Private and Social Learning in Timing Games

Pauli Murto and Juuso Välimäki

Aalto University

Cowles March 31, 2010

Motivation

- Learning in Dynamic Decision Problems
    - Learning form own experience: Private
    - Learning by observing what others do: Observational or Social
    - Learning by communicating with others: Word-of Mouth, Cheap Talk
- In this presentation, focus on the first two modes

Motivation

- Decision Theory:
- Literature on Bayesian learning
- In general a hard problem. Solutions are available in special cases such as Multi-Armed Bandits and Optimal Stopping Problems.
- With many players:
- Herding and observational learning
- Strategic experimentation
- Today: Privately observed signals and publicly observed actions in a timing game with informational externalities
- Timing games are generalizations of optimal stopping problems

Questions

- Is information used efficiently?
- How are the inefficiencies realized?
  - Delays?
  - Herds on wrong actions?
  - Who gains who loses relative to private learning only?
- What are the qualitative features of the outcomes?
  - Does randomness persist in large games?
  - Can we obtain analytical solutions to the problem?

Model

- N players consider an irreversible investment opportunity
- The opportunities are (ex ante) identical
- Each investment is either a success or a failure
- The success probabilities are correlated:
  - Two states of the world ω, in H, fraction ρ^H of the investments fail, in L, fraction $\rho^L < \rho^H$ fail
  - Conditional on the state, success is i.i.d.
- Social learning is about state
- Private learning is about type of own investment (and indirectly about state)

Decision Problem

- Infinite (discrete) time horizon, period length Δ and discount factor $e^{-r\Delta}$
- In each period, decide whether to invest or not, $a_i^t \in \{1, 0\}$
- Investment decision irreversible
- Successful investment yields a (net) payoff stream $v\Delta > 0$ per period
- A failed investment gives the same revenue, but also a catastrophic failure of cost C with probability $\lambda\Delta$ per period
- Assume that $\lambda C > v$ so that expected profit from failed investment negative

Information

- Symmetric prior q^0 on $\{\omega = H\}$
- Prior on own type being failure

$$p^0 = q^0 \rho^H + (1 - q^0)\rho^L.$$

- Prior to investment can experiment (costlessly)
- Observe whether any failures have arrived
- Observe also what other players have done
- Other players' actions correlated with their signals, hence the state, hence player i's investment type

Timing Game

- Private history: own signals on failure
- If failure observed, dominant strategy to never invest.
- Hence concentrate on the uninformed, i.e. those who have seen no failures.
- Public history in period t:
- Strategy gives an investment probability for each (public) history:

$$\sigma_i^t : H^t \to [0, 1].$$

- In a Perfect Bayesian Equilibrium, each i plays a best response to the strategies of other players after all histories.

Related Literature

- Herding literature
  - Bikchandhani, Hirschleifer & Welch (1992)
  - Banerjee (1992)
  - Smith and Sorensen (2000)
- Observational learning with endogenous timing:
  - Chamley and Gale (1994)
  - Rosenberg, Solan and Vieille (2007)

Analysis: Isolated Player

- Start with an isolated investor acting only on private signals
- If no failures observed, then posterior on a failing project falls.

$$p_t = \frac{(1-\lambda\Delta)p_{t-1}}{(1-\lambda\Delta)p_{t-1} + (1-p_{t-1})}.$$

- Optimal policy: Invest as soon as p_t falls below p^*.
- Denote by p_t^ω posterior on investment being failure conditional on state ω and being uninformed in t
- Denote isolated value (decreasing, convex) function by $V_m(p)$
- What does the fact of not having seen a failure tell about state of world?
- π_t^ω probability of being uninformed conditional on state ω

Analysis: Learning From Others

- In each period, history h^t observed
- Each player computes $\Pr\{h^t \mid \omega\}$ using knowledge of strategy profile σ
- Compute belief $q_i(h^t)$ on state being H by combining information in $\Pr\{h^t \mid \omega\}$ and π_t^ω in Bayes' rule.
- Compute belief on own type by

$$p_i(h^t) = q_i(h^t)p_t^H + (1 - q_i(h^t))p_t^L.$$

- What matters for learning from others is the difference in investment probabilities across states

$$X_j^\omega(h^t) = \sigma_j(h^t)\pi_j^\omega(h^t).$$

- Information on state increases (in the sense of Blackwell) in σ

Symmetric Equilibrium

Lemma

In symmetric equilibrium, $V(h^t) = V_m(p(h^t))$*.*

- This is really the key lemma
- No symmetric pure strategy equilibria

Theorem

The game has a unique symmetric equilibrium.

Asymmetric equilibria are closer in spirit to the models of exogenous timing

Properties of the Symmetric Equilibrium

- Under full information on state, denote by τ^ω the optimal time to invest.
- An immediate consequence of the lemma above is that investment is delayed if $\omega = L$.
- For large games, we have approximate efficiency for state H.

Theorem

For all $\epsilon > 0$ and $\tau < \tau^H$, there are $K < \infty$ and Δ' such that the probability that more than K players invest before τ is less than ϵ if period length is less than Δ'.

- Notice that K is independent of the number of players N so almost all players invest at the optimal time.

Further Properties for Large N

- If there was no investment in t then probability of investment in $t + 1$ is proportional to Δ.
- Well defined hazard rate in the limit for first investments.
- If there was investment in t, then there is a discrete probability of investment in $t + 1$
- We call sequences of consecutive investments investment waves
- Conclusion: For small Δ, investments take place in randomly occurring bursts.
- Information is aggregated quickly during these waves.

Taking Stock

- What is important for the model to work?
  - Irreversibility, i.e. that we have a game of timing
  - Bounded strength of private learning
- What can be generalized?
  - The exact form of information revelation (e.g. different positive Poisson rates)
  - Many states
  - Distribution of privately known costs from failures (to avoid mixed strategies)
  - Exact form of game payoffs (could be changing over time)

How Do We Know?

- Murto and Välimäki (2009): "Delay and Information Aggregation in Stopping Games with Private Information"
  - Optimal investment game
  - Signals given at the beginning
  - Optimal time to invest depends on state
  - Affiliated signals model and supermodularity
  - Hence marginal incentives to invest change over time
  - This plays the same role as flow of private information in today's paper
  - Results shave same flavor: investment takes place in random bursts
  - Efficient investment in state H (in today's model) corresponds to a result of no early investment

Conclusions and Further Work

- Information aggregates in randomly occurring quick bursts
- Number of bad investments is decreased as a result of social learning at the cost of delaying good ones
- The 'market' cannot learn that it acted too early
- What next?
- Generalize to include information flows during the game and time dependent payoffs
- Expand the scope by considering payoffs that do not depend on time directly, but on a state variable that changes over time
- This would allow us to deal with real options in a model with private learning