

1 Introduction

Penguins are social animals. They live in groups above the water from which they get fish for food. Unfortunately, there is more than fish in the water. From time to time, killer whales (orcas) roam under the surface waiting for some prey. Penguins are aware of the danger and would like to have some information before taking a plunge. Indeed, any sensible penguin thinks that it would be very nice if some other penguin would dive first to test the water. So what is a penguin to do? Wait. Possibly some other member of the colony who is more hungry, or has other information, will go first. Is it possible that no penguin will ever go? No, because waiting becomes more costly as hunger increases. Eventually, one or more penguins will take the plunge, and, depending on the outcome, the others will either stay put or follow *en masse*. This *waiting game* is socially inefficient. It would be better if the first individual would decide to go at least a bit earlier: the first to go is, on the margin, as well off going just a little earlier; but others strictly prefer him to go a little earlier. Actually, the penguins are well aware of this social inefficiency, which they try to remedy by pushing some poor fellow off the cliff.

First-Cousin Marriages

There is a long-standing taboo against marriages between first cousins in some parts of the world. Such taboos may entail significant costs. In the United States, about thirty states have laws forbidding first cousins to marry, but on some other continents marriages between cousins are well regarded. Recent, and remarkably late, evidence shows that the risk of defects for children is marginally higher in such marriages than in the general population (Motulsky et al., 2002):

Dr. Motulsky said that medical geneticists had known for a long time that there was little or no harm in cousins' marrying and having children. 'Somehow, this hasn't become general knowledge,' Dr. Motulsky said. 'Among the public and physicians there's a feeling it's real bad and brings a lot of disease, and there's a lot of social and legal disapproval.' Dr. Motulsky said the American laws against cousin marriage should be abolished, because they are based in part on the mistaken belief that the children of such parents will suffer from terrible physical and mental illnesses.¹

¹ This quotation is from the *New York Times*, April 3, 2002.

The White-Van Frenzy

On October 26, 2002, one could read this in the *New York Times*:

Until the final moments of the three-week reign of sniper terror in the Washington area, there was one image, one all-consuming clue, seared into the minds of a panicked public: the white van.

White vans were pictured on wanted posters everywhere. Their drivers were stopped four, five, six times a day. They were forsaken by some, who rented or borrowed other vehicles because they simply could not take it anymore: the constant traffic stops, the swooping in of gun-wielding police officers, the stares from other drivers and pedestrians, the snooping around their license plates and sneaky peeks into their rear windows to check for rifles and other signs of sniping. Each shooting and each day that passed without an arrest brought a new flood of tips and witness accounts involving white vans and trucks, seemingly fed by the earliest witness accounts.

Witnesses to some of the first shootings were able to describe only a glimpse at a fleeing white vehicle, information that the police quickly released and eventually used to put together composite sketches. So there was nothing but a white van and before that, a white box truck to look for. With no description of the killer, many people, gripped with fear of another attack and seizing on any detail in their personal lookout for the sniper, were in a kind of white-van haze. [...]

For a time, the police themselves were so focused on white vans and trucks – there are roughly 100,000 on the roads of Washington and its suburbs – that they may have overlooked the vehicle they really needed to find. In fact, officers in Washington stopped the blue Caprice, ran a check on its license plates and then allowed the two men to proceed only a couple of hours before one of the sniper killings. [...]

“Darn right I’m glad it’s over,” said Sinclair Skinner, who drives a white van for his dry-cleaning delivery service but parked it two weeks ago after being stopped by the police twice in one day, and rented a car. “The police were stopping people, people were asking me all kinds of questions. Finally I said: ‘Look, I’m not involved in anything. I’m not the Taliban. I wasn’t in the gulf war. I’m a pacifist. I listen to Pacifica Radio.’” [...]

Meanwhile, the police, including all 1,000 officers on the Montgomery County force, the lead agency in the investigation, combed the area for white vans. When a van was stopped, several police cars often surrounded it, and officers drew their guns, sometimes telling the drivers to come out with their hands up and then ordering them to lie face down on the ground.

The white van, said Officer Baliles, “was the best lookout we had.”

Umbrellas

The actions of others convey some information about the state of the world. When I see other people going out with an umbrella, I take an umbrella along without checking the weather forecast. I do so because I know the decision model of others. I can then infer their private information from their action without the need to talk to them. This herding is rational. There is, however, the possibility that everyone carries an umbrella because someone carries an umbrella. The herd may be wrong.

1.1 Overview

3

1.1 Overview

The learning from others can operate through different channels of information (choices or results of actions, words), and in different contexts (with or without delay between actions, real or financial investment) and may be affected by externalities between the actions. The chapters are designed to study one issue at a time.

PART I: SOCIAL LEARNING

The Tools

All models will have a key random variable, which will not be observed directly: the *state of nature*. The state of nature may be the profitability of a set of new techniques, the mass of agents who are prepared to topple a government in a revolution, or the reserves of a central bank in a speculative attack against a regime of fixed exchange rate. Rational learning in this book means Bayesian learning about the state of nature.

The essential tools of analysis are presented in Chapter 2. The private information is represented by private signals, which depend on the state of nature: for example, if it will rain (if the sun will shine), two-thirds (one-third) of the agents think that rain is more likely than sunshine. To think that rain is more (less) likely is the realization of a private signal. These private signals are correct, statistically, and a large number of them, if they could be pooled directly, would bring perfect or near-perfect information.

We will use restrictive models, which will be special cases. Attempts at generality would end in futile formalism, but we will need to have a good idea whether the results of a specific model are robust or not. In most cases, these results depend on some critical assumptions. It is therefore important to know the implications of the modeling choices. Two models of private information will be used recurrently (but not exclusively): the binary model with two states (“good” or “bad”), two informations, and two actions (“investment” or “no investment”), and the Gaussian–quadratic model where all random variables are normal and the payoff is a quadratic function of the level of action (a real number). These two models are discussed in detail in Chapter 2.

The Martingale Property

The beliefs of agents are defined as their subjective probabilities about the state of nature. An agent observes a variable (aggregate investment, success or failure of an oil drilling) that depends on the state of nature. Following an observation, he updates his belief using Bayes’s rule. This rule is an application of the calculus of conditional probabilities, and it generates the remarkable property of a martingale: an agent knows that he may change his belief after an observation, but the expected value of the change is nil. If that value were different from zero, the agent would change his belief right away. The martingale property, which is similar to

an efficient-market equation in finance (and for good reasons), is both simple and powerful; it implies that the agent cannot change his belief forever. The martingale convergence theorem is one of the most beautiful theorems in probability theory. Its mechanism is presented here intuitively. The theorem is the most important one in Bayesian learning. It will apply in many models (but not in all of them) and facilitate the analysis of the limit properties of learning. The theorem implies that unending fluctuations of individual beliefs cannot be compatible with rational learning.

Social Learning

When an individual learns from a person's behavior, (i.e., her choice of action or her words), he learns because the person's behavior is motivated by some information that she has about a state of nature of interest to all (e.g., the example of the umbrellas). This *private information*, or private belief, is like an endowment. The process of social learning is the diffusion of the private beliefs to all individuals through the interactions of observations, learning, and action choices. The structure of the model will be dictated by the context. Actions may be more reliable than words. Some definitions of social learning restrict the learning to the observation of actions themselves (physical actions or words), but we will consider also the observation of the outcomes of actions.

In models of social learning, a large amount of (private) information may be hidden. When agents do not act, this information remains hidden. We will encounter situations where, because of some small changes, some agents will “come out of the woods,” take an action, and thereby release their information and induce others to act. Long periods of low activity may be followed by a sudden boom. Inversely, a regime of high activity may be brought to a halt by a crash.

Actions as Words

When agents learn from actions, these actions are the “words” for communication. As in any language, communication is easier if there are many words. When actions can be taken from a wide set (e.g., the set of real numbers), the “message” sent by an agent can reflect his private information perfectly. When the set of actions is discrete (e.g., a farmer choosing the type of crop), the message about the private information is obviously more coarse. We begin with the case of continuous action sets.

Clustering

The model of social learning is introduced in Chapter 3: Agents are placed in an exogenous sequence of rounds; each agent chooses his level of *investment*, a real number, in his round and is observed by others. As the number of observations increases, the importance of history for individual choices grows and agents rely less, rationally, on their private information. Individual choices are *clustering*. However, any agent's action can be observed with perfect precision – in this model – and it reveals perfectly his private information. The learning from others' actions is equivalent to

1.1 Overview

5

the direct observation of their private information. There is no information loss in social learning.

The Weight of History and Slow Learning from Others

The assumption of perfect observability requires also perfect knowledge of the decision process of each agent. However, we have all our idiosyncrasies, which make the inference from actions imperfect. When history is short, this problem is not important. When history is long, however, and it induces agents to cluster, the observation becomes dominated by the noise. This slowing down in the rate of social learning is quantified by Vives (1995) in the Gauss–quadratic model. We will find this important property of social learning in other contexts: the memory of past actions may reduce – or completely prevent – social learning.

Herds and Cascades

In Chapter 4, the set of actions is reduced to two elements (to take or not an action: fixed investment, medical procedure, new crop). This is the celebrated model of Bikhchandani, Hirschleifer, and Welch (1992), hereafter BHW. In a simple case (with binary information), they show that social learning stops completely and rapidly (at an exponential rate): Suppose two agents invest and thus reveal that they have a good signal; if the third agent has a bad signal, that signal cancels only one of the previous two signals; he should invest like the first agent, who had only one good signal to rely on; in this case, he invests even if he has a bad signal. The third agent is *herding*. His action conveys no information. In the next round, the public information is unchanged. Nothing is learned. The fourth agent is in the same position as the third. Social learning stops completely, and a *cascade* takes place.²

Cascades are a spectacular example of the failure of social learning due to the weight of history. Is this property robust? Careful examination shows that, technically, the property is not robust at all. Cascades depend on discrete private signals that are not generic. When private beliefs are all distinct – as in random drawings from a continuum – cascades do not exist.

On the other hand, *herds* take place for sure, eventually. A herd is defined as a situation in which all agents take the same action after some date. If all agents turn out to take the same action after some date T , there is still the possibility that an agent has a private signal that induces him to take a different action after date T . This possibility is not realized, but the very fact that it is not realized yields information, which is incorporated in the social learning (Smith and Sørensen, 2001).

The existence of a herd is an elegant application of the martingale convergence theorem: when a “black sheep” breaks away from the herd, he reveals a piece of strong private information. That information is then incorporated in the public belief, which

² See Section 1.2 for a discussion of the term “cascade”.

makes a quantum jump. This event cannot occur an infinite number of times, because that would prevent the convergence of the public belief, which is a martingale.

A herd is not a cascade. However, as a herd goes on, the set of private beliefs in which an agent follows the herd grows. Social learning is very slow, because the probability of breaking the herd must be vanishingly small in order for the herd to be realized. The time profile of the evolution of beliefs is not the same as that of a cascade, where learning stops completely, but it is not very different.

An application of the model shows that a firm that introduces a new product of quality unknown to consumers should target a group willing to pay a high price, in order to establish its reputation, and then lower the price to reach a mass market.

To Forget in Order to Reduce the Weight of History

Because the weight of the information in history induces paralysis and prevents agents from conveying their information to others through their actions, the remedy is simple: partly forget the past. Indeed, Chapter 5 shows that when agents have an incomplete sample of the actions of past agents, social learning is more efficient! The inference from another individual's action is now more difficult, because the observations on which that action is based are not observable. (In Chapter 4 these observations were summarized by the public belief, which is a martingale.)

Delays Ending with a Bang or a Whimper

So far agents have made a choice in an exogenous order. This assumption cannot apply to the penguins. In Chapter 6, the setting is the same as in BHW (with arbitrary private beliefs) with the sole difference that any agent can make a decision (like plunging) at any time: each agent has an option to make one investment. This setting induces a *waiting game*, in which the more optimistic agents go first. The others know this and wait to observe how many agents plunge. A large number means that the number of optimists is large. Because private beliefs are statistically correct, the large number is a signal that the state of nature is good, and agents who were initially reluctant join the fray. A bang may take place. Events can evolve either way: there is also the possibility that the number of investors is small in an initial phase, after which investment stops completely. The game ends with a whimper (Chamley and Gale, 1994).

The waiting game has some powerful properties. The main one is the arbitrage between the cost of delay (from lost dividends) and the option value of delay (from not making a bad irreversible investment). Suppose that actions are observed through some noise. The cost of delay is not affected by the noise. Hence the information generated by the equilibrium is unaffected. The equilibrium strategy is adjusted so that more agents take action, and the net information generated by all actions is the same as without noise. The model may generate multiple equilibria with sharply different information: if most agents delay and only the greatest optimists invest, then the information generated by the equilibrium must be high (in order to induce most agents to wait). If most agents do not delay, there is little information in the

1.1 Overview

7

equilibrium. In the cascades of BHW, there was no information. These situations are found again in the setting with endogenous delays.

In Chapter 7, the model with delays is investigated under the assumption of continuous time: agents can make a decision at any point in time. If the distribution of private beliefs admits agents with a high degree of optimism, the model does not have an equilibrium, generically. It is shown that there are essential differences between a model with periods and a model in continuous time. The properties of a model with periods do not converge to those of a continuous-time model as the period length becomes vanishingly short.

Outcomes

When agents observe the outcomes of others' actions, learning may still be incomplete if actions converge to a point where the random outcome can be explained by different states, as shown in Chapter 8. For example, the individual level of effort toward a successful economic career may converge, generically, to a value where agents cannot discriminate perfectly between the contributions of noise and chance in the outcome; or the price of a monopoly may be compatible with two types of demand schedule: the first with a high level and a high price elasticity, the second with a low level and a low elasticity.

The observation of outputs does not prevent failures in social learning: if an action is deemed superior to others, it is chosen by all agents who do not try other actions. Hence, no information can be gathered on other actions to induce a switch. A herd may take place unless some agents have very strong private beliefs that other actions may be superior. These strong believers may provide significant information benefits to others.

Networks

When agents are watching the evolution of the price of an asset, the information has a sequential structure. When they learn how to fertilize apple trees in Ghana, they rely on a network of information contacts. When a new crop (say, a high-yield variety of wheat or rice) is introduced in a region, agents make simultaneous decisions in each crop cycle, and learn from their information contacts at the end of each cycle. Networks and diffusion processes are analyzed in Chapter 9. As in previous chapters, social learning may be more efficient if the information of agents is restricted. For example, the observation of a small group (e.g., a royal family) by *all* agents could induce a cascade and a herd. The model of Bala and Goyal (1998) formalizes the description of Tarde (1900), which will be presented below.

Words

The most obvious way to communicate private information seems to be to talk. Yet, economists prefer to trust what people do. Can a financial advisor be trusted? Yes, if he has sufficient incentive to build a *reputation* for his future business. In

Chapter 10, it is shown that the reputation motive can be sufficient for truth-telling within some limits. The reputation is earned by being proven right: the receiver of the advice can verify, after a while, whether the stock that was recommended actually went up or down. However, if the public belief that the stock will go up is strong, a rational advisor who has contrary private information will maximize the chance of being right and will issue a “buy” recommendation. He is herding like any agent in the BHW model. Indeed, there is no difference between reputation herding and the herding in the BHW model, and the conditions for its occurrence are identical (Ottaviani and Sørensen, 2000). The reputation can induce people to herd and say the “politically correct” thing (Morris, 2001) or to please the boss (Prendergast, 1993). Recall the witnesses of the white van, or how the lonely juror was able to sway the others in *Twelve Angry Men* (Lumet, 1957). Who should speak first in a jury where people may herd on what others say? The analysis shows that it is not clear whether the first to speak should be among the more or the less experienced (Ottaviani and Sørensen, 2001).

PART II: COORDINATION

The issues of Part I are extended in Part II with the addition of payoff externalities between different individual actions. Agents have to coordinate their expectations about an equilibrium.

One Period without Learning

Rousseau (1762) described the problem of social cooperation in the model of the stag hunt: a hunter has a higher individual payoff if he participates in a stag hunt than if he goes on his own chasing a hare. However, the success of the stag hunt depends on the participation of all. There are two equilibria in this game with *strategic complementarity*: either all hunters participate in the stag hunt or none do. This story can be adapted to a number of situations in economics, from business cycles to speculative attacks against a fixed exchange rate, and in other social situations (e.g., revolutions).

How do agents “choose” between different equilibria? A natural first step for the analysis is to consider a one-period setting in which agents make a simultaneous decision without observing what others do. Each agent has to guess, if possible, which equilibrium agents coordinate on. Carlsson and van Damme (1993a), in their *global-game* method, build on the insight that agents do not have perfect information on each other. The removal of the assumption of *common knowledge* enables one to solve the coordination problem: agents with a low cost of investment invest, no matter what others do. With that information commonly known, agents with a slightly higher cost also invest, which is again common knowledge. A process of contagion takes place, which solves the decision problem for all agents with a cost lower than average. A similar process of *iterative dominance* induces agents with a cost higher than average not to invest. This process takes place in the heads of agents in “virtual time” and is called *eductive learning*.

1.1 Overview

9

A similar process of iterative dominance had been proposed previously in the context of *strategic substitutability* when, for example, farmers in a population choose their individual supplies independently and the expectations about future prices are inversely related to the total supply. In that case, there is a unique Nash equilibrium under perfect information, but the issue is the *coordination of expectations*. The analysis in Chapter 11 highlights the similarities and differences between strategic substitutability and complementarity.

A new tool, the *cumulative value function*, provides a simple intuition for the global-game method (with strategic complementarities) and enables one to solve applied models rapidly. Speculative attacks against currencies, which will reappear at the end of the book, have become a popular topic recently and are presented as an illustration.

Switching Regimes

In 1989, countries in the former Soviet bloc switched their political regimes abruptly. The opportunities for change were not apparent in the preceding winter.³ It is the essence of oppressive regimes that a large amount of opinions remains hidden (Kuran, 1995). A turning point was the series of demonstrations in Leipzig, where protesters realized that they were a large mass. After the fall, few people seemed surprised. The story's features include a strong complementarity between the actions of agents and a large discrepancy between the common knowledge before and after the fact. These features are the main ones in the model presented in Chapter 12, which is based on Chamley (1999). The model is suggestive of regime switches in a variety of contexts of social interactions (political regimes, business cycles).

In each period, there is a new population where each agent decides whether to invest (*protest*, in a political context), or not. Investment entails a fixed cost, and the return is positive only if a sufficient mass of agents invest in the same period. The costs of agents have some unobserved distribution. Each agent knows only his own cost and observes the history of aggregate activities in the previous periods. The distribution of costs evolves slowly and randomly. The model exhibits the properties described in the previous paragraph: there are random switches between regimes of low and high activity. Most of the time, the structure of costs is such that under perfect information there would be two equilibria, with high and low activity. Under perfect information, however, the level of aggregate activity moves between a high and a low value with significant *hysteresis*, or inertia.

The inertia can be described precisely. If an outside observer had perfect knowledge about the structure of the economy, he would observe that in any period, agents

³ We may tend to forget the *ex ante* beliefs before the switch of a regime. A useful reminder is found in Halberstam (1991) with an account of Henry Kissinger addressing, on February 26, 1989, the governors of the fifty states of the union: "He was condescending about what Gorbachev was doing and he was even more condescending about those poor Americans who were taking it all so seriously."

coordinate on the equilibrium that is closest to the one in the previous period. If the structure of the economy is such that low activity is an equilibrium under perfect information in periods t and $t + 1$ and is an equilibrium under imperfect information in period t , then low activity is also an equilibrium under imperfect information in period $t + 1$ (even if a switch to high activity would be possible under perfect information).

The equilibrium with random switches between high and low activity is determined by a process of iterative dominance in what is so far the only model in the literature that provides a nontrivial extension of iterative dominance to many periods. A numerical simulation illustrates an evolution of public belief that is analogous to the evolution in 1989.

Delays and Externalities

Firms that enter a new sector generate negative externalities on each other because they lower the price of the good. At the same time, they provide outsiders some information about the cost of production or the size of the demand. The prospect of more information induces delays. When the externality is positive, a small number of agents may induce others to act, and coordination may be achieved (Gale, 1995). If the number of agents is large, however, the possibility of delay may prevent coordination under imperfect information.

PART III: FINANCIAL HERDS

Since the tulip mania in seventeenth-century Holland, spectacular rises and falls have been observed in financial markets. It is easy to dismiss these as follies. An evaluation of the “irrationality” of markets can be made only with respect to properties that may be observed in a “rational” market. Part III provides an introduction to herds in financial markets with rational agents.

In the standard model of social learning (Part I), agents learn about the state of the world. The specifications of the set of actions and of the payoff externalities play a critical role in determining the properties of social learning. In financial markets, the state of the world is the fundamental value of an asset, the actions are the trades, payoff externalities arise because the gains of some are the losses of others, and timing is essential. Two market structures are considered.

Sequential Trades

The model of social learning with individuals in an exogenous sequence (Part I) becomes in financial markets the model of sequential trades between agents with different information (Glosten and Milgrom, 1985). If the learning is about a fundamental value that is a real number, beliefs updated by financial transactions with asymmetric information converge rapidly to the truth. This property is not very surprising in view of the results in Part I. Because a price can take any value in the set of positive numbers, it is a fine signal about the agents’ private informations (as the