Chapter 1

Bayesian Learning

(02/02)

A witness with no historical knowledge

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

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

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

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

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

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

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

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

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

A witness with historical knowledge

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

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

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

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

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

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

Looking for your phone as a Bayesian

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

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

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

[4] Exercise: prove it.

1.1 The binary model

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

The sighting by the witness is equivalent to the reception of a signal s that can be 0 or 1. A signal that takes one of two value is called a *binary signal*. The uncertainty about 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 1.1.

Observation (signal)

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

Table 1.1: Binary symmetric signal

The table represent a *symmetric binary signal* (SBS) for which the probability of being right is the same in the two states, $P(s = \theta|\theta) = q$. For a SBS, q can be called the *precision* of the signal. Note that if $q = 1/2$, the signal provides no information. If $q = 1$, the signal reveals the state perfectly. If the elements in the diagonal are not the same, the signal is not symmetric. See exercise 1.2.

The previous Bayesian inference is repeated in this formal setting:

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

Likewise, $$P(\theta = 0|s = 1) = \frac{P(s = 1|\theta = 1)P(\theta = 0)}{P(s = 1)}.$$

Bayes' formula is unwieldy. It is often useful to express it 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):

$$\underbrace{\frac{P(\theta = 1|s = 1)}{P(\theta = 0|s = 1)}}_{\text{posterior LR}} = \underbrace{\left(\frac{P(s = 1|\theta = 1)}{P(s = 1|\theta = 0)}\right)}_{\text{signal factor}} \cdot \underbrace{\left(\frac{P(\theta = 1)}{P(\theta = 0)}\right)}_{\text{prior LR}}. \tag{1.2}$$

The signal factor depends only on the properties of the signal, as specified in Table 1.1. This multiplicative Bayesian rule illustrates how the power, or the precision of the signal, can overcome a prior LR. Here,

$$\frac{P(\theta = 1|s = 1)}{P(\theta = 0|s = 1)} = \frac{q}{1-q}.\frac{P(\theta = 1)}{P(\theta = 0)}.$$

In the previous case 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.

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

$$\lambda' = \lambda + a, \qquad \text{with the signal term} \qquad a = \frac{P(s = 1|\theta = 1)}{P(s = 1|\theta = 0)}. \tag{1.3}$$

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!

1.2 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

$$w'_i = A(1 - p_i) w_i, \quad \text{with} \quad A = \frac{1}{\sum_{i=1}^N (1 - p_i) w_i}. \tag{1.4}$$

An example: the search for AF447

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

[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.

[6] See L. Stone **.

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

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.

   (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 1.1.

   (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).

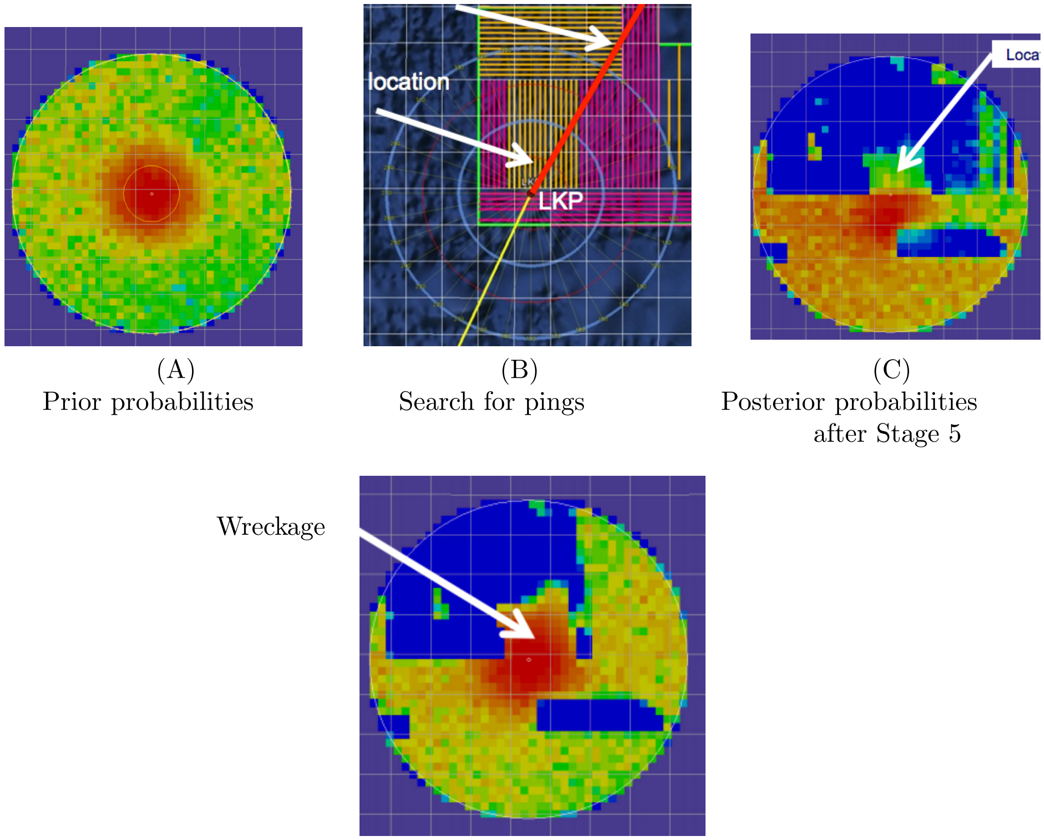

(A) Prior probabilities

(B) Search for pings

(C) Posterior probabilities after Stage 5

(D) Posterior assuming beacons failed

Source: Keller (2015).

Figure 1.1: Probability distributions in Bayesian search

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.

1.3 The Gaussian model

So far, we considered a finite number of states of nature. Now we look at a *continuum* of states. One example is the price of a crop that will be realized in the summer while farmers have to decide in the spring how much acreage to seed. A simple representation of the states and the information is provided by the *Gaussian model*. That model has technical properties that will facilitate the analysis of learning. As usual, a model is a representation that embodies more general features, and is also incompatible with some other features. That will be discussed after the presentation of the model.

- Nature's parameter θ is chosen randomly according to a *prior distribution* that is normal with mean $\bar{\theta}$ and variance $1/\rho_\theta$, $\mathcal{N}(\bar{\theta}, 1/\rho_\theta)$. It will be convenient to deal with the *precision* that is the inverse of the variance. The value of θ can therefore be negative but in the example of the crop price, we will assume that the normal distribution is an approximation. (We will *not* use a Log-normal at this stage).
- A signal s is received which is equal to the true value θ plus a noise :

$$s = \theta + \epsilon. \tag{1.5}$$

  The noise term ϵ has a normal distribution $\mathcal{N}(0, 1/\rho_\epsilon)$ and independent[8] of θ. The precision ρ_ϵ defines the quality of the information that is provided by the signal.

The update the prior distribution after the reception of the signal s is extraordinary simple: the posterior distribution is also normal with mean m and precision (inverse of the variance) ρ. We only need to know how *two* parameters are determined, and furthermore, they are determined by the following two equations that are very intuitive.

- The posterior's precision, ρ, is the sum of the precisions of the prior and of the signal:
- The posterior's mean, m is a weighted average of the prior's mean and of the signal value. Each of the two is weighted by its own precision.

$$\begin{cases} \rho = \rho_\theta + \rho_\epsilon, \\ m = (1-\alpha)\bar{\theta} + \alpha s, \quad \text{with} \quad \alpha = \dfrac{\rho_\epsilon}{\rho_\epsilon + \rho_t}. \end{cases} \tag{1.6}$$

1.4 Model properties

The Gaussian model has the following important properties.

[8] That is not a restriction since one can always "take out" of ϵ a non-zero mean and any component that is correlated with θ.

- First, it may be appropriate if there is a continuum of values of the state of nature
- The entire distribution is characterized by two parameters, which is actually the smallest possible number of parameters to define a distribution, the mean and the variance.
- The Bayesian updating rule could not be more intuitive or simpler.
- The signal is unbounded which is an important property that prevents the occurrence of an informational cascade.
- New information always reduces the variance of the distribution. Any news reduces the uncertainty about the state of nature. That property obviously does not hold in some important cases.
- The parameters of the posterior distribution are independent of the realization of the signal s. They can be computed before receiving the signal s. That property simplifies the Bayesian updates. In addition, it simplifies the decision of an agent who has to decide whether getting some signal will provide sufficient information.

The binary signal may useful for the following properties

- The number of the states may be finite, in which case the assumption of two states is often sufficient.
- The Bayesian updating rule is very simple.
- The signal is bounded, and this property is related to the occurrence of informational cascades.
- News may increase the uncertainty about the state of nature. If there are two states, the uncertainty is the greatest when both states are equally probably. Bad (good) news about a state increase the uncertainty when the prior of that state is high (low).

Bounded and unbounded updates

Consider the binary model and Bayes' rule in its LLR form. A good signal ($s = 1$) increases the LLR by $Log(1/(1-q)$ and a bad signal has a negative impact of the same magnitude. The magnitude of the update in LLR is bounded. We will say that in this case the *signal is bounded*. That property will be important later for the occurrence of informational cascades.

On example of unbounded signal is given by the combination of the binary and the Gaussian model. Take Θ with two elements, $\Theta = \{\theta_0, \theta_1\}$, and the Gaussian private signal

$$s = \theta + \epsilon, \qquad \text{with } \epsilon \sim \mathcal{N}(0, 1/\rho_\epsilon^2). \tag{1.7}$$

If λ is the LLR between states θ_1 and θ_0, Bayes' rule takes the form:

$$\lambda' = \lambda + \rho_\epsilon(\theta_1 - \theta_0)(s - \frac{\theta_1 + \theta_0}{2}). \tag{1.8}$$

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.

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

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

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

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

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.

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

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

EXERCISE 1.1. (Simple probability computation, searching for a wreck)

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

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

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

EXERCISE 1.2. (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 1.3. (Bayes' rule with a continuum of states)

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

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

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

Chapter 2

Sequences of information

2.1 Sequence of information with perfect memory

Suppose that $\mathcal{A}$ is a set of possible states, *i.e.* a subset of the space set. 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.

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 1.1. 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 $m_1 = m_3$, $mm_2 = m_4 = m_6$, $m5 = m_7$. 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.

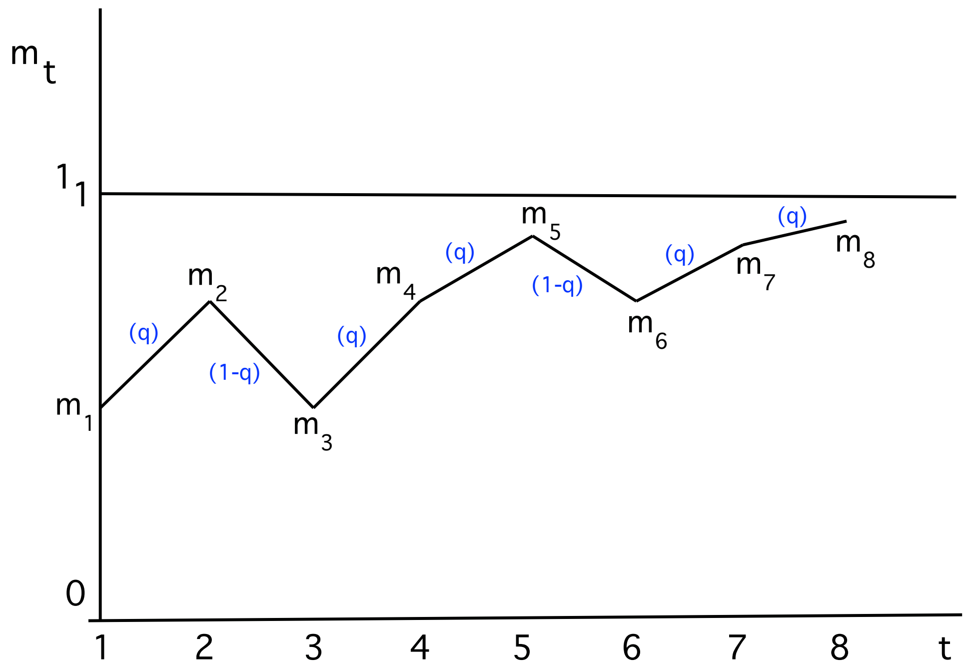

Figure 2.1: The evolution of belief as a stochastic process

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.

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?

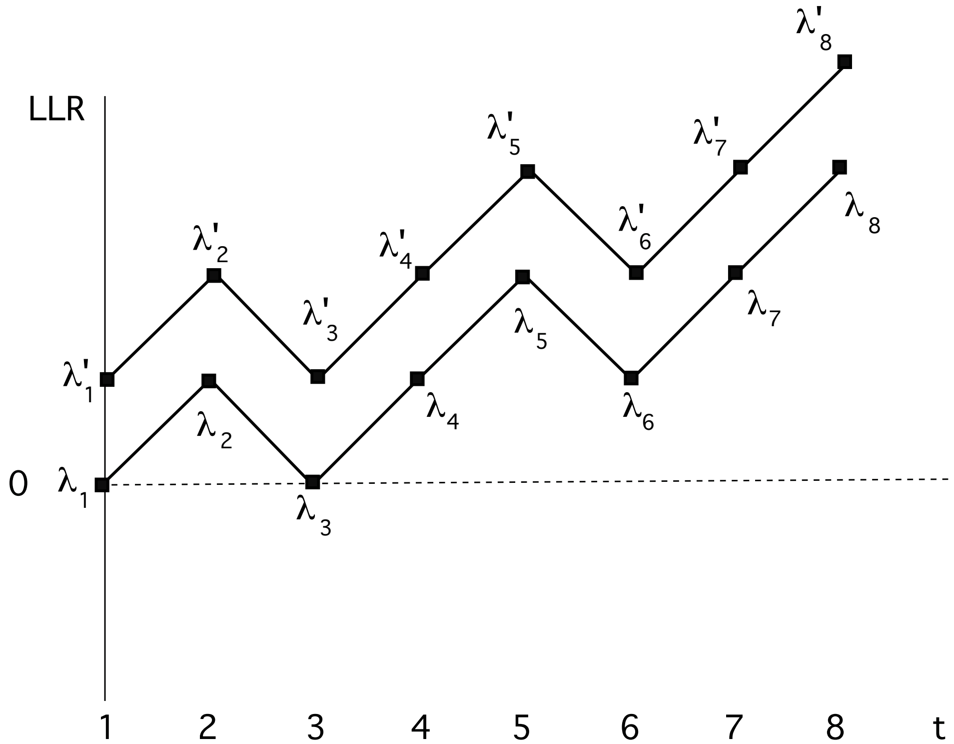

Figure 2.2: The evolution of LLs

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
Let there be two states $\theta \in \{0, 1\}$. Then $E[\theta] = P(\theta = 1)$. Let P and P' be the prior and the posterior (after the signal) of a state θ:

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

In general the martingale property of rational learning rests on the 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*

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

informations before period t; then $\mu_t = E[\mu_{t+1}|h_t]$.

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

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

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

2.3 Convergence of beliefs

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

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

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

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

Bayesian beliefs converge because of the Martingale Convergence Theorem.

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

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

[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.

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.

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

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

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

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

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

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

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

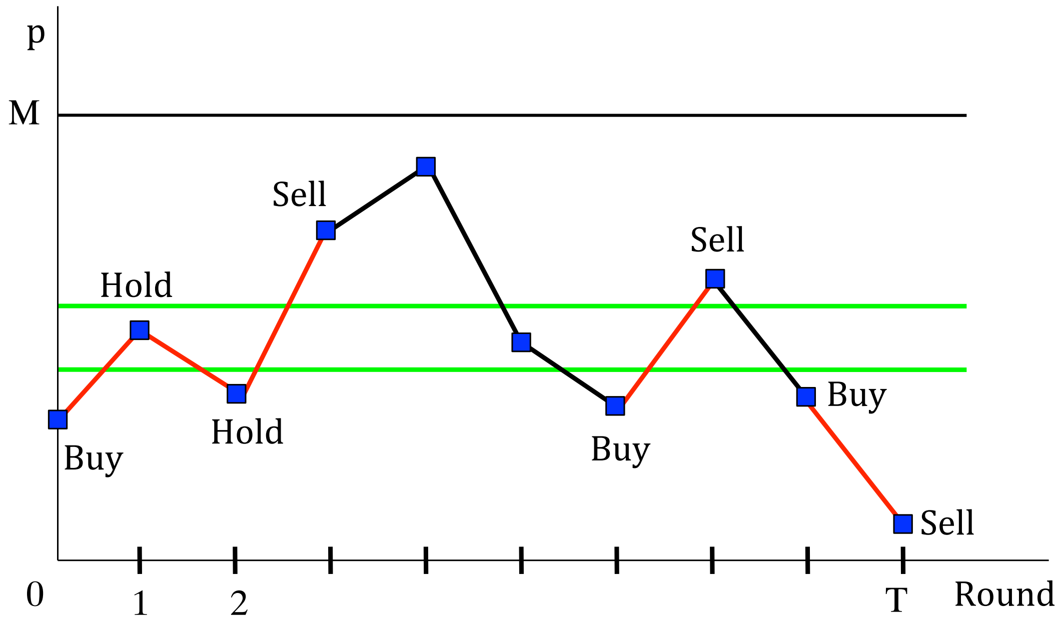

The agent holds one unit of the asset on the red segments.

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

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. 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]

`To be revised`

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

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

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

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

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,

[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.

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.

BIBLIOGRAPHY

EXERCISES