Chapter 1

Bayesian Inference

(09/17/17)

A witness with no historical knowledge
There is a town where cabs come in two colors, yellow and red.[1] Ninety percent of the cabs are yellow. One night, a taxi hits a pedestrian and leaves the scene without stopping. The skills and the ethics of the driver do not depend on the color of the cab. An out-of-town witness claims that the color of the taxi was red. The out-of town witness does not know the proportion of yellow and red cabs in the town and makes a report on the sole basis of what he thinks he has seen. Since the accident occurred during the night, the witness is not completely reliable but it has been assessed that such a witness makes a correct statement is four out of five (whether the true color of the cab is yellow or red). How should one use the information of the witness? Because of the uncertainty, we should formulate our conclusion in terms of probabilities. Is it more likely then that a red cab was involved in the accident? Although the witness reports red and is correct 80 percent of the time, the answer is no.

Recall that there are many more yellow cabs. The red sighting can be explained either by a yellow cab hitting the pedestrian (an event with high *prior* probability) which is incorrectly identified (an event with low probability), or a red cab (with low probability) which is correctly identified (with high probability). Both the *prior* probability of the event and the precision of the signal have to be used in the evaluation of the signal. Bayes' rule

[1]The example is adapted from Salop (1987)

provides the method to compute probability updates. Let $\mathcal{R}$ be the event "a red cab is involved", and $\mathcal{Y}$ the event "a yellow cab is involved". Likewise, let r (y) be the report "I have seen a red (yellow) cab". The probability of the event $\mathcal{R}$ conditional on the report r is denoted by $P(\mathcal{R}|r)$. By Bayes' rule,[2]

$$P(\mathcal{R}|r) = \frac{P(r|\mathcal{R})P(\mathcal{R})}{P(r)} = \frac{P(r|\mathcal{R})P(\mathcal{R})}{P(r|R)P(\mathcal{R}) + P(r|\mathcal{Y})(1 - P(\mathcal{R}))}. \tag{1.1}$$

The probability that a red cab is involved before hearing the testimony is $P(\mathcal{R}) = 0.10$. $P(r|\mathcal{R})$ is the probability of a correct identification and is equal to 0.8. $P(r|\mathcal{Y})$ is the probability of an incorrect identification and is equal to 0.2. Hence,

$$P(\mathcal{R}|r) = \frac{0.8 \times 0.1}{0.8 \times 0.1 + 0.2 \times 0.9} = \frac{4}{13} < \frac{1}{2}.$$

Note that this probability is much less than the precision of the witness, 80 percent, because a "red" observation is more likely to come from a wrong identification of a yellow cab than from a correct identification of a red cab.

The example reminds us of the difficulties that some people may have in practical circumstances. Despite these difficulties,[3] all rational agents in this book are assumed to be Bayesians. The book will concentrate only on the difficulties of learning from others by rational agents.

A witness with historical knowledge

Suppose now that the witness is a resident of the town who knows that only 10 percent of the cabs are red. In making his report, he tells the color which is the most likely according to his rational deduction. If he applies the Bayesian rule and knows his probability of making a mistake, he knows that a yellow cab is more likely to be involved. He will report "yellow" even if he thinks that he has seen a red cab. If he thinks he has seen a yellow one, he will also say "yellow". His private information (the color he thinks he has seen) is ignored in his report.

The omission of the witness' information in his report does not matter if he is the only witness and if the recipient of the report attempts to assess the most likely event: the witness and the recipient of the report come to the same conclusion. But suppose there is a second witness with the same sighting skill (correct 80 percent of the time) and who also thinks he has seen a red cab. That witness who attempts to report the most likely event

[2] Using the definition of conditional probabilities, $P(\mathcal{R}|r)P(r) = P(\mathcal{R} \text{ and } r) = P(r|\mathcal{R})P(\mathcal{R})$.

[3] The ability of people to use Bayes' rule has been tested in experiments, with mixed results (Holt and Anderson, 1993).

says also “yellow”. The recipient of the two reports learns nothing from the reports. For him the accident was caused by a yellow cab with a probability of 90 percent.

Recall that when the first witness came from out-of-town, he was not informed about the local history and he gave an informative report, “red”. That report may be inaccurate, but it provides information. Furthermore, it triggers more information from the second witness. After the report of the first witness, the probability of $\mathcal{R}$ increased from 0.1 to 4/13. When that probability of 4/13 is conveyed to the second witness, he thinks that a red car is more likely.[4] He therefore reports “red”. The probability of the inspector who hears the reports of the two witnesses is now raised to the level of the last (second) witness.

Looking for your phone as a Bayesian

You live in a two room apartment with two rooms, one that you keep orderly, one that is messy. After stepped out with a friend, you realize that you have left your cell phone behind. The phone is equally likely to be in one of the two rooms. You tell your friend: please looking for my phone that I have left in the apartment while I fetch the car that is parked in the next block. Your friend comes back without having found the phone. Which room is the more probable for the phone. Answer before reading the next paragraph.

You may think that your friend has looked into the two rooms. In the orderly room, it is harder to miss the phone. Therefore, no seeing the phone in that room makes it unlikely (compared to the other room) that the phone is there. You increase the probability of the messy room. You are a Bayesian.

In the formalization of this story, we can that there are two rooms 1 (orderly) and 2 (messy). There are two states of the nature: the phone is in room 1 or room 2. A search in room i, $i = 1$ or 2 produces a signal that is 1 (finding the phone) or 0 (not finding the phone. Each signal has a probability q_i to be equal to 1 if the phone is in room i. The probability of not finding the phone in room i when the phone is actually in room i is $1 - q_i$ is positive. If the phone is in room $3 - i$, (the room other than i), the signal s_i is zero. When you do not find the phone in Room 1, you think, rationally, you increase your probability that the phone is in Room 2. If you search in Room 2 for about the same time, then you think that the probability of a mistaken signal $s_2 = 0$ is higher than $s_1 = 0$ if the phone is in Room 1. Comparing the two rooms, you increase the probability of the phone in Room 2. The precise Bayesian calculus will be done later in this chapter.

[4] Exercise: prove it.

1.1 The standard Bayesian model

1.1.1 General remarks

The main issue is to learn about *something*. In the Bayesian framework, the "something" is a possible fact, which can be called a *state of nature*. That fact may take place in the future or it may already have taken place with an uncertain knowledge about it. Actually, in a Bayesian framework, there is no difference between a future event and a past event that are both uncertain. The future event may be "rain" or "shine", to occur tomorrow. For a Bayesian, nature chooses the weather today (with some probability, to be described below), and that weather is *realized* tomorrow.

The list of possible states is fixed in Bayesian learning. There is no room for learning about states that are not on the list of possible states before the learning process. That is an important limitation of Bayesian learning. There is no "unknown unknown", to use the famous characterization of secretary of state Rumsfeld, only "known unknown". In other words, one knows what is unknown.

The Bayesian process begins by putting weights on the unknowns, probabilities on the possible states of nature. These probabilities may be objective, such as the probability of "tail" or "face" in throwing a coin, but that is not important. What matters is that these probabilities are the ones that the learner uses at the learning process. These probabilities will be called *belief*. A "belief" will be a distribution of probabilities over the possible states. By an abuse of language, a belief will sometimes be the probability of a particular state, especially in the case of two possible states: the "belief" in one state will obviously define the probability of the other state. The belief before the reception of information is called the *prior belief*.

Learning is the processing of information that comes about the state. This information comes in the form of a *signal*. Examples are the witness report of the previous section, a weather forecast, an advice by a financial advisor, the action of some "other" individual, etc... In order to be informative, that signal must depend on the state. But that signal is imperfect and does not reveal exactly the state (otherwise there would be nothing interesting to think about). A natural definition of a signal is therefore a random variable that can take different values with some probabilities and the distribution of these probabilities depend on the actual state. The processing of the information of the signal is the use of the signal to update the prior belief into the posterior belief. That step is the core of the Bayesian learning process and its mechanics are driven by Bayes' rule. In that process, the learner knows the mechanics of the signal, *i.e.*, the probability of receiving a particular signal value conditional on the true state. Bayes' rule combines that knowledge with the prior distribution of the state to compute the posterior distribution.

Examples

1. The binary model
   - States of nature $\theta \in \Theta = \{0, 1\}$
   - Signal $s \in \{0, 1\}$ with $P(s = \theta|\theta) = q_\theta$.
2. Financial advising (*i.e.*, Value Line):
   - States of nature: a stock will go up 10% or go down 10% (two states).,
   - Advice {Strong Sell, Sell, Hold, Buy, Strong Buy}.
3. Gaussian signal:
   - Two states of nature $\theta \in \Theta = \{0, 1\}$
   - Signal $s = \theta + \epsilon$, where s has a normal distribution with mean zero and variance σ^2.
4. Gaussian model:
   - The state θ has a normal distribution with mean $\bar{\theta}$ and variance $sigma^2_\theta$.
   - Signal $s = \theta + \epsilon$, where s has a normal distribution with mean zero and variance σ^2_ϵ.

Note how in all cases, the (probability) distribution of the signal depends on the state. These are just examples and we will see later how each of them is a useful tool to address specific issues. We begin with the simplest model, the binary model.

1.1.2 The binary model

In all models of rational learning that are considered here, there is a *state of nature* (or just "state") that is an element of a set. We will use the notation θ for this state. In the previous story, the states $\mathcal{R}$ and $\mathcal{Y}$ can be defined by $\theta \in \{0, 1\}$ or $\theta \in \{\theta_0, \theta_1\}$.

The sighting by the witness is equivalent to the reception of a signal s that can be 0 or 1. A signal that takes one of two value is called a *binary signal*. The uncertainty about

Observation (signal)

States of Nature

	$s = 1$	$s = 0$
$\theta = \theta_1$	q_1	$1 - q_1$
$\theta = \theta_0$	$1 - q_0$	q_0

Table 1.1.1: Binary signal

the sighting is represented by the assumption that s is the realization of a random variable that depends on the true state. One possible dependence is given by Table 1.

Using the definition of conditional probability,

$$P(\theta = 1|s = 1) = \frac{P(\theta = 1 \cap s = 1)}{P(s = 1)} = \frac{P(s = 1|\theta = 1)P(\theta = 1)}{P(s = 1)},$$

which yields Bayes' rule

$$P(\theta = 1|s = 1) = \frac{q_1 P(\theta = 1)}{q_1 P(\theta = 1) + (1 - q_1)(1 - P(\theta = 1)}. \tag{1.2}$$

The signal 1 is "good news" about the state 1 (it increases the belief in state 1), if and only if $q_1 > 1 - q_0$, or

$$q_1 + q_0 > 1.$$

A signal can be informative about a state because it is likely to occur in that state, with q_1. But one should be aware that it may be even more informative when it is very unlikely to occur in the other state, when $1 - q_0$ is low. If one is looking for piece of metal, a good detector responds to an actual piece. But a better detector may be one that does not respond at all when there is no metal in front of it.

When $q_1 = q_0 = q$, the signal is a symmetric binary signal (SBS) and in this case, we will call q the precision of the signal. (The precision will have a different definition when the signal is not a SBS). Note that q could be less than $1/2$, in which case we could switch the roles of $s = 1$ and $s = 0$. The inequality $q > 1/2$ is just a convention, which will be kept here for any SBS.

Useful expressions of Bayes' rule

The formula in (1.2) is unwieldy. When the space state is discrete, it is often more useful to express Bayes' rule in terms of likelihood ratio, *i.e.*, the ratio between the probabilities of two states, hereafter LR. (There can be more than two states in the set of states). Here we have only two states, but LR is also useful for any finite number of states, as will be seen in the search application below.

$$\underbrace{\frac{P(\theta=1|s=1)}{P(\theta=0|s=1)}}_{\text{posterior LR}} = \underbrace{\left(\frac{P(s=1|\theta=1)}{P(s=1|\theta=0)}\right)}_{\text{signal factor}} \times \underbrace{\left(\frac{P(\theta=1)}{P(\theta=0)}\right)}_{\text{prior LR}}. \tag{1.3}$$

The signal factor depends only on the properties of the signal. With the specification of Table 1,

$$\frac{P(\theta=1|s=1)}{P(\theta=0|s=1)} = \frac{q_1}{1-q_0} \times \frac{P(\theta=1)}{P(\theta=0)}. \tag{1.4}$$

The expression of Bayes' rule in (1.3) is much simpler than the original formula because it takes a multiplicative form that has a symmetrical look.

State one is more likely when the LR is greater than 1. In the previous example of the car incident, say that "1" is "red". The prior for red cab is 1/10. The signal factor $P(s=1|\theta=1)/P(s=1|\theta=0)$ (correct / mistake) is .8/0.2=4. It is not sufficient to reverse the belief that yellow is more likely.

For some applications of rational learning, it will be convenient to transform the product in the the previous equation into a sum, which is performed by the logarithmic function. Denote by λ the prior *Log likelihood ratio* between the two states, and by λ' is posterior, after receiving the signal s. Bayes' rule now takes the form

$$\lambda' = \lambda + Log(q_1/(1-q_0)). \tag{1.5}$$

Both the multiplicative form in (1.3) and the additive form in (1.5) are especially when there is a sequence of signal. For example, with two signals s_1 and s_2,

$$\frac{P(\theta=1|s_1,s_2)}{P(\theta=0|s_1,s_2)} = \left(\frac{P(s_2|\theta=1)}{P(s_2|\theta=0)}\right) \times \left(\frac{P(s_1|\theta=1)}{P(s_1|\theta=0)}\right) \times \left(\frac{P(\theta=1)}{P(\theta=0)}\right).$$

One can repeat the updating for any number of signal observations. It is also obvious that the final update does not depend on the order of the signal observations.

Bounded signals and belief updates

The signal takes here only two values and is therefore bounded. The same is true if the number of signal values is more than two but finite. The implication is that values of the posterior probabilities cannot be arbitrarily close to one or zero. They are bounded away from zero and one. This will have profound implications later one. At this stage, one can just state that the binary signal (or any signal with finite values) is bounded.

1.1.3 Multiple binary signals: search on the sea floor

Some objects that have been lost at sea are extremely valuable and have stimulated many efforts for their recovery: submarines, nuclear bombs dropped of the coast of Spain, airline wrecks. In searching for the object under the surface of the sea, different informations have been used: last sight of the object, surface debris, surveys of the area by detecting instruments. The combination of these informations through Bayesian analysis led to the findings of the USS Scorpion submarine (2009), the USS Central America with its treasure (1857-1988), the wreck of AF 447 (2009-2011).

Assume that the search area is divided in N cells. The prior probability distribution is such that w_i is equal to the probability that the object is in cell i. Using previous notation, $w_i = P(\theta = \theta_i)$. If the detector is passed over cell i, the probability of finding the object is p_i, which may depend on the cell because of variations in the conditions for detection (depth, type of soil, etc..). The question is how after a fruitless search over an area, the probability distribution is updated from w to w'. Let θ_i be the state that the wreck is in cell i, and $\mathcal{Z}$ the event that no detection was made.

$$P(\theta = \theta_i|\mathcal{Z}) = \frac{1}{P(\mathcal{Z})}P(\mathcal{Z}|\theta = \theta_i)P(\theta = \theta_i).$$

$$P(\mathcal{Z}|\theta = \theta_i) = \begin{cases} 1 - p_i, \text{ if there if the detector is passed over cell } i, \\ \\ 1, \text{ if the detector is not passed over cell } i. \end{cases}$$

Defining $p_i = 0$ if there is no search in cell I (a search may not be over all the cells), the posterior distribution is given by

$$w'_i = A(1 - p_i)w_i, \quad \text{with} \quad A = \frac{1}{\sum_{i=1}^{N}(1 - p_i)w_i}. \tag{1.6}$$

An example: the search for AF447

In the early hours of June 1, 2009, with 228 passengers and crew, Air France Flight 447 disappeared in the celebrated "pot au noir".[5] No message had been sent by the crew but both "black boxes"–they are red– were retrieved after a two years. They have provided a

[5] This part of the *Intertropical Convergence Zone* (ITCZ) between Brazil and Africa is well known to aviators. It has been a special challenge for all sailboats, merchant ships in the 19th century and racers today.

gripping transcript of a failure of social learning in the cockpit during the last ten minutes of the flight. We focus here on the learning process during the search for the wreck, 3000 meters below the surface of the ocean. It provides a fascinating example of information gathering and learning.

First, a prior probability distribution (PD) has to be established. At each stage the probability distribution should orient the next search effort the result of which should be used to update the PD, and so on. That at least is the theory. [6] It will turn out that the search for AF447 did not follow the theory. Following Keller (2015), the search which lasted almost two years before a complete success, proceeded in stages.

1. The aircraft had issued an automated signal on its position at regular time intervals. From this, it was established that the object should be in a circle of 40 nautical miles[7] (nmi) centered at the last known position (LKP). That disk was endowed with a probability distribution, hereafter PD, that was chosen to be uniform.

2. Previous studies on crashes for similar conditions showed a normal distribution around the LKP with standard deviation of 8 nmi.

3. Five days after the crash, began a period during which debris were found, the first of them about 40 nmi from the LKP. A numerical model was used for "back drifting" to correct for currents and wind. That process, which is technical and beyond the scope of this analysis, led to another PD.

4. The three previous probability distributions were averaged with weights of $0.35, 0.35$ and 0.3, respectively. These weights are guesses and so far, the updating is not Bayesian. It's not clear how a Bayesian updating could have been done at this stage. The PD is now the prior distribution represented in the panel A of Figure 1.1. The Bayesian use of that PD will come only after Step 5.

5. Three different searches were conducted, with no result, between June and the end of 2010.

    (a) First, the black boxes of the aircraft are supposed to emit an audible sound for forty days. That search for a beacon is represented in the panel B of Figure 1.1. It produced nothing. There has been no Bayesian analysis at this stage, but all the steps in the search are carefully recorded and this data will be used later.

    (b) One had to turn to other methods. In August 2009, a sonar was towed in a rectangular area SE of the LKP because of a relatively flat bottom. Still nothing.

[6] See L. Stone **.

[7] One nautical mile =1.15 miles (one minute arc on a grand circle of the Earth).

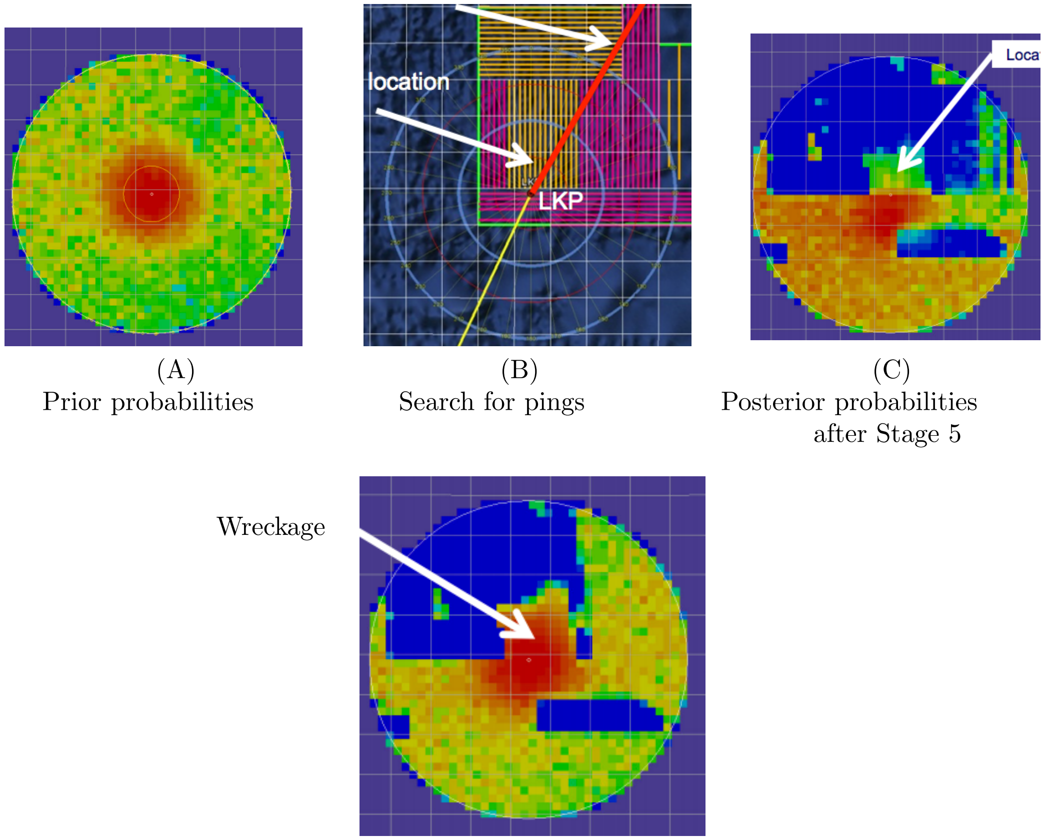

(A) Prior probabilities

(B) Search for pings

(C) Posterior probabilities after Stage 5

(D) Posterior assuming beacons failed

Source: Keller (2015).

Figure 1.1: Probability distributions in Bayesian search

(c) Two US ships from the Woods Hole Oceanographic Institute and from the US Navy searched an area that was a little wider than the NW quadrant of the 40 nmi disk. By the end of 2010, there were still no results.

6. Enters now Bayesian analysis. Each of the previous three steps, was used to update the prior PD (which, your recall, was an average of the first three PDs). The disc was divided in 7500 cells. Each search step is equivalent to 7500 binary signals s_i equal to 0 or 1 that turn out to be 0. The probabilities go according to the color spectrum, from high (red) to low (blue).

   (a) In step (a), the probability of survival for each bacon was set at 0.8. (More about this later). Conditional of survival, the probability of detection was estimated at

0.9. The probability of detection in that step was therefore 0.92. The updating is described in Exercise 1.2.

(b) In step (b), the probability of detection was estimated at 0.9 and the no find led to another Bayesian update of the PD.

(c) In step (c), the searches that were conducted in 2010 had another estimated probability of detection equal to 0.9 that was used in the third Bayesian update. The result of these three updates is represented in the panel C of Figure 1.1. The areas that have been searched have a low probability (in blue).

7. At this point, the results may have been puzzling. It was then decided, to assume that both the beacons in the black boxes had failed. The search in Panel B of the Figure was ignored and the distribution goes from Panel C to Panel D. See how the density of probability in the center part of the disc is now restored to a high level. The search was resumed in the most likely area and the wreck was found in little time (April 3, 2011).

In conclusion, the search relied on a mixture of educated guesses and Bayesian analysis. In particular, the failure of the search for pings should have led to a Bayesian increase of the probability of the failure of both beacons. The jump of the probability of failure from 0.1 to 1 in the final stage of the search seems to have been somewhat subjective, but it turned out to be correct.

1.1.4 The Gaussian model

The distributions of the prior θ and of the signal s (conditional on θ) are normal ("Gaussian", from Carl Friedrich Gauss). In this model, the learning process has nice properties. Using standard notation,

- $\theta \sim \mathcal{N}(\bar{\theta}, \sigma^2)$.
- $s = \theta + \epsilon, \quad \text{with} \quad \epsilon \sim \mathcal{N}(0, \sigma_\epsilon^2)$.

The first remarkable property of a normal distribution is that it is characterized by two parameters only, the mean and the variance. The inverse of the variance of a normal distribution is called the precision, for obvious reasons. Here the notation is such that $\rho_\theta = 1/\sigma^2$ and $\rho_\epsilon = 1/\sigma_\epsilon^2$.

The joint distribution of two normal distribution is also normal (with a density proportional to the exponential of the a quadratic form). Hence, the posterior distribution (the

These learning rules will be used repeatedly.

distribution of θ conditional on s) is also normal and the learning rule will be on two parameters only. First, the variance :

$$\sigma'^2 = \frac{\sigma^2 \sigma_\epsilon^2}{\sigma^2 + \sigma_\epsilon^2}.$$

This equation is much simpler when we use the precision, which is updated from ρ to ρ' according to

$$\rho' = \rho + \rho_\epsilon.$$

Admire the simple rule: to find the precision of the posterior we just add the precision of the signal to the precision of the prior.

Using the precisions, the updating rule for the mean is also very intuitive:

$$m' = \alpha s + (1-\alpha)m, \qquad \text{with} \qquad \alpha = \frac{\rho_\epsilon}{\rho}.$$

The posterior's mean is an average between the signal and the mean of the prior, each weighted by the precision of their distribution! It could not be more intuitive. And that rule is linear, which will be very useful.

$$\begin{cases} \rho' = \rho + \rho_\epsilon, \\ m' = \alpha s + (1-\alpha)m, \quad \text{with} \quad \alpha = \dfrac{\rho_\epsilon}{\rho}. \end{cases} \tag{1.7}$$

The Gaussian model is very popular because of the simplicity of this learning rule which which is recalled: (i) after the observation of a signal of precision ρ_ϵ, the precision of the subjective distribution is augmented by the same amount; (ii) the posterior mean is a weighted average of the signal and the prior mean, with weights proportional to the respective precisions. Since the *ex post* distribution is normal, the learning rule with a sequence of Gaussian signals which are independent conditional on θ is an iteration of (1.7).

The learning rule in the Gaussian model makes precise some general principles. These principles hold for a wider class of models, but only the Gaussian model provides such a simple formulation.

1. The normal distribution is summarized by the two most intuitive parameters of a distribution, the mean and the variance (or its inverse, the precision).
2. The weight of the private signal s depends on the noise to signal ratio in the most intuitive way. When the variance of the noise term σ_ϵ^2 tends to zero, or equivalently

its precision tends to infinity, the signal's weight α tends to one and the weight of the *ex ante* expected value of θ tends to zero. The expression of α provides a quantitative formulation of the trivial principle according to which *one relies more on a more precise signal*.

3. The signal s contributes to the information on θ which is measured by the increase in the precision on θ. According to the previous result, the increment is exactly equal to the precision of the signal (the inverse of the variance of its noise). The contribution of a set of independent signals is the sum of their precisions. This property is plausible, but it rules out situations where new information makes an agent less certain about θ, a point which is discussed further below.

4. More importantly, the increase in the precision on θ is *independent of the realization of the signal s*, and can be computed *ex ante*. This is handy for the measurement of the information gain which can be expected from a signal. Such a measurement is essential in deciding whether to receive the signal, either by purchasing it, or by delaying a profitable investment to wait for the signal.

5. The Gaussian model will fit particularly well with the quadratic payoff function and the decision problem which will be studied later.

1.1.5 Comparison of the two models

In the binary model, the distinction good/bad state is appealing. The probability distribution is given by one number. The learning rule with the binary signal is simple. These properties are convenient when solving exercises. The Gaussian model is convenient for other reasons which were enumerated previously. It is important to realize that each of the two models embodies some deep properties.

The evolution of confidence

When there are two states, the probability distribution is characterized by the probability μ of the good state. This value determines an index of confidence: if the two states are 0 and 1, the variance of the distribution is $\mu(1-\mu)$. Suppose that μ is near 1 and that new information arrives which reduces the value of μ. This information increases the variance of the estimate, *i.e.*, it reduces the confidence of the estimate. In the Gaussian model, new signals cannot reduce the precision of the subjective distribution. They always reduce the variance of this distribution.

Bounded and unbounded private informations

Another major difference between the two models is the strength of the private information. In the binary model, a signal has a bounded strength. In the updating formula (**??**), the multiplier is bounded. (It is either $p/(1-p')$ or $(1-p)/p'$). When the signal is symmetric, the parameter p defines its precision. In the Gaussian model, the private signal is unbounded and the changes of the expected value of θ are unbounded. The boundedness of a private signal will play an important role in social learning: a bounded private signal is overwhelmed by a strong prior. (See the example at the beginning of the chapter).

Binary states and Gaussian signals

If we want to represent a situation where confidence may decrease and the private signal is unbounded, we may turn to a combination of the two previous models.

Assume that the state space Θ has two elements, $\Theta = \{\theta_0, \theta_1\}$, and the private signal is Gaussian:

$$s = \theta + \epsilon, \qquad \text{with } \epsilon \sim \mathcal{N}(0, 1/\rho_\epsilon^2). \tag{1.8}$$

The LLR is updated according to

$$\lambda' = \lambda + \rho_\epsilon(\theta_1 - \theta_0)(s - \frac{\theta_1 + \theta_0}{2}). \tag{1.9}$$

Since s is unbounded, the private signal has an unbounded impact on the subjective probability of a state. There are values of s such that the likelihood ratio after receiving s is arbitrarily large.

1.1.6 Learning may lead to opposite beliefs: polarization

Different people have often different priors. The *same* information may lead to a convergence or a divergence of their beliefs. Assume first that there are only two states. In this case, without loss of generality, we can assume that the information takes the form of a binary signal as in Table 1. If two individuals observe the same signal s, their LR are multiplied by the same ratio $P(s|\theta_1)/P(s|\theta_0)$ that they move in the same direction.

In order to observe *diverging* updates, there must be more than two states. Consider the example with three states. these could be that the economy needs a reform to the left (state 1), to the center (state 2) or to the right (state 3). A signal s is produced either by a study or the implementation of a particular policy and provides an information on the state that is represented by the next table. (The signal $s = 1$ is a strong indication that

	$s=0$	$s=1$
$\theta=1$	0.3	0.7
$\theta=2$	0.9	0.1
$\theta=3$	0.3	0.7

the center policy is not working).

Two individuals, Alice and Bob, have their own prior on the states. Alice thinks that a policy on the right will not work and Bob thinks that a policy on the left will not work. Both have equal priors between the center and the right or the left. An example is presented in the next table.

	Alice	Bob
1	0.47	0.06
2	0.47	0.47
3	0.06	0.47

Priors

	Alice	Bob
1	0.79	0.1
2	0.11	0.11
3	0.1	0.79

Posteriors

After the signal $s=1$, Alice leans more on the left and Bob more on the right. The signal generates a *polarization* For Alice and Bob, the belief in the center decreases and for both of them, the beliefs in states 1 and 3 increase, but the increase is much higher for the state that has a higher prior, state 1 for Alice and state 2 for Bob. When θ is measured by a number, Alice and Bob draw opposite conclusions from the expected value of θ.

BIBLIOGRAPHY

* Anderson, Lisa R., and Charles A. Holt (1996). "Understanding Bayes' Rule," *Journal of Economic Perspectives*, 10(2), 179-187.

Salop, Steven C.. 1987. "Evaluating Uncertain Evidence with Sir Thomas Bayes: A Note for Teachers," *Journal of Economic Perspectives*, 1(1): 155-159.

* Keller, Colleen M. (2015). "Bayesian Search for Missing Aircraft," slides.

A superb presentation of four famous examples of Bayesian searches by a player in that field. Highly recommended.

Stone, Lawrence D., Colleen M. Keller, Thomas M. Kratzke and Johan P. Strumprer (2014). "Search for the Wreckage of Air France Flight AF 447," *Statistical Science*, 29 (1), 69-80.

Presents the search for AF 447. The next item, by a member of the team, is a conference presentation that discusses Bayesian searches for the USS Scorpion, the USS Central America, AF 447, and the failed search for MH 370. These slides are highly recommended, especially after reading the relevant section in this chapter.

Dixit, Avinash K. and Jörgen Weibull (2007). "Political polarization," *PNAS*, 104 (18), 7351-7356.

Williams, Arlington W., and James M. Walker (1993). "Computerized Laboratory Exercises for Microeconomics Education: Three Applications Motivated by the Methodology of Experimental Economics," *Journal of Economic Education*, 22, 291-315.

Jern, Alan, K-m I. Chang and C. Kemp (2014). "Belief Polarization is not always irrational," *Psychological Review*, **121**, 206-224.

EXERCISE 1.1. (The MLRP)

Construct a signal that does not satisfy the MLRP.

EXERCISE 1.2. (Simple probability computation, searching for a wreck)

An airplane carrying "two blackboxes" crashes into the sea. It is estimated that each box survives (emits a detectable signal) with probability s. After the crash, a detector is passed over the area of the crash. (We assume that we are sure that the wreck is in the area). Previous tests have shown that if a box survives, its signal is captured by the detector with probability q.

1. Determine algebraically he probability p_D that the detector gets a signal. What is the numerical value of p_D for $s = 0.8$ and $q = 0.9$?

2. Assume that there are two distinct spots, A and B, where the wreck could be. Each has a *prior* probability of 1/2. A detector is flown over the areas. Because of conditions on the sea floor, it is estimated that if the wreck is in A, the detector finds it with probability 0.9 while if the wreck is in B, the probability of detection is only 0.5. The search actually produces no detection. What are the *ex post* probabilities for finding the wreck in A and B?

EXERCISE 1.3. (non symmetric binary signal)

There are two states of nature, θ_0 and θ_1 and a binary signal such that $P(s = \theta_i|\theta_i) = q_i$. Note that q_1 and q_0 are not equal.

1. Let $q_1 = 3/4$ and $q_0 = 1/4$. Does the signal provide information? In general what is the condition for the signal to be informative?

2. Find the condition on q_1 and q_0 such that $s = 1$ is good news about the state θ_1.

EXERCISE 1.4. (Bayes' rule with a continuum of states)

Assume that an agent undertakes a project which succeeds with probability θ, (fails with probability $1 - \theta$), where θ is drawn from a uniform distribution on $(0, 1)$.

1. Determine the *ex post* distribution of θ for the agent after the failure of the project.

2. Assume that the project is repeated and fails n consecutive times. The outcomes are independent with the same probability θ. Determine an algebraic expression for the density of θ of this agent. Discuss intuitively the property of this density.

Chapter 2

Sequences of information and beliefs

2.1 Sequence of information with perfect memory

Suppose that $\mathcal{A}$ is a subset of the set Θ of all possible states. An example is one of two states, but there could be more than two states. There could also be a continuum of states and A could be, for example, an interval of real numbers. Let m_1 be the probability of $\mathcal{A}$. There are N rounds, or periods, of information and N can be infinite. In each round, a signal s_t is received. That signal may be, but does not have to be, a binary signal. It's probability distribution depends on the state. It therefore provides information on the state. The *history*, h_t, at the beginning of period t is defined as the sequence of signal before t:

$$\text{History in period } t: \qquad h_t = \{s_1, \dots, s_{t-1}\}. \tag{2.1}$$

We assume here perfect memory of the past signals.

After the reception of each signal s_t, the probability of $\mathcal{A}$ is revised from m_t to m_{t+1}. In formal notation,

$$m_{t+1} = P(\mathcal{A}|s_t, h_t).$$

In many cases, the information of history h_t will be summarized in m_t which is the probability of $\mathcal{A}$ given the history h_t. However, in some cases past history cannot be summarized in the current belief, in particular when the signals s_t are not independent (Exercise 2.1).

Stochastic path representations in probabilities

There are two states θ is equal to 1 or 0. There is a sequence of symmetric binary signals s_t, $(t \geq 1)$ as defined in Table 1 with a symmetric signal, $q_0 = q_1$. For a given state, the signals are independent. In each period t, the signal s_t is a random variable. Hence, the sequence of values m_t is a random sequence, a stochastic process. It can be represented by a trajectory, which is random, as on Figure 2.1. In the figure, we assume that the realization of the signals is the sequence $\{1, 0, 1, 1, 0, 1, 1, ...\}$. After each signal equal to 1, the belief increases and it decreases after each 0 signal. The signals 1 and 0 cancel each other and $m_1 = m3$, $mm_2 = m_4 = m_6$, $m5 = m7$. Note that the belief increase is smaller at m_4 than $m3$. That is because at m_4, the belief from history is higher and the impact of a good signal is smaller. (All the beliefs on the figure are greater than 1/2).

The probabilities of the branches are presented in blue under the assumption that the true state is 1. We could have other trajectories with different probabilities for their branches.

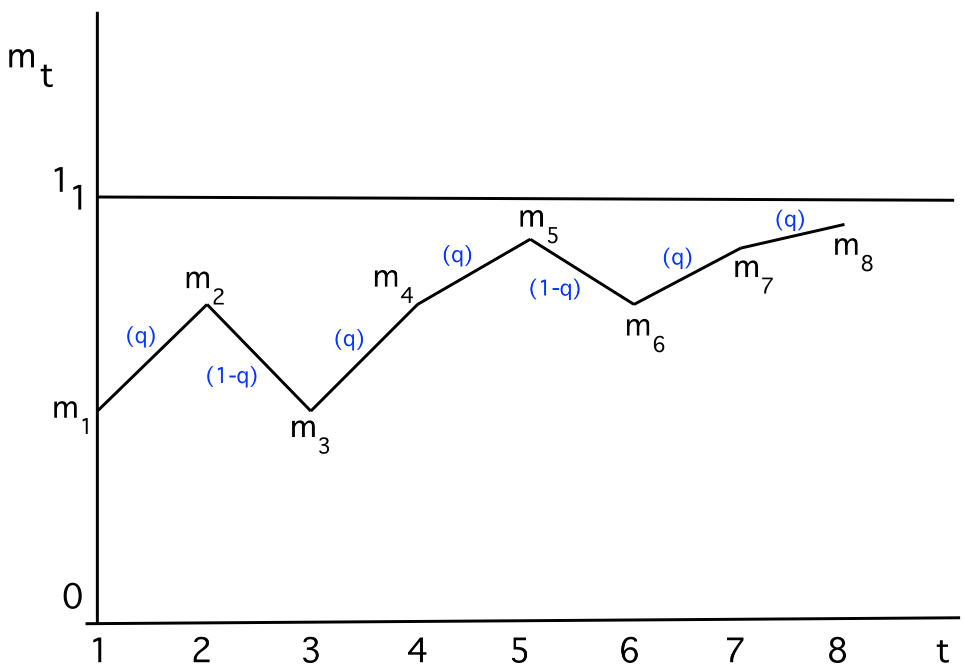

Figure 2.1: The evolution of belief as a stochastic process

Stochastic path representations in LLR

Bayes' rule in LR is simpler than the standard formula. For some applications, we can do even better with the Log Likelihood ratio (LLR). Define the prior LLR by

$$\lambda = \frac{P(\theta = 1)}{P(\theta = 0)},$$

and, likewise, the posterior LLR, λ'. Equation (1.2) becomes

$$\lambda' = \lambda + a, \qquad \text{with the signal term} \qquad a = \frac{P(s = 1|\theta = 1)}{P(s = 1|\theta = 0)}. \tag{2.2}$$

This expression has *two* useful properties: first the updating is additive; second the updating term is *independent* of the prior LLR. After some new information, agents with different prior LLRs have the *same* updating of their LLR. In the process of receiving information, different LLRs move in parallel!

In some cases, it will be useful to measure a belief by the Log likelihood (LLR). Recall that Θ is the space of all possible states. It has a probability equal to 1. Let λ_1 be the LLR of the subset of states $\mathcal{A}$ with respect to Θ:

$$\lambda_1 = Log\Big(\frac{P(\theta \in \mathcal{A})}{P(\theta \in \Theta)}\Big) = Log(P(\theta \in \mathcal{A})).$$

We have seen (equation 2.2) that the Bayesian updating after some signal s_t is such that

$$\lambda_{t+1} = \lambda_t + a_t, \tag{2.3}$$

where a_t depends on the properties of the signal s_t and on the signal value that was received in round t. Using the *parallel updating* of the LLRs, we have an elegant geometric representation of the beliefs for a population of agents with different prior beliefs. Suppose for example, that there are two agents, one with a higher private belief than the other, the "optimist" and the "pessimist", and that they receive the same sequence of informative signals. The evolution of their LLRs is illustrated in Figure 2.2.

Note that upwards and downwards moves have the same magnitude. The LLR is obviously not bounded. In the figure a LLR of 0 means equal probabilities for the two states. If the LLR is negative, the state 0 is more likely.

We can generalize this to a model with a continuum of agents, of total mass that can be taken equal to 1, each characterized by a prior belief. The distribution of prior beliefs (measured in LLR) is characterized by a density function with support **, which is assumed here to be a bounded interval of real numbers. When new information is received, the evolution of the beliefs of the population is represented by (random) translations of the support. For each of these supports, the density of the beliefs is the same as in the prior distribution.

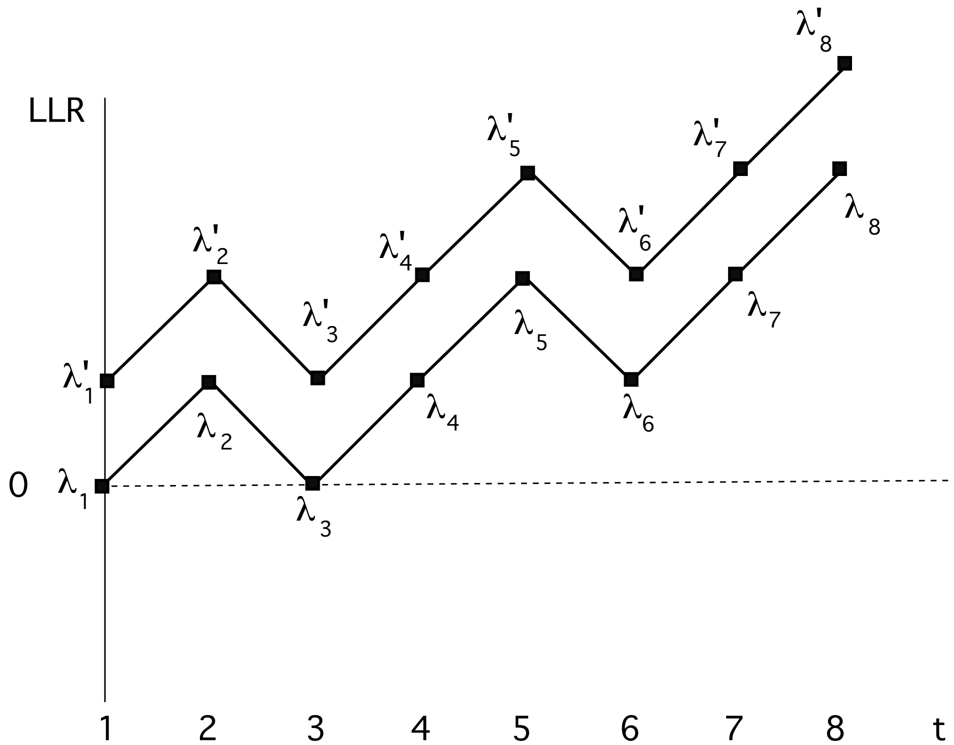

Figure 2.2: The evolution of LLs

Bounded and unbounded private informations

Definition: *When there exists M such that in the equation (2.3) for the updating of the LLR, $|a_t| \leq M$ for any t, the signal is bounded.*

When there is no such upper-bound, the signal is unbounded.

Examples:

- In the binary model, a signal has a bounded strength. In the updating formula (1.2), the multiplier is bounded. (It is either $p/(1-p')$ or $(1-p)/p'$).
- Assume that the state space Θ has two elements, $\Theta = \{\theta_0, \theta_1\}$, and the private signal is Gaussian:

$$s = \theta + \epsilon, \qquad \text{with } \epsilon \sim \mathcal{N}(0, 1/\rho_\epsilon^2). \tag{2.4}$$

Bayes' rule in log likelihood ratio (LLR) takes the form:

$$\lambda' = \lambda + \rho_\epsilon(\theta_1 - \theta_0)(s - \frac{\theta_1 + \theta_0}{2}). \tag{2.5}$$

Since s is unbounded, the private signal has an unbounded impact on the subjective probability of a state. There are values of s such that the likelihood ratio after receiving s is arbitrarily large.

2.2 Martingales

Bayesian learning satisfies the martingale property: changes of beliefs are not predictable.

Bayesian learning satisfies a strong property on the revision of the distribution of the states of nature. Suppose that before receiving a signal s, our expected value of a real number θ is $E[\theta]$. This expectation will be revised after the reception of s. Question: given the information that we have before receiving s, what is the expected value of the revision? Answer: zero. If the answer were not zero, we would incorporate it in the expectation of θ *ex ante*. This property is the *martingale property*. It is a central property of rational (Bayesian) learning. The martingale property separates rational from non rational learning.

The martingale property with learning from a binary signal

Assume that there are two signal values, $s =\in \{0, 1\}$. Let $P(\theta)$ be the probability that θ is equal to some value (or is in some set). P and P' denote prior (before the signal s) and posterior probabilities.

$$\begin{aligned}
E[P'(\theta)] &= P(s = 1)P'(\theta|s = 1) + P(s = 0)P'(\theta|s = 0),\\
&= P(s = 1)\frac{P(\theta \cap s = 1)}{P(s = 1)} + P(s = 0)\frac{P(\theta \cap s = 0)}{P(s = 0)},\\
&= P(\theta \cap s = 1) + P(\theta \cap s = 0),\\
&= P(\theta \cap (s = 1 \cup s = 0)) = P(\theta).
\end{aligned}$$

An equivalent result is

$$E[P'(\theta) - P(\theta)] = 0.$$

Note that $P(\theta)$ is not a random variable: it is the probability of θ before the signal is received. Before that reception, the expected value of the change of $P(\theta)$ (caused by the observation of the signal), is equal to 0! $P(\theta)$ is a martingale. If there are two states $\theta \in \{0, 1\}$, then $E[\theta] = P(\theta = 1)$ and $E[\theta]$ satisfies the martingale property.

The martingale property holds in general for any form of signal and if θ takes arbitrary values because it rests on the the property of conditional probabilities. Assume for example that θ has a density $g(\theta)$, and that s has a density $\phi(s|\theta)$ conditional on θ. Let $\psi(\theta|s)$ be the density of θ conditional on s. By Bayes' rule, $\psi(\theta|s) = \phi(s|\theta)g(\theta)/\phi(s)$, with $\phi(s) = \int \phi(s|\theta)g(\theta)d\theta$. Using $\int \phi(s|\theta)ds = 1$ for any θ,

$$E\big[E[\theta|s]\big] = \int\Big(\int \theta\psi(\theta|s)d\theta\Big)\phi(s)ds = \int\int \phi(s|\theta)\theta g(\theta)dsd\theta = \int \theta g(\theta)d\theta = E[\theta].$$

The similarity of this property with that of an efficient financial market is not fortuitous: in a financial market, updating is rational and it is rationally anticipated. Economists have often used martingales without knowing it.

A little formalism is helpful at this point. Assume that information comes as a sequence of signals s_t, one signal per period. Assume further that these signals have a distribution which depends on θ. They may or may not be independent, conditional on θ, and their distribution is known. Define the *history* in period t as $h_t = (s_1, \dots, s_t)$. The martingale property is defined for a sequence of real random variables as follows.[1]

DEFINITION 2.1. *The sequence of random variables Y_t is a martingale with respect to the history $h_t = (s_1, \dots, s_{t-1})$ if and only if*

$$Y_t = E[Y_{t+1}|h_t].$$

Expanding on the example with a binary signal, denote $\mu_t = E[\theta|h_t]$. Because the history h_t is random, μ_t is a sequence of random variables. The proof of the next result is the same as for the simple example

PROPOSITION 2.1. *Let $\mu_t = E[\theta|h_t]$ with $h_t = (s_1, \dots, s_{t-1})$. It satisfies the martingale property: $\mu_t = E[\mu_{t+1}|h_t]$.*

Let $\mathcal{A}$ be a set of values for θ, $\mathcal{A} \subset \Theta$, and consider the indicator function $I_{\mathcal{A}}$ for the set $\mathcal{A}$ which is the random variable given by

$$I_{\mathcal{A}}(\theta) = \begin{cases} 1 \text{ if } \theta \in \mathcal{A}, \\ 0 \text{ if } \theta \notin \mathcal{A}. \end{cases}$$

Using $P(\theta \in \mathcal{A}) = E[I_{\mathcal{A}}]$ and applying the previous proposition to the random variable $I_{\mathcal{A}}$ gives the next result.

PROPOSITION 2.2. *The probability assessment of an event by a Bayesian agent is a martingale: for an arbitrary set $\mathcal{A} \subset \Theta$, let $\mu_t = P(\theta \in \mathcal{A}|h_t)$ where h_t is the history of informations before period t; then $\mu_t = E[\mu_{t+1}|h_t]$.*

The likelihood ratio between two states θ_1 and θ_0 cannot be a martingale given the information of an agent. However, if the state is assumed to take a particular value, then the

[1] A useful reference is Grimmet and Stirzaker (1992).

likelihood ratio may be a martingale. Proving it is a good exercise.

PROPOSITION 2.3. *Conditional on $\theta = \theta_0$, the likelihood ratio*

$\dfrac{P(\theta = \theta_1|h_t)}{P(\theta = \theta_0|h_t)}$ *is a martingale.*

2.3 Convergence of beliefs

Probabilities will be equivalent to "beliefs". When more information comes in, does a belief (the probability estimate of a particular state) converge to some value. (We postpone the question whether it converges to the truth). We first need a definition of convergence. In this book, any convergence of a random variable (for example, a belief) is a convergence in probabiity[2]:

DEFINITION 2.2. *Let $X_1, X_2, \ldots, X_n, \ldots$ be random variables on some probability space $(\Omega, \mathcal{F}, P)$. X_n tends to a limit X in probability if*

- *for any given $\epsilon > 0$, $P(|X_n - X| \geq \epsilon) \to 0$ as $n \to \infty$.*

Note that the limit X is a random variable. For example, X_t may be a belief at history h_t. The sequence of beliefs converges but we don't know to which value it will converge.

Bayesian beliefs converge because of the Martingale Convergence Theorem.

A great property of any rational learning process is that beliefs converge. This convergence occurs because the sequence of beliefs is a martingale that is bounded (between 0 and 1 by definition of a probability) and the *martingale convergence theorem* (MCT) states that any bounded martingale converges.

The convergence of a bounded martingale, in a sense which will be made explicit, is a great result which is intuitive. The essence of a martingale is that its changes cannot be predicted, like the walk of a drunkard in a straight alley. The sides of the alley are the bounds of the martingale. If the changes of direction of the drunkard cannot be predicted, the only possibility is that these changes gradually taper off. For example, the drunkard cannot bounce against the side of the alley: once he hits the side, the direction of his next move is predictable.

[2] There are other criteria of convergence, for example the convergence almost sure (on a set of measure one in Ω, or convergences of the expected value of $|X_n|^r$, $r \geq 1$), but these are not useful at this stage for the analysis of the convergence of beliefs in a learning process. At this stage, there is no study of social learning with an example of convergence in probability and no convergence almost surely.

THEOREM 2.1. *(Martingale Convergence Theorem)*[3]

If X_t is a martingale with $|X| < M < \infty$ for some M and all t, then there exists a random variable X such that X_t converges to X.

Most of the social learning in this book will be about probability assessments that the state of nature belongs to some set $\mathcal{A} \subset \Theta$. We have seen that probability assessments satisfy the martingale property. They are obviously bounded by 1. Therefore they converge to some value.

PROPOSITION 2.5. *Let $\mathcal{A}$ be a subset of Θ and μ_t be the probability assessment $\mu_t = P(\theta \in \mathcal{A}|h_t)$, where h_t is a sequence of random variables in previous periods. Then there exists a random variable μ^* such that $\mu_t \to \mu^*$.*

Proof (hint): ("buy low, sell high")

There are various proofs of the MCT. Recall that the martingale property is the same as the efficient market equation. If a market is efficient, there is not strategy that has a positive expected gain. One proof of the MCT rests on the fact that the strategy "buy low, sell high" cannot generate a positive expected profit. Economists should have discovered the MCT.

We want to show that a belief, the probability of a state, or of an event, converges. Call that belief in round t, p_t. The stock is traded for T periods and new information is coming between periods. The truth is known in round $T + 1$. The stock pays 1 if the event takes place and 0 otherwise. The sequence of prices p_t is a martingale.

Take two numbers b and a with $0 < b < a < 1$. The difference $a - b$ may be small, but this is not important right now. The trading strategy is to buy one unit of the stock if the price is smaller than b, hold the stock until the price is higher than b, and sell the stock as soon as the price is higher than a. A new stock is bought when the price goes below b. In the strategy "buy low and sell high". "Low" and "high" are defined by the two values b and a.

If in period T, you hold the stock, you sell it at whatever the price in that period, p_T. The strategy is illustrated by Figure 2.3.

Define by N_T the number of times you buy a stock until round T, that is the number of upwards crossings of the band (b, a) in the trajectory of the price, p_t. The maximum loss

[3] Recall that we use only the convergence in probability. The theorem shows, under weaker conditions, the stronger property that the martingale converges almost everywhere.

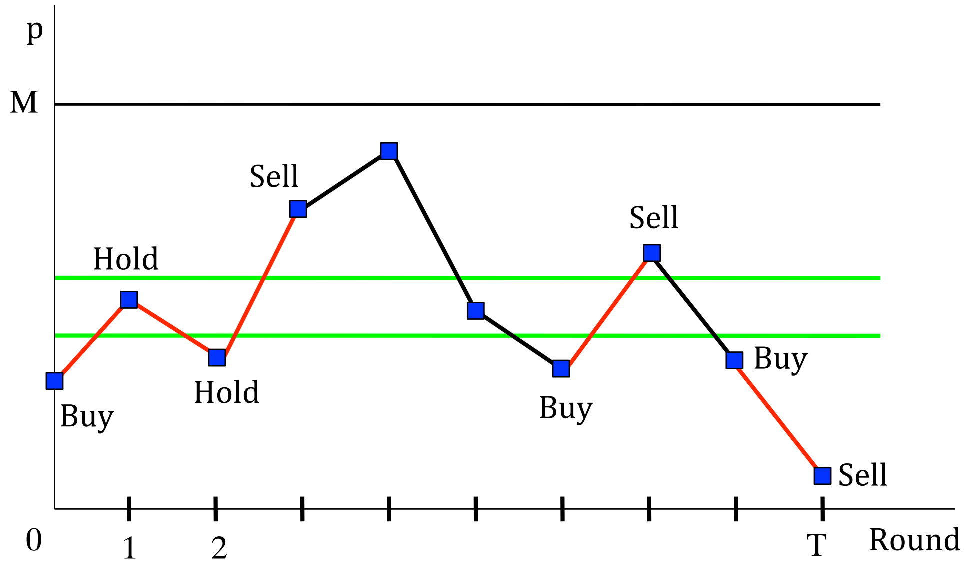

The agent holds one unit of the asset on the red segments.

Figure 2.3: A strategy of "buy low, sell high"

is b (if he has a stock that he sells in the last period). The gain is $N_T(a-b)$. Since $b < 1$, the net profit is not smaller than

$$V = N_T(a-b) - 1.$$

Because of the martingale property, the expected gain from the trading strategy cannot be positive. Hence, for any T,

$$E[N_T] \leq \frac{1}{(a-b)}.$$

The expectation of the number of upward crossing is bounded. From this, one can show that the probability of an upward crossing after period t tends to zero if t tends to infinity. One can then divide the interval $[0,1]$ in n intervals, each of with $1/n$ and iterate the previous argument for the finite number n. That means that for any ϵ, the stochastic process stays within one of these bands except with probability ϵ. Since the number n can be take as large as one wants, that proves the convergence in probability.[4]

A heuristic remark on another proof of the Martingale Convergence Theorem

The main intuition of the proof is important for our understanding of Bayesian learning. It is a formalization[6] of the metaphor of the drunkard. In words, the definition of a martingale

[4]From these intuitive hints, the reader can construct a formal proof. For verification, see Williams (1991).

[6]The proof is given in Grimmet and Stirzaker (1992). The different notions of convergence of a random

states that agents do not anticipate systematic errors. This implies that the updating difference $\mu_{t+1} - \mu_t$ is uncorrelated with μ_t. The same property holds for more distant periods: conditional on the information in period t, the random variables $\mu_{t+k+1} - \mu_{t+k}$ are uncorrelated for $k \geq 0$.

Since
$$\mu_{t+n} - \mu_t = \sum_{k=1}^{n} \mu_{t+k} - \mu_{t+k-1},$$

conditional on h_t, $Var(\mu_{t+n}) = \sum_{k=1}^{n} Var(\mu_{t+k} - \mu_{t+k-1})$.

Since $E[\mu_{t+n}^2]$ is bounded, $Var(\mu_{t+n})$ is bounded: there exists A such that

$$\text{for any } n, \qquad \sum_{k=1}^{n} Var(\mu_{t+k} - \mu_{t+k-1}) \leq A.$$

Since the sum is bounded, truncated sums after date T must converge to zero as $T \to \infty$: for any $\epsilon > 0$, there exists T such that for all $n > T$,

$$Var(\mu_{T+n} - \mu_T) = \sum_{k=1}^{n} Var(\mu_{T+k} - \mu_{T+k-1}) < \epsilon.$$

The amplitudes of all the variations of μ_t beyond any period T become vanishingly small as $t \to 0$. Therefore μ_t converges[7] to some value μ_∞. The limit value is in general random and depends on the history.

Rational (Bayesian) beliefs cannot cycle forever

Another way to look at the convergence of rational beliefs is to ask why they cannot have random cycles. If such cycles take place, there are random peaks and troughs, since the beliefs are between 0 and 1. But then how can the belief evolve when, say, it is close to 1. There is not much "room" to move up. Hence there cannot be much room to move down. If the belief could move down by a large amount, then, since it cannot move up by much, it should be have been adjusted right now. Of course, all this is in a probabilistic sense. The belief may move down by a large amount, but the larger the jump down, the smaller its probability. From this, we see that if the belief is close to 1, or to 0, it does not move up or down very much between periods.

variable are recalled in the Appendix.

[7]The convergence of μ_t is similar to the Cauchy property in a compact set for a sequence $\{x_t\}$: if $Sup_k(|x_{t+k} - x_t|) \to 0$ when $t \to \infty$, then there is x^* such that $x_t \to x^*$. The main task of the proof is to analyze carefully the convergence of μ_t.

One could also comment that if a belief, which has been generated by history is close to 1 , that means that history has provided convincing information that the event is highly probable. Any new information is rationally combined with history but the "weight" of this "convincing" history is such that new information can generate only a small change of belief.

Rational beliefs converge while non rational beliefs may not.

This deep property distinguishes rational Bayesian learning from other forms of learning. Many adaptative (mechanical) rules of learning with fixed weights from past signals are not Bayesian and do not lead to convergence. In Kirman (1993), agents follow a mechanical rule which can be compared to ants searching for sources of food, and their beliefs fluctuate randomly and endlessly.

The evolution of confidence

When there are two states, the probability distribution is characterized by the probability μ of the good state. This value determines an index of confidence: if the two states are 0 and 1, the variance of the distribution is $\mu(1-\mu)$. Suppose that μ is near 1 and that new information arrives which reduces the value of μ. This information increases the variance of the estimate, *i.e.*, it reduces the confidence of the estimate.

EXERCISE 2.1. (Non independent signals)

Construct an example with non independent signals where the history at time t cannot be summarized by the belief at time t.

BIBLIOGRAPHY

Williams, David (1991-2004). *Probability with Martingales*, Cambridge University Press.

Grimmett, Geoffrey an David Stirzaker (1982-2001). *Probability and random Processes*, Oxford University Press.

Park, Andreas and Hamid Sabourian (2011). "Herding and contrarian behavior in financial markets," *Econometrica*, **79**, 973-1026.

Chapter 3

Social learning

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

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

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

3.1 A canonical model of social learning

3.1.1 Structure

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

1. *The information endowments*: The *state of nature* is what the information is about. It is denoted by θ and is randomly chosen by nature before the learning process in a set Θ that can be finite or in a continuum. The probability distribution of nature is the *prior distribution* and is known to all agents.

2. The *private information* of an agent i, $i = 1, \ldots, N$, where N can be infinite, is what provides a value to others when they observe his action. That private information is modeled here by a random signal s_i. That signal has a probability distribution that is known by others in most cases (to make some inference possible), but by definition of *private*, the realization of the signal s_i cannot be observed by others. The signal provide some information on the state θ because its distribution depends on the true value of the state of nature θ. Any agent updates the prior on θ with the signal s_i to form a private distribution of probability of θ.

3. The action x_i of agent i is taken in round i, $(i \geq 1)$ and belongs to a set Ξ. (Without loss of generality, Ξ is the same set for all agents. The action will be the "message". We can assume here that this action is such that

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

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

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

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

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

3.1.2 The process

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

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

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

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

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

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

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

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

3.2 The Gaussian model

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

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

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

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

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

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

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

(i) The belief of agent t

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

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

(ii) The private decision

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

$$x_t = (1 - \alpha_t)\mu_t + \alpha_t s_t. \tag{3.4}$$

(iii) Social learning

Social learning is efficient when actions reveal perfectly private informations.

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

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

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

Convergence

The precision of the public belief increases linearly with time:

$$\rho_t = \rho_\theta + (t-1)\rho_\epsilon, \tag{3.6}$$

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

The weight of history and imitation

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

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

3.3 Observation noise

In the previous section, an agent's action conveyed perfectly his private information. An individual's action can reflect the slightest nuances of his information because: (i) it is chosen in a sufficiently rich menu; (ii) it is perfectly observable; (iii) the decision model of each agent is perfectly known to others.

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

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

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

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

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

An intuitive description of the critical mechanism

Period t brings to the public information the observation

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

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

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

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

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

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

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

The evolution of beliefs

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

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

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

A standard normalization will be used for most Gaussian signals.

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

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

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

Convergence

When there is no observation noise, the precision of the public belief ρ_t increases by a *constant* value ρ_ϵ in each period, and it is a linear function of the number of observations (equation (3.6)). When there is an observation noise, equation (3.11) shows that as $\rho_t \to \infty$, the increments of the precision, $\rho_{t+1} - \rho_t$, becomes smaller and smaller and tend to zero. The precision converges to infinity at a rate slower than a linear rate. The convergence of the variance σ_t^2 to 0 takes place at a rate slower than $1/t$.

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

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

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

3.3.1 Large number of agents

The previous model is modified to allow for a continuum of agents. Each agent is indexed by $i \in [0,1]$ (with a uniform distribution) and receives one private signal *once* at the beginning of the first period[3], $s_i = \theta + \epsilon_i$, with $\epsilon_i \sim \mathcal{N}(0, \sigma_\epsilon^2)$. Each agent takes an action $x_t(i)$ in each period[4] t to maximize the expected quadratic payoff in (**??**). At the end of period t, agents observe the aggregate action Y_t which is the sum of the individuals' actions and of an aggregate noise η_t:

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

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

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

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

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

By the law of large numbers[5], $\int \epsilon_i di = 0$. Therefore, $\alpha_t \int s_i di = \alpha_t \theta$. The level of endogenous aggregate activity is

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

and the observed aggregate action is

$$Y_t = (1-\alpha_t)\mu_t + \alpha_t\theta + \eta_t. \tag{3.12}$$

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

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

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

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

3.3.2 Application: a market equilibrium

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

$$p = a + \eta - bX. \tag{3.14}$$

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

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

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

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

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

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

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

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

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

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

which, using (3.15), is equivalent to

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

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

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

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

3.4 Extensions

Endogenous private information

See exercise 3.1.

Policy against mimetism

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

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

where x_t is the action of agent t. All agents observe the action plus a noise, $y_t = x_t + \epsilon_t$. The function W is interpreted as a loss function as long as θ is not revealed by a random exogenous process. In any period t, conditional on no previous revelation, θ is revealed perfectly with probability $1 - \pi \geq 0$. Assuming a discount factor $\delta < 1$, the value of β is $\beta = \pi\delta$. If the value of θ is revealed, there is no more loss.

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

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

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

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

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

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

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

EXERCISES

EXERCISE 3.1. (Endogenous private information)

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

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

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

1. Show using words and no algebra, that there is a date T after which no agent buys a private signal. What happens to information and actions after that date T?
2. Provide now a formal proof of the the previous statement. For this compute the welfare gain that an agent gets by buying a signal.
3. Assume now that the cost of a signal with precision ρ is an increasing function,[6] $c(\rho)$. Prove the following result:
   - Suppose that $c'(\rho)$ is continuous and $c(0) = 0$. If the marginal cost of precision $c'(\rho)$ is bounded away from 0, (for any $\rho \geq 0$, $c'(\rho) \geq \gamma > 0$), no agent purchases a signal after some finite period T and social learning stops in that period.
4. Assume now that $c(q) = q^\beta$ with $\beta > 0$. Analyze the rate of convergence of social learning.

REFERENCES

Burguet, R. and X. Vives, (2000). "Social Learning and Costly Information Acquisition," *Economic Theory*, **15**, 185-205, (first version 1995).

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

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

Lee, In Ho (1992). "On the Convergence of Informational Cascades," *Journal of Economic Theory*, **61**, 395-411.

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

Vives, X. (1996). "Social Learning and Rational Expectations," *European Economic Review*, **40**, 589-601.

Vives, X. (1997). "Learning From Others: a Welfare Analysis," *Games and Economic Behavior*, **20**, 177-200.

Chapter 4

Cascades and herds

A tale of two restaurants

Two restaurants face each other on the main street of a charming alsatian village. There is no menu outside. It is 6pm. Both restaurants are empty. A tourist comes down the street, looks at each of the restaurants and goes into one of them. After a while, another tourist shows up, evaluates how many patrons are already inside by looking through the stained glass windows—these are alsatian *winstube*—and chooses one of them. The scene repeats itself with new tourists checking on the popularity of each restaurant before entering one of them. After a while, all newcomers choose the same restaurant: they choose the more popular one irrespective of their own information. This tale illustrates how rational people may herd and choose one action because it is chosen by others. Among the many similar stories, two are particularly enlightening.

High sales promote high sales

In 1995, management gurus Michael Reacy and Fred Wiersema secretly purchased 50,000 copies of their business strategy book *The Discipline of Market Leaders* from stores which were monitored for the bestseller list of the *New York Times*[1]. The authors must have been motivated by the following argument: people observe the sales, but not the payoffs of the purchases (assuming they have few opportunities to meet other readers). Of course, if the manipulation had been known it would have had no effect, but people rationally expect that for any given book, the probability of manipulation is small, and that the high sales must be driven by some informed buyers.

[1] See Bikhchandani, Hirshleifer and Welch (1998), and Business Week, August 7, 1995. Additional examples are given in Bikhchandani, Hirshleifer and Welch, (1992).

The previous story illustrates one possible motivation for using the herding effect but it is only indicative. For an actual measurement, we turn to Hanson and Putler (1996) who conducted a nice experiment which combines the control of a laboratory with a "real situation". They manipulated a service provided by America Online (AOL) in the summer of 1995. Customers of the service could download games from a bulletin board. The games were free, but the download entailed some cost linked to the time spent in trying out the game. Some games were downloaded more than others.

The service of AOL is summarized by the window available to subscribers which is reproduced in Figure **??**: column 1 shows the first date the product was available; column 2 the name of the product, which is informative; column 4 the most recent date the file was downloaded. Column 3 is the most important and shows the number of customers who have downloaded the file so far. It presents an index of the "popularity" of the product. The main goal of the study is to investigate whether a high popularity increases the demand *ceteris paribus*.

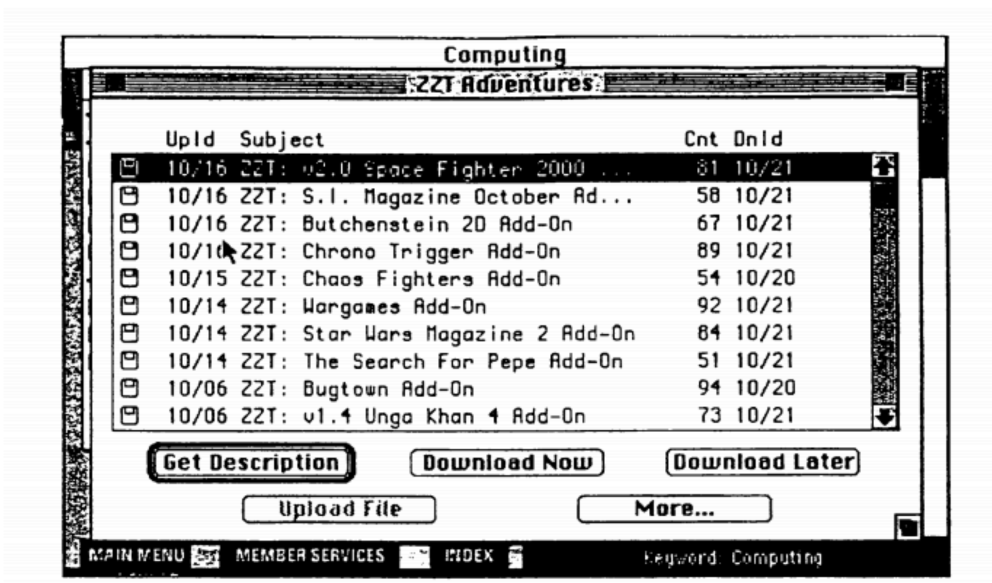

Figure 4.1: Applications for downloads

The impact of a treatment is measured by the increase in the number of downloads per day, after the treatment, as a fraction of the average daily download (for the same product) before the treatment. The results are reported in Figure ??. All treatments have an impact and the impact of the heavy treatment (100 percent) is particularly remarkable. The experiment has an obvious implication for the general manipulation of demand through advertisements.

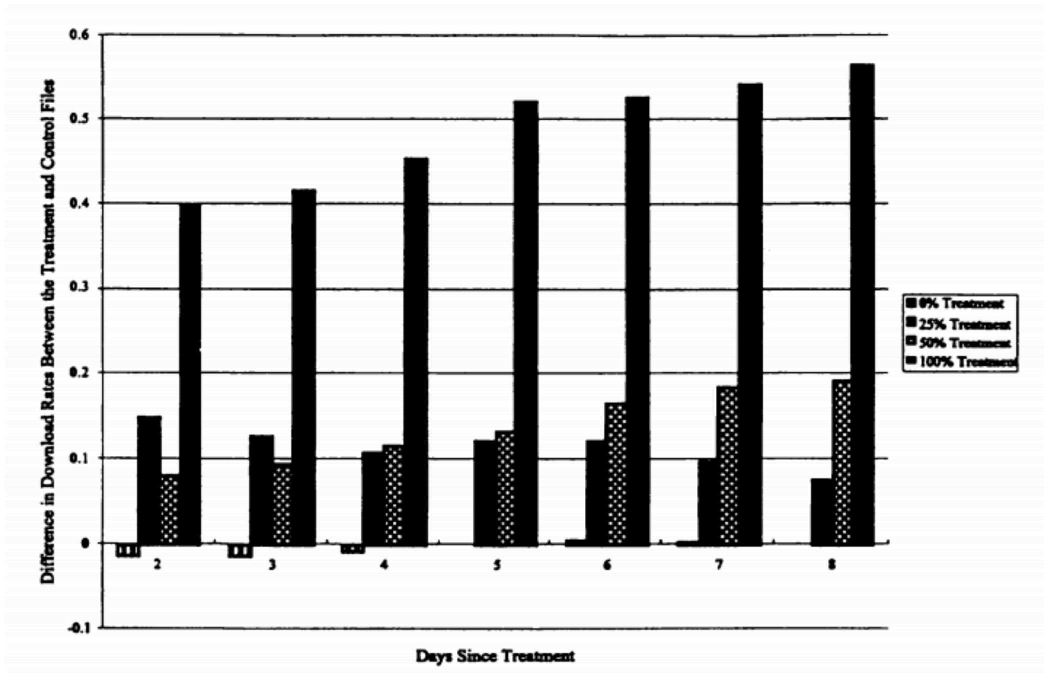

Figure 4.2: Results

To ensure *ceteris paribus*, Hanson and Putler selected pairs of similar files which were offered by AOL. Similarity was measured by characteristics and "popularity" at a specific date. Once a pair was selected, one of the files was kept as the "control", the other was the "treatment". The authors boosted the popularity index of the treatment file by downloading it repeatedly. The popularity indexed was thus increased in a short session by percentage increments of 25, 50 and 100. Customers of the service were not aware that they were manipulated.

The essential issue and the framework of analysis

The previous examples share a common feature which is essential: individuals observe the actions of others (and not their satisfaction), and the space of actions is discrete. The actions are the words for the communication of information between agents. In the previous chapter, agents chose an action in a rich set (*e.g.*, the set of real numbers), where even the smallest differences between beliefs could be expressed. Here the finite number of actions exerts a strong restriction on the vocabulary of social communication.

Bikchandani, Hirshleifer and Welch (1992), hereafter BHW, introduced the definition of informational cascades in models of Bayesian learning.[2] In a cascade, the public belief, which is gathered from the history of observations, dominates the private information of any individual: the action of any agent does not depend on his private information and all agents are herding. Since actions do not convey information, nothing is learned and the cascade goes on forever, possibly with an incorrect action. In the previous chapter we saw how observation noise slows down the process of social communicatioin. Here, it eventually comes to a complete stop.

There is an essential difference between a cascade and a herd.

A cascade generates a herd, but the concepts of cascade and herd are distinct. In a cascade, all agents ignore their information and take the same action. That behavior is known by all. Hence any agent can predict the behavior of others, *before* they take action and nothing is learned from the observation of others. In some way, the informational cascade is an *ex ante* concept.

In a herd, all agents *turn out* to take the same action. The action of an agent cannot be predicted before it is taken. Hence, something is learned from the observation of that action. For example, after an "investment", the belief in a "good state" is reinforced because there was the possibility that the agent would have not invested. In a herd with investment, the belief in the good state gradually increases. One should stress that here, agents can never know that a herd takes place. In this sense, a herd is an *post* concept.

Of course, a informational cascade generates a herd, and in this case, agents do know that they are in a herd. But the previous description hints at herds with no cascade. A cascade implies a herd, but the reverse is not true. In fact, we will see that in a setting where agents take discrete actions and no cascade occurs, a herd *must* eventually take place.

The simplest model of cascades is presented in Section 4.1. No formal arithmetics are required for that section which presents the important properties. The general model is analyzed in Section 4.2. The conditions for informational cascades are shown to be discrete actions and bounded private beliefs. If private beliefs are unbounded, then there is always the possibility that some agent with sufficiently strong and contrarian belief to diverge

[2] The expression "cascade" may be inappropriate for the description of a frozen behavior. Compare with the description of a cascade in the prologue of Faust II. One should also point that the expression "cascade" was already used by Gabriel Tarde (1890), in the same sense as BHW. (See the bibliographical note).

from "the crowd". In this case, no cascade ever takes place. However, in this case, the public belief converges to the truth. The public belief of the state tends to 1 or 0 and the bar for the strength of a contrarian belief is higher and higher. As the probability of a contrarian agent becomes smaller and smaller–a property that generates the *ex post* herd–one learns less and less from the observation of others. Although no informational cascade takes place, the social learning becomes vanishing small, asymptotically, which, from a welfare point of view, is not very different from an informational cascade where learning completely stops after finite time.

Section 4.4 presents a detailed analysis of herds and the convergence of beliefs[3]. Herds always take place eventually, as a consequence of the Martingale Convergence Theorem. There is in general some learning in a herd, but that learning is very slow. The conclusions of the simple model of BHW are shown to be extraordinarily robust. They reinforce the central message of the models of learning from others which is the self-defeating property of social learning when individuals use rationally the public information.

4.1 The basic model of cascades

A textbook case on how to build a model

Students sometimes wonder how to build a model. Bikhchandani, Hirshleifer and Welsh (1992), hereafter BHW, provide an excellent lesson of methodology[4]: (i) a good story simplifies the complex reality and keeps the main elements; (ii) this story is translated into a set of assumptions about the structure of a model (information of agents, payoff functions); (iii) the equilibrium behavior of rational agents is analyzed; (iv) the robustness of the model is examined through extensions of the initial assumptions.

We begin here with the tale of two restaurants, or a similar story where agents have to decide whether to make a fixed size investment. We construct a model with two states (defining which restaurant is better), two signal values (which generate different beliefs), and two possible actions (eating at one of two restaurants)[5].

[3] For this section, I have greatly benefited from the insights of Lones Smith.

[4] Banerjee (1992) presented at the same time another paper on herding, but its structure is more idiosyncratic and one cannot analyze the robustness of its properties.

[5] The example of the restaurants at the beginning of this chapter is found in Banerjee (1992). The model in this section is constructed on this story. It is somewhat mistifying that Banerjee after introducing herding through this example, develops an unrelated model which is somewhat idiosyncratic. A simplified version is presented in Exercise 4.2.

The 2 by 2 by 2 model

As in any model of Bayesian social learning, the structure has three blocks, the state of nature, the private informations, the private decisions for action and the observation of these actions.

1. The state of nature θ has two possible values, $\theta \in \Theta = \{0, 1\}$, and is set randomly once and for all at the beginning of the first period[6] with a probability μ_1 for the state $\theta = 1$. The value of θ could be the payoff of making a fixed size investment of 1. It could also be defined as "restaurant A is better than B" (with $\theta = 0$ representing the opposite proposition).

2. N or a countable number of agents are indexed by the integer t. Each agent's private information takes the form of a SBS (symmetric binary signal) with precision $q > 1/2$: $P(s_t = \theta \mid \theta) = q$. The signal represents the private information of an investor. It could also represent the information of a travel book or a friend's recommendation about the quality of restaurant.

3. Agents take an action in an *exogenous order* as in the previous models of social learning. The notation can be chosen such that agent t can make a decision in period t and in period t only. An agent chooses his action x in the discrete set $X = \{0, 1\}$. The action $x = 1$ may represent entering a restaurant, hiring an employee, or in general making an investment of a fixed size. The yield of the action x depends on the state of nature and is defined by

$$u(x, \theta) = \begin{cases} 0, & \text{if } x = 0, \\ \\ \theta - c, & \text{if } x = 1, \text{ with } 0 < c < 1. \end{cases}$$

   Since $x = 0$ or 1, another representation of the payoff is $u(x, \theta) = (\theta - c)x$. The cost of the investment c is fixed.[7] The yield of the investment is positive in the good state and negative in the bad state. Under uncertainty, the payoff of the agent is the expected value of $u(x, \theta)$ conditional on the information of the agent.

4. As in the previous models of social learning, the information of agent t is his private signal and the *history* $h_t = (x_1, \ldots, x_{t-1})$ of the actions of the agents who precede him in the exogenous sequence. The *public belief* at the beginning of period t is the probability of the good state conditional on the history h_t which is public information.

[6] The value of θ does not change because we want to analyze the changes in beliefs which are caused only by endogenous behavior. Changes of θ can be analyzed in a separate study (see the bibliographical notes).

[7] In the tale of two restaurants, c could be taken as $1/2$.

It is denoted by μ_t:

$$\mu_t = P(\theta = 1|h_t).$$

Without loss of generality, μ_1 is the same as nature's probability of choosing $\theta = 1$.

Let us use the representation of the social learning in Log likelihood as in Figure 2.2. The Log likelihood ratio between states 1 and 0 in the public information at the beginning of period t is λ_t. We call optimists the agents with signal $s = 1$ and pessimists the other ones. Agent t combines the public belief with his private signal to form his belief. Let λ_t^+ and λ_t^- the belief of an optimist or a pessimist in period t. We have seen (2.3) that

$$\lambda_t^+ = \lambda_t + a, \quad \lambda_t^- = \lambda_t - a, \quad \text{with} \quad a = Log(q/(1-q)),$$

and that the LLR "distance" between agents is constant in the process of social learning.

A geometric representation

Agent t takes action 1 if his belief (in state 1) is greater than c, which is equivalent to a LLR greater than $Log(c/(1-c))$. Let us denote $\gamma = Log(c/(1-c))$. Note that in the example of the two restaurants, $c = 1/2$ and $\gamma = 0$. The social learning is now represented in Figure 4.1 where the LLR is measured along the vertical axis.

In each period, a segment represents the distribution of beliefs: the top of the segment represents the belief of an optimist, the bottom the belief of a pessimist and the mid-point the public belief. The segments evolve randomly over time according to the observations.

In the first period, the belief of an optimist, λ_1^+, is above γ while the belief of a pessimist, λ_1^-, is below γ. The action is equal to the signal of the agent and thus reveals that signal. In the figure, $s_1 = 0$, and the first agent does not invest. His information is incorporated in the public information: the public belief in the second period, λ_2, is identical to the belief of the first agent: $\lambda_2 = \lambda_1^-$. The sequence of the signal ndowments is indicated in the figure. *When there is social learning, the signal of agent t is integrated in the public information of period $t + 1$.*

Consider now period 5 in the figure: agent 5 is an optimist, invests and eveals his signal since he could have been a pessimist who does not invest. His information is incorporated in the public belief of the next period and $\lambda_6 = \lambda_5^+$. The belief of a pessimist in period 6 is now higher than the cost c (here, it is equal to the public belief λ_5). In period 6, the belief of an agent is higher than the cost of investment, whatever his signal. He invests, nothing is learned and the public belief is the same in period 7: a cascade begins in period 6. The

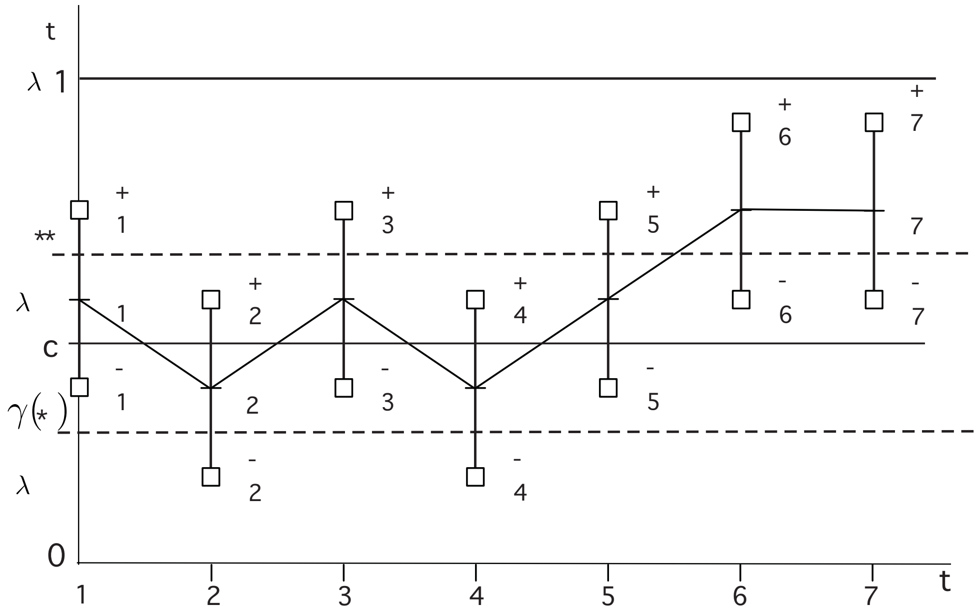

In each period, the middle of the vertical segment is the public belief, while the top and the bottom of the segment are the beliefs of an optimist (with a private signal $s = 1$) and of a pessimist (with signal $s = 0$). The private signals are $s_1 = 0, s_2 = 1, s_3 = 0, s_4 = 1, s_5 = 1$. `(Ignore the horizontal line at 1)`.

Figure 4.3: Cascade representation

cascade takes place because all the beliefs are above the cut-off level c. This condition is met here because the public belief λ_6 is strictly higher than λ^{**}. Since λ_6 is identical to the belief of an optimist in period 5, the cascade occurs because the beliefs of all investing agents are strictly higher than λ^{**} in period 5. A cascade takes place because of the high belief of the last agent who triggers the cascade.

Proposition 4.1 formalizes the previous discussion. It is expressed in beliefs $\mu_t = P(\theta = 1|h_t)$.

PROPOSITION 4.1. *In any period t, given the public belief μ_t:*

if $\mu^* < \mu_t \leq \mu^{**}$, agent t invests if and only if his signal is good ($s_t = 1$);

if $\mu_t > \mu^{**}$, agent t invests independently of his signal;

if $\mu_t \leq \mu^*$, agent t does not invest independently of his signal.

4.2 Beyond the simple model

We keep the structure with two states $\theta \in \{0, 1\}$, two actions $x \in \{0, 1\}$, with a payoff $(E[\theta] - c)x$, $0 < c < 1$. The states 1 and 0 will be called "good" and "bad". We extend the previous model to admit any distribution of private beliefs that is generated by private signals. Such a distribution is characterized by the *c.d.f.* $F^{\theta}(\mu)$ which depends on the state θ. $F^{\theta}(\mu)$ denotes the *c.d.f.* of a distribution of the beliefs measured as the probability of θ_1, and $F^{\theta}(\lambda)$ denotes the *c.d.f.* of a distribution of the LLR between θ_1 and θ_0. A graphical representation of the distributions of beliefs in LLR is given in Figure (See also exercise **??**).

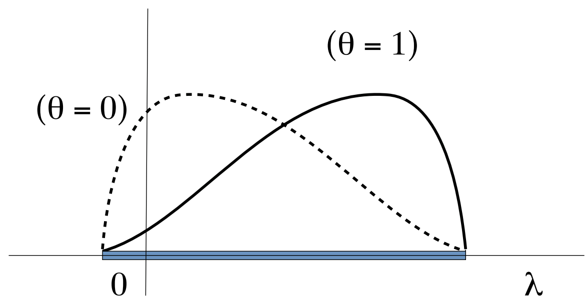

The distribution of private beliefs is represented by the curves with $\theta = 1$ and $\theta = 0$ in the two states.

Figure 4.4: Distribution of beliefs (in LLR)

Note that the mass of probabilities is shifted upwards (downwards) when $\theta = 1$ ($\theta = 0$): in the good state, more agents are optimistic. This can be observed if private signals are symmetric binary with precision q and q is drawn from a distribution. The shape of the distribution of beliefs in LLR is invariant to the news. We have seen in ** that any new information has the same additive impact (positive or negative) on the LLR of any agent. The news just translate the distribution.

In Figure 4.2, the distribution of beliefs, as in Figure 4.2, is be placed in the diagram of Figure **??** that illustrates the process of social learning. In any period, the acting agent's belief is drawn from the distribution that is determined by the state. An example is represented by the red arrows with LLR λ_1, λ_2 and λ_3 in the first three rounds.

If the belief is above the cutoff level (0 in the figure), then the agent takes $x = 1$. Otherwise, he takes $x = 0$. The probability that this belief in LLR is positive is greater in the good state. It is proportional to the area in blue in the good state and in red in the other state. Therefore, when other agents observe $x = 1$, they increase their belief in the good state. By how much?

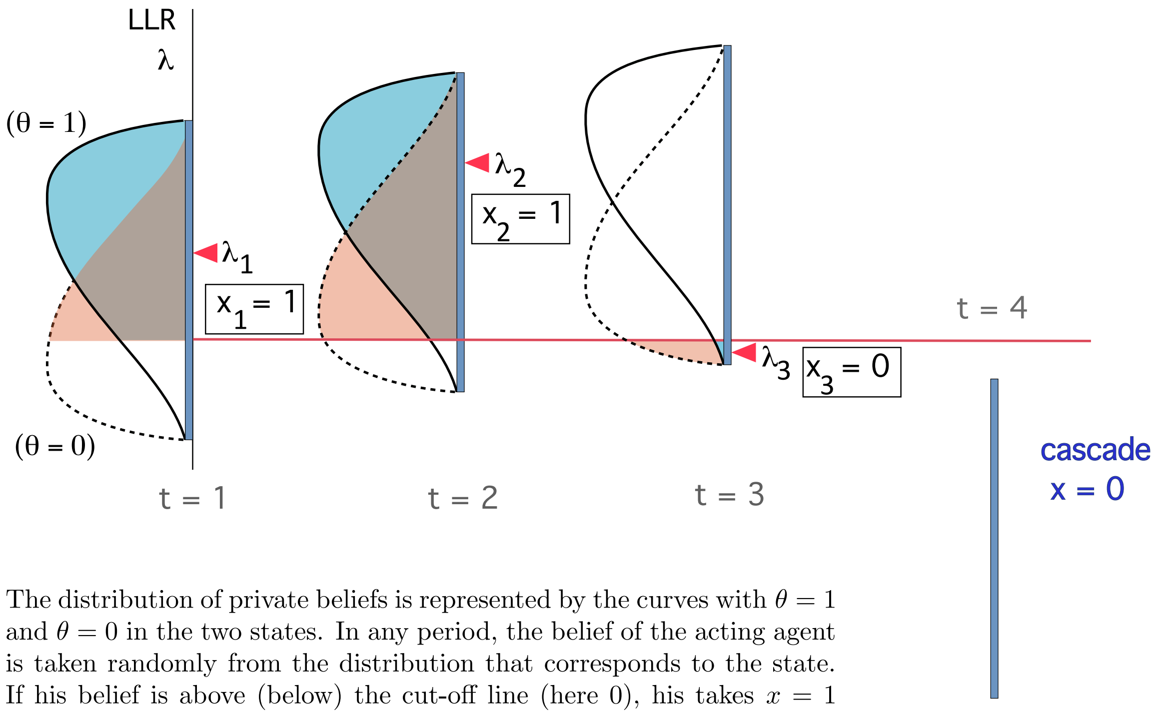

The distribution of private beliefs is represented by the curves with $\theta = 1$ and $\theta = 0$ in the two states. In any period, the belief of the acting agent is taken randomly from the distribution that corresponds to the state. If his belief is above (below) the cut-off line (here 0), his takes $x = 1$ $(x = 0)$. See the discussion in the text.

Figure 4.5: Representation with a general bounded distribution of beliefs

Take the point of view of an outside observer. For him, x_1 is just a binary signal on the state. That signal is characterized by the probabilities $x_1 = 1$ for each of the two states. Here, the matrix of the probabilities takes the following form.

		Observations	
		$x_t = 1$	$x_t = 0$
States of Nature	$\theta = \theta_1$	$1 - F_t^{\theta_1}(\gamma)$	$F_t^{\theta_1}(\gamma)$
	$\theta = \theta_0$	$1 - F_t^{\theta_0}(\gamma)$	$F_t^{\theta_0}(\gamma)$

For example, if $\theta = \theta_1$, the probability that the acting agent has a LLR below γ is $F_1^{\theta_1}(\gamma)$. One can then fill the other elements of the matrix.

The update of the LLR after the observation of x_t is, for an agent with LLR equal to λ_t (any value), is determined as in (**??**) which is repeated here.

$$\lambda_{t+1} = \lambda_t + \nu_t, \quad \text{with} \quad \nu_t = Log\Big(\frac{P(x_t|\theta_1)}{P(x_t|\theta_0)}\Big). \tag{4.1}$$

As we have seen before, the updating term ν_t is independent of the belief λ_t. Therefore, the distribution of beliefs is translated by a random term ν_t from period t to period $t+1$. Agent t invests if and only if his probability of the good state is greater than his cost, *i.e.* if his LLR, λ, is greater than $\gamma = Log(c/(1-c))$. The probability that agent t invests depends on the state and is equal to $\pi_t(\theta) = 1 - F_t^{\theta}(\gamma)$, with $\gamma = Log\Big(\frac{c}{1-c}\Big)$.

Given the above matrix of the signal x_t,

$$\lambda_{t+1} = \lambda_t + \nu_t, \quad \text{with} \quad \nu_t = \begin{cases} Log\Big(\frac{1-F_t^{\theta_1}(\gamma)}{1-F_t^{\theta_0}(\gamma)}\Big), & \text{if } x_t = 1, \\ Log\Big(\frac{F_t^{\theta_1}(\gamma)}{F_t^{\theta_0}(\gamma)}\Big), & \text{if } x_t = 0. \end{cases} \tag{4.2}$$

In this equation, $\nu_t \geq 0$ if $x_t = 1$ and $\nu_t \leq 0$ if $x_t = 0$. The observation of x_t conveys some information on the state as long as $F_t^{\theta_1}(\gamma) \neq F_t^{\theta_0}(\gamma)$.

In period 2 of the figure, the beliefs are higher. Note that in either state, the probability of $x_2 = 1$ is higher. Both are closer to 1. Hence, the ratio between the blue and the red areas (greater than 1) is closer to 1 and there is "less news" after the observation of $x_2 = 1$. The LLR increases by $\nu_2 < \nu_1$.

In the figure, the increase of the LLR after the observation $x_2 = 1$ is such that the support of the distribution is above the cutoff line: the most pessimistic person now believes that the state $\theta = 1$ is more likely. All agents take the action 1, in an informational cascade.

Suppose now that $x_3 = 0$. We can see in Figure 4.2 that the ratio of the two areas below the cutoff line, for $\theta = 1$ and $\theta = 0$, respectively is now much smaller than one, perhaps of he order of $1/10$. That means that the term ν_2 is now negative with a large absolute value. The observation x_2 is followed by a large decrease of all beliefs. It is really news.

One verifies properties that we have seen before. At the beginning of period 2, the probability of $x_2 = 1$ is high, near one. If one observes $x_2 = 1$, then this is good news, it increases the beliefs, but by a small amount. The probability of $x_2 = 0$ is small, but this event if followed by a large decrease of the beliefs. Remember that the expected change of beliefs is zero. There is a large probability of a small upward change which is balanced a small probability of a large downward change.

4.3 Cascades and herds

The core of a social learning process is how agents convey their information through actions. For this, their actions has to be affected by their information. But the public information information may be so strong that it overwhelms private information and in this case, private information is revealed through actions. For example, in a simple model with two possible outcome of an investment decision, say, 1 or 0, the public belief (probability of state 1) may be so high that even if an agent has a low signal ($s = 0$), he still makes the investment. He invests both with a bad, and obviously, a good signal. His investment reveals nothing about his private information. The following definitions, which apply to any model of social learning will help to clarify the analysis.

DEFINITION 4.1. *An agent is herding if his action is independent of his private information.*

DEFINITION 4.2. *An information cascade takes place in period t when the agents' actions generate no information on the state.*

If all agents are herding in period t, there is no information in that period. The public belief is the same in the next period. If the environment is the same (the structure of the private signals is the same), then the action of an agent in the next period will also be independent of his private information. All agents will be herding again. The public belief stays the same and so on.

PROPOSITION 4.2. *If all agents are herding in period t, in a stationary environment, there is an informational cascade from period t on.*

When a cascade takes place, one knows at the beginning of a period the action of a rational agent *before* that action is taken. It also might happen that there is some uncertainty about the action taken in, say, period t, and that this action turns out to be the same as the one that was taken in the previous period.

DEFINITION 4.3. *A herd takes place at date T if all actions after date T are identical: for all $t > T$, $x_t = x_T$.*

A cascade obviously generate a herd. Can there be a herd *without* a cascade? The answer is yes! Even more so, when there are only two possible actions, or a finite number of actions, a herd must eventually take place with probability one even if there is never a cascade!

An important distinction: an informational cascade is sufficient for a herd, but a herd may occur without a cascade.

The converse of Proposition 4.2 is not true. *Herds and cascades are not equivalent.* In a herd, all agents turn out to choose the same action—in all periods— although some of them could have chosen a different action. The following result, which is due to Smith and Sørensen (2001) shows that in the standard model of social learning with discrete actions, herds always take place eventually even in a setting where cascades cannot occur!

THEOREM 4.1. *On any path $\{x_t\}_{t\geq 1}$ with social learning, a herd begins in finite time. With probability* 1*, on any path of actions $\{x_1, \ldots\}$ all actions are identical after some finite time.*

The result holds for any type of distribution of beliefs. If the distribution is bounded, then we know that the path of actions ends in a cascade and therefore in a herd. If the distribution is not bounded, then agents keep learning in every period and the public belief converges to the truth. In this case, the public belief changes in every period, but nevertheless, after some finite time, all actions are identical. Of course agents do not know that there is a herd and that is why they keep learning.

The proof of the theorem is based on the MCT. The intuition for the proof is straightforward. If the theorem were not true, "contrarian actions" would be observed an infinite number of time. Each such contrarian action would trigger a significant change of the public belief. That would contradict the convergence of the public belief according to the MCT.

4.4 The convergence of beliefs

When private beliefs are bounded, beliefs never converge to perfect knowledge. If the public belief would converge to 1 for example, in finite time it would overwhelm any private belief and a cascade would start thus making the convergence of the public belief to 1 impossible. This argument does not hold if the private beliefs are unbounded because in any period the probability of a "contrarian agent" is strictly positive.

PROPOSITION 4.3. *Assume that the initial distribution of private beliefs is unbounded. Then the belief of any agent converges to the truth: his probability assessment of the good state converges to* 1 *in the good state and to* 0 *in the bad state.*

Does convergence to the truth matter?

A bounded distribution of beliefs is necessary for a herd on an incorrect action, as emphasized by Smith and Sørensen (1999). Some have concluded that the properties of the simple model of BHW are not very robust: cascades are not generic and do not occur for sensible distributions of beliefs; the beliefs converge to the truth if there are agents with

To focus on whether social learning converges to the truth or not can be misleading.

sufficiently strong beliefs. In analyzing properties of social learning, the literature has often focused on whether learning converges to the truth or not. This focus is legitimate for theorists, but it is seriously misleading. What is the difference between a slow convergence to the truth and a fast convergence to an error? From a welfare point of view and for many people, it is not clear.

The focus on the ultimate convergence has sometimes hidden the central message of studies on social learning: the combination of history's weight and of self-interest slows down the learning from others. The beauty of the BHW model is that it is non generic in some sense (cascades do not occur under some perturbation), but its properties are generic.

If beliefs converge to the truth, the speed of convergence is the central issue. This is why the paper of Vives (1993) has been so useful in the previous chapter. We learned from that model that an observation noise reduces the speed of the learning from others. Since the discreteness of the action space is a particularly coarse filter, the slowing down of social learning should also take place here. When private beliefs are bounded, the social learning does not converge to the truth. When private beliefs are unbounded, we should observe a slow rate of convergence.

We saw that cascades do not occur for sensible distributions of beliefs because the signal of the action (investment or no investment) is vanishingly weak when the public belief tends to the cascade set corresponding to the action. This argument applies when the distribution of beliefs is unbounded, since the mass of atoms at the extreme ends of the distribution must be vanishingly small. Hence, there is an immediate presumption that social learning must be slow asymptotically. The slow learning is first illustrated in an example and then analyzed in detail.

A numerical example

The private signals are defined by $s = \theta + \epsilon$ where ϵ is normally distributed with variance σ^2. An exercise shows that if μ tends to 0, the mass of agents with beliefs above $1 - \mu$ tends to zero faster than any power of μ. A numerical example of the evolution of beliefs is presented in Figure 4.4. One observes immediately that the pattern is similar to a cascade in the BHW model with the occurrence of "black sheeps".

For this example only, it is assumed that the true state is 1. The initial belief of the agent is $\mu_1 = 0.2689$, (equivalent to a LLR of -1), and $\sigma = 1.5$. The actions of individuals in each period are presented by the lower schedule (equal to 0.1 if $x_t = 1$ and to 0 otherwise). For the first 135 periods, $x_t = 0$ and μ_t decreases monotonically from around 0.27 to around

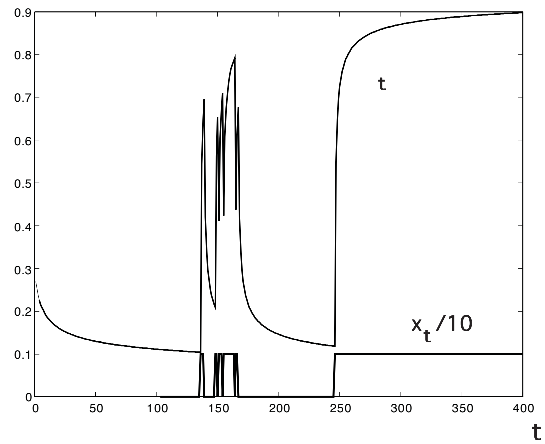

The lower curve represents the action of an agent. If positive, the action is equal to 1, otherwise it is equal to 0.

Figure 4.6: Social learning with two actions and unbounded private beliefs (Gaussian signals)

0.1. In period 136, the agent has a signal which is sufficiently strong to have a belief $\tilde{\mu}_{136} > c = 0.5$ and he invests. Following this action, the public belief is higher than 0.5 (since 0.5 is a lower bound on the belief of agent 135), and $\mu_{137} > 0.5$. In the example, $\mu_{137} = 0.54$. The next two agents also invest and $\mu_{139} = 0.7$. However, agent 139 does not invest and hence the public belief must fall below 0.5: $\mu_{140} = 0.42$. Each time the sign of $\mu_{t+1} - \mu_t$ changes, there is a large jump in μ_t.

Figure 4.4 provides a nice illustration of the herding properties found by BHW in a model with "black sheeps" which deviate from the herds. The figure exhibits two properties which are standard in models of social learning with discrete decisions:

(i) when μ_t eventually converges monotonically to the true value of 1 (after period 300 here), the convergence is very slow;

(ii) when a herd stops, the public belief changes by a quantum jump.

The slow learning from others

Assume now a precision of the private signals such that $\sigma_\epsilon = 4$, and an initial public belief $\mu_1 = 0.2689$ (with a LLR equal to -1). The true state is good. The model was simulated for 500 periods and the public belief was computed for period 500. The simulation was repeated 100 times. In 97 of the 100 simulations, no investment took place and the public belief decreased by a small amount to a value $\mu_{500} = 0.2659$. In only three cases did some investment take place with μ_{500} equal to 0.2912, 0.7052 and 0.6984, respectively. Hardly a fast convergence!

By contrast, consider the case where agents observe directly the private signals of others and do not have to make inferences from the observations of private actions. From the specification of the private signals and Bayes' rule,

$$\lambda_{t+1} = \lambda_1 + t(\frac{\theta_1 - \theta_0}{\sigma_\epsilon^2})(\frac{\theta_1 - \theta_0}{2} + \eta_t), \quad \text{with} \quad \eta_t = \frac{1}{t}\sum_{k=1}^{t} \epsilon_k.$$

Given the initial belief $\mu_1 = 0.2689$, $\theta_0 = 0$, $\theta_1 = 1$, $t = 499$ and $\sigma_\epsilon = 4$,

$$\lambda_{500} = -1 + (31.2)(0.5 + \eta_{500}),$$

where the variance of η_{500} is $16/499 \approx (0.18)^2$. Hence, λ_{500} is greater than 5.33 with probability 0.95. Converting the LLR in probabilities, μ_{500} belongs to the interval $(0.995, 1)$ with probability 0.95. What a difference with the case where agents observed private actions! The example—which is not particularly convoluted—shows that the convergence to the truth with unbounded private precisions may not mean much practically. Even when the distribution of private signals is unbounded, the process of social learning can be very slow when agents observe discrete actions. The cascades are a better stylized description of the properties of social learning through discrete actions than the convergence result of Proposition 4.3. The properties of the example are confirmed by the general analysis of the convergence that is provided in the Appendix.

REFERENCES

Battacharya, G. and G. Simons (1996). "Informational Cascades in Informal Markets," *Journal of Economics* - (MVEA), **22**, 47-55.

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

Becker, G. (1991). "A Note on Restaurant Pricing and Other Social Influences on Prices, " *Journal of Political Economy*, **99**, 1109-1116.

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**, 992-1026.

Chari, V.V. and P. Kehoe (2000). "Financial Crises as Herds", mimeo, University of Minnesota.

Conlisk, J. (1980). "Costly Optimizers versus Cheap Imitators," *Journal of Economic Behavior and Organization*, **1**, 275-93.

Huck, S. and J. Oechssler (1998). "Informational Cascades with Continuous Action Spaces," *Economics Letters*, **60**, 163-166.

Lee, In Ho (1992). "On the Convergence of Informational Cascades," *Journal of Economic Theory*, **61**, 395-411.

Smith, L. and P. Sørensen (1999). "Pathological Outcomes of Observational Learning," *Econometrica*, **68**, 371-398.

Welch, I. (1992). "Sequential Sales, Learning, and Cascades," *Journal of Finance*, **47**, 695-732.

EXERCISE 4.1. (Probability of a wrong cascade)

Consider the $2 \times 2 \times 2$ model that we have seen in class (2 states 1 and 0, 2 actions and symmetric binary signal), where μ_1 is the prior probability of the state 1, $c \in (0, 1)$ the cost of investment, and q the precision of the binary signal. There is a large number of agents who make a decision in a fixed sequence and who observe the actions of past agents. Assume that $\mu_1 < c$ and that the difference $c - \mu_1$ is small. Let $x_t \in \{0, \}$ the action of agent t. We assume that the true state (unknown by agents) is $\theta = 0$.

1. Represent on a diagram with time (horizontal axis) and the probability of state 1 in the public information (vertical axis), different examples of trajectories of the public belief that end in a cascade with investment, which is a "wrong" cascade (since the state is 0). We want to compute the probability of all these wrong cascades.

2. What is the probability that a cascade begins immediately after $x_1 = 1$. What do agents do in that cascade?

3. Call $\mathcal{I}$ the outcome that a cascade begins in some period in which all agents take action 1. Show that the probability of $\mathcal{I}$ before the decision of the first agent, call it π_0 is the same as before the decision of the third agent after a history of actions $x_1 = 1$, $x_2 = 0$.

4. Let π_1 the probability of $\mathcal{I}$ after the history $x_1 = 1$. Determine π_1 as a function of π_0.

5. What is the probability that a cascade with investment begins after $x_1 = 0$?

6. Using the previous questions, find another relation between π_0 and π_1.

7. Determine the probability π_0 of a wrong cascade.

EXERCISE 4.2. (The model of Banerjee, 1992)

Assume that the state of nature is a real number θ in the interval $(0, 1)$, with a uniform distribution. There is a countable set of agents, with private signals equal to θ with probability $\beta > 0$, and to a number uniformly distributed on the interval $(0, 1)$ with probability $1 - \beta > 0$. (In this case the signal is not informative). The agent observes only the value of his private signal. Each agent t chooses in period t an action $x_t \in (0, 1)$. The payoff is 1 if $x_t = \theta$, and 0 if $x_t \neq \theta$. Agent t observes the history of past actions and maximizes his expected payoff. If there is more than one action which maximizes his expected payoff, he chooses one of these actions with equal probability.

1. Analyze how herds occur in this model.

2. Can a herd arise on a wrong decision?

EXERCISE 4.3. (Action set is bounded below, Chari and Kehoe, 2000)

Consider a variation on the model of this chapter with two states, $\theta = 1$ or 0 and symmetric private signals s such that $P(s = \theta|\theta) = q$. Assume that any agent t chooses the action x_t that can be any real positive number: $x_t \geq 0$. The purpose of the exercise is to analyze social learning when investment cannot be negative. Agents maximize the expected value of the payoff function

$$u(x, \theta) = 2(\theta - c)x - x^2, \quad \text{with} \quad x \geq 0.$$

1. Analyze the decision rule in this model.
2. Can an informational cascade take place with positive investment? Can there be an informational cascade with no investment?
3. Show that there is a strictly positive probability of under-investment in the following sense: when the state is ba ($\theta = 0$), agents eventually do not invest, which is the right action: when the state is good ($\theta = 1$), there can also be a cascade with no investment.

EXERCISE 4.4. Confounded learning, Smith and Sørensen, 2001

There is a countable population of agents. A fraction α of this population is of type A and the others are of type B. In period t, agent t chooses between action 1 and action 0. There are two states of nature, 1 and 0. The actions' payoffs are specified in the following table.

Type A

	$x = 1$	$x = 0$
$\theta = 1$	1	0
$\theta = 0$	0	u_A

Type B

	$x = 1$	$x = 0$
$\theta = 1$	0	u_B
$\theta = 0$	1	0

Each agent has a SBS with precision p (on the state θ) which is independent of his type. Let μ be the belief of an agent about state 1: $\mu = P(\theta = 1)$.

1. Show that an agent of type A takes action 1 if and only if he has a belief μ such that $\mu > (1 - \mu)u_A$. When does a type B take action 1?
2. Let λ be the public LLR between state 1 and state 0. Use a figure similar to the figure in the text for the representation of the evolution of the public belief.
3. Using the figure, illustrate the following cases:
   (i) an informational cascade where all agents take action 1.
   (ii) an informational cascade where all agents take action 0.
   (iii) an informational cascade where agents A take action 1, agents B take action 0.

EXERCISE 4.5.

The exercise analyzes a first example with social learning where the payoff of action depends not only on the state of nature but also on the actions of others (through a payoff externality). Consider an investment project that requires two units of investment. The cost of each unit of investment is $c \in (0, 1)$. The payoff of the investment is 2θ where $\theta \in \{0, 1\}$ is the state of nature The prior probability of the good state is μ with $\mu < c$. Take $\mu = 1/2$ and $c = 0.55$

1. An agent receives two independent symmetric binary signal with precision q. Take $q = 3/4$. (The exact value is not important for the problem, but a given numerical value facilitates the writing of the answer.) Compute the probability that the investment is made in each of the two states of nature and the expected payoff of the agent before he receives any signal.

2. It is now assumed that each unit of investment is made by a different agent, agent 1 first, and then agent 2. Each agent receives one independent symmetric binary signal on θ with the same precision q as before. Agent 2 cannot observe the signal of agent 1 but observes the action of agent 1 (which is 1 or 0). Remember that the project pays off only in the good state and if both agents invest. Determine the probability of the realization of the project in this setting. Compare with the solution in the first setting and discuss. Extend your discussion beyond the exercise which should be only a step stone for more general remarks.

EXERCISE 4.6. ("Hot Money", Chari and Kehoe, 2001)

The exercise expands on the previous one. Consider a small open economy in which a government borrows from foreign lenders to fund a project. There are M risk-neutral agents who are ordered in an exogenous sequence. Agent there are N agents who make the investment. There are two states for the developing country, $\theta = 0$ or 1. The prior probability of the good state is μ_0. Each loan pays a return 1 if the project is funded, after M periods, *and* the state of the economy is good ($\theta = 1$). Each agent has a symmetric binary signal with precision q about θ. The cost of making an investment, for each agent, is $c, (0 < c < 1)$. Each agent i observes the actions of agents j with $j < i$.

Define $\mu^* = c$. Nature's probability of state 1 is μ_0. By assumption,

$$\frac{1-q}{q}\frac{\mu_0}{1-\mu_0} < \frac{c}{1-c} < \frac{\mu_0}{1-\mu_0}.$$

1. Assume $N = 3$ and $M = 5$. Analyze the equilibrium. (Show that if there is no herding, agents with a good signal invest and that agents with a bad signal do not invest. Note that the sequence $(0, 1, 0, 1, 0)$ does not lead to funding).

2. Show the same property for $N = 2M - 1$ for any M.

EXERCISE 4.7. Discontinuity of the Markov process of social learning

Take the standard model of Section 4.2 where the investment cost is $1/2$ (with payoff $(E[\theta]-1/2)x$), and each agent has a SBS with precision drawn from the uniform distribution on $(1/2, 1)$. Each agent knows his precision, but that precision is not observable by others.

1. Determine explicitly the Markov process defined by (**??**) when $\theta = 0$.

2. Show that 0 is the unique fixed point in μ if $\theta = 0$.

3. Show that $B(\cdot, 1)$ is not continuous in the first argument at the fixed point $\mu = 0$, and that therefore the partial derivative of B with respect to the second argument does not exist at the fixed point.

4. From the previous question, show that the condition of Theorem 4 in Smith and Sørensen (2001) does not apply to the standard model of social learning with discrete actions.

5. Assume that in each period, with probability $\alpha > 0$, the agent is a noise agent who invests with probability $1/2$. With probability $1-\alpha$, the agent is of the rational type described before. The type of the agent is not publicly observable. Is your answer to Question 3 modified?

Solution to Exercise 4.4

1. All the agents of type A take action 1. They are herding. Social learning takes place only because agents B choose $x = 1$ if and only if they have a signal 0. Let x be the action taken by agent 1 who is of type A with probability α and of type B with probability $1 - \alpha$. We have
$$\phi(0) = \frac{P(x = 0|\theta = 1)}{P(x = 0|\theta = 0)} = \frac{q}{1 - q}, \quad \phi(1) = \frac{P(x = 1|\theta = 1)}{P(x = 1|\theta = 0)} = \frac{\alpha + (1 - \alpha)(1 - q)}{\alpha + (1 - \alpha)q}.$$
Social learning
$$\lambda_{t+1} = \lambda_t + Log\big(\phi(x_t)\big).$$

2. In case 2, the precision q is sufficiently large and the support contains the interval (γ_A, γ_B)[11]. No agent is herding.
$$\phi(1) = \frac{\alpha q + (1 - \alpha)(1 - q)}{\alpha(1 - q) + (1 - \alpha)q} = \frac{1}{\phi(0)}.$$
The observation of x is informative. One verifies that $\phi(1) > 1$ if and only if $\alpha > 1/2$. The intuition for the inequality is straightforward.

3. In case 3, all agents are herding, but they do not take the same action: the agents of type A choose action 1 and the agents of type B choose action 0. In any period, the probability of observing $x = 1$ is α.

4. In case 4, all agents are herding on the same action 1. There is no herding on the action 0, because $\gamma_B > \gamma_A$. Herding on the action 0 may occur if $\gamma_A > \gamma_B$ and the precision q is sufficiently small.

[11] A necessary condition is $2Log(q/(1 - q)) > Log(u_B) - Log(u_A)$.

4.5 Appendix

4.5.1 The asymptotic rate of convergence is zero

When beliefs are bounded, they may converge to an incorrect value with a wrong herd. The issue of convergence speed makes sense only if beliefs are unbounded. This section provides a general analysis of the convergence in the binary model. Without loss of generality, we assume that the cost of investment is $c = 1/2$.

Suppose that the true state is $\theta = 0$. The public belief μ_t converges to 0. However, as $\mu_t \to 0$, there are fewer and fewer agents with a sufficiently high belief who can go against the public belief if called upon to act. Most agents do not invest. The probability that an investing agent appears becomes vanishingly small if μ tends to 0 because the density of beliefs near 1 is vanishingly small if the state is 0. It is because no agent acts contrary to the herd, although there could be some, that the public belief tends to zero. But as the probability of contrarian agents tends to zero, the social learning slows down.

Let f^1 and f^0 be the density functions in states 1 and 0. From the proportional property (Section **??**), they satisfy

$$f^1(\mu) = \mu\phi(\mu), \;\; f^0(\mu) = (1-\mu)\phi(\mu), \tag{4.6}$$

where $\phi(\mu)$ is a function. We will assume, without loss of generality, that this function is continuous.

If $\theta = 0$ and the public belief converges to 0, intuition suggests that the convergence is fastest when a herd takes place with no investment. The next result which is proven in the Appendix characterizes the convergence in this case.

PROPOSITION 4.5. *Assume the distributions of private beliefs in the two states satisfy (4.6) with $\phi(0) > 0$, and that $\theta = 0$. Then, in a herd with $x_t = 0$, if $t \to \infty$, the public belief μ_t satisfies asymptotically the relation*

$$\frac{\mu_{t+1} - \mu_t}{\mu_t} \approx -\phi(0)\mu_t,$$

and μ_t converges to 0 like $1/t$: there exists $\alpha > 0$ such that if $\mu_t < \alpha$, then $t\mu_t \to a$ for some $a > 0$.

If $\phi(1) > 0$, the same property applies to herds with investment, mutatis mutandis.

The previous result shows that in a herd, the asymptotic rate of convergence is equal to 0.

The domain in which $\phi(\mu) > 0$ represents the support of the distribution of private beliefs. Recall that the convergence of social learning is driven by the agents with extreme beliefs. It is therefore important to consider the case where the densities of these agents are not too small. This property is embodied in the inequalities $\phi(0) > 0$ and $\phi(1) > 0$. They represent a property of a *fat tail* of the distribution of private beliefs. If $\phi(0) = \phi(1)$, we will say that the distributions of private beliefs have *thin tails*. The previous proposition assumes the case of fat tails which is the most favorable for a fast convergence.

We know from Theorem 4.1 that a herd eventually begins with probability 1. Proposition 4.5 characterized the rate of convergence in a herd and it can be used to prove the following result[12].

THEOREM 4.2. *Assume the distributions of private beliefs satisfy (4.6) with $\phi(0) > 0$ and $\phi(1) > 0$. Then μ_t converges to the true value $\theta \in \{0, 1\}$ like $1/t$.*

The benchmark: learning with observable private beliefs

When agents observe beliefs through actions, there is a loss of information which can be compared with the case where private beliefs are directly observable. In Section **??**, the rate of convergence is shown to be exponential when agents have binary private signals. We assume here the private belief of agent t is publicly observable. The property of exponential convergence in Section **??** is generalized by the following result.

PROPOSITION 4.6. *If the belief of any agent t is observable, there exists $\gamma > 0$ such that $\mu_t = e^{-\gamma t} z_t$ where z_t tends to 0 almost surely.*

The contrast between Theorem 4.2 and Proposition 4.6 shows that the social learning through the observation of discrete actions is much slower, "exponentially slower[13]", than if private informations were publicly observable.

Proofs

[12] See Chamley (2002).

[13] Smith and Sørensen (2001) provide a technical result (Theorem 4) which states that the Markov process defined in (**??**) exhibits exponential convergence of beliefs to the truth under some differentiability condition. Since the result is in a central position in a paper on social learning, and they provide no discussion about the issue, the reader who is not very careful may believe that the convergence of beliefs is exponential in models of social learning. Such a conclusion is the very opposite of the central conclusion of all models of learning from others' actions. The ambiguity of their paper on this core issue is remarkable. Intuition shows that beliefs cannot converge exponentially to the truth in models of social learning. In all these models, the differentiability condition of their Theorem 4 is not satisfied (Exercise 4.7).

Proposition ??

Let $\underline{\mu}$ and $\bar{\mu}$ be the lower and upper bounds of the distribution of beliefs in period 1. We assume that if $\underline{\mu} < \mu < \bar{\mu}$, then $F_1^{\theta_1}(\mu) < F_1^{\theta_0}(\mu)$. This property holds for any period. By the Martingale Convergence Theorem, λ_t converges to some value λ_∞ almost surely. By contradiction, assume $\lambda_\infty \in (\gamma - \delta, \gamma + \delta)$. Since $F_t^{\theta_1}(\lambda_\infty) < F_t^{\theta_0}(\lambda_\infty)$, there exist $\epsilon > 0$ and $\alpha > 0$ such that if $|\lambda - \lambda_\infty| < \epsilon$, then

$$Log\Big(\frac{1 - F_t^{\theta_1}(\lambda)}{1 - F_t^{\theta_0}(\lambda)}\Big) > \alpha, \quad \text{and} \quad Log\Big(\frac{F_t^{\theta_1}(\lambda)}{F_t^{\theta_0}(\lambda)}\Big) < \alpha.$$

Since $\lambda_t \to \lambda_\infty$, there is T such that if $t > T$, $|\lambda_t - \lambda_\infty| < \alpha/3$. Take $t > T$. If $x_t = 1$, then by Bayes' rule in (4.2), $\lambda_{t+1} > \lambda_t + \alpha$, which is impossible since $\lambda_t - \lambda_{t+1} < 2\alpha/3$. A similar contradiction arises if $x_t = 0$. □

Proposition 4.5

An agent chooses action 0 (he does not invest) if and only if his belief $\tilde{\mu}$ is smaller than 1/2, *i.e.* if his private belief is smaller than $1 - \mu$, where μ is the public belief. In state θ, the probability of the event $x = 0$ is $F^\theta(1 - \mu)$. Since $F^1(\mu) < F^0(\mu)$, the observation $x = 0$ is more likely in state 0. It is "bad news" and induces the lowest possible public belief at the end of the period. The sequence of public beliefs in a herd with no investment satisfies

$$\mu_{t+1} = \frac{\Big(1 - \int_{1-\mu_t}^1 f^1(\nu)d\nu\Big)\mu_t}{\Big(1 - \int_{1-\mu_t}^1 f^1(\nu)d\nu\Big)\mu_t + \Big(1 - \int_{1-\mu_t}^1 f^0(\nu)d\nu\Big)(1 - \mu_t)}. \tag{4.7}$$

Taking an approximation for small μ_t,

$$\mu_{t+1} \approx \frac{\Big(1 - f^1(1)\mu_t\Big)\mu_t}{\Big(1 - f^1(1)\mu_t\Big)\mu_t + \Big(1 - f^0(1)\mu_t\Big)(1 - \mu_t)}.$$

Using the condition of the proposition for the initial beliefs,

$$\frac{\mu_{t+1} - \mu_t}{\mu_t} \approx (f^0(1) - f^1(1))\mu_t = -\phi(0)\mu_t.$$

For the second part of the result, we use the previous approximation and consider the sequence $\{z_k\}$ defined by

$$z_{k+1} = z_k - a z_k^2. \tag{4.8}$$

This sequence tends to 0 like $1/k$. Let y_k be such that $z_k = (1 + y_k)/(ak)$. By substitution in (4.8),

$$1 + y_{k+1} = (k+1)\Big(\frac{1 + y_k}{k} - \frac{(1 + y_k)^2}{k^2}\Big).$$

A straightforward manipulation[14] shows that $y_{k+1} < y_k$. Hence z_k tends to 0 like $1/k$ when $k \to \infty$.$\square$

Proposition 4.6

The evolution of the public belief is determined by Bayes' rule in LLR:

$$\lambda_{t+1} = \lambda_t + \zeta_t, \quad \text{with} \quad \zeta_t = Log(\hat{\mu}_t/(1-\hat{\mu}_t)) \tag{4.9}$$

Since $\theta = 0$, the random variable ζ_t has a bounded variance and a strictly negative mean, $-\bar{\gamma}$, such that

$$\bar{\gamma} = -\int_0^1 Log\Big(\frac{\nu}{1-\nu}\Big) f^0(\nu)d\nu > 0. \tag{4.10}$$

Choose γ such that $0 < \gamma < \bar{\gamma}$. Let $\nu_t = \lambda_t + \gamma t$. We have $\nu_{t+1} = \nu_t + \zeta'_t$ with $E[\zeta'_t] = -(\bar{\gamma} - \gamma) < 0$. Therefore, $\nu_t = \nu_0 + \sum_{k=1}^{t-1} \zeta'_k$ where $\sum_{k=1}^{n} \zeta'_k/n$ tends to $-(\bar{\gamma} - \gamma) < 0$ almost surely. Hence, $\sum_{k=1}^{t-1} \zeta'_k$ tends to $-\infty$ almost surely. Therefore, ν_t tends to $-\infty$ and e^{ν_t} tends to 0, almost surely. By definition of ν_t, $\mu_t \leq e^{-\gamma t} e^{\nu_t}$.

$\square$

Theorem 4.3

A herd takes place after period t if $x_{t+k} = 0$ for any $k \geq 1$. The complement of this event is contained in the union of the events A_k where A_k is defined as the herd's stop in period $t+k$ with the history $(x_{t+1} = 0, \dots, x_{t+k-1} = 0, x_{t+k} = 1)$. The probability of that event, conditional on the state $\theta = 0$, is

$$P(A_k) = (1-\pi_t)\dots(1-\pi_{t+k-1})\pi_{t+k} \leq \pi_{t+k},$$

$$\text{with} \quad \pi_{t+k} = \int_{1-\underline{\mu}_{t+k}}^1 f^0(\nu)d\nu,$$

and where $\underline{\mu}_{t+k}$ is the path of beliefs generated in a herd with no investment (Proposition 4.5). Using the proportional property (**??**), $f^0(\nu) \approx \nu f^1(1)$ for $\nu \approx 0$. Therefore, when μ_t is near 0,

$$\pi_{t+k} \approx \frac{f^1(1)}{2}\underline{\mu}_{t+k}^2 \approx \frac{a}{(t+k)^2} \qquad \text{for some constant } a.$$

The probability of the union of the A_k is smaller than the sum of the probabilities $P(A_k)$ which is of the order of $\sum_{k\geq 0} 1/(t+k)^2$, *i.e.*, of the order of $1/t$. Hence, the probability that a herd is broken once after date t tends to 0 like $1/t$.

[14]
$$1 + y_{k+1} = 1 + \frac{1}{k} - \frac{1}{k} - \frac{1}{k^2} + y_k + \frac{y_k}{k} - 2y_k\frac{k+1}{k^2} - y_k^2\frac{k+1}{k^2} < 1 + y_k.$$

The key step here is not that the belief μ_t tends to zero at a constant (strictly positive) rate, as alleged in Smith and Sørensen (2001), but that the probability that a contrarian agent shows up at date t tends to 0 like $1/t^2$. The square term arises because of condition (4.6): the integral of beliefs above $1-\mu$ is of the order of the area of a triangle proportional to μ if $\mu \to 0$.

Let $\mathcal{C}$ be the set of histories in which the public belief μ_t tends to zero. The complement of $\mathcal{C}$ is the intersection of the sets $\mathcal{A}_m = \cup_{k\geq m} A_k$ for all m. From the previous computation, $P(\mathcal{A}_m)$ tends to zero like $1/m$ and the sequence $\mathcal{A}_m$ is monotone decreasing. It follows that a herd begins almost surely. Furthermore, the probability that μ_t is different from the sequence of most pessimistic beliefs after date t, $\underline{\mu}_t = B(\underline{\mu}_{t-1}, 0)$, tends to 0 like $1/t$.

$\square$

4.5.2 Why do herds occur?

Herds must eventually occur as shown in Theorem 4.1. The proof of that result rests on the Martingale Convergence Theorem: the break of a herd induces a large change of the beliefs which contradicts the convergence. Lones Smith has insisted, quite rightly, that one should provide a direct proof that herds take place for sure eventually. This is done by computing the probability that a herd is broken in some period after time t. Such a probability tends to zero as shown in the next result.

THEOREM 4.3. *Assume the distributions of private beliefs satisfy (4.6) with $\phi(0) > 0$ and $\phi(1) > 0$. Then the probability that a herd has not started by date t tends to 0 like $1/t$.*

4.5.3 Discrete actions and the slow convergence of beliefs

The assumption of a "fat tail" of the distribution of beliefs, $\phi(0) > 0, \phi(1) > 0$, is easy to draw mathematically but it is not supported by any strong empirical evidence.

The thinner the tail of the distribution of private beliefs, the slower the convergence of social learning. However, if private signals are observable, the convergence is exponential for any distribution. The case of a thin tail provides a transition between a distribution with a thick tail and a bounded distribution where the convergence stops completely in finite time, almost surely (Chamley, 2002).

It is reasonable to consider the case where the density of beliefs is vanishingly small when the belief approaches perfect knowledge. We make the following assumption. For some $b > 0$, $c > 0$,

$$f^1(1) = 0, \quad \text{and} \quad Lim_{\mu \to 0}\Big(f^1(\mu)/(1-\mu)^b\Big) = c > 0. \tag{4.11}$$

The higher is b, the thinner is the tail of the distribution near the truth. One can show that the sequence of beliefs with the history of no investment tends to 0 like $1/t^{1/(1+b)}$ (Exercise **??**).

The main assumption in this chapter is, as emphasized in BHW, that actions are discrete. To simplify, we have assumed two actions, but the results could be generalized to a finite set of actions. The discreteness of the set of actions imposes a filter which blurs more the information conveyed by actions than the noise of the previous chapter where agents could choose action in a continuum. Therefore, the reduction in social learning is much more significant in the present chapter than in the previous one.

Recall that when private signals can be observed, the convergence of the public belief is exponential like $e^{-\alpha t}$ for some $\alpha > 0$. When agents choose an action in a continuum and a noise blurs the observation, as in the previous chapter, the convergence is reduced to a process like $e^{-\alpha t^{1/3}}$. When actions are discrete, the convergence is reduced, at best, to a much slower process like $1/t$. If the private signals are Gaussian, (as in the previous chapter), the convergence is significantly slower as shown in the example of Figure **??**. The fundamental insight of BHW is robust.

4.6 Bibliographical notes

Social learning in a changing world

Throughout this chapter and the next, the state of nature is invariant. This assumption is made to focus on the learning of a given state and it applies when the state does not change much during the phase of learning. Assume now, following Moscarini, Ottaviani and Smith (1998)that the value of θ switches between θ_0 and θ_1 according to a random Markov process: the set of states of nature $\Theta = \{\theta_0, \theta_1\}$ is fixed but between periods, θ switches to the other value with probability π.

Suppose that all agents are herding in period t. Does the public belief stay constant as in the previous sections of this chapter? Agents learn nothing from the observation of others, but they know that θ evolves randomly. Ignoring the actions from others, the public belief (probability of state θ_1) regresses to the mean, $1/2$. Therefore, after a finite number of periods, the public belief does not dominate the belief of some agents in which case not all agents herd. The herding by all agents stops. This property is interesting only if π is neither too small nor too high: if π is very small, the regression to the mean is slow and the herding behavior may last a long time; if π is sufficiently large, the expectation of the exogenous change between periods is so large that the learning from others' actions which is driven by their information about past values of θ bears no relation with the current value of θ. No cascade can occur.

Experiments

The BHW model has been experimented in the laboratory by Anderson and Holt (1996), (1997). Such experiments raise the issues of the actual understanding of Bayesian inference by people (Holt and Anderson, 1996), and of the power of the tests. A important difficulty is to separate the rational Bayesian learning from *ad hoc* rules of decision making after the observations of others' actions (such as counting the number of actions of a given type in history, or taking into account the last observed action)[15]. Huck and Oechssler (1998) find that the tests of Anderson and Holt are not powerful against simple rules. More recent

[15]This issue is raised again in empirical studies on the diffusion of innovations (Section **??**).

experimental studies include Çelen and Kariv (2002b), (2002c), or Holt (2001).

Chapter 5

Delays

Does the waiting game end with a bang or a whimper?

> Each agent chooses when to invest (if at all) and observes the number of investments by others in each period. That number provides a signal on the private information of other agents about the state of nature. The waiting game has in general multiple equilibria. An equilibrium depends on the intertemporal arbitrage between the opportunity cost of delay and the value of the information that is gained from more observations. The informational externality generates strategic substitutabilities and complementarities. Multiple equilibria appear which exhibit a rush of activity or delays, and generate a low or high amount of information. The convergence of beliefs and the occurrence of herds are analyzed under a variety of assumptions about the boundedness of the distribution of private beliefs, the number of agents, the existence of an observation noise, the length of the periods, and the discreteness of investment decisions.

In 1993, the US economy was in a shaky recovery from the previous recession. The optimism after some good news was dampened by a few bad news, raised again by other news, and so on. In the trough of the business cycle, each agent is waiting for some "good news" about an upswing. What kind of news? Some count occupancy rates in the first class section of airplanes. Others weigh the newspapers to evaluate the volume of ads. Housing starts, expenditures on durables are standard indicators to watch. The news are the actions of other agents. Everyone could be waiting because everyone is waiting in an "economics of wait and see" (Sylvia Nasar, 1993).

In order to focus on the problem of how a recession may be protracted by the waiting game for more information, we have to take a step back from the intricacies of the real world and the numerous channels of information. In this chapter, agents learn from the observation of the choices of action taken by others but not from the payoffs of these actions. This assumption is made to simplify the analysis. It is also justified in the context of the business cycle where lags between the initiation of an investment process and its payoff can be long (at least a year or two). The structure of the model is thus the same as in Chapter 3 but each agent can make his investment in any period: he has one option to make a fixed size investment. The central issue is when to exercise the option, if at all.

When the value of the investment is strictly positive, delay is costly because the present value of the payoff is reduced by the discount factor. The *opportunity cost of delay* for one period is the product of the net payoff of investment and the discount rate. Delay enables an agent to observe others' actions and infer some information on the state of nature. These observations may generate good or bad news. Define the bad news as an event such that the agent regrets *ex post* an irreversible investment which he has made, and would pay a price to undo it (if it were possible). The expected value of this payment in the next period after observing the current period's aggregate investment, is the option value of delay. The key issue which commands all results in this chapter is the trade-off, in equilibrium, between the opportunity cost and the option value of delay.

Consider the model of Chapter *** with two states of nature and assume that agents can choose the timing of their investment. If all beliefs (probability of the good state) are below the cost of investment, the only equilibrium is with no investment and there is a herd as in the BHW model. If all beliefs are higher than the cost of investment, there is an equilibrium in which all agents invest with no delay. This behavior is like a herd with investment in the BHW model and it is an equilibrium since nothing is learned by delaying. The herds in the BHW model with exogenous timing are equilibria in the model with endogenous timing.

However, the model with endogenous timing may have other equilibria with an arbitrage between the option value and the opportunity cost of delay. For a general distribution of private beliefs, the margin of arbitrage may occur at different points of the distribution. Generically, there are at least two equilibrium points, one in the upper tail of the distribution and another in the lower tail. In the first equilibrium, only the most optimistic agents invest; in the second, only the most pessimistic delay. The two equilibria in which most agents delay or rush, respectively, are not symmetric because of the arbitrage mechanism. In the first, the information conveyed by the aggregate activity must be large in order to keep the agents at the high margin of beliefs (with a high opportunity cost) from investing. In the second, both the opportunity cost of relatively pessimistic agents and the information conveyed by the aggregate activity are low. In the particular case of a bounded distribution, the rush where few agents delay may be replaced by the corner solution where no agent delays.

Multiple equilibria are evidence of strategic complementarities (Cooper and John, 1988). These complementarities arise here only because of informational externalities. There is no payoff externality. As in other models with strategic complementarities, multiple equilibria may provide a support for sudden switches of regime with large fluctuations of economic activity (Chamley, 1999).

The main ideas of the chapter are presented in Section 5 with a simple two-agent model based on Chamley and Gale (1994). The unique equilibrium is computed explicitly.

The general model with heterogeneous beliefs is presented in Section 5. It is the full extension of the BHW model to endogenous timing. Heterogeneous beliefs is a plausible assumption *per se* and it generates non random strategies. The model has a number of players independent of the state of nature and generalizes Chamley and Gale (1994) who assume identical beliefs. In the model with identical beliefs, the endowment of an option is the private signal and the number of players thus depends on the state of nature. This case is particularly relevant when the number of players is large.

When private beliefs are not identical, the analysis of the symmetric sub-game perfect Bayesian equilibria (PBE) turns out to be simple due to an intuitive property which is related to the arbitrage condition: an agent never invests before another who is more optimistic. Therefore, the agent with the highest belief among those who delay must be the "first" to invest in the next period if there is any investment in that period (since he has the highest belief then). All equilibria where the arbitrage condition applies can be described as sequences of two-period equilibria.

Some properties of the model are presented in Section **??**. Extensions will be discussed in the next chapter. When the public belief is a range (μ^*, μ^{**}), the level of investment in each period is a random variable and the probability of no investment is strictly positive. If there is no investment, the game stops with a herd and no investment takes place in any subsequent period. Hence the game lasts a number of periods which is at most equal to the number of players in the game. If the period length tends to zero, the game ends in a vanishingly short time. Since an agent can always delay until the end of the game, and the cost of delay tends to zero with the length of the period, the information generated by the game also tends to zero with the period length: another effect of arbitrage.

The game is illustrated in Section **??** by an example with two agents with normally distributed private signals (unbounded), which highlights the mechanism of strategic complementarity. When the time period is sufficiently short, there cannot be multiple equilibria, under some specific conditions. The presence of time lags between observation and action is thus necessary for the existence of multiple equilibria.

The case of a large number of agents (Section **??**) is interesting and illustrates the power of the arbitrage argument. When the number of agents tends to infinity, the distribution of the levels of investment tends to a Poisson distribution with a parameter which depends on the public belief, and on the discount rate. This implies that as long as the public belief μ is in the interval (μ^*, μ^{**}), the level of investment is a random variable which is small compared to the number of agents. The public belief evolves randomly until it exits the interval: if $\mu < \mu^*$, investment goes from a small random amount to nil forever; if $\mu > \mu^{**}$, all remaining agents invest with no further delay. The game ends with a whimper or a bang.

The Appendix presents two extensions of the model which show the robustness of the results: (i) with a very large number of agents (a continuum) and an observation noise, there are multiple equilibria as in the model with two agents; the equilibrium with high aggregate activity generates an amount of information which is significantly smaller than the equilibrium with low activity and delays; (ii) multiple equilibria also appear when individual investments are non-discrete.

The simple model is another example of how to start the analysis of general issues as presented in the introduction. One should stylize as much as possible. The investigation of robustness and extensions will be easier once the base model is firmly understood.

5.1 The simplest model

There are two players and time is divided in periods. There are two states of nature, $\theta \in \{0, 1\}$. In state 0, only one of two players (chosen randomly with equal probability) has one option to make an investment of a fixed size in any period. In state 1, both players have one option. To have an option is private information and is not observable by the other agent. Here, the private signal of the agent is the option. The number of players in the game depends on the state of nature[1]. As an illustration, the opportunities for productive investment may be more numerous when the state of the economy is good.

[1]One could also think that the cost of investment is very high for one or zero agent thus preventing the investment. Recall that in the BHW model, the number of players does not depend on the state of nature.

For an agent with an option, the payoff of investment in period t is

$$U = \delta^{t-1}(E[\theta] - c), \quad \text{with} \quad 0 < c < 1,$$

where E is the expectation conditional on the information of the agent and δ is the discount factor, $0 < \delta < 1$.

All agents in the game have the same private information (their own option), and observe the same history. They have the same belief (probability of state $\theta = 1$). Let μ_t be the belief of an agent at the beginning of period t. The belief in the first period is given[2] and satisfies the next assumption in order to avoid trivialities.

Assumption 5.1. $\quad 0 < \mu - c < \delta\mu(1 - c).$

Agents play a game in each period and the strategy of an agent is his probability of investment. We look for a symmetric perfect Bayesian equilibrium (PBE): each agent knows the strategy z of the other agent (it is the same as his own); he anticipates rationally to receive a random amount of information at the end of each period and that the subgame which begins next period with a belief updated by Bayes' rule has an equilibrium.

Let z be the probability of investment in the first period by an agent with an option. Such an agent will be called a player. We prove that there is a unique symmetric equilibrium with $0 < z < 1$.

- $z = 1$ cannot be an equilibrium. If $z = 1$, both agents "come out" with probability one, the number of players and therefore the state is revealed perfectly at the end of the period. If an agent deviates from the strategy $z = 1$ and delays (with $z = 0$), he can invest in the second period if and only if the true state is good. The expected payoff of this delay strategy is $\delta\mu(1-c)$: in the first period, the good state is revealed with probability μ in which case he earns $1-c$. The discount factor is applied because the investment is made in the second period. The payoff of no delay is $\mu - c$, and it is smaller by Assumption 5.1. The strategy $z = 1$ cannot define a PBE. Note that the interpretation of the right-hand side inequality is now clear: the payoff of investment, $\mu - c$, should be smaller than the payoff of delay with perfect information in the next period.
- $z = 0$ cannot be an equilibrium either. The argument is a bit more involved and proceeds by contradiction. If $z = 0$, there is no investment in the first period for

[2] One could assume that agents know that nature chooses state $\theta = 1$ with probability μ_0. In this case, by Bayes' rule, $\mu = 2\mu_0/(1 + \mu_0)$.

any state, no information and therefore the same game holds at the beginning of period 2, with the same belief μ. Indefinite delay cannot be an equilibrium strategy because it would generate a zero payoff which is strictly smaller than the payoff of no delay, $\mu - c > 0$ (Assumption 5.1). Let T be the first period in which there is some investment with positive probability. Since $z = 0$, $T \geq 2$. In period T, the current value of the payoff of investment is $\mu - c > 0$ because nothing has been learned before. The present value of this payoff is strictly smaller than the payoff of immediate investment, $\mu - c$. Hence, $T \geq 2$ is impossible and $z = 0$ cannot be an equilibrium strategy.

The necessity of investment in every period

We have shown that in an equilibrium, agents randomize with $0 < z < 1$. The level of total investment is a random variable. We will see that the higher the level of investment, the higher the updated belief after the observation of the investment. In this simple model, one investment is sufficient to reveal to the other player (if there is one), that the state is good. No investment in the first period is bad news. Would anyone invest in the second period after this bad news? The answer is no, and the argument is interesting.

If anyone delays in the first period and expects to invest in the second period after the worst possible news (zero investment), his payoff in the subgame of period 2 is the same as that of investing for sure in period 2. (He invests if he observes one investment). That payoff, $\delta(\mu - c)$, is inferior to the payoff of immediate investment because of the discount. The player cannot invest after observing no investment. Hence, *if there is no investment in the first period, there is no investment in any period after*. We will see in this chapter that this property applies in more general models. The argument shows that: (i) if there is no investment, the *ex post* belief of any agent must be smaller than the cost of investment c; (ii) since agents randomize in the first period, the event of no investment has a positive probability. There is a positive probability of an incorrect herd.

Using the previous argument, we can compute the payoff of delay. If an agent delays, he invests in period 2 if and only if he sees an investment (by the other agent) in period 1, in which case he is sure that the state is good and his second period payoff is $1 - c$. The probability of observing an investment in the first period is μz, (the product of the probability that there is another agent and that he invests). The payoff of delay (computed at the time of the decision) is therefore $\delta \mu z(1 - c)$.

Arbitrage and the existence of a unique PBE

Since $0 < z < 1$, agents randomize their investment in the first period and are indifferent between no delay and delay. This arbitrage condition between the value of investment and the value of the option to invest is essential in this chapter and is defined by

$$\mu - c = \delta\mu z(1-c). \tag{5.1}$$

By Assumption 5.1, this equation in z has a unique solution in the interval $(0,1)$. The analysis of the solution may be summarized as follows: first, the arbitrage condition is necessary if a PBE exists; second, the existence of a unique PBE follows from the arbitrage condition by construction of the equilibrium strategy. This method will be used in the general model.

Interpretation of the arbitrage condition

A simple manipulation shows that the arbitrage equation can be restated as

$$\begin{aligned} \frac{1-\delta}{\delta}(\mu - c) &= \Big(\mu z(1-c) - (\mu - c)\Big) \\ &= P(x=0|\mu)\Big(c - P(\theta_1|x=0,\mu)\Big) \end{aligned} \tag{5.2}$$

where $P(x=0|\mu)$ is the probability for an agent with belief μ that the other agent does not invest in period 1, *i.e.* the probability of bad news. The term $\mu - c$ has the dimension of a stock, as the net present value of an investment. The left-hand side is the *opportunity cost of delay*: it is the value of investment multiplied by the interest rate between consecutive periods. (If $\delta = 1/(1+r)$, then $(1-\delta)/\delta = r$). The right-hand side will be called the *information value of delay*. It provides the measurement of the value of information obtained from a delay. To interpret it, note that the term $P(\theta_1|x=0,\mu)$ is the value of an investment after the bad news in the first period. If an agent could reverse his decision to invest in the first period (and get the cost back), the associated value of this action would be $c - P(\theta_1|x=0,\mu)$. The option value of delay is the expected "regret value" of undoing the investment when the agent wishes he could do so. The next properties follow from the arbitrage condition.

In an equilibrium, the cost of delay is equal to the information value of delay---the expected regret value. This arbitrage is the linchpin of all equilibria in this chapter.

Information and time discount

The power of the signal which is obtained by delay increases with the probability of investment z in the strategy. If $z=0$, there is no information. If $z=1$, there is perfect information.

The discount factor is related to the length of the period, τ, by $\delta = e^{-\rho\tau}$, with ρ the discount rate per unit of time. If δ varies, the arbitrage equation (5.1) shows that the product δz is constant. A shorter period (higher δ) means that the equilibrium must generate less information at the end of the first period: the opportunity cost of delay is smaller and by arbitrage, the information value of delay decreases. Since this information varies with z, the value of z decreases. From Assumption 5.1, $0 < z < 1$ only if δ is in the interval $[\delta^*, 1)$, with $\delta^* = (\mu - c)/(\mu(1-c))$.

If $\delta \to \delta^*$, then $z \to 1$. If $\delta \leq \delta^*$, then $z = 1$ and the state is revealed at the end of the first period. Because this information comes late (with a low δ), agents do not wait for it.

If $\delta \to 1$ and the period length is vanishingly short, information comes in quickly but there is a positive probability that it is wrong. The equilibrium strategy z tends to δ^*. If the state is good, with probability $(1-\delta^*)^2 > 0$ both agents delay and end up thinking that the probability of the good state is smaller than c and that investment is not profitable. There is a trade-off between the period length and the quality of information which is revealed by the observation of others. This trade-off is generated by the arbitrage condition. The opportunity cost of delay is smaller if the period length is smaller. Hence the value of the information gained by delay must also be smaller.

A remarkable property is that the waiting game lasts one period, independently of the discount factor. If the period is vanishingly short, the game ends in a vanishingly short time, but the amount of information which is released is also vanishingly short. In this simple model with identical players, the value of the game does not depend on the endogenous information which is generated in the game since it is equal to the payoff of immediate investment. However, when agents have different private informations, the length of the period affects welfare (as shown in the next chapter).

Investment level and optimism

In the arbitrage equation (5.1), the probability of investment and the expected value of investment are increasing functions of the belief μ: a higher μ entails a higher opportunity cost and by arbitrage a higher option value of delay. The higher information requires that players "come out of the wood" with a higher probability z. This mechanism is different from the arbitrage mechanism in the q-theory of Tobin which operates on the margin between the financial value μ and an adjustment cost.

Observation noise and investment

Suppose that the investment of an agent is observed with a noise: if an investment is made, the other agent sees it with probability $1-\gamma$ and sees nothing with probability γ, (γ small). The arbitrage operates beautifully: the information for a delaying agent is unaffected by the noise because it must be equal to the opportunity cost which is independent of the noise. Agents compensate for the noise in the equilibrium by increasing the probability of investment (Exercise **??**).

Large number of agents

Suppose that in the good state there are N agents with an option to invest and that in the bad state there is only one agent with such an option. These values are chosen to simplify the game: one investment reveals that the state is good and no investment stops the game. For any N which can be arbitrarily large, the game lasts only one period, in equilibrium, and the probability of investment of each agent in the first period tends to zero if $N \to \infty$. Furthermore, the probability of no investment, conditional on the good state, tends to a positive number. The intuition is simple. If the probability of investment by a player remains higher than some value $\alpha > 0$, its action (investment or no investment) is an signal on the state with a non vanishing precision. If $N \to \infty$, delay provides a sample of observations of arbitrarily large size and perfect information asymptotically. This is impossible because it would contradict the arbitrage with the opportunity cost of delay which is independent of N. The equilibrium is analyzed in Exercise **??**.

Strategic substitutability

Suppose an agent increases his probability of investment from an equilibrium value z. The option value (in the right-hand side of (5.1) or (5.2)) increases. Delay becomes strictly better and the optimal response is to reduce the probability of investment to zero: there is strategic substitutability between agents. In a more general model (next section) this property is not satisfied and multiple equilibria may arise.

Non symmetric equilibrium

Assume there are two agents, A and B, who can see each other but cannot see whether the other has an option to invest. It is common knowledge that agent B always delays in the first period and does not invest ever if he sees no investment in the first period.

Agent A does not get any information by delaying: his optimal strategy is to invest with no delay, if he has an option. Given this strategy of agent A, agent B gets perfect information at the end of period 1 and his strategy is optimal. The equilibrium generates perfect information after one period. Furthermore, if the state is good, both agents invest. If the period length is vanishingly short, the value of the game is $\mu - c$ for agent A, and $\mu(1 - c)$ for agent B which is strictly higher than in the symmetric equilibrium. If agents could "allocate the asymmetry" randomly before knowing whether they have an option, they would be better off *ex ante*.

5.2 A general model with heterogeneous beliefs

The structure of the model extends the canonical model in Section 4.2 by allowing each agent to make his fixed size investment in any period of his choice. There are N agents each with one option to make one irreversible investment of a fixed size. Time is divided in periods and the payoff of exercising an option in period t is $\delta^{t-1}(\theta - c)$ with δ the discount factor, $0 < \delta \leq 1$, and c the cost of investment, $0 < c < 1$. The payoff from never investing is zero. Investment can be interpreted as an irreversible switch from one activity to another[3].

The rest of the model is the same as in the beginning of Section 4.2. The productivity parameter θ which is not observable is set randomly by nature once and for all before the first period and takes one of two values: $\theta_0 < \theta_1$. Without loss of generality, these values are normalized at $\theta_1 = 1$ for the "good" state, and $\theta_0 = 0$ for the "bad" state. As in Section **??**, each agent is endowed at the beginning of time with a private belief which is drawn from a distribution with *c.d.f.* $F_1^\theta(\mu)$ depending on the state of nature θ. For simplicity and without loss of generality, it will be assumed that the cumulative distribution functions have derivatives[4]. The support of the distribution of beliefs is an interval $(\underline{\mu}_1, \bar{\mu}_1)$ where the bounds may be infinite and are independent of θ. The densities of private beliefs satisfy the Proportional Property (**??**). Hence, the cumulative distribution functions satisfy the property of first order stochastic dominance: for any $\mu \in (\underline{\mu}_1, \bar{\mu}_1)$, $F_1^1(\mu) < F_1^0(\mu)$.

After the beginning of time, learning is endogenous. In period t, an agent knows his private belief and the history $h_t = (x_1, \dots, x_{t-1})$, where x_k is the number of investments in period k.

The only decision variable of an agent is the period in which he invests. (This period is postponed to infinity if he never invests). We will consider only symmetric equilibria. A strategy in period t is defined by the *investment set* $I_t(h_t)$ of beliefs of all investing agents: an agent with belief μ_t in period t invests in that period (assuming he still has an option) if and only if $\mu_t \in I_t(h_t)$. In an equilibrium, the set of agents which are indifferent between investment and delay will be of measure zero and is ignored. Agents will not use random strategies.

As in the previous chapters, Bayesian agents use the observation of the number of investments, x_t, to update the distribution of beliefs F_t^θ into the distribution in the next period F_{t+1}^θ. Each agent (who has an option) chooses a strategy which maximizes his expected

[3] The case where the switch involves the termination of an investment process (as in Caplin and Leahy, 1994) is isomorphic.

[4] The characterization of equilibria with atomistic distributions is more technical since equilibrium strategies may be random (*e.g.*, Chamley and Gale, 1994).

payoff, given his information and the equilibrium strategy of all agents for any future date and future history. For any period t and history h_t, each agent computes the value of his option if he delays and plays in the subgame which begins in the next period $t+1$. Delaying is optimal if and only if that value is at least equal[5] to the payoff of investing in period t. All equilibria analyzed here are symmetric subgame perfect Bayesian equilibria (PBE).

As in the model with exogenous timing (Section **??**), a belief can be expressed by the Log likelihood ratio (LLR) between the two states, $\lambda = Log(\mu/(1-\mu))$ which is updated between periods t and $t+1$ by Bayes' rule

$$\begin{aligned} \lambda_{t+1} &= \lambda_t + \zeta_t, \quad \text{where} \quad \zeta_t = Log\Big(\frac{P(x_t \mid I_t, \theta_1)}{P(x_t \mid I_t, \theta_0)}\Big), \\ \text{and } P(x_t \mid I_t, \theta) &= \frac{n_t!}{x_t!(n_t - x_t)!}\pi_\theta^{x_t}(1-\pi_\theta)^{n_t - x_t}, \;\; \pi_\theta = P(\lambda_t \in I_t \mid \theta). \end{aligned} \tag{5.3}$$

All agents update their individual LLR by adding the *same* value ζ_t. Given a state θ, the distribution of beliefs measured in LLRs in period t is generated by a translation of the initial distribution by a random variable ζ_t.

5.2.1 Characterization and existence of equilibria

The incentive for delay is to get more information from the observation of others. Agents who are relatively more optimistic have more to loose and less to gain from delaying: the discount factor applies to a relatively high expected payoff while the probability of bad news to be learned after a delay is relatively small. This fundamental property of the model restricts the equilibrium strategies to the class of *monotone strategies*. By definition, an agent with a monotone strategy in period t invests if and only if his belief μ_t is greater than some value μ_t^*. The next result, which is proven in the appendix, shows that equilibrium strategies must be monotone.

LEMMA 5.1. (monotone strategies) *In any arbitrary period t of a PBE, if the payoff of delay for an agent with belief μ_t is at least equal to the payoff of no delay, any agent with belief $\mu_t' < \mu_t$ strictly prefers to delay. Equilibrium strategies are monotone and defined by a value μ_t^*: agents who delay in period t have a belief $\mu_t \leq \mu_t^*$.*

Until the end of the chapter, strategies will be defined by their minimum belief for investment, μ_t^*. Since no agent would invest with a negative payoff, $\mu_t^* \geq c$. The support of the distribution of μ in period t is denoted by $(\underline{\mu}_t, \bar{\mu}_t)$. If all agents delay in period t, one can define the equilibrium strategy as $\mu_t^* = \bar{\mu}_t$.

[5] By assumption, an indifferent agent delays. This tie breaking rule applies with probability zero and is inconsequential.

The existence of a non trivial equilibrium in the subgame which begins in period t depends on the payoff of the most optimistic agent[6], $\bar{\mu}_t - c$. First, if $\bar{\mu}_t \leq c$, no agent has a positive payoff and there is no investment whatever the state θ. Nothing is learned in period t (with probability one), or in any period after. The game stops. Second, if $\bar{\mu}_t > c$, the next result (which parallels a property for identical beliefs in Chamley and Gale, 1994) shows that in a PBE, the probability of some investment is strictly positive. The intuition of the proof, which is given in the appendix, begins with the remark that a permanent delay is not optimal for agents with beliefs strictly greater than c (since it would yield a payoff of zero). Let T be the first period after t in which some agents invest with positive probability. If $T > t$, the current value of their payoff would be the same as in period t (nothing is learned between t and T). Because of the discount factor $\delta < 1$, the present value of delay would be strictly smaller than immediate investment which is a contradiction.

LEMMA 5.2. (condition for positive investment) *In any period t of a PBE:*

(i) if $c < \bar{\mu}_t$ (the cost of investment is below the upper-bound of beliefs), then any equilibrium strategy μ_t^* is such that $c \leq \mu_t^* < \bar{\mu}_t$; if there is at least one remaining player, the probability of at least one investment in period t is strictly positive;

(ii) if $\bar{\mu}_t \leq c$ (the cost of investment is above the upper-bound of beliefs), then with probability one there is no investment for any period $\tau \geq t$.

The decision to invest is a decision whether to delay or not. In evaluating the payoff of delay, an agent should take into account the strategies of the other agents in all future periods. This could be in general a very difficult exercise. Fortunately, the property of monotone strategies simplifies greatly the structure of equilibria. A key step is the next result which shows that any equilibrium is a sequence of two-period equilibria each of which can be determined separately.

LEMMA 5.3. (one-step property) *If the equilibrium strategy μ_t^* of a PBE in period t is an interior solution ($\underline{\mu}_t < \mu_t^* < \bar{\mu}_t$), then an agent with belief μ_t^* is indifferent between investing in period t and delaying to make a final decision (investing or not) in period $t+1$.*

Proof Since the Bayesian updating rules are continuous in μ, the payoffs of immediate investment and of delay for any agent are continuous functions of his belief μ. Therefore, an agent with belief μ_t^* in period t is indifferent between investment and delay. By definition

[6] Recall that such an agent may not actually exist in the realized distribution of beliefs.

of μ_t^*, if he delays he has the highest level of belief among all players remaining in the game in period $t+1$, *i.e.*, his belief is $\bar{\mu}_{t+1}$. In period $t+1$ there are two possibilities: (i) if $\bar{\mu}_{t+1} > c$, then from Lemma 5.2, $\mu_{t+1}^* < \bar{\mu}_{t+1}$ and a player with belief $\bar{\mu}_{t+1}$ invests in period $t+1$; (ii) if $\bar{\mu}_{t+1} \leq c$, then from Lemma 5.2 again, nothing is learned after period t; a player with belief $\bar{\mu}_{t+1}$ may invest (if $\bar{\mu}_{t+1} = c$), but his payoff is the same as that of delaying for ever. $\square$

In an equilibrium, an agent with belief μ compares the payoff of immediate investment, $\mu - c$, with that of *delay for exactly one period*, $W(\mu, \mu^*)$, where μ^* is the strategy of others. (For simplicity we omit the time subscript and other arguments such as the number of players and the *c.d.f.* F^θ). From Lemma 5.3 and the Bayesian formulae (5.3) with $\pi^\theta = 1 - F^\theta(\mu^*)$, the function W is well defined. An interior equilibrium strategy must be solution of the arbitrage equation between the payoff of immediate investment and of delay:

$$\mu^* - c = W(\mu^*, \mu^*).$$

The next result shows that this equation has a solution if the cost c is interior to the support of the distribution of beliefs.

LEMMA 5.4. *In any period, if the cost c is in the support of the distribution of beliefs, i.e., $\underline{\mu} < c < \bar{\mu}$, then there exists $\mu^* > c$ such that $\mu^* - c = W(\mu^*, \mu^*)$: an agent with belief μ^* is indifferent between investment and delay.*

Proof Choose $\mu^* = \bar{\mu}$: there is no investment and therefore no learning during the period. Hence, $W(\bar{\mu}, \bar{\mu}) = (1-\delta)(\bar{\mu} - c) < \bar{\mu} - c$. Choose now $\mu^* = c$. With strictly positive probability, an agent with belief c observes $n-1$ investments in which case his belief is higher (n is the number of remaining players). Hence, $W(c, c) > 0$. Since the function W is continuous, the equation $\mu^* - c = W(\mu^*, \mu^*)$ has at least one solution in the interval $(c, \bar{\mu})$. $\square$

The previous lemmata provide characterizations of equilibria (PBE). These characterizations enable us to construct all PBE by forward induction and to show existence.

THEOREM 5.1. *In any period t where the support of private beliefs is the interval $(\underline{\mu}_t, \bar{\mu}_t)$:*

(i) if $\bar{\mu}_t \leq c$, then there is a unique PBE with no agent investing in period t or after;

(ii) if $\underline{\mu}_t < c < \bar{\mu}_t$, then there is at least one PBE with strategy $\mu_t^* \in (c, \bar{\mu}_t)$;

(iii) if $c \leq \underline{\mu}_t$, then there is a PBE with $\mu_t^* = \underline{\mu}_t$ in which all remaining players invest in period t.

In case (ii) and (iii) there may be multiple equilibria. The equilibrium strategies $\mu_t^* \in (\underline{\mu}_t, \bar{\mu}_t)$ *are identical to the solutions of the arbitrage equation*

$$\mu^* - c = W(\mu^*, \mu^*), \tag{5.4}$$

where $W(\mu, \mu^*)$ *is the payoff of an agent with belief* μ *who delays for one period exactly while other agents use the strategy* μ^**.*

The only part which needs a comment is (ii). From Lemma 5.4, there exists μ_t^* such that $c < \mu_t^*$ and $\mu^* - c = W(\mu^*, \mu^*)$. From Lemma 5.1, any agent with belief $\mu_t > \mu_t^*$ strictly prefers not to delay and any agent with belief $\mu_t < \mu_t^*$ strictly prefers to delay. (Otherwise, by Lemma 5.1 an agent with belief μ_t^* would strictly prefer to delay which contradicts the definition of μ_t^*). The strategy μ_t^* determines the random outcome x_t in period t and the distributions F_{t+1}^{θ} for the next period, and so on.

5.3 Properties

5.3.1 Arbitrage

Let us reconsider the trade-off between investment and delay. For the sake of simplicity, we omit the time subscript whenever there is no ambiguity. If an agent with belief μ delays for one period, he foregoes the implicit one-period rent on his investment which is the difference between investing for sure now and investing for sure next period, $(1-\delta)(\mu - c)$; he gains the possibility of "undoing" the investment after bad news at the end of the current period (the possibility of not investing). The expected value of this possibility is the option value of delay. The following result, proven in the appendix, shows that the belief μ^* of a marginal agent is defined by the equality between the opportunity cost and the option value of delay.

PROPOSITION 5.1. (arbitrage) *Let* μ^* *be an equilibrium strategy in a game with* $n \geq 2$ *remaining players,* $\underline{\mu} < \mu^* < \bar{\mu}$*. Then* μ^* *is solution of the arbitrage equation between the opportunity cost and the option value of delay*

$$(1-\delta)(\mu^* - c) = \delta Q(\mu^*, \mu^*), \qquad \textit{with}$$

$$Q(\mu, \mu^*) = \sum_{k=0}^{n-1} P(x = k | \mu, \mu^*, F^{\theta}, n) Max\Big(c - P(\theta = \theta_1 | x = k; \mu, \mu^*, F^{\theta}, n), 0\Big), \tag{5.5}$$

where x *is the number of investments by other agents in the period.*

The function $Q(\mu, \mu^*)$ is a "regret function" which applies to an agent with belief μ. It depends on the strategy μ^* of the other agents and on the *c.d.f.*s F^θ at the beginning of the period. Since the gain of "undoing" an investment is c minus the value of the investment after the bad news, the regret function $Q(\mu, \mu^*)$ is the expected value of the amount the agent would be prepared to pay to undo his investment at the beginning of next period.

At the end of that period, each agent updates his LLR according to the Bayesian formula (5.3) with $\pi_\theta = 1 - F^\theta(\mu_t^*)$. A simple exercise shows that the updated LLR is an increasing function of the level of investment in period t and that the lowest value of investment $x_t = 0$ generates the lowest level of belief at the end of the period. Can the game go on after the worst news of no investment? From Proposition 5.1, we can deduce immediately that the answer is no. If the agent would invest after the worst news, the value of $Q(\mu^*, \mu^*)$ would be equal to zero and would therefore be strictly smaller than $\mu^* - c$ which contradicts the arbitrage equation (5.5).

PROPOSITION 5.2. (the case of worst news) *In any period t of a PBE for which the equilibrium strategy μ_t^* is interior to the support $(\underline{\mu}_t, \bar{\mu}_t)$, if $x_t = 0$, then $\bar{\mu}_{t+1} \leq c$ and the game stops at the end of period t with no further investment in any subsequent period.*

The result shows that a game with N players lasts at most N periods. If the period length τ is vanishingly short, the game ends in a vanishingly short time. This case is analyzed in Section **??**.

5.3.2 Representation of beliefs

An example of the evolution of beliefs is illustrated in Figure **??**. The reader may compare with the equivalent Figure **??** in the case of exogenous timing. Beliefs are measured by the LLR and are bounded, by assumption. The support of their distribution at the beginning of a period is represented by a segment. Suppose that the state is bad: $\theta = 0$. At the beginning of period 1, the private beliefs of the N players are the realizations of N independent drawings from a distribution with density $f^0(\cdot)$ which is represented by a continuous curve. (The density in state $\theta = 1$ is represented by a dotted curve).

In period 1, agents with a belief above λ_1^* exercise their option to invest. The number of investments, x_1, is the number of agents with belief above λ_1^*, which is random according to the process described in the previous paragraph.

Each agent who delays knows that x_1 is generated by the sum of $N-1$ independent binary

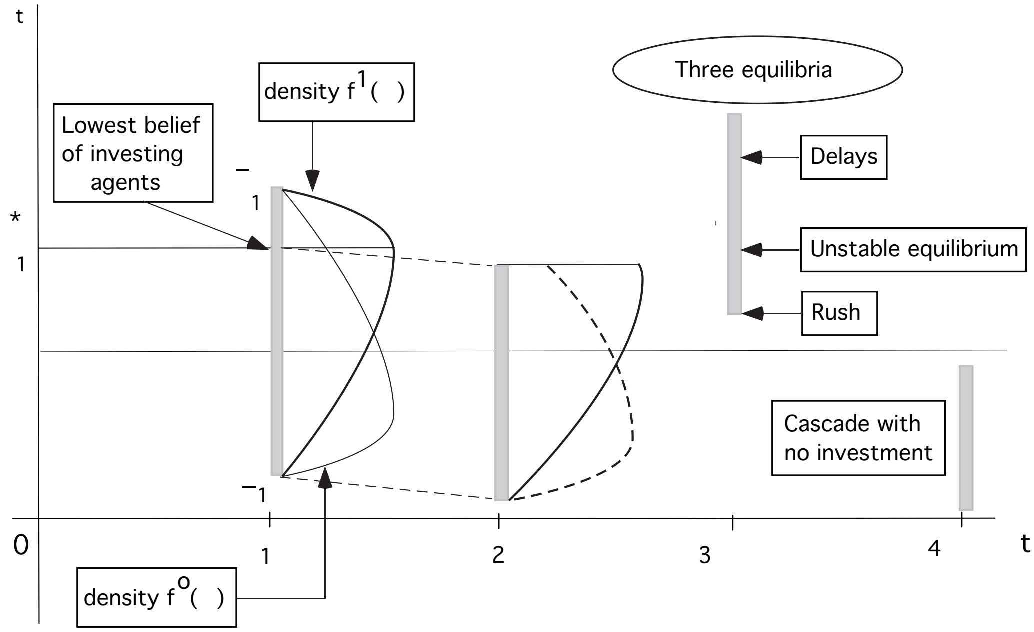

The number of investments in a period t depends on the number of agents with a belief higher than λ_t^*. At the end of a period, the updated distributions in the two states are truncated, translated and rescaled. Period 3 (in which the representation of the densities is omitted) corresponds to a case with three equilibria. In period 4, there is no investment since all beliefs are smaller than the cost of investment.

Figure 5.1: An example of evolution of beliefs

variables equal to 1 with a probability π^θ that depends on θ: $\pi^\theta = 1 - F^\theta(\lambda_1^*)$. The probability is represented in Figure **??** by the lightly shaded area if $\theta = 0$ and the darker area if $\theta = 1$.

From the updating rule (5.3), the distribution of LLRs in period 2 is a translation of the distribution of the LLRs in period 1, truncated at λ_1^*, and rescaled (to have a total measure of one): $\lambda_1^* - \underline{\lambda}_1 = \bar{\lambda}_2 - \underline{\lambda}_2$. An agent with LLR equal to λ_1^* in period 1 and who delays has the highest belief in period 2. The news at the end of period 1 depend on the random number of agents with beliefs above λ_1. In Figure **??**, the observation of the number of investments in period 1 is bad news: the agent with highest belief has a lower belief in period 2 compared to period 1.

There are two critical values for the LLR in each period: (i) an agent who has a LLR below the break-even value $\gamma = Log(c/(1-c))$ does not invest; (ii) no agent who has an LLR above some value λ^{**} delays. The value λ^{**} is defined such that if $\lambda > \lambda^{**}$, the payoff of no delay is higher than that of delay with perfect information one period later. Since the latter yields $\delta\mu(1-c)$ to an agent with belief μ, we have

$$\lambda^{**} = Log\Big(\frac{\mu^{**}}{1-\mu^{**}}\Big), \quad \text{with} \quad \mu^{**} - c = \delta\mu^{**}(1-c). \tag{5.6}$$

Note that λ^{**} (or μ^{**}) depends essentially on the discount rate. If the discount rate is vanishingly small, the opportunity cost of delay is vanishingly small and only the super-optimists should invest: if $\delta \to 1$, then $\lambda^{**} \to \infty$.

5.3.3 Herds: a comparison with exogenous sequences

Case (iii) in Theorem 5.1 is represented in period 3 of Figure **??**. The lower bound of the distribution of beliefs is higher than the cost of investment, with $\underline{\lambda}_3 > \gamma = Log(c/(1-c))$. There is an equilibrium called a rush, in which no agent delays. In that equilibrium, nothing is learned by delay since the number of investments is equal to the number of remaining players, whatever the state of nature. This outcome occurs here with endogenous delay under the same condition as the "cascade" or herd of BHW, in which all agents invest, regardless of their private signal[7].

For the distribution of beliefs in period 3, there may be another equilibrium with an interior solution λ_3^* to the arbitrage equation (5.4). Since agents with the lowest LLR $\underline{\lambda}_3$ strictly prefer to invest if all others do, there may be multiple equilibria with arbitrage, some of them unstable. This issue is reexamined in the next subsection.

For the case of period 4, all beliefs are below the break-even point: $\bar{\lambda}_4 < \gamma$. No investment takes place in period 4 or after. This equilibrium appears also in the BHW model with exogenous timing, as a cascade with no investment. From Proposition 5.2, this equilibrium occurs with positive probability if agents coordinate on the equilibrium λ_3^* in period 3.

The present model integrates the findings of the BHW model in the setting with endogenous timing. We could anticipate that the herds of the BHW model with exogenous timing are also equilibria when timing is endogenous because they generate no information and therefore no incentive for delay.

The cascades of the BHW model are also equilibria when timing is endogenous.

A rush where all agents invest with no delay can take place only if the distribution of beliefs (LLR) is bounded below. However, if beliefs are unbounded, the structure of equilibria is

[7] In the BHW model, distributions are atomistic, but the argument is identical.

very similar to that in Figure ??. In a generic sense, there are multiple equilibria and one of them may be similar to a rush. This issue is examined in an example with two agents and Gaussian signals. The Gaussian property is a standard representation of unbounded beliefs.

EXERCISES

EXERCISE 5.1.

Consider the model of Section 5. Determine the belief (probability of the good state) after the bad news of no investment. Determine the limit of this value when $\delta \to 1$.

EXERCISE 5.2. Observation noise

Consider the model of Section 5 with observation noise. Assume that if an agent invests, he is seen as investing with probability $1 - \gamma$ and not investing with probability γ, where γ is small. Determine the equilibrium strategy. Show that for some interval $\gamma \in [0, \gamma^*)$ with $\gamma^* > 0$, the probability of the revelation of the good state and the probability of an incorrect herd are independent of γ.

EXERCISE 5.3.

Consider the simple model of delay in Section 5 where there are two possible states 1 and 0. In state 1, there are two agents each with an option to make an investment equal to 1 at the cost $c < 1$. In state 0, there is only one such agent. The gross payoff of investment is θ. The discount factor is $\delta < 1$ and the initial probability of state 1 is μ such that $0 < \mu - c < \mu\delta(1 - c)$.

1. A government proposes a policy which lowers the cost of investment, through a subsidy τ which is assumed to be small. Unfortunately, due to lags, the policy lowers the cost of investment by a small amount in the *second* period, and only in the second period. This policy is fully anticipated in the first period. Analyze the impact of this policy on the equilibrium and the welfare of agents.

2. Suppose that in addition (in each state) one more agent with an option to invest (and discount factor δ), and a belief (probability of the good state) $\underline{\mu} < c$. How is your previous answer modified?

EXERCISE 5.4.

Consider the model of Section 5 with N players in the good state and one player in the bad state. Solve for the symmetric equilibrium. Show that the probability of a herd with no investment converges to $\pi^* > 0$ if $N \to \infty$. Analyze the probability of investment by any agent as $N \to \infty$.

EXERCISE 5.5.

Show that there is strategic substitutability at an equilibrium with the strategy μ^* if

$$\mu^* > \frac{\sqrt{c/(1-c)}}{1+\sqrt{c/(1-c)}}.$$

EXERCISE 5.6.

In the model of Section **??**, assume $n \to \infty$ and the period length converges to zero, ($\delta \to 1$), at a rate slower than n. Assume that not all agents invest in the equilibrium (there is no rush).

1. Determine the payoff of an agent with private belief μ as a function of μ, $\bar{\mu}$ and c.
2. Is there a measurement of the externality of information which an agent with private belief μ receives from the agents in the upper tail of the distribution of beliefs?

REFERENCES

Bar-Ilan, A. and A. Blinder (1992). "Consumer Durables: Evidence on the Optimality of Doing Nothing," *Journal of Money Credit and Banking*, **24**, 253-272.

Caplin, A. and J. Leahy (1994). "Business as Usual, Market Crashes and Wisdom after the Fact," *American Economic Review*, **84**, 547-564.

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

Chamley, C. (2000). "Delays and Equilibria with Large and Small Information in Social Learning," mimeo, Boston University.

Chari, V. V. and P. Kehoe (2000). "Financial Crises as Herds," mimeo, Federal Reserve Bank of Minneapolis.

Hendricks, K. and D. Kovenock (1989). "Asymmetric Information, Information Externalities, and Efficiency: the Case of Oil Exploration," *RAND Journal of Economics*, **20**, 164-182.

Nasar, S. (1993). "The Economics of Wait and See," *The New York Times*, May 12.

Smith, L. and P. Sørensen (1994). "Pathological Outcomes of Observational Learning," *Econometrica*, **68**, 371-398.

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

Appendix: continuum of agents with observation noise

In macroeconomics, aggregate data are reported at discrete intervals, quarterly or monthly. These data (*e.g.* GDP growth, housing starts, durable expenditures) pertain to a large number of agents. They are also affected by noise and imperfection, and may be subject to revisions. The theoretical model of this section should be viewed in that context.

By assumption, there is a continuum of agents of total mass equal to one. As in the two-agent model, each rational player gets his private information in the form of a signal $s = \theta + \epsilon$ where the noise ϵ is independent from any other private noise or other variables in the economy and is normally distributed $\mathcal{N}(0, \sigma_\epsilon^2)$. This process of private information generates in the first period an unbounded support of the distribution of private beliefs. At the end of each period, each agent observes the level of aggregate activity

$$Y_t = y_t + \eta_t,$$

where y_t is the integral of the investments by the rational agents, and η_t is a random term which is exogenous, independent from all the other variables in the economy and normally distributed $\mathcal{N}(0, \sigma_\eta^2)$. The history h_t is now defined by $h_t = (Y_1, \ldots, Y_{t-1})$.

The analytical method of Section 5 applies. In any period t of a PBE, the strategy is monotone. It is defined by the marginal value of the signal s_t^* which depends on h_t: an agent delays if and only if his signal[17] is smaller than s_t^*. The value of s_t^* is determined by the arbitrage between the payoff of immediate investment and that of delay for one period only. The equilibrium with an infinite number of periods is thus reduced to a sequence of two-period equilibria. As long as the learning phase proceeds, agents in the interval of beliefs $(s_t^*, \bar{s}_t)$ invest in period t and are taken away from the game at the end of period t. If an agent with signal s_t^* delays in period t, he has the highest belief in period $t+1$. Note that *the distribution of beliefs is bounded above in each period after the first.*

Let F be the cumulative distribution function of the normal distribution $\mathcal{N}(0, \sigma_\epsilon^2)$. Since the mass of agents is equal to one, the observation in period t is equal to

$$Y_t = \underbrace{Max\Big(F(s_{t-1}^* - \theta) - F(s_t^* - \theta), 0\Big)}_{\text{endogenous activity } y_{\theta,t} = y(\theta, s^*)} + \underbrace{\eta_t}_{\text{noise}},$$

with $s_1^* = \infty$ by convention.

The variable Y_t is a signal on θ through the arguments of the cumulative distribution functions. If s_t^* is either large or small, the endogenous level y_t is near zero or near the

[17] It is simpler to work here with signals rather than with beliefs.

mass of remaining players, for any value of θ. In this case, the signal of the endogenous activity y_t is dwarfed by the noise η_t, and the information content of Y_t becomes vanishingly small.

Consider an agent with LLR equal to λ_t at the beginning of period t. Conditional on the observation Y_t, his LLR at the end of the period is equal to λ_{t+1} with

$$\begin{aligned}\lambda_{t+1} &= -\frac{(Y_t - y_{1,t})^2 - (Y_t - y_{0,t})^2}{2\sigma_\epsilon^2} + \lambda_t,\\ &= \frac{y_{1,t} - y_{0,t}}{\sigma_\epsilon^2}(Y_t - \frac{y_{1,t} + y_{0,t}}{2}) + \lambda_t.\end{aligned}$$

An agent with a marginal belief for investment who delays in period t has the highest belief in period $t+1$. He does not invest in the next period $t+1$ if and only if his *ex post* observation LLR is smaller than $Log(c/(1-c))$. We have the following result which is analogous to Proposition 5.2.

PROPOSITION 5.6. *In any period t of a PBE, if the observation Y_t is such that*

$$\frac{y_{1,t} - y_{0,t}}{\sigma_\epsilon^2}(Y_t - \frac{y_{1,t} + y_{0,t}}{2}) < Log\Big(\frac{c(1-s_t^*)}{s_t^*(1-c)}\Big),$$

where s_t^* is the equilibrium strategy in period t, then there is no endogenous investment after period t. All activity is identical to the noise and provides no information.

A numerical example

Figure **??** represents the option value $Q(s,s)$ and the opportunity cost of delay $((1-\delta)/\delta)(\mu(s) - c)$ as functions of the signal value s in the first period. Three graphs are represented for different standard errors of the observational noise. The diagram is very similar to Figure **??** for the case with two agents[18]. There are multiple equilibria if the discount rate and the variance of the observation noise are not small. These properties are intuitive.

The speed of learning

Recall that in the model of Section 5 with a bounded distribution of beliefs, there may be multiple equilibria with delay or no delay, respectively. An equilibrium with delay generates significant information when the marginal belief for investment is high (because of the opportunity cost), while a rush generates no information. We will now see that the rush is a stylized representation of an equilibrium in the model with a continuum of agents and observation noise in which few agents delay.

[18] The values are functions of μ in Figure **??** and functions of s in Figure **??**.

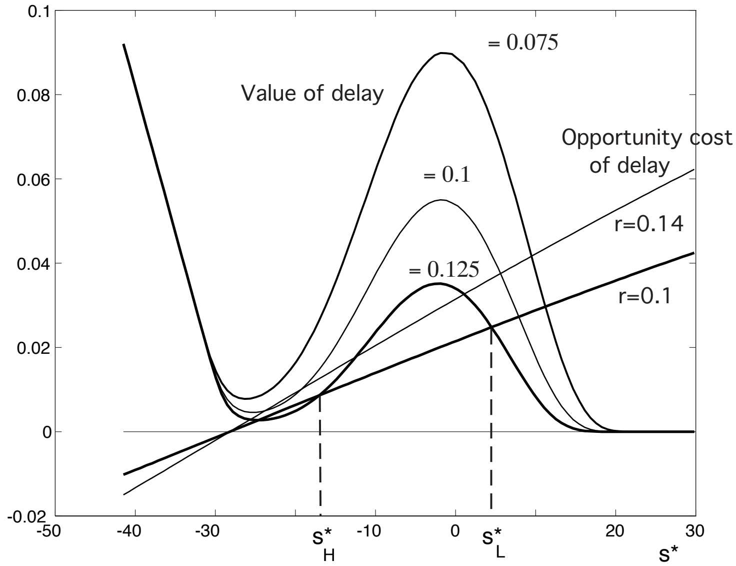

Other parameters: $\mu = 0.6$, $\theta_0 = 0$, $\theta_1 = 2$, $\sigma_\epsilon = 12$, $\delta = 1 - r$.

Figure 5.2: Equilibria with a continuum of agents

Consider in the first period an agent with a belief measured by a LLR equal to λ_1. Denote by $f(\cdot;\sigma)$ the density of the distribution $\mathcal{N}(0, \sigma^2)$, and by s^* the equilibrium strategy in the first period. Following the observation of aggregate investment in the period, $Y = 1 - F(s^* - \theta; \sigma_\epsilon) + \eta$, the agent updates his LLR from λ_1 to $\lambda_2(\theta, \eta, s^*)$ defined by

$$\lambda_2(\eta, \theta; s^*) = \lambda_1 + Log\left(\frac{f(Y - 1 + F(s^* - \theta_1, \sigma_\epsilon); \sigma_\eta)}{f(Y - 1 + F(s^* - \theta_0, \sigma_\epsilon); \sigma_\eta)}\right).$$

If the true state is good ($\theta = \theta_1$), this equation becomes

$$\lambda_2(\eta, \theta; s^*) - \lambda_1 = \frac{\Big(F(s^* - \theta_0, \sigma_\epsilon) - F(s^* - \theta_1, \sigma_\epsilon) + \eta\Big)^2}{2\sigma_\eta^2} - \frac{\eta^2}{2\sigma_\eta^2}. \tag{5.11}$$

The expectation, or the *ex post* average, of this updating over all realizations of the observation noise η is

$$\Delta(\theta_1; s^*) = E[\lambda_2(\eta, \theta; s^*) - \lambda_1] = \frac{\Big(F(s^* - \theta_0, \sigma_\epsilon) - F(s^* - \theta_1, \sigma_\epsilon)\Big)^2}{2\sigma_\eta^2}. \tag{5.12}$$

Let $\Delta(s^*) = \Delta(\theta_1; s^*)$ be the certainty equivalent of the updating expression (5.11). If the true state is bad, using the same notation one finds

$$\Delta(\theta_0; s^*) = -\Delta(\theta_1; s^*) = -\Delta(s^*).$$

The two expected values of the updates of the LLR conditional on the good and the bad states are just opposite of each other. The positive value $\Delta(s^*)$ provides an indicator of the learning process in the period and depends on the equilibrium strategy s^*.

In the example of Figure **??**, for $\sigma_\eta = 0.125$ and $r = 0.10$, there are two stable equilibria with strategies $s^*_H < s^*_L$. Investment is higher in the s^*_H-equilibrium than in the s^*_L-equilibrium. The respective mean values of the changes of beliefs are

$$\Delta(s^*_H) = 0.0015, \qquad \Delta(s^*_L) = 0.129.$$

The difference in information between the two equilibia is significant. In the equilibrium with low investment in the first period (s^*_L), the variation of the LLR is 80 times[19] higher than in the H-equilibrium.

In the equilibrium with high investment (s^*_H), a large fraction of agents invest with no delay. In that period and the periods after, agents do not learn much. The equilibrium is remarkably similar to the rush equilibrium of the model with bounded beliefs of Section 5 (in which they learned nothing). The rush is a stylized property of the s^*_H-equilibrium.

Learning in multiple periods
After the first period, the support of private beliefs has a finite upper-bound. This is important: it means that agents never learn with certainty whether the state is good. Furthermore, in each period after the first, with a strictly positive probability investment stops completely in a cascade with no investment: assuming a marginal value s^*_τ in the support of beliefs[20] for each $\tau \leq t$, then $s^*_{t+1} \geq \bar{s}_{t+1}$ with some strictly positive probability. The game and the evolution of beliefs proceed as in the model of Section 5 with a finite number of agents. In each period, the possible equilibria are of the types described in Theorem 5.1.

5.3.9 Investments of variable size

We have seen how the model with discrete actions and exogenous timing can be extended to a model In the models considered so far, individual actions are an imperfect filter of individual information because they are discrete. Lee (199 actions were taken in a continuum, they would reveal perfectly a one dimensional private information (Lee (1993)).

[19] Other simulations have shown similar results.

[20] The marginal value is not close to the upper-bound of the support as in Section **??**, because the mass of endogenous investment would be dwarfed by the observation noise and would not convey significant information.

Such a critique is similar to the argument of an unbounded distribution in Smith and Sørensen (Chapter 3) and is subject to the same counter-argument: the perfect information on individuals' actions and on their decision models are required, which is not realistic. The problem of social learning with individual actions in a continuum has to be analyzed in the context of imperfect observability. The previous setting is now extended to include a variable investment size and an observation noise. Each agent has one option to invest and the investment which is made only once (if ever) is chosen in the set of real numbers. For example, agents decide both the period in which to purchase a new car and the amount spent on the car, (number of accessories, etc...). Each agent has therefore two choice variables, the time of the investment and its scale. As before, investment is irreversible. Following the previous results, one can assume without loss of generality that there are two periods. Since the scale of investment is taken in a continuum, we redefine the payoff function.

Any agent who has not yet exercised his option to invest receives a payoff equal to $(1-\delta)b$ per period where δ is the discount factor. An agent who never invests receives a payoff equal to b. The difference $1-\delta$ corresponds to the rate of return between two periods.

For tractability, the payoff of investment is a quadratic function[21]. If the agent invests in period t, he foregoes in that period the payoff of never investing and gets a payoff with a current value equal to $E[2az-(\theta-z)^2]$, where the expectation operator E depends on the information of the agent, and a is a constant parameter. The scale of investment z is chosen in the set of real numbers, and θ is the productivity parameter which is determined as in the previous sections.

The payoff of investing in period 1 is

$$U_1 = 2az - E\Big[(\theta-z)^2\Big] - b,$$

and the payoff of investing in the second period is

$$U_2 = (1-\delta)b + \delta E\Big[2az - (\theta-z)^2 - b\Big].$$

By assumption, nature's distribution of θ is $\mathcal{N}(\bar{\theta}, \omega_0)$. θ is not directly observable, but each agent receives once, before the first period, a signal

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

In this section the symbol s denotes the private signal of an agent (not his belief). The

[21]The model presented here is inspired by Bar-Ilan and Blinder (1992).

private noise ϵ is normally distributed and independent from any other random variable in the model.

As in Section **??**, each agent is infinitesimal and the total mass of agents is equal to one. At the end of period 1, the observed level of aggregate investment is equal to

$$Y = y + \eta, \quad \text{with} \quad \eta \sim \mathcal{N}(0, \sigma_\eta),$$

where y is the integral of the individual investments z. The variable η is an exogenous random term which is independent from the other variables in the economy.

It can be shown that for some parameter values, there are multiple PBEs with monotone strategies such that agents delay if and only if they have a private signal smaller than some value s^*. The signal at the end of the first period is the aggregate investment

$$Y = z_1(\theta; s^*) + z_2(\theta; s^*) + \eta.$$

Each of the two terms $z_1(\theta; s^*)$ and $z_2(\theta; s^*)$ is an increasing function of θ, for given s^*, and thus contributes to the information on θ. The two terms represent two separate effects. The first is proportional to the mass of agents who invest in period 1. It is identical to the endogenous investment in a model where each investment has a fixed scale. This is the *timing effect*. The second term depends on the mean scale of investment by investing agents and is called the *level effect*.

Because of the observation noise η, the information which is conveyed through each of the two effects depends on the impact of θ on z_1 and z_2. If the impact is small, it is drowned in the noise. It can be shown that the magnitude of the level effect in $z_2(\theta; s^*)$ becomes vanishingly small if the precision of the individual signal, $1/\sigma_\epsilon^2$, tends to zero. There is a simple interpretation: if an individual has a signal of small precision, the scale of his investment does not depend much on his signal. The timing effect however remains of the same order of magnitude as the (given) mass of agents, and does not become vanishingly small when $1/\sigma_\epsilon^2$ tends to zero. The information property of Y is similar to that in a model with fixed investment scale.

A numerical example

Since there is no algebraic solution to the model, we consider a numerical example. From the previous discussion, we know that the important parameter is the precision of the private signals. The ratio σ_ϵ/ω_0 is taken to be equal to 5. It implies that if an agent could observe directly the signals of others, in order to double the precision of his estimate (as measured by the inverse of the variance), he would have to observe roughly 25 other private signals.

The option value $Q(s, s^*) = \omega_1^2 - E_{\{s,s^*\}}[\omega_2^2(Y, s, s^*)]$ and the opportunity cost of delay $c(s)$ for the marginal agent $s = s^*$ are represented in a figure that is similar to Figure 5.3.3. In particular, there are two stable equilibria, with a large and a small mass of delaying agents, respectively.

An analysis

Individual decisions

An agent with a signal s updates the public information distribution on θ with his own signal s. His subjective distribution is therefore $\mathcal{N}(m_1(s), \omega_1)$, with

$$m_1(s) = \bar{\theta} + \gamma(s - \bar{\theta}), \quad \gamma = \frac{\omega_0^2}{\omega_0^2 + \sigma_\epsilon^2} \quad \text{and} \quad \frac{1}{\omega_1^2} = \frac{1}{\omega_0^2} + \frac{1}{\sigma_\epsilon^2}. \tag{5.13}$$

If he invests in the first period, he chooses a level $z(s)$ which depends on his information:

$$z(s) = a + m_1(s) = a + (1 - \gamma)\bar{\theta} + \gamma s, \tag{5.14}$$

and the payoff of investing in the first period is

$$U_1(s) = -\omega_1^2 + 2am_1(s) + a^2 - b.$$

An agent with signal s who delays while others use the strategy s^* invests in period 2 and has a payoff

$$U_2(s, s^*) = (1 - \delta)b + \delta E_{\{s,s^*\}}\Big[-\omega_2^2(Y, s, s^*) + 2am_2(Y, s, s^*) + a^2 - b\Big],$$

where the expectation is computed over $\omega_2^2(Y, s, s^*)$ and $m_2(Y, s, s^*)$ which are the mean and the standard error of θ, respectively, after the observation of Y.

Since $m_2(Y, s, s^*)$ is an updating of $m_1(s)$, then $E_{\{s,s^*\}}[m_2(Y, s, \mu^*)] = m(s)$, and the difference between the payoffs of delay and investment in the first period is

$$U_2(s, s^*) - U_1(s) =$$
$$\delta\Big(\omega_1^2 - E_{\{s,s^*\}}[\omega_2^2(Y, s, s^*)]\Big) - \Big(1 - \delta\Big)\Big(-\omega_1^2 + a^2 + 2am_1(s) - 2b\Big).$$

This difference can be rewritten as the difference between the option value $Q(s, s^*)$ and the opportunity cost $c(s)$ of delay:

$$U_2(s, s^*) - U_1(s) = \delta\Big(Q(s, s^*) - c(s)\Big), \quad \text{with}$$

$$\begin{cases} Q(s, s^*) = \omega_1^2 - E_{\{s,s^*\}}[\omega_2^2(Y, s, s^*)], \\ c(s) = \dfrac{1-\delta}{\delta}\Big(-\omega_1^2 + a^2 + 2am_1(s) - 2b\Big). \end{cases} \tag{5.15}$$

In models with normal distributions and linear decision rules, the learning rules are linear and the *ex post* variance, ω_2, is independent of the observation and can be computed *ex ante*. This very nice property does not hold in the present model because the endogenous investment y is not a linear function of the random variables.

Equilibrium and information

A symmetric equilibrium in monotone strategies is defined by a value s^* which satisfies the arbitrage equation between the option value and the opportunity cost:

$$Q(s^*, s^*) = c(s^*).$$

Using the updating rule (5.13) and the expression of the individual level of investment $z(s)$ in (5.14), the level of endogenous aggregate activity is equal to

$$\begin{aligned} y(\theta; s^*) &= \int_{s^*-\theta} \Big(a + (1-\gamma)\bar{\theta} + \gamma(\theta + \epsilon)\Big) f(\epsilon; \sigma_\epsilon) d\epsilon \\ &= \Big(a + (1-\gamma)\bar{\theta} + \gamma\theta\Big)\Big(1 - F(s^* - \theta; \sigma_\epsilon)\Big) + \gamma \int_{s^*-\theta} \epsilon f(\epsilon; \sigma_\epsilon) d\epsilon. \end{aligned}$$

We can normalize $\bar{\theta} = 0$, (or incorporate $(1-\gamma)\bar{\theta}$ in the definition of a).

Since $\displaystyle\int_{s^*-\theta} \epsilon f(\epsilon; \sigma_\epsilon) d\epsilon = \sigma_\epsilon^2 f(s^* - \theta; \sigma_\epsilon)$, and $1 - F(z; \sigma) = F(-z; \sigma)$,

$$\begin{aligned} y(\theta; s^*) &= \Big(a + \frac{\omega_0^2\theta}{\omega_0^2 + \sigma_\epsilon^2}\Big) F(\frac{\theta - s^*}{\sigma_\epsilon}; 1) + \frac{\omega_0^2\sigma_\epsilon^2}{\omega_0^2 + \sigma_\epsilon^2} f(\frac{\theta - s^*}{\sigma_\epsilon}; 1) \\ &= \qquad z_1(\theta; s^*) \qquad\qquad + \qquad z_2(\theta, s^*). \end{aligned} \tag{5.16}$$

The aggregate activity which is observed is

$$Y = z_1(\theta; s^*) + z_2(\theta; s^*) + \eta.$$

Suppose that $\sigma_\epsilon \to \infty$. Since

$$\sigma_\epsilon^2 f(\frac{\theta - s^*}{\sigma_\epsilon}; 1) = \frac{\sigma_\epsilon}{\sqrt{2\pi}} exp\Big(-\frac{(\theta - s^*)^2}{2\sigma_\epsilon^2}\Big),$$

one can see in equation (5.16) that the magnitude of the level effect in $z_2(\theta; s^*)$ becomes vanishingly small.

5.3.10 Proofs

Lemma 5.1

We first prove the following: in any arbitrary period t of a PBE, if an agent with belief μ_t delays, then any agent with belief $\mu_t' < \mu_t$ strictly prefers to delay. Let the arbitrary period be the first one. Consider an agent with belief μ who has a strategy with delay: this is a rule to invest in period t, (with $t \geq 2$), if and only if the history h_t in period t belongs to some set H_t. For this agent the difference between the payoff of the strategy of delay and the payoff of immediate investment is

$$\begin{aligned}
W(\mu) &= \textstyle\sum_{t\geq 2, h_t \in H_t} \delta^{t-1} P(h_t|\mu)\Big(P(\theta=\theta_1|\mu, h_t) - c\Big) - (\mu - c)\\
&= \textstyle\sum_{t\geq 2, h_t \in H_t} \delta^{t-1} P(h_t|s)\Big(\dfrac{P(h_t|\theta=\theta_1)}{P(h_t|\mu)}\mu - c\Big) - (\mu - c)\\
&= \textstyle\sum_{t\geq 2, h_t \in H_t} \delta^{t-1} \Big(\mu(1-c)P(h_t|\theta=\theta_1) - c(1-\mu)P(h_t|\theta=\theta_0)\Big)\\
&\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad -(\mu - c)\\
&= as - b - (\mu - c),
\end{aligned}$$

where a and b are independent of μ:

$$\begin{aligned}
a &= \textstyle\sum_{t\geq 2, h_t \in H_t} \delta^{t-1}\Big((1-c)P(h_t|\theta=\theta_1) + cP(h_t|\theta=\theta_0)\Big),\\
b &= c\textstyle\sum_{t\geq 2, h_t \in H_t} \delta^{t-1} P(h_t|\theta=\theta_0).
\end{aligned}$$

For $\mu = 0$, because $t \geq 2$, $\delta < 1$ and $\sum_{t\geq 2, h_t\in H_t} P(h_t|\theta=\theta_0) \leq 1$,

$$W(0) = c\Big(1 - \sum_{t\geq 2, h_t \in H_t} \delta^{t-1} P(h_t|\theta=\theta_0)\Big) > 0.$$

Since an agent with belief μ delays, $W(\mu) \geq 0$. Since W is linear in s, $W(\mu') > \mu - c$ for any $\mu' < \mu$.

Consider now an agent with belief μ' who mimicks an agent with belief μ: he invests at the same time as the agent with belief μ (*i.e.*, in period t if and only if $h_t \in H_t$). For such an agent, the difference between the payoff of this strategy and that of investing with no delay is $W(\mu')$, which by the previous argument is strictly positive if $\mu' < \mu$. The agent with belief μ' strictly prefers to delay.

The set of beliefs for delay is not empty since it includes all values below c. The value of μ_t^* in the lemma is the upper-bound of the set of beliefs of delaying agents. The previous result in this proof shows that any agent with $\tilde{\mu}_t < \mu_t^*$ delays. □

Proposition 5.1

Denote by $W(\mu, \mu^*)$ the payoff of an agent with belief μ who delays for one period while other agents follow the strategy μ^*. By (5.4), μ^* is solution of

$$\mu^* - c = W(\mu^*, \mu^*).$$

Denote by $P(x_t = k|\mu, \mu^*, f^j, n)$ the probability that $x_t = k$ for an agent with belief μ when all other agents use the strategy μ^*, the density functions are f^j and the number of remaining players is n. Using Bayes' rule and the sum of probabilities equal to one,

$$\begin{aligned}
\mu^* - c &= \sum_k P(x_t = k|\mu^*, \mu^*, f^j, n)\Big(P(\theta = \theta_1|x = k; \mu^*, \mu^*, f^j, n) - c\Big) \\
&= \sum_k P(x_t = k|\mu^*, \mu^*, f^j, n) Max\Big(P(\theta = \theta_1|x = k; \mu^*, \mu^*, f^j, n) - c, 0\Big) \\
&\quad - \sum_k P(x_t = k|\mu^*, \mu^*, f^j, n) Max\Big(c - P(\theta = \theta_1|x = k; \mu^*, \mu^*, f^j, n), 0\Big).
\end{aligned}$$

An agent who delays invests in the next period only if his payoff is positive. Therefore, the payoff of delay is

$$\begin{aligned}
&W(\mu^*, \mu^*) = \\
&\delta \sum_k P(x = k|\mu^*, \mu^*, f^j, n) Max\Big(P(\theta = \theta_1|x = k; \mu^*, \mu^*, f^j, n) - c, 0\Big).
\end{aligned}$$

We conclude the proof by comparing the two previous equations and using the decomposition $\mu^* - c = (1 - \delta)(\mu^* - c) + \delta(\mu^* - c)$. $\square$

Chapter 6

Regime switches

In the winter of 1989, despite of the simmering of future events in the Soviet Union, Kissinger delivered another Cold-War rhetoric in a speech to US governors (Halberstam, 1991). In the Spring and the Summer, the simmering led to ebullition with growing demonstrations in East Germany. people."[1] Twelve days later, Honecker resigned. On November 9 the Berlin Wall fell. "Western observers were initially stunned at the speed of the economic and political collapse of the East German regime. With hindsight however, the regime's economic collapse seems to have been inevitable."[2]. Events that had been hard to imagine in the sphere of public information acquired an aspect of obvious inevitability.

Later, "springs" of various colors, orange or green, would bring surprises in the Ukraine and in Arab countries. The subsequent fading of the flowers do not foreclose the possibility of similar surprises in the future. This chapter focuses on the mechanism by which such events can take rational agents by surprise and on the contrast between the low expectations *ex ante* and a feeling of obvious determination *ex post*.

Sudden and unexpected changes in political regimes, economic activity, financial crises, share a fundamental underlying property. The payoff of individuals' actions, (*e.g.*, street demonstration, investment) increases with other agents taking the same action. The collective behavior generates strategic complementarities.

In the previous chapter, the coordination game with strategic complementarities took place in one period. All individuals were thinking simultaneously without learning from the past.

[1] Lohman, 1994, p. 42).
[2] Lohmann (1994, 43)

The process of equilibrium selection between a high and a low level of aggregate activity rested on the agents' imperfect information about others' payoffs and the possibility that the fundamentals of the economy took "extreme values" where one action (*e.g.*, investment or no investment) was optimal independently of others' actions. In the one-period setting, all individuals were thinking by induction without the possibility of learning. Learning from the observation of others' actions is the central issue in this chapter.

6.1 Contexts and issues

How do business cycles, demonstrations toward a revolution and conventional discourse (what the French call "wooden speak") share the property that the payoffs of individual action is augmented by the number of other individuals taking the same action? Three different contexts are first presented to justify a canonical model.

Business cycles

The profitability of individual investment increases with the level of activity in the economy, which itself increases with the level of individual investments. This feature has been represented in models with imperfect competition by Blanchard and Kiyotaki (1987), Schleifer (1986), Cooper (1993), and others. In such models, more aggregate investment increases the productivity of the economy and the demand curve of each firm shifts upwards, thus generating more profits which in turn stimulate more investment by each individual firm. The strategic complementarity between individuals' investments, if sufficiently strong, generates multiple Nash equilibria that are indicative of business cycles.

A canonical model is a simple analytical representation that focuses on a particular effect and is abstracted from the clutter of non essential features of the reality. Assume that there is a large number of agents, that is a continuum with a mass normalized to 1. There is one period and each agent has to make a 1 or 0 decision, for example whether to make a fixed size investment or not. Each agent is characterized by his own cost of investment, c, that is taken from a distribution with a density f, as represented by the graph (f) in the lower panel of Figure 6.1). The cumulative distribution for that density is represented by the curve (F) in the upper panel of the figure.

The positive impact on the payoff of anyone's investment by the aggregate investment, X, is represented here by an increasing function of X. Without much loss of generality,[3] we can assume that this function is linear. We are thus led to the payoff function for

[3] Some non linear gross payoff functions can be transformed into a linear payoff by changing the distribution of the costs c.

investment x by an agent with cost c:

$$w(x, X, c) = \begin{cases} X - c, & \text{if } x = 1, \\ 0, & \text{if } x = 0. \end{cases} \tag{6.1}$$

Suppose that all agents follow the monotone strategy to invest when their cost is less some some cutoff $\bar{c}$. The value $\bar{c}$ defines the strategy. By definition of the c.d.f., the gross payoff of any investing agent is $F(\bar{c}) - c$. For an agent with a cost less than $F(\bar{c})$, investment has a positive payoff. We can thus define a *reaction function*. A strategy is defined as investing when the cost is less than c. (It is monotone). When others have the strategy $\bar{c}$, the optimal response is $F(\bar{c})$. A Nash equilibrium strategy c^* is a fixed point of the cumulative distribution function: $F(c^*) = c^*$. In Figure 6.1, a fixed point is represented by an intersection of the graph of F with the 45^o line. Here, there are three such points. The middle point can be discarded by a loose argument on stability. The points L and H represent low and high levels of aggregate activity.

Financial crises and speculative attack

Consider a bank for which the probability of bankruptcy increases with the quantity of deposits withdrawals, X. Using a previous argument, assume that this function is linear. Let c the cost for depositors to withdraw their deposits, and for example invest them in projects of lower return. One can normalize the costs and the gross payoff from avoiding the capital loss in case of the bank failure such that the payoff for withdrawing ($x = 1$), and not withdrawing ($x = 0$) are given by

$$w(x, X, c) = \begin{cases} X - c, & \text{if } x = 1, \\ a(1 - X), & \text{if } x = 0, \end{cases} \tag{6.2}$$

where a measures the payoff if the bank does not go bankrupt, an event with probability $1 - X$. The payoff difference between the two actions is $(1 + a)X - c - a$. It has the same form as (6.1). A speculative attack against a central bank that manages a regime of fixed exchange rate is the formally same as an attack against a commercial bank that manages a fixed exchange rate between its deposits and the legal currency.

The Leipzig demonstrations

The fall of the Berlin wall was preceded by a wave of increasing demonstrations in Leipzig, beginning in September 1989.[4] Suppose that the individual benefit from participating in a demonstration (action $x = 1$) increases with the size of the demonstration, X, and depends on the individual cost c that increases with approval for the regime. Such a payoff can be represented by the same function as (6.1). The framework that is now presented provides

[4]Other waves of demonstrations took place after the fall of the wall, See Lohmann (1994).

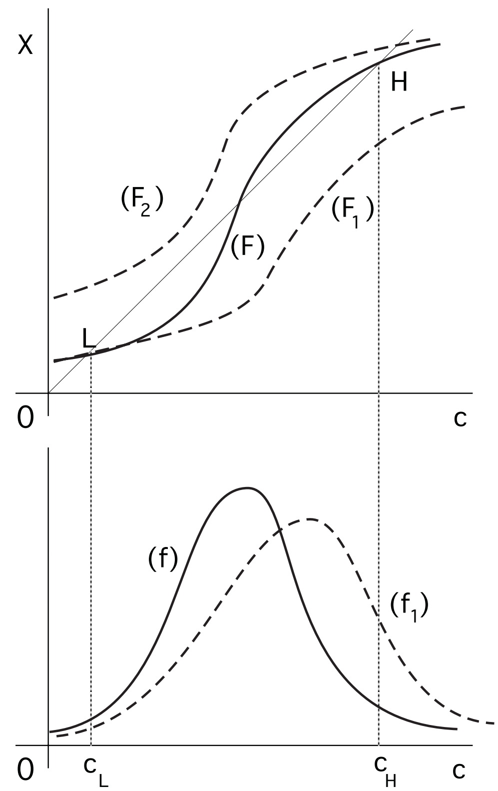

Figure 6.1: Cumulative distribution functions are represented in the upper part and associated density functions in the lower part.

an analytical representation of the sudden change of beliefs from the unpredictable to the "inevitable" of the Western observers at that time.

Social changes and revolutions

Why do sudden changes of opinions or revolutions which were not anticipated with high probability seem anything but surprising in hindsight? This question was asked by Kuran (1995). The gap between the *ex ante* and the *ex post* views is especially striking when no important exogenous event occurs (*e.g.*, the fall of the communist regimes)[5].

These social changes depend essentially on the distribution of individuals' payoffs, on which each agent has only partial information. According to Kuran, "historians have systematically overestimated what revolutionary actors could have known". If a revolution were to be fully anticipated, it would probably run a different course. The July 14th entry in the

[5]For a common view before the fall, read the speeches of H. Kissinger in Halberstam (1991).

diary of Louis XVI was "today, nothing"[6]. Before a social change, individuals who favor the change do not have perfect information on the preferences of others *ex ante*, but they are surprised to find themselves in agreement with so many *ex post*, and this common view in hindsight creates a sense of determinism.

Following Kuran (1988), (1995), suppose that individuals decide in each period between two actions or "expressed opinions" as revealed by some behavior: action 1 is to speak against a given political regime, while action 0 is to speak in favor. Each individual is characterized by a preference variable c which is distributed on the interval $[0, 1]$ with a cumulative distribution function $F(c)$. The preference for the regime increases with c. There is a continuum of individuals with a total mass equal to one. For an individual with parameter c, the payoff of his action x (which is either 0 or 1), is a function which is (i) decreasing in the "distance" between his action and his preference, (ii) increasing in the mass of individuals who choose the same action. For example, in talking to someone, the probability to find a person speaking against the regime increases with the mass X of people speaking against the regime. Assume that speaking against the regime yields a payoff $X - c$. Likewise, speaking for the regime has a payoff $1 - X - (1 - c)$. The difference between speaking against and speaking for is thus

$$u(c) = X - c - (1 - X) + (1 - c) = 2(X - c). \tag{6.3}$$

It has the same form as the previous utility for demonstrating.

The model of "Private Truths and Public Lies" of Kuran is thus a special case of the canonical model with strategic complementarities. For a suitable distribution of individual preferences, the model has multiple equilibria under perfect information. Kuran follows the *ad hoc* rule of selection and assumes that a regime stays in power as long as the structure of preferences allows it. When this structure evolves such that the regime is no longer a feasible equilibrium, society jumps to the other equilibrium regime. But it is obvious that for the analysis of sudden changes of beliefs, such an *ad hoc* rule in a static model, with perfect information, is not appropriate. The previous discussion points to a dynamic approach and an explicit formulation of expectations in a setting of imperfect information and learning. These features have a central place in the dynamic models of this chapter.

In such a model, we will see that until the very end of the old regime, the public information is that a large fraction of the population supports the old regime, whereas the actual distribution could support a revolution. When the regime changes, beliefs change in two

[6]However, the entry may mean "no hunting". The quote at the beginning of the chapter is from a conversation between Louis XVI and the duke of La Rochefoucault-Liancourt. In the numerous stages of the French revolution, the actors did not seem to have anticipated well the subsequent stages, especially when they manipulated the crowds.

ways: first, the perceived distribution of preferences shifts abruptly towards the new regime; second, the precision of this perception is much more accurate. The high confidence in the information immediately after the revolution may provide all individuals with the impression that the revolution was deterministic.

6.2 Analysis in a canonical model

Following the previous discussion, the canonical model is defined by the continuum of heterogenous agents, each characterized by his cost of "investment" c, and the payoff function in (6.1). As with many analyses of strategic complementarities, we begin with the case of perfect information, both on the *structure* of the economy, and on the *strategies* of all agents. We then move on to imperfect information.

6.2.1 Perfect information

Suppose that in some period, the actual distribution of costs is represented by the *c.d.f.* F_1 as in Figure 6.5. Under perfect information about F_1, there are two equilibria in monotone strategies (to act when the cost is less than some cutoff), L_1 and H_1. In a setting of perfect information within one period, there is no criterion for choosing a high or a lower cutoff point for an equilibrium strategy. And recall that either equilibrium requires that all agents have no uncertainty on the strategies of others and coordinate on the same strategy.

Hysteresis as a device for equilibrium selection

Suppose that agents coordinate on the equilibrium L_1 and that the game is reproduced in another period with a structure of costs (the density function of the costs) that is slightly different. In the figure, the new c.d.f. is represented by the curve (F_2). When agents know in period 2 that the equilibrium L_1 has been achieved in the first period, it is reasonable to assume that with the two possible equilibria L_2 and H_2, they choose L_2 which is closest to L_1. This selection device may be loosely defined as inertia or hysteresis. Between the two periods, a small change in the structure of the economy generates a small change in the aggregate activity.

Suppose now that the structure of economy moves a little more from (F_2) to (F_3): the low equilibrium vanishes and H_3, at a high level, is the unique equilibrium. A small change in the structure of the economy generates a large jump of the aggregate activity. After the jump, further small changes of the structure do not generate another jump. For example, if the c.d.f. returns to (F_1), the level of activity stays high at H_1 and does not jump down

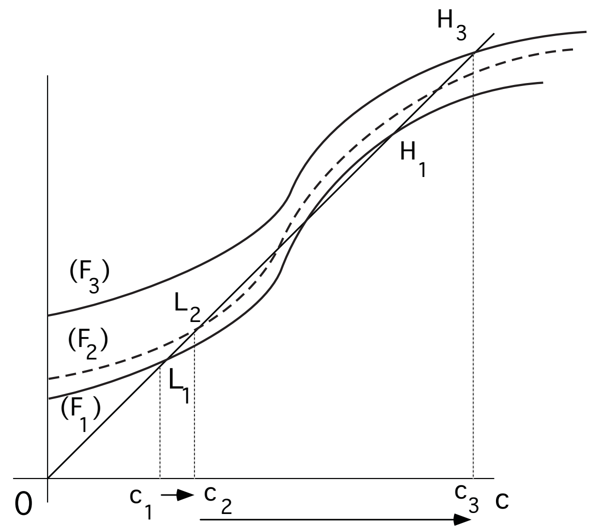

F_1, F_2 and F_3 are the realizations of the *c.d.f.* F for three periods. The *c.d.f.* evolves slowly between consecutive periods and agents coordinate on the equilibrium strategy that is closest equilibrium strategy in the previous period. Between the second and the third period, the strategy jumps to a higher equilibrium level.

Figure 6.2: Evolutions of a cdf (or the reaction function)

to L_1. In this setting with perfect information and the selection through inertia, when the structure evolves by small steps, the aggregate activity in any period is strongly dependent on the its level in the previous period. There is *hysteresis*. The aggregate activity evolves by small steps during extended regimes that are separated by large jumps.

The assumptions about perfect information on the structure and the strategies, together with the *ad hoc* criterion of inertia are somewhat problematic in a setting with a large number of diverse agents. We will now see that the pattern of hysteresis and regime switches will be robust when agents have imperfect information and learn from past level of aggregate activity.

6.2.2 Coordination with imperfect information

The distribution of the costs is not directly observable. It is perceived by agents through probability distributions that are agent specific. These distributions are updated after the observations of aggregate activity. The distribution of costs evolves randomly by small steps from period to period. Nature does not make jumps.[7] In each period, agents play a one period game under imperfect information with a payoff equal to the expected value

[7] *Natura non facit saltus* (Leibnitz).

of the payoff in (6.1). For practicality, the pool of agents is new in each period.[8] As in a one-period setting with no learning from the actions of others (in a global game), imperfect information on the structure of fundamentals will enable us to solve the problem of strategic coordination.

Learning from activity with strategic complementarities and the tail property

Consider again the point L in Figure 6.1. That point is compatible with the functions (F) and (F_1). For the first, there is another equilibrium (under perfect information) with a higher activity, at the point H. For the second, (F_1), there is no such equilibrium. The level of activity at the point L is determined by the mass of agents with a cost lower than c_L, in the left tail of the density function f. When the cost c_L is low, the left tail of the distribution should, in a realistic model, provide little information on the rest of the distribution. We call this the *tail property*.

The tail property is important when agents learn about the structure of an economy with strategic complementarities. In such a setting, the strategic complementarity operates like a critical mass. Either few agents take action because that critical mass is not reached and the payoffs of action are low, except for these few agents, or that critical mass of active agent is reached and that is why a large mass is acting, except for the few that have a cost much higher than the average. In the present setting where agents are differentiated by their own cost of action, the strategic complementarity imposes that in an equilibrium the cutoff point of the cost for investment is in the tail of the distribution, to the left or to the right. In that case, agents learn little from the observation of others. We can thus expect that under strategic complementarity when the structure of the economy evolves randomly by small steps, the level of activity hovers around successive plateaus where little is learn, which are separated by abrupt changes that generate a large amount of information.

Modeling the essential property of leaning under strategic complementarity

In the construction of a canonical model we have to keep it simple and yet to embody the complexity of the possible states of nature that individuals face under imperfect information. For tractability, one has to choose a family of possible distributions that are indexed by some parameters. However, the reduction of the states of the world to a few parameters may also trivialize the inference problem. If for example, the distribution of the costs is normal with mean m and variance σ^2, two observations of the mass in a tail of the distribution, no matter how far from the mean, are sufficient to identify perfectly the

[8] If agents live more than one period, the evolution of their cost provides additional information on the evolution of the distribution of costs, and the inference problem becomes very complex.

two parameters, thus providing perfect information on the distribution. That is obviously not a plausible representation.

One example of a family of distribution that keeps the tail property is presented in the Appendix, following Chamley (1999). It has the shape of a square hat with the central part moving randomly left or right. The observation of the mass in tail on the side of the "hat" provides on information on the position of the central part as long as that central part does not "bite" on the tail. The properties of the model can be investigated analytically. In particular, provided that the variance of the distribution is within some bounds, there is a unique equilibrium that is strongly rationalizable. The model is effectively a global game model. At this stage, this may be the only global game model with an infinite number of periods where the fundamental evolves random in small steps.[9]

The tail property can also be modeled by the combination of a simple family of cost distributions, normal distributions with fixed variance, and the observation of aggregate activity subject to a noise with a fixed variance. In this setting, when the cutoff for taking action is very low, the mean of the fundamental distribution has a small impact on the mass of agents taking action, and that impact is drowned by the noise. In this case, the observation of aggregate activity provides little information on the mean of the distribution. When the cutoff point is near the center of the distribution, small variations of the mean have an impact that dwarfs the noise and the observation of aggregate activity is highly informative.

Observing the activity of others through noise

We assume that the population is the sum of two groups. In the first, the distribution of costs is normal $\mathcal{N}(\theta_t, \sigma_\theta^2)$, where σ_θ is a publicly known constant, and θ_t follows a random walk that will be discussed below. The second is the sum of a fixed mass a that always invest, and a population with a uniform distribution of costs with density β on the interval (b, B). At the end of any period t, agents observe the variable Y_t defined by

$$Y_t = a + \beta(b + c_t^*) + F(c_t^*; \theta_t) + \eta_t, \quad \text{with} \quad \eta_t \sim \mathcal{N}(0, \sigma_\eta^2). \tag{6.4}$$

The noise η_t may arise from imperfect data collection or from the activity of "noise agents" who act independently of the level of the aggregate activity.

Since individuals follow the strategy to invest when their cost is lower than c_t^*, that value

[9] Other models with multiple periods either assume that the fundamental is subject to unbounded shocks between periods, thus generating a sequence of one period global game models (Carlsson and Van Damme ****), or have a global game with a unique equilibrium only in the first period, after which equilibria are multiple (Angeletos and Hellwig (****).

is publicly known and the observation of aggregate activity is informationally equivalent to the observation of

$$Z_t = F(c_t^*; \theta_t) + \eta_t. \tag{6.5}$$

As discussed above, when $|c^* - \theta|$ is large, $F(c^*; \theta)$ does not depend much on θ and it is near 0 or 1. In that case, the noise η dwarfs the impact of θ on $F(c^*; \theta)$, and the observation of Y conveys little information on θ. Learning is significant only if $|c^* - \theta|$ is relatively small, *i.e.*, when the associated density function $f(c^*; \theta)$ is sufficiently high. But the strength of the strategic complementarity is positively related to $f(c^*; \theta)$ (which is identical to the slope of the reaction function under perfect information). We thus verify that *learning and strategic complementarity are positively related.* Agents only learn a significant amount of information when the density of agents near a critical point is sufficiently large to push the economy to the other regime.[10]

"Natura non facit saltus"[11]

Following the discussion around Figure 6.5, there imperfect information because the structure of (the costs in) the economy evolves randomly over time. In all known cases, aggregate productivity (the inverse of the cost) does not jump but evolves only in small steps. This restriction has an important implication for multi-period models with strategic complementarity and imperfect information.[12] For computation, the mean of the distribution, θ_t is assumed to take a value on the grid

$$\Theta = \{\omega_1, \dots, \omega_K\}, \quad \text{with} \quad \omega_1 = \gamma, \omega_K = \Gamma. \tag{6.6}$$

The distance between consecutive values is equal to ϵ, which can be small. Between consecutive periods, the value of θ evolves according to a symmetric random walk: it randomly either stays constant or move to a set of a small number of adjacent grid points. If θ is on a reflecting barrier (γ or Γ), it moves away from that barrier with some probability.

6.3 The behavior of the canonical model

In each period, t, learning and decision making proceed in the following steps.

1. Let $\{\pi_{k,t-1}\}$ be the public distribution of probabilities on the grid Ω at the beginning of period $t-1$ when agents determined the strategy c_{t-1}^*. This belief is updated in

[10] This property has a strong form in the model with a rectangular distribution that is sketched in the Appendix.

[11] Leibnitz

[12] For example, it rules out multi-period global games with an aggregate parameter that is subject to unbounded random shocks (Carlsson and Van Damme ***), and for which a new global game takes place in each period.

two steps, first using the knowledge of the strategy in the previous period, c^*_{t-1}, second, using the law of the random evolution of θ between period $t-1$ and period t. Using the observation of the aggregate activity Y_{t-1} in the previous period, which, as we have seen, is equivalent to $Z_{t-1} = F(c^*_{t-1}; \omega_k) + \eta_{t-1}$, and Bayes' rule, the first updating leads to the distribution $\{\pi_{k,t}\}$ with

$$Log(\hat{\pi}_{k,t}) = Log(\pi_{k,t-1}) - \frac{\big(Z_t - F(c^*_t, \omega_k)\big)^2}{2\sigma^2_\eta} + \alpha, \tag{6.7}$$

where α is a constant such that the sum of the probabilities $\pi_{k,t}$ is equal to one.[13]

2. The second updating, from $\{\pi_{k,t}\}$ to the public belief $\{\pi_{k,t}\}$ at the beginning of period t, is straightforward. For example is θ_t follows a random walk with equal probabilities of $1/3$ for staying constant, or moving up or down by one step on the grid, for all points away from the boundaries,

$$\pi_{k,t} = (\hat{\pi}_{k-1,t} + \hat{\pi}_{k,t} + \hat{\pi}_{k+1,t})/3. \tag{6.8}$$

3. Each agent with a cost c knows that c is drawn from the true distribution with mean θ_t. He updates the public distribution $\{\pi_{k,t}\}$ into $\{\pi_{k,t}(c)\}$ as in (**??**):

$$Log(\pi_{k,t}(c)) = Log(\pi_{k,t}) - \frac{(c-\omega_k)^2}{2\sigma^2_\epsilon} + \alpha', \tag{6.9}$$

where α' is a constant such that the sum of the probabilities is equal to one. Note that each agent "pulls" the distribution of θ_t towards his own cost c.

4. Each agent computes for his own cost c, the *cumulative distribution function* (CVF). By definition of the CVF, the agent assumes that all the agents with a cost not greater than his own c make the investment, of equivalently that the strategy of others is c. Given this assumption, the agent computes the expected value of the mass of investment according to his probability estimates of θ_t. The CVF is therefore defined by

$$V_t(c) = E[F_{\theta_t}(c) | \{\pi_{k,t}(c)\}]. \tag{6.10}$$

5. The function $V_t(c)$ is increasing in c. In the analytical model with a rectangular density, under some parametric conditions, this function is proven to have a slope smaller than one and its graph has a unique intersection with the 45^o line. Hence there is a unique equilibrium strategy c^*_t such that $V_t(c^*_t) = c^*_t$). (This equilibrium is much stronger than a Nash equilibrium because it is strongly rationalizable).

[13] In this model, agents could use the fact that θ_t takes discrete values in order to obtain more information from the observation of Y_t. However, this feature is spurious. The random changes of θ_t could be defined such that the distribution of θ_t has a piecewise linear density function in every period. The previous updating formula should therefore be understood as the relevant formula for the "nodes" of the density function of θ_t, (at integer values of θ_t). The entire distribution of θ_t could be recovered through a linear interpolation.

However, the model with observational noise cannot exhibit a unique equilibrium for *all* values of the random noise. Suppose for example that the economy is in a low state and that the distribution of costs is such that there are two equilibria under perfect information. A very high value of the noise in some period may induce a large mass of agents to act in the next period. This could reveal a large amount of information, and generate two equilibria for the next period.

The main purpose of the model in this section is not to show that there is a unique equilibrium for all realizations of (θ_t, η_t). It is to show that the properties of the analytical model apply for most of these realizations: under the types of uncertainty and heterogeneity which are relevant in macroeconomics or in other contexts of social behavior, the model generates a SREE for most periods. In the numerical model below, there is a SREE in each of the 600 periods which are considered.

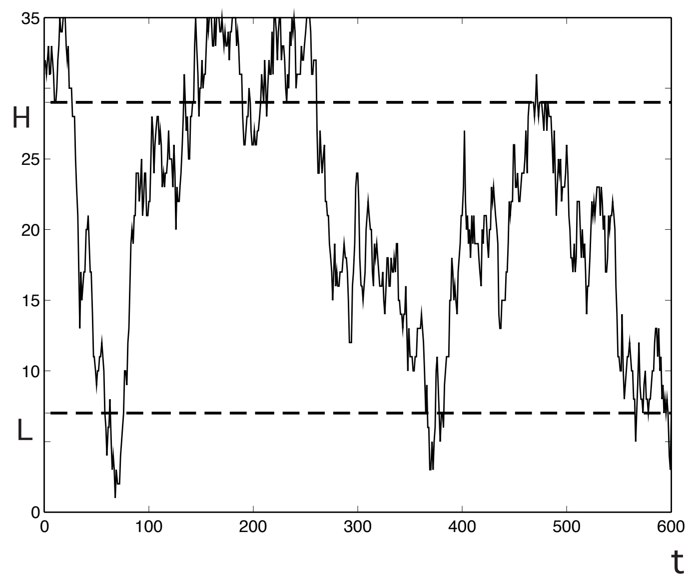

Under perfect information, if $\theta > \theta_H$ ($\theta < \theta_H$), the equilibrium is unique with a low (high) level of activity. In the middle band there are two equilibria with high and low activity.

Figure 6.3: The realization of the random path of θ

The numerical example

The properties of the model are illustrated for a particular realization of the random walk of θ_t that is represented in Figure 6.3. In the region $\theta \leq 7$, there is only one equilibrium under perfect information with low activity. In the region $\theta \geq 29$, there is only one equilibrium under perfect information, with high activity. The sum of the stationary probabilities of these two events is less than 1/2. In the simulation, the values of η_t are set to zero but of course, unbeknownst to the agents. Note that in the first period of the simulation in Figure 6.3, θ is in the high region with a unique regime of low activity.[14]

The first regime switch, from high to low, takes place in period 61. The public beliefs and the CVF just before and after the switch are represented in figure 6.4. On the left panel, the vertical line indicates the true value of θ_t and the curve is the graph of the probability distribution of θ_t in the public information. The right panel presents the graph of the CVF.

Just before the switch, in period 60, the public belief is completely off the mark: the actual value of the fundamental θ_t is very low while the public belief puts strong probabilities on high values of the fundamental. Because the public believes that the individual costs are high, the CVF is low in the range $(0, 20)$ and there is a unique equilibrium, which in this case is a SRE, with a low aggregate activity.

Just after the switch, in period 61, the public belief has completely changed while the fundamental has barely moved. The CVF has shifted up. In the equilibrium, which is also a SRE

[14] The parameters of the model are chosen such that the random walk is symmetric with $p = 1/3$, and has five independent steps within each period (which is defined by the observation of the aggregate activity). There is a mass of agents equal to 2 who have negative private costs. The first sub-population has a uniform density equal to $\beta = 0.5$. The other parameters are $\sigma_\theta = 1.5$, $\sigma_\eta = 1$ and $K = 35$. The mass of the cluster is equal to 14.

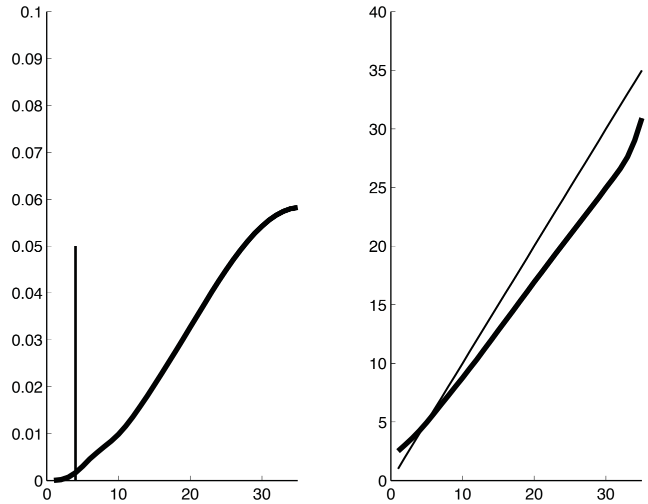

t=60

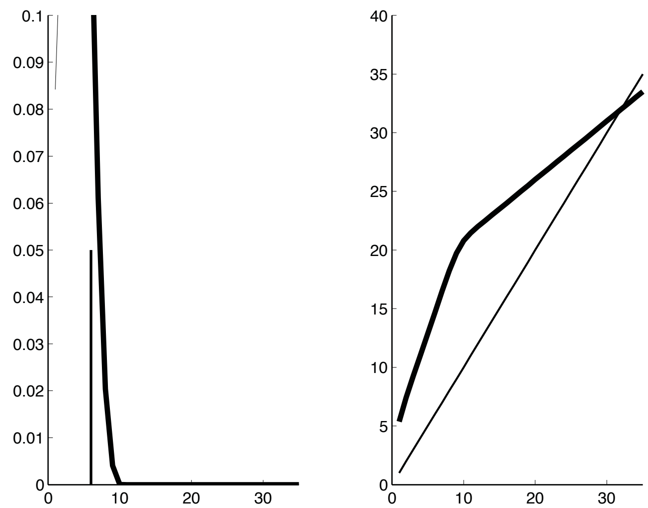

t=61

On the left panel, the vertical line indicates the true value of $\theta 60$ and the curve represents the probability distribution of θ_{60} according to the public information at the beginning of period 60. The right panel presents the graph of the CVF.
in this case, the level of activity jumps up to a new regime.
Figure 6.4: Public belief on θ and CVF before and after a switch

APPENDIX

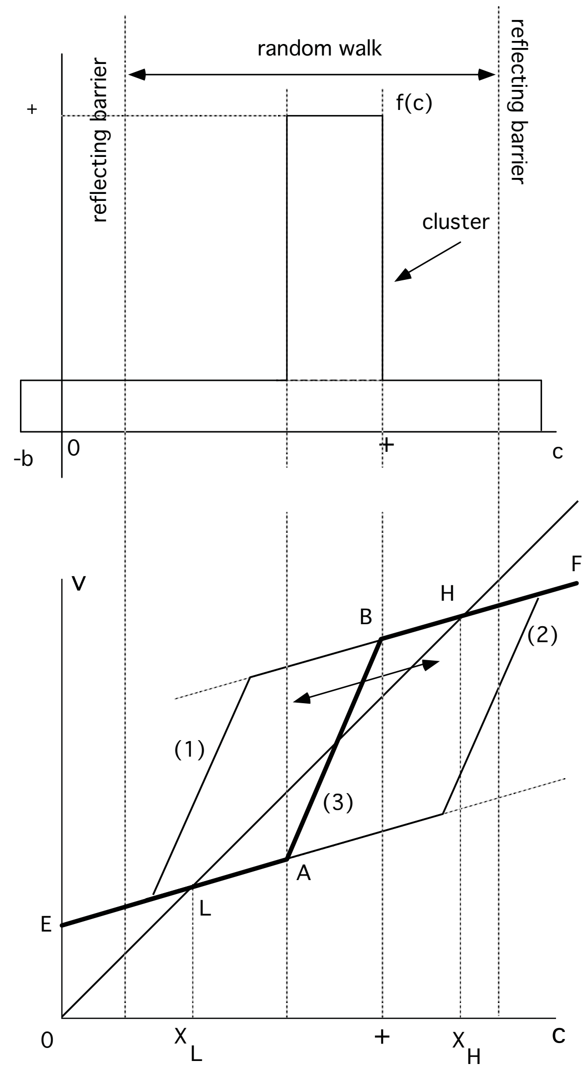

From the observation of the activity at the point L, the only information is that the left border of the high density (point A) is to the right of L.

Figure 6.5: A square distribution