Social Learning with Endogenous Timing

Daniel N. Hauser, Pauli Murto, Juuso Välimäki

Aalto University

Social Learning

We often learn from observing decisions of others.

Timing is an important part of this learning

- Should I adopt a new product early?
- Or wait to see if it's just a fad?
- How much I learn if I wait and see depends on the choices of others
- If you can choose your timing freely, information from choices determined in equilibrium

How much does allowing for free timing of actions aid or hurt social learning?

Sequential Observational Learning

Literature on sequential observational learning highlights failures of information aggregation.

- Boundedly informative signals + coarse action = Inefficient herds.
- Initial mistakes persist, no one agent has enough information to overturn them.
- Information aggregation can be facilitated by:
  - Unbounded Signals (Smith and Sorensen 2000)
  - Refined action space (Lee 1993, Ali 2018)
  - Censored history (Acemoglu et al. 2008)

Literature has understated the effectiveness of simply allowing for endogenous timing.

Endogenous Timing

We take the standard social learning game w/ boundedly informative signals and allow agents to freely choose when to move.

- There's an unknown state of the world.
- At the start of the game, each agent receives a single private signal.
- Agents choose when to make an irreversible decision (try to guess the state).
- There's a pure informational externality, no way for players to internalize the benefit to other players from revealing their private information.

Endogenous Timing

A player can now benefit from waiting to see others' act. Both the decision itself and its timing are informative. Is this enough to avoid herds?

We describe two very different equilibria

- The rush equilibrium: Players herd on the action taken by the first mover.
  - Fast, imprecise decisions
- An informative equilibrium
  - Players eventually learn the state, but with potentially significant delays.

Endogenous Timing

In the informative equilibrium, the free timing of actions can be used to prevent herding.

- In equilibrium, there are long periods of inactivity followed by bursts of action.
- With positive probability, an action choice is followed immediately by many more action choices.
- But in real time, these waves happen almost instantaneously if decisions can be taken frequently
- These exit waves end either with:
  - Almost all players choose almost simultaneously and adopt an action that is very likely to be correct.
  - Many players remain waiting, and the remaining players are very uncertain about the state.

The inefficiency in this equilibrium manifests as delays, instead of herds.

Continuous-Time Idealization

- Agents have boundedly informative private signals on binary state $\omega \in \{L, R\}$.
- They choose a time $t \in [0, \infty)$ to pick an action $a \in \{l, r\}$.
- Payoff is higher when the action matches the state and time is discounted.
- The state is perfectly revealed by a signal s_ω arriving with hazard rate λ_ω in state ω.
- Can we find an *'equilibrium outcome'* with two properties:
  - As long as the state is uncertain, agents with the most favorable signal for state ω are indifferent between choosing the action and delaying their action.
  - All other agents delay until the state is revealed.

Continuous-Time Idealization

Proposition

Such an outcome is possible only if $\lambda_R = \lambda_L$ and posterior belief p_t on state R based on public information up to t is fixed at p^* until the state is revealed.

- It is straightforward to compute the required hazard rates λ_ω (using martingale property of beliefs and Bellman Equation) for given payoffs, signal precision, and discounting.
- If $p_0 \neq p^*$, then an initial randomization revealing the state with positive probability and resulting in $p_{t+} = p^*$ with positive probability is needed.
- Payoff is higher when the action matches the state.

Continuous-Time Idealization

Why is this not the end of the story?

- We have not specified how the informative signals relate to the strategies of the players, i.e. we do not have a game yet.
- Defining the game in continuous time directly is problematic (more on this later).
- Our approach: Look for limits of discrete-time games that satisfy the above two properties as the number of players gets large.

Related Literature

Observational learning

- Banerjee (1992), Bikhchandani,Hirshleifer and Welch (1992), Smith and Sorensen (2000),...

Endogenous timing with action space Invest/Delay.

- Chamley and Gale (1994), Caplin and Leahy (1994), Chamley (2004), Levin and Peck (2008), Murto and Välimäki (2011, 2013)...

Endogenous timing with A/B/Delay

- Zhang (1997), Aoyagi (1998), Rogers (2005), Gul and Lundholm (1995)

Model

- Unknown state of the world $\omega \in \{L, R\}$, each equally likely.
- N players, $i \in \{1, 2, \ldots, N\}$
- Each player has a private signal $\theta_i \in \Theta$, drawn from F_ω
  - Θ finite.
  - Signals are ordered by likelihood ratio, normalize so $\theta = Pr(\theta|R)/Pr(\theta|L)$
  - Signals are informative ($\theta \neq 1$) but not perfectly ($\theta \neq 0$ or ∞).
  - For this talk, environment is symmetric.
- Discrete time $t \in \{0, \Delta, 2\Delta, 3\Delta, ...\}$.
  - Most interested in the case: $\Delta \to 0$.

Model

- At each t players either wait or choose an action $a_i^t \in \{\ell, r\}$.
- Stage payoff 0 if wait
- If a_t^i chosen, game ends at payoff $e^{-rt}u(\omega, a_t^i)$.
- For notational simplicity, let $u(\omega, a) = 1$ if action matches state zero otherwise.
- All past actions publicly observable: $h^t = (\mathbf{a}_0, ..., \mathbf{a}_{t-1})$ is the public history.
- Public belief: $p_t = Pr(\omega = R|h^t)$.
- We characterize different symmetric Public Perfect Equilibria of this game

Individual Decision Problem

Given public belief p_t, a player's myopically optimal action is ℓ if

$$\frac{p_t}{1 - p_t}\theta_i < 1.$$

and r with reverse inequality.

- Divide types into L and R types.
  - $\Theta_L := \{\theta : \theta < 1\}$
  - $\Theta_R := \{\theta : \theta > 1\}$

Dynamics

Lemma (Skimming Property)

In any symmetric equilibrium if type θ weakly prefers stopping and playing $\ell(r)$, then all lower (higher) type strictly prefer stopping and taking the same action to continuing.

This follows from strict single crossing of type and stopping time (in pointwise order on stopping times).

Dynamics: Terminology for $\Delta \to 0$

- Flow randomization: At time t, exits occur with probability proportional to Δ.
- Waves: At time t, exits occur with discrete positive probability. Wave ends either with all players exiting or return to Flow randomization.
- Rush: All types exit and play the same action: actions reveal no information.

Equilibrium

There must be inefficiencies in this setting.

- Information is revealed only through exits
- If players believe lots of information is revealed quickly, then players want to wait.
- If all wait, no information is revealed and it is better to exit.
- Information must either be:
  - Revealed slowly.
  - Very imprecise.

Rush Equilibrium

One possibility, not much information is transmitted in the game.

- If p_t is such that all types have the same myopically optimal action, it is sequentially rational to rush.
- An equilibrium for large enough N
  - Most extreme type exits first (w/ delays)
  - All types then immediately mimic first exiter's action.
  - Equilibrium features fast, but inefficient adoption

Informative Equilibrium

Instead of reducing the amount of information conveyed in equilibrium, it could be transmitted slowly.

To build some intuition, consider as a thought experiment the game both action and information are publicly revealed when exiting.

- Clearly, the rush equilibrium is no longer an equilibrium.
- Rushes are impossible, unless all types have dominant action.
- An equilibrium in this game must have delays.
- Is it possible to mimic this equilibrium if only actions are revealed?

Equilibrium

A natural equilibrium to look for in this game

- Alternation between Flow and Waves:
- Flow randomization until the next exit
  - Extreme types exit slowly.
  - Intermediate types wait.
- In Waves: Extreme types exit very rapidly in (real time)
- Waves end:
  - Either with all players exiting and choosing the same action
  - or with a return to Flow mode

In such an equilibrium, if N is large and Δ is small, actions reveal information almost perfectly

Existence: Why not Continuous Time Directly?

Unfortunately, existence is an issue in this game. Suppose players choose optimally from a continuum.

- Let $N = 2$, and $p_t = 1/3$ (so $\frac{p_t}{1-p_t} = 1/2$).
- Two types, one with likelihood ratio $\theta_1 = 1/3$, the other with likelihood ratio $\theta_2 = 3$
- If a player has type θ_1 they have a dominant action to exit and play ℓ immediately.

Existence: Why not Continuous Time Directly?

Unfortunately, existence is an issue in this game. Suppose players choose optimally from a continuum.

- Let $N = 2$, and $p_t = 1/3$ (so $\frac{p_t}{1-p_t} = 1/2$).
- Two types, one with likelihood ratio $\theta_1 = 1/3$, the other with likelihood ratio $\theta_2 = 3$
- If a player has type θ_1 they have a dominant action to exit and play ℓ immediately.
- If a player has type θ_2 their optimal action is l iff the other player's type is θ_1.

Existence: Why not Continuous Time Directly?

Unfortunately, existence is an issue in this game. Suppose players choose optimally from a continuum.

- Let $N = 2$, and $p_t = 1/3$ (so $\frac{p_t}{1-p_t} = 1/2$).
- Two types, one with likelihood ratio $\theta_1 = 1/3$, the other with likelihood ratio $\theta_2 = 3$
- If a player has type θ_1 they have a dominant action to exit and play ℓ immediately.
- If a player has type θ_2 their optimal action is l iff the other player's type is θ_1.
- So a player with type θ_2 wants to react immediately after seeing the opponent's choice,
- How do you formulate the best replies then?

How to Get Existence

- Tweak the extensive form and construct an equilibrium by backward induction on exits and beliefs on remaining signals:
  - Can be done with binary signals.
  - With richer signals, backward induction with continuous timing yields new problems.
  - Not sufficient monotonicity in the model at the end game (i.e. small N).

Existence

Here we take a different approach: Discrete time, with period length $\Delta = 1/N$.

With finite signal space, Nash's theorem guarantees the existence of symmetric equilibria for finite horizon T.

With discounting, the limit to $T \to \infty$ not problematic.

We call the limit process tracking the number of exits up to $\tau = t\Delta$ the Poisson process of equilibrium exits.

Can we find an equilibrium outcome without rushes?

Existence

Theorem

There exist a sequence of symmetric equilibria in the discrete time game where information aggregates in the limit as $N \to \infty$ in the sense that:

1. *Almost all agents choose the same action (with arbitrarily high probability).*
2. *The action is optimal given the state.*

Existence

To find such an equilibrium, we define an auxiliary game:

- Recall that Θ_ω is the set of signals where privately optimal actions match state ω.
- Fix n and modify the game so that before at least n players have exited, types in Θ_L can't play r and similarly for Θ_R.
- This game is symmetric and thus has a symmetric equilibrium, by standard Nash.

The Restricted Game

For large enough N (relative to n), equilibrium of restricted game are also equilibria of the unrestricted game.

Sketch

- Skimming property holds for the restricted game.
- If the restriction binds (say against choosing ℓ), then by skimming property all types in Θ_L have a unique BR to choose ℓ.
- By waiting for a period, the agent either has the restriction lifted or learns against choosing ℓ.
- Losses are thus bounded by the cost of waiting (vanishing as $\Delta \to 0$).
- For large N, waiting gives a lot of information if public belief is not too close to boundary already.
- For fixed n, public belief must lie in $(\delta_n, 1 - \delta_n)$ for some $\delta_n > 0$.

As $n \to \infty$, the n^{th} action choice reveals the state with arbitrarily high accuracy.

Continuous time limit in the symmetric case

We construct a sequence of equilibrium that converge to a continuous time limit where information aggregates.

In this limiting equilibrium:

- Exits occur at a Poisson rate ρ.
- Whenever an exit occurs, it triggers a wave that either
  - Pushes beliefs to certainty.
  - Returns beliefs to a unique interior belief p^* where both extreme types are indifferent between exiting and waiting.
- From eqm construction possible limiting belief $\lim_{t\to\infty} p_t \notin (0,1)$.
- Moreover, $\lim_{t\to\infty} p_t = 1_{\{\omega=R\}}$ with probability 1

In this limit, $p_t \in \{0, 1, p^*\}$ at any $t > 0$.

Martingales

Fix belief p_t and consider p_{t+dt}. Let $\gamma = \max_{\theta \in \Theta} Pr(R|\theta)$.

From construction, if $p_t < p^*$, p_{t+dt} is either 0 or p^*. Moreover

$$E(p_{t+dt}|p_t) = \pi \cdot 0 + (1-\pi)p^* = p_t$$

where π is the (unconditional) probability of learning the state. By symmetry $p^* = 1/2$ and

$$\pi_t = 1 - 2p_t.$$

Rate of Learning

Using this, we calculate the equilibrium value for the most extreme R type.

$$V_R(p) = \begin{cases} \frac{\gamma p}{\gamma p + (1-\gamma)(1-p)} \text{ for } p \geq p^* \\ \frac{\gamma p + (1-\gamma)(1-2p)}{\gamma p + (1-\gamma)(1-p)} \text{ for } p < p^* \end{cases}$$

Bellman equation at $p = 1/2$ gives us

$$rV_R(1/2) = \rho \underbrace{2\gamma(1-\gamma)}_{\text{probability of L exit}} (V_R(1-\gamma) - V_R(1/2)).$$

So exit rate is

$$\rho = \frac{r\gamma}{(2\gamma - 1)(1-\gamma)^2}.$$

Moreover, it follows from martingale property state is revealed at Poisson rate

$$\Gamma = \rho(2\gamma - 1).$$

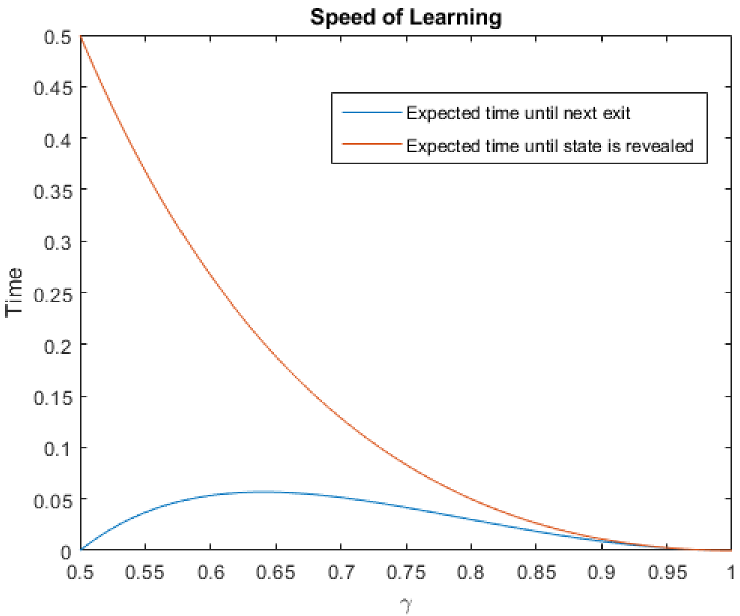

Speed of Learning
Expected time until next exit
Expected time until state is revealed
Time
0
0.05
0.1
0.15
0.2
0.25
0.3
0.35
0.4
0.45
0.5
0.5
0.55
0.6
0.65
0.7
0.75
0.8
0.85
0.9
0.95
1
γ

Conclusion

We've constructed an equilibrium where information aggregates.

- Our characterization of rates immediately identifies payoffs for interior types.
- Players gain by being able to wait for more confident types.
- But, learning is slow.
- Extreme types do not gain.

Conclusion

We relax the sequential structure of the canonical social learning.

- There still is a fundamental inefficiency due to the externality. Manifests as:
  - Herding
  - Delay
- Purely driven by the informational externality – not coarseness of actions – unlike with exogenous timing
- Next steps: exploring the correspondence of games where actions **and** signals are observed to models with pure social learning.