6. Learning from others and herding

In the preceding chapters we have considered static models of market interaction in which agents would learn from prices and their private signals. This is not the only possibility for agents to learn about the parameters they are uncertain about. To start with, learning is a dynamic process. Furthermore, agents can learn also from other agents and, in particular, from the actions of other agents. Actions speak louder than words, the dictum says. Examples abound: consumers purchase the most popular brands, tourists patronize well-attended restaurants, and readers buy best-sellers. Learning from others can use informal routes like word-of-mouth. There has been a surge of interest in social learning, the process by which certain non-price mechanisms in society aggregate the information of individuals. In this chapter we introduce dynamic models of social learning as a stepping stone to address the dynamics of information in full-fledged market environments, which will be the object of study in Chapters 7, 8, and 9. The present chapter builds a bridge between social learning and dynamic rational expectations models.

The social learning literature has emphasized the possibility of inefficient outcomes in contexts with fully rational agents. For example, agents may "herd" on a wrong action disregarding valuable private information (Banerjee (1992), Bikhchandani, Hirshleifer and Welch (1992), and the monograph by Chamley (2004a)). This literature stresses market failure as an outcome of rational decisions taken by Bayesian agents. The market failure theme has a parallel in the study of excess volatility and crashes in financial markets, fads, and coordination failures (see, for example, Shiller (1981, 1984, 1989), and De Long *et al.* (1990) who, however, emphasize behavioral explanations). The literature on rational expectations and market efficiency provides a striking contrast with its emphasis on the market mechanism as an aggregator of the dispersed information of agents as we have seen in previous chapters.

In this chapter, in the tradition of the social learning literature, we restrict attention mostly to pure prediction/information externality models. That is, for the most part we obviate payoff externalities among agents to concentrate on interactions arising from information. It is well-known that in games of strategic complementarities, in which the marginal payoff of one agent is increasing in the actions of the other players, players tend to act alike or "herd". The most notable example are coordination games where multiple equilibria are the norm.[1] We introduce payoff externalities in Chapters 7, 8 and 9. The present chapter builds connections between social learning and rational expectations models. The common ground of both models is an information externality problem. The difference between the herding and the rational expectations model lies more in the "discrete" versus "smooth" nature of each model than in a difference in underlying driving forces.

The chapter introduces the dynamics of Bayesian updating and learning and central results like the self-correcting property of learning from others by Bayesian agents. It characterizes the learning process and its speed, the basic information externality present in dynamic learning environments and provides a welfare benchmark for dynamic incomplete information economies (extending the analysis in Chapters 1, 2 and 3 based on team efficiency a dynamic environment). The chapter studies questions such as the dynamics of beliefs and volatility, the effects of endogenous information acquisition, and whether more public information may hurt welfare. The latter may have important policy implications since in many instances public agencies, like a central bank, have to ponder whether to release information.

The plan of the chapter is as follows. The herding/informational cascades model as well as several extensions are presented in Section 6.1. A central smooth and noisy version model of learning from others is examined in Section 6.2. Section 6.3 explores applications to market environments and some links with dynamic rational expectations equilibria. Section 6.4 characterizes the information externality present in models of learning from others from the point of view of welfare analysis. Section 6.5 performs a welfare analysis in a static version of the model with a rational expectations flavor.

6.1. Herding, informational cascades, and social learning

Consider the following stylized example reported in Banerjee (1992). Suppose that there is a population of 100 people, each person having to choose between two unknown restaurants A and B. It is known that restaurant A is slightly better than

---

[1] See Section 4 in the Technical Appendix for brief introduction to the theory of games of strategic complementarities and Vives (Ch.2, 1999 and 2005) for more complete treatments. See Section 8.4.2 for an application to the modeling of crises.

restaurant B. More concretely, there is a common prior probability of .51 that the restaurant A is better than B. People arrive in sequence at the restaurants and each person has a private assessment of the quality of each restaurant (a noisy private signal of the same precision for everyone) and observes the choices of the predecessors. The signal provides good or bad news about restaurant A. A signal favorable to A combined with the prior knowledge would make a single person to choose this restaurant. Each person, therefore, has to decide whether to go to restaurant A or B on the basis of his signal and the previous choices of other restaurant patrons.

Suppose that 99 people have received bad news about A and the remaining person good news about A. However, it so happens that the patron with good news about A is first in line. He will decide to go to A. Then the second person in line infers that the news about B of the first in line are bad and also goes to A because this restaurant is a priori better with a slightly higher probability. Indeed, his good signal cancels with the bad draw of the first in line and he goes with the prior. The second person in line "herds" and does not follow his private information. The actions of the second person conform to those of the first one. *Herding* is typically understood as conformity of actions. When there is herding an agent may imitate the predecessor against the information contained in his own signal.

The second person in line chooses A irrespective of his signal. This means that his choice conveys no information about his signal to the third person in line. His problem is exactly the same as the second person in line and therefore he will go to restaurant A. This will happen also for the rest of the line and everyone ends up going to A although with high probability B is better. We say that the second person in line starts an informational cascade, where no further information accumulates.

In this context with a sequence of decision makers, who are imperfectly informed and each of whom observes the actions of predecessors, we say that an *informational cascade* arises when an agent, as well as all successors, make a decision independently of the private information received. Then the actions of predecessors do not provide any information to successors and therefore any learning stops. After the informational cascade starts the beliefs of the successor do not depend on the action of the predecessor. A cascade implies herding but a herd can arise even with no cascade (and in a herd there may be learning).[2] As we will see later on (Section 6.5) we could also take a more general view and define *herding* as the situation where agents put too little weight on their private signals with respect to an appropriately defined welfare benchmark.

Putting together all the private signals would indicate that B is better with probability close to one. This sequential decision making process does not aggregate information and leads to an inefficient outcome.

The insight of the example extends to a model with two states of the world, two signals and two actions ($2 \times 2 \times 2$ model). Suppose that agents have to decide in sequence whether they adopt or reject a project with unknown value $\theta \in \{0,1\}$, each with equal probability (without loss of generality). The cost of adoption is $c = 1/2$. Each agent i = 1, ..., t (t can be infinite) decides to adopt or reject ($x_i \in \{\text{adopt}, \text{reject}\}$) on the basis of a private binary, conditionally independent, signal $s_i \in \{s_L, s_H\}$ with $\Pr(s_H | \theta = 1) = \Pr(s_L | \theta = 0) = \ell > 1/2$, and the history of past actions $x^i \equiv \{x_1, ..., x_{i-1}\}$. Denote by θ_i the posterior probability that the state is high given public information $x^i : \theta_i = \Pr(\theta = 1 | x^i)$. We can call θ_i the public belief. Note that $E[\theta | x^i] = 1 \times \theta_i + 0 \times (1 - \theta_i) = \theta_i$. It follows then that there is an interval of public beliefs $(1 - \ell, \ell)$ such that for beliefs above the upper threshold ℓ everyone adopts, independently of the realized signal, and for beliefs below the lower threshold $1 - \ell$ everyone rejects, independently of the realized signal. The reason is that when the public belief θ_i is strictly above ℓ, even after receiving a bad signal, according to Bayes's formula the private belief of the agent is strictly larger than ½. The case of $\theta_i <$ $1 - \ell$ is similar (see Exercise 6.1). This means that learning takes place only when beliefs are in the interval $(1 - \ell, \ell)$, in which case an agent adopts only if he receives good news. Otherwise, the agent will herd (follow the public belief independently of his private signal) and an informational cascade will ensue. Once the cascade has started all agents choose the same action from this period on.

---

2 Çelen and Kariv (2004a) distinguish experimentally herds and cascades.

It can be shown that the probability that a cascade has not started when is the turn of i to move converges to zero exponentially as i increases, and there is a positive probability that agents herd on the wrong action. [3] The results extend to a sequential decision model where each agent moves at a time, choosing among a finite number of options, having observed the actions of the predecessors and receiving an exogenous discrete signal (not necessarily binary) about the uncertain relative value of the options (Bikhchandani, Hirshleifer and Welch (1992) or BHW for short).[4]

In the models in this family the payoff to an agent depends on the actions of others only through the information they reveal. That is, these are models of pure information externalities. In those models informational cascades occur, with agents eventually disregarding their private information and relying only on public information. Furthermore, it is possible that all agents "herd" on a wrong choice despite the fact that the pooled information of agents reveals the correct choice.

What is at the root of the extreme potential inefficiency of incorrect herds? It is a combination of an information externality, an agent when acting does not take into account the informational effects of his actions on successors, and two assumptions of the BHW model: discrete actions and signals of bounded strength.

With continuous action spaces (containing potentially optimal actions) and agents being rewarded according to the proximity of their action to the full-information optimal action convergence to the latter obtains (see Lee (1993) and Section 6.2.1). In this case the actions of agents are always sufficient statistics for their information, all information of agents is aggregated efficiently and the correct choice eventually identified. With a discrete action space (and discrete signals) there is always a positive probability of herding in a non-optimal action since agents can not fine-tune their actions to their information and actions cannot be sufficient statistics for agents' posteriors. As the set

---

3 The reader is referred to Bikhchandani et al (1998) for a complete introduction to the model and its applications.

4 Banerjee (1992) presents a similar model with a continuous action space and degenerate payoffs according to which only by hitting the right choice agents obtain a positive payoff.

of possible actions becomes richer cascades on average take longer to form and aggregate more information.

The second assumption is that signals are imperfect, identically distributed, and discrete and this implies that they are of bounded strength, which is necessary for a cascade to occur. Smith and Sørensen (2000) show, in the context of the BHW model, that if signals are of unbounded strength, then (almost surely) eventually all agents learn the truth and take the right action.[5] That is, eventually a herd on the right action occurs. With signals of unbounded strength incorrect herds are overturned by the action of an agent with a sufficiently informative contrary signal (and this individual eventually appears). With signals of (uniformly) bounded strength herding occurs (almost surely) and it may be on the wrong action. An example of signals of unbounded strength is provided by the model with two states of the world $\theta \in \{0,1\}$, two actions, and normally distributed signals $s_i = \theta + \varepsilon_i$, with error terms independently distributed across agents $\varepsilon_i \sim N\left(0, \sigma_\varepsilon^2\right)$. Then if an agent receives a very high signal he may believe that the state is high ($\theta = 1$) despite a strong public belief to the contrary.[6]

In fact, except when signals are discrete informational cascades need not arise. Chamley (2003) argues that for reasonable distributions for the signals cascades will not occur. Convergence to the correct action, however, will be slow. He shows that in the two states of the world - two action model with unbounded beliefs the public belief converges (in probability) to the truth no faster than $1/t$ while if signals were observable convergence would be exponentially faster. The reason for the slow convergence is the self-correcting property of learning from others (due to Vives (1993) and examined in detail in Section 6.3). Suppose that the state of the world is high. Then the public belief converges to $\theta = 1$. However, as the public belief tends to 1, and most agents adopt, it is increasingly unlikely that an agent appears with a sufficiently low signal so that it induces this agent to reject adoption. Since there is some probability that this agent appears the herd is informative and the public belief tends to one.

---

5 See Section 1.6 in the Technical Appendix for a formal definition of signals or beliefs of bounded strength.

6 See Chamley (2003) for a simulation of this example.

Nonetheless, because the probability of such an agent appearing tends to zero the informativeness of the herd and the rate of learning diminish.

In summary, either with continuous action spaces and regular payoffs, or with discrete action spaces and signals of potentially unbounded strength incorrect herds cannot arise and convergence to the optimal action obtains. In this sense incorrect herds are not a robust phenomenon. What is robust is the self-correcting property of learning from others and the fact that learning from others is slow when there are frictions, be it in the form of discrete actions (with a dramatic effect) or because of noise. In Section 6.3 we will see how noise slows down learning when agents can choose from a continuum of actions. Common to all models is the presence of the information externality and the associated inefficiencies it generates.

6.2 Extensions of the herding model

Several extensions of the basic model have been considered in the literature. We will consider in turn partial informational cascades, endogenous order of moves, learning from neighbors and/or reports of predecessors, and payoff externalities and reputational herding.

6.2.1 Partial informational cascade

Gale (1996) presents a simple model with the property that, although a full informational cascade can never occur, outcomes may be inefficient. Suppose that each of n agents, $i = 1, ..., n$, has to make a binary choice, to invest or not to invest in a project, and receives an independent signal s_i uniformly distributed on $[-1,1]$. The payoff to investing is given by $\theta = \sum_{i=1}^{n} s_i$. The optimal investment is achieved if all agents invest if and only if $\sum_{i=1}^{n} s_i > 0$.

If agents decide in sequence (exogenously given by $i = 1, ..., n$) then $i = 1$ invests if and only if $s_1 > 0$; $i = 2$ invests if and only if $s_2 + E[s_1 | \text{action of } i = 1] > 0$ and so on. If $i=1$ has invested then $i = 2$ invests if and only if $s_2 + E[s_1 | s_1 > 0] = s_2 + 1/2 > 0$. If $i=1$, 2 invest then $i = 3$ invests if and only if $s_3 + 3/4 > 0$. The result is that the more agents have invested the more extreme must a signal be to overturn the "partial" informational cascade (similarly as when we have discussed above the role of signal strength). Obviously, the outcome need not be efficient. For example, for n = 2, we may have both agents investing with $s_1 + s_2 < 0$. This model highlights the difference between cascades and herds. A herd may occur even if there is no cascade.

6.2.2 Endogenous order of moves

In the basic model the order in which individuals act is exogenously given. If the order of moves is endogenous then agents learn both from the actions and the delay (no-action) of other agents. There is a trade-off between the urgency of acting (impatience) and the benefit of waiting and acting with superior information.

According to Gul and Lundholm (1995) this trade-off creates clustering (similarity of agents' actions) by allowing first movers to infer some of the information of later movers and by allowing agents with more extreme signals to act first. The authors consider a continuous time model where agent i, i = 1, 2, receives an independent signal s_i uniformly distributed on $[0,1]$. Agents need to predict $\theta = s_1 + s_2$. The utility of agent i making prediction q_i at time t_i is given by $-(\theta - q_i)^2 - \alpha\theta t_i$, where $\alpha > 0$. A strategy for player i can be described in this context by a function $t_i(s_i)$ which gives the (latest) time at which the player will move given that other players have not moved and that the player has received s_i. It can be shown that at the unique symmetric equilibrium $t(s_i)$ is (strictly) decreasing and continuous and $t(1) = 0$. This means that when the first player moves it reveals his signal. Then the second agent moves immediately since there is no longer any benefit of waiting.

Clustering is explained by two factors: "anticipation" and "ordering". An agent now learns not only from predecessors but also from successors. The reason is anticipation: an agent learns something from the lack of action of another agent about the signal this agent has and this makes the prediction of the first agent similar to the successor's prediction. Furthermore, agents with extreme signals have a higher cost of waiting and will act first, revealing their signals. The forecasts of agents tend to cluster together then because of the higher impact of extreme signals on the forecasted variable. The basic results of the model still hold with n agents. Despite the fact that information is used "efficiently" in Gul and Lundholm's context the informational externalities present are not internalized and there is room for Pareto improvements.

Chamley and Gale (1994) explore the forms of market failure involved in delaying action in an investment model and how they depend on the speed of reaction of agents. Gale (1996) presents a simpler model, of which we have already seen the exogenous sequencing case. Consider now a discrete time two-agent version of the Gale (1996) model with endogenous sequencing and discount factor δ. An agent can invest in any period and his decision is irreversible. The agent with a higher signal will be more impatient to invest. The reason is that the expected value of investing in the first period equals s_i and the cost of delay is $(1-\delta)s_i$. It is possible to show that there is a unique equilibrium in which agent i invests in the first period if and only if his signal is above a certain threshold $s_i > \overline{s}$. If agent i waits he will invest in the second period if and only if someone invested in the first period (that is, $s_j > \overline{s}$ $j \neq i$). The equilibrium $\overline{s}$ must balance the cost $(1-\delta)\overline{s}$ and the option value of delaying. The latter is computed as follows. If agent i does not delay and agent j does not invest in the first period, agent i will regret its decision if $s_i + E\left[s_j \middle| s_j < \overline{s}\right] < 0$. This happens with probability $\Pr\left(s_j < \overline{s}\right)$. The equilibrium $\overline{s}$ is the unique solution then to $(1-\delta)\overline{s} = -\delta \Pr\left(s_j < \overline{s}\right)\left(\overline{s} + E\left[s_j \middle| s_j < \overline{s}\right]\right)$. In this equilibrium there is no complete information aggregation and agents may ignore their own information. The result may be that an inefficient outcome obtains (for example, it may be that $s_1, s_2 > 0$ but there is no investment because $s_i < \overline{s}$, i = 1, 2). It is worth noting that the game ends in two periods even if potentially there are many. If there has been no investment by the end of the second period this is the end of the story: investment collapses. In fact, if there is no investment in the first period this means there will never be investment because no new information will be revealed. The results generalize to more than two agents and then the game finishes in finite time.

In Chamley and Gale (1994) time is also discrete, there is discounting with factor δ, and each agent, i = 1, ..., n, receives a binary signal that provides ($s_i = 1$) or not ($s_i = 0$) an investment opportunity. The payoff to the investment $\pi(\tilde{n})$ is increasing in the realized number of investment opportunities $\tilde{n} = \sum_{i=1}^{n} s_i$. A player that invests at date t gets a payoff $\delta^{t-1}\pi(\tilde{n})$. A player that does not invest gets a zero payoff. An agent has to decide whether to invest or wait. By investing early the agent reveals that he had an investment opportunity. For any history of actions three things may happen in a symmetric PBE (in behavioral strategies).[7] If the beliefs about $\tilde{n}$ are pessimistic enough no one invests and the game ends (in the next period the situation will not change). If the beliefs about $\tilde{n}$ are optimistic enough, everyone invests and the game also ends; if beliefs about $\tilde{n}$ are intermediate, then an agent randomizes between investing now and waiting. This balances the incentives to invest if no one else does (because then by not investing nothing is learned and in the next period the agent faces the same situation) and the option value of waiting (if other people with option to invest do so then it is better to wait and learn from them). Market failure tends to lead to too little information revelation and it may involve "collapse" in the sense that no information is revealed at all and there is no investment.

The model of social learning is extended by Chamley (2004b) considers a general social learning model with irreversible investment of a fixed size for every of finite number of agents, endogenous timing, and any distribution of private information. The payoff of exercising the option in period t is given by $\delta^{t-1}(\theta - c)$ where δ is the discount factor, θ is a productivity parameter fixed by nature, not observable, and which can take a high or a low value, and $c > 0$ the cost of investment. If the agent never invests he gets a payoff of zero. At any period agents still in the market have available the history of the number of investments in each past period. The author finds that generically there may be multiple equilibria which generate very different amounts of information. In one equilibrium information revealed by aggregate activity is large and most agents delay investment (only the most optimistic invest), and in the other information revealed by aggregate activity is low and most agents rush to invest (only the most pessimistic delay). Equilibrium strategies are of the cut-off type where an agent will invest if only if his private belief about the good state is larger than some value. The presence of multiple equilibria is linked to the presence of strategic complementarities in certain regions (where larger cut-offs for the strategies of rivals imply as a best response a large

7 A behavioral strategy involves randomization over possible actions at each information set of a player (see Fudenberg and Tirole (1991)).

cut-off for the strategy of the player).[8] This is a model where strategic complementarities may arise solely out of informational externalities. The author also considers the model with a continuum of agents and noisy observation of the aggregate activity along the lines of Vives (1993). The previous results are shown to be robust to this case.

Zhang (1997) introduces also heterogeneity in the precision of the signals received by agents. This precision is also private information to the agents. There are two types of investment projects (two states) and each agent receives a two-point support signal with random precision. Time is continuous, there is discounting, and when an agent moves the choice is observed. At the unique equilibrium in pure strategies of the game there is an initial delay and then the agent with the highest precision moves first and everyone else follows immediately. A delayed investment cascade always occurs. The agent with the highest precision moves first because he has the highest expected return from investing as well as cost of waiting. The second agent will do as the first irrespective of his information because he has a signal of lower quality (given the binary signal). The rest of the agents will follow because the action of the second agent is not informative and waiting is costly. The result is inefficient for two reasons. First because of strategic delay, which is costly. Second, because of incomplete information aggregation (herding) as investment depends only on the signal of highest precision.

6.2.3 Learning from neighbors and/or reports from predecessors

In the basic model it is assumed that each agent observes the entire sequence of the actions of his predecessors. Smith and Sørensen (1995) assume that agents observe imperfect signals ("reports") of some number of predecessors' posterior beliefs and consider two cases: learning from aggregates (the aggregate number of agents taking each action, for example) and learning from samples of individuals. The latter encompasses word-of-mouth learning and bounded memory. They find in both cases that complete learning obtains eventually with unbounded informativeness of private signals. The analysis is complicated because public beliefs need not be martingales. The authors find systematic biases in the forecasts agents make of future beliefs held by

---

[8] See Vives (2005). Section 8.4.2 spells out a model with strategic complementarities.

successors, with "mean reversion" in the sampling case and the converse "momentum" in the aggregate statistic case.

Banerjee and Fudenberg (2004) have studied word-of-mouth learning in a model of successive generations making choices between two options. They find that convergence to the efficient outcome obtains if each agent samples at least two other agents, each person in the population is equally likely to be sampled, and signals are sufficiently informative. Convergence is obtained without agents observing the popularity or "market shares" of each choice.

Caminal and Vives (1996, 1999) consider a model of consumer learning about quality and firm competition where consumers learn both from word-of-mouth and market shares. (A stylized version of the consumer side of the model is presented in Section 6.4.)

Ellison and Fudenberg (1993, 1995) depart from rational learning by examining the consequences of agents using exogenously specified decision rules to learn from neighbors and with word-of-mouth communication. In Ellison and Fudenberg (1993) agents use a weighted average of past experience and popularity of choice. An early model of social memory with popularity weighting is provided in Smallwood and Conlisk (1979).

Gale and Kariv (2003) extend the social learning model to consider learning in a network and agents are allowed to choose a different action in each date. Convergence of actions is studied as well as the impact of network architecture. Çelen and Kariv (2004b) build on the model of Gale (1996) to study the case where each of a sequence of agents observes only the (binary) action of the predecessor. With imperfect information it is found that beliefs and actions cycle forever: long periods of herding can be observed but switches to the other action do occur. As time passes the herding periods become longer and longer and the switches increasingly rare.

Callander and Hörner (2006) consider a variant of the BHW model where agents are differentially informed and they do not observe the entire sequence of decisions but only the number of agents having chosen each option. Now it is not known what earlier agents knew when making their choices. This has important implications. Majorities may be wrong as in BHW, and in fact, they are more likely to be wrong than right when there is enough heterogeneity in information accuracy. Take the restaurant example (Section 6.1) but suppose that in the town there are tourists and locals. Locals have better information than tourists about the quality of the restaurant but otherwise they are indistinguishable from tourists. People arrive at the restaurant one at a time but in random order. An agent has to decide then based on the number of patrons seated at the restaurant (but, unlike in Banerjee (1992) not knowing in which order they arrived) and his private signal. The inference from the observation of a lone dissenter would vary if he were to be the first (a tourist arrived at random?) or the last (a local that knows better?). The authors show that the lone dissenter is more likely to be right if there are more tourists than locals and if their information is sufficiently more accurate. Then the wisdom of the minority holds. This is a finding of anti-herding that is purely informational (instead of reputational as we will see in the next section).

6.2.4 Payoff externalities and reputational herding

Payoff externalities can lead easily to agents taking similar actions. This is the case for example in coordination games or, more in general, games of strategic complementarities where the incremental benefit of the action of a player is increasing in the actions of other players. A typical example would be an adoption game with network externalities.

It is well-known that payoff externalities can be an obstacle to communication. In a cheap talk game an informed sender sends a message to a receiver who chooses an action that affects the payoff of the sender. Although information is not verifiable there can be communication provided that the preferences of sender and receiver have some degree of congruency. However, if preferences are completely opposed no communication is possible (see Crawford and Sobel (1982) and Vives (Section 8.3.4, 1999)).

Reputational herding models introduce informational externality considerations in principal-agent models. Typically the action of an agent affects the beliefs of a principal as well as his payoff. The payoff to the agent depends on the beliefs of the principal. Suppose that agents are of low or high ability and they want to impress the principal (but neither knows the type of an agent). Scharfstein and Stein (1990) and Graham (1999) find that if the signals of the high ability agents are positively correlated then they tend to choose the same investment projects and therefore there is an incentive for second movers to imitate first movers. This happens in a context where agents do not learn about their type when receiving their signals. Herding may occur even if signals are conditionally independent if agents learn about their type (Ottaviani and Sørensen (2000)). The same occurs if agents receive an additional signal about their type (Trueman (1994) and Avery and Chevalier (1999)). Other models in which the agents know their type and herding arises are Zwiebel (1995) and Prendergast and Stole (1996). Effinger and Polborn (2001) find that anti-herding occurs (that is, the second expert always opposes the report of the first expert) if the value of being the *only* high ability agent is sufficiently large.[9]

6.2.5 Evidence

There are several papers trying to find evidence for or against herding behavior. Evidence of herding-type phenomena in analysts' forecasts – as well as some anti-herding evidence – can be found in Graham (1999), Hong *et al.* (2000), Welch (2000), Zitzewitz (2001), Lamont (2002), Bernhardt *et al.* (2006), and Chen and Jiang (2006). For evidence in mutual fund performance see Hendricks *et al.* (1993), Grinblatt *et al.* (1995), Wermers (1999), Chevalier and Ellison (1999). See also Foster and Rosenzweig (1995) for evidence of learning from others in agriculture.

However a main problem of empirical work is that there is typically no data on the private information of agents and that the estimation of herding is not structural (and

---

[9] The authors offer the following illustration in the economics profession:
"A real-life example of antiherding can be found in the economic research industry, and in particular in the empirical branch. After the publication of an applied econometric paper, the next paper concerned with the same topic seems much more likely to find different results than the same results; moreover, the effect seems to be too great as to be explained purely by chance, even if econometric results were completely independent of the available data. A possible explanation for the phenomenon is the same as for antiherding in our model. The payoff for the second econometrician (in terms of expected reputation, publication possibilities and so on) is much higher if he finds a different result than if he finds the same result as his predecessor" (pp.386-387, Effinger and Polborn (2001)).

therefore not linked to the theory).[10] This difficulty is overcome in experimental designs. Anderson and Holt (1997) and Hung and Plott (2001) find experimental evidence in favor of the herding and informational cascades as predicted by the Bayesian theoretical model. Some of this evidence is disputed by Huck and Oechssler (2000), Nöth and Weber (2003), Kübler and Weizsäcker (2004) and Çelen and Kariv (2004a) –although the latter find that over time agents tend to Bayesian updating. The papers that dispute the predictions of the theoretical model present explanations of the experimental results based on the bounded rationality of participants. In particular, Kübler and Weizsäcker (2004) look at the case where signals in the basic herding model are costly to acquire. The theoretical prediction then is that only the first player buys a signal and makes a decision based on his information. The others herd behind the first agent. However, in the experimental results too many signals are bought. This is interpreted in terms of the limited depth of the reasoning process of players. Goeree *et al.* (2007) find that observed behavior with long sequences of decision makers does not conform to the theoretical predictions in an environment theoretically prone to information cascades. The authors develop a quantal response equilibrium model with implications that find support in the experimental data.

Guarino, Harmagart and Huck (2007) study a BHW-type of model in which only one of two possible actions is observable to others. Agents decide in some random order that they do not know. When called upon, agents are only informed about the total number of other agents who have chosen the observable action before them. The result then is that only the aggregate cascade on the observable action arises in equilibrium. A cascade on the unobservable action never arises. The authors test experimentally the model and find support in the data.

6.2.6 Conclusion

The basic model and the extensions considered are still very rough approximations to the phenomenon of social learning. Indeed, the interaction of agents is constrained to a rigid sequential procedure in which individuals take decisions in turn having observed past decisions. Although many examples have been given to apply the basic model, ranging from choice of investments, stores, technologies, candidates for office, number

[10] An exception is Cipriani, Gale and Guarino (2005) who propose a structural estimation of herd behavior with transaction data of the NYSE.

of children, drugs, medical decisions, and religion, to all kinds of fads, it is not immediately obvious that the model actually fits well any of those situations. A fortiori, the model is still far from capturing the functioning of markets in which there is an explicit price formation mechanism, agents have a large flexibility in terms of actions (quantities and/or prices, for example), interact both simultaneously and sequentially, observe aggregate statistics of the behavior of others, and the system is subject to shocks. In the next section we consider a stylized statistical prediction model which seems closer to these stylized features of markets and helps to understand some robust principles of learning from others. At the center stage there is the inefficiency of equilibrium in the presence of informational externalities. Furthermore, the model is a step in the direction of bridging the gap with dynamic rational expectations models (to be dealt with in Chapters 7, 8 and 9).

6.3. A smooth and noisy model of learning from others[11]

In this section the basic model of learning from others with noisy observation is presented, a particular case of which is the sequential decision model with smooth objective and continuous action sets. The model is extended to allow for endogenous information acquisition, and short-lived or long-lived agents.

6.3.1 Slow learning with noisy public information

Consider a model where in each period t = 0, 1,... there is a generation (a continuum) of short-lived agents, each trying to predict a random variable θ, unobservable to them. The expected loss to agent i in period t when choosing an action/prediction q_{it} is the mean squared error: $L_{it} = E\left[\left(\theta - q_{it}\right)^2\right]$.

Agent i in period t has available a *private signal* $s_{it} = \theta + \varepsilon_{it}$, where $\varepsilon_{it} \sim N\left(0, \sigma_\varepsilon^2\right)$, $\operatorname{cov}\left[\varepsilon_{it}, \varepsilon_{jt}\right] = \varsigma\sigma_\varepsilon^2$, i≠j, $\varsigma \in [0,1]$. Similarly as in Section 1.4 the convention is made that the average $\varepsilon_t = \int_0^1 \varepsilon_{it}\, di$ is a normal random variable with zero mean, and variance

[11] This section follows Sections 2 and 3 of Vives (1997) and Vives (1993).

and covariance with ε_{it} both equal to $\varsigma\sigma_\varepsilon^2$.[12] All random variables are jointly normally distributed. Let $\bar{\theta} = 0$. When $\varsigma = 0$ there is no correlation between the error terms of the signals and, again by convention, $\varepsilon_t = 0$ (a.s.). The agent has also available a public information vector $p^{t-1} = \{p_0, ..., p_{t-1}\}$, where p_t is the average action of agents in period t plus noise, $p_t = \int_0^1 q_{it} di + u_t$, and $\{u_t\}_{t=0}^{\infty}$ a white noise process. In short, agent i in period t has available the information vector $I_{it} = \{s_{it}, p^{t-1}\}$.

Agent i in period t solves then

$$\text{Min}_q E\left[(\theta - q)^2 | I_{it} \right]$$

and sets $q_{it} = E[\theta | I_{it}]$. It follows that the period expected loss L_{it} is given by $L_{it} = E\left[E\left[(\theta - E[\theta | I_{it}])^2 | I_{it} \right]\right] = E[\text{var}[\theta | I_{it}]]$. (Given the information structure L_{it} will be symmetric across players and, because of linearity and normality $\text{var}[\theta | I_{it}]$ will be nonrandom.)

Remark: It is worth noting that the formal analysis of the model would be unchanged if agent i in period t had an *idiosyncratic* expected loss function $L_{it} = E\left[(\theta + \eta_{it} - q_{it})^2 \right]$ with η_{it} being a random variable with zero mean and finite variance σ_η^2, independently distributed with respect to the other random variables of the model. In this case $L_{it} = E\left[(\theta - q_{it})^2 \right] + \sigma_\eta^2$ and the agent faces the same minimization problem as before. With idiosyncratic loss functions private signals can be thought to be *endogenous* and represent *word-of-mouth* communication. When agent i in period t obtains his payoff we may suppose he learns $\theta + \eta_{it}$ and communicates it to a friend of the next generation. That is, the signal received by i in period $t+1$ is $\theta + \eta_{it}$. This

12 This convention is in accord with the finite dimensional version of the stochastic process of the error terms. See Section 2.3 in the Technical Appendix.

corresponds to the model provided η_{it} has the same properties as the error terms ε_{it}. In this case the realized payoffs in period t generate information for other agents in period $t+1$.

If the signals of agents of the same generation are perfectly correlated ($\varsigma = 1$), then we have sequential decision making as in the basic herding model with minor variations but with transmission noise. In this case there is a *representative agent* each period and the model is purely *sequential* with agents taking actions in turn.

Agents act simultaneously in every period and noise avoids that their average action fully reveals θ. The dynamics of public information are easily characterized. Let us posit that the strategies of agents in period t are linear (and symmetric given the symmetric information structure): $q_{it} = a_t s_{it} + \varphi_t\left(p^{t-1}\right)$, with a_t the weight to private information and φ_t a linear function of p^{t-1}. Now, the current public statistic is given by

$$p_t = \int_0^1 q_{it} di + u_t = a_t\left(\theta + \varepsilon_t\right) + u_t + \varphi_t\left(p^{t-1}\right)$$

Where $\varepsilon_t = \int_0^1 \varepsilon_{it}\, di$ is normally distributed. The public signal at period t, p_t is a linear function of $z_t = a_t\left(\theta + \varepsilon_t\right) + u_t$ and past public signals p^{t-1} and is normally distributed since it is a linear combination of normal random variables. Consequently $q_{it+1} = E\left[\theta \middle| s_{it+1}, p^t\right]$ is again a linear function of s_{it+1} and p^t. Further, it is clear that for t = 0 the induction process to claim linearity of the solution can be started since $E\left[\theta \middle| s_{i0}\right]$ is linear in s_{i0} because of normality.

Letting $z_t = a_t\left(\theta + \varepsilon_t\right) + u_t$ we have that $p_t = z_t + \varphi_t\left(p^{t-1}\right)$ and that the vector of public information p^t can be inferred from the vector z^t and vice versa. The variable z_t is the new information about θ in p_t. From normality of the random variables it is immediate that the conditional expectation $\theta_t \equiv E\left[\theta \middle| p^t\right] = E\left[\theta \middle| z^t\right]$ is a sufficient statistic for public information p^t in the estimation of θ. It follows that the sequence of public beliefs $\{\theta_t\}$ follows a martingale (with $\theta_{-1} \equiv \bar{\theta}$ and relative to the history of public information z^t)

$$E\left[\theta_t \middle| \theta_{t-1}\right] = E\left[E\left[\theta \middle| z^t\right] \middle| z^{t-1}\right] = E\left[\theta \middle| z^{t-1}\right] = \theta_{t-1}$$

since when conditioning the coarser information set (i.e. the one with less information) dominates.[13] Since the conditional expectation is a sufficient statistic for normal random variables it follows also that $\theta_t = E\left[\theta \middle| \theta_t\right]$ and therefore $\mathrm{var}\left[\theta\right] = \mathrm{var}\left[\theta \middle| \theta_t\right] + \mathrm{var}\left[E\left[\theta \middle| \theta_t\right]\right] = \mathrm{var}\left[\theta \middle| \theta_t\right] + \mathrm{var}\left[\theta_t\right]$.

Let $\tau_t \equiv \left(\mathrm{var}\left[\theta \middle| \theta_t\right]\right)^{-1}$ denote the informativeness (precision) of public information $\theta_t \equiv E\left[\theta \middle| z^t\right]$, with $z_k = a_k\left(\theta + \varepsilon_k\right) + u_k$, k = 0, …, t, in the estimation of θ. Noting that we can condition equivalently on the variables $\hat{z}_k = \theta + \varepsilon_k + \left(a_t\right)^{-1} u_k$, and recalling that $\mathrm{var}\left[\varepsilon_t\right] = \varsigma\sigma_\varepsilon^2$ and $\mathrm{var}\left[a_t^{-1} u_k\right] = a_t^{-2}\sigma_u^2$, we obtain (see Section 2.1 in the Technical Appendix) that

$$\tau_t = \tau_{t-1} + \left(\varsigma\tau_\varepsilon^{-1} + \left(a_t^2\tau_u\right)^{-1}\right)^{-1} = \tau_\theta + \sum_{k=0}^{t}\left(\varsigma\tau_\varepsilon^{-1} + \left(a_k^2\tau_u\right)^{-1}\right)^{-1}.$$

The result is that the random vector $\left(s_{it}, \theta_{t-1}\right)$ is sufficient in the estimation of θ based on $I_{it} = \left\{s_{it}, z^{t-1}\right\}$ (that is, $E\left[\theta \middle| s_{it}, z^{t-1}\right] = E\left[\theta \middle| s_{it}, \theta_{t-1}\right]$). The posterior mean of θ with information $I_{it} = \left\{s_{it}, \theta_{t-1}\right\}$ is a weighted average of the signals received with weights according to their precisions (the private signal with precision τ_ε and the public with precision τ_{t-1}):

$$E\left[\theta \middle| s_{it}, \theta_{t-1}\right] = a_t s_{it} + \left(1 - a_t\right)\theta_{t-1}, \text{ with } a_t = \tau_\varepsilon / \left(\tau_\varepsilon + \tau_{t-1}\right).$$

---

13 See Section 3 in the Technical Appendix. It is in fact a general result that posterior (Bayesian) beliefs have the martingale property.

From the martingale property of $\{\theta_t\}$ we have that (i) $\text{cov}[\Delta\theta_t, \Delta\theta_{t-1}] = 0$ where $\Delta\theta_t = \theta_t - \theta_{t-1}$ and (ii) $\text{var}[\Delta\theta_t] = \text{var}[\theta_t] - \text{var}[\theta_{t-1}]$. Furthermore, because of normality, $\text{var}[\theta_t] - \text{var}[\theta_{t-1}] = \text{var}[\theta_t | \theta_{t-1}]$. (See Section 3.3 of the Technical Appendix.)

We have then that the total volatility of public information up to period t $\sum_{k=0}^{t} \text{var}[\Delta\theta_k]$ equals $\sum_{k=0}^{t} \text{var}[\theta_k | \theta_{k-1}]$ and adds up to current volatility $\text{var}[\theta_t]$:

$$\sum_{k=0}^{t} \text{var}[\Delta\theta_k] = \sum_{k=0}^{t} \text{var}[\theta_k | \theta_{k-1}] = \text{var}\left[\sum_{k=0}^{t} \Delta\theta_k\right] = \text{var}[\theta_t],$$

since $\sum_{k=0}^{t} \Delta\theta_k = \theta_t - \bar{\theta}$ and from the martingale properties (i) and (ii). Furthermore, since $\text{var}[\theta] = \text{var}[\theta | \theta_t] + \text{var}[\theta_t]$, we obtain that $\text{var}[\theta_t] = \tau_\theta^{-1} - \tau_t^{-1}$, and therefore $\text{var}[\theta_t | \theta_{t-1}] = \text{var}[\theta_t] - \text{var}[\theta_{t-1}] = \tau_{t-1}^{-1} - \tau_t^{-1}$.

As a benchmark let us take the case where there is no noise in the public statistic ($\tau_u = \infty$). Then the order of magnitude of $\tau_t = \tau_\theta + (t+1)\tau_\varepsilon \varsigma^{-1}$ is t for $\varsigma > 0$. In this case the new information in p_t, $a_t(\theta + \varepsilon_t)$, reveals the relevant information of agents. Indeed, even though $a_t \xrightarrow{t} 0$ as t tends to infinity, $a_t > 0$ and therefore the sequence of noisy signals $\{\theta + \varepsilon_t\}$ can be inferred from the sequence of new information $\{a_t(\theta + \varepsilon_t)\}$. Therefore, there is learning about θ as t grows and learning is at the standard rate $1/\sqrt{t}$ (as in the usual case of i.i.d. noisy observations of θ) since the order of magnitude of τ_t is t. When $\tau_u = \infty$ there is no information externality since public information is a sufficient statistic for the information of agents.

However, when there is noise in public information a notable property of the information dynamics of the system is that public precision is accumulated unboundedly but at a slow rate. The result is a manifestation of 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, 1997)). Indeed, the weight given to private information a_t is decreasing in the precision of public information τ_{t-1}, and the lower a_t is the less information is incorporated in the public statistic p_t because $\tau_t = \tau_{t-1} + \left(\varsigma\tau_\varepsilon^{-1} + \left(a_t^2\tau_u \right)^{-1} \right)^{-1}$. A higher inherited precision of public information τ_{t-1} induces a low current response to private information a_t, which in turn yields a lower increase in public precision $\tau_t - \tau_{t-1}$. In this sense learning from others is self-defeating. Conversely, a lower inherited precision of public information τ_{t-1} induces a high current response to private information a_t, which in turn yields a higher increase in public precision. In this sense learning from others is self-enhancing.

The self-enhancing aspect means that public precision τ_t will be accumulated unboundedly. If this were not the case the weight given to private precision $a_t = \tau_\varepsilon / \left(\tau_\varepsilon + \tau_{t-1} \right)$ would be bounded away from zero, necessarily implying that $\tau_t = \tau_{t-1} + \left(\varsigma\tau_\varepsilon^{-1} + \left(a_t^2\tau_u \right)^{-1} \right)^{-1}$ grows unboundedly, a contradiction.[14] As τ_t tends to infinity, $a_t = \tau_\varepsilon / \left(\tau_\varepsilon + \tau_{t-1} \right)$ tends to zero.

The self-defeating aspect means that accumulation is slow. Note that $a_t\tau_t \xrightarrow{t} \tau_\varepsilon$. Indeed, $\lim_{t\to\infty} a_t\tau_t = \lim_{t\to\infty} a_t\tau_{t-1} = \lim_{t\to\infty} \left(\tau_\varepsilon^{-1} + \tau_{t-1}^{-1} \right)^{-1} = \tau_\varepsilon$ since $a_t \xrightarrow{t} 0$ and $\tau_t \xrightarrow{t} \infty$. It is possible to show (Vives (1993, 1997) that

$$t^{-1/3}\tau_t \xrightarrow{t\to\infty} \left(3\tau_u\tau_\varepsilon^2 \right)^{1/3}.$$

The result implies that τ_t grows at the rate of $t^{1/3}$. This slow learning result is quite remarkable since it means that for a large number of observations, if to attain a certain level of public precision (approximately) 10 rounds more are needed when there is no noise in the public statistic $\left(\tau_u = \infty \right)$ we need (approximately) 1000 additional rounds to obtain the same precision in the presence of noise $\left(\tau_u < \infty \right)$. This is the difference

14 Similarly, in Banerjee and Fudenberg (2004) 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.

between the usual linear rate of learning t and the cubic concave rate of $t^{1/3}$ The result demonstrates also that the social learning model with perfect observation of the actions of others is not robust.

The rate of learning (accumulation of public precision) is independent of the level of noise. However, the asymptotic precision (or constant of convergence) $\left(3\tau_u\tau_\varepsilon^2\right)^{1/3}$ increases with less noise (higher τ_u) and more precise signals (higher τ_ε). This asymptotic precision influences the "slope" of convergence. (See Figure 6.2 -case $\delta = 0$- below for the shape of τ_t when $\tau_u < \infty$; when $\tau_u = \infty$ τ_t would grow linearly with t.) More noise in the public statistic or in the signals slows down learning of θ by decreasing the asymptotic precision but it does not alter the convergence rate. Indeed, the asymptotic precision (or inverse of the asymptotic variance) is a refined meaure of the speed of convergence for a given convergence rate.[15]

The proof of the result is somewhat involved.[16] An heuristic argument to show that τ_t is of the order of $t^{1/3}$, $\tau_t \approx t^{1/3}$ for short, that $a_t \approx t^{-1/3}$, and that $t^{-1/3}\tau_t \xrightarrow{t} \left(3\tau_u\tau_\varepsilon^2\right)^{1/3}$ runs as follows. Let $\tau_t \approx Kt^\upsilon$ for some $K > 0$ and $\upsilon > 0$. Then $a_t \approx \tau_\varepsilon K^{-1}t^{-\upsilon}$ (because $a_t \approx \tau_\varepsilon\tau_t^{-1}$ from $a_t = \tau_\varepsilon / \left(\tau_\varepsilon + \tau_{t-1}\right))$ and therefore $\tau_t = \tau_\theta + \sum_{k=0}^{t}\left(\varsigma\tau_\varepsilon^{-1} + \left(a_k^2\tau_u\right)^{-1}\right)^{-1} \approx \tau_u\tau_\varepsilon^2K^{-2}\sum_{k=0}^{t}k^{-2\upsilon}$. The term $\sum_{k=0}^{t}k^{-2\upsilon}$ can be seen to be of the order of $t^{1-2\upsilon}/\left(1-2\upsilon\right)$.[17] Now the equality $\upsilon = 1-2\upsilon$ implies that $\upsilon = 1/3$. Furthermore, $K = 3\tau_u\tau_\varepsilon^2K^{-2}$ and therefore $K = \left(3\tau_u\tau_\varepsilon^2\right)^{1/3}$. The result obtains because as t grows unboundedly, a_t tends to zero and so does the amount of new information incorporated into p_t, which is represented by $z_t = a_t\left(\theta + \varepsilon_t\right) + u_t$.

---

15 See Section 3.2 in the Technical Appendix for a discussion of measures of speed of convergence.

16 It follows from an extension of Lemma A1 of the Appendix of Chapter 7 (from Vives (1993) to the case $\varsigma > 0$ and the fact that $a_t\tau_t \xrightarrow{t} \tau_\varepsilon$. A version of the proof of the result (adapted from Chamley (2004a) is provided in the Appendix.

The fact that τ_t tends to infinity as t grows implies that θ_t tends to θ with probability one (almost surely) and in mean square. The latter is easily checked because, given that $E[\theta_t]=E[\theta]$, $E\left[(\theta-\theta_t)^2\right]=\mathrm{var}[\theta-\theta_t]$ and, given that $\theta_t=E\left[\theta\middle|\theta_t\right]$, $\mathrm{var}[\theta-\theta_t]=\mathrm{var}\left[\theta\middle|\theta_t\right]=\tau_t^{-1}$, and $\tau_t^{-1}\xrightarrow{t}0$. In fact, convergence (a.s and in mean square) of the sequence of public beliefs to a limit random variable follows directly from the martingale convergence theorem because $\{\theta_t\}$ is a bounded martingale (since $\mathrm{var}[\theta_t]\le\mathrm{var}[\theta]<\infty$, see Section 3.3 in the Technical Appendix). Furthermore, the speed of learning is slow, at the rate of $1/\sqrt{t^{1/3}}$: $\sqrt{t^{1/3}}(\theta_t-\theta)$ tends in distribution to a normal random variable with zero mean and variance $(3\tau_u\tau_\varepsilon^2)^{-1/3}$. This is immediate because $t^{1/3}\mathrm{var}[\theta-\theta_t]=t^{1/3}\tau_t^{-1}\to(3\tau_u\tau_\varepsilon^2)^{-1/3}$ as t tends to infinity. Proposition 6.1 summarizes the slow learning results. Denote by $\xrightarrow{L}$ convergence in law (distribution).

<u>Proposition 6.1 </u>(Vives (1993, 1997)). Slow learning. As t tends to infinity:

(i) a_t tends to 0 and τ_t tends to infinity;

(ii) θ_t tends to θ almost surely and in mean square;

(iii) $t^{-1/3}\tau_t\xrightarrow{t}(3\tau_u\tau_\varepsilon^2)^{1/3}$; and

(iv) $\sqrt{t^{1/3}}(\theta_t-\theta)\xrightarrow{L}N\left(0,(3\tau_u\tau_\varepsilon^2)^{-1/3}\right)$

It is interesting to analyze the implications of the results for the evolution of the *cross section of beliefs* in the population. Identify the belief of agent i with his prediction

$$q_{it}=E\left[\theta\middle|s_{it},\theta_{t-1}\right]=a_ts_{it}+(1-a_t)\theta_{t-1}\text{, with } a_t=\tau_\varepsilon/(\tau_\varepsilon+\tau_{t-1}).$$

The distribution of beliefs will be normal and characterized by its average $q_t=\int_0^1 q_{it}di$ and dispersion $\int_0^1(q_{it}-q_t)^2di$. Let $\varsigma=0$ to illustrate. Noting that $q_t=a_t\theta+(1-a_t)\theta_{t-1}$, using our convention $\int_0^1 s_{it}di=\theta$ (a.s.), $\theta_t=E\left[\theta\middle|z^t\right]=\left(\tau_u\sum_{k=1}^t a_kz_k\right)/\tau_t$ and

---

[17] See Lemma A3 in the Appendix to Chapter 7.

therefore $E\left[\theta_t \left|\theta\right.\right] = \left(1-\tau_\theta/\tau_t\right)\theta$ (see Section 2.1 in the Technical Appendix) we obtain easily that

$$E\left[q_t \left|\theta\right.\right] = \left(1-\tau_\theta/\left(\tau_\varepsilon+\tau_{t-1}\right)\right)\theta$$

and

$$E\left[\int_0^1 \left(q_{it}-q_t\right)^2 di\right] = E\left[\left(q_{it}-q_t\right)^2\right] = \tau_\varepsilon/\left(\tau_\varepsilon+\tau_{t-1}\right)^2.$$

Given that τ_t grows as $t^{1/3}$ we obtain that $E\left[q_t \left|\theta\right.\right]$ increases monotonically in a concave way to θ (since the function $\varphi(\tau) = -\tau_\theta/\left(\tau_\varepsilon+\tau\right)$ is concave) and $E\left[\left(q_{it}-q_t\right)^2\right]$ decreases monotonically in a convex way to 0 (since the function $\phi(\tau) = \tau_\varepsilon/\left(\tau_\varepsilon+\tau\right)^2$ is convex).

With perfect correlation of the signals of agents ($\varsigma = 1$) we have a representative agent each period and we are in the context of the herding models of last section with the representative agent of each generation taking actions in sequence. With no noise in public information then actions are fully revealing of the information of agents and public precision τ_t grows at the rate of n. With noisy observation of the actions the self-correcting property of learning from others implies that τ_t grows much more slowly, at the rate of $t^{1/3}$ despite the continuous action space.

With an agent at every period we may think also that observational noise comes from the very action that an agent takes because of an idiosyncratic element. Denote the agent in period t, agent t, then with the idiosyncratic expected loss function $E\left[\left(\theta+u_t-q_t\right)^2\right]$ if agent t observes the random shock u_t then his action will be $q_t = E\left[\theta\left|s_t,\theta_{t-1}\right.\right]+u_t$ and successors of t by observing the history of actions have only noisy observations of signals of the predecessors.

Herd behavior or informational cascades are extreme manifestations of the self-defeating aspect of learning from others. With discrete action spaces and signals of bounded informativeness public information may end up overwhelming the private signals of the agents, who may (optimally) choose not to act on their information.

At the other extreme, an instance where convergence obtains at the standard rate $1/\sqrt{t}$ is when at every period there is a positive mass of perfectly informed agents. The reason is that perfectly informed agents can not learn from public information and therefore put constant weight on their (perfect) signals. The consequence is that at every round the amount of information incorporated into the public statistic is bounded away from zero ($a_t \xrightarrow{t} a_\infty > 0$) and learning is not self-defeating. We will examine this situation in Section 7.4. Exercise 7.4 generalizes Proposition 6.1 to the case of a general distribution of private precisions in the population (including the case that a positive mass have infinite precision).

Larson (2006) analyzes the representative agent version of the model ($\varsigma = 1$) but where in each period some agents may die and the representative agent observes the average action of all agents still alive with no noise. It is found that if the population size is stable (with agents departing at the same rate they arrive) then public precision accumulates at the fast rate of t (as in the representative agent version of the model with no noise). However, if the population grows then the rate of accumulation of public precision slows down to $\log(t)$. The combination of population growth and the fact that agents cannot perfectly observe the order in which predecessors acted explains the result. The reason is that by observing the average action the error terms in early private signals persist for a long time since early, less informed, decisions cannot be distinguished from later, better informed ones.

6.3.2 Endogenous information acquisition[18]

Suppose now that the error terms in the signals are uncorrelated ($\varsigma = 0$) and that private signals have to be purchased at a cost, increasing and convex in the precision τ_ε of the

[18] The analysis follows Burguet and Vives (2000).

signal, according to a smooth function $C(\cdot)$ that satisfies C(0) = 0, $C'>0$ for $\tau_\varepsilon>0$, and $C''\geq 0$. There are, thus, nonincreasing returns to information acquisition. The model is otherwise as before, in particular, there are t = 0, 1, 2, … generations of short-lived agents and the information structure is the same. Each agent is interested in minimizing the sum of the prediction loss and the information costs:

$$\text{Min}_{(q,\tau_\varepsilon)}\left(E\left[(\theta-q)^2 \middle| I\right]+C(\tau_\varepsilon)\right)$$

where I the information set of the agent with a private signal of precision τ_ε -the choice variable- and a public signal which summarizes public information history. The linear-normal structure means that, as in Section 6.3.1, public information follows also a normal distribution. Denote the precision of the public signal by τ. Notice that the solution to the problem is $q = E\left[\theta\middle|I\right]$ and, given τ_ε and τ, the prediction loss is just the inverse of the total precision of the information of the agent: $L(\tau_\varepsilon,\tau)=\text{var}\left[\theta\middle|I\right]=(\tau_\varepsilon+\tau)^{-1}$. Thus, for given inherited precision of public information τ, the representative agent will minimize over τ_ε

$$\Lambda(\tau_\varepsilon,\tau)=L(\tau_\varepsilon,\tau)+C(\tau_\varepsilon).$$

The expected loss $\Lambda(\tau_\varepsilon,\tau)$ is strictly convex in τ_ε and there will be a unique solution to the minimization problem $\gamma^m(\tau)$ as a function of τ. Let $g^m(\cdot)$ be the market policy function that yields the dynamics of public precision public precision τ. Given τ and private precision purchase $\gamma^m(\tau)$ the weight to private information is $a=\dfrac{\gamma^m(\tau)}{\gamma^m(\tau)+\tau}$, and therefore public precision in the following period is given by $g^m(\tau)=\tau+\tau_u\left[\dfrac{\gamma^m(\tau)}{\gamma^m(\tau)+\tau}\right]^2$. The following proposition characterizes $\gamma^m(\cdot)$ and $g^m(\cdot)$.

<u>Proposition 6.3</u> (Burguet and Vives (2000)). If $\tau\geq\left[C'(0)\right]^{-1/2}$, then $\gamma^m(\tau)=0$. Otherwise, $\gamma^m(\tau)>0$ and $\gamma^m(\cdot)$ is a (strictly) decreasing, differentiable function of τ.

If $C'(0)=0$, then $\gamma^m(\cdot)\to 0$ as $\tau\to\infty$. The market policy function g^m is increasing for τ large enough (for τ close enough to $\left[C'(0)\right]^{-1/2}$ if $C'(0)>0$).

Proof.: From $\left.\frac{\partial\Lambda}{\partial\tau_\varepsilon}\right|_{\tau_\varepsilon=0}=C'(0)-\tau^{-2}$ we have that $\gamma^m=0$ whenever $\tau\geq\left[C'(0)\right]^{-1/2}$. Otherwise, the solution is interior, $\gamma^m(\tau)>0$, and the first order condition (F.O.C.) of the minimization problem yields $\gamma-\left[C'(\gamma)\right]^{-1/2}=-\tau$. The left hand side is strictly increasing and ranges from $-\infty$ to $+\infty$. Therefore this F.O.C. defines implicitly γ^m. From the implicit function theorem, $\frac{d\gamma^m}{d\tau}=-\frac{1}{1+\frac{1}{2}C''(\gamma)\left[C'(\gamma)\right]^{-3/2}}$, and therefore, $-1\leq\frac{d\gamma^m}{d\tau}<0$. Now, from $g^m(\tau)=\tau+\tau_u\left[\frac{\gamma^m(\tau)}{\gamma^m(\tau)+\tau}\right]^2$ we have that

$\frac{dg^m}{d\tau}=1+2\tau_u\frac{\gamma^m}{\left(\gamma^m+\tau\right)^3}\left(\frac{d\gamma^m}{d\tau}\tau-\gamma^m\right)$, which is positive for γ^m close enough to 0 (and/or τ large enough). This shows that g^m is increasing for τ large enough and for τ close enough to $\left[C'(0)\right]^{-1/2}$ when $C'(0)>0$. ◊

It is worth noting that the policy function need not be increasing for any τ. (See Figure 6.1 where a case where the function is quasiconvex is depicted) For example, $dg^m(0)/d\tau=1-2\tau_u/\gamma^m(0)<0$ for τ_u large enough. This means that starting with a larger public precision it is possible that at the following iteration we end up with a lower public precision than if we have started from a lower base. This is possible because private (τ_ε) and public (τ) precisions are strategic substitutes in the minimization of the expected period loss Λ. That is, $\partial^2\Lambda/\partial\tau_\varepsilon\partial\tau>0$. It implies that the purchase of private information is decreasing in the amount of (inherited) public precision. An increase in τ reduces the incentives to purchase private information and moderates the increase in τ since the weight to private information $a=\left(1+\tau/\tau_\varepsilon\right)^{-1}$ is reduced. This is a manifestation of the self-correcting property of learning from others.

As an important consequence of this there are instances where more public information hurts. More precisely, aggregate the losses across generations of agents with discount factor $\delta > 0$ and let the aggregate loss be $\ell^m(\tau) \equiv \sum_{t=0}^{\infty} \delta^t L_t$ where L_t is evaluated at the market solution for an initial public precision τ. Then $\ell^m(\tau)$ can be increasing in τ. For instance, it can be checked that this happens if $C(\gamma) = \gamma$, $\tau_u = 2$, and $\delta = .98$ for $\tau \leq .61$. (Exercise 6.2 examines the case where the cost function C is linear.)[19] The result is akin to the remark by Smith and Sørensen (1995) that the observation of larger samples of predecessors does not necessarily improve welfare at the market solution. In a somewhat related vein, Banerjee (1993) in a model of the economics of rumors finds that speeding up the transmission of information (the rumor) has no welfare effect since then the rumor must be received sooner to be trusted (in his model after a certain, endogenously determined time, a rumor is not trusted).

Note that even in the model of Section 6.3.1 with exogenous private signals the same effect is potentially present. Indeed, when the private precision is fixed at $\tau_\varepsilon > 0$ (and $\varsigma = 0$) then we have that the policy function is $g(\tau) = \tau + \tau_u \left(\frac{\tau_\varepsilon}{\tau_\varepsilon + \tau} \right)^2$, which is strictly convex, can be easily seen to be increasing for $\tau \geq 8\tau_u / 27$, but it may be decreasing for τ_u large. Note also that $g(\tau) > \tau$ and $g' \to 1$ as $\tau \to \infty$. With $(8\tau_u / 27) > \tau_\theta > \tau_\theta'$ it may happen that $g(\tau_\theta) < g(\tau_\theta')$ and starting from a lower base more public precision is accumulated since, because of the convexity of g, after the first period we are in the increasing portion of g.

It is remarkable that more public information may hurt even in an environment where there are no payoff externalities and only information externalities matter. Morris and Shin (2002, 2005), for example, find in a static model that more information may hurt because of a special "beauty contest" form of the payoff of agents that induces agents to have a private incentive to coordinate which is socially wasteful. (See Section 1.4.3.).

---

19 However, for these parameter values it is easily seen that $dg^m / d\tau < 0$ and therefore the range $\tau \leq .61$ is never "visited" after the initial period for any market sequence τ_t^m (see Figure 6.1).

We will confirm robustly in the next section that informational externalities alone suffice to find that more public information may hurt.

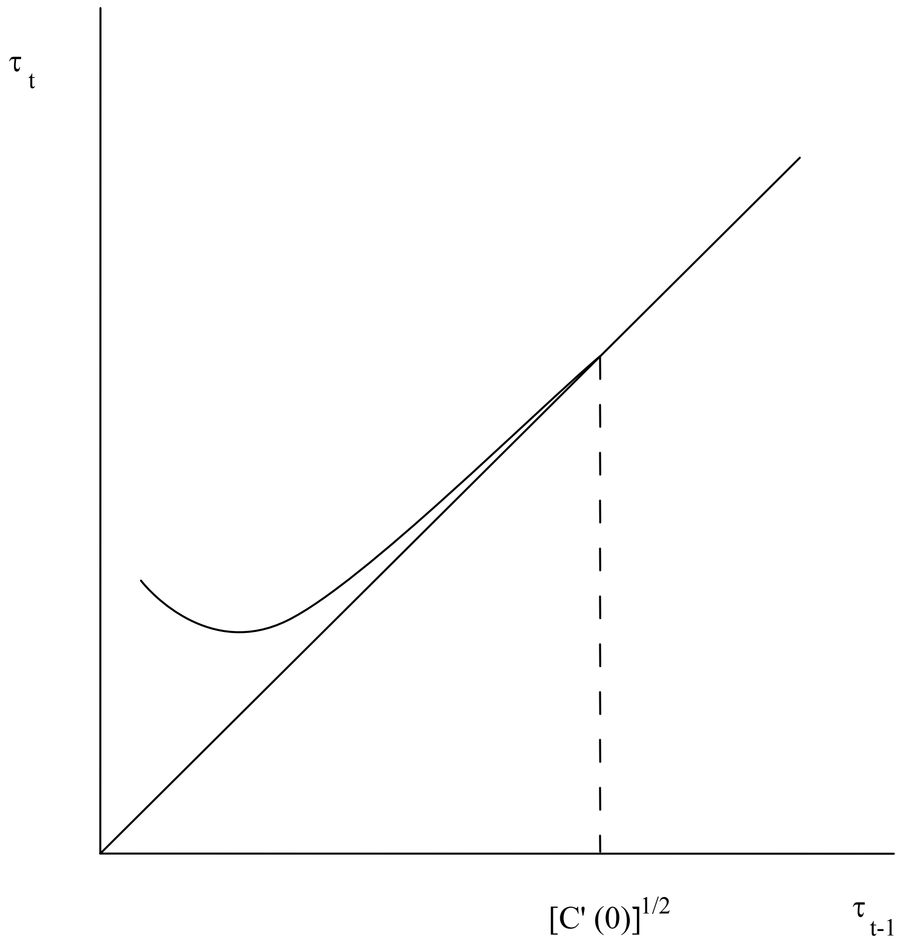

Figure 6.1. The Market Policy Function $g^m_{\cdot}(\cdot)$

Full revelation of θ obtains with endogenous private precisions if and only if the marginal cost of acquiring information when there is no information is zero ($C'(0)=0$). Indeed, if this marginal cost is bounded away from zero, public precision cannot accumulate without bound since then the marginal benefit of acquiring private information would tend to zero and would be strictly inferior to the marginal cost of information acquisition.

The stated facts provide a hint for a dynamic resolution of the Grossman-Stiglitz paradox (see Section 4.2.2) when $C'(0)=0$. In that case, the market is asymptotically strongly efficient in informational terms -that is, fully revealing of θ- but agents have an incentive to purchase private information all along. The contradiction between eventual full revelation of θ and incentives to acquire information disappears.

We can see also that if $C'(0)=0$, the speed of learning decreases as we move away from the exogenous signals situation. This can be done considering a family of cost functions parameterized by their degree of convexity λ: $C(\tau_\varepsilon)=c\lambda^{-1}(\tau_\varepsilon)^\lambda$, with $\lambda \geq 1$, $c > 0$. For this example it is easily seen that $g^m(\tau)$ is convex. The relative degree of convexity of C is $\tau_\varepsilon C''/C'=\lambda-1$, and as λ tends to infinity agents are given signals of precision 1 at no cost. Then for $\lambda > 1$, it can be shown that τ_t is of the order of t^r with $r=(\lambda-1)/(3\lambda+1)$ and a_t of the order of $t^{-\upsilon}$ with $\upsilon=(\lambda+1)/(3\lambda+1)$. We have that $\upsilon > 1/3 > r > 0$ and r is increasing in λ. As λ tends to infinity r tends to 1/3, which is the order of magnitude of public precision when private signals are given exogenously (Proposition 6.1). Convergence is faster the more decreasing returns to information acquisition are (the more convex is C and the closer we are to the exogenous information case as parameterized by λ). (See Exercise 6.3)

It is worth to remark also that with our natural specification of costly information acquisition, and contrary to the results of Radner and Stiglitz (1984), the value of information need not be nonpositive at zero. Indeed, the value of precision τ_ε is $\phi(\tau_\varepsilon)=-\left((\tau_\varepsilon+\tau)^{-1}+C(\tau_\varepsilon)\right)$ and $\phi'(0)=\tau^{-2}-C'(0)$, which can be positive or negative. With our specification the nonconcavity in the value of information of Radner and Stiglitz does not arise.[20] Indeed, $\phi''<0$ for all $\tau_\varepsilon \geq 0$.

6.3.3 Long-lived agents

The model in Section 6.3.1 admits an interpretation in terms of a continuum of long-lived agents interacting repeatedly in the market t = 0, 1,... . The agents are rewarded according to the proximity of their prediction to some random variable θ, unobservable to them. At any period there is an independent probability $1-\delta$, with $1>\delta>0$, that θ is realized and the payoffs up to this period collected. The expected loss to agent i in period t when choosing an action/prediction q_{it} is the mean squared error: $L_{it}=E\left[(\theta-q_{it})^2\right]$. The agent has available in period t a private signal $s_i=\theta+\varepsilon_i$ (the same for every period) and a public information vector $p^{t-1}=\{p_0,\dots,p_{t-1}\}$ as before. Agent i in period t has available the information vector $I_{it}=\{s_i,p^{t-1}\}$. Signals are conditionally independent with the same precision $\tau_\varepsilon=1/\sigma_\varepsilon^2$, and as usual we make the convention that errors on average cancel out (so that in this case the average signal reveals θ).

Myopic behavior is optimal for an agent, solving $\text{Min}_q E\left[(\theta-q)^2 \mid I_{it}\right]$ and setting $q_{it}=E\left[\theta \mid I_{it}\right]$. Indeed, an agent is infinitesimal and can not affect the public statistics. Agents act simultaneously in every period and noise avoids that their average action fully reveals θ. The model is formally identical to the model in Section 6.3.1 when $\varsigma=0$. In this case public precision is given by $\tau_t=\tau_\theta+\tau_u\sum_{k=0}^{t}a_k^2$.

Amador and Weil (2007) provide a continuous time extension of the model (for the case $\varsigma=0$). Agent i receives the payoff $\int_0^T(\theta-q_{it})^2 di$ at the random time T when θ is realized. T has an exponential distribution with intensity $\kappa>0$: θ is realized at time t with probability density $\kappa e^{-\kappa t}$. This means that the expected payoff of agent i is $E\left(\int_0^T(\theta-q_{it})^2 di\right)=E\left(\int_0^\infty e^{-\kappa t}(\theta-q_{it})^2 dt\right)$. Note that this equivalent to our formulation according to which the probability that θ is realized at time t is δ^t by letting $\delta=e^{-\kappa}$. The nice feature of continuous time is that a closed-form solution can be provided. It is found that

---

20 See Section 3 in Radner and Stiglitz (1984) to see how a crucial assumption of their theorem fails in our case. See also Kihlstrom (1984) for further analysis of the potential nonconcavity of the value of information.

$$\tau_t = \left(3\tau_u\tau_\varepsilon^2 t + (\tau_\varepsilon + \tau_\theta)^3\right)^{1/3} - \tau_\varepsilon .$$

From this result it is immediate then that τ_t is of the order of $t^{1/3}$ and that $t^{-1/3}\tau_t \xrightarrow{t} \left(3\tau_u\tau_\varepsilon^2\right)^{1/3}$. Note that now $\partial\tau_t / \partial\tau_\theta > 0$ for all t while in the discrete time version we have seen that this need not be the case in the first round.

Amador and Weil (2007) extend the model by allowing agents to receive at time zero a private exogenous signal with precision τ_ε and a public exogenous signal of precision τ_o about the unknown θ plus a private and a public signal about the average action with, respectively, precisions τ_e and τ_u. That is, it adds to the original formulation an exogenous public signal at t = 0 as well as a private signal about the average action for any t. The private learning channel may arise, for example, when an agent samples from a distribution centered on the aggregate action. Now, the precision of private information evolves according to γ_t and the precision of public information according to τ_t yielding a response to private information of $\gamma_t / (\gamma_t + \tau_t)$ and a mean squared prediction error at t of $(\gamma_t + \tau_t)^{-1}$.

The author obtains two interesting results. The first is that the path of the average prediction $q_t = \int_0^1 q_{it} di$ conditional on θ, $E[q_t | \theta] = \left(1 - \tau_\theta / (\gamma_t + \tau_t)\right)\theta$, has an S-shaped diffusion pattern if private information is sufficiently dispersed initially (i.e. τ_ε is low enough). Then learning accelerates at the beginning with agents putting more weight on their private information and the public signal becoming more and more informative. Eventually as agents learn θ the effect tappers off. Furthermore, due to learning from the private channel the dispersion of beliefs $E\left[(q_{it} - q_t)^2\right] = \gamma_t / (\gamma_t + \tau_t)^2$ can be seen to be hump-shaped. At the beginning beliefs are concentrated close to the common prior. As agents learn privately beliefs become dispersed only to converge later as θ is revealed.

The second result is that an increase in initial public precision τ_θ increases the total precision $\gamma_t + \tau_t$ in the short run but decreases it in the long run since it decreases uniformly the endogenous private precision γ_t. In fact, the same is true if the noise in the public signal is reduced (i.e. τ_u increases). These results hold provided that the private learning channel is active (i.e., provided that $\tau_e > 0$). This opens the door to an increase in total disounted expected loss $E\left(\int_0^\infty e^{-\kappa t}\left(\theta - q_{it}\right)^2 dt\right)$ as a result of an increase in τ_θ. The authors show that a marginal increase in public information hurts as long as the payoff is realized in a sufficiently long time (i.e. the mean of the random end T κ^{-1} is large enough) and/or the amount learned by unit of time (as measured by a common scaling factor in τ_e and τ_u) high enough. However, as expected, a sufficiently large increase in public precision would be good. As in the basic models of Sections 6.3.1 and 6.3.2 welfare eventually increases with public precision starting from a high enough base.

6.3.4 Summary

We have provided in this section a central result on slow learning from others in the presence of frictions (i.e. noise). This is the outcome of the self-correcting property of learning from others. With no frictions the public statistic is sufficient for the signals of predecessors and learning is at the usual rate. However, this result is not robust to the presence of noise in the public statistic. If there is noise the public statistic is not sufficient. If the public statistic is more informative then agents rely less on their private information and less of this private information is incorporated into the public statistic. This slows down learning. We could say that the weight of history slows down learning. However, if the public statistic is less informative then the opposite happens and this speeds up learning and ensures eventual full revelation. The results obtain both with short-lived and long-lived agents.

An important associated result is that public and private precision are strategic substitutes from the point of view of the decision maker. This means that increases in the initial stock of public information may discourage the use of private information and end up in a lower path of accumulation of public precision and welfare. This is particularly so when agents acquire private signals about the fundamental since then better public information discourages private information acquisition. Futhermore, if agents have to acquire their private signals then full revelation will be precluded if the marginal cost of acquiring information is positive at 0. Otherwise, full revelation obtains but the speed of learning decreases as we move away from the exogenous signal case.

Even more so, the potential damaging effect of public information is more pronounced when there is an active private learning channel about aggregate activity. We find therefore that even in a world without payoff externalities and where agents are interested only in predicting the fundamentals, more public information may hurt.

6.4 Applications and examples

This section presents several examples and applications of the basic smooth learning model of Section 6.3 and a brief discussion of the relationship of the smooth herding model with dynamic rational expectations models.

6.4.1 Examples

The first two examples involve short-lived agents, the others involve long-lived agents.

Consumers learning about quality

As we have already observed the formal analysis of the model would be unchanged if agents had an *idiosyncratic* expected loss function $L_{it} = E\left[\left(\theta + \eta_{it} - q_{it}\right)^2\right]$ with η_{it} 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 $L_{it} = E\left[\left(\theta - q_{it}\right)^2\right] + \sigma_\eta^2$.

Suppose now that in each period there are many consumers of two types: "rational" and "noise" or "random". All consumers are endowed with a utility function that is linear with respect to money. Consumers only differ in their information and are one-period lived. Generation t consumer i's utility when consuming q_{it} is given by:

$$U_{it} = \left(\theta + \eta_{it}\right) q_{it} - \frac{1}{2} q_{it}^2 .$$

The willingness to pay of consumer i in period t is $\theta + \eta_{it}$. Consumers are uncertain about $\theta + \eta_{it}$ 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 i in period t receives a signal s_{it}

about θ (word of mouth from the experience of a previous consumer or an independent test of the product). Given that the idiosyncrasy η_{it} of consumer i in period t is uncorrelated with all other random variables in the environment and that the consumer learns $\theta+\eta_{it}$ only after consuming the good we have that $E\left[\theta+\eta_{it}\middle|I_{it}\right]=E\left[\theta\middle|I_{it}\right]$. 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 $q_{it}=E\left[\theta\middle|I_{it}\right]$. If u_t denotes the purchases of the random consumers, then aggregate demand will be: $p_t=\int_0^1 E[\theta|I_{it}]\,di+u_t$. Consumers active in period t have access to the history of past sales $p^{t-1}=\{p_0,p_1,...,p_{t-1}\}$. Consumer i's information set in period t is therefore $I_{it}=\left\{s_{it},p^{t-1}\right\}$. The model is thus formally identical to the one presented in Section 6.3 (under the same distributional assumptions). Indeed, the expected welfare loss with respect to the full-information first best (where θ is known and $q_{it}=\theta$) is easily seen to be $E\left[\left(\theta-q_{it}\right)^2\right]/2$. The results imply that consumers will learn slowly quality from quantities consumed or market shares. Furthermore, slow learning by consumers enhances the possibilities of firms of manipulating consumer beliefs (for example, signal-jam the inferences consumers make from market shares, see Caminal and Vives (1996, 1999)).

<u>Location decisions and information acquisition</u>

Consider a world where an earthquake (the "big one" in California) may strike at any period with probability $1-\delta$. The location θ is the safest from the point of view of the earthquake, the problem is that θ is unknown and it will not be known until the earthquake happens! Agents have to make (irrevocable) location decisions based on their private (costly) assessment of θ and any public information available. The latter consists of the average location decisions of past generations. These average locations contain an element of noise since for every generation there are agents who locate randomly independently of any information. The private assessment of an agent is based on geological research he conducts. The higher the effort the agent spends on this research the better estimate he obtains. This example would correspond then to the endogenous information acquisition case considered in Section 6.3.2.

Macroeconomic forecasting and investment

Consider competitive firms deciding about investment in the presence of macroeconomic uncertainty, which determines profitability, represented by the random variable θ. At each period there is an independent probability $1-\delta$ that the uncertainty is resolved. Firms invest taking into account that the profits of their accumulated investment depend on the realization of θ. The investment of a firm is directly linked to its prediction of θ. To predict θ each firm has access to a private signal as well as to public information, aggregate past investment figures compiled by a government agency. Data on aggregate investment incorporates measurement error.[21] In consequence, at each period a noisy measure of past aggregate investment of the past period is made public. The issue is whether, and if so how fast, the repeated announcement of the aggregate investment figures reveals θ.

Reaching consensus and common knowledge

At a more abstract level, consider the reaching of consensus starting from disparate expectations. It has been shown that repeated public announcements of a stochastically monotone aggregate statistic of conditional expectations, which need not be common knowledge, leads to consensus (McKelvey and Page (1986) and Nielsen *et al.* (1990) following up on Aumann (1976)). In the iterative process, individuals compute conditional expectations with the information they have available and the aggregate statistic is announced. Individuals compute again then 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 nevertheless market interaction will provide agents only with 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. In the model of Section 6.3.3 repeated public

[21] For example, quarterly data on national accounts are subject to measurement error. See Rodríguez-Mora and Schulstald (2007) for the consequences of this measurement errors on output.

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 can not disagree forever but we can disagree for a long time."

Learning by doing

Models of *learning by doing* 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. A typical model of learning by doing assumes that the unit cost of production with an accumulated production of t is of the form $C(t) = kt^{-\lambda}$ with λ between 0 and 1 and k a constant. A rate of cost reduction of $t^{-1/3}$ 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%.[22] 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 coordination problem takes a very simple (and extreme) form in the model: Costs are lower the closer the actions of workers are to an unknown parameter θ. The total expected cost of output in production round t is proportional to $\int_0^1 E\left[(\theta - q_{it})^2\right] di$ where q_{it} is the action of worker i in period t. Worker interaction reveals the statistic p_t. The model of Section 6.3.3 predicts that the rate of learning, as given by the precision of public information τ_t, and consequently the period loss (expected unit

22 See Scherer and Ross (1990, p. 98-99). The following quotation is from Arrow's seminal paper on learning by doing (1962): "It was early observed by aeronautical engineers, particularly T.P. Wright, that the number of labor-hours expected in the production of an airframe (airplane body without engines) is a decreasing function of the total number of airframes of the same type previously produced. Indeed, the relation is remarkably precise; to produce the nth airframe of a given type, counting from the inception of production, the amount of labor required is proportional to $t^{-1/3}$".

cost), will be of the order of $t^{-1/3}$. This means that expected cost will decline at the rate $t^{-1/3}$. So we would have λ = 1/3 in $C(t) = kt^{-\lambda}$.

6.4.2 Relation to dynamic rational expectations and preview of results

Do the results obtained in Section 6.3 extend to more complex economic situations, closer to the actual functioning of markets? In those situations payoff externalities will matter and will interact with the informational externalities examined in our prediction model. In this section we overwiew briefly some of the models with payoff and informational externalities that we will deal with in the following chapters. One issue of interest, as we will see in Chapter 7, is whether convergence to full-information equilibria will obtain at a (relatively) fast rate in markets environments. Otherwise, the Walrasian model may not be a good approximation of the behavior of agents in the economy even when repeated interaction has given the opportunity to prices to reveal information.

The smooth noisy model of learning from others of Section 6.3 is close to classical dynamic rational expectations models. In the latter prices are noisy aggregators of dispersed information and agents choose from a continuum of possible actions with smooth payoffs as rewards. Nevertheless, the slow learning results can not be applied mechanistically. Indeed, markets are more complex than the simple models of the previous sections. A general reason is that, unlike in the pure learning/prediction model, in market models the payoff of an agent depends directly on the actions of other agents. That is, there are payoff externalities. For example, the profit of a firm depends on the outputs of rivals firms.

Furthermore, learning need not be always from others, agents can learn also from the environment. An example is provided by the classical learning in rational expectations partial equilibrium model with asymmetric information (as developed by Townsend (1978) and Feldman (1987), see Sections 7.2 and 7.3). In this model long-lived firms, endowed with private information about an x_t uncertain demand parameter θ, compete repeatedly in the marketplace. Inverse demand in period t is given by $p_t = \theta + u_t - x_t$, with x_t being average output, and production costs are quadratic. The result is that learning θ and convergence to the full-information equilibrium (or shared-information equilibrium in which θ is revealed) occur at the standard rate ($1/\sqrt{t}$).[23] The reason is that public information (prices) depend directly (independently of the actions of agents) on the unknown parameter θ. Even if agents had no private information they would learn from prices at the standard rate because, for a given output , the price observation p_t corresponds to an i.i.d. noisy signal $\theta + u_t$ of θ.

In contrast, in a variation of the classical model (Vives (1993), see Section 7.4) where the unknown θ is a cost parameter (a pollution damage, for example) prices will be informative about θ only because they depend on the actions of firms, and the strength of the dependence will vanish a t grows large due to the self-defeating facet of learning from others. Indeed, expected profit maximization yields an optimal production which is linear in the conditional expectation of θ and in past prices, and the result is that, as in the purely statistical model of Section 6.3, firms learn slowly about θ and convergence to the full-information equilibrium is also slow.

The speed at which prices reveal information is particularly important in financial markets where price or value discovery mechanisms are in place. An example is provided by the information tâtonnement designed to decrease the uncertainty about prices after a period without trade (overnight) in the opening batch auction of some continuous stock trading systems. At the beginning of trade there is a period where agents submit orders to the system and *theoretical* prices are quoted periodically as orders accumulate. No trade is made until the end of the tâtonnement and at any point agents may revise their orders. It is important that this information aggregation mechanism performs its value discovery purpose fast. A stylized version of this mechanism is considered in Section 9.1.3 where it is shown that information is aggregated at a fast rate in the presence of a competitive market making sector while without it convergence is slow. The reason why market makers (or agents using limit orders/demand schedules and supplying liquidity to the market) speed up convergence is as follows. Market makers, by expanding market depth, induce risk averse agents not to respond less to their private information as prices become more informative about the

23 As we will see in Chapter 7 (following Jun and Vives (1996)), in general, learning θ and converging to a full-information equilibrium are not equivalent.

fundamental value of the risky asset. The presence of market makers prevents the self-defeating facet of learning from others from settling in. The outcome is that price quotations converge to the underlying value of the asset at a rate of $1/\sqrt{t}$ where t is the number of periods.

In a related vein, Avery and Zemsky (1998) show how introducing a competitive market making sector in the basic herding model (by adapting the Glosten and Milgrom (1985) model of sequential trading in financial markets, see Section 9.1.1) convergence of the price to the fundamental value obtains. Indeed, the basic herding model (Section 6.1) does not allow for price changes once agents buy or sell the good. Furthermore, and interestingly, the authors show that when traders are uncertain about the quality of information (precision) held by other traders then "booms and crashes" based on herd behavior are possible. The reason is that a poorly informed market may not be easily distinguishable from a well informed one from the point of view of competitive market makers. A similar explanation of market crashes has been offered by Romer (1993). (See Section 8.4.1.)

In summary, in the market examples considered there is learning from others and we see how changes in the market microstructure have consequences for convergence and the speed of learning. This shows that caution must be exercised when applying herding and slow learning results to market situations. A fuller development of these models is provided in Chapters 7, 8 and 9.

6.5 The information externality and welfare[24]

In this section we perform a welfare analysis of the prediction model in Section 6.3. We are particularly interested on whether agents put too little weight on their private information with respect to a well-defined welfare benchmark. The welfare analysis in herding models has been typically neglected in the literature. Here we address the issue in the smooth learning from others model developed in Section 6.3.

[24] The analysis in this section follows Vives (1997) and Burguet and Vives (2000).

At the root of the inefficiencies detected in models of learning from others lies an information externality. An agent when making its decision (prediction) does not take into account the benefit to other agents. Consider the basic model of Section 6.3.3 (with long-lived agents). The (ex-ante expected) loss of a representative agent in period t, $L_t = E\left[\left(\theta - E\left[\theta | I_{it}\right]\right)^2\right] = \mathrm{var}\left[\theta | I_{it}\right] = \left(\tau_\varepsilon + \tau_{t-1}\right)^{-1}$, is decreasing in public precision τ_{t-1}. Note that L_t is independent of i since the information structure is symmetric. A larger response of agents to their private signals in period t will lead to a larger precision in period t, τ_t, and consequently a lower loss in period $t+1$ (and in subsequent periods). This is not taken into account by an individual agent.

The analysis of the information externality leads naturally to a *welfare-based definition of herding* as an excessive reliance on public information with respect to a well-defined welfare benchmark. This, as in Chapter 1, is the team solution (Radner (1962) which assigns to each agent a decision rule so as to minimize the discounted sum of period losses. The planner or team manager, however, cannot manipulate the information flows (in particular, it has no access to their private information). This solution internalizes the externality respecting the decentralized information structure of the economy. Optimal learning at the team solution then trades off short-term losses with long-run benefits. That is, it involves experimentation.

Welfare losses are discounted with discount factor δ, 1 > δ ≥ 0 (consistently with the unknown θ having an independent probability of $1-\delta$ of being revealed in any period in the model with long-lived agents). The planner is restricted to use linear rules (for agent i, $q_{it}\left(I_{it}\right)$ where, as before, $I_{it} = \left\{s_i, p^{t-1}\right\}$ and $p_k = \int_0^1 q_{ik} di + u_k$); the same family of simple rules that agents in the market use. This means that the planner needs to convey to the agents only the weight they should put on their private information. This corresponds to a low level of complexity in the instructions of the team manager to the agents.

The team manager has an incentive to depart from the myopic minimization of the short term loss and "experiment" to increase the informativeness of public information. This is accomplished imposing a response to private information above the market response.

We look first at a two-period model, move on to the infinite horizon model, and conclude with results when information is costly to acquire.

6.5.1 A two-period example

Let us illustrate the team problem with a *two-period* example, t = 0, 1. The team has to choose linear decision rules $q_0(s_i)$ and $q_1(s_i, p_0)$ to minimize $L_0 + \delta L_1$, where $L_t = E\left[(\theta - q_t)^2\right]$. This is easily solved by backwards recursion and a unique linear team solution is found. Posit $q_0(s_i) = a_0 s_i$ (since by assumption $\overline{\theta} = 0$). It is immediate that $L_0 = E\left[(\theta - q_0(s_i))^2\right] = \frac{(1-a_0)^2}{\tau_\theta} + \frac{a_0^2}{\tau_\varepsilon}$ and that the precision of public information for period 1 is $\tau_0 = \tau_\theta + \tau_u a_0^2$. In period 1 for a given τ_0, the team solution is just the market (Bayesian) solution and therefore $L_1 = (\tau_\varepsilon + \tau_0)^{-1}$.

The first order condition to minimize $L_0 + \delta L_1$ which turns out to be sufficient (since $L_0 + \delta L_1$ is quasiconvex in a_0), is given by

$$\frac{\partial L_0}{\partial a_0} + \delta \frac{\partial L_1}{\partial a_0} = 2\left(-\tau_\theta^{-1} + a_0\left(\tau_\theta^{-1} + \tau_\varepsilon^{-1}\right) - \delta \frac{a_0 \tau_u}{(\tau_\varepsilon + \tau_0)^2}\right) = 0$$

This yields a unique solution a_0^o between the market solution $a_0^m = \tau_\varepsilon / (\tau_\varepsilon + \tau_\theta)$ and 1.

It is clear that $\frac{\partial L_1}{\partial a_0} < 0$ and consequently $\frac{\partial L_0}{\partial a_0} > 0$ at the optimum. The market solution a_0^m involves $\frac{\partial L_0}{\partial a_0} = 0$ and therefore, given convexity, $a_0^o > a_0^m$.

Comparative static properties can be derived in the standard way obtaining: $\frac{\partial a_0^o}{\partial \delta} > 0$, $\frac{\partial a_0^o}{\partial \tau_\theta} < 0$ (using the fact that $1 > a_0$) and $\frac{\partial a_0^o}{\partial \tau_\varepsilon} > 0$. Furthermore, $\text{sign}\left(\frac{\partial a_0^o}{\partial \tau_u}\right) = \text{sign}\left(\tau_\theta + \tau_\varepsilon - \tau_u (a_0^o)^2\right)$ and $\frac{\partial a_0^o}{\partial \tau_u}$ is positive if τ_u is small ($\tau_\theta + \tau_\varepsilon > \tau_u$) and negative if τ_u is large ($\tau_\theta + \tau_\varepsilon < \tau_u\left(\tau_\varepsilon / (\tau_\theta + \tau_\varepsilon)\right)^2$).

The information externality implies that there is underinvestment in public information at the market solution $a_0^o > a_0^m \equiv \tau_\varepsilon / (\tau_\theta + \tau_\varepsilon)$. The comparative static results are intuitive: the optimal response to private information a_0^o increases with the discount factor (the weight given to the future) and with the precision of private information, and it decreases with the prior precision about θ. That is, it is optimal to experiment more whenever the future matters more, the quality of private information is better or the precision of prior information worse.

An increase in the noise of public information induces agents to experiment less (more) if the starting level of noise is low (high). Recall that the market solution in period 0, $\tau_\varepsilon / (\tau_\theta + \tau_\varepsilon)$, is independent of τ_u. The optimal solution does depend on τ_u. When τ_u is small there is a lot of noise in public information and reducing the noise induces more experimentation with a higher a_0^o; when τ_u is large there is little noise in public information and reducing more the noise induces less experimentation because it is costly and what matters is the informativeness of the public statistic.

It should be obvious also that welfare increases at the team solution with increases of either τ_θ or τ_u. An increase in τ_θ reduces L_1 and L_0 for any a_0; an increase in τ_u reduces L_1 and does not affect L_0. The situation is potentially different at the market solution since there is an indirect negative effect of an increase in τ_θ on $a_0^m \equiv \tau_\varepsilon / (\tau_\theta + \tau_\varepsilon)$ and therefore on τ_0.

6.5.2 <u>The infinite horizon model</u>

Let us turn now to the infinite horizon problem. Given decision rules $q_{it}(I_{it})$, where I_{it} is the information of agent i in period t, the average expected loss in period t is $\int_0^1 E\left[\left(\theta - q_{it}(I_{it})\right)^2\right] di$. Furthermore, it can be checked that there is no loss of generality in restricting attention to symmetric rules (see Exercise 6.4): $q_{it}(I_{it}) = q_t(I_{it})$ for all i. The expected loss in period t is then $L_t = E\left[\left(\theta - q_t(I_{it})\right)^2\right]$. The objective of the team is to minimize $\sum_{t=0}^{\infty}\delta^t L_t$ choosing a sequence of linear functions $\{q_t(\cdot)\}_{t=0}^{\infty}$ where q_t is a function of $I_{it}=\{s_i,p^{t-1}\}$ and $p_k=\int_0^1 q_k\left(s_i,p^{k-1}\right)di+u_k$.

The analysis proceeds as in the market case. As before, we can write the strategy of agent i in period t as $q_k\left(s_i,\theta_{t-1}\right)=a_t s_i+c_t\theta_{t-1}$, with a_t and c_t the weights, respectively, to private and public information (with θ_{t-1} the sufficient statistic for public information). The precision of public information θ_t equals $\tau_t=\tau_\theta+\tau_u\left(\sum_{k=0}^{t}a_k^2\right)$, and therefore only the coefficients a_t matter in the accumulation of information (the coefficients c_t have no intertemporal effect). At period t, c_t is chosen, contingent on a_t, to minimize L_t. This is accomplished setting $c_t=1-a_t$. It is easily seen then that $L_t=E\left[\left(\theta-q_k\left(s_i,\theta_{t-1}\right)\right)^2\right]=\frac{\left(1-a_t\right)^2}{\tau_{t-1}}+\frac{a_t^2}{\tau_\varepsilon}$ where $\tau_{t-1}=\tau_\theta+\tau_u\sum_{k=0}^{t-1}a_k^2$.

The reduced form team minimization problem is then to choose a sequence of real numbers $\{a_t\}_{t=0}^{\infty}$ to solve

$$\text{Min }\sum_{t=0}^{\infty}\delta^t L_t\text{ , with }L_t=\frac{\left(1-a_t\right)^2}{\tau_{t-1}}+\frac{a_t^2}{\tau_\varepsilon}$$

where $\tau_{-1}=\tau_\theta$ and $\tau_{t-1}=\tau_\theta+\tau_u\sum_{k=0}^{t-1}a_k^2$ for $t=1,2,...$

The team problem can be posed in a classical dynamic programming framework taking the sequence $\{\tau_t\}_{t=0}^{\infty}$ as control (noting that $\tau_t-\tau_{t-1}=\tau_u a_t^2$). It can be checked that the value function $\Lambda(\cdot)$ associated to the control problem is the solution to the functional equation

$$\Lambda(\tau)=\min_{\tau'}\left\{L(\tau',\tau)+\delta\Lambda(\tau')\right\},$$

where $L(\tau',\tau)=\frac{(1-\sqrt{(\tau'-\tau)/\tau_u})^2}{\tau}+\frac{\tau'-\tau}{\tau_u\tau_\varepsilon}$.

It can be shown that $\Lambda : \mathbb{R}_+ \to \mathbb{R}$ is strictly convex, twice-continuously differentiable and strictly decreasing with $\Lambda' < 0$. As τ tends to infinity $\Lambda'(\tau)$ and $\Lambda(\tau)$ tend to 0. Indeed, Λ is strictly decreasing in the accumulated precision of public information τ because a higher accumulated precision today generates uniformly (strictly) lower period losses for all feasible sequences from then on. At the team solution, which takes care of the dynamic information externality, increasing the precision of public information is unambiguously good. This need not be the case at the market solution because a higher initial public precision might discourage the use of private information and subsequent accumulation of public precision.

The policy function $g^o(\cdot)$ gives the unique solution to the team problem: next's period public precision as a function of the current one. It can be shown (see proof of Proposition 4.4 in Vives (1997)) that g^o is quasiconvex, continuously differentiable, with $dg^o/d\tau > 0$ for τ large enough, and $dg^o/d\tau < 1$ (with $dg^o/d\tau$ tending to 1 as τ tends to infinity). Given that $g^o(\tau) > \tau$, it is immediate then that τ tends to infinity and that a decreases over time (since $a_t^2 = (\tau_t - \tau_{t-1})/\tau_u = (g^o(\tau_{t-1}) - \tau_{t-1})/\tau_u$).

The results for $\delta > 0$ are as follows. First of all, there is herding, in the sense that for any given accumulated precision of public information τ, the team optimal response to private information a^o is strictly larger than the market solution $a^m = \tau_\varepsilon/(\tau_\varepsilon + \tau)$ (in particular, $a_0^o > a_0^m$). Given the function $L(\tau', \tau)$ defined above denote by $\partial_1 L$ the partial derivative with respect the first argument and by $\partial_{11} L$ the second partial with respect the first argument. The result follows from the first order condition of the team problem $\partial_1 L(g^o(\tau), \tau) + \delta\Lambda'(g^o(\tau)) = 0$ because $\Lambda' < 0$ and therefore $\partial_1 L(g^o(\tau), \tau) > 0$, while at the market solution τ^m (with $\delta = 0$) $\partial_1 L(\tau^m, \tau) = 0$. Given that $\partial_{11} L > 0$ we conclude that $g^o(\tau) - \tau = \tau_u (a^o)^2 > \tau_u (a^m)^2 = \tau^m - \tau$. Secondly, the market underinvests in information: the team solution at any period has accumulated more public precision than the market ($\tau_t^o > \tau_t^m$). The second result follows from the properties of the policy function g^o.

For $\delta = 0$ the market solution is obtained ($a_t^o = a_t^m$ for all t). With $\delta > 0$ and given a certain accumulated public precision the optimal program calls for a larger response to private information because it internalizes the benefit of a larger response today in lowering future losses. Agents herd and rely too little on their private information at the market solution. However, this does not mean that the optimal program involves a uniformly larger response to private information overtime because it accumulates more public precision. Simulations[25] for $\delta > 0$ show that there is a critical $\overline{t}$, increasing in δ and τ_ε, and decreasing in τ_u, after which the optimal program calls for a lower response to private information than the market to collect the benefits of the initial accumulation/experimentation phase. That is, there is a $\overline{t}$ is such that $a_t^o > a_t^m$ for $t < \overline{t}$, and $a_t^o < a_t^m$ for $t > \overline{t}$. Comparative statics with the discount factor δ are as follows:

(i) For any t, τ_t is increasing in δ. (See Figure 6.2.)

(ii) For t large (small) L_t is decreasing (increasing) in the discount factor δ.

This is in accordance with the "investment" and "consumption" phases of the accumulation program.

---

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

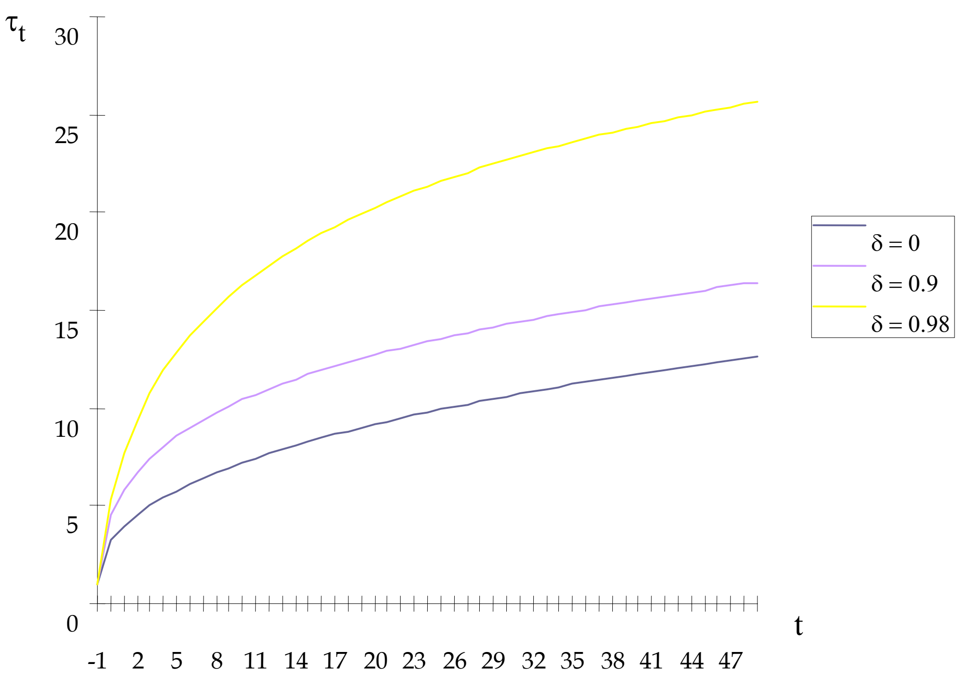

Figure 6.2 - Precision of public information and discount factors.
$\tau_\theta = 1, \tau_\varepsilon = 2, \tau_u = 5$

Let $\ell^m = \sum_{t=0}^{\infty} \delta^t L_t$, with L_t the period loss at the market solution and $\ell^o = \Lambda(\tau_\theta)$, the value at the team solution. Then, the relative welfare loss of the market solution with respect to the team solution $(\ell^m - \ell^o)/\ell^m$ can be quite high. For example, $(\ell^m - \ell^o)/\ell^m$ is around 25% for $\tau_u = 5$, $\tau_\varepsilon = .5$, $\delta = .95$, and $\tau_\theta = 1$. The relative welfare loss is increasing in δ and nonmonotonic in τ_ε and τ_u. 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 there is almost perfect information and almost no loss. For $\tau_\varepsilon = 0$ there is no private information. For τ_u very large public information is almost fully revealing and for $\tau_u = 0$ it is uninformative.

However, slow learning at the market solution is not suboptimal. The team solution has exactly the same asymptotic properties as the market. The reason is that for t large the market and the optimal program look similar because the value function of the latter is almost flat for large public precision. In both, the responsiveness to private information tends to zero as t grows. The result follows as in the market case since at the optimal solution $a_t \tau_{t-1}$ tends to τ_ε as t tends to infinity. This is so because at the optimal solution: $\partial_1 L(\tau_t, \tau_{t-1}) + \delta \Lambda'(\tau_t) = 0$ for any t. As t tends to infinity so does τ_t and consequently $\Lambda'(\tau_t)$ tends to 0. Therefore $\partial_1 L(\tau_t, \tau_{t-1}) = \tau_\varepsilon^{-1} - (1 - a_t)(a_t \tau_{t-1})^{-1}$ tends to zero as t tends to infinity and $a_t \tau_{t-1} \xrightarrow{t} \tau_\varepsilon$ and $a_t \xrightarrow{t} 0$. This is the only way to enjoy the benefits of the accumulated public precision because to accumulate public precision faster would entail too large a departure from the minimization of current expected losses. In particular, to put asymptotically a constant positive weight on private information and consequently obtain an increase in precision which is linear in t is not optimal.

Proposition 6.3 summarizes the results obtained.

<u>Proposition 6.3</u> (Vives (1997)). Optimal learning. Let $\delta > 0$, the team solution:

(i) Responds more to private information, for any given τ, than the market solution $\tau_\varepsilon / (\tau + \tau_\varepsilon)$

(ii) Period by period accumulates more public precision than the market ($\tau_t^o > \tau_t^m$).

(iii) Has the same asymptotic properties as the market; in particular, the same rate of (slow) learning.

The properties of the optimal learning program when agents are short-lived and signals potentially correlated are similar, and the same results hold. The analysis of optimal learning proceeds as before with minor variations. The presence of correlation in the signals tends to decrease the optimal weight to private information. For example, in the two-period optimal learning problem it is easily seen that a_0^o decreases with ς. The reason is that with correlated error terms in the signals increasing the weight in private signals also increases the weight in the aggregate error, which now is not zero. Obviously, when $\tau_u = \infty$ there is no information externality since public information is a sufficient statistic for the information of agents (although with $\varsigma > 0$ public information is not perfectly revealing of θ) and the market solution is optimal.

In terms of the applications presented in Section 6.4 the optimal learning results imply the following. In the learning by doing example the team problem is to coordinate workers to minimize the expected costs of production. Suppose that the team manager can impose decision rules on the workers which are measurable in the information they have. Then the theory predicts that independently of whether the team manager behaves myopically or as a long-run optimizer (taking into account the learning externality), the rate of learning as given by the precision of public information τ_t, and consequently the period loss (expected unit cost), will be of the order of $t^{-1/3}$. This means that expected cost will decline at the rate $t^{-1/3}$. In the consumer learning example the results imply that consumers are too cautious with respect to the welfare benchmark in responding to their private information.

6.5.3 Costly information acquisition

When private information is costly to acquire as in Section 6.3.2 the effects of the information externality are accentuated. Consider the same model as in Section 6.3.2 and a second best welfare benchmark in which private information purchases can be controlled, via tax-subsidy mechanisms, but otherwise agents are free to take actions (make predictions). Given a sequence of private precisions $\{\gamma_t\}$ chosen by the planner, an agent at period t will choose the action that minimizes his loss, taking τ_{t-1} and γ_t as given. The agent will put a weight on his private signal equal to $a_t = \gamma_t / (\gamma_t + \tau_{t-1})$ inducing a period expected loss of $(\gamma_t + \tau_{t-1})^{-1}$. The problem of the planner is to choose a sequence of nonnegative real numbers $\{\gamma_t\}_{t=0}^{\infty}$ to solve

$$\text{Min} \sum_{t=0}^{\infty} \delta^t \left(\frac{1}{\gamma_t + \tau_{t-1}} + C(\gamma_t) \right)$$

where $\tau_{-1} = \tau_\theta$ and $\tau_t = \tau_{t-1} + \tau_u \left(\dfrac{\gamma_t}{\gamma_t + \tau_{t-1}} \right)^2$ for $t = 1, 2, \ldots$

Burguet and Vives (2000) characterize the solution to the program. Similarly to the market solution if $C'(0) > 0$ public precision does not accumulate unboundedly. We have seen in Section 6.3.2 how an increase in initial public precision may hurt welfare at the market solution. Indeed, private information acquisition at the market solution is decreasing in inherited public precision and therefore the reduction of the prediction losses today may make more costly to decrease future losses. The same is true for the second best benchmark. For instance, an increase in τ_θ increases the discounted expected losses in the range $\tau \in (.28,.54)$ for $\tau_u = 2$, $C(\gamma) = 2\gamma$ and $\delta = .98$. The reason is the self-defeating aspect of learning from others. Increasing public precision today reduces the weight agents put on private information (for a given precision τ_ε) tomorrow. This means that to maintain a certain weight on private information (which is what determines the increase in the next period public precision) more effort has to be devoted to acquire private information (to raise τ_ε).

However, it can be shown that, for a large enough initial public precision τ, more public precision is always good at the second best solution. The reason is that for τ large enough, the weight to private information has to be small and the indirect effect of an increase in τ in discouraging private information acquisition cannot be large while the direct effect is the same. Denote by $\gamma^{sb}(\tau)$ the optimal purchase of private precision at the second best solution for a given public precision τ. It follows that for large τ there is herding at the market solution, whenever $\gamma^{m}(\tau) > 0$, in relation to the second best . That is, less private information purchase $\gamma^{sb}(\tau) > \gamma^{m}(\tau)$, and subaccumulation of public information . This subaccumulation may be very severe when C'(0) > 0 and the discount factor δ large.

Another relevant welfare benchmark is the *team efficient* solution. This corresponds to the solution where the planner can assign decision rules to agents as well as control information purchases. It is a first best solution with decentralized strategies. The team solution internalizes the dynamic information externalities.

Concentrate attention on linear and symmetric (and this is without loss of generality) decision rules. At period t the team manager has to choose parameters a_t , b_t and γ_t

so that each agent will buy a private signal with precision γ_t and then take a decision $q_{it} = a_t s_{it} + b_t \theta_{t-1}$ where, as before, θ_t is the summary public statistic with precision $\tau_t = \tau_\theta + \tau_u \sum_{k=0}^{t} a_k^2$. The parameter b_t has no intertemporal effect and, given the values γ_t and a_t, it is set $b_t = 1 - a_t$ to minimize the one period prediction loss. This yields a period loss of

$$L_t = \frac{a_t^2}{\gamma_t} + \frac{(1-a_t)^2}{\tau_{t-1}} + C(\gamma_t)$$

Notice that, given τ_{t-1} and γ_t, choosing $a_t < 0$ is dominated by $|a_t| > 0$ since τ_t remains unchanged and L_t is lower. Thus, we only need to consider $a_t \geq 0$. Furthermore, given a_t, γ_t has no intertemporal effect either, and therefore γ_t is chosen to minimize the one period loss, again. Therefore, at any interior solution γ_t and given the strict convexity of L_t in γ_t, the F.O.C. to minimize L_t has to hold:

$$\left(\frac{a_t}{\gamma_t}\right)^2 = C'(\gamma_t)$$

If $a_t = 0$, then $\gamma_t = 0$. We have therefore for a given a, a unique solution $\gamma^o(a)$ with $\gamma^o(0) = 0$ and it is strictly increasing. In summary, the team problem can be written as

$$\min_{\{a_t \geq 0\}} \sum_{t=0}^{\infty} \delta^t L_t \text{ where}$$

$$L_t = \frac{a_t^2}{\gamma^o(a_t)} + \frac{(1-a_t)^2}{\tau_{t-1}} + C(\gamma^o(a_t)) \text{ for } a_t > 0,$$

$$L_t = (\tau_{t-1})^{-1} \text{ for } a_t = 0$$

with $\tau_t = \tau_{t-1} + \tau_u a_t^2$, and $\tau_{-1} = \tau_\theta$.

The team's problem can be written equivalently as finding a value function $\Lambda^o(\cdot)$ such that

$$\Lambda^{o}(\tau)=\inf_{a\geq 0}\left[L(a,\tau)+\delta\Lambda^{o}\left(\tau+\tau_{u}a^{2}\right)\right], \text{ where } L(a,\tau)=\frac{a^{2}}{\gamma^{o}(a)}+\frac{(1-a)^{2}}{\tau}+C\left(\gamma^{o}(a)\right).$$

The characterization of the solution is similar to the second best solution but now the value function is strictly decreasing always. Indeed, this is so since $\partial L(a,\tau)/\partial\tau<0$. Indeed, the team manager controls both γ and a and therefore it does not have to worry about agents giving less weight to private information when public precision increases. A consequence is that at the team optimum it never pays to add noise to public information. More public information is always good once the dynamic information externalities are controlled for at the team solution. A consequence is that the information purchase is higher at the team solution for any τ than at the market solution. Denote by $a^{o}(\tau)$ the team solution, then $\gamma^{o}\left(a^{o}(\tau)\right)>\gamma^{m}(\tau)$ for any τ and the team solution accumulates more public precision. This can be heuristically checked as follows. Suppose that $\Lambda^{o}(\tau)$ is differentiable and $\partial\Lambda^{o}/\partial\tau<0$. The F.O.C. of the team control problem yields $\partial L/\partial a+2a\tau_{u}\delta\partial\Lambda^{o}/\partial\tau=0$ and therefore, since a > 0, at the team solution $\partial L/\partial a>0$. At the market solution ($\delta=0$) we have that $\partial L/\partial a=0$ and using the envelope condition for γ we obtain $\partial^{2}L/(\partial a)^{2}>0$ and therefore $a^{o}(\tau)>a^{m}(\tau)$ and consequently $\gamma^{o}\left(a^{o}(\tau)\right)>\gamma^{m}(\tau)$.

Moreover, using the parameterized example, it can be checked that the speed of learning and asymptotic properties are the same in the market, the second best benchmark, and the team solution.

Simulations of the model with the cost function $C(\tau_{\varepsilon})=c\lambda^{-1}\tau_{\varepsilon}^{\lambda}$, with $\lambda\geq 1$, c > 0 show the following:[26]

a. For any t, $\tau_{t}^{o}>\tau_{t}^{sb}>\tau_{t}^{m}$. Typically, the second best is much closer to the market than to the first best. (See Figure 6.3)

---

26 The simulations have been performed using the following approximation to the infinite horizon minimization problem:

$$\mathrm{Min}\left(\Sigma_{t=0}^{T}\ \delta^{t}L_{t}+\delta^{T+1}\left((1-\delta)\tau_{T}\right)^{-1}\right)$$

with the appropriate controls in each case, and T = 75 for $\delta\leq .9$, T = 100 for δ = .95, and T = 150 for δ = .98. The range of parameter values examined has been: δ in [.5, .98], τ_{u} in [.01, 20], c in [.01, 50], λ in [1, 50], fixing τ_{θ} = .5.

b. Underinvestment in public precision at the market solution, both with respect to the team and the second best, is increasing in the distance to the exogenous signals case (as parameterized by $1/\lambda$).

c. The relative welfare loss of the market solution, both with respect to the team and the second best, is increasing in the distance to the exogenous signals case $1/\lambda$ and in the discount factor δ but nonmonotonic in the precision of noise in the public signal τ_u. The relation between improvements in information transmission (increases in τ_u) and relative welfare losses is not monotonic. This should not be surprising since the information externality bites in an intermediate range of τ_u. If it is very low public information is very imprecise anyway and if it is very high public information is almost fully revealing. The relative welfare losses can be very large (up to 35% for the first best and to 12% for the second best, even for moderate values of δ, and λ -like, for example, .9 and 2, respectively).

d. The team solution may display a weight for the private precision a well above 1, implying a negative weight to public information. Furthermore, the team solution $a^o(\tau)$ is not always monotone in τ.

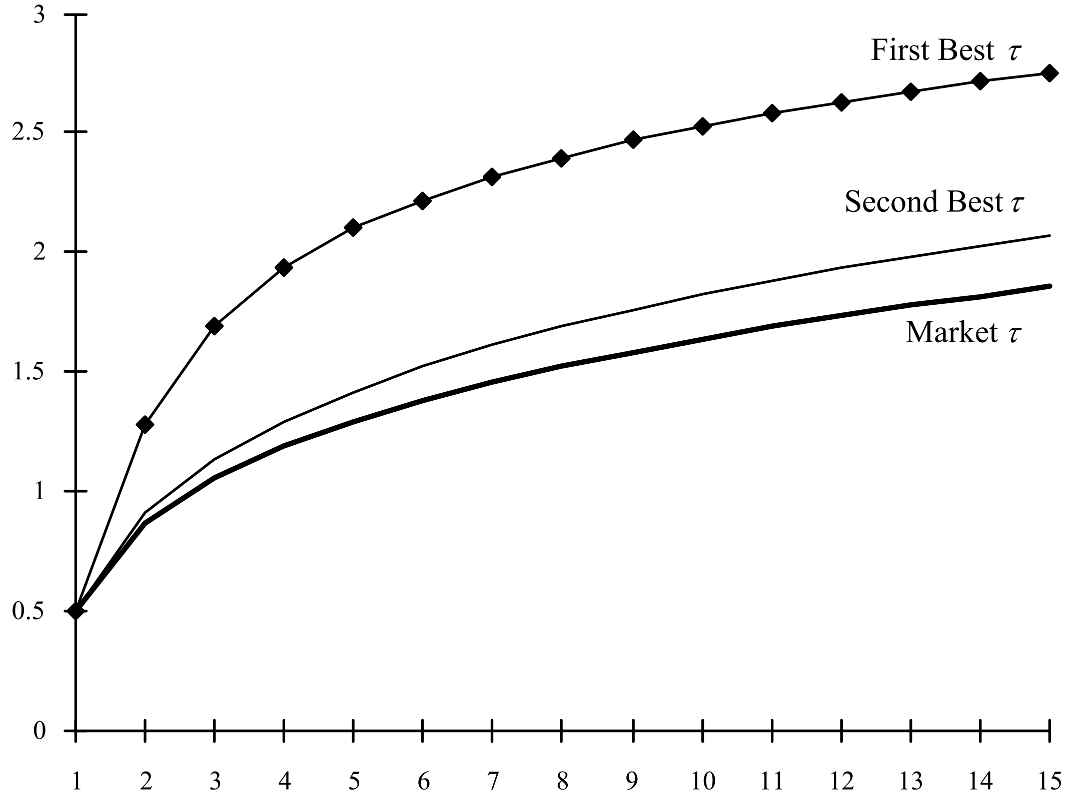

Figure 6.3. Evolution of the public precision τ_t in the three problems ($\tau_\theta = .5$, $\tau_u = c = 1$, $\delta = .9$, $\lambda = 3$).

6.5.4 Summary

Learning from others involves a basic information externality that induces herding and under-accumulation of public information. A welfare benchmark that internalizes the information exteranlities is needed to compare with the market solution. This is accomplished by the team efficient solution which internalizes the information externality and provides an appropriate benchmark to compare with the market solution. The welfare loss due to the information externality may be important but the (slow) rate of learning in the market is not suboptimal. The information externality and the associated welfare loss may be substantial and are aggravated with costly information acquisition. The strategic substitutability between private and public information means that public information may hurt welfare except in the team efficient solution.

6.6 Rational expectations, herding and information externalities

As stated in the introduction of this chapter and in Section 6.1, herding or, more in general, the insufficient reliance of agents on their private information has been put forward as an explanation for different phenomena like financial crisis, fashion, and technology adoption. The herding literature has put the finger on the welfare consequences of information externalities in a very stark statistical prediction model.

We have seen in Sections 6.3 to 6.5 how the root of inefficiency in herding models is an informational externality not taken into account by agents when making decisions. We consider here a static version of the model in Section 6.3 with a rational expectations flavor. In this section agents when making predictions can condition on the current public statistic (similarly to the strategies considered in Chapters 3 and 4). The rational expectations equilibrium in the prediction model can be compared then with the team efficient solution in which the informational externality is taken into account.

We compare the team with the market solution similarly to the analysis of REE in Section 3.4.[27] The analysis is basically a simplified version of Sections 3.2 and 3.4 where there is an information externality but not a payoff externality.

6.6.1 A model with a rational expectations flavor

The model is a static version of the model in Section 6.3 where the public information signal is given by $p = \int_0^1 q_i di + u$ with q_i being the prediction of agent i and u normally distributed noise, $u \sim N(0, \sigma_u^2)$. Agent i receives a private signal s_i about θ and solves the problem

$$\mathrm{Min}_q E\left[(\theta - q)^2 \,|\, I_i\right], \text{ where } I_i = \{s_i, p\}.$$

This information structure corresponds to a rational expectations solution. Indeed, think of agent i submitting a schedule contingent on the realizations of $p, q_i(s_i, \cdot)$, where p

27 This section is based on Bru and Vives (2002).

solves the equation $p=\int_0^1 q_i(s_i,p)\,di+u$. The forecast of agent i is contingent on his private signal and the (noisy) average forecast. We assume that all random variables are normally distributed with the same properties as in Section 6.3. In particular, $\theta \sim N(\bar{\theta},\sigma_\theta^2)$. The solution to the agent's problem is $q_i = E[\theta|I_i]$.

The minimization of the square loss function may arise from agents having quadratic utility functions. Suppose agent i has a utility function given by $U_i = (\theta+\eta_i)q_i - \frac{1}{2}q_i^2$ where η_i is an idiosyncratic random term (uncorrelated with everything else). For example, let q_i be the capacity decision of firm i, where $\theta+\eta_i$ indexes the marginal (random) value of capacity and let investment costs be quadratic. A firm decides about capacity based on its private information and the aggregate capacity choices in the industry, which includes some firms that invest for exogenous reasons, irrespective of information signals, with aggregate value u. Alternatively, competitive firms decide about investment with macroeconomic uncertainty represented by the random variable θ which determines average profitability. Firms invest taking into account that the profits of investment will depend on the realization of θ. To predict θ each firm has access to a private signal as well as to public information, the aggregate investment figures compiled by a government agency. Data on aggregate investment incorporates measurement error. The agent could be also a consumer facing a good of random quality as in the example in Section 6.4. Finally, agent i may be a retailer facing a random price $\theta+\eta_i$.

In any case the expected welfare loss with respect to the full-information first best (where θ is known and $q_i=\theta$) is easily seen to be $E\left[(\theta-q_i)^2\right]/2$.

Given the structure of the model Bayesian equilibria will necessarily be symmetric. Let a be the coefficient of s_i in the candidate linear equilibrium strategy of agent i. Then from the normality assumption and $p=\int_0^1 q_i\,di+u$, it follows that p will be a linear transformation of $z \equiv a\theta+u$ and that $E[\theta|p]=E[\theta|z]$. Let $\theta^* \equiv E[\theta|z]$. We can write the equilibrium strategy as

$$q_i(s_i, z) = E[\theta | s_i, z] = as_i + (1-a)\theta^*$$

where $a = \tau_\varepsilon / (\tau + \tau_\varepsilon)$ and $\tau \equiv (\mathrm{var}[\theta | p])^{-1} = \tau_\theta + \tau_u a^2$.[28]

It is not difficult to check that there is a unique linear Bayesian equilibrium (the market solution). The equilibrium strategy is given by $q_i = a^m s_i + (1-a^m)\theta^*$ where a^m is the unique positive real solution to the cubic equation $a = \tau_\varepsilon / (\tau_\varepsilon + \tau_\theta + \tau_u a^2)$. We have that $0 < a^m < \tau_\varepsilon / (\tau_\varepsilon + \tau_\theta)$ and $L^m \equiv E\left[(\theta - E[\theta | s_i, z])^2\right] = (\tau_\varepsilon + \tau^m)^{-1} = a^m / \tau_\varepsilon$ where $\tau^m = \tau_\theta + \tau_u (a^m)^2$. Furthermore, a^m and L^m decrease with τ_θ and τ_u; a^m increases and L^m decreases with τ_ε. In this context public information does not hurt: welfare increases with either better prior information τ_θ or a less noisy transmission channel τ_u. However, the result is not trivial since there are two effects. An increase in τ_θ or τ_u has a direct positive impact on τ^m but an indirect negative one on a^m which tends to reduce τ^m. The direct one prevails and public information reduces the prediction loss. The impact of an increase in private precision is positive on both accounts reducing directly the prediction loss and increasing a^m and τ^m.

6.6.2. The team efficient solution

Given decision rules $q_i(I_i)$ for the agents the average expected loss is $\int_0^1 E\left[(\theta - q_i(I_i))^2\right] di$. Restrict the planner to impose linear rules and let, without loss of generality, $q_i(I_i) = a_i s_i + c_i \theta^*$. Then it is optimal to set $c_i = 1 - a_i$ since otherwise public information would not be exploited efficiently (which in our context means that

---

28 It is immediate from normal theory that $E[\theta | s_i, z] = (\tau_\varepsilon s_i + \tau_u az + \tau_\theta \bar{\theta}) / (\tau_\varepsilon + \tau)$ and $\tau = \tau_\theta + \tau_u a^2$. Equivalently, $E[\theta | s_i, z] = as_i + (1-a)E[\theta | z]$ where $a = \tau_\varepsilon / (\tau + \tau_\varepsilon)$ because $E[\theta | z] = (\tau_u az + \tau_\theta \bar{\theta}) / \tau$.

$\text{cov}\left[\left(\theta - q_i\right), \theta^*\right] = 0$). As in Section 6.5, there cannot be any gain from using asymmetric rules. With a symmetric rule $q_i\left(I_i\right) = as_i + \left(1-a\right)\theta^*$, the expected loss is given by

$$L\left(a\right) = \frac{\left(1-a\right)^2}{\tau_\theta + \tau_u a^2} + \frac{a^2}{\tau_\varepsilon}.$$

It is easily seen that $L' = -2\tau^{-1}\left(1 - a\left(\tau_\varepsilon + \tau\right)\tau_\varepsilon^{-1} + \left(1-a\right)^2 \tau_u \tau^{-1}\right)$ from which it follows that $L'\left(1\right) > 0$ and $L'\left(a^m\right) < 0$ since $a^m = \tau_\varepsilon / \left(\tau_\varepsilon + \tau^m\right)$.

Denote by a^o the (unique) team solution[29] and let $L^o \equiv L\left(a^o\right)$. It follows that at the unique linear team solution $1 > a^o > a^m$ and a^o is increasing in τ_ε and decreasing in τ_θ. This follows immediately from $\partial^2 L\left(a^o\right)/\partial a \partial \tau_\theta > 0$ and $\partial^2 L/\partial a \partial \tau_\varepsilon < 0$. However, $\text{sign}\left(\partial^2 L/\partial a \partial \tau_u\right)$ is ambiguous. L^o is decreasing in τ_ε, τ_θ and τ_u. This follows directly from $\partial L/\partial \tau_u < 0$, $\partial L/\partial \tau_\varepsilon < 0$, and $\partial L/\partial \tau_\theta < 0$.

As expected, the weight to private information is too low and the weight to public information is too high at the market solution. The reason is that agents at the market solution do not internalize the positive effect on others of their response to private information because of the increase in the informativeness of the public statistic p. Here public (or private) information cannot hurt because the information externality is taken care of.

The proposition summarizes the results.

Proposition 6.4 (Bru and Vives (2002)). Market and team solutions.

(i) We have that $1 > a^o > a^m$.

29 L is strictly quasiconvex in a.

(ii) All coefficients are increasing in τ_ε and decreasing in τ_θ; a^m is also decreasing in τ_u.

(iii) L^o and L^m are decreasing in $\tau_\theta, \tau_\varepsilon$ and τ_u.

For extreme values of the parameters τ_ε and τ_u the information externality disappears. Indeed, as τ_ε tends to 0 (uninformative signals), a^o and a^m tend to 0, and as τ_ε tends to infinity (perfectly informative signals), a^o and a^m tend to 1. As τ_u tends to 0 (no public information), a^o and a^m tend to $\tau_\varepsilon / (\tau_\theta + \tau_\varepsilon)$, and as τ_u tends to infinity (perfect public information), a^o and a^m tend to 0.

Summary

The chapter has presented the basic models of learning from others developing the results obtained in the social learning and herding literature and confronting the apparent contradiction between those and the ones obtained by the dynamic rational expectations literature, providing a common frame for the analysis. The chapter has provided also an appropriate welfare benchmark with incomplete information: the team efficient solution introduced in Chapter 1. The most important conclusions are the following:

- The disparate results obtained in those two strands of the literature basically reflect the underlying assumption in the herding literature that does not allow an agent to fine-tune his action to his information.
- The basic information externality problem, the fact that an agent when taking an action today does not take into account the informational benefit that other agents will derive from his action, underlies the discrepancy between market and team-efficient solutions.
    - With discrete actions spaces (and signals of bounded strength) the inefficiency may take a very stark form with agents herding on the wrong action.
    - In more regular environments learning from others ends up revealing the uncertainty, although it will do so slowly if there is observational noise or friction.

- Learning from others with noisy observation by Bayesian agents has a self-correcting property. A higher inherited precision of public information will lead a Bayesian agent to put less weight in his private signal and consequently less of his information will be transmitted by the public statistic. This slows down learning.

- However, slow learning is team-efficient; that is, it is optimal as long as attention is restricted to decentralized strategies. This does not mean that the market solution is team-efficient. Because of the information externality, the market solution under-accumulates public precision and there is herding in the sense that agents put too little weight on their private signals with respect to the team solution that internalizes the externality.
- Endogenous and costly information acquisition accentuates the effect of the information externality, slowing down learning and increasing the relative welfare loss at the market solution.
- The strategic substitutability between public and private information is behind the possibility that more public information may decrease welfare in a world without payoff externalities. This canot happen at the team-efficient solution since then the information externality is internalized.

Appendix

Proof of (iii) in Proposition 6.1: $t^{-1/3}\tau_t \xrightarrow{t} \left(3\tau_u\tau_\varepsilon^2\right)^{1/3}$. (It is an adaptation of the proof provided by Chamley (2004a) for the case $\varsigma = 1$.) From the analysis in Section 6.3.1 we have that $\tau_{t+1} = \tau_t + \left(\varsigma\tau_\varepsilon^{-1} + \left(a_{t+1}^2\tau_u\right)^{-1}\right)^{-1} = \tau_t + \left(\varsigma\tau_\varepsilon^{-1} + \tau_u^{-1}\left(1+\tau_\varepsilon^{-1}\tau_t\right)^2\right)^{-1}$ since $a_{t+1} = \tau_\varepsilon/\left(\tau_\varepsilon + \tau_t\right)$. This defines the policy function $\tau_{t+1} = g(\tau_t) > 0$, $\tau_t > 0$. Let $\nu_t = 1/\tau_t$ and consider the function $\varphi : \mathbb{R}_+ \to \mathbb{R}$, $\nu_{t+1} = \varphi(\nu_t) \equiv 1/g(1/\nu_t)$. It is easily seen that

$$\varphi(\nu_t) = \nu_t - \frac{\nu_t^4}{\nu_t^3 + \left(\varsigma\sigma_\varepsilon^2 + \sigma_u^2\right)\nu_t^2 + 2\sigma_u^2\sigma_\varepsilon^2\nu_t + \sigma_u^2\sigma_\varepsilon^4}.$$

It follows that $\varphi(0) = 0$, and for $\nu > 0$, $0 < \varphi(\nu) < \nu$. Furthermore, $\varphi'(\cdot) > 0$ for ν small. We know that as $t \to \infty$, then $\tau_t \to \infty$ and $\nu_t \to 0$. We obtain now the convergence rate of $\nu_t \to 0$. Given that $\nu_t \to 0$ and $\tau_u\tau_\varepsilon^2 = \left(\sigma_u^2\sigma_\varepsilon^4\right)^{-1}$ we can write $\nu_{t+1} = \nu_t - \nu_t^4\tau_u\tau_\varepsilon^2\left(1 + O(\nu_t)\right)$ where $O(\nu_t)$ is a term of order of at most ν_t (see Section 3.2 in the Technical Appendix).

Define the sequence $\{b_t\}$ by $b_t = t^{1/3}\nu_t$. We obtain

$$\frac{b_{t+1}}{(t+1)^{1/3}} = \frac{1}{t^{1/3}}\left(b_t - \frac{b_t^4}{t}\tau_u\tau_\varepsilon^2\left(1 + O\left(\frac{b_t}{t^{1/3}}\right)\right)\right)$$

or

$$b_{t+1}\left(\frac{t+1}{t}\right)^{-1/3} = b_t - \frac{b_t^4}{t}\tau_u\tau_\varepsilon^2\left(1 + O\left(\frac{b_t}{t^{1/3}}\right)\right).$$

Therefore,

$$b_{t+1}\left(1 - \frac{1}{3t} + O\left(\frac{1}{t^2}\right)\right) = b_t - \frac{b_t^4}{t}\tau_u\tau_\varepsilon^2\left(1 + O\left(\frac{b_t}{t^{1/3}}\right)\right).$$

From the latter equation it can be checked that the sequence $\{b_t\}$ must be bounded. We show now that it converges to $\left(3\tau_u\tau_\varepsilon^2\right)^{-1/3}$. Extract first a subsequence of $\{b_t\}$ which converges to some limit $\bar{b}$. Extract then from this subsequence another subsequence such that b_{t+1} converges to a limit $\hat{b}$. It must hold then that

$$\hat{b}\left(1-\frac{1}{3t}+O\left(\frac{1}{t^2}\right)\right)=\bar{b}-\frac{\bar{b}^4}{t}\tau_u\tau_\varepsilon^2\left(1+O\left(\frac{b_t}{t^{1/3}}\right)\right).$$

Taking the limit as t goes to infinity we have that $\bar{b}=\hat{b}=b$. Equating the terms of the order 1/t and disregarding those of smaller order we obtain

$$\frac{b}{3}=b^4\tau_u\tau_\varepsilon^2$$

and therefore

$$b=\left(3\tau_u\tau_\varepsilon^2\right)^{-1/3}.$$

Exercises

6.1 The 2 x 2 x 2 herding model. Consider the $2\times 2\times 2$ model of Section 6.1 and assume that an agent does not adopt when the net value of adoption is 0. Let the public belief be $\theta_i=\Pr\left(\theta=1\middle|x^i\right)$. Use Bayes' rule to show that $\Pr\left(\theta=1\middle|x^i,s_L\right)<\theta_i<\Pr\left(\theta=1\middle|x^i,s_H\right)$. Show that:

(i) If $\theta_i\le 1-\ell$ agent i rejects no matter the value of his signal.

(ii) If $\theta_i>\ell$ agent i adopts no matter the value of his signal.

(iii) If $1-\ell<\theta_i\le\ell$ then agent i adopts if and only if his signal is high ($s=s_H$).

(iv) Provide an argument to show that the probability that an informational cascade has not started when is the turn of i to move converges to zero exponentially as i increases.

Solution: (i), (ii) and (iii) are immediate. Hint for (iv): a necessary condition to avoid a cascade is that signals alternate between high and low values.

6.2 Endogenous private information with linear acquisition costs. Let $C(\gamma)=c\gamma$ with $c>0$. Show that for a given precision of the inherited public information, τ, agents buy signals with precision

$$\gamma=\begin{cases}c^{-1/2}-\tau & \text{if } c^{-1/2}>\tau\\ 0 & \text{otherwise}\end{cases}$$

and therefore,

$$g^m(\tau) = \begin{cases} \tau + \tau_u(1-c^{1/2}\tau)^2 & \text{if } c^{-1/2} > \tau \\ \tau & \text{otherwise} \end{cases}$$

Show that if $\tau_\theta \le \frac{\tau_u c^{1/2} - 1}{\tau_u c}$ then $\gamma = 0$ and $\tau = \tau_\theta + \tau_u(1-c^{1/2}\tau_\theta)^2 \ge c^{-1/2}$ for all $t \ge 1$. Otherwise, $\gamma_t > 0$ and $\tau_t < c^{-1/2}$ for all t, whereas $\tau_t \xrightarrow{t} c^{-1/2}$. Interpret the results.

Answer: Follow the steps in Section 6.3.2 for the general case. If τ_θ is low enough and $c^{1/2} > 1/\tau_u$, only the first generation buys private signals, whereas later on, agents simply rely on past, public information when taking their decisions. This means that the precision of the information on which these decisions are taken will remain constant over time. Otherwise, first generation agents buy less precise signals, and they continue buying private precision at every point in time, although at a decreasing rate. The result is an increasing precision of the public information, which will converge to $c^{-1/2}$.

6.3 <u>The rate of learning when information is endogenous</u>. Consider the model in Section 6.3.2 and let $C(\gamma) = c\lambda^{-1}\gamma^\lambda$ with $\lambda > 1$. Show that $\tau_t \approx t^r$ with $r = (\lambda-1)/(3\lambda+1)$ and $a_t \approx t^{-\upsilon}$ with $\upsilon = (\lambda+1)/(3\lambda+1)$. Note that $\upsilon > 1/3 > r > 0$ and that r is increasing in λ.

Answer: From the F.O.C. $\gamma^* - \left[C'(\gamma^*)\right]^{-1/2} = -\tau$ it is immediate that

$\frac{1}{c} = \tau_{t-1}^{1+\lambda} \frac{a_t^{\lambda-1}}{(1-a_t)^{1+\lambda}}$ and therefore $\tau_{t-1}^{1+\lambda} a_t^{\lambda-1} \to 1/c$ as t tends to infinity since then a_t tends to zero. Let $a_t \approx t^{-\upsilon}$ for some $\upsilon > 0$, then $\tau_t \approx \sum_{k=0}^t a_k^2 \approx \sum_{k=0}^t k^{-2\upsilon} \approx t^{1-2\upsilon}$. Therefore $t^{(1-2\upsilon)(1+\lambda)} t^{-\upsilon(\lambda-1)}$ is of the order of a constant and consequently $(1-2\upsilon)(1+\lambda) - \upsilon(\lambda-1) = 0$. It follows that $\upsilon = (1+\lambda)/(3\lambda+1)$ and $r = 1 - 2\upsilon = (\lambda-1)/(3\lambda+1)$.

6.4 <u>Optimal learning with asymmetric rules</u>? Show that asymmetric rules cannot improve upon symmetric ones in the optimal learning program of Section 6.5.2.

Solution: First of all, using asymmetric weights a_{it} for the signals of different agents s_i it is still optimal to set the weights to public information so that $c_{it} = 1 - a_{it}$ (the reason is as before that the coefficients c_{it} do not have any intertemporal effect). This means that

$$L_{it} = E\left[\left(\theta - q_{it}\left(I_{it}\right)\right)^2\right] = \frac{\left(1-a_{it}\right)^2}{\tau_{t-1}} + \frac{a_{it}^2}{\tau_\varepsilon}.$$

Let $L(a) = \frac{a^2}{\tau_\varepsilon} + \frac{(1-a)^2}{\tau}$. Given that $L(a)$ is convex in a, $\int_0^1 L(a_i)\, di \geq L\left(\int_0^1 a_i di\right)$ and therefore there can not be any static gain from asymmetric rules. Furthermore, there can not be any dynamic gain either since if in period t different weights a_{it} were given to the signals of different agents, the signal-to-noise ratio in the new information of the public statistic p_t would depend only on $\int_0^1 a_i di$.

6.5 The market solution is incentive compatible. Show that the market solution in Section 6, $q_i = E\left[\theta \middle| s_i, \theta^*\right]$, makes privately efficient use of private information (i.e. $E\left[\theta - q_i\right] = 0$ and $E\left[\left(\theta - q_i\right)s_i\right] = 0$) and efficient use of public information (i.e. $E\left[\left(\theta - q_i\right)\theta^*\right] = 0$). In contrast the team optimum is bound only by the second restriction. Depict the team and the market solutions in the space of weights to private (a) and public (c) information.

Answer: (i) From $E\left[\theta - q_i \middle| s_i, \theta^*\right] = 0$ for all s_i and θ^*, and using the projection theorem for normal random variables (according to which $\text{cov}\left[\left(\theta - q_i\right) - E\left[\theta - q_i \middle| s_i, \theta^*\right], E\left[\theta - q_i \middle| s_i, \theta^*\right]\right] = 0$) we obtain $\text{cov}\left[\theta - q_i, s_i\right] = \text{cov}\left[\theta - q_i, \theta^*\right] = 0$ (equivalently, $E\left[\left(\theta - q_i\right)s_i\right] = E\left[\left(\theta - q_i\right)\theta^*\right] = 0$). (See Bru and Vives (2002).)

6.6 "Business as usual" and "wisdom after the fact". Consider the model of Section 6.3.1 where both the state of the world $\theta \in \{0,1\}$ and actions $x_{it} \in \{0,1\}$ are discrete. The payoff to take the action $x = 1$ (investing) is $\theta - c$ for some small c > 0 and the payoff of not investing ($x = 0$) is 0. The rest is as in Section 6.3.1. Aggregate investment in period t is thus $\int_0^1 x_{it} di$ and the public statistic is $p^{t-1} = \{p_0, ..., p_{t-1}\}$.

(i) Show that an agent will choose to invest if and only if his private signal is larger than $\frac{1}{2}-\sigma_\varepsilon^2\lambda_t$, where λ_t is the likelihood ratio between states $\theta = 1$ and $\theta = 0$ based on public information p^{t-1}. We have then that $p_t = 1-\Phi\left(\frac{\frac{1}{2}-\sigma_\varepsilon^2\lambda_t-\theta}{\sigma_\varepsilon}\right)+u_t$ where Φ is the cumulative distribution of the standard normal.

(ii) Argue that when the likelihood ratio λ_t is large or small (in algebraic terms) then the public statistic does not provide a very accurate signal of θ. This means that when an extremal amount of agents invest (i.e., most of them or very few) the informativeness of p_t is low.

(iii) Argue how can (ii) can explain that after a period of "business as usual" (say with many agents investing), a crash occurs (as agents learn that the true state is bad), and then this belief seems evident ("wisdom after the fact").

Solution: The exercise is a streamlined version of Caplin and Leahy (1994)). For (i) and (ii) see Section 1 in the Technical Appendix if you have trouble with likelihood ratios; for (iii) imagine what would happen if the true state is bad but that due to some realizations of noise in the past the public belief is positive about the state of the world.

6.7 Overconfidence. Consider the learning model in Section 6.3.1. How would you model the case in which agents are overconfident with respect to the precision of their information? What are the implications of overconfidence for aggregation of information?

Hint: See Nöth and Weber (2003) for inspiration and experimental results.

References

Anderson, L. and Holt, C. (1997), "Information Cascades in the Laboratory", *American Economic Review*, 87, 5, 847-862.

Arrow, K.J. (1962), "The Economic Implications of Learning by Doing", *Review of Economic Studies*, 155-173.

Amador, M. and P-O. Weill (2007), "Learning from Private and Public Observations of Others' Actions", mimeo.

Aumann, R. (1976), "Agreeing to Disagree", *Annals of Statistics*, 4, 1236-1239.

Avery, C. and J. Chevalier (1999), "Herding Over the Career", *Economics Letters*, 63, 327-333.

Avery, C. and P. Zemsky (1998), "Multi-Dimensional Uncertainty and Herd Behavior in Financial Markets", *American Economic Review*, 88, 4, 724-748.

Baloff, N. (1966), "The Learning Curve. Some Controversial Issues", *Journal of Industrial Economics*, 275-283.

Banerjee, A. (1992), "A Simple Model of Herd Behavior", *Quarterly Journal of Economics*, 107, 3, 797-817.

Banerjee, A. (1993), "The Economics of Rumors", *Review of Economic Studies*, 60, 309-327.

Banerjee, A. and D. Fudenberg (2004), "Word-of-mouth Learning", *Games and Economic Behavior*, 46, 1, 1-22.

Bernhardt, D., M. Campbell and E. Kutsoati (2006), "Who Herds?", *Journal of Financial Economics*, 80, 3, 657-675.

Bikhchandani, S., D. Hirshleifer and I. Welch (1992), "A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades", *Journal of Political Economy*, 100, 51, 992-1026.

Bikhchandani, S., D. Hirshleifer and I. Welch (1998), "Learning form the Behavior of Others", *Journal of Economic Perspectives*, 12, 3, 151-170.

Bru, L. and X. Vives (2002), "Informational Externalities, Herding, and Incentives," *Journal of Institutional and Theoretical Economics*, 158, 1, 91-105.

Burguet, R. and X. Vives (2000), "Social Learning and Costly Information Acquisition", *Economic Theory*, 15, 1,185-205.

Callander, S. and J. Hörner (2006), "The Wisdom of the Minority", WP Northwestern University.

Caminal, R. and X. Vives (1996), "Why Market Share Matters: An Information-based Theory", *RAND Journal of Economics*, 27, 2, 221-239.

Caminal, R. and X. Vives (1999), "Price Dynamics and Consumer Learning", *Journal of Economics and Management Strategy*, 8, 1, 95-131.

Caplin, A. and J. Leahy (1994), "Business as Usual, Market Crashes, and Wisdom after the Fact", *American Economic Review*, 84, 3, 548-65.

Çelen, B. and S. Kariv (2004a), "Distinguishing Informational Cascades from Herd Behavior in the Laboratory", *American Economic Review*, 94, 3, 484-498.

Çelen, B. and S. Kariv (2004b), "Observational Learning under Imperfect Information", *Games and Economic Behavior*, 47, 72-86.

Chamley, C. and D. Gale (1994), "Information Revelation and Strategic Delay in a Model of Investment", *Econometrica*, 62, 5, 1065-85.

Chamley, C. (2003), "Cascades and Slow Social Learning", Working Paper Boston University.

Chamley, C. (2004a), *Rational Herds*, Cambridge: Cambridge University Press.

Chamley, C. (2004b), "Delays and Equilibria with Large and Small Information in Social Learning", *European Economic Review*, 48, 477-501.

Chen, Q. and W. Jiang (2006), "Analysts' Weighting of Private and Public Information", *Review of Financial Studies*, 19, 1, 319-355.

Chevalier, J. and G. Ellison (1999), "Career Concerns of Mutual Fund Managers", *Quarterly Journal of Economics*, 114, 2, 389-432.

Cipriani, M., D. Gale, and A. Guarino (2005), "Estimating Herd Behavior in Financial Markets: A Structural Approach", mimeo.

Crawford, V. and J. Sobel (1982), "Strategic Information Transmission", *Econometrica*, 50, 6, 1431-1451.

De Long, J. B., A. Shleifer, L. H. Summers and R. J. Waldmann (1990), "Noise Trader Risk in Financial Markets", *Journal of Political Economy*, 98, 703-738.

Effinger, M. and M. Polborn (2001), "Herding and Anti-Herding: A Model of Reputational Differentiation", *European Economic Review*, 45, 385-403.

Ellison, G. and D. Fudenberg (1993), "Rules of Thumb for Social Learning", *Journal of Political Economy*, 101, 4, 612-643.

Ellison, G. and D. Fudenberg (1995), "Word-mouth Communication and Social Learning", *The Quarterly Journal of Economics*, 110, 1, 93-125.

Feldman, M. (1987), "An Example of Convergence to Rational Expectations with Heterogeneous Beliefs", *International Economic Review*, 28, 635-650.

Fellner, W. (1969), "Specific Interpretations of Learning by Doing", *Journal of Economic Theory*, 1, 2, 119-140.

Foster, A. and M. Rosenzweig (1995), "Learning by Doing and Learning from Others: Human Capital and Technical Change in Agriculture", *Journal of Political Economy*, 103, 6, 1176-1209.

Fudenberg, D. and J. Tirole (1991), *Game Theory*, Cambridge, Mass: MIT Press.

Gale, D. (1996), "What Have We Learned from Social Learning?", *European Economic Review*, 40, 617-628.

Gale, D. and S. Kariv (2003), "Bayesian Learning in Social Networks", *Games and Economic Behavior*, 45, 329-346.

Geanakoplos, J. and H. Polemarchakis (1982), "We Can't Disagree Forever", *Journal of Economic Theory*, 28, 192-200.

Glosten, L. and P. R. Milgrom (1985), "Bid, Ask and Transaction Prices in a Specialist Market with Heterogeneously Informed Traders", *Journal of Financial Economics*, 17, 71-100.

Goeree, J., T. Palfrey, B. Rogers and R. McKelvey (2007), "Self-Correcting Information Cascades", *Review of Economic Studies*, 74, 733-762.

Graham, J. (1999), "Herding among Investment Newsletters: Theory and Evidence", *Journal of Finance*, 54, 1, 237-268.

Grinblatt, M., S. Titman and R. Wermers (1995), "Momentum Investment Strategies, Portfolio Performance, and Herding: A Study of Mutual Fund Behavior", *American Economic Review*, 85, 5, 1088-1105.

Guarino, A., H. Harmgart and S. Huck (2007), "When Half the Truth is Better than the Truth: Theory and Experiments on Aggregate Information Cascades", Working paper UCL.

Gul, F. and R. Lundholm (1995), "Endogenous Timing and the Clustering of Agents' Decisions", *Journal of Political Economy*, 103, 5, 1039-66.

Hendricks, D., J. Patel and R. Zeckhauser (1993), "Hot Hands in Mutual Funds: Short-run Persistence of Performance", *Journal of Finance*, 48, 93-130.

Hong, H., J. Kubik and A. Solomon (2000), "Security Analysts' Career Concerns and Herding of Earnings Forecasts", *RAND Journal of Economics*, 31, 1, 121-144.

Huck, S. and J. Oechssler (2000), "Informational Cascades in the Laboratory: Do they Occur for the Right Reasons?", *Journal of Economic Psychology*, 21, 661-671.

Hung, A. and C.R. Plott (2001), "Information Cascades: Replication and an Extension to Majority Rule and Conformity-Rewarding Institutions", *American Economic Review*, 91, 5, 1508-1520.

Jun, B. and X. Vives (1996), "Learning and Convergence to a Full-Information Equilibrium are not Equivalent", *Review of Economic Studies*, 63, 653-674.

Kihlstrom, R. (1984), "A Simple Example of the Radner-Stiglitz Nonconcavity in the Value of Information", *Bayesian Models in Economic Theory*, pp. 53-61, Studies in Bayesian Econometrics, vol. 5. New York; Amsterdam and Oxford: North-Holland.

Kübler, D. and G. Weizsäcker (2004), "Limited Depth of Reasoning and Failure of Cascade Formation in the Laboratory", *Review of Economic Studies*, 71, 425-441.

Lamont, O. (2002), "Macroeconomic Forecasts and Microeconomic Forecasters", *Journal of Economic Behavior & Organization*, 48, 3, 265-280.

Larson, N. (2006), "In with the New or Out with the Old? Bottlenecks in Social Learning", Working paper University of Virginia.

Lee, I. (1993), "On the Convergence of Informational Cascades", *Journal of Economic Theory*, 61, 395-411.

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

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

Morris, S. and H. Shin (2002), "The Social Value of Public Information", *American Economic Review*, 92, 1521-1534.

Morris, S. and H. Shin (2005), "Central Bank Transparency and the Signal Value of Prices", *Brookings Papers on Economic Activity*, 2 1-66.

Nöth, M. and Weber, M. (2003), "Information Aggregation with Random Ordering: Cascades and Overconfidence", *Economic Journal*, 113, 166-189.

Ottaviani, M. and P. Sørensen (2000), "Herd Behavior and Investment: Comment", *American Economic Review,* 90, 3, 695-704.

Prendergast, C. and L. Stole (1996), “Impetuous Youngsters and Jaded Old-Timers: Acquiring a Reputation for Learning”, *Journal of Political Economy*, 104, 6, 1105-1134.

Radner, R. (1962), “Team Decision Problems”, *Annals of Mathematical Statistics*, 33, 857-881.

Radner, R. and J. E. Stiglitz (1984), “A Nonconcavity in the Value of Information,” *Bayesian Models in Economic Theory*, 1984, 33-52, Studies in Bayesian Econometrics, vol. 5. New York; Amsterdam and Oxford: North-Holland.

Rodríguez-Mora, J.V. and P. Schulstad (2007), "The Effect of GNP Announcements on Fluctuations of GNP Growth", *European Economic Review*, 51, 8, 1922-1940.

Romer, D. (1993), “Rational Asset-Price Movements Without News”, *American Economic Review*, 83, 1112-1130.

Scharfstein, D. and J. Stein (1990), “Herd Behavior and Investment”, *American Economic Review*, 83, 3, 465-479.

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

Shiller, J. (1981), “Do Stock Prices Move too much to be Justified by Subsequent Changes in Dividends?”, *American Economic Review*, 71, 421-436.

Shiller, J. (1984), “Stock Prices and Social Dynamics”, *Brookings Papers of Economic Activity*, 2, 457-498.

Shiller, J. (1989), *Market Volatility*, Cambridge, MA: MIT Press.

Smallwood, D. and J. Conlisk (1979), “Product Quality in Markets Where Consumers are Imperfectly Informed”, *Quarterly Journal of Economics*, 93, 1, 1-23.

Smith, L. and P. Sørensen (1995), “Rational Social Learning by Random Sampling”, mimeo.

Smith, L. and P. Sørensen (2000), “Pathological Outcomes of Observational Learning”, *Econometrica*, 68, 2, 371-398.

Townsend, R. M. (1978), “Market Anticipations, Rational Expectations, and Bayesian Analysis”, *International Economic Review*, 19, 481-494.

Trueman, B. (1994), “Analyst Forecasts and Herding Behavior”, *Review of Financial Studies*, 7, 97-124.

Vives, X. (1993), “How Fast Do Rational Agents Learn?”, *Review of Economic Studies*, 60, 2, 329-347.

Vives, X. (1997), “Learning from Others: A Welfare Analysis”, *Games and Economic Behavior*, 20, 177-200.

Vives, X. (1999), *Oligopoly Pricing: Old Ideas and New Tools*, Cambridge: MIT Press.

Vives, X. (2005), “Complementarities and Games: New Developments”, *Journal of Economic Literature*, 2005, 43, 437-479

Welch, I. (2000), “Herding Among Security Analysts”, *Journal of Financial Economics*, 58, 369-396.

Wermers, R. (1999), “Mutual Fund Herding and the Impact on Stock Prices”, *Journal of Finance*, 2, 581-622.

Zhang, J. (1997), “Strategic Delay and the Onset of Investment Cascades”, *RAND Journal of Economics*, 28, 188-205.

Zitzewitz, E. (2001), “Measuring Herding and Exaggeration by Equity Analysts and Other Opinion Sellers”, Stanford GSB Working Paper No. 1802.

Zwiebel, J. (1995), “Corporate Conservatism and Relative Compensation”, *Journal of Political Economy*, 103, 1-25.