (Harsanyi and Selten, 1988) and global games (Carlsson and van Damme, 1993). Furthermore, Bayesian games with binary actions and a threshold equilibrium have wide applications that are related to Harsanyi's (1973) purification theorem, public good provision and voting games following Palfrey and Rosenthal (1988), and global games (Carlsson and van Damme, 1993 and see Morris and Shin, 2003 for various examples with a threshold equilibrium).[3]

Since each player's monitoring is private, each player's estimate about the population distribution is also private. Hence, each player's strategy is a function of both the player's type and observations. We call player i's history of observations up to $t-1$ player i's *experience* at t, and a vector of player i's type and experience player i's *characteristic*, since a type (resp. an experience) of player i can be defined as his or her innate (resp. acquired) characteristic. Given this private type and estimate, each player anticipates what samples and data the opponent has observed in the past as well as the opponent's type.

This paper provides three main results: *existence*, *monotonicity* and *convergence*. For each period, every two-paired players play a Bayesian game with their characteristics. Since a Bayesian game even with a *one*-dimensional type space may not have an equilibrium,[4] we adopt a simple model in which a Bayesian equilibrium strategy is parameterized by a threshold type. In addition, by assuming that players have the same *initial* beliefs, we enable them to construct expectations without invoking the hierarchy of beliefs problem, which extends Harsanyi's ingenious idea (Theorem 1).

Second, the monotonicity result establishes that if given each period, a player believes that a stochastic dominant distribution is more likely, then the player's equilibrium strategy shows a certain monotonic pattern (Theorem 2).[5] In particular, if a player has an experience that generates beliefs such that a stochastic dominant distribution is more likely, then his or her optimal strategy is to make more types take an action that *were* chosen under the stochastic dominant one. Hence, given each period, the same type, which is determined in the beginning of a game, can make different decisions depending on different experiences, or sample paths. For example, if a person takes a "bad" action, it may be because the person is bad by nature or because the person has observed many bad actions on the part of others in the past.[6]

This also allows us new interpretations about experimental results, which often do not support theoretical predictions, especially, in two different ways: before an experiment, different "subjects" may have different preconceptions, given their past *private* experiences, about "objects" of the experiment, and during an experiment, subjects can learn even with a random matching; for example, Van Huyck et al. (1990) find, using minimum effort games, that in random-pair experiments, the subjects' dynamic behavior shows learning features similar to those in *fixed*-pair experiments,[7] and this also can explain theoretically why such a (high) variance in subject beliefs is observed even with random matching in Nyarko and Schotter (2002).[8]

Third, we show that each player's initial belief about the true distribution converges *almost surely* to a correct belief (Theorem 3). This implies that the limit of equilibrium strategies from a sequence of observations is equal to the equilibrium strategy with common knowledge. In other words, Harsanyi's common distribution assumption is justifiable in the long run.[9] On the other hand, in the real world, people can observe a large but only *finite* number of samples, so their experiences will influence their beliefs about population distributions, such as the above second result, which provides a different perspective on "almost common knowledge" situation studied by Rubinstein (1989).

To the author's knowledge, this is the first paper to attempt to study Bayesian learning of an unknown type distribution through players' interactions. Players' beliefs about a true population distribution are updated using Bayes' rule, and the outcome of their strategic interactions is derived as a Bayesian equilibrium, which makes this model differ from agent-based models including reinforcement learning (see Duffy, 2006 for survey on agent-based modeling and related experiments). This paper also departs from papers in Bayesian learning (for surveys, see Marimon, 1997; Fudenberg and Levine, 1998;

[2] As observed by Fudenberg and Levine (1998), this environment is most frequently used in experiments for game theory: "*Random-matching model* Each period all players are randomly matched. At the end of each round, each player observes only the play in his own match…… This is the treatment most frequently used in game theory experiments" (p. 6).

[3] Note that this model can be extended to n-person games under a certain condition. See concluding remarks.

[4] There are two main approaches to tackling the existence of a Bayesian game with general action and multidimensional type spaces. One is by McAdams (2003), extending (Athey, 2001), who suggests the single crossing condition (p. 866) for a *one*-dimensional type space, and the other is by Vives (1990) and Van Zandt and Vives (2007), who utilize supermodular *payoffs*.

[5] See Marimon (1997) for arguments that convergence alone is not sufficient to make learning theory interesting.

[6] Hence, the world one believes one knows can be just the reflection of what one perceives.

[7] Van Huyck et al. (1990) write that "Experiments six and seven randomly paired subjects with an unknown partner. Hence, experiments six and seven test whether the results obtained in experiments four and five were due to subjects repeating the period with the same opponent…. Moreover, the subjects' dynamic behavior was similar to that found in the fixed pair C treatment" (p. 244). However, they do not provide an explanation for how this occurs.

[8] Nyarko and Schotter (2002) point out that "At this point in time we have no explanation for why the volatility of beliefs did not settle down when subjects were randomly matched. As stated above, in the random matching experiments, it would make sense to treat one's opponent at any point in time as a time-average of what one has observed in one's own experience in the game." (p. 996).

[9] This reports an optimistic prediction; if people can be matched infinitely often in a random manner, the *negative* biases or stereotypes they hold about others will disappear. For instance, if members of a group can have certain beliefs about the others' characteristics, then the infinite random matching between them yields the convergence to a correct belief.