GAMES AND ECONOMIC BEHAVIOR **20,** 177–200 (1997)
ARTICLE NO. GA970562

Learning from Others: A Welfare Analysis*

Xavier Vives

Institut d'Anàlisi Econòmica, CSIC, Barcelona, Spain

Received June 20, 1995

This paper considers a smooth and noisy version of the statistical prediction model studied in the herding/informational cascades literature and compares market and optimal learning. The latter is characterized by defining a decentralized welfare benchmark as the solution to an infinite horizon *team* problem. Market behavior involves herding, in the sense that agents put too little weight on their private information for any given precision of public information, and yields underinvestment in the production of public information. However, both market and optimal learning involve *slow* learning. Examples of the model include learning by doing, reaching consensus, and consumer learning about quality. *Journal of Economic Literature* Classification Numbers: D82, D83. © 1997 Academic Press

1. INTRODUCTION

Recent work has studied learning from others in the context of a simple statistical model where agents have to predict in sequence as accurately as possible a random variable (a "correct" option). Apparently surprising outcomes may emerge from the decisions of rational (Bayesian) agents who may end up disregarding their own private information and taking actions based only on public information (a situation referred to as an *informational cascade* by Bikhchandani *et al.*, 1992). A result is that the whole sequence of agents may "herd" and choose a "wrong" action (Banerjee, 1992).[1]

*A preliminary version of this paper was circulated under the title: "Learning from Others." I am grateful to Jordi Caballé, Ramon Caminal, David Easley, Guy Laroque, Meg Meyer, an anonymous referee, and an associate editor, as well as seminar participants at the universities of Chicago, Harvard, MIT, Northwestern and Pennsylvania, for helpful comments. The support of the Spanish Ministry of Education through DGICYT Grants PB90-0132 and PB93-0679 and from the Generalitat de Catalunya, CIRIT, is gratefully acknowledged. Lluís Bru provided excellent research assistance.

[1] For related literature see Scharferstein and Stein (1990), Ellison and Fudenberg (1993), Kirman (1993), and Smith and Sørensen (1995 and 1996).

0899-8256/97 $25.00

Herd behavior needs payoffs to be nonregular (degenerate as in Banerjee, 1992) or action spaces to be not rich enough to allow for appropriate responses to information (discrete as in Bikhchandani *et al.*, 1992). Otherwise, with continuous action spaces (containing potentially optimal actions) and regular payoffs convergence to the correct action obtains (see Banerjee, 1992, and Lee, 1993). Further, in a smooth and noisy context, work on "learning in rational expectations" shows that prices *slowly* reveal an unknown parameter in a model where firms learn from the aggregate information of other firms (Vives, 1993).[2] In other words, convergence obtains but at a slow rate.

The received literature leaves open the analysis of the welfare issue. It is easy to see that the outcomes may be inefficient from the welfare point of view given the externality present in learning from others. Indeed, an agent when responding to her private information will not take into account the future benefit to other agents of her action. This leads naturally to a welfare-based definition of herding as a departure of the weight agents put on private information from optimal behavior according to an appropriately defined welfare benchmark. There will be herding if agents rely too much on public information (putting too little weight on their private information).

In the present paper we consider a model of learning from others which is a *smooth* and *noisy* version of the statistical prediction model used in the herding literature. The presence of noise in the public signals means that the latter are not sufficient statistics of the information of agents. The model is infinite-horizon and the interaction of agents in the market produces a public aggregate noisy statistic of the agents' actions. Agents, based on private and public information, strive to minimize the (mean square) prediction error of a random variable θ.

Equipped with this model we first review market (Bayesian) learning behavior and then we analyze socially optimal learning. We take as the social optimum benchmark the *team solution* which assigns decision rules to agents to minimize the discounted sum of prediction errors taking as given the decentralized private information structure of the economy. This decentralized benchmark, in the tradition of team theory (Radner, 1962), ignores possible incentive problems and focuses on coordinating decisions to minimize the discounted prediction errors.

[2] This work is a continuation of the tradition of studying aggregation of information by prices and rational expectations models (Grossman, 1976; Grossman and Stiglitz, 1980; Hellwig, 1980) and their dynamic extensions (see, for example, Townsend, 1978; Kyle, 1985; Feldman, 1987; and Jordan, 1992a). A rate of convergence result for a class of myopic Bayesian processes has been obtained by Jordan (1992b).

The market solution involves an *information externality*: An agent, when responding to his private information, does not take into account the benefit produced in terms of increased informativeness of public information in the future (which translates into lower future prediction errors or period losses).[3] The team solution does internalize this externality and involves experimentation. The team manager knows that by increasing the response of agents to private information short-term losses also increase, but there is the benefit of the additional public information gained.[4]

The analysis of the market solution confirms previous results (Vives, 1993) on slow convergence in the context of the prediction model. Indeed, it is shown that predictions do converge to the correct action θ, but they do it *slowly*.

The solution to the welfare program (the infinite horizon stochastic team optimization problem) is characterized using dynamic programming methods under the assumption that the team manager assigns linear decision rules to agents. The result is that, provided the discount rate is finite, the *market underinvests in information* in two precise senses. First, for a given accumulated informativeness of public information, the team solution calls for a larger response of agents to private information than the market. That is, *agents "herd" or rely too much on public information* according to the welfare benchmark. Second, at any period the informativeness of public information is larger at the team solution (that is, the team solution accumulates more public information than the market). Furthermore, simulations uncover that the relative welfare loss of the market with respect to the team solution can be quite large. Nevertheless, *the market does not distort the (asymptotic) speed at which agents learn*.

In order to motivate the model think of a situation of investment with macroeconomic uncertainty represented by the random variable θ which will determine the profitability of investment. At each period there is an independent probability $1 - \delta$ that the uncertainty is resolved. Firms invest by taking into account that the profits of their accumulated investment will depend on the realization of θ. The investment of a firm will be

[3] The literature contains many instances of analysis of information externalities. Some examples are Stein (1987), Rob (1991), Caplin and Leahy (1994), and Chamley and Gale (1994).

[4] The trade-off between short-term profit maximization and information accumulation is typical of optimal learning models. There are many examples in the literature where optimal learning procedures, precisely because of this trade-off, lead only to partial or no learning (see, for example, Rothschild, 1974; Grossman *et al*., 1977; McLennan, 1984; Kihlstrom *et al*., 1984; Easley and Kiefer, 1988; Nyarko, 1991; Aghion *et al*., 1991; and Bolton and Harris, 1993).

directly linked to its prediction of θ. To predict θ each firm will have access to a private signal as well as to public information, which is formed by aggregate past investment figures compiled by a government agency. Data on aggregate investment incorporates measurement error.[5] In consequence, at each period a noisy measure of aggregate investment of the last period is made public. The issue is whether, and if so, how fast, the repeated announcement of the aggregate investment figures reveals θ. The result according to our paper is that information is revealed slowly and that individual firms put too little weight on their private information from a social point of view.

At a more abstract level we could think about the reaching of consensus in the *common knowledge* literature. It has been shown that repeated public announcements of a stochastically monotone aggregate statistic of conditional expectations, which need not be common knowledge, lead to consensus (McKelvey and Page, 1986, and Nielsen *et al*., 1990, following up on Aumann, 1976).[6] In the model presented in this paper, repeated public announcements of a linear *noisy* function of agents' conditional expectations leads to consensus, but slowly. We could say, rephrasing a result in the literature (Geanakoplos and Polemarchakis, 1982), that in the presence of noisy public information "we cannot disagree forever but we can disagree for a long time."

The stylized statistical model presented can also provide a statistical microfoundation in a world of private information of the *learning curve* and can be extended to encompass *short-lived agents*. An application of the latter model to consumers learning about the quality of a product is presented. The model with short-lived agents has a direct link with the usual sequential herding/informational cascades model.

The plan of the paper is as follows. The model is presented in Section 2. Section 3 reviews market learning. Section 4 studies optimal learning and compares it to the market solution. Section 5 considers short-lived agents. Concluding remarks follow.

[5] For example, quarterly data on national accounts are subject to measurement error.

[6] In the iterative process, individuals compute conditional expectations with the information they have available and the aggregate statistic is announced. Individuals then recompute their expectations on the basis of their private information plus the new public information, and the process continues. The aggregate statistic is supposed to represent the outcome of the interaction of agents in a reduced form way. In many instances market interaction will nevertheless provide agents with only a noisy version of an aggregate statistic of individual conditional expectations due to the presence of noise in the communication channels, random traders, demand or supply shocks, etc.

2. THE MODEL

Consider an infinite horizon model $n = 0, 1, \ldots,$ with a continuum of long-lived agents indexed in the unit interval $[0, 1]$ (endowed with the Lebesgue measure). Let θ be a random variable, unobservable to the agents, normally distributed with mean $\bar{\theta}$ (set equal to 0 to simplify notation) and finite variance σ_θ^2 $(\theta \sim N(\bar{\theta}, \sigma_\theta^2))$. The payoff to an agent in period n when choosing an action q_{in} is $-(\theta - q_{in})^2$. In period n the agent has available an information vector I_{in} (to de described shortly). At any period there is an independent probability $1 - \delta$, with $1 > \delta \geq 0$, that θ is realized and the payoffs up to this period collected. Let $L_{in} = E(\theta - q_{in}(I_{in}))^2$.

The objective of agent i is then to minimize $\sum_{n=0}^{\infty} \delta^n L_{in}$, choosing a sequence of functions $\{q_{in}(\,.\,)\}_{n=0}^{\infty}$ dependent (respectively) on the information vectors $\{I_{in}\}_{n=0}^{\infty}$.

Agent i has available at the start a *private signal* $s_i = \theta + \varepsilon_i$, where $\varepsilon_i \sim N(0, \sigma_\varepsilon^2)$. The convention is made that errors on average cancel out: $\int_0^1 \varepsilon_i \, di = 0$ almost surely (a.s.).[7] In consequence, the average signal reveals θ a.s. Apart from their private information agents in period n observe also a *public information* vector: $p^{n-1} = \{p_0, \ldots, p_{n-1}\}$, where p_t is a noisy version of the average action of agents in period $t = 0, 1, \ldots, n - 1$, $p_t = \int_0^1 q_{it} \, di + u_t$. The sequence of random variables $\{u_n\}_{n=0}^{\infty}$ is a normal white noise process. In summary, agent i in period n has available the information vector $I_{in} = \{s_i, p^{n-1}\}$.

Myopic behavior, that is, agent i in period n trying to minimize the mean square error in predicting the random variable θ, is (individually) optimal in the present framework. This is so because the agent is infinitesimal and he cannot affect the public statistics. That is, agent i in period n solves the problem

$$\text{Min}_q \; E\left\{(\theta - q)^2 | I_{in}\right\}.$$

As is well known, the solution to this problem is $q_{in} = E(\theta | I_{in})$.

3. A REVIEW OF MARKET LEARNING

In this section I analyze the convergence properties of public information in the model presented. The strategy of the analysis is to find a sufficient statistic for public information and characterize its convergence

[7] See Feldman and Gilles (1985) for the measure-theoretic issues involved.

properties. An outline of the analysis is presented here; for details see the Appendix and for related results and proofs see Vives (1993).

Normality and the recursive structure of the model imply that $q_{in} = E(\theta | I_{in})$ is a linear function of s_i and $p^{n-1} = \{p_0, \ldots, p_{n-1}\}$. It follows that $\theta_n = E(\theta | p^n)$ is a sufficient statistic for public information and follows a martingale: $E(\theta_n | \theta_{n-1}) = \theta_{n-1}$. Let τ_n denote the *informativeness* (precision) *of public information* in the estimation of θ. That is, $\tau_n = (\mathrm{Var}(\theta | p^n))^{-1}$. The optimal action $q_n(I_{in})$ given information $I_{in} = \{s_i, p^{n-1}\}$ is then equal to $E(\theta | s_i, \theta_{n-1}) = \alpha_n s_i + (1 - \alpha_n)\theta_{n-1}$, with $\alpha_n = \tau_\varepsilon / (\tau_\varepsilon + \tau_{n-1})$ and $\tau_n = \tau_\theta + \tau_u \sum_{t=0}^{n} \alpha_t^2$.

The posterior mean of θ is a weighted average of the signals of the agents with weights according to their precisions (the private signal with precision τ_ε and the public with precision τ_{n-1}). Note that the "discount" δ plays no role in the market solution since myopic behavior is individually optimal.

The induced dynamics in τ_n are driven by a *self-correcting property* of learning from others whenever agents are imperfectly informed and public information is not a sufficient statistic of the information agents have (Vives, 1993, and Smith and Sørensen, 1995). Indeed, the weight given to private information α_n is decreasing in the precision of public information τ_{n-1}, and the lower α_n is the less information is incorporated in the public statistic p_n. Learning is self-defeating: a higher inherited precision of public information τ_{n-1} induces a low current response to private information α_n, which in turn yields a lower increase in public precision. Conversely, learning is self-enhancing: a lower inherited τ_{n-1} induces a high current α_n, which in turn yields a larger increase in τ_n. The self-enhancing aspect means that public precision τ_n will be accumulated unboundedly. If this were not the case the weight given to private precision would be bounded away from zero, necessarily implying that τ_n grows unboundedly, a contradiction.[8] The self-defeating aspect means that accumulation is slow. As τ_n tends to infinity, α_n tends to zero and the signal-to-noise ratio in the new information about θ in p_n, $\alpha_n \theta + u_n$, worsens. The amount of new information incorporated in p_n is asymptotically vanishing. The result is that τ_n tends to infinity with n not at the usual linear rate but more

[8] Similarly, in Banerjee and Fudenberg (1995) convergence to efficiency is obtained when people use samples larger than one because this allows the possibility of "mixed" samples which are relatively uninformative and consequently induces agents to rely on their private information and enhance the information flow into the system. This is again the self-enhancing aspect of learning from others.

slowly, at the rate of $n^{1/3}$. (See Fig. 1 below with $\delta = 0$.)[9] A linear rate for increase in precision obtains in the benchmark case of i.i.d. noisy observations of θ.

Let L_n stand for the (ex-ante expected) loss of a representative agent in period n. That is, $L_n = E(\theta - E(\theta | I_{in}))^2$. Note that L_n is independent of i since both the information structure and the strategies $q_n(I_{in}) = E(\theta | I_{in})$ are symmetric. It follows that $L_n = \mathrm{Var}(\theta | I_{in}) = (\tau_\varepsilon + \tau_{n-1})^{-1}$. The period loss is decreasing in the precision of private (τ_ε) and public (τ_{n-1}) information. The following proposition states the results so far.

PROPOSITION 3.1. (i) *The response to private information* $\alpha_n \underset{n}{\to} 0$ *and the precision of public information* $\tau_n \underset{n}{\to} \infty$ as $n \to \infty$.

(ii) *The precision* τ_n *is of the order of* $n^{1/3}$ *and* α_n *is of the order of* $n^{-1/3}$. *The asymptotic precision is given by* $A\tau_\infty = \lim_{n\to\infty} n^{-1/3}\tau_n = (3\tau_u)^{1/3} (\tau_\varepsilon)^{2/3}$.

COROLLARY. *The order of magnitude of the loss* L_n is $n^{-1/3}$.

From Proposition 3.1 it is easy to infer that public information θ_n converges almost surely to θ (because its precision τ_n goes to infinity) at the rate of $1/\sqrt{n^{1/3}}$ (because τ_n is of the order of $n^{1/3}$). Furthermore, for a given rate, convergence is faster the larger τ_u and τ_ε are, since the

[9] The asymptotic precision $A\tau_\infty$ (the inverse of the asymptotic variance of $\theta_n - \theta$) controls the "slope" of increase of τ_n for a given order $n^{1/3}$. For n large, τ_n is approximately $\tau_\theta + A\tau_\infty n^{1/3}$. A larger precision of noise τ_u or of private information τ_ε decreases the "slope" of τ_n and slows down the increase in τ_n.

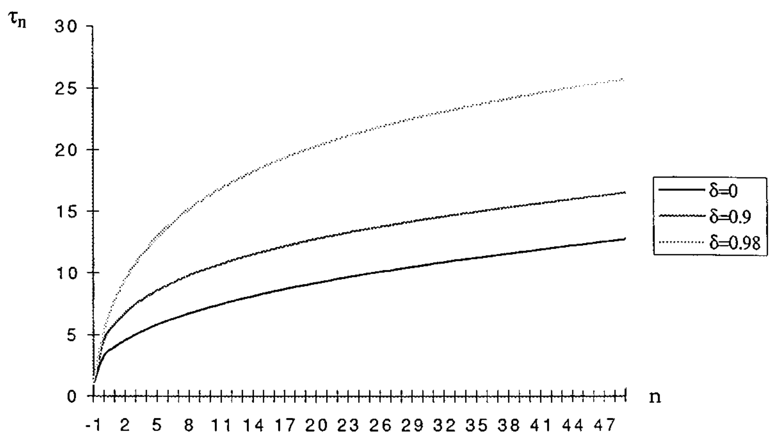

FIG. 1. Precision of public information and discount factors ($\tau_\theta = 1$, $\tau_\varepsilon = 2$, $\tau_u = 5$).

asymptotic variance of $\theta_n - \theta$, $(A\tau_\infty)^{-1}$, is decreasing in those parameters. The next proposition states the result formally.[10] Denote by $\xrightarrow{L}$ convergence in law (distribution).

PROPOSITION 3.2. (i) *$\theta_n \xrightarrow[n]{} \theta$ a.s., and*

$$\text{(ii) } \sqrt{n^{1/3}}(\theta_n - \theta) \xrightarrow{L} N(0,(A\tau_\infty)^{-1}), \text{ with } A\tau_\infty = (3\tau_u)^{1/3}(\tau_\varepsilon)^{2/3}.$$

4. OPTIMAL LEARNING

The market solution involves an *information externality*. Indeed, a larger response to private information at period n induces a larger precision of public information τ_n and a lower welfare loss next period $L_{n+1} = (\tau_\varepsilon + \tau_n)^{-1}$ (and in all subsequent periods for given responses to private information). Agents when deciding what action to choose do not take into account that the response to their private information determines the precision of the public information next period. This information externality induces agents to respond too little to their private information.

An appropriate welfare benchmark has to respect the private information structure of the economy and has to weigh the welfare losses of the different periods.

Welfare losses are discounted with discount factor $\delta, 1 > \delta \geq 0$ (consistently with the unknown θ having an independent probability of $1-\delta$ of being revealed in any period). I will take as a decentralized welfare benchmark the *team solution* where the planner is able to impose decision rules on agents but cannot manipulate the information flows (in particular, it has no access to their private information). Furthermore, the planner will be restricted to using linear rules, the same family of simple rules that agents in the market use. This means that the planner needs to convey to the agents only a number, the weight they should put on their private information. The requirement is appealing since it corresponds to a low level of complexity in the instructions of the team manager to the agents, communicating the value of one behavioral parameter. In summary, it is assumed that the planner can impose a *linear* decision rule on agent i as a function of her information vector, $q_{in}(I_{in})$, where, as before, $I_{in} = \{s_i, p^{n-1}\}$.

[10] The results extend to the case in which agents have different precisions of signals provided that there is no positive mass of agents perfectly informed (see Vives, 1993). If this were the case then convergence would obtain at the standard rate $1/\sqrt{n}$. This is so since perfectly informed agents do not learn from public information and therefore put a constant weight on their (perfect) signals. The outcome is that a constant amount of information is incorporated into the public statistic.

The planner has an incentive to depart from the myopic minimization of the period loss and "experiment" to increase the informativeness of public information. The experimentation comes in the form of imposing an increased response to private information, over and above the market response. This experimentation yields a long-run benefit but comes at the cost of increasing short-term losses.

Given decision rules $q_{in}(I_{in})$, the average expected loss in period n is then $\int_0^1 E(\theta - q_{in}(I_{in}))^2\, di$. Furthermore, given the structure of our problem there is no loss of generality in restricting attention to symmetric rules (see Footnote 11 below): $q_{in}(I_{in}) = q_n(I_{in})$ for all i. The loss in period n is then $L_n = E(\theta - q_n(I_{in}))^2$.

The objective of the team is to minimize $\sum_{n=0}^{\infty} \delta^n L_n$, choosing a sequence of functions $\{q_n(\,.\,)\}_{n=0}^{\infty}$ linear (respectively) in the information vectors $\{I_{in}\}_{n=0}^{\infty}$, $I_{in} = \{s_i, p^{n-1}\}$. The analysis can proceed as in the market case. Given the linear strategies imposed by the center, $q_{in} = a_n s_i + \varphi_n(p^{n-1})$, with φ_n a linear function of p^{n-1}, public information is $p_n = \int_0^1 q_{in}\, di + u_n = z_n + \varphi_n(p^{n-1})$, where $z_n = a_n\theta + u_n$ is the new information in p_n. As before we can then write $q_{in} = a_n s_i + c_n \theta_{n-1}$, with c_n a real number and θ_{n-1} the sufficient statistic for public information. Recall that the precision of public information θ_n equals $\tau_n = \tau_\theta + \tau_u A_n$, with $A_n = \sum_{t=0}^{n} a_t^2$, and therefore only the coefficients a_t matter in the accumulation of information, the coefficients c_t having no *intertemporal* effect. The consequence is that at an optimal solution the coefficient c_n is a function of a_n. At period n, c_n is chosen, contingent on a_n, to minimize L_n. This is accomplished by setting $c_n = 1 - a_n$. The reduced form of the team minimization problem is then

$$\text{Min} \sum_{n=0}^{\infty} \delta^n L_n, \qquad \text{with } L_n = \frac{(1-a_n)^2}{\tau_\theta + \tau_u A_{n-1}} + \frac{a_n^2}{\tau_\varepsilon},$$

choosing a sequence of real numbers $\{a_n\}_{n=0}^{\infty}$, with $a_{-1} = A_{-1} = 0$.[11]

It is immediate that the team solution calls for the weight of private information to be no less than the market weight (for a given accumulated public precision) and that it accumulates public precision in an unbounded way.

[11] Asymmetric rules cannot improve upon the symmetric case. First of all, using asymmetric weights a_{in} to the signals of different agents s_i, it is still optimal to set the weights to public information so that $c_{in} = 1 - a_{in}$ (the reason is as before that they do not have any intertemporal effect). This means that $L_{in} = E(\theta - q_{in}(I_{in}))^2 = ((1-a_{in})^2/(\tau_\theta + \tau_u A_{n-1})) + (a_{in}^2/\tau_\varepsilon)$. Let $L(a) = ((1-a)^2/\tau) + (a^2/\tau_\varepsilon)$. Given that $L(a)$ is convex in a it is immediate that $\int_0^1 L(a_i)\, di \geq L(\int_0^1 a_i\, di)$ and therefore there cannot be any static gain from asymmetric rules. Furthermore, there cannot be any dynamic gain either since if in period n different weights a_{in} were given to the signals of different agents, the signal-to-noise ratio in the new information z_n of the public statistic p_n would depend only on the average a_{in}, $\int_0^1 a_{in}\, di$.

PROPOSITION 4.1. *Let* $\{a_n\}_{n=0}^{\infty}$ *be a solution to the team problem, then* $a_n \geq \tau_\varepsilon/(\tau_\varepsilon + \tau_{n-1})$ and $\tau_n \underset{n}{\rightarrow} \infty$ as $n \rightarrow \infty$.

Proof. If for some n, $a_n < \tau_\varepsilon/(\tau_\varepsilon + \tau_{n-1})$ then by letting $a_n = \tau_\varepsilon/(\tau_\varepsilon + \tau_{n-1})$ (the market solution for a given τ_{n-1}) the solution can be improved upon since $a_n = \tau_\varepsilon/(\tau_\varepsilon + \tau_{n-1})$ minimizes $L_n = (1 - a_n)^2/\tau_{n-1} + a_n^2/\tau_\varepsilon$, and furthermore subsequent losses will be no larger (since τ_t for $t \geq n$ will be no smaller). Therefore, at an optimal solution $\{a_n\}_{n=0}^{\infty}$, $a_n \geq \tau_\varepsilon/(\tau_\varepsilon + \tau_{n-1})$. Now, let us show that $\tau_n \underset{n}{\rightarrow} \infty$. The sequence $\{\tau_n\}_{n=0}^{\infty}$ is increasing. If it is bounded above then necessarily $a_n \underset{n}{\rightarrow} 0$. Let τ be an upper bound on τ_n, then $a_n \geq \tau_\varepsilon/(\tau_\varepsilon + \tau) > 0$, a contradiction. ■

The team problem can be posed in a classical dynamic programming framework by taking the sequence $\{A_n\}_{n=0}^{\infty}$ as control (note that $A_n - A_{n-1} = a_n^2$). Given that $\tau_n = \tau_\theta + \tau_u A_n$, this is equivalent to taking the sequence of precisions of public information as controls. The reformulated problem is then

$$\operatorname{Min} \sum_{n=0}^{\infty} \delta^n H(A_n, A_{n-1}), \qquad \text{with } A_{-1} = 0 \text{ and } A_n \geq A_{n-1},$$

where

$$H(A_n, A_{n-1}) = \frac{\left(1 - \sqrt{A_n - A_{n-1}}\right)^2}{\tau_\theta + \tau_u A_{n-1}} + \frac{A_n - A_{n-1}}{\tau_\varepsilon}.$$

The team problem is equivalent also to the solution of the functional equation

$$V(A) = \operatorname{Min}_{B \geq A}\{H(B, A) + \delta V(B)\}.$$

The result follows since attention can be restricted (for an appropriate $K > 0$) to a feasibility correspondence $\Gamma(A) = \{B\colon A + K \geq B \geq A\}$ compact-valued and continuous, and H is bounded on the graph of Γ (see Theorems 4.2 and 4.3 in Stokey *et al.*, 1989, noting that their Assumptions 4.3 and 4.4 are fulfilled). Indeed, an upper bound on B given A can be constructed as follows. The candidate V will fulfill $V(A) \leq ((1 - \delta)(\tau_\theta + \tau_u A))^{-1}$ for any $A \geq 0$. This is so since by letting $B = A$ in the functional equation we have $V(A) \leq (\tau_\theta + \tau_u A)^{-1} + \delta V(A)$ and therefore $V(A) \leq ((1 - \delta)(\tau_\theta + \tau_u A))^{-1}$. At any solution therefore we will have $H(B, A) \leq ((1 - \delta)(\tau_\theta + \tau_u A))^{-1}$ and this yields the desired upper bound on B.

The solution $V(\cdot)$ to the functional equation is the *value function* associated to the team problem. It is easily seen that at an optimal solution, as τ_n tends to ∞ with n, a_n tends to 0. Indeed, the optimal solution requires the short-run loss to tend to zero as n grows and consequently public precision must be accumulated unboundedly and the weight given to private information must correspondingly tend to zero.

PROPOSITION 4.2. *Let* $\{a_n\}_{n=0}^{\infty}$ *be a solution to the team problem, then* $a_n \to 0$ as $n \to \infty$.

Proof. As A tends to infinity, $V(A)$ and a^2 $(= B - A)$ tend to 0. Indeed, we have seen that $V(A) \leq ((1 - \delta)(\tau_\theta + \tau_u A))^{-1}$ and therefore $V(A)$ tends to 0 as A tends to infinity. It follows then that $a^2 = B - A$ tends to 0 also from the functional equation since this is necessary for $H(B, A) = ((1 - \sqrt{B - A})^2/(\tau_\theta + \tau_u A)) + ((B - A)/\tau_\varepsilon)$ to tend to 0 as A tends to infinity. We then have from Proposition 4.1, given that $\tau_n = \tau_\theta + \tau_u A_n$ tends to infinity, that any solution a_n tends to zero as n tends to infinity. ■

In order to characterize the optimal solution further I consider the dynamic programming problem restricting attention to coefficients a_n lying in the unit interval. As we have seen this is not restrictive since at any solution $a_n \underset{n}{\to} 0$, and in any case we are interested in the dynamics for n large.[12] Denote by $g(\cdot)$ the *policy function* of our dynamic problem. The characterization of $V(\cdot)$ and $g(\cdot)$ is given in Propositions 4.3 and 4.4.

Given that $A_n - A_{n-1} = a_n^2$, we can write the feasibility correspondence as $\Gamma(A) = \{B: A + 1 \geq B \geq A)$. Let $V(\cdot)$ denote the *value function* associated to the team problem,

$$V(A) = \text{Min}_{B \in \Gamma(A)} \{H(B, A) + \delta V(B)\},$$

and denote the solution to the minimization problem, the *policy function*, by $g(\cdot)$. The characterization of $V(\cdot)$ and $g(\cdot)$ is given in Propositions 4.3 and 4.4 (with proofs in the Appendix).

PROPOSITION 4.3. *The value function* $V: R_+ \to R$ *is strictly convex and twice-continuously differentiable with* $V' < 0$. *As* A *tends to infinity* $V'(A)$ (and $V(A)$) *tend to* 0.

The key fact that V is strictly decreasing in the accumulated precision of public information should be clear since a higher accumulated precision today generates uniformly (strictly) lower period losses for all feasible sequences from then on.

[12] For all simulations performed—the results of which are summarized below—a_n lies in the unit interval for all n.

PROPOSITION 4.4. *There is a continuous policy function* $g: R_+ \to R_+$ *which yields the unique solution to the team problem* ($A_0 = g(0)$, $A_1 = g(A_0), \ldots$). *We have that* $g(A) > A$, g *is quasiconvex and continuously differentiable, with* $g' > 0$ *for* A *large enough, and* $g' < 1$ (*with* $g' \to 1$ *as* A *tends to infinity*).

From the shape of the policy function (Fig. 2) it is clear that public precision accumulates and tends to infinity (A_n tends to infinity with n). Further, the fact that the slope of g is strictly less than 1 implies that $g(A) - A$ is decreasing in A and therefore a_n is strictly decreasing in t (recall that $g(A_{n-1}) - A_{n-1} = A_n - A_{n-1} = a_n^2$). That is, the optimal response to private information monotonically decreases over time.

The characterizations of the value and policy function yield immediately the market underinvestment result.

PROPOSITION 4.5. *Let* $\delta > 0$, *then*

(i) *For any given accumulated precision of public information* τ_{n-1}, *the team optimal response to private information* a_n *is strictly larger than the market solution* $\alpha_n = \tau_\varepsilon/(\tau_\varepsilon + \tau_{n-1})$ (*in particular*, $a_0 > \alpha_0$).

(ii) *At any period* n *the team solution* τ_n^* *has accumulated more precision than the market* τ_n: $\tau_n^* > \tau_n$.

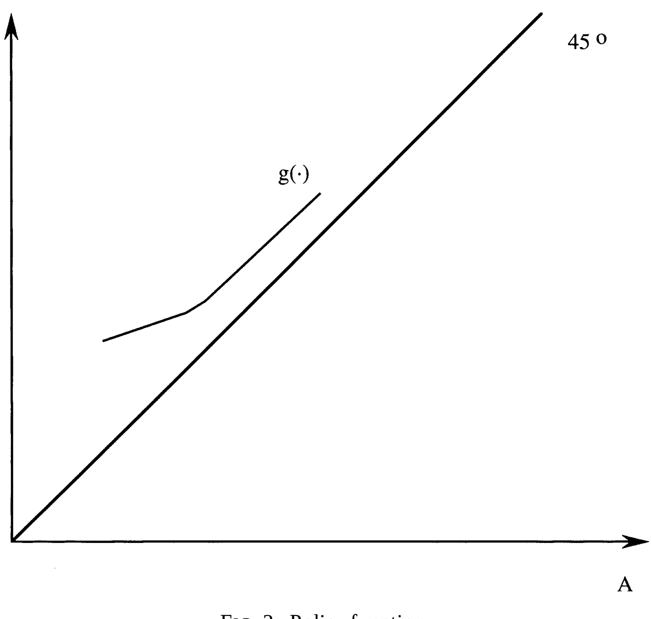

FIG. 2. Policy function.

Proof. (i) Given A, at the (interior) solution:

$$H_1(g(A), A) + \delta V'(g(A)) = 0.$$

From Proposition 4.3 we know that $V' < 0$ and therefore $H_1(g(A), A) > 0$. The market solution B satisfies $H_1(B, A) = 0$. Given that $H_{11} > 0$, $g(A) > B$. That is, for a given inherited precision $\tau = \tau_\theta + \tau_u A$, the optimal response $\tau_\theta + \tau_u g(A)$ is larger than the market response $\tau_\theta + \tau_u B$. We have then that $g(A) - A = a^2 > \alpha^2 = B - A$.

(ii) Denote by A_n^* (and τ_n^*) the team optimal solution and by A_n (and τ_n) the market solution. I show that $A_n^* > A_n$ by induction. Note first that $A_0^* > A_0$ from (i). Posit that $A_{n-1}^* > A_{n-1}$, then $A_n^* = g(A_{n-1}^*) > g(A_{n-1}) > A_n$. The last inequality follows from (i). The first holds if $g(\cdot)$ is increasing at A_{n-1}, which, given the quasiconvexity of $g(.)$, means that it is increasing for larger A's, since $A_{n-1}^* > A_{n-1}$. If $g(\cdot)$ were to be decreasing at A_{n-1}, then $A_n < g(A_{n-1}) < A_0^* \leq A_{n-1}^* < A_n^*$. I conclude that $\tau_n^* > \tau_n$. ■

For $\delta = 0$ the market solution is obtained. At period n, $a_n = \alpha_n$ is chosen to minimize L_n without regard to future losses. With $\delta > 0$ and given a certain accumulated precision the optimal program always calls for a larger response to private information because it internalizes the information externality, the benefit of a larger response today in lowering future losses. In other words, agents herd and rely too little on their private information at the market solution. This does not mean that the optimal program involves a uniformly larger response to private information overtime ($a_n > \alpha_n$ for all n) since the optimum accumulates more public precision. In fact, as we will see from the results of simulations, there is a critical n^* after which the optimal program calls for a lower response to private information than to the market in order to reap the benefits of an initial accumulation phase.

Simulations[13] produce the following results:

(1) For any n, the accumulated precision τ_n is increasing in the discount factor δ. (See Fig. 1.)

(2) For $\delta > 0$ there is a n^* such that $a_n > \alpha_n$ for $n < n^*$ and $a_n < \alpha_n$ for $n > n^*$. The length n^* represents the "investment" or experimentation phase of the team program and it is increasing in δ (larger weight to the

[13] Simulations have been performed for a horizon of N periods approximating the infinite horizon with the discounted loss implied by a constant public precision from period $N + 1$ on at the level of public precision in period N. That is, the total present ($n = 0$) discounted loss from period $N + 1$ on when the public precision in period N is τ equals $\delta^{N+1}/(1 - \delta)\tau$. The range of parameter values explored has been δ in [.75, .98], τ_u in [.1, 100], and τ_ε in [.05, 20], fixing $\tau_\theta = 1$ and with N up to 50.

future) and τ_ε (more precise private information) and decreasing in τ_u (noisier public signals).[14]

(3) For n large (small) L_n is decreasing (increasing) in the discount factor δ. The period loss for n large displays the opposite pattern to τ_n: L_n is smaller for larger δ. For n small the ranking is reversed since in the first periods investment in public information, with a_n larger than the market solution α_n for positive discounts, implies a larger loss.

(4) The relative welfare loss of the market solution with respect to the team solution (that is, the difference between the total discounted loss at the market and the team solutions divided by the total discounted loss at the market solution) can be quite high. This is $(v - v^o)/v$, where $v = \sum_{n=0}^{\infty} \delta^n L_n$, with L_n the period loss at the market solution and v^o the same expression at the team solution (the value function for $\tau = \tau_\theta$). For example, $(v - v^o)/v$ is around 25% for $\tau_u = 5$, $\tau_\varepsilon = .5$, $\delta = .95$, and $\tau_\theta = 1$. The relative welfare loss is increasing in the discount factor, and nonmonotonic in τ_ε and τ_u. Indeed, for extreme values of τ_ε or τ_u there is no information externality and the market and team solutions coincide. For τ_ε very large (infinite in the limit) there is perfect information and no loss; for $\tau_\varepsilon = 0$ there is no private information. For τ_u very large (infinite in the limit) public information is fully revealing and for $\tau_u = 0$ it is uninformative. (See Table I.)

We have characterized up to now the information externality arising in the market solution as compared with the team optimal solution. Furthermore, from Propositions 4.1 and 4.2, we know that as n tends to infinity at the team solution public precision accumulates unboundedly and the response to private information tends to zero. The question arises then as to whether the market solution distorts the rate at which public information is accumulated. Proposition 4.6 makes clear that this is not the case.

PROPOSITION 4.6. *The market and the team solution have exactly the same asymptotic properties.*

Proof. We show that at the optimal solution τ_n is of the order of $n^{1/3}$, a_n is of the order of $n^{-1/3}$, and $A\tau_\infty = (3\tau_u)^{1/3}(\tau_\varepsilon)^{2/3}$, exactly as in the market solution. The result follows as in the market case since I claim that at the optimal solution $a_n \tau_{n-1}$ tends to τ_ε as n tends to infinity. In order to prove the claim note that the optimal solution fulfills $H_1(g(A), A) +$

[14] Note that with a finite horizon the existence of n^* as claimed is clear. In the last period, say N, there is no future and therefore the optimal solution is just Bayesian updating given accumulated precision τ^*_{N-1}: $a_N = \tau_\varepsilon/(\tau_\varepsilon + \tau^*_{N-1})$. Since $\tau^*_{N-1} > \tau_{N-1}$ (the market precision) we have that $a_N < \alpha_N$. On the other hand, for n low the optimal solution will call for an increased response to private information.

TABLE I
Comparative Statics of the Relative Welfare Loss

(a) $\tau_\theta = 1$, $\tau_\varepsilon = 0.5$, $\delta = 0.95$		(b) $\tau_\theta = 1$, $\tau_\varepsilon = 2$, $\delta = 0.95$	
τ_u	$\frac{v - v^o}{v}$	τ_u	$\frac{v - v^o}{v}$
0.1	1.78	0.1	1.28
0.5	8.52	0.5	5.04
1	13.22	1	7.74
2	18.50	2	11.04
5	25.61	5	15.82
20	33.84	20	20.23
100	30.20	100	11.93

(c) $\tau_\theta = 1$, $\tau_u = 0.5$, $\delta = 0.95$		(d) $\tau_\theta = 1$, $\tau_u = 5$, $\delta = 0.95$	
τ_ϵ	$\frac{v - v^o}{v}$	τ_ϵ	$\frac{v - v^o}{v}$
0.05	0.93	0.05	23.91
0.1	3.67	0.1	29.22
0.5	8.52	0.5	25.61
1	7.13	1	20.76
2	5.04	2	15.82
5	2.66	5	10.43
20	0.76	20	4.76

(e) $\tau_\theta = 1$, $\tau_\varepsilon = 0.5$, $\tau_u = 0.5$		(f) $\tau_\theta = 1$, $\tau_\varepsilon = 2$, $\tau_u = 5$	
δ	$\frac{v - v^o}{v}$	δ	$\frac{v - v^o}{v}$
0.75	1.04	0.75	4.38
0.9	4.10	0.9	10.24
0.95	8.52	0.95	15.82
0.98	19.62	0.98	26.93

$\delta V'(g(A)) = 0$ or $H_1(A_n, A_{n-1}) + \delta V'(A_n) = 0$ for any n. As n tends to infinity so does A_n and consequently $V'(A_n)$ tends to 0 (see Proposition 4.3). Therefore $H_1(A_n, A_{n-1}) = (1/\tau_\varepsilon) - ((1 - a_n)/(a_n \tau_{n-1}))$ tends to zero as n tends to infinity. The claim follows since a_n tends to zero with n (Proposition 4.2). ∎

Slow learning at the market solution cannot be blamed on the information externality. Indeed, it is optimal to learn slowly even when the information externality is accounted for. For a large enough accumulated public precision the optimal program looks very similar to the market problem (in technical terms, V' tends to zero or V "flattens out" as public precision increases unboundedly). The reason is that to accumulate public precision faster would entail too large a departure from the minimization of current expected losses. Note in particular that to put asymptotically a constant positive weight on private information and consequently obtain an increase in precision which is linear in n is not optimal. The weight given to private information a_n must tend asymptotically to zero as public information precision increases without bound.[15] Proposition 3.2 on the convergence of θ_n to θ holds also for optimal learning. Given that the above intuition does not rely on linearity it seems plausible to conjecture that the optimality of the rate of learning result is not an artifact of the requirement that the team uses linear rules, as in the market case.

In summary, the team solution involves the same rate of convergence of public information to θ (Proposition 4.6) although the market uniformly accumulates less information (Proposition 4.5). In other words, the information externality present in the market solution leads to underinvestment in public precision but it does not distort the rate of learning.

EXAMPLE. Models of *learning by doing* usually assume that unit production costs decrease with the total accumulated production. There is empirical evidence of learning by doing on production processes which involve complex coordinated labor operations like aircraft assembly and, more recently, in computers. The applied literature emphasizes the importance of group effort and ‘‘integrated adaptation effort’’ in the explanation of the learning curve (see, for example, Baloff, 1966). Improved coordination seems to be at the root of improved productivity. The model presented can be interpreted as a team problem to coordinate workers to minimize the expected costs of production. The problem of the team is to coordinate workers to minimize the expected costs of production. The coordination problem takes a very simple (and extreme) form: Costs are lower the closer the actions of workers are to an unknown parameter θ. The total expected cost of output in production round n is proportional to $\int_0^1 E(\theta - q_{in})^2\, di$ where q_{in} is the action of worker i in period n. Worker interaction reveals the statistic p_n. The team manager can impose decision rules on the workers which are measurable in the information they have. Our theory predicts that independently of whether the team manager behaves myopically or as a long-run optimizer (taking into account the learning externality and solving the discounted welfare problem) the rate of learning as given by the precision of public information τ_n, and consequently the period loss (expected unit cost), will be of the order of $n^{-1/3}$.

[15] As in the market case, the new information in the public signal p_n, $z_n = a_n\theta + u_n$, is asymptotically pure noise and convergence of public information θ_n to θ is slowed down.

This means that expected cost will decline at the rate $n^{-1/3}$. The value $\lambda = 1/3$ for $C(n) = kn^{-\lambda}$ is typical for airframes.[16]

5. A VARIATION: SHORT-LIVED AGENTS

The model presented can be easily extended to encompass short-run (one period lived) agents who want to minimize the mean square error of predicting θ. It is assumed that their signals (given θ) are uncorrelated across generations and have correlation ρ among members of the same generation. If this correlation is perfect there is a representative agent each period and the model is purely *sequential* (like most models of herding) with agents taking actions in turn.

Agent i in period n has available a *private signal* $s_{in} = \theta + \varepsilon_{in}$, where $\varepsilon_{in} \sim N(0, \sigma_\varepsilon^2)$, $\mathrm{Cov}(\varepsilon_{in}, \varepsilon_{jn}) = \rho\sigma_\varepsilon^2$, $i \neq j$, $\rho \in [0,1]$. Similarly to before, the convention is made that $\varepsilon_n = \int_0^1 \varepsilon_{in}\, di$ is a normal random variable with zero mean, and variance and covariance with ε_{in} are both equal to $\rho\sigma_\varepsilon^2$.[17] When $\rho = 0$ there is no correlation between the error terms of the signals and, again by convention, $\varepsilon_n = 0$ (a.s.). Public information is as before and therefore agent i in period n has available the information vector $I_{in} = \{s_{in}, p^{n-1}\}$.[18] Note that when $\rho = 0$ the model is *formally*

[16] A typical model of learning by doing assumes that the unit cost of production with an accumulated production of n is of the form $C(n) = kn^{-\lambda}$ with λ between 0 and 1 and k a constant. The $n^{-1/3}$ rate is typical for airframes and corresponds to a 20% "progress ratio" (that is, the proportionate reduction of per-unit labor input when the cumulated output doubles; see Fellner, 1969). Progress ratios oscillate in empirical studies between 20 and 30%. (See Scherer and Ross, 1990, pp. 98–99.)

[17] This convention is in accord with the finite-dimensional version of the stochastic process of the error terms. Let $(\varepsilon_1, \ldots, \varepsilon_k)$ be Normal random variables with zero mean, Var $\varepsilon_i = \sigma_\varepsilon^2$, for all i, and $\mathrm{Cov}(\varepsilon_i, \varepsilon_j) = \rho\sigma_\varepsilon^2$, with $\rho \in [0,1]$ for $i \neq j$. Then $\mathrm{Var}(k^{-1}\sum_{i=1}^k \varepsilon_i) = \mathrm{Cov}(k^{-1}\sum_{i=1}{}^k \varepsilon_i, \varepsilon_j) = \sigma_\varepsilon^2(1 + (k-1)\rho)/k$, for any $j = 1, \ldots, k$. That is, the average of the error terms is normally distributed with mean zero, and variance and covariance with any individual error $\sigma_\varepsilon^2(1 + (k-1)\rho)/k$. As k goes to infinity $(1 + (k-1)\rho)/k$ tends to ρ.

[18] It is interesting to note also that the formal analysis of the model would be unchanged if agents had an *idiosyncratic* expected loss funtion $L_{in} = E(\theta + \eta_{in} - q_{in})^2$ with η_{in} being a random variable with finite variance σ_η^2 independently distributed with respect to the other random variables of the model. In this case $L_{in} = E(\theta + q_{in})^2 + \sigma_\eta^2$. With idiosyncratic loss functions private signals can be thought to be *endogenous* and represent *word-of-mouth* communication (see Banerjee and Fudenberg, 1995, for a Bayesian model of word-of-mouth communication). When agent i in period n obtains his payoff he learns $\theta + \eta_{in}$ and communicates it to a friend of the next generation of agents. That is, the signal received by i in period $t + 1$ is $\theta + \eta_{in}$. This corresponds to the model provided η_{in} has the same properties as the error terms ε_{in}. In this case the realized payoffs in period n generate information for other agents in period $n + 1$.

identical to the model with long-lived agents. Let now

$$\tau_n = \tau_\theta + \sum_{t=0}^{n} \left(\rho\tau_\varepsilon^{-1} + \left(\alpha_t^2 \tau_u\right)^{-1} \right)^{-1}.$$

Propositions 3.1 and 3.2 hold.[19] When $\tau_u = \infty$, that is, when there is no noise in public information (and signals are correlated, $\rho > 0$) the order of τ_n is n ($\tau_n = \tau_\theta + (n+1)\tau_\varepsilon \rho^{-1}$) and the order of L_n is n^{-1}. In this case the new information in p_n, $\alpha_n(\theta + \varepsilon_n)$, reveals the relevant information of agents. Indeed, even though $\alpha_n \to 0$, α_n is always positive and therefore the sequence of noisy signals $\{\theta \overset{n}{+} \varepsilon_n\}$ can be inferred from the sequence of new information $\{\alpha_n(\theta + \varepsilon_n)\}$. The analysis of optimal learning proceeds as before with minor variations.[20] Obviously, when $\tau_u = \infty$ there is no information externality since public information is a sufficient statistic for the information of agents (although with $\rho > 0$ public information is not perfectly revealing of θ) and the market solution is optimal.

An example is provided by *consumers learning about quality*. In each period there are many consumers of two types: "rational" and "noise" or "random." All consumers are endowed with a utility function which is linear with respect to money. Consumers only differ in their information and are one-period lived. Generation t, consumer i's utility consuming q_{in}, is given by $U_{in} = (\theta + \eta_{in})q_{in} - \frac{1}{2}q_{in}^2$. The willingness to pay of consumer i in period n (consumer *in*) is $\theta + \eta_{in}$. Consumers are uncertain about $\theta + \eta_{in}$ and only learn it after consuming the good. The parameter θ represents the average component of the willingness to pay and will depend on the matching between product and population characteristics. Consumer *in* receives a signal s_{in} about θ (word of mouth from the experience of a previous consumer or an independent test of the product). Given that consumer *in*'s idiosyncrasy η_{in} is uncorrelated with all other random variables in the environment and that the consumer learns $\theta + \eta_{in}$ only after consuming the good, we have that $E(\theta + \eta_{in} | I_{in}) = E(\theta | I_{in})$. Assume for simplicity that firms produce at zero cost and that prices are fixed at marginal cost. Expected utility maximization plus price taking behavior (at zero price) imply that $q_{in} = E(\theta | I_{in})$. If u_n denotes the purchases of the random consumers, then aggregate demand will be $p_n = \int_0^1 E(\theta | I_{in})\, di + u_n$. Consumers active in period n have access to the history of past sales $p^{n-1} = (p_0, p_1, \dots, p_{n-1})$. Consumer i's information

[19] The proof of Proposition 3.1(ii) needs now an extension of Lemma AI in Vives (1993) for the case $\rho > 0$.

[20] The presence of correlation in the signals tends to decrease the optimal weight to private information. For example, in a two-period optimal learning problem it is easily seen that a_0 decreases with ρ.

vector in period n is therefore $I_{in} = \{s_{in}, p^{n-1}\}$. The model is thus formally identical to the one presented (under the same distributional assumptions). Indeed, the expected welfare loss with respect to the full-information first best (where θ is known and $q_{in} = \theta$) is easily seen to be $E(\theta - q_{in})^2/2$. The results imply that consumers will learn slowly quality from quantities consumed or market shares and that they will be too conservative, with respect to the welfare benchmark, in responding to their private information. Further, slow learning by consumers enhances the possibilities of firms manipulating consumer beliefs (signal-jam the inferences consumers make from market shares; see Caminal and Vives, 1996).

6. CONCLUDING REMARKS

Learning from others via noisy public statistics in a ‘‘smooth’’ environment (continuous payoffs and actions spaces) is *successful*, agents end up discovering the truth, but *slow* even if learning is socially optimal. Market learning nevertheless involves an *information externality* which induces a too low accumulation of public precision, that is, underinvestment in public information. The results have been obtained in a purely statistical model where agents try to minimize the mean square error of prediction. Several economic problems which fit the framework have been presented.

A question may arise in regard to the robustness of results to less restrictive assumptions on the formulation of the model. For example, relaxing the assumptions of a continuum of agents, mean square error prediction (in favor of more general concave preferences), and normal distributions. The extension of the results to finite populations should present no problems. Furthermore, it seems plausible to conjecture that convergence to the correct action will obtain under general conditions and even that convergence will be ‘‘slow’’ in the presence of noise (although certainly it need not be the $n^{-1/3}$ rate which arises in the linear-normal framework). By ‘‘slow’’ I mean strictly less than the usual $n^{-1/2}$. The intuition is that with smooth payoffs and continuous strategy spaces actions can be fine-tuned to information and eventual learning of θ should obtain, but convergence has to be slow since otherwise the benefits of accumulated public information would never be enjoyed. Both in market and optimal learning schemes imperfectly informed agents have to put decreasing weight on their private information as public information improves, indeed, zero in the limit when public information reveals θ. The rate of convergence $n^{-1/2}$ requires that agents put a weight to private information bounded away from zero. Obviously, to characterize precisely the rate of convergence under general circumstances is a much more difficult task which is left for future research.

Two other topics deserve further research: the effects of endogenous and costly information acquisition on the learning process, and learning in a changing environment where the unknown parameter follows a stochastic process.

APPENDIX

Market Learning

Normality and the recursive structure of the model imply that $q_{\text{in}} = E(\theta | I_{\text{in}})$ is a linear function of s_i and $p^{n-1} = \{p_0, \dots, p_{n-1}\}$,

$$E(\theta | I_{\text{in}}) = \alpha_n s_i + \phi_n(p^{n-1}),$$

where α_n is a constant and ϕ_n is a linear function of p^{n-1}. Indeed, if the statement is true for n then (and according to our convention on the period average signal)

$$p_n = \int_0^1 E(\theta | I_{\text{in}})\, di + u_n = \alpha_n \theta + \phi_n(p^{n-1}) + u_n.$$

The public signal at period n, p_n, is a linear function of $\theta + u_n$ and past public signals p^{n-1}, and consequently $E(\theta | I_{\text{in}+1})$ is again a linear function of s_i and p_n. Further, it is clear that for $n = 0$ the induction process can be started since $E(\theta | s_i)$ is linear in s_i.

Defining $z_n = \alpha_n \theta + u_n$, we have that $p_n = z_n + \phi_n(p^{n-1})$. Note that the random variable z_n represents the *new information* in the public signal p_n. Due to the presence of noise z_n does not fully reveal the aggregate of the private information of agents in period n, θ. Nevertheless it is clear that the vector of public information p^n can be inferred from the vector z^n and vice versa. The following lemmata characterize sufficient statistics for public and individual information at stage n. Their proof is standard and is omitted (for related results and proofs see Vives, 1993).

LEMMA A.1. *The random variable $\theta_n = E(\theta | z^n)$ is sufficient in the estimation of θ based on z^n and follows a martingale, $E(\theta_n | \theta_{n-1}) = \theta_{n-1}$. Further, $\tau_n = \tau_\theta + \tau_u \sum_{t=0}^{n} \alpha_t^2$.*

LEMMA A.2. *The random vector (s_i, θ_{n-1}) is sufficient in the estimation of θ based on $I_{\text{in}} = \{s_i, z^{n-1}\}$, and individual i's action in period n is*:

$$\begin{aligned} q_n(I_{\text{in}}) &= E(\theta | s_i, \theta_{n-1}) \\ &= \alpha_n s_i + (1 - \alpha_n)\theta_{n-1}, \qquad \textit{with } \alpha_n = \tau_\varepsilon / (\tau_\varepsilon + \tau_{n-1}). \end{aligned}$$

In order to characterize the asymptotic behavior of α_n, τ_n, and L_n, let us define the order and the asymptotic value of a sequence of real numbers. I will say that the sequence B_n *is of the order* $(\approx)$ $n^v, v \in R$, and I will write $B_n \approx n^v$, whenever $n^{-v} B_n \underset{n}{\to} k$, for some nonzero constant k.

Proof of Proposition 3.1. (i) I show first that $\tau_n \underset{n}{\to} \infty$. If this is not true, since τ_n is increasing in n, we would have $\tau_n \underset{n}{\to} \tau < \infty$ and $\alpha_n \underset{n}{\to} \alpha > 0$ (recall that $\alpha_n = \tau_\varepsilon/(\tau_\varepsilon + \tau_{n-1})$). This is a contradiction since $\tau_n = \tau_\theta + \tau_u \sum_{t=0}^{n} \alpha_t^2 \approx \sum_{t=0}^{n} \alpha_t^2$ whenever τ_u is finite, and consequently τ_n would be of the order n and tend to ∞. It follows then that $\alpha_n \to 0$.

(ii) An heuristic argument to show that $\tau_n \approx n^{1/3}$ and $\alpha_n \approx n^{-1/3}$ runs as follows. Let $\tau_n \approx n^v$ for some $v > 0$. Then $\alpha_t \approx \tau^{-v}$ (because $\alpha_t \approx \tau_t^{-1}$ from $\alpha_t = \tau_\varepsilon/(\tau_\varepsilon + \tau_{t-1})$) and $\tau_n \approx \sum_{t=0}^{n} t^{-2v}$, which can be seen to be of the order of n^{1-2v}. Now, the equality $v = 1 - 2v$ implies that $v = 1/3$. For a formal proof it is sufficient to show that $n^{-1/3}\tau_n \underset{n}{\to} (3\tau_u)^{1/3}(\tau_\varepsilon)^{2/3}$. This follows from Lemma AI in Vives (1993) and the fact that $\alpha_n \tau_n \underset{n}{\to} \tau_\varepsilon$. ■

Proof of Proposition 4.3. The function H is continuously differentiable. It can be checked that for interior solutions $(a_n \in (0,1))$ the Hessian matrix of H is positive definite. It follows then that H is (strictly) convex. Recall that according to Proposition 4.1 at any solution $a_n > 0$. The feasibility correspondence Γ has a convex graph. It follows then (see Stokey *et al.*, 1989, Theorems 4.8 and 4.11) that $V(\cdot)$ is strictly convex and continuously differentiable. I claim that $V(\cdot)$ is strictly decreasing. It follows then from strict convexity and differentiability that $V' < 0$. The claim follows from the equivalence between the problem with $\{a_n\}_{n=0}^{\infty}$ and with $\{A_n\}_{n=0}^{\infty}$ as controls. Consider the problem with $\{a_n\}_{n=0}^{\infty}$ as controls. A higher A_{n-1} translates into a higher accumulated precision $\tau_{n-1} = \tau_\theta + \tau_u (\sum_{h=0}^{n-1} a_h^2)$ and a uniformly strictly lower period loss $L_n = (1 - a_n)^2/\tau_{n-1} + a_n^2/\tau_\varepsilon$ for all feasible sequences from period n on. Therefore $V(\cdot)$ has to be strictly decreasing. Twice-continuous differentiability of V follows from Santos (1991, Theorem 2.1 under Assumption B.2) since it can be checked that the ratio of the eigenvalues of the Hessian matrix of H is between zero and one along any optimal path. Now, according to the proof of Proposition 4.2, as A tends to infinity $V(A)$ tends to 0 and this implies given strict convexity of V and $V' < 0$ that $V'(A)$ tends also to 0 as A tends to infinity. ■

Proof of Proposition 4.4. Existence of a continuous (optimal) policy function follows as in Proposition 4.2 (from Stokey *et al.*, 1989, Chap. 4).

The fact that $g(A) > A$ follows from the feasibility constraint (public information accumulates). Continuous differentiability follows as in Proposition 4.3 above from Santos (1991). Let H_i denote the derivative of H with respect to the ith argument and H_{ij} the cross derivative. Consider the first order condition:

$$H_1(g(A), A) + \delta V'(g(A)) = 0.$$

It follows then (since $H_{11} + \delta V'' > 0$) that $g' = -H_{12}/(H_{11} + \delta V'')$, and therefore $\mathrm{Sign}\{g'\} = \mathrm{Sign}\{-H_{12}\}$. Now

$$-H_{12} = \Big((2(g(A) - A))^{-1} - \Big(1 - \sqrt{g(A) - A}\Big)\tau_u/(\tau_\theta + \tau_u A)\Big)\Big/\Big((\tau_\theta + \tau_u A)\sqrt{g(A) - A}\Big).$$

This is positive for A large since as A tends to infinity $g(A) - A$ tends to zero (Proposition 5.1). Further, it is easily seen that

$$-H_{12}/H_{11} = 1 - \Big(\Big(\sqrt{g(A) - A}\Big)^{-1} - 1\Big)\tau_u/(\tau_\theta + \tau_u A)^2.$$

The expression is strictly less than 1 provided $(\sqrt{g(A) - A})^{-1} - 1)$ is positive. This holds at any interior solution $(g(A) - A < 1)$. The fact that $-H_{12}/H_{11} < 1$ is sufficient, given that $V'' \leq 0$, to conclude that $g' < 1$. From the proof of Proposition 4.6 at an optimal solution $A \to \infty$, and $\sqrt{g(A) - A}\,(\tau_\theta + \tau_u A) \to \tau_\varepsilon$ as $A \to \infty$. This implies that $-H_{12}/H_{11} \to 1$ as $A \to \infty$. Since V flattens out as $A \to \infty$, with $V'' \to 0$, we can conclude that $g' \to 1$ as $A \to \infty$. Quasiconvexity follows from the stated properties of g plus the fact that if H_{12} (and g' therefore) becomes 0 at some A then $-H_{12} > 0$ (and consequently $g' > 0$) for larger A's. ■

REFERENCES

Aghion, P., Bolton, P., Harris, C., and Jullien, B. (1991). "Optimal Learning by Experimentation," *Rev. Econ. Stud.* **58**, 621–654.

Aumann, R. (1976). "Agreeing to Disagree," *Ann. Statist.* **4**, 1236–1239.

Baloff, N. (1966). "The Learning Curve. Some Controversial Issues," *J. Indust. Econ.* 275–283.

Banerjee, A. (1992). "A Simple Model of Herd Behavior," *Quart. J. Econ.* **107** (3), 797–818.

Banerjee, A., and Fudenberg, D. (1995). "Word-of-Mouth Learning," Discussion Paper 22. Harvard University, Cambridge, MA.

Bikhchandani, S., Hirshleifer, D., and Welch, I. (1992). "A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades," *J. Polit. Econ.* **100,** 992–1026.

Bolton, P., and Harris, C. (1993). "Strategic Experimentation," mimeo.

Caminal, R., and Vives, X. (1996). "Price Dynamics and Consumer Learning," mimeo. Institut d'Analisi Econòmica.

Caplin, A., and Leahy, J. (1994). "Business as Usual, Market Crashes, and Wisdom after the Fact," *Amer. Econ. Rev.* **84** (3), 548–565.

Chamley, C., and Gale, D. (1994). "Information Revelation and Strategic Delay in a Model of Investment," *Econometrica* **62** (5), 1065–1085.

Easley, D., and Kiefer, N. (1988). "Controlling Stochastic Process with Unknown Parameters," *Econometrica* **56** (5), 1045–1064.

Ellison, G., and Fudenberg, D. (1993). "Rules of Thumb for Social Learning," *J. Polit. Econ.* **101** (4), 612–643.

Feldman, M. (1987). "An Example of Convergence to Rational Expectations with Heterogeneous Beliefs," *Int. Econ. Rev.* **28** (3), 635–650.

Feldman, M., and Gilles, C. (1985). "An Expository Note on Individual Risk without Aggregate Uncertainty," *J. Econ. Theory* **35**, 26–32.

Fellner, W. (1969). "Specific Interpretations of Learning by Doing," *J. Econ. Theory* **1** (2), 119–140.

Geanakoplos, J., and Polemarchakis, H. (1982). "We Can't Disagree Forever," *J. Econ. Theory* **28**, 192–200.

Grossman, S. (1976). "On the Efficiency of Competitive Stock Markets where Trades Have Diverse Information," *J. Finance* **31** (2), 573–585.

Grossman, S., and Stiglitz, J. (1980). "On the Impossibility of Informationally Efficient Markets," *Amer. Econ. Rev.* **70**, 393–408.

Grossman, S., Kihlstrom, R., and Mirman, L. (1977). "A Bayesian Aproach to the Production of Information and Learning by Doing," *Rev. Econ. Stud.* 533–547.

Hellwig, M. F. (1980). "On the Aggregation of Information in Competitive Markets," *J. Econ. Theory* **22,** 477–498.

Jordan, J. S. (1992a). "Convergence to Rational Expectations in a Stationary Linear Game," *Rev. Econ. Stud.* **59** (1), 109–123.

Jordan, J. S. (1992b). "The Exponenetial Convergence of Bayesian Learning in Normal Form Games," *Games and Econ. Behav.* **4** (2), 202–217.

Kihlstrom, R., Mirman, L., and Postlewaite, A. (1984). "Experimental Consumption and the 'Rothschild Effect'," in *Bayesian Models of Economic Theory* (M. Boyer and R. Kihlstrom, Eds.). Amsterdam: Elsevier.

Kirman, A. (1993). "Ants, Rationality and Recruitment," *Quart. J. Econ.* **108** (1), 137–156.

Kyle, A. S. (1985). "Continous Auctions and Insider Trading," *Econometrica* **53**, 1315–1335.

Lee, In-Ho (1993). "On the Convergence of Informational Cascades," *J. Econ. Theory* **61** 395–411.

McKelvey, R. D., and Page, T. (1986). "Common Knowledge, Consensus, and Aggregate Information," *Econometrica* **54** (1), 109–127.

McLennan, A. (1984). "Price Dispersion and Incomplete Learning in the Long Run," *J. Econ. Dynamics and Control* **7**, 331–347.

Nielsen, L., Branderburger, A., Geanakoplos, J., McKelvey, R. D., and Page, T. (1990). "Common Knowledge of an Aggregate of Expectations," *Econometrica* **58** (5), 1235–1239.

Nyarko, Y. (1991). "On the Convergence of Bayesian Posterior Processes in Linear Economic Models: Counting Equations and Unknowns," *J. of Econ. Dynamics and Control* **15** (4), 687–713.

Radner, R. (1962). ''Team Decision Problems,'' *Ann. Math. Statist.* **33**, 857–881.

Rob, R. (1991). ''Learning and Capacity Expansion under Demand Uncertainty,'' *Rev. Econ. Stud.* **58**, 655–675.

Rothschild, M. (1974). ''A Two-Armed Bandit Theory of Market Pricing,'' *J. Econ. Theory* **9**, 185–202.

Santos, M. (1991). ''Smoothness of the Policy Function in Discrete Time Economic Models,'' *Econometrica* **59** (5), 1365–1382.

Scharfstein, D., and Stein, J. (1990). ''Herd Behavior and Investment,'' *Amer. Econ. Rev.* **80**, 465–479.

Scherer, F., and Ross, D. (1990). *Industrial Market Structure and Economic Performance*, pp. 98–99. Boston: Houghton–Mifflin Co.

Smith, L., and Sørensen, P. (1995). ''The Theory of Rational Social Learning,'' mimeo.

Smith, L., and Sørensen, P. (1996). ''Pathological Outcomes of Observational Learning,'' Discussion Paper 115. Oxford: Nuffield College.

Stein, J. (1987). ''Informational Externalities and Welfare-reducing Speculation,'' *J. Polit. Econ.* **95** (6), 1123–1145.

Stokey, N., Lucas, R., and Prescott, E. (1989). *Recursive Methods in Economic Dynamics*. Cambridge, MA: Harvard Univ. Press.

Townsend, R. M. (1978). ''Market Anticipations, Rational Expectations, and Bayesian Analysis,'' *Int. Econ. Rev.* **19** (2), 481–494.

Vives, X. (1993). ''How Fast Do Rational Agents Learn ?,'' *Rev. Econ. Stud.* **60**, 329–347.