Chapter 6

Words

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

Trust but verify.

Communication with words is the subject of a vast literature. This chapter will be selective and focus on the relations between models of communication through words (models of "cheap talk") and issues of social learning that are addressed in other parts of this book. For example, herding may arise in financial markets because of the observation of others' actions or because of the behavior of financial advisors who are influenced by others' predictions. We will see that herding on actions and herding on words occur under similar conditions.

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

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

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

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

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

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

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

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

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

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

6.1 Advice by one expert

You asked for advice expert about the direction of the market. Say there are two possible future events, up or down; in mathematical language, the state of the world, to be realized later is $\theta \in \{0, 1\}$. Suppose that there is some general opinion about the future, from the press and other news, which is quantified by the prior probability of the event "up"

($\theta = 1$), to be μ, between 0 and 1. Your expert has some additional information (that is why he is called "expert"). The quality of this information is that the expert's prediction is correct with a probability q ($1/2 < q < 1$). This information is equivalent to a signal s that is equal to the state with probability q: $P(s = \theta|\theta) = q$).[1] f What you would like to know is the information of the expert, his signal value. But whether he will truthfully tell you his information depends on his reward.

We suppose him that you reward the expert by comparing his advice with the actual performance of the market that you observe later. Shouldn't the expert just tell you his information? You will compare the message of the expert with the state of nature that is revealed later. The expert knows that and has some information on the probabilities of the states in the future. If you reward him for being correct, his objective is to match his message with the most probable state in the future, and truth telling may not be the best strategy. If the "general consensus" is that the market will go up, the expert may follow the consensus even if his private information points the other way.

6.1.1 Evaluation payoff after verification

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

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

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

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

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

The truthtelling strategy is optimal if it yields to the expert a payoff which is not strictly

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

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

smaller than that obtained from deviating. For each signal value of the expert, there is an incentive compatibility constraint to tell the truth ($m_i = s_i$):

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

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

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

The condition for truthtelling is the same as the no herding condition in the BHW model.

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

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

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

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

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

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

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

6.1.2 Equilibrium with an evaluation based on reputation

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

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

The babbling equilibrium

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

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

The truthtelling equilibrium

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

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

Let $\mathcal{H}$ and $\mathcal{L}$ be the events that his signal has high or low precision. Suppose that the

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

expert tells the truth. $m(s) = s$). By Bayes' rule, the *ex post* reputation is

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

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

$$v_{11} = v_{00} > v_{10} = v_{01}. \quad (6.8)$$

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

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

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

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

The type of the expert is known

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

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

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

What have we learned so far?

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

A contrarian facing biased experts

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

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

Bayesian computation as before, the evaluation function is now defined by

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

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

Note that

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

where the inequality are strict if and only the bad expert lies with some probability $\nu > 0$. The inequality between $v_{1,1}$ and $v_{1,0}$ appears because the probability of a lie in the message 1 is higher when the true state is 0 compared to the state 0. When the message is 0, the probability of lying is the same in both states.

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

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

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

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

- Should the state be 0, it is always better to have lied all the time with a message 0 that turns out to be equal with the state.

[4]

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

Note that

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

- Should the state to be 1, we see from the inequality in (6.12) that his reputation is higher if he gives the wrong message 0 instead of 1.

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

PROPOSITION 6.3. *Assume that there is some probability that the expert is bad in which case he gives with probability $\nu > 0$ the advice $m = 1$ while having the signal $s = 0$, then a good expert with a signal 1 who cares only about this reputation to be good should lie all the time and give the message 0.*

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

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

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

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

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

---

[5]

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

6.2 Panel of experts

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

6.2.1 A sequence of experts with a pre-established order

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

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

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

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

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

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

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

The goal of the receiver is to choose the state which is most likely once he has listened to

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

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

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

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

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

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

Possible cases with a panel of two experts

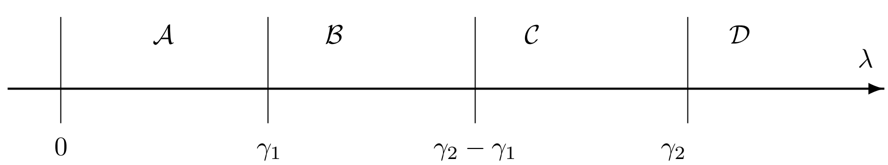

1. Suppose first that λ is in the interval $\mathcal{A}$: $0 \leq \lambda < \gamma_1$. If Junior speaks first, his signal is stronger than the public belief ($\gamma_1 > \lambda$) and he speaks the truth[10]. But since $\lambda + \gamma_1 < \gamma_2$, the public belief once he has spoken is smaller ("weaker")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

2. Suppose that λ is in the interval $\mathcal{B}$: $\gamma_1 < \lambda < \gamma_2$[11]. If Junior speaks first, he babbles. (His signal is weaker than the public belief). If Junior speaks second, he also babbles (as

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

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

can be verified). Junior is irrelevant. In region $\mathcal{B}$, the receiver never gets to observe the signal of Junior whatever the rule.

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

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

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

6.3 Bibliographical notes

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

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

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

Welch (2000) develops in a remarkable study an econometric methodology to estimate imitation when choices are discrete[12]. He analyzes how the probabilities of analysts' revisions of the recommendations (which take place in a set of 5 values from "strong buy" to "strong sell"), depend on the established consensus. His results indicate that some herding takes place, especially in a bull market. A next step in this research could be the construction of a structural model with both an exogenous process of information diffusion and learning from others, and the analysis of its empirical properties. (See also Grinblatt, Titman and Wermers, 1995).

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

EXERCISES

EXERCISE 6.1. Imperfect verification of the expert's message

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

1. Using the notation of Section 6.1 for the reward function v_{my}, determine the payoff function $V(s, m)$.
2. Establish the condition for truthtelling by the expert.
3. Show that if the reward function is such that $v_{00} = v_{11}$ and $v_{10} = v_{01}$, the condition for truthtelling is independent of $q \in (1/2, 1]$. Provide an intuitive interpretation.

EXERCISE 6.2. The value of reputation

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

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

EXERCISE 6.3. Computation of the reputation function

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

EXERCISE 6.4. (Partial truth telling)

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

EXERCISE 6.5. A continuum of beliefs

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

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

2. Determine the family of admissible functions.

3. Is the condition in question 1 sufficient?

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

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

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

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

EXERCISE 6.7. (The white van)

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

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

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

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

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

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

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

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

EXERCISE 6.8. (‘Yes men” for a partially informed receiver)

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

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

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

where ϵ_z is normal with zero mean and precision ρ_z.

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

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

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

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

Solve the problem of the expert.