Games and Economic Behavior 104 (2017) 392–410

Contents lists available at ScienceDirect

Games and Economic Behavior

www.elsevier.com/locate/geb

Payoff externalities and social learning ☆

Itai Arieli [1]

Faculty of Industrial Engineering and Management, Technion—Israel Institute of Technology, Israel

ARTICLE INFO

Article history:
Received 5 January 2016
Available online 19 May 2017

JEL classification:
C72
C73
D80
D83

Keywords:
Social learning
Payoff externalities
Asymptotic learning

ABSTRACT

We study a social learning model with payoff externalities in which one of two state-dependent games is chosen at random and then played repeatedly by a different group of agents. Each "generation" observes the history of actions and receives conditionally independent private signals about the realized game. We show that with probability one, the play converges to the set of equilibria of an appropriate convex combination of the two state games. We provide a necessary and sufficient condition on the private signal distribution for asymptotic learning and show that in some cases asymptotic learning holds for a wide range of bounded private signals.

1. Introduction

Social learning is an integral part of social interaction. The seminal papers on social learning by Banerjee (1992) and Bikhchandani, Hirshleifer, and Welch (BHW) (1992) offer a rational explanation for herding, i.e., the tendency to neglect private information and instead act as others do. In the canonical model, countably many, fully rational agents make a decision under uncertainty while observing the choices of their predecessors. Agents learn by making inferences about the private information of their predecessors based on the latter's actions. The striking result by Banerjee (1992) and Bikhchandani et al. (1992) shows that all but a group of finitely many agents disregard their own private information and, as a result, end up herding.

Following these works, the literature has focused on characterizing sufficient conditions on the signal distribution or the observational structure such that asymptotic learning occurs (see, among others, Smith and Sørensen, 2000; Acemoglu et al., 2010; Lobel and Sadler, forthcoming). The study of payoff externalities and their implications, however, is still lacking. The main goal of this work is to integrate payoff externalities into the canonical model and study their implications on phenomena such as learning and herding.

Payoff externalities are common in many social learning interactions and often arise in various contexts. For example, an increasing number of investors purchasing a given product may result in an increase in prices.[2] In addition, the choice of software or hardware of consumers often exhibits some degree of payoff externalities that are based upon network effects.[3]

☆ I thank Nageeb Ali, Yakov Babichenko, Manuel Mueller-Frank, and Peter Sørensen for helpful comments and suggestions. I also thank two anonymous referees and an associate editor for many valuable comments.

E-mail address: iarieli@tx.technion.ac.il.

[1] The author thanks the Israel Science Foundation (grant number 2018889) for research support.

[2] See Eyster et al. (2014) for an extended discussion.

[3] See Bresnahan (2004).

http://dx.doi.org/10.1016/j.geb.2017.05.005
0899-8256/

When individuals are engaged in a strategic interaction there are two potential sources of relevant information: their own private information about the fundamentals of the game and the available public information that is achieved by observing other agents. In this work we aim to study these effects in a model that allows for payoff externalities.

We start with the canonical sequential social learning model, except that instead of having a single agent who has to decide on a binary action, we allow for multiple agents who play a state-dependent game. As in the canonical model, an unknown state of the world out of a binary set is realized in accordance with a commonly known prior probability distribution. Each state determines a game, and the two corresponding games have the same structure and may differ only in the payoff functions of the players. Thereafter a countable number of disjoint groups (generations) of agents act sequentially. Every member of each generation plays the role of a certain player in a game and has to choose an action out of a finite set. The underlying state of the world and the choices of all agents within a specific generation determine the payoffs for all agents. Every agent in each generation observes the realized actions of the previous generations, and prior to his play receives a private signal regarding the true state of the world.

This model provides a natural extension to the canonical model. Similarly, agents are ignorant of the true state of the world and act sequentially and irreversibly. The distinctions are that within every period of time, the current generation acts simultaneously and the payoff of every member of each generation is also affected by the actions that are played by the other members of his generation.

We shall characterize asymptotic equilibrium play according to two features: (i) the signal distribution, which determines how informative the private information received by the individuals is, and (ii) the payoff structure of the two possible underlying games.

Our first result shows, as one would expect, that if the public belief converges to the value p, then the play converges to an equilibrium of the game $\Gamma(p)$, which is the game that is obtained from a weighted convex combination of the two state-dependent games.

Our main contribution provides a necessary and sufficient condition on the state-dependent games and the bounds of the private signals under which agents asymptotically learn the true state. Our characterization relies on the simple logic that in order for asymptotic learning to hold the action distribution of the agents must remain informative as a function of the public belief throughout the play. This in turn holds if and only if there exists no prior probability $\mu \in [0,1]$ for which the agents can correlate their play with a unique equilibrium action of the games $\Gamma(p)$ across all possible posterior beliefs.

In order to formalize this we define the map Eq that assigns to every value p the set of all mixed Nash equilibria of the game $\Gamma(p)$. We say that Eq is *flat* on an interval $I \subseteq [0,1]$ if there exists a mixed action distribution that comprises a Nash equilibrium of the game $\Gamma(p)$ for every value $p \in I$. Asymptotic learning is shown to hold if and only if for any such interval and a prior in I there exists a positive probability that the posterior probability of the agents lies outside I. We characterize this using two parameters. One is the flatness parameter θ^* that measures the *injectivity* of the map Eq by taking into consideration all intervals I on which Eq is flat. The other is a parameter that determines the strength (or level of informativeness) of the signal distribution. Our main theorem (Theorem 2) shows that asymptotic learning holds if and only if the strength of the signal distribution exceeds the flatness parameter θ^*.

1.1. Related literature

The study of sequential social learning started with the seminal papers of Banerjee (1992) and BHW Bikhchandani et al. (1992). Variations and applications of their models have been extensively studied. A seminal contribution by Smith and Sørensen (2000) establishes a striking asymptotic learning characterization using properties of the signal distribution. They show that asymptotic learning holds in any Bayesian equilibrium if private signals have unbounded strength, but fails in any Bayesian equilibrium if the strength of the signals is bounded. This result has given rise to extensions such as Acemoglu et al. (2010) who consider a random social network, Arieli and Mueller-Frank (2015) who consider unbounded signal distribution on a general state space, and others. Similarly to these results, our contribution also relates the strength of the signal to the quality of learning, but in our case the characterization of asymptotic learning is determined also by the two state-dependent games. As we shall show, in contrast to the canonical model, asymptotic learning can hold for a wide range of bounded signals.

The class of games that we analyze in the present paper is that of *recurring games.* Like a repeated game, a recurring game is played repeatedly, but each stage is played by a new group of agents. These games were first introduced by Jackson and Kalai (1997) and serve as a natural extension to the canonical sequential social learning game. As argued by Jackson and Kalai, recurring games often arise in social contexts, and allow one to model the evolution of behavior in multigenerational games. Recurring games capture short-term strategic interactions in which players can draw inferences about the game from observing the past behavior of other players who were engaged in a similar interaction.

There are several real-life examples that can be modeled as a recurring game (see also Jackson and Kalai, 1997). Real-estate transactions occur each year but involve different participants. Students choose an academic institution based on the quality of the institution, which to some degree can be inferred from the past choices of other students. This choice, however, may be subject to payoff externalities that can arise from the current choices of the students' friends. When people decide which movie to watch, or which restaurants to dine at, they base their decision both on the preferences of their potential partners and on the experience of others. In all these examples the choices of the agents are irreversible: there are *information externalities* from past choices and payoff externalities from current choices.