
Chapter 3

Social learning

Why learn from others' actions? Because these actions reflect something about their information. Why don't we exchange information directly using words? People may not be able to express their information well. They may not speak the same language. They may even try to deceive us. What are we trying to find? A good restaurant, a good movie, a tip on the stock market, whether to delay an investment or not,... Other people know something about it, and their knowledge affects their behavior which, we can trust, must be self-serving. By looking at their behavior, we will infer something about what they know. This chain of arguments will be introduced here and developed in other chapters. We will see how the transmission of information may or may not be efficient and may lead to herd behavior, to sudden changes of widely believed opinions, etc...

For actions to speak and to speak well, they must have a sufficient vocabulary and be intelligible. In the first model of this chapter, individuals are able to fine tune their action in a sufficiently rich set and their decision process is perfectly known. In such a setting, actions reflect perfectly the information of each acting individual. This case is a benchmark in which social learning is equivalent to the direct observation of others' private information. Social learning is efficient in the sense that private actions convey perfectly private informations.

Actions can reveal perfectly private informations only if the individuals' decision processes are known. But surely private decisions depend on private informations and on personal parameters which are not observable. When private decisions depend on unobservable idiosyncracies, or equivalently when their observation by others is garbled by some noise, the process of social learning can be much slower than in the efficient case (Vives, 1993).

3.1 A canonical model of social learning

3.1.1 Structure

The purpose of a canonical model is to present a structure which is sufficiently simple and flexible to be a tool of analysis for a number of issues. Many models of rational social learning are built with the following three blocks:

1. *The information endowments:* The *state of nature* is what the information is about. It is denoted by θ and is randomly chosen by nature before the learning process in a set Θ that can be finite or in a continuum. The probability distribution of nature is the *prior distribution* and is known to all agents.
2. The *private information* of an agent i , $i = 1, \dots, N$, where N can be infinite, is what provides a value to others when they observe his action. That private information is modeled here by a random signal s_i . That signal has a probability distribution that is known by others in most cases (to make some inference possible), but by definition of *private*, the realization of the signal s_i cannot be observed by others. The signal provide some information on the state θ because its distribution depends on the true value of the state of nature θ . Any agent updates the prior on θ with the signal s_i to form a private distribution of probability of θ .
3. The action x_i of agent i is taken in round i , ($i \geq 1$) and belongs to a set Ξ . (Without loss of generality, Ξ is the same set for all agents. The action will be the “message”. We can assume here that this action is such that

$$x_i^* = E_i[\theta], \tag{3.1}$$

where E_i is the expectation of agent i when the action is taken.

One can explain the decision rule in (3.1) by the optimization of the agent. For example, it is the decision rule if the agent maximizes the expected value of the payoff function $-(x - \theta)^2$ or the function $\theta x - x^2/2$, which both have a simple intuitive interpretation. However, this “structural foundation” of the behavioral rule is not required here for the analysis of the social learning. Note that for these two functions, the optimal payoff is equal to minus the variance of θ (up to a constant). That may be convenient in evaluating the benefit of information.

What is essential at this stage, is that agents other than i know that (3.1) is the decision rule. We will deal later with the important case of an imperfect or imperfectly known decision rule. One can also have other payoff functions but they may lead to a more complex inference problem without additional insight.

Since agents “speak” through their actions, the definition of the action set Ξ is critical. A language with many words may convey more possibilities for communication than a language with few words. Individuals will learn more from each other about a parameter θ when the actions are in an interval of real numbers than when the actions are restricted to be either zero or one.

3.1.2 The process

In this chapter and the next, agents are ordered in an *exogenous sequence*. Agent t , $t \geq 1$, chooses his action in period t . We define the *history* of the economy in period t as the sequence

$$h_t = \{x_1, \dots, x_{t-1}\}, \quad \text{with } h_0 = \emptyset.$$

Agent t knows the history of past actions h_t before making a decision.

To summarize, at the beginning of period t (before agent t makes a decision), the *knowledge which is common to all agents* is defined by

- the distribution of θ at the beginning of time,
- the distributions of private signals and the payoff functions of all agents,
- the history h_t of previous actions.

We will assume that agents cannot observe the payoff of the actions of others. Whether this assumption is justified or not depends on the context. It is relevant for investment over the business cycle: given the lags between investment expenditures and their returns, one can assume that investment decisions carry the sole information. Later in the book, we will analyze other mechanisms of social learning. For the sake of clarity, it is best to focus on each one of them separately.

Agent t combines the public belief on θ with his private information (the signal s_t) to form his belief which has a *c.d.f.* $F(\theta|h_t, s_t)$. He then chooses the action x_t to maximize his payoff $E[u(\theta, x_t)]$, conditional on his belief.

All remaining agents know the payoff function of agent t (but not the realization of the payoff), and the decision model of agent t . They use the observation of x_t as a signal on the information of agent t , *i.e.*, his private signal s_t . The action of an agent is a message on his information. The social learning depends critically on how this message conveys information on the private belief. The other agents update the public belief on θ once the observation x_t is added to the history h_t : $h_{t+1} = (h_t, x_t)$. The distribution $F(\theta|h_t)$ is updated to $F(\theta|h_{t+1})$.

3.2 The Gaussian model

Social learning is efficient when an individual's action reveals completely his private information. This occurs when the action set which defines the vocabulary of social learning is sufficiently large. We begin with the Gaussian model (Section ??) that provides a simple and precise case for discussion.

The prior distribution on θ is normal, $\mathcal{N}(m_1, 1/\rho_1)$, with mean m_1 and precision ρ_1 . Since we focus on the social learning of a given state of nature, the value of θ does not change once it is set.

There is a countable number of individuals, indexed by $i \geq 1$, and each individual i has one private signal s_i such that

$$s_i = \theta + \epsilon_i, \quad \text{with } \epsilon_i \sim \mathcal{N}(0, 1/\rho_\epsilon).$$

Individual t chooses his action $x_t \in \mathcal{R}$ once and for all in period t : the order of the individual actions is set exogenously.

The public information at the beginning of period t is made of the initial distribution $\mathcal{N}(\bar{\theta}, 1/\rho_\theta)$ and of the history of previous actions $h_t = (x_1, \dots, x_{t-1})$.

Suppose that the public belief on θ in period t is given by the normal distribution $\mathcal{N}(\mu_t, 1/\rho_t)$. This assumption is obviously true for $t = 1$. By induction, we now show that it is true in every period.

(i) The belief of agent t

The belief is obtained from the Bayesian updating of the public belief $\mathcal{N}(\mu_t, 1/\rho_t)$ with the private information $s_t = \theta + \epsilon_t$. Using the standard Bayesian formulae with Gaussian distributions, the belief of agent t is $\mathcal{N}(\tilde{\mu}_t, 1/\tilde{\rho}_t)$ with

$$\begin{cases} \tilde{\mu}_t = (1 - \alpha_t)\mu_t + \alpha_t s_t, & \text{with } \alpha_t = \frac{\rho_\epsilon}{\rho_\epsilon + \rho_t}, \\ \tilde{\rho}_t = \rho_t + \rho_\epsilon. \end{cases} \quad (3.3)$$

(ii) The private decision

From the specification of $\tilde{\mu}_t$ in (3.3),

$$x_t = (1 - \alpha_t)\mu_t + \alpha_t s_t. \quad (3.4)$$

(iii) Social learning

Social learning is efficient when actions reveal perfectly private informations.

The decision rule of agent t and the variables α_t, μ_t are known to all agents. From equation (3.4), the observation of the action x_t reveals perfectly the private signal s_t . This is a key property. The public information at the end of period t is identical to the information of agent t : $\mu_{t+1} = \tilde{\mu}_t$, and $\rho_{t+1} = \tilde{\rho}_t$. Hence,

$$\begin{cases} \mu_{t+1} = (1 - \alpha_t)\mu_t + \alpha_t s_t, & \text{with } \alpha_t = \frac{\rho_\epsilon}{\rho_\epsilon + \rho_t}, \\ \rho_{t+1} = \rho_t + \rho_\epsilon. \end{cases} \quad (3.5)$$

In period $t + 1$, the belief is still normally distributed $\mathcal{N}(\mu_{t+1}, 1/\rho_{t+1})$ and the process can be iterated as long as there is an agent remaining in the game. The history of actions $h_t = (x_1, \dots, x_{t-1})$ is informationally equivalent to the sequence of signals (s_1, \dots, s_{t-1}) .

Convergence

The precision of the public belief increases linearly with time:

$$\rho_t = \rho_\theta + (t - 1)\rho_\epsilon, \quad (3.6)$$

and the variance of the estimate on θ is $\sigma_t^2 = 1/(\rho_\theta + t\rho_\epsilon)$, which converges to zero like $1/t$. This is the rate of the efficient convergence.

The weight of history and imitation

Imitation increases with the weight of history, but does not slow down social learning if actions reveal private informations.

Agent t chooses an action which is a weighted average of the public information μ_t from history and his private signal s_t (equation (3.4)). The expression of the weight of history, $1 - \alpha_t$, increases and tends to 1 when t increases to infinity. The weight of the private signal tends to zero. Hence, agents tend to “imitate” each other more as time goes on. This is a very simple, natural and general property: a longer history carries more information. Although the differences between individuals’ actions become vanishingly small as time goes on, the social learning is not affected because these actions are perfectly observable: no matter how small these variations, observers have a magnifying glass which enables them to see the differences perfectly. In the next section, this assumption will be removed. An observer will not “see” well the small variations. This imperfection will slow down significantly the social learning.

3.3 Observation noise

In the previous section, an agent’s action conveyed perfectly his private information. An individual’s action can reflect the slightest nuances of his information because: (i) it is

chosen in a sufficiently rich menu; (ii) it is perfectly observable; (iii) the decision model of each agent is perfectly known to others.

The extraction of information from an individual's action relies critically on the assumption that the decision model is perfectly known, an assumption which is obviously very strong. In general, individuals' actions depend on a common parameter but also on private characteristics. It is the essence of these private characteristics that they cannot be observed perfectly (exactly as the private information is not observed by others). To simplify, assume that the observation of the action of agent i is given by

$$x_i = E_i[\theta] + \eta_i, \quad \text{with} \quad \eta_i \sim \mathcal{N}(0, 1/\rho_\eta). \quad (3.7)$$

The noise η_i is independent of other random variables and it can arise either because there is an observation noise or because the payoff function of the agent is subject to an idiosyncratic variable.¹

Since the private parameter η_i is not observable, the action of agent i conveys a *noisy signal* on his information $E_i[\theta]$. Imperfect information on an agent's private characteristics is operationally equivalent to a noise on the observation of the actions of an agent whose characteristics are perfectly known.

The model of the previous section is now extended to incorporate an observation noise, along the idea of Vives (1993)². We begin with a direct extension of the model where there is one action per agent in each period. The model with many agents is relevant in the case of a market and will be presented in Section 3.2.

An intuitive description of the critical mechanism

Period t brings to the public information the observation

$$x_t = (1 - \alpha_t)\mu_t + \alpha_t s_t + \eta_t, \quad \text{with} \quad \alpha_t = \frac{\rho_\epsilon}{\rho_t + \rho_\epsilon}. \quad (3.8)$$

The observation of x_t does not reveal perfectly the private signal s_t because of a noise $\eta_t \sim \mathcal{N}(0, \sigma_\eta^2)$. This simple equation is sufficient to outline the critical argument. As time goes on, the learning process increases the precision of the public belief on θ , ρ_t , which tends to infinity. Rational agents imitate more and reduce the weight α_t which they put on their private signal as they get more information through history. Hence, they reduce the multiplier of s_t on their action. As $t \rightarrow \infty$, the impact of the private signal s_t on x_t becomes vanishingly small. The variance of the noise η_t remains constant over

¹For example if the payoff is $-(x_i - \theta - \eta_i)^2$.

²Vives assumes directly an observation noise and a continuum of agents. His work is discussed below.

Imitation increases with the weight of history and reduces the signal to noise ratio of private actions.

time, however. Asymptotically, *the impact of the private information on the level of action becomes vanishingly small relative to that of the unobservable idiosyncrasy*. This effect reduces the information content of each observation and slows down the process of social learning.

The impact of the noise cannot prevent the convergence of the precision ρ_t to infinity. By contradiction, suppose that ρ_t is bounded. Then α_t does not converge to zero and the precision ρ_t increases linearly, asymptotically (contradicting the boundedness of the precision). The analysis now confirms the intuition and measures accurately the impact of the noise on the rate of convergence of learning.

The evolution of beliefs

Since the private signal is $s_t = \theta + \epsilon_t$ with $\epsilon_t \sim \mathcal{N}(0, \sigma_\epsilon^2)$, equation (3.8) can be rewritten

$$x_t = (1 - \alpha_t)\mu_t + \alpha_t\theta + \underbrace{\alpha_t\epsilon_t + \eta_t}_{\text{noise term}} \quad (3.9)$$

A standard normalization will be used for most Gaussian signals.

The observation of the action x_t provides a signal on θ , $\alpha_t\theta$, with a noise $\alpha_t\epsilon_t + \eta_t$. We will encounter in this book many similar expressions of noisy signals on θ . We use a simple procedure to simplify the learning rule (3.9): the signal is normalized by a linear transformation such that the right-hand side is the sum of θ (the parameter to be estimated), and a noise:

$$\frac{x_t - (1 - \alpha_t)\mu_t}{\alpha_t} = z_t = \theta + \epsilon_t + \frac{\eta_t}{\alpha_t}. \quad (3.10)$$

The variable x_t is *informationally equivalent* to the variable z_t . We will use similar equivalences for most Gaussian signals. The learning rules for the public belief follow immediately from the standard formulae with Gaussian signals (3.3). Using (3.8), the distribution of θ at the end of period t is $\mathcal{N}(\mu_{t+1}, 1/\rho_{t+1}^2)$ with

$$\begin{cases} \mu_{t+1} = (1 - \beta_t)\mu_t + \beta_t \left(\frac{x_t - (1 - \alpha_t)\mu_t}{\alpha_t} \right), & \text{with} \\ \beta_t = \frac{\sigma_t^2}{\sigma_t^2 + \sigma_\epsilon^2 + \sigma_\eta^2/\alpha_t^2}, \\ \rho_{t+1} = \rho_t + \frac{1}{\sigma_\epsilon^2 + \sigma_\eta^2/\alpha_t^2} = \rho_t + \frac{1}{\sigma_\epsilon^2 + \sigma_\eta^2(1 + \rho_t\sigma_\epsilon^2)^2}. \end{cases} \quad (3.11)$$

Convergence

When there is no observation noise, the precision of the public belief ρ_t increases by a *constant* value ρ_ϵ in each period, and it is a linear function of the number of observations (equation (3.6)). When there is an observation noise, equation (3.11) shows that as $\rho_t \rightarrow \infty$,

the increments of the precision, $\rho_{t+1} - \rho_t$, becomes smaller and smaller and tend to zero. The precision converges to infinity at a rate slower than a linear rate. The convergence of the variance σ_t^2 to 0 takes place at a rate slower than $1/t$.

The slowing down of the convergence when actions are observed through a noise has been formally analyzed by Vives (1993). In a remarkable result (Proposition 3.1 in the Appendix), he showed that the precision of the public information, ρ_t increases only like the cubic root of the number of observations, $At^{1/3}$. The value of the constant A depends on the observation noise, but the rate $1/3$ is independent of that variance. Recall that with no noise, the precision increases linearly with t .

When the number of observations is large, 1000 additional observations with noise generate the same increase of precision as 10 observations when there is no observation noise.

Proposition 3.1 shows that the standard model of social learning where agents observe perfectly others' actions and know their decision process is not robust. When observations are subject to a noise, the process of social learning is slowed down, possibly drastically, because of the weight of history. That weight reduces the signal to noise ratio of individual actions. The mechanism by which the weight of history reduces social learning will be shown to be robust and will be one of the important themes in the book.

3.3.1 Large number of agents

The previous model is modified to allow for a continuum of agents. Each agent is indexed by $i \in [0, 1]$ (with a uniform distribution) and receives one private signal *once* at the beginning of the first period³, $s_i = \theta + \epsilon_i$, with $\epsilon_i \sim \mathcal{N}(0, \sigma_\epsilon^2)$. Each agent takes an action $x_t(i)$ in each period⁴ t to maximize the expected quadratic payoff in (??). At the end of period t , agents observe the aggregate action Y_t which is the sum of the individuals' actions and of an aggregate noise η_t :

$$Y_t = X_t + \eta_t, \quad \text{with} \quad X_t = \int x_t(i) di, \quad \text{and} \quad \eta_t \sim \mathcal{N}(0, 1/\rho_\eta).$$

At the beginning of any period t , the public belief on θ is $\mathcal{N}(\mu_t, 1/\rho_t)$, and an agent with signal s_i chooses the action

$$x_t(i) = E[\theta | s_i, h_t] = \mu_t(i) = (1 - \alpha_t)\mu_t + \alpha_t s_i, \quad \text{with} \quad \alpha_t = \frac{\rho_\epsilon}{\rho_t + \rho_\epsilon}.$$

³If agents were to receive more than one signal, the precision of their private information would increase over time.

⁴One could also assume that there is a new set of agents in each period and that these agents act only once.

By the law of large numbers⁵, $\int \epsilon_i di = 0$. Therefore, $\alpha_t \int s_i di = \alpha_t \theta$. The level of endogenous aggregate activity is

$$X_t = (1 - \alpha_t)\mu_t + \alpha_t \theta,$$

and the observed aggregate action is

$$Y_t = (1 - \alpha_t)\mu_t + \alpha_t \theta + \eta_t. \quad (3.12)$$

Using the normalization introduced in Section ??, this signal is informationally equivalent to

$$\frac{Y_t - (1 - \alpha_t)\mu_t}{\alpha_t} = \theta + \frac{\eta_t}{\alpha_t} = \theta + \left(1 + \frac{\rho_t}{\rho_\epsilon}\right)\eta_t. \quad (3.13)$$

This equation is similar to (3.10) in the model with one agent per period. (The variances of the noise terms in the two equations are asymptotically equivalent). Proposition 3.1 applies. The asymptotic evolutions of the public beliefs are the same in the two models.

Note that the observation noise has to be an aggregate noise. If the noises affected actions at the individual level, for example through individuals' characteristics, they would be "averaged out" by aggregation, and the law of large numbers would reveal perfectly the state of nature. An aggregate noise is a very plausible assumption in the gathering of aggregate data.

3.3.2 Application: a market equilibrium

This setting is the original model of Vives (1993). A good is supplied by a continuum of identical firms indexed by i which has a uniform density on $[0, 1]$. Firm i supplies x_i and the total supply is $X = \int x_i di$. The demand for the good is linear:

$$p = a + \eta - bX. \quad (3.14)$$

Each firm (agent) i is a price taker and has a profit function

$$u_i = (p - \theta)x_i - \frac{c}{2}x_i^2,$$

where the last term is a cost of production and θ is an unknown parameter. Vives views this parameter as a pollution cost which is assessed and charged after the end of the game.

As in the canonical model, nature's distribution on θ is $\mathcal{N}(\mu, 1/\rho_\theta)$ and each agent i has a private signal $s_i = \theta + \epsilon_i$ with $\epsilon_i \sim \mathcal{N}(0, 1/\rho_\epsilon)$. The expected value of θ for firm i is

$$E_i[\theta] = (1 - \alpha)\mu + \alpha(\theta + \epsilon_i), \quad \text{with} \quad \alpha = \frac{\rho_\epsilon}{\rho_\theta + \rho_\epsilon}. \quad (3.15)$$

⁵A continuum of agents of mass one with independent signals is the limit case of n agents each of mass $1/n$ where $n \rightarrow \infty$. The variance of each individual action is proportional to $1/n^2$ and the variance of the aggregate decision is proportional to $1/n$ which is asymptotically equal to zero.

The optimal decision of each firm is such that the marginal profit is equal to the marginal cost:

$$p - E_i[\theta] = cx_i.$$

Integrating this equation over all firms and using the market equilibrium condition (3.14) gives

$$p - \int E_i[\theta] di = cX = \frac{c}{b}(a + \eta - p),$$

which, using (3.15), is equivalent to

$$(b + c)p - ac - (1 - \alpha)\mu = \alpha\theta + c\eta.$$

Dividing both sides of this equation to normalize the signal, the observation of the market price is equivalent to the observation of the signal

$$Z = \theta + c\frac{\eta}{\alpha}, \quad \text{where} \quad \alpha = \frac{\rho_\epsilon}{\rho_\theta + \rho_\epsilon}.$$

The model is isomorphic to the canonical model of the previous section.

3.4 Extensions

Endogenous private information

See exercise 3.2.

Policy against mimetism

A selfish agent who maximizes his own welfare ignores that his action generates informational benefits to others. If the action is observed without noise, it conveys all the private information without any loss. But if there is an observation noise, the information conveyed by the action is reduced when the response of the action is smaller. When time goes on, the amplitude of the noise is constant and the agent rationally reduces the multiplier of his signal on his action. Hence, the action of the agent conveys less information about his signal when t increases. A social planner may require that agents overstate the impact of their private signal on their action in order to be “heard” over the observation noise. Vives (1997) assumes that the social welfare function is the sum of the discounted payoffs of the agents

$$W = \sum_{t \geq 0} \beta^t \left(-E_t[(x_t - \theta)^2] \right),$$

where x_t is the action of agent t . All agents observe the action plus a noise, $y_t = x_t + \epsilon_t$. The function W is interpreted as a loss function as long as θ is not revealed by a random exogenous process. In any period t , conditional on no previous revelation, θ is revealed

perfectly with probability $1 - \pi \geq 0$. Assuming a discount factor $\delta < 1$, the value of β is $\beta = \pi\delta$. If the value of θ is revealed, there is no more loss.

As we have seen in (3.3) and (3.4), a selfish agent with signal s_t has a decision rule of the form

$$x_t - \mu_t = (1 + \gamma) \frac{\rho_\epsilon}{\rho_t + \rho_\epsilon} (s_t - \mu_t), \quad (3.16)$$

with $\gamma = 0$. Vives assumes that a social planner can enforce an arbitrary value for γ . When $\gamma > 0$, the action to noise ratio is higher and the observers of the action receive more information.

Assume that a selfish agent is constrained to the decision rule (3.16) and optimizes over γ : he chooses $\gamma = 0$. By the envelope theorem, a small first order deviation of the agent from his optimal value $\gamma = 0$ has a second order effect on his welfare. We now show that it has a first order effect on the welfare of any other individual who make a decision. The action of the agent is informationally equivalent to the message

$$y = (1 + \gamma)\alpha s + \epsilon, \quad \text{with} \quad \alpha = \frac{\rho_\epsilon}{\rho_t + \rho_\epsilon}.$$

The precision of that message is $\rho_y = (1 + \gamma)^2 \alpha^2 \rho_\epsilon$.

Another individual's welfare is minus the variance after the observation of y . The observation of y adds an amount ρ_y to the precision of his belief. If γ increases from an initial value of 0, the variation of ρ_y is of the order of $2\gamma\alpha^2\rho_\epsilon$, *i.e.*, of the first order with respect to γ . Since the variance is the inverse of the precision, the impact on the variance of others is also of the first order and dwarfs the second order impact on the agent. There is a positive value of γ which induces a higher social welfare level.

EXERCISES

EXERCISE 3.1. (history cannot be summarized by a number)

Assume that (i) the distribution of the state of nature θ has a support in the set of real numbers (which does not have to be bounded); (ii) there is an infinite sequence of agents each with a private signal that is binary and symmetric such that $P(s = 1) = q$ with $q = \phi(\theta)$ for some monotone function ϕ which maps the set of real numbers to the open interval $(1/2, 1)$. You may take the example $\phi(\theta) = \frac{1}{4} \left(3 + \frac{\theta}{1 + |\theta|} \right)$; (iii) each agent t knows the history of the actions of the previous $t - 1$ agents and chooses the real number x_t to maximize his payoff function $-E[(\theta - x_t)^2]$.

1. Show, using words and no algebra that the action of an agent reveals perfectly his private signal.
2. Can the history h_t be summarized by $\sum_{i \leq t-1} x_i$?

EXERCISE 3.2. (Endogenous private information)

In the standard Gaussian model of social learning, each agent has to pay a fixed cost c to get a signal with precision ρ which is

$$s = \theta + \epsilon, \quad \text{with } \epsilon \sim \mathcal{N}(0, 1/\rho).$$

The cost c is assumed to be small. Agent t makes a decision in period t (both on the signal and on the action), and his action is assumed to be perfectly observable by others. The payoff function of each agent is quadratic: $U(x) = E[-(x - \theta)^2]$.

1. Show using words and no algebra, that there is a date T after which no agent buys a private signal. What happens to information and actions after that date T ?
2. Provide now a formal proof of the previous statement. For this compute the welfare gain that an agent gets by buying a signal.
3. Assume now that the cost of a signal with precision ρ is an increasing function,⁶ $c(\rho)$. Prove the following result:

⁶Suppose for example that the signal is generated by a sample of n independent observations and that each observation has a constant cost c_0 . Since the precision of the sample is a linear function of n , the cost of the signal is a step function. For the sake of exposition, we assume that ρ can be any real number.

- Suppose that $c'(\rho)$ is continuous and $c(0) = 0$. If the marginal cost of precision $c'(\rho)$ is bounded away from 0, (for any $\rho \geq 0$, $c'(\rho) \geq \gamma > 0$), no agent purchases a signal after some finite period T and social learning stops in that period.
4. Assume now that $c(q) = q^\beta$ with $\beta > 0$. Analyze the rate of convergence of social learning.
-

REFERENCES

- Burguet, R. and X. Vives, (2000). "Social Learning and Costly Information Acquisition," *Economic Theory*, **15**, 185-205, (first version 1995).
- Jun, B. and X. Vives (1996). "Learning and Convergence to a Full-Information Equilibrium are Not Equivalent," *Review of Economic Studies*, **63**, 653-674.
- Lee, In Ho (1992). "On the Convergence of Informational Cascades," *Journal of Economic Theory*, **61**, 395-411.
- Vives, X. (1993). "How Fast Do Rational Agents Learn?," *Review of Economic Studies*, **60**, 329-347.
- Vives, X. (1996). "Social Learning and Rational Expectations," *European Economic Review*, **40**, 589-601.
- Vives, X. (1997). "Learning From Others: a Welfare Analysis," *Games and Economic Behavior*, **20**, 177-200.

3.5 APPENDIX

PROPOSITION 3.1. (Vives, 1993) *In the Gaussian-quadratic model with an observation noise of variance σ_η^2 and private signals of variance σ_ϵ^2 , the variance of the public belief on θ , σ_t^2 , converges to zero as $t \rightarrow \infty$ and*

$$\frac{\sigma_t^2}{\left(\frac{\sigma_\eta^2 \sigma_\epsilon^4}{3t}\right)^{\frac{1}{3}}} \rightarrow 1. \quad (3.17)$$

Proof

Since we analyze a rate of convergence, it is more convenient to consider a variable which converges to zero than a variable which converges to infinity. (We will use Taylor expansions). Let $z_t = \sigma_t^2 = 1/\rho_t$. The third equation in (3.11) is of the form

$$z_{t+1} = G(z_t). \quad (3.18)$$

A standard exercise shows that $G(0) = 0$, and for $z > 0$, $0 < G(z) < z$ and $G'(z) > 0$. This implies that as $t \rightarrow \infty$, then $z_t \rightarrow 0$ which is a fixed point of F . The rest of the proof is an exercise on the approximation of (3.18) with the particular form (3.11) near the fixed point 0. Equation (3.11) can be rewritten:

$$z_{t+1} = \frac{z_t \left((\sigma_\epsilon^2 + \sigma_\eta^2) z_t^2 + 2\sigma_\eta^2 \sigma_\epsilon^2 z_t + \sigma_\eta^2 \sigma_\epsilon^4 \right)}{\sigma_t^6 + (\sigma_\epsilon^2 + \sigma_\eta^2) z_t^2 + 2\sigma_\eta^2 \sigma_\epsilon^2 z_t + \sigma_\eta^2 \sigma_\epsilon^4},$$

or

$$z_{t+1} = z_t - \frac{z_t^4}{z_t^3 + (\sigma_\epsilon^2 + \sigma_\eta^2) z_t^2 + 2\sigma_\eta^2 \sigma_\epsilon^2 z_t + \sigma_\eta^2 \sigma_\epsilon^4}.$$

$$\text{Since } z_t \rightarrow 0, \quad z_{t+1} = z_t - \frac{z_t^4}{A} (1 + O(z_t)), \quad \text{with } A = \sigma_\eta^2 \sigma_\epsilon^4,$$

where $O(z_t)$ is a term of order smaller than or equal to 1: there is $B > 0$ such that if $z_t \rightarrow 0$, then $O(z_t) < Bz_t$. Let b_t be such that $z_t = b_t/(t^{1/3})$. By substitution in the previous equation,

$$b_{t+1} \left(\frac{1+t}{t} \right)^{-\frac{1}{3}} = b_t - \frac{b_t^4}{At} \left(1 + O\left(\frac{b_t}{t^{\frac{1}{3}}} \right) \right),$$

or

$$b_{t+1} \left(1 - \frac{1}{3t} + O\left(\frac{1}{t^2} \right) \right) = b_t - \frac{b_t^4}{At} \left(1 + O\left(\frac{b_t}{t^{\frac{1}{3}}} \right) \right). \quad (3.19)$$

This equation is used to prove that b_t converges to a non zero limit. The proof is in two steps: (i) the sequence is bounded; (ii) any subsequence converges to the same limit.

(i) The boundedness of b_t :

First, from the previous equation, there exists T_1 such that if $t > T_1$, then

$$b_{t+1} < b_t \left(1 + \frac{1}{2t} \right). \quad (3.20)$$