Economic Models of Social Learning*

Giuseppe Moscarini Marco Ottaviani

September 1995

Abstract

The theory of rational social learning studies how individual decision makers are influenced by the actions taken by others when information is dispersed. We present and discuss some models of Bayesian social learning which recently appeared in the economics literature. We focus on the problems of information acquisition, stationarity of the environment and endogenous pricing, and we propose some applications.

Social learning; Herding; Informational cascades; Experimental consumption; Endogenous information acquisition; Price competition.

fi D83

The financial support of Università Bocconi is gratefully acknowledged. The paper is the result of joint work. Sections 1, 3, and 7 have been written by Giuseppe Moscarini; Sections 2, 4, 5, 6 and 8 by Marco Ottaviani. Affiliation for both authors: Department of Economics, Massachusetts Institute of Technology. Address for both authors: Department of Economics, MIT, 50 Memorial Drive, E52-243d, Cambridge MA 02139, USA. Tel: (617) 742-8652. Fax: (617) 253-1330. E-mail: and .

Contents

1 Introduction **3**

2 Canonical Model of Sequential Social Learning **5**

3 Social Learning in a Changing World **8**

1 Introduction

A central theme in economics is the aggregation of information dispersed among agents. As Hayek wrote in 1945: "The peculiar character of the problem of a rational economic order is determined precisely by the fact that the knowledge of the circumstances of which we must make use never exists in concentrated or integrated form but solely as the dispersed bits of incomplete and frequently contradictory knowledge which all the separate individuals possess". In order to make decisions, rational agents use all available information. In particular, they look at the decisions of other agents. The theory of rational social learning studies how individual decision makers are influenced by the actions taken by others.

The canonical Bayesian social learning model of Banerjee (1992) and Bikhchandani, Hirshleifer and Welch (1992) describes the decision problem faced by a countable number of individuals, who must each take an action sequentially (in an exogenous order) under uncertainty as to the payoff-relevant state of the world. Each individual decides after having observed both an informative private signal and the entire history of decisions made by her predecessors. The individual cannot observe directly the private signals received by predecessors or their realized payoffpherpfcac(a)0.5(545(T)8)1.1(f)g(a-)7.3(tco)6ize-7

consider the implications of these models for economic behavior. We report extensively on our own research program aimed at applying these models to economic situations. In particular, we focus on the problem of information acquisition and of stationarity of the environment, and we propose some applications with endogenous pricing. We refer to the survey of Vives (1996) for a discussion of the links of the smooth and noisy model of social learning proposed by Vives (1995) and the rational expectations literature.

In Section 2 we introduce the sequential social learning model of Banerjee (1992) and Bikhchandani, Hirshleifer and Welch (1992). We define informational cascades, we review the main proposition that herding occurs with positive probability, and we point out the undesirable implications of the model.

In Section 3 we discuss the issue of the non stationarity of the environ-

2 Canonical Model of Sequential Social Learning

In this section we will introduce a simple version of the model otionigin

observed by individual q. A simple application of Bayes' rule yields

$$i_l(\ ^q) = \frac{\Pr(\$_1 \cap \ _l \mid k^q)}{\Pr(\ _l \mid k^q)} = \frac{\Pr(\ _l \mid k^q > _1)\$\Pr(\$_1 \mid k^q)}{\sum_{m=0}^{1} \Pr(\ _l \mid k^q > _m)\$\Pr(\$_m \mid k^q)} >$$

so that

$$i_0(\ ^q) = \frac{(1 - \ )\ ^q}{(1 - \ ^q) + (1 - \ )\ ^q} > \tag{2}$$

$$i_1(\ ^q) = \frac{\ ^q}{\ ^q + (1 - \ )(1 - \ ^q)} = \tag{3}$$

These posterior probabilities are used to compute the expected payoff from taking the two different actions in the two states. The following table summarizes the expected valuation of the consumer depending on the signal received:

EXPECTED VALUATION for action	d_0	d_1
if signal received		
$^q = $	í^a•˜"„ò•c9	

$$\text{if } \theta_q = 0 \text{ and } \begin{cases} s_q \leq & \Rightarrow d_q = d_0 \\ s_q \; A & \Rightarrow d_q = d_1 \end{cases} \qquad \text{if } \theta_q = 1 \text{ and } \begin{cases} s_q \; ? 1 - & \Rightarrow d_q = d_0 \\ s_q \geq 1 - & \Rightarrow d_q = d_1 \end{cases} \tag{5}$$

where the choice in the case of indifference between the two actions is assumed to be the one which minimizes the possibility of herding.

We will say that there in an (or) on action d_i at time q whenever good d_i is chosen by consumer q regardless of the individual's own private signal s_q. A cascade on action d_i is if the state is θ_m with $m \neq i$.

Can an incorrect cascade arise? With initial belief $p_1 = \frac{1}{2}$, a cascade starts whenever two individuals in a row take the same action. Let α_q be the probability that an incorrect cascade starts by period q with $q \geq 2$.

3 Social Learning in a Changing World

In this paragraph we consider the implications of the possibility that the state of the world changes during the social learning process, based on Moscarini and Ottaviani (1994b). For a similar extension in the experimentation lit-

$$^{q+1} = (1 - t)\ ^{q} + s(1 - \ ^{q}) = \qquad (8)$$

The dynamics of the public beliefs prior to the observation of the private signal is determined by (7) as long as $1- \ \leq \ ^{q} \leq$ (when not in a cascade) and by (8) when either ^{q} A or ^{q} ? $1-$ (during the cascade), with given initial condition ^{1}.

Since $H(\ ^{q+1}) = (1 - t)\ ^{q} + s(1 - \ ^{q})$ the public belief is a martingale unconditional on the state of the world, being a belief in a changing state. Similarly, the likelihood ratio $\frac{}{1\ 3}$ is not a martingale conditional on the true state. This finding can be reconciled with the usual result in Bayesian learning theory that the public belief is a martingale once the state of the world is defined as the sequence of states { $\$^{1} >^{2}\$ = =^{q} => }=.$ We welfare analysis requires the study of the conditional process that satisfies the Markov property.

It can easily be established that any cascade will eventually stop, provided that s and t are strictly positive. As shown in Moscarini and Ottaviani (1994b) if both s and t are strictly positive, cascades on d_{1}ÞUR$cU@

Similarly a necessary and sufficient condition for the existence of temporary alternating cascades can be derived. When $s = t$, it can be shown that the condition that guarantees existence of alternating cascades is $s \quad A \; 1 - \quad (1 - \quad)$.

Interestingly enough, the non-stationarity of the environment destroys

sible rules of thumb to determine their decisions. Banerjee and Fudenberg (1995) propose instead a model of rational word-of-mouth communication. They show that the system converges to the efficient outcome if each player samples two or more others. For a discussion of the numerous differences between these word-of-mouth models and the herding ones – like the one constructed in this section – see the appendix of Banerjee and Fudenberg (1995). Finally, the model of strategic experimentation of Bolton and Harris (1993) extends a continuous-time version of the classic two-armed bandit problem to a many-agent setting. In their setting agents observe the outcome of the experiments of other players, whereas in the model proposed in this section others can only infer information revealed through behavior.

Two varieties of an indivisible good, 0 and 1, are available for sale at fixed prices (set to zero for convenience of notation). There is a countable number of risk-neutral Bayesian decision makers. Though this might be any decision problem, we will adopt the language of a consumption decision. Each consumer lives and consumes for two consecutive periods, and then she leaves the market. One and only one consumer in each period demands one unit of one of two varieties of the good. The order of the consumers does not represent a choice variable, being either fixed or random.

The true relative quality of each good is unknown to the consumers, and corresponds to the fixed state of nature. To simplify notation and keep a symmetric structure, assume that the payoff from buying good 0 is known to be equal to 0, while good 1 gives a payoff of 1 in state $\$_1$ and -1 in state $\$_0$. Since a consumer buys for two periods, she has the opportunity to perform experiments in the first period, that can be interpreted as a period of trial. Moreover the consumer, before deciding which variety to buy, observes also the decisions made by the previou 3 T D 0 T c s t u rv t s e - 1 . 1 T f . 3 (i 3 0 . 2

For $\in$ £ i$_0$ i $\frac{1}{2}$ ¢ >$_1$ i $\frac{1}{2}$ ¢¤, this becomes

$$2 \quad -1+ \quad \Pr(\ _1|\ )[2\, i_1(\ ) - 1] = 2 \quad -1+ \quad [\ -(1-\ )] >$$

so that it is optimal to experiment with good 1 for

$$\geq e^2 \equiv \frac{1+ \quad (1-\ )}{2+} =$$

The signal received after the first-period consumption is revealed to the next consumer for a prior that belongs to the interval £ e^2 >$_1$ i $\frac{1}{2}$ ¢¤. Whenever the consumer is taking a fi

Our analysis uncovers the effect of property rights and competitive structure on the efficiency of social learning. The motivation of our research program is twofold. First, we wish to provide a stylized model of real mar-

Bergemann and Valimaki (1993) study strategic pricing by two firms selling to a single consumer who experiments optimally by purchasing one of two goods. On the demand side they have an individual agent, instead of our sequence of short-run consumers who learn from each other. The crucial difference of our formulation stems from the fact that in the model of social learning the buyer in the stage game has private information on the quality of the good, so that the stage game is a Bayesian game with asymmetric information.

5.1 Supply

In each period the prices of the two goods are posted simultaneously. Each good is produced by a different sector. Each sector can be either monopolized by a unique firm or perfectly competitive with free entry and exit. A monopolist can change prices in order to take advantage of the social learning of the consumers. The firms play at each stage a Bertrand game, quoting simultaneously prices to which they precommit to sell to the current consumer. They cannot condition the price on the signal received by the current customer because the signal is not observable to them. Each firm is risk-neutral and maximizes discounted expected profits, being a long-run player. The canonical model reported in section 2 corresponds to the outcome of competition between two competitive sectors. The outcome of competition of a monopolist against a competitive sector is described in section 6. In Section 7 we report on the model of duopoly, with one monopolist for each sector.

5.2 Demand

Two varieties of the good, 0 and 1

All the consumers have the same preferences, and they would all like to buy the good of higher quality. The payoff matrix is as specified in (1), with action d_l corresponding to purchase of good l. The consumer gets a payoff of 0 if she does not buy any good. Each consumer is risk-neutral and maximizes her expected valuation net of the price paid. Before deciding which good to buy, each consumer q observes a private signal $s_q \in \{s_0, s_1\}$ and the public history of action decisions of all preceding individuals. The distribution of the private signal depends on the state of the world. As in the previous sections, we consider a binary signal distribution: the probability that the signal s_l is realized conditional on the state being $\$_m$ is α if $l = m$ and $1-\alpha$ if $l \neq m$, with both l and m belonging to $\{0, 1\}$. It is assumed that the quality of the private signal is bounded, i.e. $\frac{1}{2} < \alpha < 1$. For $q \geq 2$

initial prior ^{1}, will be denoted from now on by

$$\Pr(\ \|\ ^{q}) =$$

The superscript to the belief will be suppressed when possible.

The social learning dynamics is:

$$^{q+1} = \begin{array}{l} ^{1/2} \\ i_0(\ ^{q}) \quad \text{if } d^{q} = d_0 \\ i_1(\ ^{q}) \quad \text{if } d^{q} = d_1 \end{array} \tag{9}$$

if not in a cascade, and

$$^{q+1} = \ ^{q} \tag{10}$$

during a cascade.

5.3 Information and Timing

We now summarize the timing of the game and the assumption made on information. In each period the prices of the two goods are posted simultaneously. Therefore the firms play at each stage a Bertrand game, quoting simultaneously prices to which they precommit to sell to the current consumer. The input of the stage game is the public belief . Each firm simultaneously posts the price for the sale of its good to the current consumer. Nature determines the signal received by the consumer, or the type of the consumer, according to the conditional probability distribution

$$\Pr(\ \|\ \$_m) = \begin{array}{ll} ^{1/2} & \\ & \text{for } m = l \\ 1- & \text{for } m \neq l \end{array}$$

with $l > \min\{0,1\}$. The true state of the world is not known by the firms, so that the probability assessment that signal l is received is $\Pr_{\ }(\ )$. Firms cannot condition the price on the signal received by the current customer because the signal is not observable to them. The customer observes the signal, updates her private belief, which represents, given the 0-1 payoff, his expected valuation of good 1 in monetary terms. The prior belief of each seller about customer qs type is the unconditional probability of the

which firm to buy. Her decision is publicly observed, and the game goes on to the next stage. The payoff to a firm when selling is equal to the price charged minus the marginal cost of the good (= 0). The payoff of the consumer is equal to the valuation for the good bought minus the price paid for it.

6 Monopoly Pricing

In this section, we consider the pricing strategy of a producer of one of the two goods, when she alone can act strategically on prices. Good 0 is produced by a competitive sector and good 1 by a monopolist. This corresponds to the case in which the property rights for this "new" good 1 are in the hands of a single producer, while the other "older" good 0 is produced by a sector with free entry and exit in any period. The model can easily be modified to consider the choice between buying good 1 and not buying at all, and the same qualitative results obtain. We prefer to formulate the problem in terms of the choice between two goods in order to allow for easier comparisons with the model of duopoly presented in Section 7. To facilitate the exposition, we report the results obtained in a two-period version of the monopoly model of Ottaviani (1995c) and discuss the implications of the multi-period analysis for price dynamics.

This dynamic model of monopoly predicts that in a first phase prices will be set high enough so that consumers with different signals buy different goods. This high price strategy by the monopolist allows next consumers to

own variety of the good and supplies it at zero marginal cost whenever the buyer demands it. The price quoted by the monopolist firm 1 that supplies good 1 in period q is denoted by S_1^q. The price of a firm m in sector 0 that supplies good 0 in period q is denoted by S_{0m}^q. The price for buyer q of good 0 is equal to the minimum price quoted by any supplier of this good, $S_0^q \equiv \min_m S_{0m}^q$.

The competitive sector is not able to effectively act on the price in any Perfect Bayesian Equilibrium, as shown in the next subsection, because of a free rider effect: the future gains from the appreciation of the consumers will be dissipated by the Bertrand competition among producers of this same good. In other words, this sector cannot internalize the externality due to the social learning of the consumers, because the appreciation of the consumers

period and exited immediately afterwards. In this way this firm would have made strictly positive expected profits. This possible "deviation" implies

backward from the last period. We first consider the one-period problem, that is also the problem faced by the monopolist in the last period. The solution of the next-to-last period concludes the analysis of the two-period problem and sheds light on the mechanics of the model.

6.1.1 Last Period

The solution of the last period problem gives the myopically optimal strategy. Not selling at all gives 0, posting the separating price yields $-(1-)$, and posting the pooling price $2 i_0( ) - 1$. For $\frac{1}{2} -$ the monopolist decides not to sell since in this region $2 i_0(\ ) - 1 \ 0$ and this is the last period. The separating price Sir6 8.9