Chapter 1

Bayesian Inference

(11/11/24 - minor edits)

A witness with no historical knowledge

There is a town where taxis come in two colors, yellow and red.[1] Ninety percent of the taxis are yellow. One night, a taxi hits a pedestrian and leaves the scene without stopping. The skill and ethics of the driver are not dependent on the color of the cab. An out-of-town witness claims that the color of the cab was red. The out-of-town witness does not out-of-town witness does not know the proportion of yellow and red cabs in the city and makes a report based solely on what he thinks he saw. Since the accident occurred at night, the witness is not completely reliable, but it has been estimated that such a witness makes a correct statement four out of five times (whether the true color of the taxi is yellow or red). How should we use the information from the witness? Because of the uncertainty, we should formulate our conclusion in terms of probabilities. Is it more likely that a red taxi was involved in the accident? Even though the witness says red and is correct 80 percent of the time, the answer is no.

Remember that there are many more yellow cabs. The red sighting can be explained either by a yellow cab hitting the pedestrian (a high *prior* event) being misidentified (a low probability event), or a red taxi (low probability) being correctly identified (high probability). Both the *prior* probability of the event and the precision of the signal must be used in evaluating the signal. Bayes' rule provides the method for computing probability updates. Let $\mathcal{R}$ be the event "a red taxi is involved" and $\mathcal{Y}$ be the event "a yellow taxi is involved". Similarly, let r (y) be the report "I saw a red (yellow) cab". The probability of

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

the event $\mathcal{R}$ conditional on the report r is denoted by $P(\mathcal{R}|r)$. By Bayes' rule,[2]

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

The probability that a red taxi 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 0.8. $P(r|\mathcal{Y})$ is the probability of an incorrect identification and is equal to 0.2. Hence,

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

This probability is less than half because the probability of a false "red" report (in the denominator, 0.2×0.9) is less than that of a correct sighting ($0.8\, times 0.1$). This is because there are so many yellow cabs (90 percent), and the observer makes a false report with a probability 0.2 that is not small.

A witness with historical knowledge

Now suppose the witness is a resident of the city who knows that only 10 percent of the taxis are red. In his testimony, he states the color that is most likely according to his rational inference. If he applies the Bayesian rule and knows his probability of making a mistake, he knows that it is more likely to be a yellow cab. He will report "yellow" even if he thinks he saw 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 saw) 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 receiver of the is trying to judge the most likely event: the witness and the receiver 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) who also thinks he saw a red taxi. This witness, trying to report the most likely event, also says "yellow". The receiver of the two reports learns nothing from them. For him, the accident was caused by a yellow cab with a probability of 90 percent.

Remember that when the first witness came from out of town, he was not informed of the local history and gave an informative report, "red. This report may be inaccurate, but it provides information. It also triggers more information from the second witness. After the first witness' report, the probability of $\mathcal{R}$ increased from 0.1 to 4/13. When this probability of 4/13 is communicated to the second witness, he thinks that a red car is more likely. likely.[3] So he 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.

[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] Exercise: prove it.

1.1 Remarks on the Bayesian model of learniingl

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 the learning problem would be trivial). An informative signal can be defined as 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.

We focus here on two types of Bayesian models. In the first, both the number of states and the number of signal values is finite. The model is discrete. For example

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 σ_θ^2.
   - Signal $s = \theta + \epsilon$, where s has a normal distribution with mean zero and variance σ_ϵ^2.

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.

1.2 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 report 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 the

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 1.2.1: Binary signal

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.

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

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,

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

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.

1.3 Multiple binary signals: searches 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".[4] No message had been sent by the crew but

[4] 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.

both “black boxes”–they are red– were retrieved after a two years. They have provided a 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. [5] 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[6] (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.

[5] See L. Stone **.

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

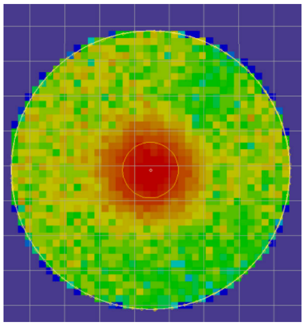

(A)
Prior probabilities

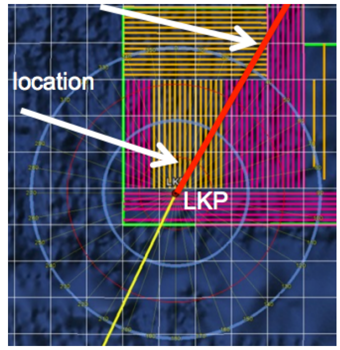

(B)
Search for pings

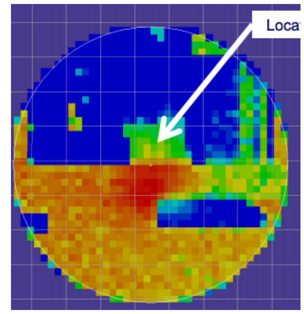

(C)
Posterior probabilities
after Stage 5

Wreckage

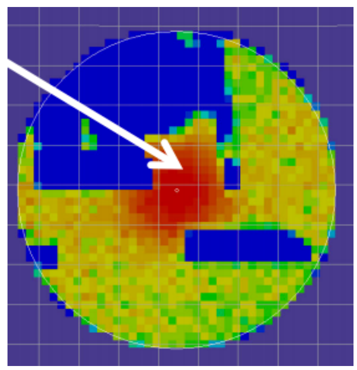

(D)
Posterior assuming beacons failed

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.

1.4 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.

1.5 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.

1.6 Learning may lead to opposite beliefs: polarization

(to be. revised)

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 the center policy is not working).

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

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.

1.7 Learning with a sequence of information and 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:

$$\text{History in period } t: \qquad h_t = \{s_1, \ldots, s_{t-1}\}. \tag{1.10}$$

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 **??**).

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 1.2. 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.

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

$$\lambda' = \lambda + a, \qquad \text{with the signal term} \quad a = \frac{P(s = 1|\theta = 1)}{P(s = 1|\theta = 0)}. \tag{1.11}$$

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!

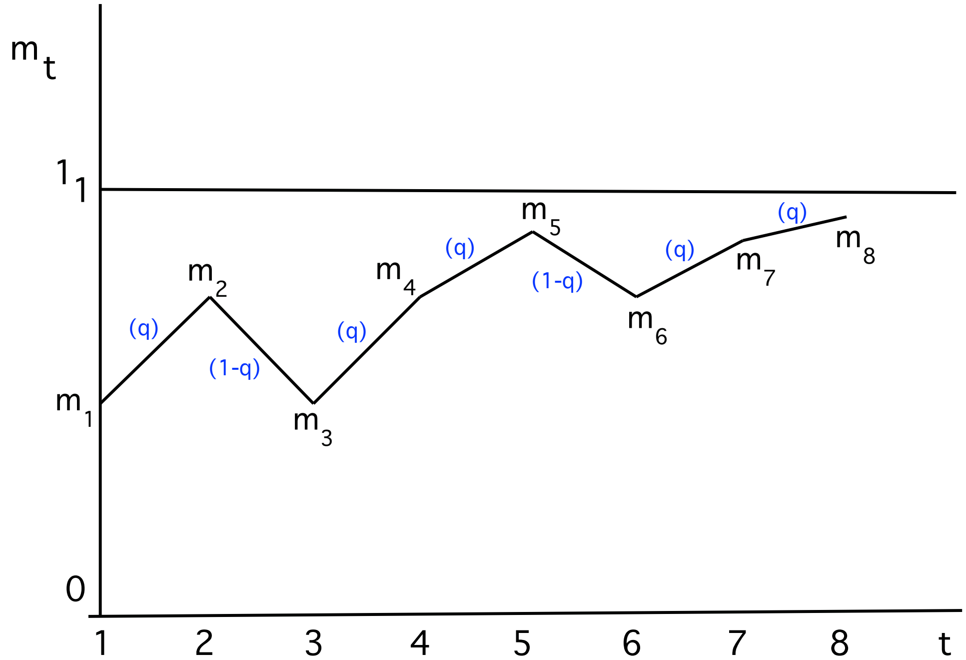

Figure 1.2: The evolution of belief as a stochastic process

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 1.11) that the Bayesian updating after some signal s_t is such that

$$\lambda_{t+1} = \lambda_t + a_t, \tag{1.12}$$

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 1.3.

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.

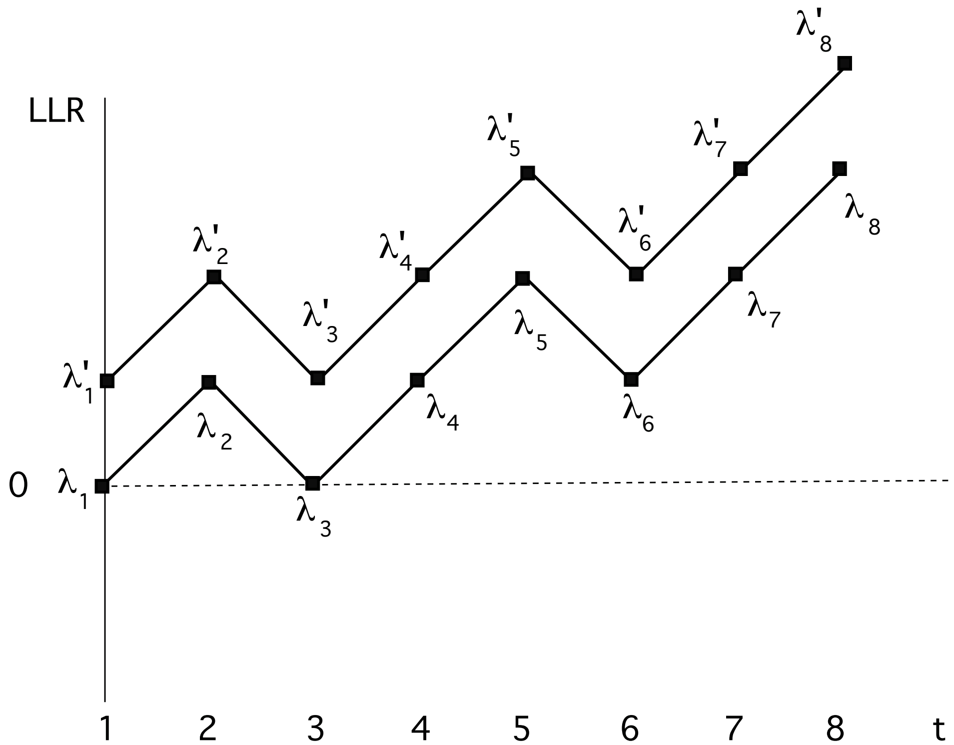

Figure 1.3: The evolution of LLs

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.

Bounded and unbounded private informations

Definition: *When there exists M such that in the equation (1.12) 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:

$$s = \theta + \epsilon, \qquad \text{with } \epsilon \sim \mathcal{N}(0, 1/\rho_\epsilon^2). \tag{1.13}$$

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

$$\lambda' = \lambda + \rho_\epsilon(\theta_1 - \theta_0)(s - \frac{\theta_1 + \theta_0}{2}). \tag{1.14}$$

  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.

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.

Stone, L.D. (1992). "Search for SS 'Central America': Mathematical Treasure Hunting," *Interfaces*, 22, 32-54.

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 that the wreck is 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

Bayesian Learning and Martingales

Bayesian learning is an application of the calculus of conditional probabilities and it generates often but not always, the remarkable property of a martingale (Proposition 1): an agent knows that he may change his belief after an observation, but the expected value of the change is nil: if that value were different from zero, the agent would change his belief right away, before getting the signal.

The previous argument applies to an agent who receives some signal about the state of nature and updates his information. In the context of social learning, if the observer shares the same information as the agent, before the private signal is revealed to the agent, for example the history of past actions, then the public belief (from the public information) will also be a martingale. However, if the observer does not know the pre-signal information of the agent, for example when the agent observes a random sample of previous actions, then the public belief will not be a martingale. This will be a source of significant difficulties.

The martingale property is similar to the efficient market equation in finance (and for good reasons). This simple property has a very powerful implication: in the process of Bayesian learning from a history h_t, the belief of a state of nature must converge, which does not mean that it converges to the truth.

The Martingale Convergence Theorem (MCT) is one of the most beautiful theorems in probability theory and very important in Bayesian learning. The theorem implies that unending fluctuations of individual beliefs cannot be compatible with rational learning.

The contradiction with unending cycles is actually at the core of its original proof which is due to Doob. (There have been others since). This proof is based on the fact that one cannot make a profit in an efficient market. Economists should have discovered and proven the theorem (as they should not have left the invention of optimal control to Pontryagin).

2.0.1 Martingales

Assume that information comes as a sequence of signals s_t, one signal per period, and that these signals have a distribution which depends on ω. They may or may not be independent, conditional on ω, and their distribution is known.

The expected value of ω, conditional on h_t is denoted by $\mu_t = E[\omega|h_t]$. Because the history h_t is random, μ_t is a stochastic process, a sequence of random variables.

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

$$X_t = E[X_{t+1}|h_t]. \tag{2.1}$$

Suppose that an agent has a distribution on ω with mean $E[\omega]$ and receives a signal s with a distribution which depends on ω. By the rules for conditional expectations, $E\Big[E[\omega|s]\Big] = E[\omega]$, and the next result follows

PROPOSITION 2.1. *Let $\omega \in \Omega$ with a probability distribution and h_t the history at time t defined as a sequence of signals $\{s_1, ..., s_{t-1}\}$, each with a probability distribution that depends on ω. For any $\mathcal{A} \subset \Omega$, $\mu_t = P(\omega \in \mathcal{A}|h_t)$ is a martingale.*

These results depend on the assumption that ω and the signals belong to the same probability space. (ω and the signals belong to one σ-algebra with a probability measure). In the above proof for the discrete case, the joint values of ω and the signals determine a partition of their space in subsets that each have a probability.

2.1 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

[1] Useful references: Grimmet and Stirzaker (1992), Omer Tamuz: Notes on Probability Theory.

this book, any convergence of a random variable (for example, a belief) is a convergence in probabiity[2]:

DEFINITION 2.2. *A sequence of random variables $\{X_t\}$ converges in probability to X^* if for any $\alpha > 0$, $Lim_{t\to\infty} P(|X_t - X^*| > \alpha) = 0$.*

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.

The martingale property is a wonderful tool in Bayesian learning because of the Martingale Convergence Theorem (MCT). Consider a Bayesian rational agent who receives a sequence of signals. Let his belief be his subjective probability assessment of an event, $\{\omega \in \mathcal{A}\}$, for some fixed $\mathcal{A} \subset \omega$. Can the agent keep changing his belief in endless random fluctuations? Or does this belief converge to some value (possibly incorrect)? The answer is simple: it must converge.

The belief must converge because the probability assessment is a bounded martingale. The convergence of a bounded martingale is intuitive. This intuition can take two forms, each of which leads to a proof of the theorem. First, the martingale property is equivalent to the efficient market property for an Arrow-Debreu security. Therefore a strategy of "buy low and sell high" (in a sense to be defined), cannot make money. In the second intuition, the essence of the 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. The drunkard cannot bounce against the side of the alley!

THEOREM 2.1. *(Martingale Convergence Theorem) If X_t is a martingale with $|X_t| < M < \infty$ for some M and all t, then there exists a random variable X_∞ such that X_t converges to X_∞.*

The convergence holds under conditions that are weaker than in the previous statement. $E|X_|] < M$ is sufficient, the theorem applies to supermartingales, $(E[X_{t+1}|h_t \leq X_t)$, and the convergence is almost surely.

headiIntuition of Doob's proof The martingale property is the same as the efficient market condition with a zero interest rate. Therefore one should not be able to expect a profit

[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.

from trading. (If there were cycles, even random cycles, one may hope to expect a profit).

The first proof of MCT, by Doob in the early 50s, was based on an efficient market argument.[3] There have been other proofs since. (See the next result). Suppose that the martingale X_t is the price of an Arrow-Debreu security that delivers 1 if a specific state of nature is realized (after the trading). The number of trading periods, T, is arbitrarily large. WLOG, X_t is in $[0, 1]$. You are the manager of a fund and you have n agents in your firm. Agent i has the following instructions: buy one unit of the security at the market price when this price is less than $(i-1)/n$ and hold it until the first period when the price is higher than i/n. Then wait and buy again one unit in the first period when the price is below $i-1$. In period T, if you hold the asset, sell it. The process is illustrated in the next figure. Figure 2.1.

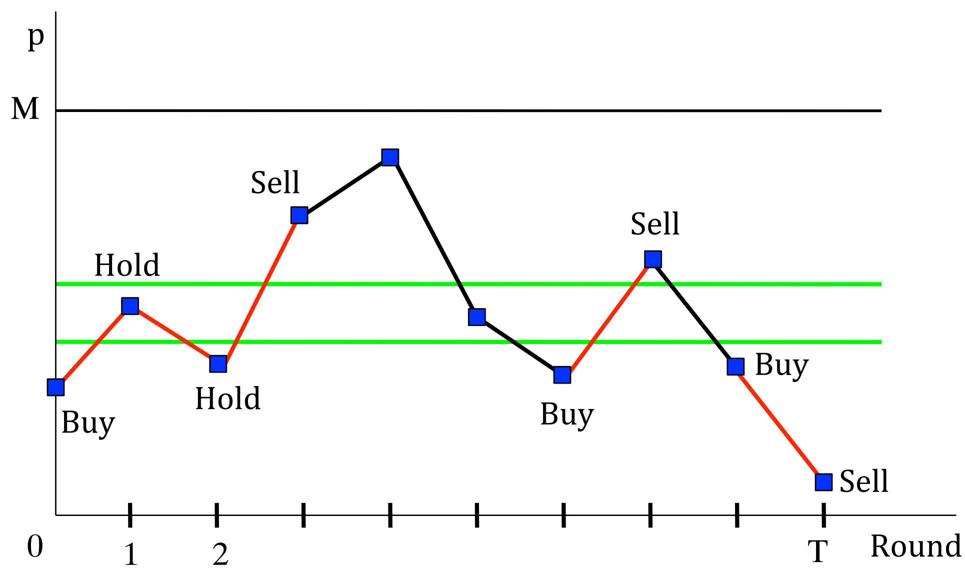

The agent holds one unit of the asset on the red segments.

Figure 2.1: A strategy of "buy low, sell high"

Define by N_T the number of times you buy the security until round T, that is the number of upwards crossings of the band $((i-1)/n, i/n)$ in the trajectory of the price, p_t. The maximum loss is 1 (if he has a stock that he sells in the last period). The net profit is not smaller than

$$V = N_T/n - 1.$$

Because of the martingale property, the expected gain from the trading strategy cannot be

[3]For a rigorous exposition, see the excellent notes of Omer Tamuz: Notes on Probability Theory.

positive. Hence, for any T,

$$E[N_T] \leq n.$$

For a given n, the expectation of the number of upward crossing of a band $((i-1)/i/n)$ is bounded. Since n is finite, the same applies to the set of all the bands. From this, one can show that the probability that X_t stays within the same band for $t > R$ tends to 1 when $T \to \infty$. Since n can be taken arbitrarily large, X_t converges to some value X^* in probabiiity.[4]

A heuristic remark on another proof of the Martingale Convergence Theorem

(To be edited)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 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 $k \geq 0$.

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

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.

---

[4] See Williams (1991).

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

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.

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.

PROPOSITION 2.5. *(Learning cannot be totally wrong, asymptotically)*

Let $\Omega = \{\omega, \dots, \omega_K\}$ be the finite set of states of nature, $\mu_t = \{\mu_t^1, \dots, \mu_t^K\}$ the probability assessment of a Bayesian agent in period t, and $\mu_1^1 > 0$ where ω_1 is the true state. Then for any $\epsilon > 0$,

$$P(\mu_t < \epsilon) < \epsilon/\mu_1^1.$$

If $\bar{\mu}^1$ is the limit value of μ_t^1, $P(\bar{\mu}^1 = 0) = 0$.

Under Bayesian learning, if the subjective distribution on ω converges to a point, it must converge to the truth.

Proof

For any history h_t, $P(h_t|\omega = \omega_1) = P(\omega = \omega_1|h_t)\dfrac{P(h_t)}{P(\omega = \omega_1)}$.

Let H_t be the set of histories h_t such that $\mu_t^1 < \epsilon$. By definition,

$$P(h_t \in H_t|\omega = \omega_1) < \epsilon\frac{P(h_t \in H_t)}{P(\omega = \omega_1)} \leq \epsilon\frac{1}{P(\omega = \omega_1)}$$

Q.E.D.

The likelihood ratio between two states ω_1 and ω_0 cannot be a martingale given the information of an agent. However, if the state is assumed to take a particular value, then the likelihood ratio may be a martingale. Proving it is a good exercise.

PROPOSITION 2.6. *Conditional on $\omega = \omega_0$, the likelihood ratio*

$\dfrac{P(\omega = \omega_1|h_t)}{P(\omega = \omega_0|h_t)}$ *is a martingale.*

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

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

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 $h_t = (x_1, \dots, x_{t-1})$.

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

(i) The belief of agent t

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

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

(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 $h_t = (x_1, \ldots, x_{t-1})$ is informationally equivalent to the sequence of signals $(s_1, \ldots, s_{t-1})$.

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 (Proposition 3.1 in the Appendix), he showed that the precision of the public information, ρ_t increases only like the cubic root of the number of observations, $At^{1/3}$. The value of the constant A depends on the observation noise, but the rate 1/3 is independent of that variance. Recall that with no noise, the precision increases linearly with t.

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

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

3.3.1 Large number of agents

The previous model is modified to allow for a continuum of agents. Each agent is indexed by $i \in [0,1]$ (with a uniform distribution) and receives one private signal *once* at the beginning of the first 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

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

[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 3.1 applies. The asymptotic evolutions of the public beliefs are the same in the two models.

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

3.3.2 Application: a market equilibrium

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

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

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 3.2.

Policy against mimetism

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

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

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. (history cannot be summarized by a number)

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

1. Show, using words and no algebra that the action of an agent reveals perfectly his private signal.

2. Can the history h_t be summarized by $\sum_{i \leq t-1} x_i$?

EXERCISE 3.2. (Endogenous private information)

In the standard Gaussian model of social learning, each agent has to pay 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:

[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.

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

4. Assume now that $c(q) = q^{\beta}$ with $\beta > 0$. Analyze the rate of convergence of social learning.

REFERENCES

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

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

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

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

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

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

****

3.5 APPENDIX

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

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

Proof

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

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

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

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

or

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

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

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

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

or

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

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

(i) The boundedness of b_t:

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

$$b_{t+1} < b_t\Big(1 + \frac{1}{2t}\Big). \tag{3.20}$$

Using (3.19) again, there exists $T > T_1$ such that for $t > T$,

$$b_{t+1} < b_t\Big(1 + \frac{1}{t}\Big)\Big(1 - \frac{b_t^3}{2At}\Big).$$

From this inequality, there is some value M such that if $b_t > M$ and $t > T$, then

$$b_{t+1} < b_t\Big(1 - \frac{1}{t}\Big). \tag{3.21}$$

We use (3.20) and (3.21) to show that if $t > T$, then $b_t < 2M$. Consider a value of $t > T$. If $b_{t-1} < M$, then by (3.20),

$$b_{t+1} < M\Big(1 + \frac{1}{t}\Big) < 2M.$$

If $b_{t-1} > M$, then by (3.21), $b_{t+1} < b_t$. It follows that b_t is bounded by the maximum of b_T and $2M$:

$$\text{for } t > T, \quad b_t < Max(b_T, 2M). \tag{3.22}$$

(ii) To show the convergence of b_t, one can extract a subsequence of b_t which converges to some limit ℓ_1. Then one can extract from this subsequence another subsequence such that b_{t+1} (defined by the previous equation) converges to a limit ℓ_2. Taking the limit,

$$\ell_2\left(1 - \frac{1}{3t} + O\Big(\frac{1}{t^2}\Big)\right) = \ell_1 - \frac{\ell_1^4}{At}\left(1 + O\Big(\frac{\ell_1}{t^{\frac{1}{3}}}\Big)\right).$$

We must have

$$\ell_1 = \ell_2, \quad \text{and} \quad \frac{\ell_2}{3} = \frac{\ell_1^4}{A}.$$

Therefore,

$$\ell_1 = \ell_2 = \ell = \Big(\frac{A}{3}\Big)^{\frac{1}{3}}.$$

The result follows from the definition of A.

□

Chapter 4

Cascades and herds

To be pruned (11/11/24)

Each agent observes what others do and takes a zero-one decision in a pre-ordered sequence. In a cascade, all agents herd on a sufficiently strong public belief and there is no learning. In a herd, all agents turn out to take the same decision. A cascade generates a herd but the converse is not true. Cascades are non generic for atomless distributions of beliefs while a herd always takes place, eventually! Since a herd does take place eventually, the probability that it is broken must converge to zero. Hence, there is some learning in a herd (it is not broken), but the learning is very slow. The stylization of that property is the cascade.

Beliefs converge to the truth only if the distribution of private beliefs is unbounded, but the self-defeating principle in social learning implies that the convergence is slow. Since the filter imposed by discrete actions is coarse, the slowdown of social learning is much more significant than in the previous chapter. Applications for welfare properties and pricing policies by a monopoly are discussed.

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.

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

??: 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*.

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

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

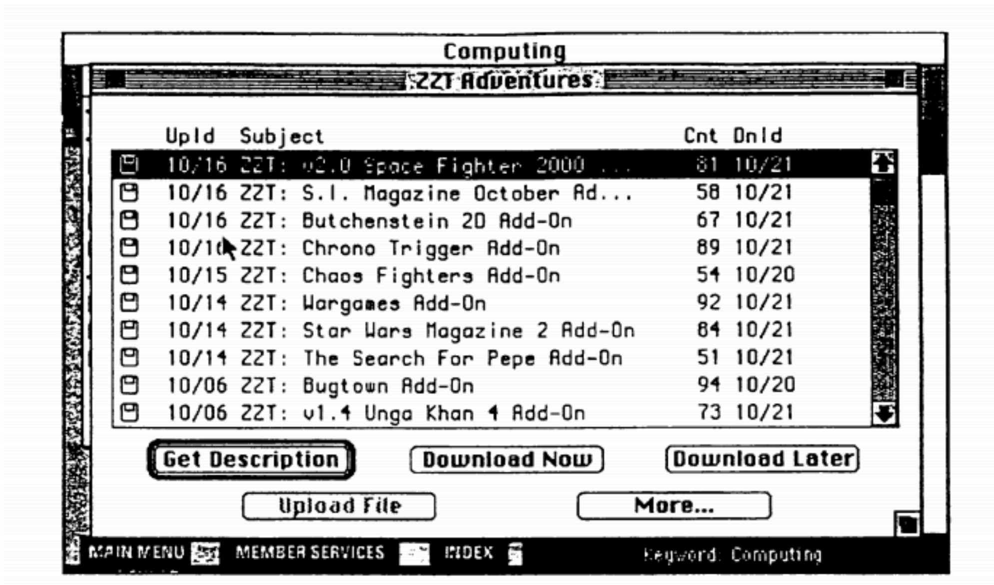

Figure 4.1: Applications for downloads

advertisements.

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 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 made of all the real numbers. Here the finite number of actions exerts a strong restriction on the vocabulary of social communication.

If there is a seminal study on social learning, it is the paper by Bikchandani, Hirshleifer and Welch (1992), hereafter BHW[2]. They introduced the definition of informational cascades in models of Bayesian learning. In a cascade, the public belief, which is gathered from the history of observations, dominates the private signal of any individual: the action of any agent does not depend on his private information. In a cascade, all agents are herding. Since actions do not convey private informations, nothing is learned and the cascade goes on forever, possibly with an incorrect action. The failure of social learning is spectacular.

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. A herd is defined as an outcome where all agents take the same action after some period. Which period is a random event and it is unknown at the start of the learning process. Actually, we we will see that it may never be known. In a herd not all agents may be herding. It is precisely because not all agents are herding in a herd that some learning takes place. The probability that the herd could be broken generates some information. But this probability must be vanishingly small for the herd to be sustained. Hence, the amount of social learning in a herd is very small.

**

Cascades do not occur, except in very special models which are not generic, while herds always take place eventually. The reader may think that cascades are therefore not important. Wrong: cascades are good approximations for the properties of the generic models of learning from others' actions when these actions are discrete.

**

Suppose that the set of states is finite. The *support* of a distribution of probabilities (beliefs) is the set of states that have a strictly positive belief (probability). If the set of the states is a continuum, for example a continuum of real numbers and the distribution of probabilities has a continuous distribution, the support of the the distribution is the subset where the density function is strictly positive. We will see that the process of social learning depends on whether the probability distributions, of the state or or a signal on the state, are bounded or not.

The general model is built on the models with bounded private beliefs which have been presented in Section ??. (The reader is advised to review that section if necessary). The evolution of the beliefs is presented in a diagram which will be used later in the book.

When the support is bounded, private beliefs become dominated by a public belief which is either optimistic or pessimistic, as the number of observations increases. Such a situation actually never occurs when private beliefs have a distribution without points of positive

[2] 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.

mass (which is not just a perturbation of a distribution with such points). However, the limit behavior of the model is closely approximated by cascades.

Beliefs converge to the truth, almost surely, only if the support of the distribution of beliefs is unbounded. In this respect, the results of BHW have been criticized as not robust. Such theoretical focus on the limit beliefs is misleading. What matters is the speed of convergence.

Section 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.

The social optimum

In an equilibrium, no agent takes into account the externality created by his action for the information of others. In a social optimum, this externality is taken into account (as in the model with actions in a continuum, Section 3.4). A social optimum is constrained in the sense that each agent "speaks" to others only through his action. An agent has a decision rule according to which his action depends on his private belief and the public belief. He can reveal his private belief only through his action. He departs from the selfish rule of using history for his own payoff only if the externality provided to others outweighs the personal loss.

In Section **??**, it is shown that the social optimal rule is to forget history if the belief from history—the public belief—is in some interval of values, and to herd otherwise. If the belief is outside of that "interval of experimentation", there is no social learning anymore. The socially optimal rule may be implemented by setting a price of investment contingent on the public belief.

Monopoly pricing of a new good

A monopoly who captures some consumer surplus will take into account the benefit of experimentation for the future. This problem is considered in Section **??**. A monopoly introduces on the market a new good of imperfectly known quality. The optimal strategy is divided in two phases. The first is the "elitist phase": the price of the good is relatively

[3] For this section, I have greatly benefited from the insights of Lones Smith and I am very grateful to him.

high. Only the agents with a good signal on the good buy and the volume of sales raises the estimate of the other agents. When this estimate is sufficiently high, the monopoly lowers the price to reach all customers.

The incentive to learn is inversely related to the discount rate. If the discount rate is vanishingly small, the difference between the level of social welfare and the monopoly profit converges to zero. At the limit, the monopoly follows a strategy which is socially optimal. (Monopoly profits are redistributed).

4.1 The basic model of herding

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: (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)[4].

4.1.1 The 2 by 2 by 2 model

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[5] with a probability μ_1 for the "good" state $\theta = 1$.

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 $q > 1/2$: $P(s_t = \theta \mid \theta) = q$.

[4] 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 ??.

[5] 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).

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.[6] 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. By convention, if the payoff of $x = 1$ is zero, the agent chooses $x = 0$.

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. 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 $\theta = 1$.

4.1.2 Informational cascades

Agents with a good signal $s = 1$ will be called optimists and agents with a bad signal $s = 0$ will be called pessimists. An agent combines the public belief with his private signal to form his belief. If μ is the public belief in some arbitrary period, the belief of an optimist is higher than μ and the belief of a pessimist is lower. Let μ^+ and μ^- be the beliefs of the optimists and the pessimists[7]: $\mu^- < \mu < \mu^+$.

A pessimist invests if and only if his belief μ^- is greater than the cost c, *i.e.* if the public belief is greater than some value $\mu^{**} > c$. (If $c = 1/2$, $\mu^{**} = q$). If the public belief is such that a pessimist invests, then *a fortiori*, it induces an optimist to invest. Therefore, if $\mu_t > \mu^{**}$ agent t invests whatever his signal. If $\mu_t \leq \mu^{**}$, he does not invest if his signal is bad.

[6] In the tale of two restaurants, c is the opportunity cost of not eating at the other restaurant.

[7] By Bayes' rule,

$$\mu^- = \frac{\mu(1-q)}{\mu(1-q) + (1-\mu)q} < \mu < \frac{\mu q}{\mu q + (1-\mu)(1-q)} = \mu^+.$$

Likewise, let μ^* be the value of the public belief such that $\mu^+ = c$. If $\mu_t \leq \mu^*$, agent t does not invest no matter the value of his private signal. If $\mu_t > \mu^*$ he invests if he has a good signal. The cases are summarized in the next result.

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.

Cascades and herds

Proposition 4.1 shows that if the public belief, μ_t, is above μ^{**}, agent t invests and ignores his private signal. His action conveys no information on this signal. Likewise, if the public belief is smaller than μ^*, then the agent does not invest. This important situation deserves a definition.

DEFINITION 4.1. *An agent herds on the public belief when his action is independent of his private signal.*

The herding of an agent describes a decision process. The agent takes into account only the public belief; his private signal is too weak to matter. If all agents herd, no private information is revealed. The public belief is unchanged at the beginning of the next period and the situation is identical: the agent acts according to the public belief whatever his private signal. The behavior of each agent is repeated period after period. This situation has been described by BHW as an *informational cascade*. The metaphor was used first by Tarde at the end of the nineteenth century.

DEFINITION 4.2. *If all agents herd (Definition 4.1), there is an informational cascade.*

We now have to make an important distinction between the herding of all agents in an informational cascade and the definition of a herd.

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

In a cascade, all agents are herding and make the same decision which depends only on the public belief (which stays invariant over time). Hence, all actions are identical.

PROPOSITION 4.2. *If there is an informational cascade in period t, there is a herd in the same period.*

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. We will see later that generically, cascades do not occur, but herds eventually begin with probability one! Why do we consider cascades then? Because their properties are stylized representations of models of social learning.

In the present model, an informational cascade takes place if $\mu_t > \mu^{**}$ or $\mu_t \leq \mu^*$. There is social learning only if $\mu^* < \mu_t \leq \mu^{**}$. Then $x_t = s_t$ and the action reveals perfectly the signal s_t. The public belief in period $t+1$ is the same as that of agent t as long as a cascade has not started. The history of actions $h_t = (x_1, \ldots, x_{t-1})$ is equivalent to the history of signals $(s_1, \ldots, s_{t-1})$.

Assume that there is no cascade in periods 1 and 2 and that $s_1 = 1$ and $s_2 = 1$. Suppose that agent 3 is a pessimist. Because all signals have the same precision, his bad signal "cancels" one good signal. He therefore has the same belief as agent 1 and should invest. There is a cascade in period 3.

Likewise, two consecutive bad signals ($s = 0$) start a cascade with no investment, if no cascade has started before. If the public belief μ_1 is greater than c and agent 1 has a good signal, a cascade with investment begins in period 2. If $\mu_1 < c$ and the first agent has a bad signal, he does not invest and a cascade with no investment begins in period 2.

In order *not to* have a cascade, a necessary condition is that the signals alternate consecutively between 1 and 0. We infer that

- the probability that a cascade has not started by period t converges to zero exponentially, like β^t for some parameter $\beta < 1$;
- there is a positive probability that the cascade is wrong: in the bad states all agents may invest after some period, and investment may stop after some period in the good state;
- beliefs do not change once a herd has started; rational agents do not become more confident in a cascade.

PROPOSITION 4.3. *When agents have a binary signal, an informational cascade occurs after some finite date, almost surely. The probability that the informational cascade has* *not started by date t converges to* 0 *like* β^t *for some* β^t *with* $0 < \beta < 1$.

A geometric representation

The evolution of the beliefs is represented in Figure **??**. 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, μ_1^+, is above c while the belief of a pessimist, μ_1^-, is below c. 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, μ_2, is identical to the belief of the first agent: $\mu_2 = \mu_1^-$. The sequence of the signal endowments is indicated in the figure. *When there is social learning, the signal of agent t is integrated in the public information of period $t+1$.* Using the notation of the previous chapter, $\mu_{t+1} = \tilde{\mu}_t$.

Consider now period 5 in the figure: agent 5 is an optimist, invests and reveals 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 $\mu_6 = \mu_5^+$. The belief of a pessimist in period 6 is now higher than the cost c (here, it is equal to the public belief μ_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 cascade takes place because all the beliefs are above the cut-off level c. This condition is met here because the public belief μ_6 is strictly higher than μ^{**}. Since μ_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 μ^{**} in period 5. A cascade takes place because of the high belief of the last agent who triggers the cascade. Since this property is essential for the occurrence of an informational cascade, it is important and will be discussed later in more details.

In this simple model, the public belief $\mu_t = P(\theta = 1|h_t)$ converges to one of two values (depending on the cascade). From the Martingale Convergence Theorem, we knew μ_t would necessarily converge in probability. The exponential convergence is particularly fast. The informational cascade may be incorrect however: all agents may take the wrong decision.

Black sheeps

Assume there is a cascade in some period T in which agents invest whatever their signal.

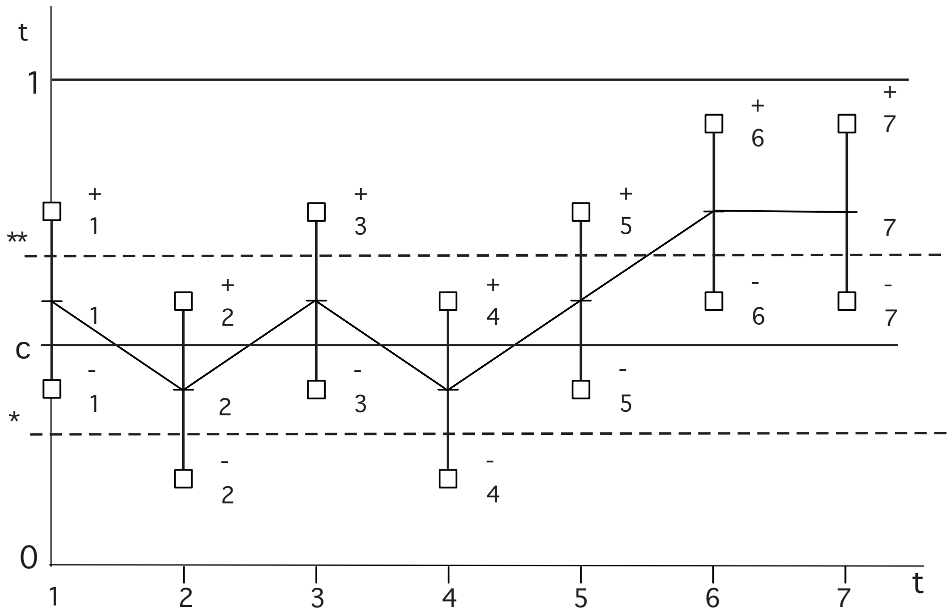

In each period, the middle of the vertical segment is the public belief, while the top and the bottom of the segment are the beliefs of an optimist (with a private signal $s = 1$) and of a pessimist (with signal $s = 0$). The private signals are $s_1 = 0, s_2 = 1, s_3 = 0, s_4 = 1, s_5 = 1$.

Figure 4.2: Cascade representation

Extend now the previous setting and assume that agent T may be of one of two types. Either he has a signal of precision q like the previous agents, or his precision is $q' > q$ and q' is sufficiently high with respect to the public belief that if he has a bad signal ($s_T = 0$), he does not invest. The type of the agent is private and therefore not observable, but the possibility that agent T has a higher precision is known to all agents.

Suppose that agent T does not invest: $x_T = 0$. What inference is drawn by others? The only possible explanation is that agent T has a signal of high precision q' and that his signal is bad: the information of agent T is conveyed *exactly* by his action.

If agent T invests, his action is like that of others. Does it mean that the public belief does not change? No! The absence of a black sheep in period T (who would not invest) increases the confidence that the state is good. Social learning takes place as long as not all agents herd. The learning may slow down however as agents with a relatively low precision begin to herd. The inference problem with heterogeneous precisions requires a model which incorporates the random endowment of signals with different precisions. A model with two types of precision is presented in the appendix.

The simple model has served two useful purposes: (i) it is a lesson on how to begin to think formally about a stylized fact and the essence of a mechanism; (ii) it strengthens the intuition about the mechanism of learning and its possible failures. These steps need to be as simple as possible. But the simplicity of the model could generate the criticism that its properties are not robust. The model is now generalized and we will see that its basic properties are indeed robust.

4.2 The standard model with bounded beliefs

We now extend the previous model to admit any distribution of private beliefs as described in Section **??**. Such a distribution is characterized by the *c.d.f.* $F^\theta(\mu)$ which depends on the state θ. Recall that the *c.d.f.*s satisfy the Proportional Property (**??**) and therefore the assumption of first order stochastic dominance: for any μ in the interior of the support of the distribution, $F^{\theta_0}(\mu) > F^{\theta_1}(\mu)$. By an abuse of notation, $F^\theta(\mu)$ will represent the *c.d.f.* of a distribution of the beliefs measured as the probability of θ_1, and $F^\theta(\lambda)$ will represent the *c.d.f.* of a distribution of the LLR between θ_1 and θ_0.

We keep the following structure: two states $\theta \in \{\theta_0, \theta_1\}$, two actions $x \in \{0, 1\}$, with a payoff $(E[\theta] - c)x$, $\theta_0 < c < \theta_1$. The states θ_1 and θ_0 will be called "good" and "bad". We may take $\theta_0 = 1$ and $\theta_0 = 0$, but the notation may be clearer if we keep the symbols θ_1 and θ_0 rather than using numerical values.

4.2.1 Social learning

At the end of each period, agents observe the action x_t. Any belief λ is updated using Bayes' rule. This rule is particularly convenient when expressed in LLR as in equation (**??**) which is repeated here.

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

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 $\pi_t(\theta) = 1 - F_t^\theta(\gamma)$.

Observations

States of Nature

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

with $\gamma = Log\Big(\frac{c}{1-c}\Big)$.

The action in period t, $x_t \in \{0, 1\}$, provides a binary random signal on θ with probabilities described in Table **??**. Since the *c.d.f.* F^{θ_1} dominates F^{θ_0} in the sense of first order stochastic dominance (Proposition **??**), there are more optimistic agents in the good than in the bad state on average. Hence, the probability of investment is higher in the good state, and the observation $x_t = 1$ raises the beliefs of all agents.

Following the observation of x_t, the updating equation (**??**) takes the particular form

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

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 $F_t^{\theta_1}(\gamma) \neq F_t^{\theta_0}(\gamma)$.

Since the distribution of LLRs is invariant up to a translation, it is sufficient to keep track of one of the beliefs. If the support of beliefs is bounded, we choose the mid-point of the support, called by an abuse of notation the public belief. If the support is not bounded, the definition of the public belief will depend on the particular case.

The Markov process

The previous process has an abstract formulation which may provide some perspective on the process of social learning. We have seen that the position of the distribution in any period can be characterized by one point λ_t. Let μ_t be the belief of an agent with LLR equal to λ_t. The Bayesian formula (4.2) takes the general form $\mu_{t+1} = B(x_t, \mu_t)$ and x_t is a random variable which takes the value 1 or 0 according to Table **??**. These values depend only on λ_t and therefore μ_t depends on θ. The process of social learning is summarized by the equations

$$\begin{cases} \mu_{t+1} = B(\mu_t, x_t), \\ \\ P(x_t = 1) = \pi(\mu_t, \theta). \end{cases} \tag{4.3}$$

The combination of the two equations defines a Markov process for μ_t. Such a definition is natural and serves two purposes. It provides a synthetic formulation of the social learning. It is essential for the analysis of convergence properties. However, such a formulation can be applied to a wide class of processes and does not highlight specific features of the structural model of social learning with discrete actions.

4.2.2 Bounded beliefs

Assume the initial distribution of private beliefs is *bounded*. Its support is restricted to a finite interval $(\underline{\lambda}_1, \overline{\lambda}_1)$. This case is represented in Figure **??**. Let λ_t be the public belief in period t, *i.e.*, the mid-point of the support: $\lambda_t = (\underline{\lambda}_t + \bar{\lambda}_t)/2$ and let $\sigma = (\bar{\lambda}_t - \underline{\lambda}_t)/2$, a constant. If λ_t is greater than the value $\lambda^{**} = \gamma + \sigma$, the support of the distribution is above γ and agent t invests, whatever his belief. Likewise, if $\lambda \leq \lambda^* = \gamma - \sigma$, no agent invests. In either case, there is an informational cascade. There is no informational cascade as long as the public belief stays in the interval $(\lambda^*, \lambda^{**}) = (\gamma - \sigma, \gamma + \sigma)$. The complement of that interval will be called the *cascade set*.

Figure **??** is drawn under the assumption of an atomless distribution of beliefs but it can also be drawn with atoms as in Figure **??**.

i We know from the Martingale Convergence Theorem that the probability of the good state, $\mu_t = e^{\lambda_t}/(1 + e^{\lambda_t})$, converges in probability. Hence, λ_t must converge to some value. Suppose that the limit is not in the cascade set. Then, asymptotically, the probability that $x_t = 1$ remains different in states θ_1 and θ_0. Hence, with strictly positive probability, the common belief is updated by some non vanishing amount, thus contradicting the convergence of the martingale. This argument is used in the Appendix to prove that λ_t must converge to a value in the cascade set.

PROPOSITION 4.4. *Assume that the support of the initial distribution of private beliefs is $I = [\lambda_1 - \sigma, \lambda_1 + \sigma]$. Then λ_t converges almost surely to a limit $\lambda_\infty \notin (\gamma - \sigma, \gamma + \sigma)$ with $\gamma = Log(c/(1 - c))$.*

Is the occurrence of a cascade generic?

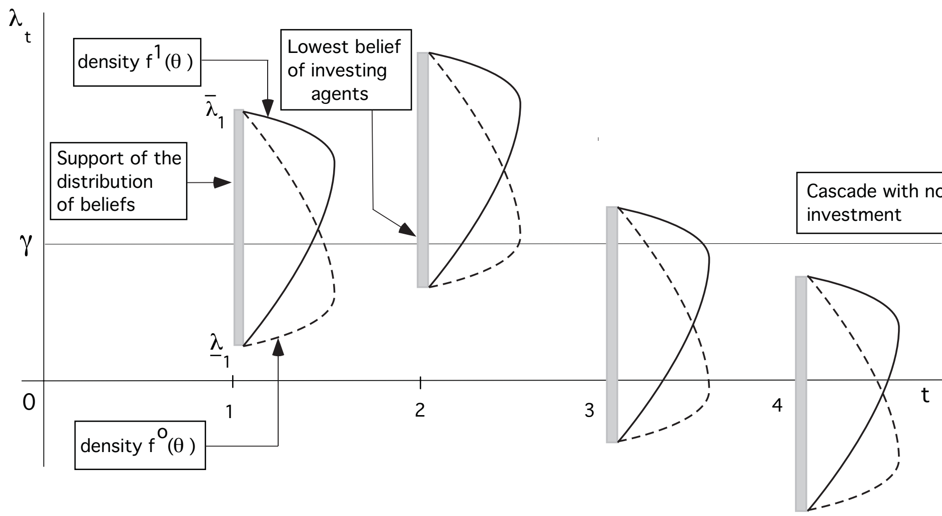

In each period, the support of the distribution of beliefs (LLR) is represented by a segment. The action is $x_t = 1$ if and only if the belief (LLR) of the agent is above γ. If agent t happens to have a belief above (below) γ, the distribution moves up (down) in the next period $t + 1$. If the entire support is above (below) γ, the action is equal to 1 (0) and the distribution stays constant.

Figure 4.3: Cascade representation

Generically, cascades do not occur!

The previous result shows that the beliefs tend to the cascade set. But for an arbitrary distribution of initial beliefs, is this convergence as fast as with the discrete beliefs of Figure **??**, or is it slow? It turns out that, for most distributions which are "smooth", the convergence is slow and cascades do not occur.

The mechanism can be explained simply. Suppose that the beliefs converge to the upper part of Figure **??** where agents invest. The probability that agent t invests is lower in the bad than in the good state, but as the beliefs move upwards, these two probabilities converge to each other with a common limit equal to one. The observation of an investment conveys a vanishingly small amount of information and the upward shift of the beliefs is also vanishingly small[8].

Assume the distribution of initial beliefs has a density $f^\theta(\mu)$ in state θ such that

$$f^1(\mu) = \mu\phi(\mu), \quad \text{and} \quad f^0(\mu) = (1-\mu)\phi(\mu), \tag{4.4}$$

[8] The argument does not apply when beliefs are high and there is no investment in the period. In that case, the probability of no investment is low in both states, but the ratio between these probabilities is not small.

for some function $\phi(\mu)$ with a support in $[a, 1-a]$, $a > 0$. This distribution is "natural" in the sense that it is generated by a two-step process in which agents draw a SBS of precision μ with a density proportional to $\phi(\mu)$. A simple case is provided by a uniform distribution of precisions where ϕ is constant. The proof of the following result is left to the reader.

PROPOSITION 4.5. *Assume that the density of initial beliefs are proportional to μ and to $1-\mu$ in the two states. If there is no cascade in the first period, there is no cascade in any period.*

The result applies if $\phi(\mu)$ does not put too much mass at either end of its support. This is intuitive: in the model with discrete beliefs (Figure **??**), all the mass is put at either end of the support. A smooth perturbation of the discrete model does not change its properties. A numerical simulation of the case which satisfies (4.4) shows that a cascade occurs if $\phi(\mu) = x^{-n}$ with n sufficiently high ($n \geq 4$ for a wide set of other parameters). In this case, the distribution puts a high mass at the lower end of the support.

Right and wrong cascades

A cascade may arise with an incorrect action: for example, beliefs may be sufficiently low that no agent invests while the state is good. However, agents learn rationally and the probability of a wrong cascade is small if agents have a wide diversity of beliefs as measured by the length of the support of the distribution.

Suppose that the initial distribution in LLR is symmetric around 0 with a support of length 2σ. We compute the probability of a wrong cascade for an agent with initial belief 1/2. A cascade with no investment arises if his LLR λ_t is smaller than $\gamma - \sigma$, *i.e.*, if his belief in level is such that

$$\mu_t \leq \epsilon = e^{\gamma-\sigma}/(1 + e^{\gamma-\sigma}).$$

When the support of the distribution in LLR becomes arbitrarily large, $\sigma \to \infty$ and ϵ is arbitrarily small. From Proposition **??** with $\mu_1 = 1/2$, we know that

$$P(\mu_t \leq \epsilon | \theta_1) \leq 2\epsilon.$$

The argument is the same for the cascades where all agents invest. The probability of a wrong cascade for a neutral observer (with initial belief 1/2) tends to zero if the support of the distribution in LLR becomes arbitrarily large (or equivalently if the beliefs measured as probabilities of θ_1 are intervals converging to $(0, 1)$).

PROPOSITION 4.7. *If the support of the initial distribution of LLRs contains the interval $[-\sigma, +\sigma]$, then for an observer with initial belief 1/2, the probability of a wrong cascade is less than 4ϵ, with $\epsilon = e^{-\sigma}c/\big(1 - c + e^{-\sigma}c\big)$.*

4.3 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.

4.3.1 Unbounded beliefs: convergence to the truth

From Proposition 4.7 (with $\sigma \to \infty$), we have immediately the next result.

PROPOSITION 4.8. *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?

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

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 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 **??**. One observes immediately that the pattern is similar to a cascade in the BHW model with the occurrence of "black sheeps".

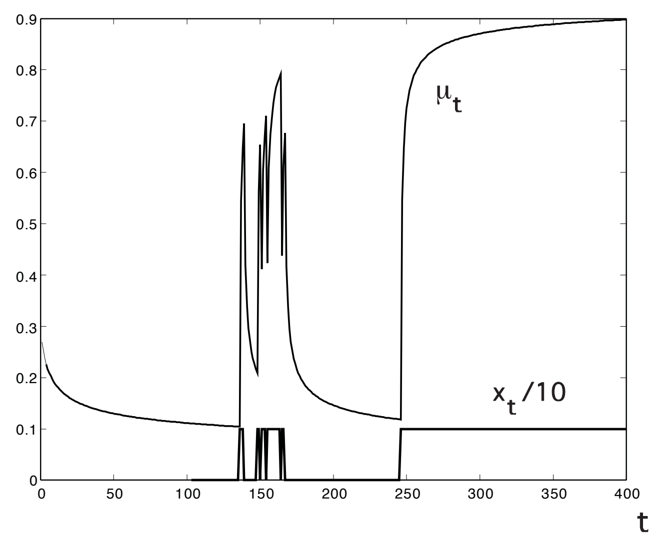

The upper graph represents the evolution of the public belief. The lower graph represents the sequence of individuals' actions. It is distinct from the horizontal axis only if $x_t = 1$.

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 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.

The figure 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. Cascades provide a better stylized description of the properties of social learning through discrete actions than the convergence result of Proposition 4.8. The properties of the example are confirmed by the general analysis of the convergence in Section **??**.

4.4 Herds and the slow convergence of beliefs

4.4.1 Herds

The Martingale Convergence Theorem implies that the public belief converges almost surely. Assume that the distribution of beliefs is bounded. At the limit, the support of the distribution must be included in one of the two cascade sets. Suppose that on some path the support of the distribution converges to the upper half of the cascade set where all agents invest: $\underline{\mu}_t \to c$. We now prove by contradiction that the number of periods with no investment is finite on this path.

Since there is a subsequence $x_n = 0$, we may assume $\underline{\mu}_n < c$. Following the observation of $x_n = 0$, Bayes' rule implies

$$\lambda_{n+1} = \lambda_n + \nu_n, \quad \text{with} \quad \nu_n = Log\Big(\frac{F^1(\underline{\lambda}_1 + z_n)}{F^0(\underline{\lambda}_1 + z_n)}\Big), \quad \text{and} \quad z_n = \gamma - \underline{\lambda}_n.$$

By the assumption of first order stochastic dominance, if $z_n \to 0$, there exists $\alpha < 0$ such that $\eta_n < \alpha$, which contradicts the convergence of λ_n: the jump down of the LLR contradicts the convergence. The same argument can be used in the case of an unbounded distribution of beliefs.

All paths with social learning end with a herd.

THEOREM 4.1. *On any path $\{x_t\}_{t\geq 1}$ with social learning, a herd begins in finite time. If the distribution of beliefs is unbounded and $\theta = \theta_1$ ($\theta = \theta_0$), there exists T such that if $t > T$, $x_t = 1$ ($x_t = 0$), almost surely.*

This result is due to Smith and Sørensen (2001). It shows that herds take place eventually although, generically, not all agents are herding in any period!

4.4.2 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 $\mu_t \rightarrow 0$, 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

$$f^1(\mu) = \mu\phi(\mu), \;\; f^0(\mu) = (1-\mu)\phi(\mu), \tag{4.5}$$

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 4.9. *Assume the distributions of private beliefs in the two states satisfy (4.5) with $\phi(0) > 0$, and that $\theta = 0$. Then, in a herd with $x_t = 0$, if $t \rightarrow \infty$, 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 $t\mu_t \rightarrow a$ 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 4.9 characterized the rate of convergence in a herd and it can be used to prove the following result[10].

THEOREM 4.2. *Assume the distributions of private beliefs satisfy (4.5) 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 is generalized by the following result.

PROPOSITION 4.10. *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 4.10 shows that the social learning through the observation of discrete actions is much slower, "exponentially slower[11]", than if private informations were publicly observable.

4.4.3 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 (4.5) with* $\phi(0) > 0$ *and* $\phi(1) > 0$*. Then the probability that a herd has not started by date* t *tends to* 0 *like* $1/t$*.*

[10] See Chamley (2002).

[11] Smith and Sørensen (2001) provide a technical result (Theorem 4) which states that the Markov process defined in (4.3) 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 **??**).

4.4.4 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 $b > 0$, $c > 0$,

$$f^1(1) = 0, \quad \text{and} \quad Lim_{\mu \to 0}\Big(f^1(\mu)/(1-\mu)^b\Big) = c > 0. \tag{4.6}$$

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 Crashes and booms

The stylized pattern of a herd which is broken by a sudden event is emblematic of a pattern of "business as usual" where at first beliefs change little, then some event generates a crash or a boom, after which the new beliefs seem "obvious" in a "wisdom after the facts". This sequence has been illustrated by Caplin and Leahy (1994). Their assumption of endogenous timing is not necessary for the property.

In each period, there is a new population of agents which forms a continuum of mass one. Each agent has a private information on θ in the form of a Gaussian signal $s_t = \theta + \epsilon_t$, where ϵ_t has a normal distribution $\mathcal{N}(0, \sigma_\epsilon^2)$ and is independent of other variables. Each agent chooses a zero-one action $x \in \{0, 1\}$.

In period t, agents know the history $h_t = \{Y_1, \ldots, Y_{t-1}\}$ of the aggregate variable $Y_t = X_t + \eta_t$, where X_t is the mass of investments by the agents in period t and η_t is a noise which is distributed $\mathcal{N}(0, \sigma_\eta^2)$.

If λ_t is the public LLR between states θ_1 and θ_0, an agent with private signal s has a LLR equal to

$$\lambda(s) = \lambda_t + \frac{\theta_1 - \theta_0}{\sigma_\epsilon^2}\Big(s - \frac{\theta_0 + \theta_1}{2}\Big).$$

Given the net payoffs in the two states, the agent invests if and only if he believes state θ_1 to be more likely than state θ_0, hence if his LLR is positive. This is equivalent to a private signal s such that

$$s > s^*(\lambda_t) = \frac{\theta_0 + \theta_1}{2} - \frac{\sigma_\epsilon^2}{\theta_1 - \theta_0}\lambda_t.$$

Let $F(\cdot; \sigma)$ be the *c.d.f.* of the Gaussian distribution $\mathcal{N}(0, \sigma^2)$. Since the mass of the population in period t is one, the level of aggregate endogenous investment is

$$X_t = 1 - F(s^*(\lambda_t) - \theta; \sigma_\epsilon).$$

The level of aggregate activity

$$Y_t = 1 - F(s^*(\lambda_t) - \theta; \sigma_\epsilon) + \eta_t.$$

is a noisy signal on θ. The derivative of Y_t with respect to θ is

$$\frac{\partial Y_t}{\partial \theta} = (s^*(\lambda_t) - \theta) exp\left(-\frac{(s^*(\lambda_t) - \theta)^2}{2\sigma_\epsilon^2}\right).$$

If the cut-off point $s^*(\lambda_t)$ is far to the right or to the left, the multiplier of θ on Y_t is small and the impact of θ on Y_t is dwarfed by the observation noise η_t, exactly as in the model of Vives (1993). Hence, the information content of the observation Y_t is small when most agents invest ($s^*(\lambda_t)$ is low), or most do not invest ($s^*(\lambda_t)$ is high).

Crash and boom

Suppose that the true state is θ_0 and that the level of optimism, as measured by the LLR, is high. Most agents invest and the aggregate activity is dominated by the noise.

However, the beliefs of agents are unbounded and the public belief converges to the true state. When the public belief decreases to the "middle range", the difference between the mass of agents in the two states becomes larger and this difference dominates the noise. The level of aggregate activity is more informative. Since the true state is θ_0, the public belief decreases rapidly and the aggregate activity falls drastically. A crash occurs.

This property is illustrated by a simulation. Two realizations of the observation shocks are considered. In the first, all realizations of the shocks are set at zero, $\eta_t \equiv 0$. The evolution of the public belief[15], measured in LLR, is represented in Part A of Figure 4. In Part B, the evolution of the public belief is represented for random realizations η_t.

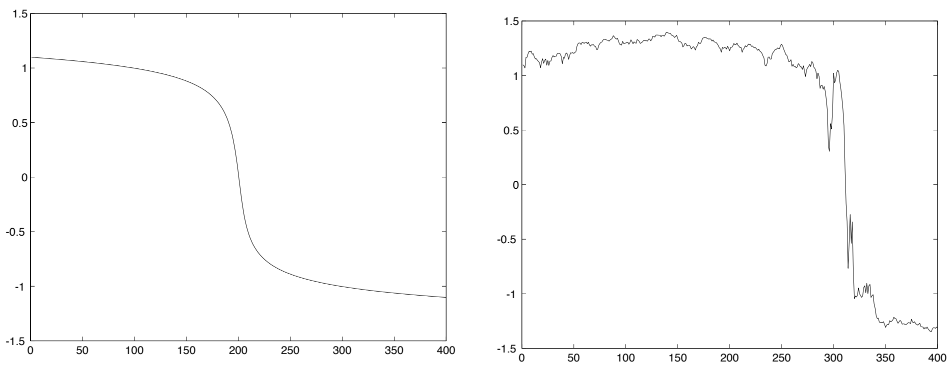

Part A　　　　Part B

Parameters: $\theta_0 = 0, \theta_1 = 1, \sigma_\eta = 0.5, \sigma_\epsilon = 2, c = 0.5,\ \theta = \theta_1,\ \mu_1 = 0.75.$
The public belief is measured in LLR on the vertical axis. (Note that the belief is not measured by the probability of state θ_1). The period is reported on the horizontal axis.

Figure 4.4: Examples of evolution of beliefs

In Part A, the public belief evolves slowly at first, then changes rapidly in a few periods and evolves slowly after. The LLR tends to $-\infty$, but this convergence is obviously very slow.

If agents could observe directly the signals of others, the curve in Part A would be replaced by a straight line. The sudden change occurs here because of the non-linearity of the

[15]The update of the public belief from λ_t to λ_{t+1} is given by Bayes' rule:

$$\lambda_{t+1} = \lambda_t + Log\left(\frac{f(x_t - \Big(1 - F(s^*(\mu_t) - \theta_1; \sigma_\epsilon)\Big); \sigma_\eta)}{f(x_t - \Big(1 - F(s^*(\mu_t) - \theta_0; \sigma_\epsilon)\Big); \sigma_\eta)}\right),$$

where f is the density function associated to the *c.d.f.* F.

information content of individual actions.

In Part B, the changes λ_t are also sudden. Note that even if the public belief is pessimistic (with $\lambda_t < 0$ at around $t = 100$), a string of positive shocks can induce a sudden upward jump of the public belief.

The model generates symmetrically crashes and booms. If the initial level of pessimism is low and the true state is high, eventually agents learn about it and the learning process goes through a phase of rapid changes of beliefs.

4.7 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)[16]. Huck and Oechssler (1998) find that the tests of Anderson and Holt are not powerful against simple rules. More recent experimental studies include Çelen and Kariv (2002b), (2002c), or Holt (2001).

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

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. Assume that $\theta = 1$. What is a wrong cascade?
3. Suppose that $x_1 = 0$. At the end of period 1 and beginning in period 2, what is the probability to have a wrong cascade?
4. Suppose that $x_1 = x_2 = 1$. What is the probability to have eventually a wrong cascade (with action 0)?
5. Let π_0 and π_1 be the probabilities to have a wrong cascade before any observation (at the beginning of period 1) and after the observation $x_1 = 1$. Remember that all this is conditional on $\theta = 1$. Using the geometric figure in the first question and your previous answers, show that after the observations $x_1 = 1$ and $x_2 = 0$, the probability of to end with a wrong cascade is equal to π_0.
6. Using your previous answers, find two linear equations between π_0 and π_1, and solve for π_0.
7. Comment on the relation between π_0 and μ.

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)

In the standard model of this chapter, assume that agent t chooses an investment level x_t which can be any real positive number. All agents have a binary private signal with precision $p > 1/2$ and a payoff function

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

1. Can an informational cascade take place with positive investment? Can there be an informational cascade with no investment?

2. Show that there is a strictly positive probability of under-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.

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 Figure **??** to represent 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 and agents B take action 0.

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

EXERCISE 4.5. (Discontinuity of the Markov process of social learning)

Take the standard model of Section 4 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 (4.3) 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?

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

The exercise shows how the BHW model can be applied to a case where the payoff of investment depends on the actions of others and on the state of nature.

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 i can make in period i a loan of size 1. The project is funded if there are N agents who make the investment. There are two states for the developing country, $\theta = 0$ or 1. Each loan pays a return R 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. If an agent does not make a loan he earns the market return r. Each agent i observes the actions of agents j with $j < i$.

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

$$\frac{1-q}{q}\frac{\mu_0}{1-\mu_0} < \frac{\mu^*}{1-\mu^*} < \frac{\mu_0}{1-\mu_0}.$$

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.

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.

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

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.

4.8 Appendix (proofs)

Proposition 4.4

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 4.9

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

$$\mu_{t+1} = \frac{\Big(1 - \int_{1-\mu_t}^1 f^1(\nu)d\nu\Big)\mu_t}{\Big(1 - \int_{1-\mu_t}^1 f^1(\nu)d\nu\Big)\mu_t + \Big(1 - \int_{1-\mu_t}^1 f^0(\nu)d\nu\Big)(1 - \mu_t)}. \tag{4.9}$$

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

$$z_{k+1} = z_k - az_k^2. \tag{4.10}$$

This sequence tends to 0 like $1/k$. Let y_k be such that $z_k = (1 + y_k)/(ak)$. By substitution in (4.10),

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

A straightforward manipulation[18] shows that $y_{k+1} < y_k$. Hence z_k tends to 0 like $1/k$ when $k \to \infty$.□

Proposition 4.10

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

$$\lambda_{t+1} = \lambda_t + \zeta_t, \quad \text{with} \quad \zeta_t = Log(\hat{\mu}_t/(1-\hat{\mu}_t)) \tag{4.11}$$

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

$$\bar{\gamma} = -\int_0^1 Log\Big(\frac{\nu}{1-\nu}\Big) f^0(\nu)d\nu > 0. \tag{4.12}$$

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, $\mu_t \leq e^{-\gamma t} e^{\nu_t}$.

□

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 4.9). 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 $1/t$.

[18]
$$1 + y_{k+1} = 1 + \frac{1}{k} - \frac{1}{k} - \frac{1}{k^2} + y_k + \frac{y_k}{k} - 2y_k\frac{k+1}{k^2} - y_k^2\frac{k+1}{k^2} < 1 + y_k.$$

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 (4.5): 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 $1/t$.

□

4.8.1 A model of learning with two types of agents

In each period, agent t receives his signal s_t in a sequence of two independent steps. First, the precision q_t of his private signal takes either the value $\bar{q}$ with probability π, or the value $\underline{q}$ with probability $1-\pi$. By convention, $\bar{q} > \underline{q}$. Second, the value of the signal is realized randomly and such that $P(s_t = 1 \mid \theta = j) = q_t$. Each agent t observes the realization (q_t, s_t) which is not observable by others. The parameters $\underline{q}, \bar{q}$ and π and the signaling process are known by all agents. In order to facilitate the discussion, the fraction π of the agents endowed with a signal of high precision is assumed to be small. In any period, the model is in one of three possible regimes which depend on the public belief λ.

i **A**. In the first regime, regime A, no agent herds. Define the values λ_A^* and λ_A^{**} such that

$$\begin{cases} \lambda_A^* = \gamma - Log\Big(\dfrac{\underline{q}}{1-\underline{q}}\Big), \\ \lambda_A^{**} = \gamma + Log\Big(\dfrac{\underline{q}}{1-\underline{q}}\Big), \end{cases} \quad \text{with} \quad \gamma = Log\Big(\frac{c}{1-c}\Big).$$

If $\lambda_A^* < \lambda_t \leq \lambda_A^{**}$, any agent with low precision invests if and only if the signal is good. An agent with high precision follows the same strategy *a fortiori*. Since no one herds, the observation of x_t is equivalent to the observation of an agent who does not herd and has a signal with precision equal to the average precision of signals in the population. The updating of the public belief is therefore[19]

$$\lambda_{t+1} = \begin{cases} \lambda_t + \alpha, \text{ if } x_t = 1, \\ \lambda_t - \alpha, \text{ if } x_t = 0, \end{cases} \quad \text{with} \quad \alpha = Log\Big(\frac{(1-\pi)\underline{q} + \pi\bar{q}}{(1-\pi)(1-\underline{q}) + \pi(1-\bar{q})}\Big).$$

[19] To find these expressions, note that $\dfrac{P(\theta=1|x=1)}{P(\theta=0|x=1)} = \dfrac{P(\theta=1)}{P(\theta=0)}\dfrac{P(x=1|\theta=1)}{P(x=1|\theta=0)}$.

i **B**. In the second regime, regime B, only the agents with a higher precision do not herd. The regime is bounded by the critical values λ_B^* and λ_B^{**} with

$$\begin{cases} \lambda_B^* = \gamma - Log\Big(\frac{\bar{q}}{1-\bar{q}}\Big), \\ \lambda_B^{**} = \gamma + Log\Big(\frac{\bar{q}}{1-\bar{q}}\Big). \end{cases}$$

Since $\underline{q} < \bar{q}$, one verifies that $\lambda_B^* < \lambda^*{}_A < \gamma < \lambda_A^{**} < \lambda_B^{**}$.

i This regime is divided in two sub-cases.

1. If $\lambda_A^{**} < \mu_t \leq \lambda_B^{**}$, the agents with lower precision herd and invest. Agents with high precision do not herd and reveal their signal only if that signal is bad. Bayes' rule takes the form

$$\lambda_{t+1} = \begin{cases} \lambda_t + \beta, \quad \text{with} \quad \beta = Log\Big(\frac{1-\pi+\pi\bar{q}}{1-\pi+\pi(1-\bar{q})}\Big) & \text{if} \quad x_t = 1, \\ \lambda_t - Log\Big(\frac{\bar{q}}{1-\bar{q}}\Big), & \text{if} \quad x_t = 0. \end{cases}$$

The LLR changes by a larger amount when the action $x = 0$ is taken.

2. If $\lambda_B^* < \lambda_t \leq \lambda_A^*$, the low precision agents do not invest and Bayes' rule is the symmetric of **1.**:

$$\lambda_{t+1} = \begin{cases} \lambda_t + Log\Big(\frac{\bar{q}}{1-\bar{q}}\Big), & \text{if } x_t = 1, \\ \lambda_t - \beta, & \text{if } x_t = 0. \end{cases}$$

One verifies that the change of the belief is much stronger when the action against the herd is taken.

i **C**. In the third regime, regime C, all agents herd. Either $\lambda \leq \lambda^*{}_B$ or $\lambda > \lambda^{**}{}_B$.

An example of evolution of the public belief is represented in Figure 4 which is a special case of Figure **??**. It illustrates some properties of the social learning with heterogeneous agents.

In the regime with moderate optimism where $\lambda_t \in (\lambda_A^{**}, \lambda_B^{**}]$, investment generates a relatively small increase of the belief. On the other hand, a zero investment generates a significant jump down of the belief. No matter how high λ_t may be, after an observation of no investment ($x_t = 0$), λ_{t+1} must be smaller than $Log(c/(1-c))$. Following such "bad news", the new regime may be of type A or B. Note also that the continuation of the learning phase does not depend on alternating private signals as in the simple model with identical agents.

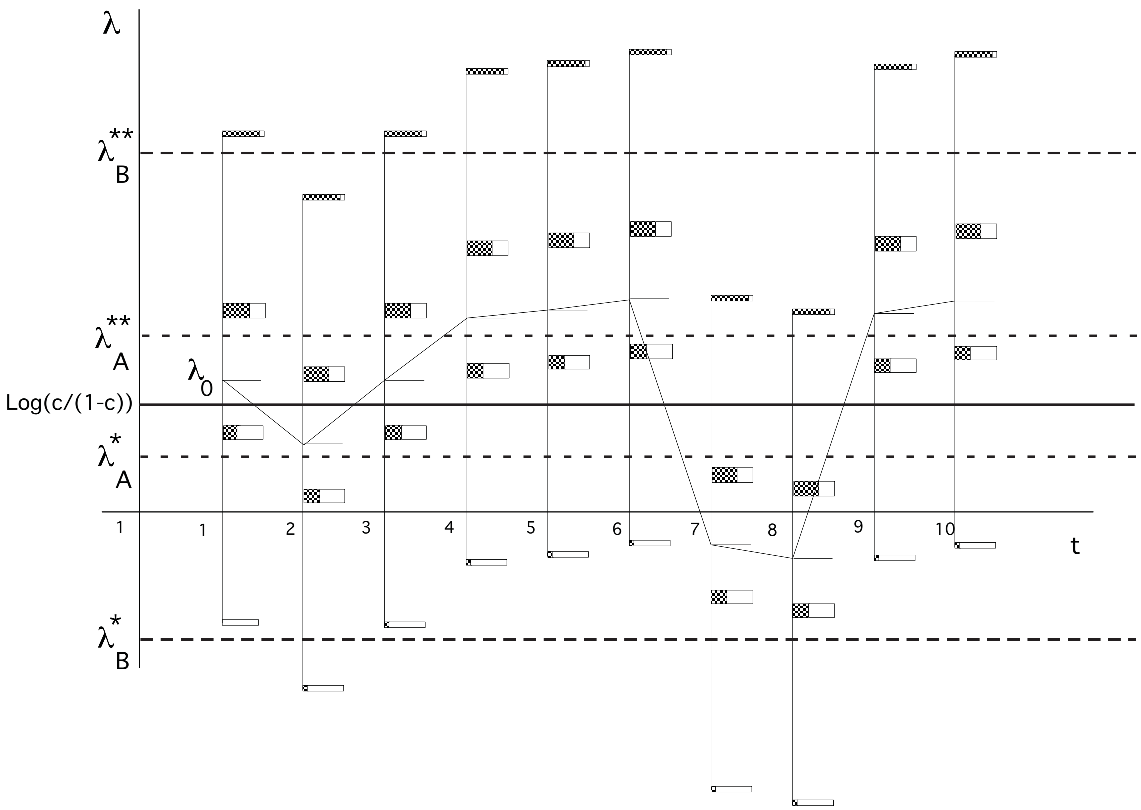

There are two precisions of the binary signals. If the public LLR is outside of the band $[\lambda_A^*, \lambda_A^{**}]$, the agent with the lower precision herds. Since there are fewer agents with higher precision, who do not herd, the evolution of the belief in the herd is slow. If the herd is broken by a black sheep, the public belief jumps.

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. In an equilibrium individual strategies are determined by the trade-off between the opportunity cost of delay and the value of the informationthat is gained from more observations. The informational externality may generate strategic substitutabilities and complementarities. Multiple equilibria may 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. The framework extends the standard cascade model of BHW.

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

[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.

productive investment may be more numerous when the state of the economy is good.

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.

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

- $z = 0$ 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 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 $\delta \rightarrow \delta^*$, then $z \rightarrow 1$. 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 $\delta \rightarrow 1$ 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 5.2).

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 5.4.

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

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. 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 payoff, given his information and the equilibrium strategy of all agents for any future date

[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).

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. □

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 $(c, \bar{\mu})$. □

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 5.1. 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

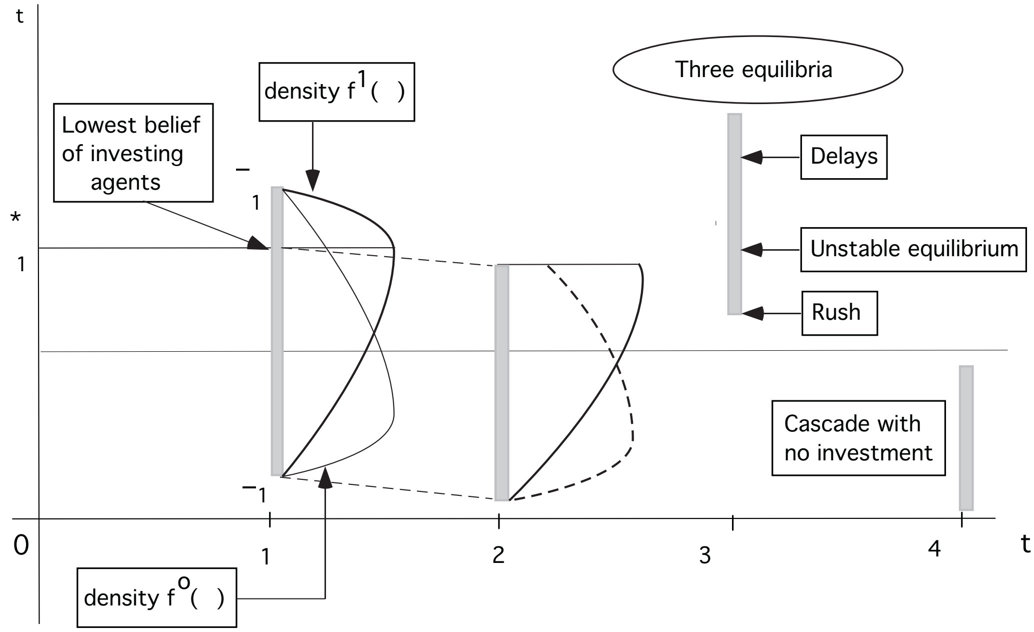

The number of investments in a period t depends on the number of agents with a belief higher than λ_t^*. At the end of a period, the updated distributions in the two states are truncated, translated and rescaled. Period 3 (in which the representation of the densities is omitted) corresponds to a case 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 5.1 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 5.1, 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 5.1. 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

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

(LLR) is bounded below. However, if beliefs are unbounded, the structure of equilibria is very similar to that in Figure 5.1. 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. (Vanishingly small period)

Consider the model of Section 5. Determine the limit of the belief (probability of he good state) after the bad news of no investment, μ^-, when $\delta \to 1$, without computing this value. Explain the result. Do agents learn something, asymptotically. Show that you can also compute μ^-.

EXERCISE 5.2. (Observation noise)

Consider the simple model of delay in this chapter with two agents and two possible states. We now introduce an observation noise. Assume that if a person invests, she 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. (An investment subsidy)

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 $\theta = 1$. 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. (delay with a large number of agents)

Consider the simple model of this chapter 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 $N \to \infty$.

Chapter 6

Words

If we all think alike, it means we do not think anymore.

Trust but verify.

Communication with words is the subject of a vast literature. In previous chapters, people sent messages to others through their actions. In this chapter, the messages are words. One critical issue will be whether "talking" is credible. We will see that herding on actions and herding on words occur under similar conditions. For example, herding may arise in financial markets because of the observation of others' actions or because of the behavior of financial advisors who are influenced by others' predictions.

In the generic setting, an agent is an expert with private information on the state of nature and his action takes the form of a message that is sent to a receiver. How can he transmit credibly his information by mere words? The key is that the receiver has some independent information on the true state, an information that he gets after, or even before, the advice from the expert. The receiver thus can *verify* the expert's message against his independent information. (The precision of the independent information of the receiver does not matter). The payoff of the expert depends on his message (his advice) and on the independent information of the receiver.

The vocabulary for advices is often limited. In financial advising, a well respected agency, Value Line, uses a dictionary with only five words. The restriction of a small number of words parallels that of the discrete actions in the basic model of informational cascades and leads to similar results. The example of financial advising justifies the basic model of advising that is presented here in the first section: there are two states of nature, the expert has binary information and speaks with two words. The state is revealed to the receiver after the advice of the expert (as an investor who experiences the fluctuations of the stock market after receiving an advice). Three types of payoff functions are considered.

(i) The payoff of the expert is a function that depends on his message and the state as verified by the receiver. The goal of the expert is to conform as much as possible to the verified state. His belief is formed from the public belief and his private signal. If the public belief in one of the two states is high, the probability of that state is high even with a private signal favoring the other state. In that case, the expert predicts the same state as the public belief: he herds on the public belief and his message is ignored by the receiver. The expert tells the truth (sends a prediction according to his private signal) only if the public belief is not too strong on one of the two states. The condition for truth telling by the expert turns out to be identical to the condition for no herding in the BHW model of Chapter 2. In this first case, the payoff of the expert is set arbitrarily by the receiver.

(ii) The payoff based on reputation to be a good (versus a bad) expert. Reputation may be valuable because of future business for the expert, or his capability to have influence in the future. This setting may be more relevant but it puts restrictions on the payoff function and on the type of equilibrium. Two types of reputation will be considered.

In this first case, a good expert has private information of a higher precision than a bad expert. There are two types of private signals for the expert, one more informed than the other. First, assume that the expert does not know the quality of his signal. A key difference with (i) is that the value of reputation, and therefore the payoff of the expert, depends on an equilibrium. If the expert sends an irrelevant message, he *babbles*, then the receiver may ignore his message. But if the receiver ignores his message, the expert has no incentive to tell the truth. There is always a *babbling equilibrium*. We focus on the condition for the existence of a *truthtelling equilibrium*. It is similar to the condition in case (i): the public belief, as expressed by the probability of one of the two states, must be neither too high nor too low. Second, if expert knows the quality of his private information, the analysis is similar but the expert with low precision herds for a wider set of public belief than the highly informed expert.

(iii) In the second case, a good expert does not manipulate the receiver. All experts have private information of the same precision, but some experts would like the agent to take a specific action. As an example, some people would like to systematically increase or decrease welfare programs. An unbiased agent may be in a position to support a particular program, but he does not want to be identified with these people. In order to enhance his reputation, he may act as a contrarian and advise against the program that is recommended by the biased expert even if he thinks that this is the better program.

In all the models considered so far, the expert does not know the independent information used by the receiver for the verification and the reward of the expert. This assumption is can be relaxed. For example, the expert may know the financial literature read by the receiver, or the consultant may know the prejudice of the boss. It is essential that the eceiver does not know what the expert knows about him, or how the expert uses that information. If the receiver knows what the expert knows about him, he can simply "factor out" the expert's information on him from the advice and still get to know the information of the expert.

The analysis of one expert provides the ground for considering the sequence of advices, say in a committee, where each expert before speaking, hears the advices of the previous speakers. How does the opinion issued by the first have an impact on the saying of the second, and so on? People influence each other in jury trials (*i.e.*, *Twelve Angry Men*), financial advising or economic forecasting.

This setting reproduces the basic model of BHW where any acting agent observes the previous actions. The public belief evolves after each expert's message and there is a herd by all remaining experts if the public belief favors sufficiently strongly one of the two states. The model is isomorphic in assumptions and properties to the BHW model.

In the Talmud, the older speaks after the young. Presumably the older is wiser and his advice could intimidate the young to assent instead of conveying truthfully the information. In a setting with two agents where the older's information is more precise than that of the younger, this proposition is shown to be false here (Section 6.2.1). The older should speak first. When we compare the two sequences where the older speaks first or second, the first sequence is never inferior to the second and is strictly better when the older is a contrarian and speaks against a strong prior consensus: the young would herd on the consensus and his advice is ignored, but if the older speaks shakes the consensus, then the young advice will have some information value and he will be listened to.

6.1 Advice by one expert

Suppose you are asking an expert for advice about the future direction of the market. Say there are two possible future events, up or down; in mathematical language, the state of the world, to be realized later is $\theta \in \{0, 1\}$. Suppose that there is some general opinion about the future, from the press and other news, which is quantified by the prior probability of the event "up" ($\theta = 1$), to be μ, between 0 and 1. Your expert has some additional information (that is why he is called "expert"). The quality of this information is that the expert's prediction is correct with a probability q $(1/2 < q < 1)$. This information is equivalent to a signal s that is equal to the state with probability q: $P(s = \theta|\theta) = q$).[1] What you would like to know is the information of the expert, his signal value. e incentive to reveal it truthfully depends on his reward.

In all this section, we assume that the receiver of the advice compares the advice with the observation of the state of the world, which occurs later. The expert knows the evaluation process and has some information on the probabilities of the states in the future. If his rewarded for being correct, his objective is to match his message with the most probable state in the future, and truth telling may not be the best strategy. If the "general consensus" is that the market will go up, the expert may follow the consensus even if his private information points the other way.

6.1.1 Evaluation payoff after verification

Here, we analyze on the incentive problem of the expert and, as usual, we consider the simplest model that focuses on this issue. There is only one type of expert and his reward, noted $v_{m,\theta}$ depends on his message m and the true state θ that is revealed after the advice. Later, we will consider heterogenous experts and the rewards will be endogenous to the evaluation of the quality of the expert.

The expert sends to a receiver a message m which is a (possibly random) function of his signal, $m(s)$. The expert cannot communicate more than his information which is in the set of values $\{0, 1\}$. Without loss of generality, the message takes values in the set $\{0, 1\}$. The truth telling strategy is defined[2] by $m(s) = s$. The expert, who has a signal s, maximizes his expected payoff computes his payoff

$$V(s, m) = P(\theta = 1|s, \mu)v_{m1} + P(\theta = \theta_0|s, \mu)v_{m0}, \tag{6.1}$$

where his belief $P(\theta = 1|s, \mu)$ depends on both his private signal and the (prior) public belief μ according to Bayes' rule. We make the common sense assumption that the reward is higher when correct:

$$v_{ii} > v_{ij} \qquad \text{if} \quad i \neq j. \tag{6.2}$$

The truthtelling strategy is optimal if it yields to the expert a payoff which is not strictly smaller than that obtained from deviating. For each signal value of the expert, there is an incentive compatibility constraint to tell the truth ($m_i = s_i$):

$$V(1, 1) \geq V(1, 0), \quad \text{and} \quad V(0, 0) \geq V(0, 1). \tag{6.3}$$

[1] One could of course consider a non symmetric signal.

[2] My son Sebastian has frequently reminded me that $m(s) = 1 - s$ is also a truth telling strategy.

Using the expression of $V(s,m)$ in (6.1), these constraints are equivalent to

$$\begin{aligned} P(\theta=1|s=0,\mu) \leq c \leq P(\theta=1|s=1,\mu), \quad \text{with} \\ c = \frac{v_{00}-v_{10}}{v_{11}+v_{00}-v_{01}-v_{10}}. \end{aligned} \tag{6.4}$$

Since the probabilities $P(\theta|s,\mu)$ are the beliefs of the expert, the incentive compatibility constraints are *the same* as the condition for no herding in the BHW model where agents have a cost of investment c. If the public belief μ is higher than some value μ^{**}, an expert with signal s_0 has a belief (based on his signal and μ) that is higher than c and he sends the message s_1. He is herding. Likewise when the public belief is below some threshold μ^*.

Without loss of generality, assume that there are only two reward values for being right and wrong, respectively:

$$v_{00} = v_{11} > v_{10} = v_{01}. \tag{6.5}$$

The value of c is now 1/2. The situation of the expert is the same as that of an agent in the BHW model with an investment cost of 1/2. He will say "1″ ("invest" in the BHW model) if and only if he thinks that given all his information (public and private), the state 1 is more likely. If the public belief is sufficiently high (low), he will say "1″ ("0″) regardless of his private information. We know from the analysis of the BHW model that the expert will tell the truth if and only if the public belief is in some intermediate range that is defined by

$$1-q < \mu < q.$$

The range of values of the public belief with truthtelling by the expert obviously increases with the precision of his information. A "poorly informed" expert herds more easily.

If the receiver can choose the reward function, he may always get the private information of the expert by choosing $v_{m\theta}$ such that the value of c in (6.4) falls between the beliefs of an optimistic and a pessimistic expert (with signal 1 and 0). The receiver may not be able to write a contract that specifies the values of the rewards. We now turn to rewards based on reputation.

6.1.2 Equilibrium with an evaluation based on reputation

There are some good and some bad experts, that is, with high and low precisions in their private information. To improve one's reputation may be a powerful incentive to send a message which gives the best possible prediction. To analyze the issue, let us build on the previous model. The symmetric binary signal of an expert is correct with high or low probability, $q_H > q_L$. The prior probability of a good expert is α. As usual in all Bayesian models, the receiver of the advice knows the structure and the parameters of the model but does not know the private information, *i.e.,* the value of the signal of the expert. After the observation of the state of nature, the Bayesian receiver updates the probability that the expert is good from α to a new value $v_{m,\theta}$. The rewards is now endogenous to the behavior of the expert. Do we have to justify why an expert would value his reputation?[3]

The evaluation by the receiver depends on the strategy of the expert, and the strategy of the expert depends on the evaluation function which can be defined as the strategy of the receiver. Both strategies have to be determined simultaneously in a game. The situation is thus different from the previous case with an exogenous payoff $v_{m\theta}$.

The babbling equilibrium

The endogenous property of the reward function is highlighted by the existence of the babbling equilibrium. If the agent sends a message which is independent of his signal, he cannot be evaluated. His message is ignored by the receiver and his reputation stays constant at α. But if the receiver does not listen, the expert has no incentive to speak the truth. No strategy can strictly improve his reputation. He can claim to be good as much as he wants. The receiver has no way to discriminate him from other experts who babble. Therefore, in a setting where reward is based on the reputation to have more accurate information, for any value of the public belief, there is a babbling equilibrium where an expert is not listened to and has no incentive to speak the truth.

`When the expert's payoff is based on reputation against some other agents in the game, there is always a babbling equilibrium.`

The truthtelling equilibrium

Let us first make the assumption that the expert has no better information than the receiver about his own type. That may be strange, but it turns out that this assumption is not restrictive. We will remove it later. We have seen (Exercise ***) that such an agent treats his signal as having the precision (probability to be correct)

$$q = \alpha q_H + (1-\alpha) q_L. \tag{6.6}$$

Let $\mathcal{H}$ and $\mathcal{L}$ be the events that his signal has high or low precision. Suppose that the expert tells the truth. $m(s) = s$). By Bayes' rule, the *ex post* reputation is

$$v_{s\theta} = P(\mathcal{H}|s,\theta) = \frac{P(s|\mathcal{H},\theta)\alpha}{P(s|\mathcal{H},\theta)\alpha + P(s|\mathcal{L},\theta)(1-\alpha)}. \tag{6.7}$$

The terms $P(s|\mathcal{H},\theta)$ and $P(s|\mathcal{L},\theta)$ are the probabilities of the realization of the expert's signal given the type of the signal and the state of nature. They depend only on the

[3]For an example where the expert would like the receiver to make the best decision according to the expert, see Exercise 6.2 that is based on Morris (2001).

structure of the agent's signals. Since the signal is symmetric,

$$v_{11} = v_{00} > v_{10} = v_{01}. \tag{6.8}$$

The truth telling condition is *the same* as in the case where the payoffs are fixed. Proposition 6.1 summarizes the previous discussion and introduces an additional result.

PROPOSITION 6.1. *For any value of the public belief $\mu = P(\theta = 1)$, there is a babbling equilibrium where the expert conveys no information.*

Assuming equal prior probabilities for the two states, if $1 - q < \mu < q = \alpha q_H + (1 - \alpha) q_L$, there is a truth telling equilibrium.

When the prior belief is strong, either high or low and μ is outside of the middle range $[\mu^*, \mu^{**}]$, babbling is the only equilibrium. Is this bad for the receiver? Not necessarily: if he would know the signal of the expert and choose the most likely state, as a rational Bayesian, he would ignore that signal.

The type of the expert is known

Consider first the case where the expert almost knows his type. The value of α is vanishingly close to one. For simplicity assume that the signal of low precision is not informative at all: $q_L = 1/2$. From the revious section, the agent tells the truth if the public belief is in the interval $(1 - q, q)$ with $q = q_H \alpha + 0.5(1 - \alpha)$. When α is vanishingly small, asymptotically, the expert gives his best possible advice given his precision q_H.

PROPOSITION 6.2. *If the type of the expert is known with a probability vanishingly close to one, there is a truthtelling equilibrium in which the expert speaks against the public belief if and only if he believes his advice is more likely to be true.*

The key assumption for a truth telling equilibrium is that there must be some experts of lower quality. In the proposition, a vanishingly small probability of a bad expert is sufficient when all experts have the same prior information about their own quality. But the same mechanism for truth telling is a work when experts are of two types and *know* their own type. In order to improve their reputation, all experts, good and bad, will try to predict the most likely state on the basis of their information. If there are two types with precision $q_H > q_L > 1/2$, respectively, the low (high) quality expert will send a message equal to his signal if the prior public belief is in the range $(1 - q_L, q_H)$, $((1 - q_L, q_H))$. For example, when the public belief is between the low and the high precision, the low quality expert herds on the public belief and the high quality expect tells the truth. But note that the low quality expert herds on the public belief because the high quality expert tells the truth and does not babble: he sends a message to conform as much as possible to the behavior of the high quality expert.

What have we learned so far?

TO BE DISCUSSED IN CLASS

In all the cases we have seen, the reputation updating rewards experts *not* for giving good advice, but for giving advice that is different from the advice of the low quality experts. In the previous cases, the low quality experts predicted less well the state of nature. Therefore, the incentive was to try to predict as well as possible. But the low quality experts may have some characteristic that is different from a low predicting accuracy. In this situation, there is no reason that the good expert should attempt to predict accurately. We now consider such a case.

A contrarian facing biased experts

There are two types of experts, the good and the bad. The good are the same as in the previous model. The bad expert always send the message 1. One may say he is biased toward action 1. So what should the good expert do? If he cares only about his reputation, then whatever his signal, he should send the message 0! This limit case may illustrate how an expert may want to give bad advice just to be seen as a "contrarian" that does not follow some biased purpose. These few lines may be a little short for a publication in a top 5 journal, but the construction of a more elaborate model does not really bring additional insight about the mechanism for contrarian advice.

In the previous paragraph, the good expert achieves is perfectly identified when he sends the message 0. To reduce the increase of reputation, assume that the bad expert always says 1 when his signal is 1 and lies with probability ν when his signal is 0. Using the same Bayesian computation as before, the evaluation function is now defined by

$$v_{1,1} = \frac{\lambda q}{\lambda q + (1-\lambda)(q + (1-q)\nu)}, \qquad v_{1,0} = \frac{\lambda(1-q)}{\lambda(1-q) + (1-\lambda)(1-q(1-\nu))}, \tag{6.10}$$

$$v_{0,0} = v_{0,1} = \frac{\lambda}{\lambda + (1-\lambda)(1-\nu)}.$$

Note that

$$v_{00} = v_{0,1} > \lambda > v_{1,1} > v_{1,0}$$

where the inequality are strict if and only the bad expert lies with some probability $\nu > 0$. The inequality between $v_{1,1}$ and $v_{1,0}$ appears because the probability of a lie in the message

1 is higher when the true state is 0 compared to the state 0. When the message is 0, the probability of lying is the same in both states.

In the introductory paragraph, the values of $v_{0,0}$ and $v_{0,1}$ where equal to 1 because the bad expert with a signal 0 was always lying ($\nu = 1$). Now these values are smaller than 1 but we obviously keep the inequality $v_{0,1} > \lambda > v_{1,1}$ and the incentive effect for the good expert to lie when his signal is 1 in order to increase his reputation of being unbiased.

Assume that any good expert with a signal 1 lies and sends the message 0 with probability ζ. We can recompute the expressions of the reputation (6.10) which now depend on the probability ζ that the good expert lies by sending the message 0 while having the signal 1. One can show that for $\nu > 0$, [4]

$$v_{00}(\zeta) > v_{01}(\zeta) > \lambda > v_{11}(\zeta) \geq v_{10}(\zeta). \tag{6.12}$$

Consider a (good) expert with a signal 1. After he sends his message, the state will be revealed to be equal to 1, with probability q, and with probability $1 - q$, equal to 0.

- Should the state be 0, it is always better to have lied all the time with a message 0 that turns out to be equal with the state.
- Should the state to be 1, we see from the inequality in (6.12) that his reputation is higher if he gives the wrong message 0 instead of 1.

If other good experts with signal 1 have a strategy to lie with probability ζ, then a good expert with such signal should deviate and lie all the time.

PROPOSITION 6.3. *Assume that there is some probability that the expert is bad in which case he gives with probability $\nu > 0$ the advice $m = 1$ while having the signal $s = 0$,*

[4]
$$v_{00} = \frac{\lambda(q + (1-q)\zeta)}{\lambda(q + (1-q)\zeta) + (1-\lambda)q(1-\nu)}$$
$$v_{11} = \frac{\lambda q(1-\zeta)}{\lambda q(1-\zeta) + (1-\lambda)(q + (1-q)\nu)}$$
$$v_{10} = \frac{\lambda(1-q)(1-\zeta)}{\lambda(1-q)(1-\zeta) + (1-\alpha)(1-q+q\nu)}$$
$$v_{01} = \frac{\lambda((1-q) + q\zeta)}{\lambda((1-q) + q\zeta) + (1-\lambda)(1-q)(1-\nu)}$$

Note that

$$\begin{cases} v_{00}(\zeta) \text{ and } v_{01}(\zeta) \text{ are increasing in } \zeta \,, \\ v_{10}(\zeta) \text{ and } v_{11}(\zeta) \text{ are decreasing in } \zeta \,. \end{cases} \tag{6.11}$$

then a good expert with a signal 1 *who cares only about this reputation to be good should lie all the time and give the message* 0*.*

So what could prevent an expert from lie all the time when his signal is 1? The cost of sending the message 0 is on the receiver and if the good expert cares about the receiver of that particular message in the same way as the receiver does, he bears the same cost.

Let $C(\zeta)$ be the cost for the expert of lying by sending the message 0 while having the signal 1. That cost function could take any shape, depending on the context. It could be a function of the distance between the action of the receiver and the true state, a function that would be perfectly known by the expert (as assumed by Morris, 2001). But the algebraic formulation is here only an exercise in algebra and does not provide additional insight. Let us just assume that cost function of lying $C(\zeta)$ for the expert. And by the way, this cost function could include the psychological cost of lying for the expert.

Likewise, the valuation by the expert of his reputation may depend on a number of factors which could include sheer pride of oneself. Let us denote this valuation by $A(U(\zeta))$, where $U(\zeta)$ is the expected evaluation of the type of the expert when he lies with probability ζ while having the signal 1.[5] The objective function of an expert with signal 1 is

$$V(\zeta) = U(\zeta) - C(\zeta).$$

Depending on the shape of the functions U and C, anything is possible, with corner solutions at 0 or 1, or on the interval $(0, 1)$ or even multiple solutions. But these are trivialities and ground for algebraic exercises. They are irrelevant for the main message that is presented in Proposition 6.3.

6.2 Panel of experts

When the advice is given by a panel of experts (a committee, a jury in a trial), members of the panel hear the advices given by other members and influence each other. Financial or medical advisors, economic forecasters, discussants of papers, are aware of the predictions of others and do take them into account. We first analyze a simple model in which each expert "speaks" once in a pre-established order. We will then compare the quality of the panel's advice for different sequences in which members speak.

5

$$U(\zeta) = q(v_{1,1}(1-\zeta) + v_{0,1}\zeta) + (1-q)(v_{1,0}(1-\zeta) + v_{0,0}\zeta) \tag{6.13}$$

6.2.1 A sequence of experts with a pre-established order

The model is the same as in Section 6.1. We add a sequence of experts with independent types and signals on the state $\theta \in \{0, 1\}$. Each expert cares for his reputation as described as in the previous section and has a symmetric binary signal of precision ρ_H with probability α and of precision ρ_L otherwise, $\rho_H > \rho_L$. The precision is not observable directly. The value of α is vanishingly close to 1. We could also assume that the reward for advice is an exogenous symmetric reward function.

Each expert speaks once and knows the messages of the experts who have spoken before him. Once all the experts have spoken, the receiver learns the true state and updates his estimate of the precision of each expert. Since the evaluation of each expert depends only on his message and the true state, each expert has no incentive to manipulate the messages of other experts. Each expert in the panel is exactly in the same situation as the unique expert in Section 6.1. An expert who speaks in round t formulates his message according to the public belief μ_t, (which depends on the history of messages $h_t = \{m_1, \ldots, m_{t-1}\}$), and his own signal s_t. Recall that in any round, babbling is an equilibrium. We will assume that whenever there is another equilibrium with no babbling (herding), both the expert and the receiver (through the evaluation function) coordinate on this equilibrium. Following the analysis in the previous section, an expert herds if and only if the public belief is outside the band $(1-\rho, \rho)$ where $\rho = \alpha\rho_H + (1-\alpha)\rho_L$ is the average precision. We assume of course that the public belief in the first period, μ_1, is in the interval $(1-\rho, \rho)$.

Given the condition $1 - \rho < \mu_1 < \rho$, the first expert reveals his signal. Because of the equivalence with the BHW model with a cost of investment c equal to $1/2$, the analysis of Chapter 4 applies. Suppose that $\mu_1 > 1/2$ (state θ_1 is more likely), and that the signal of the first expert is bad: $s(1) = 0$. He tells the truth and sends the message $m(1) = s(1) = 0$. His information is incorporated in the public belief μ_2. When two consecutive experts in the sequence have the same signal, the truthtelling condition is not met. At that point, the babbling equilibrium is the only equilibrium. Since nothing is learned, the truthtelling condition is not met in the following period, and so on. The babbling equilibrium is the only equilibrium for all subsequent periods. Learning from experts stops. One might as well assume that all experts give the same advice. The expression “herding” is appropriate here. Given the equivalence between herding and babbling, the model is isomorphic to the BHW model. The probability that a herd has not occurred by round T converges to zero at an exponential rate. Note that the behavior of the agents does not depend on the probability α of a signal with high precision.

Scharfstein and Stein (1990), in the first analysis of herding by experts, assume that the signals of experts are correlated in the following sense: if the signals of both experts are informative, they are identical. Scharfstein and Stein seem to support the following story: the first expert has no incentive to lie and he tells the truth. The second expert who learns the signal of the first expert could say: if I have a signal of high precision, it is more likely that my signal is the same as that of the first expert because signals of high precision are identical. As emphasized by Ottaviani and Sørensen (2000), such an argument is irrelevant and confuses the issue. This case is left as an exercise for the reader. The condition for babbling is modified when the experts' signal are correlated. This modification is the same as in the BHW model where agents' actions are observed.

6.2.2 Who should speak first: the strongly or weakly informed?

In a deliberating group, the order in which people voice their opinion may be critical for the outcome. The less experienced expert often speaks first while the old and wise[8] waits and speaks last. Presumably, this rule of *anti-seniority* (to use an expression of Ottaviani and Sørensen, 1999b)[9], enables the less experienced to express their opinion free of the influence by the more experienced. Can the anti-seniority rule be validated by the analysis of this chapter? The answer will be negative.

Assume N experts, indexed by $i \in \{1, \dots, N\}$, each with a SBS of precision almost equal to q_i. (Each private signal is uninformative with arbitrarily small probability). By convention, q_i is strictly increasing in i. (Expert N is the most informed, or the senior). The values of q_i are publicly known and the receiver, before receiving any message, can choose the order in which experts speak. Each expert knows which experts have spoken before him and their messages.

The goal of the receiver is to choose the state which is most likely once he has listened to each expert. This objective is equivalent to the maximization of the payoff $E[\theta]x - c$ with $c = 1/2$ where the action x is taken in the set $\{0, 1\}$. Once all the experts on the panel have spoken, the state is revealed and each expert is evaluated by comparing his message with the true state, as shown in the first section. In round t, expert t "speaks": he sends a message which maximizes his expected evaluation as in the model of Section 6.1. His message depends on the evaluation function and his belief which depends, in a Bayesian fashion, on the public belief in round t, μ_t, and on his private message s_t. We begin with the case of two experts.

The two-expert panel ($N = 2$)

[8] This expression is used as a convenient picture for the analysis.

[9] The presentation in this section is complementary to that of Ottaviani and Sørensen (1999b).

The two experts are called Junior (with a signal of precision q_1) and Senior (with a signal of precision $q_2 > q_1$). The *ex ante* public belief as expressed by the LLR between the good and the bad states is denoted by λ. Let $\gamma_i = Log(q_i/(1-q_i))$.

Without loss of generality, it is assumed that $\lambda \geq 0$ and that $\gamma_2 - \gamma_1 > \gamma_1$. (The case $\gamma_2 - \gamma_1 < \gamma_1$ is similar and it is left as an exercise). There are four possible cases which depend on the value of λ, as represented the following figure.

Possible cases with a panel of two experts

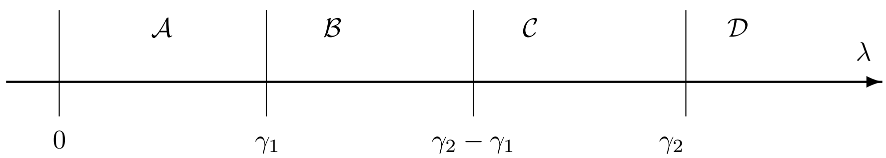

1. Suppose first that λ is in the interval $\mathcal{A}$: $0 \leq \lambda < \gamma_1$. If Junior speaks first, his signal is stronger than the public belief ($\gamma_1 > \lambda$) and he speaks the truth[10]. But since $\lambda + \gamma_1 < \gamma_2$, the public belief once he has spoken is smaller ("weaker") than the strength of the signal of Senior. For any signal of Senior, Junior is overruled and has no impact on the decision of the receiver. If Junior speaks after Senior, the only equilibrium is the babbling equilibrium. Whatever his message, he is not listened to.

2. Suppose that λ is in the interval $\mathcal{B}$: $\gamma_1 < \lambda < \gamma_2$[11]. If Junior speaks first, he babbles. (His signal is weaker than the public belief). If Junior speaks second, he also babbles (as can be verified). Junior is irrelevant. In region $\mathcal{B}$, the receiver never gets to observe the signal of Junior whatever the rule.

3. Suppose that λ is in the interval $\mathcal{C}$. If Junior speaks first, he babbles as in region $\mathcal{B}$. Suppose that Senior (who does not babble) speaks first a message $s_2 = 0$. The public belief LLR for Junior is $\lambda - \gamma_2 < 0$. Since $\lambda - \gamma_2 - \gamma_1 < 0 < \lambda - \gamma_2 + \gamma_1$, Junior reveals his signal. Junior has an impact on the decision of the receiver. The anti-seniority rule strictly dominates the seniority rule.

4. In region $\mathcal{D}$, all experts babble whatever the order in which they speak and the panel

[10] Recall that if there is a truthtelling equilibrium, this equilibrium is chosen by the expert and the receiver.

[11] The case of $\lambda = \gamma_1$ can be ignored because its *ex ante* probability is zero.

can be ignored.

PROPOSITION 6.4. (Dominance of the seniority rule) *Assume that a receiver chooses $x \in \{0,1\}$ to maximize the payoff function $E[\theta]x - 1/2$, with $\theta \in \{0,1\}$, and gets advice from a "junior" and a "senior" expert who have private signals with precision q_1 and $q_1 < q_2$, respectively. For any prior μ on state $\theta = 1$, the seniority rule (where the senior agent with higher precision speaks first) dominates the anti-seniority rule. For some values of $\mu \in (\beta_1, \beta_2)$ where $1/2 < \beta_1 < \beta_2 < 1$, the payoff with the seniority rule is strictly higher than that with the anti-seniority rule. For other values of μ, both rules generate the same outcome.*

6.3 Bibliographical notes

In Section 6.1, the case where experts know their precision corresponds to the model of Trueman (1994). This model is presented in Exercise **??**. Proposition 6.2 applies.

In Section **??**, the fundamental paper on manipulative experts is by Crawford and Sobel (1982). They assume that θ is on an interval of real numbers and that the expert has a systematic bias towards a higher (or lower) level of action by the receiver. They show that the message of the expert takes discrete values: the expert lies, but not too much. Very nice papers about the transmission of information, which unfortunately cannot be discussed here, have been written by Benabou and Laroque (1992) and Brandenburger and Polak (1996). Zwiebel (1995) analyzes how agents choose similar actions in order to be able to be evaluated by a manager.

In relation to Section 6.2, Ottaviani and Sørensen (1999b) analyze an extension of the model in which the sets of values for θ, s and m is the interval $[-1,1]$. An expert is endowed with a type t and a signal s with a density $f(s,t,\theta) = (1 + st\theta)/2$. (A higher type t means a higher precision of the signal s). They show that there is no truthtelling equilibrium. Glazer and Rubinstein (1996) propose a mechanism to prevent herding between referees.

Welch (2000) develops in a remarkable study an econometric methodology to estimate imitation when choices are discrete[12]. He analyzes how the probabilities of analysts' revisions of the recommendations (which take place in a set of 5 values from "strong buy" to "strong sell"), depend on the established consensus. His results indicate that some herding takes

[12]The estimation software is downloadable from his web site. The data comes from Zacks' Historical Recommendations Database (which is used by the *Wall Street Journal* to review the major brokerage houses).

place, especially in a bull market. A next step in this research could be the construction of a structural model with both an exogenous process of information diffusion and learning from others, and the analysis of its empirical properties. (See also Grinblatt, Titman and Wermers, 1995).

EXERCISES

EXERCISE 6.1. Imperfect verification of the expert's message

Consider the model of Section 6.1 where both states 1 and 0 have equal priors and the expert has a SBS with precision ρ. The receiver does not observe the state *ex post* but has a private symmetric binary signal y with precision $q \in (1/2, 1]$. The timing of that signal is not important if its value is not observed by the expert.

1. Using the notation of Section 6.1 for the reward function v_{my}, determine the payoff function $V(s, m)$.

2. Establish the condition for truthtelling by the expert.

3. Show that if the reward function is such that $v_{00} = v_{11}$ and $v_{10} = v_{01}$, the condition for truthtelling is independent of $q \in (1/2, 1]$. Provide an intuitive interpretation.

EXERCISE 6.2. The value of reputation

Following Morris (2001), assume that an expert gives an advice in a second period (with a new signal of the same precision) to a receiver who has a payoff function $-E[(x - \theta)^2]$ and that the expert's payoff is the same as that of the receiver. Both states are equally likely.

1. Determine the action taken by the receiver in the next period as a function of the *ex post* reputation of the expert, β.

2. Determine the value of β for the expert.

EXERCISE 6.3. Computation of the reputation function

In the model of Section 6.1, assume that with probability α, the agent has a binary signal of precision $\rho > 1/2$, and with probability $1 - \alpha$ a binary signal of precision $1/2$ (which is not informative). Determine the algebraic expression of $v_{s\theta}$ in (6.7). Show (6.8).

EXERCISE 6.4. (Partial truth telling)

Assume that $1/2 < \mu < \rho$, $(1 - \rho < \mu < 1/2)$. Show that there is an equilibrium in which the agent tells the truth if he has a good (bad) signal and lies with some probability if he has a bad (good) signal, and that this equilibrium is unstable in the same sense of stability as in Proposition 6.1.

EXERCISE 6.5. A continuum of beliefs

Assume that the private belief of the agent takes a value in the bounded interval $[\underline{\mu}, \bar{\mu}]$, $0 < \underline{\mu} < \bar{\mu} < 1$. Set $v(m, 0) = 1 - m$ and replace $v(m, 1)$ by $v(m)$.

1. Determine a necessary condition on the derivative $v'(m)$ such that the expert reveals his belief μ (and sends the message μ for any $\mu \in [\underline{\mu}, \bar{\mu}]$).

2. Determine the family of admissible functions.

3. Is the condition in question 1 sufficient?

EXERCISE 6.6. The value of reputation (Section **??**)

Let α be the reputation of the expert (probability of being of the good type). Suppose there is only one period and the expert does not care about his reputation at the end of the period; he gives an advice such that the receiver takes an action which maximizes the expert's payoff.

1. Determine the action of the receiver as a function of the message m and the reputation α.

2. Compute the *ex ante* expected payoff of the good expert, $V_G(\alpha)$, and of the bad expert, $V_B(\alpha)$, at the beginning of the period, before he gets his private information. Show that both functions are strictly increasing in α.

EXERCISE 6.7. (The white van)

This exercise white van problem which has been raised recently in the suburbs of Washington D.C.. The issue is related to Chapter 12 in the notes, but with some shortcuts, it can be addressed now. First, some very brief introductory remarks:

In most of the problems we studied, agents receive exogenous private information (or signals). The revelation of this information is endogenous—that is the essence of all the problems we had—but the private information is exogenous. It seems that in the case of white van, endogenous signals are important: people look for news; if they look for a white van, they don't look for a dark sedan (which turned out to be the true state). There is much more work there... At this stage, one should simplify as much as possible. Hence, the following model.

There are two spots, 1 and 2, and the state θ is in one of the two spots: $\theta \in \{1, 2\}$. Say, the criminal is in a white van if $\theta = 1$, and otherwise $\theta = 2$. We assume by convention that the true state is 1. (Agents do not know that however). The state is randomly fixed before the first period.

In each period, nature issues two signal $s_j \in \{0, 1\}$, one in each spot j, which is defined as follows: (i) With probability β, there is no information: with probability α one signal is equal to 1, and that signal is in spot 1 or 2, with the same probability; with probability $1 - \alpha$ both signals are 0; (ii) with probability $1 - \beta$, the signal is informative in the sense that $s_1 = 1$ with probability γ (and to 0 with probability $1 - \gamma$), and $s_2 = 0$. Recall that $\theta = 1$, by convention.

There is a sequence of N agents, N finite. Agent t is "active" in period t only. His decision is to monitor one of the two spots, *i.e.*, one of the two signals s_1 or s_2. After observing the signal, he makes a truthful report, that means he reports a sighting if and only he observes a signal 1. The incentive compatibility constraint for the report will be met if there is a reward for agents who have made a sighting in the correct spot after the true state is eventually revealed, once all agents have played.

The policy maker (the police), makes one and only one investment $x \in \{1, 2\}$ in one of the spots after all agents have played. The payoff of the investment is equal to 1 if $x = \theta$, that is if the police invests in the correct spot. The police maximizes its expected payoff.

Let μ_t be the public belief at the beginning of period t, *i.e.*, the probability that $\theta = 1$. In questions 1 and 2, the agents' reports are publicly available.

1. 1. Show that an agent monitors the spot 1 if $\mu_t > 1/2$, and the spot 0 is $\mu_t < 1/2$. (Ignore ties). Show that the monitoring choice by agent t is observable by others.
2. Assume that $\gamma = \alpha$ and assume $\mu_t > 1/2$. Determine μ_{t+1} if the observed signal is 1, and if it is 0. (Use Bayes' rule in the likelihood ratio or in the LLR). Show that the increase in LLR (if there is a sighting) is larger in absolute value than the decrease (if there is no sighting). Interpret this property as a kind of "herding" behavior. (You may take a numerical example to check some order of magnitude: $\alpha = \beta = \gamma = 0.2$).
3. Assume, to simplify the argument, that $\alpha = \beta = \gamma \leq 1/2$, μ_1 is strictly greater than 1/2 but vanishingly close to 1/2, and $N = 3$. The first agent makes his report to the police, but the police makes the official report to the other agent. Assume the first agent reports seeing a white van (a signal 1). Show that
   (a) if the police makes the report public, any report by the next two agents is irrelevant for the police's investment decision;
   (b) the police should deceive the other agents, if possible, and claim that the first agent has seen nothing.
   (c) Draw some conclusions about the social usefulness of the media in recent events.

EXERCISE 6.8. ('Yes men" for a partially informed receiver)

The exercise is based on an article by Prendergast (JPE, 1993). In all the models considered so far, the expert does not know the private information of the receiver about the state. This is critical since that information is used by the receiver for the reward function. If the expert knew this information he could change his message to manipulate the reward. Of course, the expert could do this only if the receiver did not know what the expert knows about him...

The state θ has a prior distribution that is normal with mean 0 and precision ρ_θ. The "boss" (Prendergast) has a signal on θ:

$$z = \theta + \epsilon_z,$$

where ϵ_z is normal with zero mean and precision ρ_z.

The expert has two signals, one on θ and the other on the information of the boss:

$$\begin{cases} s = \theta + \epsilon_s, \\ s_z = z + \eta, \quad \text{with} \quad \eta \sim \mathcal{N}(0, 1/\rho_\eta). \end{cases}$$

The expert's objective function is such that he attempts to fit the information of the boss. He sends a message m such that

$$m = E[z|s, s_z].$$

Solve the problem of the expert.