Games and Economic Behavior 80 (2013) 131–146

Contents lists available at SciVerse ScienceDirect

Games and Economic Behavior

www.elsevier.com/locate/geb

Biased social learning

Helios Herrera [a], Johannes Hörner [b,∗]

[a] SIPA, Columbia University, 420W 118 Street, 1306 New York, NY 10011, USA
[b] Department of Economics, Yale University, 30 Hillhouse Av., New Haven, CT 06520, USA

ARTICLE INFO

Article history:
Received 20 March 2012
Available online 16 January 2013

JEL classification:
D83

Keywords:
Informational herds
Cascades
Selection bias

ABSTRACT

This paper examines social learning when only one of the two types of decisions is observable. Because agents arrive randomly over time, and only those who invest are observed, later agents face a more complicated inference problem than in the standard model, as the absence of investment might reflect either a choice not to invest, or a lack of arrivals. We show that, as in the standard model, learning is complete if and only if signals are unbounded. If signals are bounded, cascades may occur, and whether they are more or less likely than in the standard model depends on a property of the signal distribution. If the hazard ratio of the distributions increases in the signal, it is more likely that no one invests in the standard model than in this one, and welfare is higher. Conclusions are reversed if the hazard ratio is decreasing.

1. Introduction

The process of learning in social contexts confronts the same difficulties as any other statistical analysis. The data available to an individual may be subject to selection bias. Consider the leading example used by Bikhchandani, Hirshleifer and Welch (1992) (henceforth BHW), for instance. Upon learning that a paper has been previously rejected, a referee at a second journal tilts toward rejection. But what if, as is usually the case, he did not learn about this rejection? Surely, he would nevertheless wonder about the paper's journey onto his desk, and speculate about rejections the paper might have gone through. While publications are by definition observable, rejections are not.

To wit, there are far more significant instances of such bias in academia. How easy is it to publish a paper that finds inconclusive empirical evidence? In medical and social sciences, studies whose findings are statistically insignificant get filed away, biasing the published papers toward positive results.[1]

The difficulty in interpreting the absence of negatives is encountered everywhere. Is no one waiting in this line because cabs come by all the time, or because this isn't actually a cab line? Do the low figures of tax evasion reflect the success of deterrent policies, or the success of tax evaders? Why do 90% of mutual funds truthfully claim to have performance in the first quartile of their peers? (The other three quarters of funds have closed. See Elton et al., 1996.)

This paper develops a model of biased social learning and revisits the findings of the literature. In this model, individuals arrive randomly over time. As in Smith and Sørensen (2001), each agent has some private, conditionally independent information about the relevance of taking some decision immediately upon arrival—say, making an investment. As often in the social learning literature, we assume that the payoff from investing depends on the state of the world, but not on what

∗ Corresponding author.
E-mail addresses: hhh2108@columbia.edu (H. Herrera), johannes.horner@yale.edu (J. Hörner).

[1] This phenomenon is known as the "file-drawer problem", or the "publication bias". See Scargle (2000). As a result, prominent medical journals no longer publish results of drug research sponsored by pharmaceutical companies unless that research was registered in a public database from the start. Some of them also encourage publication of study protocols in their journal.

0899-8256/$ – see front matter
http://dx.doi.org/10.1016/j.geb.2012.12.006

earlier or later individuals decide. Therefore, values are common, and externalities are purely informational.[2] As is standard as well, signal distributions satisfy the strict monotone likelihood ration property (MLRP).

What sets this model apart from standard models is the following informational assumption. While the decision to invest (but not the payoff from investing) is observable to all future individuals, the failure to do so, and in fact, the mere arrival of individuals (who do not invest), remains hidden. Individuals arriving later will observe "positives" (if and when earlier individuals invested), but not "negatives" (if and when earlier individuals chose to abstain).

Therefore, every individual faces a complex problem of statistical inference: given the observed history, and the randomness in the arrival of individuals, how likely is it that some individuals had the opportunity to invest, but chose not to? And if so, what were their private signals? Note that, in this problem, time plays a key role, as it becomes increasingly more likely, as time passes by, that some individuals must have had the opportunity to invest.

In this context, we ask whether biased social learning exacerbates or mitigates herding. Could it be the case that some investment opportunities, or lucrative projects, remain unexploited because agents considering making it suspect that others must have thought of it, or even tried it before them? How many entrepreneurs, or scientists, stumbling across a new idea, chose not to follow through this idea because of the rational belief that they were unlikely to be the first to think of it?

Our first main result shows that, qualitatively, the absence of negatives does not alter the conditions under which cascades can, or cannot occur. If the informativeness of signals is bounded, wrong herds can occur (that is, they will occur for some prior and payoff parameters). On the other hand, if signals are unbounded, learning is necessarily complete; whether the state of the world is such that investment is profitable or not, agents will eventually learn it.

On the other hand, our second main result shows how, quantitatively, the absence of negatives affects the probability of a wrong herd. Consider the case of bounded signals (so that cascades may occur). What is the probability that no agent ever invests, while agents should, in the case of biased learning, relative to this probability in the benchmark model of BHW, in which all decisions, to invest or not, are observed? As it turns out, the comparison of these probabilities hinges upon a simple statistical property of the signal distribution, the increasing hazard ratio property.[3] If signals satisfy the increasing hazard ratio property (IHRP), that is, if the ratio of the hazard rates increases in the signal, then the probability of no one ever investing is lower under biased learning, *independently* of the state of the world. Conversely, if the hazard ratio is decreasing, then this probability is lower in the benchmark model. While biased learning always leads to higher investment (relative to the benchmark model) under IHRP, it nevertheless leads to lower welfare, at least in the version of our model in which there is only one investment opportunity.[4]

The first models of sequential decisions and observational learning by Banerjee (1992) and Bikhchandani et al. (1992), and their subsequent generalization by Smith and Sørensen (2001) all assume that all actions are observed by later individuals. Namely, agents could observe the precise sequence of decisions made by all the predecessors. Later work, notably Çelen and Kariv (2004), Callander and Hörner (2009) and Smith and Sørensen (1997) relaxes this assumption and considers the case in which either a subsample of the sequence, or a statistic thereof is observable. As these authors show, the asymptotic properties of social learning may radically change. For instance, Çelen and Kariv (2004) show that, when agents only observe the action of their immediate predecessor, beliefs do not converge. Therefore, complete learning never occurs, as beliefs and actions end up cycling. Hence, limiting the information available to agents may alter the qualitative properties of learning in general, although it turns out not to do so in the case of biased learning. Callander and Hörner (2009) show that if agents can only observe the fraction of agents having taken each action, rather than the entire sequence, then it might be optimal to take the action that was adopted by the minority of predecessors. A similar observational assumption is made in Hendricks et al. (2012).

Guarino et al. (2011) also analyze a framework in which one of the two actions is not observable, assuming simpler and more essential public information: the only public state variable is the aggregate number of people that took the observable action. Interestingly, they obtain that only cascades on the observable action can occur, never on the non-observable action.

Chari and Kehoe (2004) develop an observational learning model with a similar investment bias, in the sense that more information is revealed after observing an investment than after observing a decision not to investment. Each investment amount is an observable continuous variable, so the investor's private signal can be fully inferred (when this investment amount is positive). As in usual models, in case of a non-investment is observed then only a truncation on the investor's private signal can be inferred. In a sense, their model adds information in a biased fashion to the standard model (investment decisions become more informative), while we suppress information in a biased fashion, by assuming that decisions not to invest are not observable.

Our model is also related to models of endogenous timing such as the elegant paper of Chamley and Gale (1994). In their model as in ours, whether an agent has an opportunity to invest or not is a random variable. In their model, there is a finite number of agents who are all present from the start, and may choose to wait before investing, if they wish to. Inefficiently

[2] We shall also discuss at length a version in which there is only one investment opportunity, in which case there is an obvious payoff externality.

[3] The IHRP property is introduced in the statistical literature by Kalashnikov and Rachev (1986). Properties of IHRP in the standard ("BHW") learning model are derived in Herrera and Hörner (2011). Namely, IHRP is the necessary and sufficient condition under for the absence of informational cascades, namely provided the decision of the first individual depends on his signal, the decision of all later individuals will do as well. That is, it ensures that the posterior public belief necessarily stays in the learning region provided that the prior lies in it. More precisely, IHRP guarantees that this is the case after "good news", that is, after an observed investment decision. There is a corresponding property for the case of "bad news".

[4] This version allows us to focus on the history of no observed investment.