

Cascades and Herds

Ec 717

November 4, 2024

Herds

Assumptions:

- Standard model with sequential decisions
- Actions are taken in the discrete set $X = \{0, \dots, K\}$. For simplicity, $K = 1$.
- State $\omega \in \{0, 1\}$.
- Payoff $x(\omega - c)$, $c \in (0, 1)$. For simplicity, $c = 1/2$.
- Agents have a private belief s with a distribution that depends on ω and is not bounded (to make it interesting), with cdf such that $F_0(s) > F_1(s)$, (stochastic dominance).

Definition A **herd** takes place if there exists T such that for any $t > T$, all the actions are identical.

Remark: a herd is an *ex post* concept.

Theorem (Smith and Sørensen)

With probability 1, the public belief converges to the truth and a herd takes place with probability 1.

Proof

First step:

- By contradiction, assume $\mu_t \rightarrow \mu^* \in (0, 1)$.
- For any $\epsilon > 0$, vanishingly small, $P(|\mu_t - \mu^*| < \epsilon) \rightarrow 0$. (1)
- $x_t = 1$ if the belief of the agent is at least $1/2$, or $\frac{\mu_t}{1 - \mu_t} \frac{s_t}{1 - s_t} > 1$ or $s_t > 1 - \mu_t$.
There exists $\alpha > 0$ such that if $\mu_t \rightarrow \mu^*$, $P(x_t = 1) \geq 1 - F_0(1 - \mu_t) > \alpha$. (2)
- Following $x_t = 1$, $\frac{\mu_{t+1}}{1 - \mu_{t+1}} = \frac{\mu_t}{1 - \mu_t} \frac{1 - F_1(1 - \mu_t)}{1 - F_0(1 - \mu_t)}$.
- When $\mu_t \rightarrow \mu^*$, there exists $\beta > 0$ such that following $x_t = 1$, $\mu_{t+1} - \mu_t > \beta$. (3)
- (2) and (3) contradict (1).

Proof

Second step:

- If $\omega = 0$, $\mu^* \neq 0$. (A martingale cannot be absolutely wrong). Therefore $\mu^* = 1$.
- Suppose that $\omega = 1$. $\mu_t \rightarrow 1$. If there is a strictly positive probability that a herd does not take place, then for any T , there is $t > T$ such that $\mu_{t+1} < \mu_t - 1/2$, a contradiction of the convergence to 1.

Rational learning cannot be totally wrong, asymptotically

Proposition

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

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

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

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

Proof

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

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

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

Q.E.D.

Efficient learning in the quadratic model

- Nature's parameter ω fixed for all periods before the first period according to $\mathcal{N}(\bar{\omega}, 1/\rho_\omega)$.
- Each individual t has one private signal $s_t = \omega + \epsilon_t$, with $\epsilon_t \sim \mathcal{N}(0, 1/\rho_\epsilon)$.
- All individuals have the same payoff function $U(x) = -E[(x - \omega)^2]$. Individual t chooses his action $x_t \in \mathcal{R}$, in period t : exogenous order.

Payoff function very convenient because the optimal action is the expected value of ω . The decision rule is the same when the payoff is $u(x, \omega) = 2\omega x - x^2$, where x can stand for the scale of an investment.

- The public information at the beginning of period t is made of the initial distribution $\mathcal{N}(\bar{\omega}, 1/\rho_\omega)$ and of the history of previous actions $h_t = (x_1, \dots, x_{t-1})$.

Steps in the quadratic model

- 1 The belief of agent t Using the standard Bayesian formulae for Gaussian distributions, the belief of agent t is $\mathcal{N}(\tilde{\mu}_t, 1/\tilde{\rho}_t)$ with

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

- 2 Agent t action:
$$x_t = \tilde{\mu}_t = (1 - \alpha_t)\mu_t + \alpha_t s_t. \quad (2)$$

- 3 Social learning
$$\begin{cases} \mu_{t+1} = (1 - \alpha_t)\mu_t + \alpha_t s_t, & \text{with } \alpha_t = \frac{\rho_\epsilon}{\rho_\epsilon + \rho_t}, \\ \rho_{t+1} = \rho_t + \rho_\epsilon. \end{cases} \quad (3)$$

- 4 The history of actions $h_t = (a_1, \dots, x_{t-1})$ is informationally equivalent to the sequence of signals $(s_1, \dots, s_{t-1})$.

Remarks on the quadratic model

1 The weight of history and imitation

Agent t chooses an action which is a weighted average of the public information μ_t from history and his private signal s_t (equation (2)). The expression of the weight of history, $1 - \alpha_t$, increases and tends to 1 when t increases to infinity.

Imitation increases with the weight of history, but does not slow down social learning if actions reveal private informations. The weight of the private signal tends to zero. Hence, agents tend to “imitate” each other more as time goes on. This is a very simple, natural and general property: a longer history carries more information. Although the differences between individuals’ actions become vanishingly small as time goes on, the social learning is not affected because these actions are perfectly observable: no matter how small these variations, observers have a magnifying glass which enables them to see the differences perfectly. In the next section, this assumption will be removed. An observer will not “see” well the small variations. This imperfection will slow down significantly the social learning.

Imperfect knowledge on other payoffs

- In the previous section, an agent's action conveyed perfectly his private information. An individual's action can reflect the slightest nuances of his information because: (i) it is chosen in a sufficiently rich menu; (ii) it is perfectly observable; (iii) the decision model of each agent is perfectly known to others.

- **A simple representation**

$$U(x, \eta_t) = -E_t[(x_t - \omega - \eta_t)^2], \quad \text{where } \eta_t \text{ is private to agent } t.$$

The optimal action is $x_t = E_t[\omega] + \eta_t$. Since the private parameter η_t is not observable, the action of agent i conveys a *noisy signal* on his information $E_i t \theta$.

- Imperfect information on an agent's private characteristics is operationally equivalent to a noise on the observation of the actions of an agent whose characteristics are perfectly known.

One action per period

■ Individual action $x_t = (1 - \alpha_t)\mu_t + \alpha_t s_t + \eta_t$, with $\alpha_t = \frac{\rho_\epsilon}{\rho_t + \rho_\epsilon}$. (4)

■ The evolution of beliefs $x_t = (1 - \alpha_t)\mu_t + \alpha_t \omega + \underbrace{\alpha_t \epsilon_t + \eta_t}_{\text{noise term}}$ (5)

■ Normalization $\frac{x_t - (1 - \alpha_t)\mu_t}{\alpha_t} = y_t = \omega + \epsilon_t + \frac{\eta_t}{\alpha_t}$. (6)

One action per period

■ Individual action $x_t = (1 - \alpha_t)\mu_t + \alpha_t s_t + \eta_t$, with $\alpha_t = \frac{\rho_\epsilon}{\rho_t + \rho_\epsilon}$. (4)

■ The evolution of beliefs $x_t = (1 - \alpha_t)\mu_t + \alpha_t \omega + \underbrace{\alpha_t \epsilon_t + \eta_t}_{\text{noise term}}$ (5)

■ Normalization
$$\frac{x_t - (1 - \alpha_t)\mu_t}{\alpha_t} = y_t = \omega + \epsilon_t + \frac{\eta_t}{\alpha_t}. \quad (6)$$

■
$$\begin{cases} \mu_{t+1} = (1 - \beta_t)\mu_t + \beta_t \left(\frac{x_t - (1 - \alpha_t)\mu_t}{\alpha_t} \right), \text{ with} \\ \beta_t = \frac{\sigma_t^2}{\sigma_t^2 + \sigma_\epsilon^2 + \sigma_\eta^2 / \alpha_t^2}, \quad \text{with } \sigma_t^2 = \frac{1}{\rho_t}, \\ \rho_{t+1} = \rho_t + \frac{1}{\sigma_\epsilon^2 + \sigma_\eta^2 / \alpha_t^2} = \rho_t + \frac{1}{\sigma_\epsilon^2 + \sigma_\eta^2 (1 + \rho_t \sigma_\epsilon^2)^2}. \end{cases} \quad (7)$$

Slow learning

- □ When there is no observation noise, the precision of the public belief ρ_t increases by a **constant** value ρ_ϵ in each period, and it is a linear function of the number of observations.
- When there is an observation noise, equation (7) shows that as $\rho_t \rightarrow \infty$, the increments of the precision, $\rho_{t+1} - \rho_t$, converge to zero. The precision converges to infinity at a rate slower than a linear rate. The convergence of the variance σ_t^2 to 0 takes place at a rate slower than $1/t$.
The next result measures this rate¹.
- **Proposition (Vives, 1993)** In the Gaussian-quadratic model with an observation noise of variance σ_η^2 and private signals of variance σ_ϵ^2 , the variance of the public belief on θ , σ_t^2 , converges to zero as $t \rightarrow \infty$ and

$$\frac{\sigma_t^2}{\left(\frac{\sigma_\eta^2 \sigma_\epsilon^4}{3t}\right)^{\frac{1}{3}}} \rightarrow 1. \quad (8)$$

¹The analysis of a vanishingly small variance is simpler than that of a precision which tends to infinity.

Large number of agents (1)

- At the end of period t , agents observe the aggregate action Y_t which is the sum of the individuals' actions and of an aggregate noise η_t :

$$Y_t = X_t + \eta_t, \text{ with } X_t = \int x_t(i) di, \quad \text{and} \quad \eta_t \sim \mathcal{N}(0, 1/\rho_\eta).$$

- At the beginning of any period t , the public belief on θ is $\mathcal{N}(\mu_t, 1/\rho_t)$, and an agent with signal s_i chooses the action

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

- By the law of large numbers², $\int \epsilon_i di = 0$. Therefore, $\alpha_t \int s_i di = \alpha_t \omega$. The level of endogenous aggregate activity is

$$X_t = (1 - \alpha_t)\mu_t + \alpha_t \omega,$$

- and the observed aggregate action is

$$Y_t = (1 - \alpha_t)\mu_t + \alpha_t \theta + \eta_t. \tag{9}$$

Large number of agents (2)

- Using a normalization , this signal is informationally equivalent to

$$\frac{Y_t - (1 - \alpha_t)\mu_t}{\alpha_t} = \theta + \frac{\eta_t}{\alpha_t} = \theta + \left(1 + \frac{\rho_t}{\rho_\epsilon}\right)\eta_t. \quad (10)$$

- The observation noise has to be an aggregate noise. If the noises affected actions at the individual level, for example through individuals' characteristics, they would be “averaged out” by aggregation, and the law of large numbers would reveal perfectly the state of nature. An aggregate noise is a very plausible assumption in the gathering of aggregate data.
- This equation is similar to (6) in the model with one agent per period. (The variances of the noise terms in the two equations are asymptotically equivalent). Proposition ?? applies. The asymptotic evolutions of the public beliefs are the same in the two models.

Learning with a cost of information

- Each agent can purchase a signal s of precision q which is defined by costly signal

$$s = \omega + \epsilon, \text{ with } \epsilon \sim \mathcal{N}(0, 1/q).$$

- Cost of precision q is an increasing function, $c(q)$. (Example of n observations). No observation noise. The payoff function of each agent is quadratic: $U(x) = E[-(x - \omega)^2]$.

- Gain of the signal is the difference between the *ex ante* and *ex post* variances of the subjective distribution on ω :

$$V = \sigma_\omega^2 - \frac{\sigma_\omega^2}{\sigma_\omega^2 q + 1} = \frac{\sigma_\omega^4 q}{\sigma_\omega^2 q + 1}.$$

If there is an interior solution, the first order condition for q is $\frac{\sigma_\omega^4}{(\sigma_\omega^2 q + 1)^2} = c'(q)$.

Proposition

- Suppose that $c'(q)$ is continuous and $c(0) = 0$. If the marginal cost of precision $c'(q)$ is bounded away from 0, (for any $q \geq 0$, $c'(q) \geq \gamma > 0$), no agent purchases a signal after some finite period T and social learning stops in that period.
- Note that the case of a fixed cost of information with $c(0) > 0$ is trivial. Other cases with $c(0) = 0$ are left as exercises.

Policy (1)

- Social planner WF

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

The function W is interpreted as a loss function as long as ω is not revealed by a random exogenous process. In any period t , conditional on no previous revelation, ω is revealed perfectly with probability $1 - \pi \geq 0$. Assuming a discount factor $\delta < 1$, the value of β is $\beta = \pi\delta$. If the value of ω is revealed, there is no more loss.

- With no policy, the perception of an individual's action is hampered by the observation noise. A policy, which is based on public information, may amplify the individual response to his private signal and thus create a benefit for other agents. The policy is to pay agent t a bonus $(\gamma/2)(x_t - \mu_t)^2$. μ_t is publicly known. The policy amplifies the moves away from the expected value of ω in the public information.

-

$$x_t = \mu_t + \frac{\rho_\epsilon}{\rho_t + \rho_\epsilon} (1 + \beta)(s_t - \mu_t), \text{ with } \beta = \frac{\gamma}{1 - \gamma}.$$

Policy (2)

- Observation of others: $y_t = (1 + \beta)\alpha_t s_t + \eta_t$, with $\alpha_t = \frac{\rho_\epsilon}{\rho_t + \rho_\epsilon}$.
- The precision of that message is $\rho_y = (1 + \beta)^2 \alpha_t^2 \rho_\eta$.
- When β small, first-order impact on welfare of others, second-order on agent t .
- For γ small, the impact on W is strictly positive.