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.

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

JEL Classification: 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: `mottavia@mit.edu` and `mosca@mit.edu`.

Contents

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 payoffs. Learning about the state of the world occurs in a Bayesian fashion.

Private signals are assumed to be independent and identically distributed draws from a random variable correlated with the true state of the world. If all signals were observable by one individual, the true state would eventually be revealed almost surely by the Strong Law of Large Numbers. Because signals are private, the extent to which they can be inferred from actions becomes an interesting problem. An externality arises from the information that the action of agent n conveys to her successors $n+1, n+2, ...$ who observe that action, a fact that agent n does not take into account. The striking result in this model is that the agents eventually disregard their private information and rely only on public information, so that herd behavior becomes individually rational. With positive probability eventually everyone will take the same less profitable action.

In this paper we present some models of rational social learning that appeared recently in the economics literature, and review some attempts to 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 environment over long time horizons. We report on our model of social learning in a changing world, where we assume that the payoff-relevant state of the world changes stochastically during the learning process. We show that informational cascades are (locally) robust to this perturbation. We briefly discuss the implication that this kind of nonstationarity bears on the results and methodology of the existing Bayesian learning literature.

In Section 4 we examine the interactions between individual experimentation and social learning. We review a very simple model where experimental consumption allows individuals to acquire "firsthand" information on the quality of the goods, while observation of the behavior of other consumers provides valuable "secondhand" information. We analyze the dynamics of information acquisition, and we show again the robustness of informational cascades to the possibility of choosing the amount of information to be collected.

In Sections 5, 6, and 7 we consider the effect of pricing on the social learning of consumers. We interpret actions in the canonical model as purchasing decisions, and we specify the supply side in correspondence with different market structures. We characterize the optimal dynamic pricing strategies adopted in equilibrium by sellers in this market, depending on different allocations of property rights. Building on these results, we verify how the price system can mitigate the information externality and thus the herding inefficiency. Our analysis indicates that non competitive intertemporal allocations are more efficient than the competitive allocation, a result that parallels the well known results on dynamic efficiency of monopoly in the literature on research and development.

2 Canonical Model of Sequential Social Learning

In this section we will introduce a simple version of the model of Bikhchandani, Hirshleifer and Welch (1992), that we consider the "canonical" model of sequential social learning. Suppose, for simplicity, that there are two possible actions, denoted by a_0 and a_1, and two states of the world, denoted by ω_0 and ω_1. Let η^1 be the common prior belief that the state is ω_1 at the beginning of time.

A countable number of individuals must decide sequentially which action to take. Action a_1 is more rewarding than action a_0 in state ω_1, while the opposite is true in state ω_0: the payoff if action a_i is taken when the state is ω_j is 1 if $i = j$ and 0 if $i \neq j$, with both i and j belonging to $\{0, 1\}$. The payoff matrix is

	State of the world	ω_0	ω_1
Action taken			
a_0		1	0
a_1		0	1

(1)

Before deciding which action to take, each individual n observes a private signal $\sigma^n \in \{\sigma_0, \sigma_1\}$ and the public history of action decisions of all preceding individuals $1, 2, ..., n-1$. The signals received by the predecessors cannot be observed (in this sense they are private). The distribution of the private signal depends on the current state of the world. The probability that the signal σ_i is realized conditional on the state being ω_j is α if $i = j$ and $1-\alpha$ if $i \neq j$, with both i and j belonging to $\{0, 1\}$. It is assumed that the quality of the private signal is bounded, i.e. $\frac{1}{2} < \alpha < 1$. A signal of quality $\alpha < \frac{1}{2}$ would be equivalent to one of quality $1 - \alpha > \frac{1}{2}$ after relabelling the alternatives. Different private signals are independent draws from this state-dependent Bernoulli distribution.

For $n \geq 2$ let $H^n \equiv \{a_0, a_1\}^{n-1}$ be the space of all possible period n histories of actions chosen by the $n-1$ predecessors of individual n. Let h^n denote an element of H^n. Let $\eta^n \equiv \Pr(\omega_1 | h^n)$ be the public probability belief that the state is ω_1 in period n conditional on the publicly observed history of actions chosen by the predecessors of individual n. Similarly let $f_i(\eta^n) \equiv \Pr(\omega_1 | h^n, \sigma_i)$ be the private belief that the state is ω_1 conditional on both the action history h^n and the realization σ_i of the private signal observed by individual n. A simple application of Bayes' rule yields

$$f_i(\eta^n) = \frac{\Pr(\omega_1 \cap \sigma_i | h^n)}{\Pr(\sigma_i | h^n)} = \frac{\Pr(\sigma_i | h^n, \omega_1)\Pr(\omega_1 | h^n)}{\sum_{j=0}^{1} \Pr(\sigma_i | h^n, \omega_j)\Pr(\omega_j | h^n)},$$

so that

$$f_0(\eta^n) = \frac{(1-\alpha)\eta^n}{\alpha(1-\eta^n) + (1-\alpha)\eta^n}, \tag{2}$$

$$f_1(\eta^n) = \frac{\alpha\eta^n}{\alpha\eta^n + (1-\alpha)(1-\eta^n)}. \tag{3}$$

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	a_0	a_1	
if signal received				
$\sigma^n = \sigma_1$		$1 - f_1(\eta^n)$	$f_1(\eta^n)$	(4)
$\sigma^n = \sigma_0$		$1 - f_0(\eta^n)$	$f_0(\eta^n)$	

The decision rule of the agent n is to choose the optimal action a^n which gives her the highest payoff. For instance, if the private signal received by individual n is $\sigma^n = \sigma_1$, then $a^n = a_1$ when

$$f_1(\eta^n) \geq (1 - f_1(\eta^n)),$$

i.e.

$$f_1(\eta^n) \geq \frac{1}{2},$$

which after substituting from (3) becomes

$$\eta^n \geq 1 - \alpha.$$

The decision rule can be summarized as:

$$\begin{array}{ll} \text{if } \sigma^n = 0 \text{ and} & \begin{cases} \eta^n \leq \alpha & \Rightarrow \quad a^n = a_0 \\ \eta^n > \alpha & \Rightarrow \quad a^n = a_1 \end{cases} \\ \text{if } \sigma^n = 1 \text{ and} & \begin{cases} \eta^n < 1-\alpha & \Rightarrow \quad a^n = a_0 \\ \eta^n \geq 1-\alpha & \Rightarrow \quad a^n = a_1 \end{cases} \end{array} \tag{5}$$

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 *informational cascade* (or *cascade*) on action a_i at time n whenever good a_i is chosen by consumer n regardless of the individual's own private signal σ^n. A cascade on action a_i is *incorrect* if the state is ω_j with $j \neq i$.

Can an incorrect cascade arise? With initial belief $\eta^1 = \frac{1}{2}$, a cascade starts whenever two individuals in a row take the same action. Let ψ_n be the probability that an incorrect cascade starts by period n with n even. Clearly:

$$\psi_n = \psi_{n-2} + [\alpha(1-\alpha)]^{\frac{n-2}{2}} \left[\frac{1}{2}(1-\alpha) + \frac{1}{2}(1-\alpha)^2 \right],$$

so that

$$\psi_n = \frac{\frac{1}{2}\left[(1-\alpha) + (1-\alpha)^2\right]\left[1 - \alpha(1-\alpha)\right]^{\frac{n}{2}}}{1 - \alpha(1-\alpha)}.$$

The probability that eventually an incorrect cascade starts is equal to

$$\lim_{n\to\infty} \psi_n = \frac{(1-\alpha) + (1-\alpha)^2}{2\left[1 - \alpha(1-\alpha)\right]} > 0.$$

The information externality drives this pathological outcome. When public evidence represented by past actions becomes too strong with respect to the precision of a single signal, an agent rationally prefers to ignore her own signal and to conform to the prevailing choice. All successors know that the action taken by the predecessor is uninformative, so that they are in the same situation. Because valuable private signals are wasted from now on and the social learning process stops, "herding" on the same action is both individually rational and socially wasteful. Smith and Sørensen (1994) show that this pathological outcome of "herding" holds whenever the quality of private signals is bounded. Lee (1993) obtains complete learning in the long run in a sequential social learning model with a continuous action space. In the above papers on social learning the prices and payoffs of the goods are fixed exogenously, so that the focus is only on the demand side of the market.

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 literature see Keller and Rady (1995). Assume that after the decision of each individual the state of the world changes according to the Markov transition matrix

$$\begin{pmatrix} 1-p & p \\ q & 1-q \end{pmatrix}.$$

If p and q were both equal to 0, when not in a cascade, the public belief when individual $n+1$ has to decide is equal to the posterior belief that leads the predecessor n to act according to her signal $\sigma^n = \sigma_i$. In our notation $\eta^{n+1} = f_i(\eta^n)$. A cascade would then start the first time k that $\eta^k < 1-\alpha$ or $\eta^k > \alpha$, and once started it would never end, because the individual coming after k would have the same prior information as k, which has the property that no single private signal can affect the decision.

If instead either p or q is strictly greater than zero, the analysis changes drastically. When the possibility that the state of the world has changed in the meantime is accounted for, the public prior belief of individual $n+1$ coming after individual n who chose $a^n = a_i$ not during a cascade satisfies

$$\eta^{n+1} = (1-q) f_i(\eta^n) + p(1 - f_i(\eta^n)), \tag{6}$$

which can be rewritten by (2) and (3) as

$$\eta^{n+1} = \begin{cases} \zeta_0(\eta^n) \equiv \frac{(1-q)(1-\alpha)\eta^n + p\alpha(1-\eta^n)}{(1-\alpha)\eta^n+\alpha(1-\eta^n)} & \text{if } a^n = a_0 \\ \zeta_1(\eta^n) \equiv \frac{(1-q)\alpha\eta^n + p(1-\alpha)(1-\eta^n)}{\alpha\eta^n+(1-\alpha)(1-\eta^n)} & \text{if } a^n = a_1 \end{cases} \tag{7}$$

A cascade on a_1 (or a_0) arises as soon as $\eta^k > \alpha$ (or $\eta^k < 1-\alpha$). Note that anyone can compute these conditional probabilities and determine when this happens. We will consider the case $\eta^k > \alpha$, since the other case can be treated symmetrically. The action chosen will be $a^k = a_1$, regardless of the signal σ^k. The next individual $k+1$ knows that $a^k = a_1$ is uninformative, and computes her prior as

$$\eta^{k+1} = (1-q)\eta^k + p\left(1-\eta^k\right).$$

In general the following individual $n+1$, as long as $\eta^n > \alpha$ or $\eta^n < 1-\alpha$, will update her prior belief during the cascade in the same fashion

$$\eta^{n+1} = (1-q)\,\eta^n + p\,(1-\eta^n)\,. \tag{8}$$

The dynamics of the public beliefs prior to the observation of the private signal is determined by (7) as long as $1-\alpha \leq \eta^n \leq \alpha$ (when not in a cascade) and by (8) when either $\eta^n > \alpha$ or $\eta^n < 1-\alpha$ (during the cascade), with given initial condition η^1.

Since $E\left(\eta^{n+1}\right) = (1-q)\,\eta^n + p\,(1-\eta^n)$ the public belief is *not* a martingale unconditional on the state of the world, being a belief in a changing state. Similarly, the likelihood ratio $\frac{\eta}{1-\eta}$ 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 $\{\omega^1, \omega^2, ..., \omega^n, ...\}$. 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 p and q are strictly positive. As shown in Moscarini and Ottaviani (1994b) if both p and q are strictly positive, cascades on a_1 are temporary provided that $\alpha > \frac{p}{p+q}$, i.e. that $p < \frac{\alpha}{1-\alpha}q$, and cascades on a_0 are temporary whenever $\frac{p}{p+q} > 1-\alpha$ (equivalent to $p > \frac{1-\alpha}{\alpha}q$).

Cascades on a single good can arise only if the state of the world is sufficiently persistent. If instead state changes are sufficiently unpredictable (i.e. p and q are sufficiently close to $\frac{1}{2}$) the belief η^n will be always close to $\frac{1}{2}$, and so no cascade will ever arise. It is possible to have alternating cascades if the changes of the state of the world are sufficiently predictable. A necessary and sufficient condition for existence of temporary cascades on a_1 is that the fixed point of the difference equation

$$\eta^{n+1} = \zeta_1\left(\eta^n\right)$$

is above α. This condition can be shown to be equivalent to

$$p > \tilde{p}\,(\alpha) \equiv \alpha\,(1-\alpha)\,\frac{1-2\alpha}{(1-\alpha)^2} + \frac{\alpha^2}{(1-\alpha)^2}q.$$

Similarly cascades on a_0 arise whenever

$$q > \tilde{q}\,(\alpha) \equiv \alpha\,(1-\alpha)\,\frac{1-2\alpha}{(1-\alpha)^2} + \frac{\alpha^2}{(1-\alpha)^2}p.$$

If we constrain p and q to be equal, then cascades arise for $p < \alpha\,(1-\alpha)$.

Similarly a necessary and sufficient condition for the existence of temporary alternating cascades can be derived. When $p = q$, it can be shown that the condition that guarantees existence of alternating cascades is $p > 1 - \alpha(1-\alpha)$.

Interestingly enough, the non-stationarity of the environment destroys the Martingale property of the public belief (unconditional on the true state of the world) and of the likelihood ratio (conditional on the state of the world), the cornerstone of Bayesian learning theory. We consider this result a serious weakness in standard Bayesian learning models, where asymptotic behavior is described by means of Martingale Convergence Theorems, which require both a long time horizon and a completely fixed state. Moscarini and Ottaviani (1994b) use alternative tools for establishing the existence, uniqueness and global stability of an invariant probability measure over beliefs and actions. The "true" state is a sequence of states $\{\omega^1, \omega^2, ..., \omega^n, ...\}$, so that each history is a single sample. Therefore, the convergence of aggregate behavior to the invariant distribution describing the steady state of the system is weak, and not uniform, as would be the case for Martingale processes.

4 Individual Experimentation and Social Learning

In this section we construct a simple model of individual experimentation and social learning, along the lines of Ottaviani (1995a). In most models of observational learning the decisions of others are a valuable "secondhand" source of information, although no explicit consideration is made for the acquisition of information. In the optimal experimentation literature (Rothschild (1974), Grossman, Kihlstrom and Mirman (1977), and Aghion, Bolton, Harris and Jullien (1991)), consumption allows single individuals to acquire "firsthand" information on the different goods available. In this section we are concerned primarily with the interplay between information acquisition through experimentation and learning from others. The explicit consideration of these two sources of information allows for the analysis of the dynamics of information acquisition.

Ellison and Fudenberg (1995) analyze a model of social learning in which individual players can learn from their own experience and that of others, obtained via word-of-mouth communication. Players follow intuitively plausible 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 ω_1 and -1 in state ω_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 previous consumers.

To obtain information on which good is optimal to buy, a consumer can experiment in the first period of her life (individual experimentation) and/or look at the decisions of previous consumers (social learning). Consumption of good 1 gives an imperfect payoff realization. Assume that the payoff signals have a binary distribution: $\Pr(\sigma_0|\omega_0) = 1 - \Pr(\sigma_0|\omega_0) = \alpha > \frac{1}{2}$, and $\Pr(\sigma_1|\omega_1) = 1 - \Pr(\sigma_1|\omega_1) = \alpha$.

The choice in the last period given belief $\eta = \Pr(\omega_1)$ is to buy good 1 if $\eta \geq \widetilde{\eta}^1 \equiv \frac{1}{2}$. In the next to last period consumption of good 1 allows acquisition of information, so that he total payoff from good 1 is

$$2\eta - 1 + \beta \Pr(\sigma_1|\eta) [\max \langle 2f_1(\eta) - 1, 0\rangle] + \beta \Pr(\sigma_0|\eta) [\max \langle 2f_0(\eta) - 1, 0\rangle] .$$

For $\eta \in \left[f_0\left(\frac{1}{2}\right), f_1\left(\frac{1}{2}\right)\right]$, this becomes

$$2\eta - 1 + \beta \Pr(\sigma_1|\eta)[2f_1(\eta) - 1] = 2\eta - 1 + \beta[\eta - (1-\alpha)],$$

so that it is optimal to experiment with good 1 for

$$\eta \geq \widetilde{\eta}^2 \equiv \frac{1+\beta(1-\alpha)}{2+\beta}.$$

The signal received after the first-period consumption is revealed to the next consumer for a prior that belongs to the interval $\left[\widetilde{\eta}^2, f_1\left(\frac{1}{2}\right)\right]$. Whenever the consumer is taking a first-period action which is not myopically optimal, then the second-period action changes depending on the realization, so that the next decision maker can infer the information acquired through experimental consumption. For a full treatment of the dynamics of information acquisition in a multi-period version of this model see Ottaviani (1995a). Since eventually an informational cascade occurs, and with positive probability it is on the inferior good, the fundamental results of the model of rational social learning of Bikhchandani, Hirshleifer and Welch (1992) and Banerjee (1992) appear to be robust to the endogenization of the acquisition of information through experimental consumption.

5 Social Learning and Price Competition

The focus of the previous work on social learning is on the demand side of the market, with no consideration for the supply side. In this section we consider both the demand and supply side in a model of price competition that allows us to study the effects of information externalities in different markets.

In the theory of rational social learning prices of different possible actions available to the individuals are fixed. In this section we extend this theory in order to consider the effect of prices on the decisions of the consumers and the incentive of the firm that controls prices to induce herding on its own good and to prevent herding on the other good. In markets where consumers learn from each other about product quality prices adjust along the learning process. What does social learning by the consumers imply for equilibrium prices of the products? Is the price system capable of mitigating the information externalities and thus potential inefficiencies?

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 markets where buyers are learning from each other, and suppliers compete in an attempt to exploit the social learning of the buyers. Second, we wish to see to what extent the main inefficiency result of herding of people on the same wrong action carries through when prices can change during the learning process. Do the long-run inefficiencies due to the social learning externality decrease when long-run players – the firms supplying the good – are brought into the picture?

On the demand side of the market there is a sequence of consumers exogenously ordered. Each consumer comes to the market with the same preferences. Two goods vertically differentiated by quality are sold in the market. Individual consumers have to decide which one of two goods to buy without knowing for certain which is the most desirable one. Each consumer receives an informative private signal and can look at the decisions of all the consumers who decided before her, in an attempt to recall their useful private information on the quality of the two different goods.

There are different possible specifications of the supply side of the market. Each variety of the good is produced by a corresponding sector. Suppliers are not allowed to observe the realization of the signal of its present potential buyer, and can infer – as every other consumer – the past signals only when the decisions of the consumers revealed them. We argue that the "canonical" model of Bikhchandani, Hirshleifer and Welch (1992) corresponds to the case of perfect competition with free entry and exit in both sectors. The case of duopoly, in which there is a monopolist in each sector, is studied in detail by Moscarini and Ottaviani (1994a). Ottaviani (1995c) considers the case in which a monopolist makes up one sector and competes against a perfectly competitive sector that produces the other good. The monopolist can change prices so as to influence the social learning process of the consumers.

It is not difficult to find examples of different markets that can be analyzed by these different supply specifications. The competitive case can be applied to study the adoption of a new agricultural technology. The duopoly model is appropriate to understand the competition among the firms in the aircraft industry (e.g. Boeing *vs.* McDonnell Douglas), where the purchase decisions of the airline companies are easily observable by the competitor producer and by the other airline companies. The model of monopoly that is studied in this section can be applied to the market for professional services or books.

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 are available for sale. The true relative quality of the two goods is unknown to both the buyers and the sellers, and corresponds to the underlying state of nature. A countable number of individuals, indexed by $n = 1, 2...$, must decide sequentially which one of the two available goods denoted by 0 and 1 to buy at market prices. One buyer in each period demands one unit of an indivisible good, then she leaves the market. Buyers are therefore short-run players of a dynamic game among sellers. There are two states of the world, ω_0 and ω_1, indicating which one is the better good.

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 a_i corresponding to purchase of good i. 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 n observes a private signal $\sigma^n \in \{\sigma_0, \sigma_1\}$ 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 σ_i is realized conditional on the state being ω_j is α if $i = j$ and $1 - \alpha$ if $i \neq j$, with both i and j belonging to $\{0, 1\}$. It is assumed that the quality of the private signal is bounded, i.e. $\frac{1}{2} < \alpha < 1$. For $n \geq 2$ let $H^n \equiv \{\{a_0, a_1\} \times \mathcal{R} \times \mathcal{R}\}^{n-1}$ be the space of all possible period n histories of actions chosen by the $n-1$ predecessors of individual n and prices for the two goods posted in the past. Let h^n denote an element of H^n.

Let η^1 be the common prior belief that the state is ω_1 (i.e. good 1 is better than good 0) at the beginning of time, and $\eta^n \equiv \Pr(\omega_1 | h^n)$ be the public probability belief that the state is ω_1 in period n conditional on the publicly observed history of goods bought by the predecessors of consumer n. Similarly let $f_i(\eta^n) \equiv \Pr(\omega_1 | h^n, \sigma_i)$ be the private belief that the state is ω_1 conditional on both the action history and the realization σ_i of the private signal observed by consumer n. Bayes' rule yields the posterior beliefs $f_0(\eta^n)$ and $f_1(\eta^n)$ as in (2) and (3). We will refer to customer n after receiving signal σ_i as to *type-i* customer n, as the signal modifies her valuation in a way that is not known to the sellers according to (4). The decision a^n of consumer n is to choose the good which gives her the highest expected payoff net of the price. For instance, if the private signal received by individual n is $\sigma^n = \sigma_1$, then $a^n = a_1$ when

$$f_1(\eta^n) - P_1^n \geq (1 - f_1(\eta^n)) - P_0^n,$$

i.e.

$$f_1(\eta^n) \geq \frac{1}{2} + \frac{P_1^n - P_0^n}{2}.$$

The probability assessed by the players that the consumer n has received signal σ_i, given the prior belief η^n that results from the history h^n and the initial prior η^1, will be denoted from now on by

$$\Pr(\sigma_i|\eta^n).$$

The superscript to the belief will be suppressed when possible.

The social learning dynamics is:

$$\eta^{n+1} = \begin{cases} f_0(\eta^n) & \text{if } a^n = a_0 \\ f_1(\eta^n) & \text{if } a^n = a_1 \end{cases} \tag{9}$$

if not in a cascade, and

$$\eta^{n+1} = \eta^n \tag{10}$$

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 η. 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

$$\Pr(\sigma_i|\omega_j) = \begin{cases} \alpha & \text{for } j = i \\ 1-\alpha & \text{for } j \neq i \end{cases}$$

with $i, j \in \{0, 1\}$. The true state of the world is not known by the firms, so that the probability assessment that signal i is received is $\Pr_{\sigma_i}(\eta)$. 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 n's type is the unconditional probability of the corresponding signal, which depends on the public belief η^n and is therefore variable from one stage game to the next.

The customer compares her updated valuation - "learned" after observing the private signal - with the prices quoted and decides whether and from 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 recall the private information of the preceding consumers from the observation of their purchase decisions. In this phase of active social learning, the price decreases on average over time. If the monopolist sells, then the price increases, otherwise it decreases. Eventually learning will stop, because the price-setting firm decides either to exit the market, or to capture the whole market by reducing the price. We believe that our model provides an articulate story for the "social influence on price" of Becker (1991), who just assumes that a consumer's demand for a good depends on the demands by other consumers.

The demand side of the market and the information structure is as specified in the previous section. As for the supply side of the market, good 0 is supplied by a competitive sector and good 1 by a monopolist. In any period each firm within each sector quotes simultaneously a price for the unit of its 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 n is denoted by P_1^n. The price of a firm j in sector 0 that supplies good 0 in period n is denoted by P_{0j}^n. The price for buyer n of good 0 is equal to the minimum price quoted by any supplier of this good, $P_0^n \equiv \min_j P_{0j}^n$.

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 is a public good to the sector. The producer of good 1, instead, has property rights on the future purchases of its good and can therefore effectively act on prices.

6.1 Solution

In this section we solve the model by reducing it to a problem of discounted dynamic optimization. We will first show that the particular structure of the sector that produces the alternative good 0 allows to reduce the analysis of this game to a decision problem for the monopolist. We then characterize the optimal pricing strategy that solves the one-period problem. The analysis and discussion of the solution to the two-period problem sheds light on the mechanics of the model. We then report on the characterization of the properties of the solution to the infinite horizon problem.

Consider first the price of good 0, produced by the competitive sector. Assume that anyone is allowed to enter freely into or exit from this sector and produce this same good at any point in time. Free entry and exit imply the impossibility to collect any possible future gain due to the favorable social learning of the consumers that would be induced by a price temporarily lower than the marginal cost. This implies that in any period of time the price of good 0 in equilibrium is equal to the marginal cost ($= 0$). To see this notice that in any perfect Bayesian equilibrium the instantaneous expected profit of each firm cannot be positive. If it were positive, the price would exceed the marginal cost in at least one period. But then one firm could have entered in that period, charged a price slightly lower for a unit of good 0 on sale that period and exited immediately afterwards. In this way this firm would have made strictly positive expected profits. This possible "deviation" implies that the only possible equilibrium is one in which the price of good 0 is not greater than the marginal cost in any period. Free exit implies that future expected profits are at least equal to zero. We can therefore conclude that in any period of time the price of good 0 is equal to the marginal cost $(= 0)$.

In this section we analyze the optimization problem of a new entrant with monopoly power in the supply of good 1, differentiated from and competing with good 0, that is sold at $P_0^n = 0$ in any period n. In any period the monopolist can decide not to sell. This clearly gives a payoff of zero. If instead the monopolist wishes to sell, she can restrict attention to two prices, all other prices being strictly dominated by one of these. The two prices are: the *separating price* $P_S(\eta) \equiv 2f_1(\eta) - 1$, the maximum price that the current type-1 consumer is willing to pay to buy the good; and the *pooling price* $P_P(\eta) \equiv 2f_0(\eta) - 1$, the maximum price at which also type-0 consumer buys the good.[1] Note that $P_S(\eta) > P_P(\eta)$. Different prices yield a different current expected profit and have different dynamic effects. As for current expected profit, $P_S(\eta)$ sells with probability $\Pr(\sigma_1|\eta)$, so that it yields an immediate expected payoff of $\Pr(\sigma_1|\eta)[2f_1(\eta) - 1] = \eta - (1 - \alpha)$, while $P_P(\eta)$ yields the immediate payoff of $2f_0(\eta) - 1$ with certainty. As for dynamic effects, the pooling price $P_P(\eta)$, as any price strictly lower, stops social learning, because the next consumer cannot infer the signal of the immediate predecessor who buys the good regardless of the signal received. This means that, after a pooling price is posted, the public belief in the subsequent period will remain equal to the prior η of the predecessor. The separating price $P_S(\eta)$, as any price between $P_P(\eta)$ and $P_S(\eta)$, instead, sells only to part of the market, so that it allows for revelation of signals and thus for social learning. Clearly any price strictly lower than the pooling one $P_P(\eta)$ gives an immediate profit strictly lower than $2f_0(\eta) - 1$ and the same continuation payoff. Any price between $P_P(\eta)$ and $P_S(\eta)$ gives an immediate profit strictly lower than $\eta - (1 - \alpha)$ that can be attained by the separating price $P_S(\eta)$, and the same continuation payoff. Therefore, from now on only these three undominated policies of the monopolist will be considered: not selling, posting the separating price, posting the pooling price.

To determine the optimal pricing policy of the monopolist we proceed

[1] More precisely, the tie-breaking assumption implies that the pooling price is $2f_0(\eta) - 1 - \epsilon$, with $\epsilon > 0$ arbitrarily small.

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 $\eta - (1-\alpha)$, and posting the pooling price $2f_0(\eta) - 1$. For $\eta < 1-\alpha$ the monopolist decides not to sell since in this region $2f_1(\eta) - 1 < 0$ and this is the last period. The separating price $P_S(\eta)$ will be charged if the associated expected payoff $\eta - (1-\alpha)$ is higher than $2f_0(\eta) - 1$ yielded by the pooling price $P_P(\eta)$.After substitutions it can be easily verified that the pricing strategy in the last stage is

$$P^1(\eta) = \begin{cases} \text{any } > 2f_1(\eta) - 1 & \text{for } \eta \le 1-\alpha \\ P_S(\eta) = 2f_1(\eta) - 1 & \text{for } 1-\alpha \le \eta \le \overline{\eta}^1 \\ P_P(\eta) = 2f_0(\eta) - 1 & \text{for } \eta \ge \overline{\eta}^1 \end{cases}$$

where

$$\overline{\eta}^1 = \frac{-(1-\alpha)^2 + \sqrt{(1-\alpha)^4 + (2\alpha-1)\alpha^2}}{(2\alpha-1)} \in (\alpha, 1). \tag{11}$$

is the largest root of the quadratic equation

$$\bar{k}(\eta,\alpha) = (2\alpha-1)\eta^2 + 2(1-\alpha)^2\eta - \alpha^2.$$

The probability of selling decreases with the price charged. When the belief is η, if the price $P \in [P_P(\eta), P_S(\eta)]$, the probability of selling is $\Pr(\sigma_1|\eta)$; if instead $P < P_P(\eta)$, then the good is sold for sure. The demand function has two steps: as η increases the upper step approaches the lower one. Because in the limit as η goes to 1, $f_1(\eta)$ and $f_0(\eta)$ both tend to 1, the separating price tends to the pooling one. The demand at the separating price is equal to $\Pr(\sigma_1|\eta)$ and tends to α as η goes to 1. Therefore it must be optimal to charge the pooling price that sells with probability 1, for η large enough.

The resulting value function is globally convex, being the maximum of convex functions. The last period value function

$$V^1(\eta) = \begin{cases} 0 & \text{for } \eta \le 1-\alpha \\ \eta - (1-\alpha) & \text{for } 1-\alpha \le \eta \le \overline{\eta}^1 \\ 2f_0(\eta) - 1 & \text{for } \eta \ge \overline{\eta}^1 \end{cases}$$

is flat at zero for $\eta < 1-\alpha$, linearly increasing in η for $\eta \in [1-\alpha, \overline{\eta}^1]$, and strictly convex for $\eta > \overline{\eta}^1$. The convexity of the value function can be obtained whenever the distributions of signals conditional on the state of nature satisfy the monotone hazard rate condition, as shown in Ottaviani (1995c).

6.1.2 Next-to-Last Period

In the next to last period, for $\eta < 1-\alpha$, the monopolist might be willing to bear current losses charging the price $P_S(\eta) < 0$ in order to gain by selling at a positive price in the next (and last) period in case of a good draw (i.e. if σ_1 is realized). Consider $\eta < 1-\alpha$ and the price $P_S(\eta)$ in the next to last period followed by the optimal policy in the last period. Denoting the discount factor of the firm by β, the expected discounted profit under this policy is

$$\eta - (1-\alpha) + \beta \Pr(\sigma_1|\eta) V^1(f_1(\eta)).$$

because $V^1(f_0(\eta)) = 0$, as the monopolist exits the market in the last period if unfavorable information (σ_0) is reveled in the next-to-last period, being $f_0(\eta) < \eta < 1-\alpha$. If instead the realization is σ_1, the separating price for the new belief $f_1(\eta)$ is charged in the last period yielding $V^1(f_1(\eta)) = f_1(\eta) - (1-\alpha)$, so that the total expected discounted payoff from continuing learning is

$$\eta - (1-\alpha) + \beta \Pr(\sigma_1|\eta)[f_1(\eta) - (1-\alpha)]$$

which is larger than zero, the payoff obtained by exiting the market right away, when

$$\eta \geq \frac{(1-\alpha)(1+\beta(1-\alpha))}{1+\beta[1-2\alpha(1-\alpha)]} \equiv \underline{\eta}^2. \tag{12}$$

Notice that $\underline{\eta}^2 < \underline{\eta}^1 \equiv 1-\alpha$. For $\eta \in \left[\underline{\eta}^2, \underline{\eta}^1\right]$ the monopolist is willing to bear immediate losses, that will be more than offset by the expected future profits made if favorable information is revealed to the market through a sale.

Consider $\eta > \overline{\eta}^1$, then there is a trade-off between

1. a higher instantaneous payoff with the pooling price $P_P(\eta)$ (myopically optimal in this region) than with the separating price $P_S(\eta)$,

2. a higher future expected payoff with the separating price $P_S(\eta)$ that allow for social learning, due to the convexity of $V^1(.)$.

For η close enough to, and strictly larger than $\overline{\eta}^1$, it is optimal to charge the separating price $P_S(\eta)$ since the benefit from learning is larger than the reduction in current profits (which is infinitesimal for η close enough to $\overline{\eta}^1$). Define the "expected value function" with one period to go as

$$EV^1(\eta) \equiv \Pr(\sigma_0|\eta) V^1(f_0(\eta)) + \Pr(\sigma_1|\eta) V^1(f_1(\eta)).$$

We need to compare the payoff from the pooling price

$$\underbrace{2f_0(\eta)-1}_{\substack{\text{Current profit}\\ \text{if pooling}}} + \underbrace{\beta V^1(\eta)}_{\substack{\text{Continuation value}\\ \text{if pooling}}}$$

to the one from the separating price

$$\underbrace{\eta-(1-\alpha)}_{\substack{\text{Current profit}\\ \text{if separating}\\ =\Pr(\sigma_1|\eta)[2f_1(\eta)-1]}} + \underbrace{\beta EV^1(\eta)}_{\substack{\text{Continuation value}\\ \text{if separating}}},$$

or, equivalently

$$2f_0(\eta)-\eta-\alpha \gtrless \beta\left[EV^1(\eta)-V^1(\eta)\right] \tag{13}$$

for $\eta \geq \overline{\eta}^1$. In this region the left hand side of (13) is increasing in η, while the right hand side is decreasing in η, so that there exists a unique cutoff level of the public belief $\overline{\eta}^2$ above which the monopolist strictly prefers to charge the pooling price and below which the separating price is charged.

Then the optimal pricing policy in the next to last period (2) is

$$P^2(\eta) = \begin{cases} \text{any } > 2f_1(\eta)-1 & \text{for } \eta \leq \underline{\eta}^2 \\ P_S(\eta) = 2f_1(\eta)-1 & \text{for } \underline{\eta}^2 \leq \eta \leq \overline{\eta}^2 \\ P_P(\eta) = 2f_0(\eta)-1 & \text{for } \eta \geq \overline{\eta}^2 \end{cases}$$

Notice that for $\underline{\eta}^2 \leq \eta \leq \overline{\eta}^1$ there are two motives for charging $P_S(\eta)$ instead of $P_P(\eta)$ since the separating price yields

1. a higher instantaneous payoff.
2. a higher future expected payoff, since $V^1(.)$ is a convex function.

6.1.3 Solution to the Dynamic Programming Problem

In this subsection we report on the solution of the dynamic programming problem of the monopolist.[2] The first result of the infinite horizon problem is that when the belief is larger than a threshold $\overline{\eta} \in (0,1)$, learning is stopped by the monopolist, and the pooling price is charged thereafter. To gain intuition note that as η tends to 1, the pooling price converges to the separating price, but allows firm 1 to sell with probability 1, instead of α. The value of continuing learning goes to 0 as η tends to 1, so that for large enough η the higher current profits from pooling customers dominates the future expected gain from more learning.

For beliefs lower than a threshold $\underline{\eta} \in (0, \overline{\eta})$ firm 1 exits the market. Clearly when the belief η is sufficiently low, for $\eta < \underline{\eta}$, the firm has a strictly negative value from operating, so that it will prefer to quit this market.

In a period t of the first "learning" phase, the monopolist charges the separating price $P_S(\eta^t) = 2f_1(\eta^t) - 1$, high relative to the public belief that the good is superior η^t. If still in the learning phase, the price in next period increases to $P_S(f_1(\eta^t)) = 2f_1(f_1(\eta^t)) - 1$ if the good was sold in the previous period, or decreases to $P_S(f_0(\eta^t)) = 2\eta^t - 1$ if the good was not sold. Eventually, the monopolist will almost surely stop the social learning process by either exiting the market or capturing it entirely. The monopolist will exit the first time ($\underline{t}$) that the belief η^t is below the cutoff level $\underline{\eta}$. This happens when the difference between the number of periods in which the monopolist was able to sell and those in which it was not is negative enough. Otherwise, the entire market will be captured by reducing deterministically the price to the separating level $P_S(\eta^t) = 2f_0\left(\eta^{\overline{t}}\right) - 1$, low relative to the belief, from the first time ($\overline{t}$) that the belief crosses the cutoff level $\overline{\eta}$. This happens when the difference between the number of periods in which the monopolist was able to sell and those in which it was not is large enough.

6.2 Empirical Implications

This model of monopoly predicts that in a first phase after the introduction of the new good prices will be set high enough so that consumers with different signals buy different goods. Along the learning process, prices increase at a decreasing rate as the share of sales in the past increases. If instead the

[2] For a complete derivation of these and other related results see Ottaviani (1995c).

share of sales in the past is reduced, prices decrease at an increasing rate. Eventually learning will stop, because the monopolist decides either to exit the market if it is not capable of selling for a large enough number of times, or to capture the whole market by reducing the price if has sold much in the past so that its variety of the good is perceived as superior.

Ottaviani (1995c) shows that the stochastic process of prices optimally set by the monopolist is a supermartingale, when not conditioning on the knowledge of the state of the world. This means that the price is expected to decrease, unconditional on the state of the world. Conditioning instead on the state being ω_1, the price sequence is increasing in the learning phase, and conditioning on state ω_0 the price sequence is decreasing in the learning phase. When going from the learning to the pooling phase the price decreases deterministically.

6.3 Efficiency Analysis

"Which of these systems is likely to be more efficient depends mainly on the question under which of them we can expect that fuller use will be made of the existing knowledge. Thus, in turn, depends on whether we are more likely to succeed in putting at the disposal of a single central authority all the knowledge which ought to be used but which is initially dispersed among many different individuals, or in conveying to the individuals such additional knowledge as they need in order to enable them to dovetail their plans with those of others " (Hayek, 1945).

Do superior goods prevail in the long run? In the herding model the decisions of initial consumers affect the behavior of the following ones. This information externality is at the origin of the long run inefficiency occurring when all but a finite number of consumers choose the wrong good in the model of section 2. If the best good is 1, the inefficiency arises when the belief hits the lower threshold $1-\alpha$ before reaching the upper one α. Similarly conditional on good 0 being the best one, all but a finite number of consumers buy good 1 when the belief hits α before $1-\alpha$.

In the model studied in this section the property right for one good is given to a firm that is allowed to change the price along the social learning path. Intuitively, this agent will internalize the externality when it is unfavorable to its own good, but not when it is favorable. Therefore it might seem that the allocation of property rights should reduce the inefficiency associated with that good, but increase the inefficiency that arises when good 1 is sold even if 0 is the best good. The allocation of property rights on the good brings in another inefficiency due to the monopoly position of firm 1 vis-a-vis its customers. Firm 1 extracts more rents from the consumers by delaying the capture of the entire market. This is inefficient in the short run because the high price of good 1 discourages type-0 consumers from buying good 1, even if expected valuation for good 1 is higher than for good 0. But this pricing strategy has the desired effect of increasing the region of learning, so that the other long run inefficiency tends to be reduced as well. In this subsection we follow Ottaviani (1995c) to argue that the expected inefficiency in the monopoly model is lower than in the canonical model. We also refer to that paper for the comparison of the solution of the monopolist problem with the constrained social optimum allocation that results when the social planner needs to use the price system to aggregate information.

Consider the effect of the increase in $\overline{\eta}$ on the inefficiency "favorable" to the monopolist – good 1 is sold when good 0 is superior. This inefficiency clearly decreases because it becomes more difficult that the higher barrier $\overline{\eta}$ is hit before the lower barrier. The effect of a decrease in $\underline{\eta}$ is to increase this type of inefficiency, so that the two effects go in opposite direction. Since $\overline{\eta} > \alpha$ and $\underline{\eta} < 1 - \alpha$, the effect of strategic pricing by one firm on the inefficiency favored by firm 1 is ambiguous in general. To quantify the change in this inefficiency it is necessary to consider the quantities $[\overline{\eta} - \alpha]$ and $\left[(1-\alpha) - \underline{\eta}\right]$. Similarly the effect of strategic pricing by one firm on the other inefficiency is ambiguous in general.

Nonetheless, it is possible to sign the change in expected inefficiency due to the increase in $\overline{\eta}$. Note that the "unfavorable" inefficiency is instead increased by an increase in the upper barrier $\overline{\eta}$. The effect on the favorable inefficiency is much stronger than the one on the unfavorable one, due to the drift toward higher η when the best good is 1 and toward lower η when the best good is 0. Similarly for the effect of the decrease in $\underline{\eta}$: the reduction of the unfavorable inefficiency is much larger than the increase in the favorable one due to the drift. This is the basic intuition behind the result, obtained in Ottaviani (1995c), that the ex-ante expected inefficiency is reduced by strategic pricing by one firm for initial belief η^1 close enough to $\frac{1}{2}$.

We remark that, for $\beta = 0$, $\overline{\eta} = \overline{\eta}^1 > \alpha$ and $\underline{\eta} = 1 - \alpha$. When firm 1 is completely myopic the inefficiency favorable to the firm decreases, while the other inefficiency increases because of endogenous pricing. For $\beta = 0$ firm 1 is a short run player and exploits its monopoly power to extract more surplus from the consumers.

In summary, the expected long run inefficiency is shown to decrease due to two different effects. First, the unfavorable externality is partially internalized by the firm. Second, the monopoly position allows the firm to delay the capture of the entire market in order to sell at a higher price. This second effect works to decrease the other inefficiency as well, so that the expected inefficiency is lower (when the prior belief is not too different from $\frac{1}{2}$) when compared with the standard case. The remaining unfavorable inefficiency due to incomplete learning is standard in learning models and due to the impatience of the long-run player.[3]

6.4 Extensions, Examples and Applications

The model can be easily extended to the case of foresighted consumers who can wait to purchase. The cost of waiting for the consumer derives from the loss of surplus due to the delay in the consumption of the good. In the model with fixed prices, by waiting the consumer might be able to use the information revealed by others in the meantime, and prices are decreasing on average. Ottaviani (1995c) studies conditions that guarantee that the pricing strategy of the monopolist is such that no consumer would like to delay her purchase decision.

This model provides a story for why and how market share matters. Selling can be valuable even at a loss because it might convey valuable information to the market. Some interesting implications for industrial policy can also be drawn from this model. A novel argument for why pricing below marginal cost might increase efficiency is provided. Long-run inefficiency can be reduced by entitling a monopolist to the property rights to one good. The monopolist achieves its objective by pricing below marginal cost when its good is perceived as inferior (for $\eta < 1 - \alpha$).

We now turn to a brief discussion of some possible applications of this models to real-world markets. This model is consistent with the price of books and compact disks being reduced when they become best sellers, but not before then. Popularity at separating prices (i.e. high relative to the belief) let the information flow from past consumers to future potential customers. In particular this models predicts also introductory discounts, reduced if the product is successful and increased if it is not. On average prices are predicted to be decreasing, even though the price of a superior product

[3] See Aghion, Bolton, Harris and Jullien (1991).

is predicted to be increasing on average in the learning phase.

Similarly, this model is consistent with young independent professionals (doctors and lawyers, for instance) charging fees that are high relative to the perceived quality and being willing to be underemployed, but not to reduce the price for their services. It might be argued that fees of older professionals are higher than those of younger ones, even when the different experience, that changes the characteristics of the service provided, is taken into account. But this implication is in line with our model, since not everyone survives in this markets, and higher quality professionals are more likely to succeed. For valuable professionals this model predicts increasing prices in the learning phase.

An application to information externalities in the labor market has been undertaken by Ottaviani (1995b). The monopolist of this model is a worker who is learning about her intrinsic quality and setting each period a reservation wage at which the current employer decides to hire or not. The history (past wages and employments) is recorded in the *curriculum vitae* of the worker. Future employers look at the decisions of previous employers in an attempt to infer their information.

7 Duopoly Equilibrium Prices

In this section we consider the case of competition among two firms, one for each sector. The equilibria of the one-period game are characterized and compared to the equilibria of the two-period game, following the analysis of Moscarini and Ottaviani (1994a).

Two duopolists (firm 0 and firm 1) compete to sell two different varieties of the good. At each period of time each firm quotes simultaneously a price for the unit of its own variety of the good and supplies it whenever the buyer demands it.

In such dynamic games, there is a host of equilibria with strategies that depend in a complicated way on anything observable that occurred in the past, whether or not directly relevant to the current and future play of the game. To retain predictive power we restrict players to adopt Markov strategies and look for Markov perfect Bayesian equilibria (MPBE) (see Maskin and Tirole (1993)). We also require that a firm that does not sell in equilibrium be indifferent between selling and conceding to the competitor. This natural requirement of "cautiousness" is also made by Bergemann and Valimaki (1993) to avoid aggressive non-selling prices that only reduce the profit of the competitor.

In Section 7.1 we solve for the equilibrium correspondence of the one period game with no continuation. In Section 7.2 we add one period backward in order to analyze the two-period game. At each stage, the Markov perfect Bayesian equilibria (or MPBE) are found by backward induction: firms in the first stage know the consequences of their price bids not only in terms of current sales, but also in terms of the effect on the public belief and therefore on the type of possible equilibria in the next stage.

As part of the solution we find the path of the lowest price that each firm can profitably quote in each period in order to gain the customer and pull the public belief in the direction of its own good. The continuation value of gaining one customer brings this lower bound of prices below the marginal cost, justifying current losses in terms of future market share.

7.1 One-Period Game

In this section we report the equilibria of the one-period pricing game starting with an arbitrary belief η. It can be shown that (for $\alpha \geq \frac{2}{3}$) when the belief is intermediate,

$$\eta \in \left[\underline{\eta}_S(\alpha), \overline{\eta}_S(\alpha)\right],$$

with $\underline{\eta}_S(\alpha) \equiv \frac{(1-\alpha)^2}{\alpha(2\alpha-1)}$ and $\overline{\eta}_S(\alpha) = 1 - \underline{\eta}_S(\alpha)$, the unique equilibrium is a pure strategy separating equilibrium in which each firm sells to her own customer at prices $P_0 = 1 - f_0(\eta)$ and $P_1 = f_1(\eta)$.

When the belief η is high enough, the unique cautious MPBE is a pure strategy pooling equilibrium on good 1, in which firm 1 sells to both customer types at a price equal to $P_1 = 2f_0(\eta) - 1$, and firm 0 does not sell to any customer and announces $P_0 = 0$. This happens for $\eta > \overline{\eta}_P(\alpha)$, where $\overline{\eta}_P(\alpha) = \overline{\eta}_1$ (cf. (11)). Note that $\overline{\eta}_P(\alpha) > \overline{\eta}_S(\alpha)$. Symmetrically a pooling equilibrium on good 0 exists if and only if $\eta < \underline{\eta}_P(\alpha)$ where $\underline{\eta}_P(\alpha) = 1 - \overline{\eta}_P(\alpha)$.

In the remaining regions, i.e. for

$$\eta \in (\overline{\eta}_S(\alpha), \overline{\eta}_P(\alpha)] \cup [\underline{\eta}_P(\alpha), \underline{\eta}_S(\alpha)),$$

there are no pure strategy equilibria. By endogenizing the sharing rule there exists a mixed strategy Nash equilibrium, as can be shown by applying the results of Simon and Zame (1990).

In the basic model without prices (cf. Section 2) the cascade regions are $(\alpha, 1]$ and $[0, 1-\alpha)$. Identifying pooling equilibria with cascades of the basic model, and separating and mixed strategy equilibria with no-cascades, we notice that the set of public beliefs leading to a cascade for each α shrinks in this duopoly model with respect to the basic model. Furthermore, it shrinks in a nonlinear way: as the precision of the signal (α) increases, the Lebesgue measure of the cascade regions falls at a rate which is decreasing but is initially (for $\alpha > \frac{1}{2}$) higher than 1.

7.2 Two-Period Game

We add one period backward to the game, so that the second and last period subgames, for any public belief η that follows from the decision of the first customer according to the learning dynamics (9) or (10) (with apex running backward and representing the number of periods left to go), are the one-period games $\Gamma^1(\eta, \alpha)$. In the previous subsection we have characterized the outcomes and payoffs of the equilibria of the possible continuations in the last period for any belief $f_i(\eta)$ produced in the first period by Bayesian updating. The requirement of Markov perfection implies that the continuations of a MPBE for $\Gamma^2(\eta, \alpha)$ follow an equilibrium of $\Gamma^1(\eta^1, \alpha)$ both after a deviation and along the equilibrium path.

We will number periods backward, from the last (1) to the first (2). The initial belief is η. Consider the second period subgame. If signals are not deduced in the first period (2) because customer 2 (the first customer) pooled in equilibrium on either good, then the second period subgame is $\Gamma(\eta, \alpha)$, examined in the previous section. If, instead, the first period signals are revealed, then there are two symmetric situations to be considered, depending on the realized signal σ_i^2 in the first of the two periods.

How does the first stage of the two-period game compare with the one-period game? In order to study the emergence of informational cascades and the effects of the information externality due to social learning of the consumers on dynamic pricing behavior, we study how the continuation of the game alter the equilibrium strategies of the players.

In particular we are interested in the dynamics of the pooling equilibrium region. Moscarini and Ottaviani (1994a) show that the pooling equilibrium region of the first stage of the two-period game is smaller than the pooling equilibrium region of the second period stage game. For this they need to characterize the endogenous lower bound on the price that firm 0 can cautiously post when 2 periods are left to go.

The minimum price that firm 0 can profitably quote to avoid a pooling equilibrium in the next to last period is the price $\underline{P}_0^2$ such that the no deviation condition from the pooling equilibrium for firm 0 holds with equality,

$$0 + \beta V_0^1(\eta) = \Pr(\sigma_0|\eta)\left[\underline{P}_0^2 + \beta V_0^1(f_0(\eta))\right] + \beta V_0^1(f_1(\eta)) \Pr_{\sigma_1}(\eta) \qquad (14)$$

where $V_0^1(\cdot)$ is the continuation payoff of firm 0 in the last stage. Given that firm 1 is quoting a price slightly below $P_1 = 2f_0 - 1 + P_0$, both customers buy good 1. Then, for firm 0, the only possible deviation is to a lower price. $\underline{P}_0^2$ is exactly the price at which firm 0 has no incentive to undercut further, because the expected payoffs from such a separating deviation (on the right hand side of (14)) would fall below the payoff from pooling (on the left hand side). Rearranging (14) we obtain that firm 0 prices are bounded below in equilibrium and equal to

$$\underline{P}_0^2 = \beta \frac{V_0^1(\eta) - \Pr(\sigma_1|\eta) V_0^1(f_1)}{\Pr(\sigma_0|\eta)} - \beta V_0^1(f_0(\eta)). \qquad (15)$$

Under the claim that the pooling equilibrium region enlarges as the periods left to go decrease, if there exists a pooling equilibrium at belief η at the next to last stage 2, there will still be a pooling equilibrium on the same good in the next period, because in a pooling equilibrium the belief remains unchanged. This situation corresponds to the endless cascade of the infinite horizon model. A fortiori, there will be a cascade on good 1 if there is separation and the signal σ_1 is drawn, as $f_1(\eta) > \eta$. Since firm 0 gets an immediate payoff of 0 in a pooling equilibrium, we can conclude that

$$V_0^1(f_1(\eta)) = V_0^1(\eta) = 0$$

when η is in the pooling region, so that by (15) the price of the non-selling firm becomes

$$\underline{P}_0^2 = -\beta V_0^1(f_0(\eta)). \qquad (16)$$

As expected, the lower bound to firm 0 undercutting is a negative number, as the signal σ_0 leads the belief to a level more favorable to firm 0.

Moscarini and Ottaviani (1994a) show that the no-deviation condition for a pooling equilibrium in the next to last period is not satisfied so that the set of beliefs for which there exists a pooling equilibrium on good i in the first stage of the two-period game is strictly contained in the analogous set of the one-period game. If the equilibrium in the first stage is pooling on good i, the equilibrium in the second stage is also pooling on good i.

The effect of endogenous pricing and competition is to reduce the cascade region. The issue of inefficiency can be addressed as in Section 6.3. That analysis could be extended to the more general framework of this duopoly model and allow us to conclude that endogenous pricing reduces the long run expected inefficiency associated with social learning.

8 Conclusion

The outcome of the interaction among rational agents who have dispersed information can be dramatic when information externalities are at work. We have presented some models of Bayesian social learning which recently appeared in the economics literature. We have discussed the robustness of the herding results. The focus of our analysis has been on the problems of information acquisition, stationarity of the environment and endogenous pricing.

We have shown how individual experimentation interacts with social learning. Individuals are allowed to acquire information through experimental consumption. The observation of the behavior of other consumers provides an additional source of valuable information. Once a cascade has started the valuable information acquired through experimental consumption is not transmitted to others, since experimentation is not strong enough to change the decision of the experimenter.

If the state of the world is allowed to change, then only temporary informational cascades can arise and, if the change of state is very unpredictable, no cascade ever arises. The familiar Martingale properties of the public belief (unconditional on the true state of the world) and of the likelihood ratio (conditional on the state of the world) do not hold when the world is changing.

The price system is shown to reduce the long-run inefficiencies caused by the informational externality. Free entry and perfect competition among producers of the same good prevent dynamic pricing strategies, that would enlarge the learning region and thereby would promote long-run efficiency. Assigning monopoly power to a far-sighted firm can increase efficiency. "Fundamentally, in a system in which the knowledge of the relevant facts is dispersed among many people, prices can act to co-ordinate the separate actions of different people in the same way as subjective values help the individual to co-ordinate the parts of his plan" (Hayek, 1945).

Finally, the discussion of the applications proposed sheds light on problems that have not been previously analyzed within the social learning framework.

9 References

Aghion, Philippe, Patrick Bolton, Christopher Harris and Bruno Jullien, 1991, Optimal Learning by Experimentation, *Review of Economic Studies*, 58, 621-654.

Banerjee, Abhijit V., 1992, A Simple Model of Herd Behavior, *Quarterly Journal of Economics*, 107, 797-817.

Banerjee, Abhijit V. and Drew Fudenberg, 1995, Word-of-Mouth Learning, MIT Mimeo.

Becker, Gary S. A Note on Restaurant Pricing and Other Examples of Social Influences on Price, 1991, *Journal of Political Economy*, 99, 1009-1016.

Bergemann, Dirk and Juuso Valimaki, 1993, Learning and Strategic Pricing, University of Pennsylvania Mimeo.

Bikhchandani, Sushil, David Hirshleifer and Ivo Welch, 1992, A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades, *Journal of Political Economy*, 100, 992-1026.

Bolton, Patrick and Christopher Harris, 1993, Strategic Experimentation, LSE Discussion Paper TE/93/261.

Ellison, Glenn and Drew Fudenberg, 1995, Rules of Thumb for Social Learning, *Quarterly Journal of Economics*, 110, 93-125.

Grossman, Sandford, Richard E. Kihlstrom, and Leonard Mirman, 1977, A Bayesian Approach to the Production of Information and Learning by Doing, *Review of Economic Studies*, 44, 533-547.

Hayek, Friedrich A., 1945, The Use of Knowledge in Society, *American Economic Review*, 35, 515-530.

Keller, Godfrey and Sven Rady, 1995, Optimal Experimentation in a Changing Environment, Discussion Paper, Department of Economics, The University of Edinburgh.

Lee, In Ho, 1993, On the Convergence of Informational Cascades, *Journal of Economic Theory*, 61, 395-411.

Maskin, Eric and Jean Tirole, 1993, Markov Perfect Equilibrium, Harvard University Mimeo.

Milgrom, Paul, 1981, Good News and Bad News: Representation Theorems and Applications, *The Bell Journal of Economics*, 380-391.

Moscarini, Giuseppe and Marco Ottaviani, 1994a, Social Learning and Competition, MIT Mimeo.

Moscarini, Giuseppe and Marco Ottaviani, 1994b, Social Learning in a Changing World, MIT Mimeo.

Ottaviani, Marco, 1995a, Individual Experimentation and Social Learning, MIT Mimeo.

Ottaviani, Marco, 1995b, Information Externalities in the Labor Market, MIT Mimeo.

Ottaviani, Marco, 1995c, Social Learning and Monopoly Pricing, MIT Mimeo.

Rothschild, Michael, 1974, A Two-Armed Bandit Theory of Market Pricing, *Journal of Economic Theory*, 9, 185-202.

Simon, Leo and William Zame, 1990, Discontinuous Games and Endogenous Sharing Rules, *Econometrica*, 58, 861-872.

Smith, Lones and Peter Sørensen, 1994, Pathological Outcomes of Observational Learning, MIT Mimeo.

Stokey, Nancy and Robert Lucas, Jr., 1989, *Recursive Methods in Economic Dynamics*, Harvard University Press.

Vives, Xavier, 1995, Learning from Others: A Welfare Analysis. Institut d'Anàlisi Econòmica Mimeo.

Vives, Xavier, 1996, Social Learning and Rational Expectations. Forthcoming in *European Economic Review*.