Available online at www.sciencedirect.com

ScienceDirect

Journal of Economic Theory 177 (2018) 426–460

JOURNAL OF Economic Theory

www.elsevier.com/locate/jet

Social learning and the shadow of the past ☆

Yuval Heller [a,*], Erik Mohlin [b]

[a] *Department of Economics, Bar Ilan University, Israel*
[b] *Department of Economics, Lund University, Sweden*

Received 13 September 2017; final version received 25 June 2018; accepted 10 July 2018
Available online 19 July 2018

Abstract

In various environments new agents may base their decisions on observations of actions taken by a few other agents in the past. In this paper we analyze a broad class of such social learning processes, and study under what circumstances the initial behavior of the population has a lasting effect. Our results show that this question strongly depends on the expected number of actions observed by new agents. Specifically, we show that if the expected number of observed actions is: (1) less than one, then the population converges to the same behavior independently of the initial state; (2) between one and two, then in some (but not all) environments there are decision rules for which the initial state has a lasting impact on future behavior; and (3) more than two, then in all environments there is a decision rule for which the initial state has a lasting impact.

JEL classification: C73; D83

Keywords: Social learning; Steady state; Unique limiting behavior; Path dependence

☆ Old versions of this paper were previously titled "When Is Social Learning Path-Dependent?" and "Unique Stationary Behavior." We thank Ron Peretz, Doron Ravid, Satoru Takahashi, Xiangqian Yang, Peyton Young, an associate editor, and a referee for valuable discussions and helpful suggestions. Yuval Heller is grateful to the European Research Council for its financial support (ERC starting grant #677057). Erik Mohlin is grateful to Handelsbankens forskningsstiftelser (grant #P2016-0079:1), the Swedish Research Council (grant #2015-01751), and the Knut and Alice Wallenberg Foundation (Wallenberg Academy Fellowship #2016-0156) for their financial support.

* Corresponding author.
E-mail addresses: yuval.heller@biu.ac.il (Y. Heller), erik.mohlin@nek.lu.se (E. Mohlin).

https://doi.org/10.1016/j.jet.2018.07.003

1. Introduction

Agents must often make decisions without knowing the costs and benefits of the possible choices. In such situations an inexperienced (or "newborn") agent may learn from the experience of others, by basing his decision, on observations of a few actions taken by other agents in the past (see, e.g., the social learning models of Ellison and Fudenberg, 1993, 1995; Acemoglu et al., 2011). In other environments, agents interact with random opponents, and an agent may base his choice of action on a few observations of how his current opponent behaved in the past (as first described in Rosenthal, 1979, and further developed and applied to various models of community enforcement in the Prisoner's Dilemma game in Nowak and Sigmund, 1998; Takahashi, 2010; Heller and Mohlin, forthcoming).

When analyzing such dynamic situations two related important questions arise: (1) does the initial behavior of the population have a lasting influence on the population's behavior in the long run, and (2) does the model admit a unique prediction or multiple predictions for the long-run behavior? For concreteness, consider an environment in which new agents face a choice between competing technologies with positive externalities, where the unknown state of nature determines which technology is superior (see, e.g., Banerjee and Fudenberg, 2004). A central issue when analyzing this environment is to characterize in which setups the population may be "stuck" in the long run with the inferior technology, due to the arbitrary behavior of a few early adopters, and in which setups the population will always converge to the superior technology regardless of the initial behavior.

The present paper analyzes a broad class of processes in which agents obtain information by sampling the behavior of other agents, and it shows that the above two questions strongly depend on the expected number of actions observed by new agents. Specifically, we show that: (1) if the mean sample size (expected number of actions observed by a new agent) is less than one, then the population converges to the same behavior independently of the initial state; (2) if the mean sample size is between one and two, then any environment allows for a rule with multiple steady states according to which agents learn from the experience of other agents, but only some environments allow for decision rules with multiple locally stable states; and (3) if the mean sample size is more than two, then all environments admit a decision rule with multiple locally stable states, and the initial state determines which steady states will prevail. It should be noted that none of our results in any way depend on specifying how payoff considerations enter the agent's choices.[1]

Overview of the model We consider an infinite population of agents (a continuum of mass one). Time is discrete and in every period each agent is faced with a choice among a fixed set of alternatives. The *population state* is a distribution of actions describing the aggregate behavior of agents in the population. In each period a fixed share of the agents die and are replaced with new agents. Each new agent observes a finite sequence of actions (called a *sample*) of random size. The sample may consist either of past actions of random agents in the population (as in the social learning models mentioned above) or past actions of the current, randomly drawn, opponent (as in the community enforcement models mentioned above).

A sampling process is a tuple that specifies all the above components. A decision rule specifies the distribution of actions played by new agents as a function of the observed sample. The agent

[1] This feature is a main difference compared to other models of convergence of social learning such as Schlag (1998); Oyarzun and Ruf (2014).