Chapter 1

Bayesian Inference

(01/24/2019)

A witness with no historical knowledge

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

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

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

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

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

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

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

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

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

A witness with historical knowledge

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

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

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

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

thinks he has seen a red cab. That witness who attempts to report the most likely event says also "yellow". The recipient of the two reports learns nothing from the reports. For him the accident was caused by a yellow cab with a probability of 90 percent.

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

Looking for your phone as a Bayesian

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

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

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

[4] Exercise: prove it.

1.1 The standard Bayesian model

1.1.1 General remarks

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

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

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

Learning is the processing of information that comes about the state. This information comes in the form of a *signal*. Examples are the witness report of the previous section, a weather forecast, an advice by a financial advisor, the action of some "other" individual, etc... In order to be informative, that signal must depend on the state. But that signal is imperfect and does not reveal exactly the state (otherwise 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.

1.1.2 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 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.

Observation (signal)

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

Table 1.1.1: Binary signal

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.

1.1.3 Multiple binary signals: search on the sea floor

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

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

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

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

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

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

An example: the search for AF447

In the early hours of June 1, 2009, with 228 passengers and crew, Air France Flight 447 disappeared in the celebrated "pot au noir".[5] No message had been sent by the crew but

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

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. [6] It will turn out that the search for AF447 did not follow the theory. Following Keller (2015), the search which lasted almost two years before a complete success, proceeded in stages.

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

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

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

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

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

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

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

[6] See L. Stone **.

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

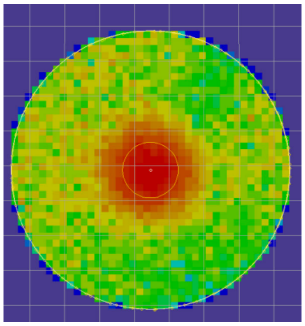

(A)
Prior probabilities

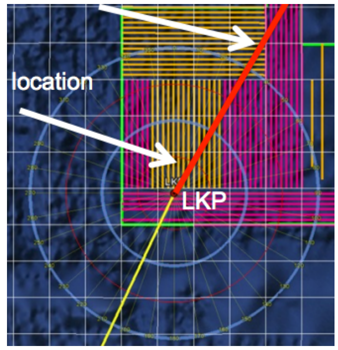

(B)
Search for pings

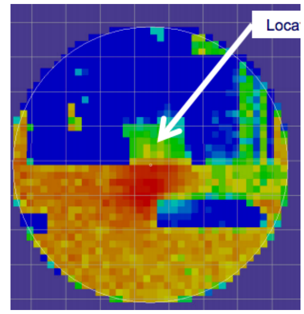

(C)
Posterior probabilities
after Stage 5

Wreckage

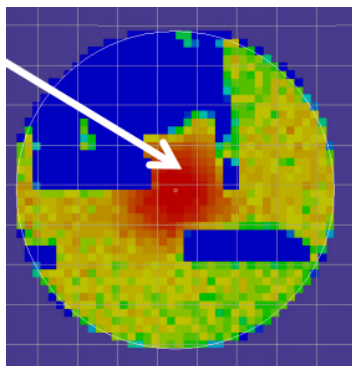

(D)
Posterior assuming beacons failed

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.

1.1.4 The Gaussian model

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

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

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

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

`These learning rules will be used repeatedly.`

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

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

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

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

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

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

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

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

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

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

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

1. The normal distribution is summarized by the two most intuitive parameters of a distribution, the mean and the variance (or its inverse, the precision).

2. The weight of the private signal s depends on the noise to signal ratio in the most intuitive way. When the variance of the noise term σ_ϵ^2 tends to zero, or equivalently

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

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

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

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

1.1.5 Comparison of the two models

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

The evolution of confidence

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

Bounded and unbounded private informations

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

Binary states and Gaussian signals

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

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

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

The LLR is updated according to

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

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

1.1.6 Learning may lead to opposite beliefs: polarization

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

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

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

the center policy is not working).

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

	Alice	Bob
1	0.47	0.06
2	0.47	0.47
3	0.06	0.47

Priors

	Alice	Bob
1	0.79	0.1
2	0.11	0.11
3	0.1	0.79

Posteriors

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

BIBLIOGRAPHY

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

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

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.

ANSWERS

Exercise 1.2

1. $p_D = 1 - (1-q)^2(1-s)^2$.

2. Let $p_A = 0.9$ and $p_B = 0.5$. Prob. that the wreck is in A is
$$\pi_A = \frac{P(d=0|A)P(A)}{P(d=0)} = \frac{0.1}{0.6} = \frac{1}{6}.$$
Obviously $P(B) = 5/6$. Since the probability of detection is so good on A and we hear nothing there (if it had been there we should have heard something), it is more likely that the wreck is in B.

Exercise 1.3

1. Denote by μ the prior on θ_1. After a signal 1, the likelihood ratio between the two states is
$$\frac{P(\theta_1|s_1)}{P(\theta_1|s_0)} = \frac{q_1}{1-q_0}\frac{\mu}{1-\mu}.$$
A signal is informative iff it changes the likelihood ratio. For this, $q_1/(1-q_0) \neq 1$, or $q_1 + q_0 \neq 1$. Note that both q_1 and $1-q_0$ can be smaller than $1/2$.

2. By the same argument, $q_1 + q_0 > q1$.

Exercise 1.4

1. Let f and $\hat{f}$ the prior and the posterior densities (on $(0,1)$).
$$\hat{f}(\theta) = A(1-\theta)f(\theta) = A(1-\theta),$$
where A is a constant such that the "sum of the probabilities" is equal to 1 and $\int_0^1 \hat{f}(\theta)d\theta = 1$. Therefore,
$$\hat{f}(\theta) = 2(1-\theta).$$

2. $$\hat{f}^n(\theta) = \frac{(1-\theta)\hat{f}^{n-1}(\theta)}{\int_0^1(1-\theta)\hat{f}^{n-1}(\theta)d\theta}$$
By recurrence,
$$\hat{f}^n(\theta) = (n+1)(1-\theta)^n.$$
As increases, the density is shifted to higher to a lower probability of success....

Chapter 2

Sequences of information and beliefs

2.1 Sequence of information with perfect memory

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

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

We assume here perfect memory of the past signals.

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

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

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

Stochastic path representations in probabilities

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

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

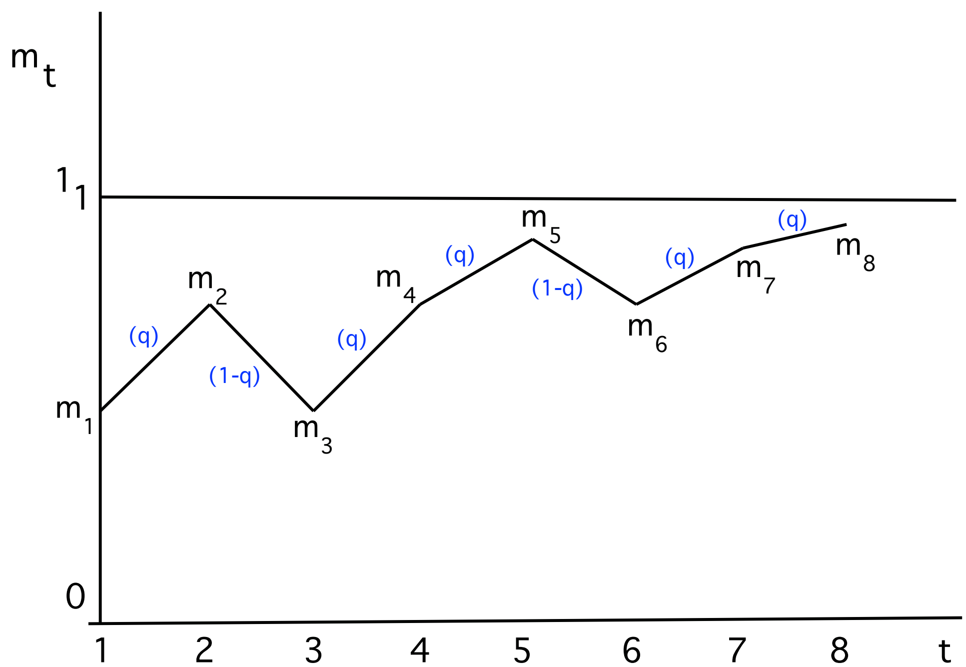

Figure 2.1: The evolution of belief as a stochastic process

Stochastic path representations in LLR

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

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

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

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

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

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

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

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

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

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

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

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

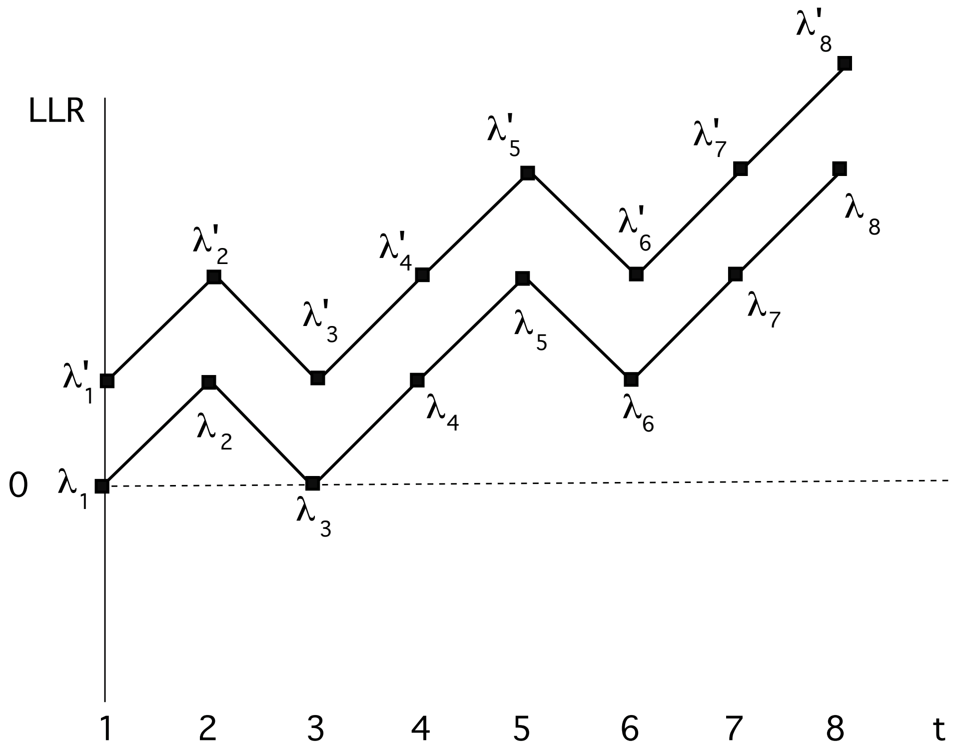

Figure 2.2: The evolution of LLs

Bounded and unbounded private informations

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

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

Examples:

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

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

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

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

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

2.2 Martingales

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

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

The martingale property with learning from a binary signal

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

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

An equivalent result is

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

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

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

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

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

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

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

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

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

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

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

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

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

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

2.3 Convergence of beliefs

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

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

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

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

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

Bayesian beliefs converge because of the Martingale Convergence Theorem.

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

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

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

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

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

[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.

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

PROPOSITION 2.4. *Let $\mathcal{A}$ be a subset of Θ and μ_t be the probability assessment $\mu_t = P(\theta \in \mathcal{A}|h_t)$, where h_t is a sequence of random variables in previous periods. Then there exists a random variable μ^* such that $\mu_t \to \mu^*$.*

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

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

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

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

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

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

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

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

The expectation of the number of upward crossing is bounded. From this, one can show that the probability of an upward crossing after period t tends to zero if t tends to infinity.

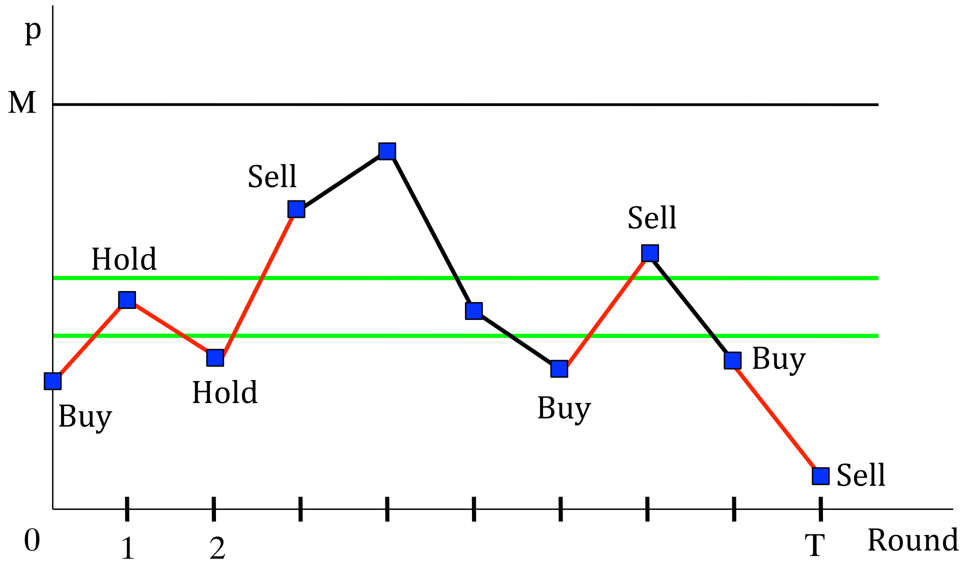

The agent holds one unit of the asset on the red segments.

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

One can then divide the interval $[0, 1]$ in n intervals, each of with $1/n$ and iterate the previous argument for the finite number n. That means that for any ϵ, the stochastic process stays within one of these bands except with probability ϵ. Since the number n can be take as large as one wants, that proves the convergence in probability.[4]

A heuristic remark on another proof of the Martingale Convergence Theorem

The main intuition of the proof is important for our understanding of Bayesian learning. It is a formalization[6] of the metaphor of the drunkard. In words, the definition of a martingale states that agents do not anticipate systematic errors. This implies that the updating difference $\mu_{t+1} - \mu_t$ is uncorrelated with μ_t. The same property holds for more distant periods: conditional on the information in period t, the random variables $\mu_{t+k+1} - \mu_{t+k}$ are uncorrelated for $k \geq 0$.

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

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

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

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

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

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

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

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

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

Rational (Bayesian) beliefs cannot cycle forever

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

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

[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.

Bayesian and do not lead to convergence. In Kirman (1993), agents follow a mechanical rule which can be compared to ants searching for sources of food, and their beliefs fluctuate randomly and endlessly.

The evolution of confidence

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

EXERCISE 2.1. (Non independent signals)

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

Answers

Exercise 2.1

$s_t = z_t + z_{t-1}$, where z_t are sequentially independent signals.

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.

Appendix

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

PROPOSITION 2.7. *Conditional on $\theta = \theta_0$, the likelihood ratio*

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