Games and Economic Behavior 87 (2014) 381–396

Contents lists available at ScienceDirect

Games and Economic Behavior

www.elsevier.com/locate/geb

Learning in society

Braz Camargo

Sao Paulo School of Economics – FGV, Rua Itapeva 474, São Paulo, SP 01332-000, Brazil

ARTICLE INFO

Article history:
Received 21 June 2011
Available online 17 June 2014

JEL classification:
C73
D82
D83

Keywords:
Multi-armed bandit
Social learning
Strategic experimentation

ABSTRACT

In an individual experimentation problem a decision maker learns only from his own experience. It is well known that an optimal experimentation strategy for such problems sometimes results in the best alternative being dropped altogether, which is the so-called "Rothschild effect." Many experimentation problems of interest, however, involve learning from both individual experience and the experience of others. This paper shows that learning in society can overcome the Rothschild effect. We consider an economy with a continuum of infinitely lived players in which each player faces a multi-armed bandit and in each period a player observes the action choice of another randomly chosen player. We show that social conformity always happens in the long run, and we use this fact to derive a condition on the distribution of prior beliefs that implies that the fraction of players who choose the best alternative always converges to one in the long run.

1. Introduction

In an experimentation problem, actions not only provide rewards but also provide information that is relevant for future decisions so that there is a trade-off between the maximization of current rewards and information acquisition. When a decision maker facing an experimentation problem can learn only from his own experience, the trade-off just described implies that an optimal experimentation strategy sometimes results in the best alternative being dropped altogether, which is the so-called "Rothschild effect."[1] Many experimentation problems of interest, however, involve individuals learning from both their own experience and the experience of others. Examples include consumers learning about product quality and doctors learning about different treatments for the same disease. This paper investigates whether social learning can overturn the Rothschild effect.

We study a discrete time economy with a continuum of infinitely lived and anonymous players in which each player faces a multi-armed bandit with independent arms. The unknown stochastic payoffs to each of the possible actions are the same for all players. The players, however, can have heterogeneous prior beliefs about the true stochastic payoffs. Besides observing the outcomes of their action choices, in every period each player observes the action choice of another randomly chosen player.[2] We refer to the latter observations as the *observations in society*.

E-mail address: braz.camargo@fgv.br.

[1] Rothschild's (1974) paper is the first one in economics to draw attention to this fact. The term "Rothschild effect" was introduced in Kihlstrom et al. (1984). Banks and Sundaram (1992) and Brezzi and Lai (2000) establish that optimal experimentation strategies can lead an individual to settle on an inferior alternative for multi-armed bandits with independent arms. Easley and Kiefer (1988) obtain the same result for a different class of experimentation problems that includes multi-armed bandits with correlated arms as a special case.

[2] It is easy to adapt our analysis to the case in which the size of the random sample that a player observes is greater than one or is itself random (but finite). We can also adapt our analysis to the case in which the points in time when a player observes the action choices of other players are random.

http://dx.doi.org/10.1016/j.geb.2014.05.014
0899-8256/

The problem the players face in our environment differs from a standard multi-armed bandit in that there are informational externalities. In fact, in each period, the likelihood that a player makes a given observation in society is determined by the behavior of the other players. Since the behavior of the other players depends (indirectly) on the true payoffs to the action choices, the observations in society reveal payoff-relevant information. Formally, our setting is a game of strategic experimentation with nonatomic players.

The first result of the paper is that regardless of the true stochastic payoffs to each of the available actions and the distribution of prior beliefs in the population, almost all players choose the same action in the long run in any equilibrium of the game. In other words, social conformity always takes place in the long run. Nevertheless, the long-run behavior of the players can be suboptimal despite the presence of social learning. We say that social learning is *efficient* if the fraction of individuals who choose the best action always converges to one. The main result of the paper is a condition on the distribution of prior beliefs that, if satisfied, implies that social learning is efficient in every equilibrium. The condition we derive requires heterogeneous prior beliefs.

It is straightforward to adapt our analysis to the case in which the observations in society include not only actions but also their outcomes. There are two reasons why we assume that only actions can be observed in society. The first reason is to emphasize that this alone provides information about payoffs. The second reason is to show that just observing action choices can be enough to overturn the Rothschild effect. Besides, the assumption that outcomes are observable is not always adequate in many settings that involve social learning. For instance, when consumers learn about product quality, they typically observe only the choices of other consumers, not the outcomes of these choices.

The rest of the paper is organized as follows. The remainder of this section discusses the related literature. Section 2 presents the main ideas of the paper in the context of a simple example. Section 3 introduces the model. Section 4 studies the experimentation problem the players face when they take the behavior of the other players as given. Section 5 establishes the result on long-run social conformity. Section 6 studies efficient social learning. Section 7 discusses some of our modeling choices and addresses the issue of the speed of social conformity. Section 8 concludes. The appendices contain omitted proofs and details, as well as a discussion about the existence of equilibria.

1.1. Related literature

This paper bridges the gap between the literatures on social learning and strategic experimentation. The literature on social learning usually considers environments in which agents are either boundedly rational or myopic. For instance, Smallwood and Conlisk (1979), Ellison and Fudenberg (1993, 1995), Schlag (1998), and Juang (2001) consider models of social learning in which agents follow simple exogenously given rules of behavior. Bala and Goyal (1998) and Gale and Kariv (2003) study social learning in networks with myopic agents. We consider an environment in which individuals are Bayesian and forward-looking.[3] Our work also differs from the literature on strategic experimentation in that while this literature is typically concerned with characterizing equilibrium behavior in some game of strategic experimentation, see, for example, Bolton and Harris (1999), Keller et al. (2005), and Rosenberg et al. (2007), the focus of this article is on the long-run behavior of agents.

The papers most closely related to ours are Aoyagi (1998), Bala and Goyal (1998), and Rosenberg et al. (2009).[4] Aoyagi (1998) analyzes a finite-player game of strategic experimentation in which action choices, but not their outcomes, are public and shows that all players eventually settle on the same action in any Nash equilibrium of the game. The proof of our result on long-run social conformity adapts the reasoning in Aoyagi (1998) to our setting. Our argument is slightly more general since we allow for heterogeneous priors, while Aoyagi (1998) assumes a common prior. Rosenberg et al. (2009) establish long-run conformity of behavior in a large class of finite-player games of strategic experimentation that includes the environment in Aoyagi (1998) as a special case.[5] Rosenberg et al. (2009) also assume a common prior. Neither Aoyagi (1998) nor Rosenberg et al. (2009) are concerned with efficient social learning, which is the main focus of our analysis.

Bala and Goyal (1998) study social learning in an environment in which each agent can observe the action choices and outcomes of a fixed set of neighbors. As in Bala and Goyal (1998), we use the long-run conformity of behavior to derive a condition on the distribution of prior beliefs that ensures efficient social learning. The analysis in Bala and Goyal (1998) differs from our analysis in important ways, though. Besides being myopic, the agents in Bala and Goyal (1998) are also not Bayesian in the sense that they do not use the observed action choices of their neighbors to make inference about the true stochastic payoffs to the available actions. In our environment, the observed action choices of the other players are the only source of information a player has besides the outcomes of his own action choices.

[3] The literature on informational cascades considers models in which, unlike in our setting, agents make irreversible decisions. Banerjee (1992), Bikhchandani et al. (1992), and Smith and Sorensen (2000) consider settings in which agents move in an exogenous order. Caplin and Leahy (1994), Chamley and Gale (1994), and Peck and Yang (2011) consider models of herd behavior in which the timing of decisions is endogenous.

[4] See also Banerjee and Fudenberg (2004), who study a rational model of word-of-mouth learning in which successive generations of individuals make irreversible choices between two alternatives.

[5] Rosenberg et al. (2009) also provide examples in which social conformity does not take place in the long run. These examples are non-generic, though.