AI in Information Cascades

Yiqun Wang

May 28, 2026

How Herding Forms

Financial analyst arrive sequentially and decide whether to recommend buying a stock. Each sees earlier actions and her own private signal of the stock's quality.

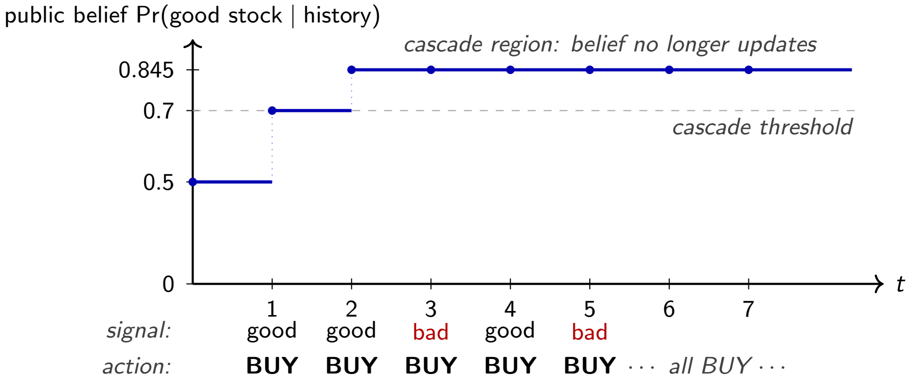

Motivation

A structural shift. As AI becomes embedded in decision-making, a growing share of agents condition on the *same* signal rather than on independent private evidence.

Trade-off.

- *Echo-herding effect:* repeated reliance on a common signal makes consecutive identical actions more likely.
- *Echo-discount effect:* rational later agents recognize that such consecutive actions may reflect a recycled common signal rather than fresh independent evidence, and discount them accordingly.

This paper. Study which force dominates, and what public histories really mean under shared AI use.

Preview of findings   Related literature

Model Setup

States and signals.

- Payoff state $\theta \in \{H, L\}$, $\Pr(\theta = H) = \frac{1}{2}$.
- Common signal $s_A \in \{h, \ell\}$ with $\Pr(s_A = h \mid \theta = H) = \Pr(s_A = \ell \mid \theta = L) = q_A$.
- At each date t: one agent arrives with type

$$m_t = \begin{cases} A & \text{w.p. } \lambda \\ I & \text{w.p. } 1 - \lambda \end{cases}$$

  - A-type ("AI user"): observes s_A.
  - I-type ("independent"): observes a fresh $s_t \in \{h, \ell\}$ with precision q_I.
- Agent picks $a_t \in \{H, L\}$ to maximize $\Pr(a_t = \theta)$.
- Types and signals are private info. Public observes only the action history a^{t-1}.

Three primitives: q_A, q_I, $\lambda \in (0, 1)$, with $q_A, q_I \in (\frac{1}{2}, 1)$.

Public Belief and Decision Rules

Hidden state: $\omega = (\theta, s_A) \in \{H, L\} \times \{h, \ell\}$, a 4-state space.

- Joint public posterior: $\mu_t(\theta, s_A) := \Pr\big(\omega = (\theta, s_A) \mid a^{t-1}\big)$
- Marginal belief on θ: $p_t := \mu_t(H, h) + \mu_t(H, \ell) = \Pr(\theta = H \mid a^{t-1}\big)$.

I-**type**. Acts on p_t and her own signal s_t:

- if $p_t \in [\,1 - q_I,\ q_I\,]$: follows her own signal s_t (*separating*);

A-**type**. Acts on the slice posterior:

$$d_t^A(h) = H \iff \mu_t(H, h) \geq \mu_t(L, h), \quad d_t^A(\ell) = L \iff \mu_t(L, \ell) \geq \mu_t(H, \ell).$$

A-types read histories better; *I*-types hold more history-independent information.

Asymptotic Behavior

Two key notions.

- *Herding:* $\exists\, T < \infty$ and $a^* \in \{H, L\}$ such that $a_t = a^*$ for all $t \geq T$.
- *Cascade:* $\exists\, T < \infty$ such that, for every $t \geq T$, the optimal action is constant in the agent's private signal.
- Cascade $\Rightarrow$ herding

Result. Almost surely, the I-type eventually enters a cascade and the A-type eventually herds on the cascading action.

$\Rightarrow$ a well-defined *eventual incorrect-cascade probability*

$$W_I(q_A, q_I, \lambda) \;:=\; \Pr\big(a^* \neq \theta\big).$$

Regime 1 (weak AI, $q_A < q_I$)

Suppose $a_1 = H$.

(a) A-type herds on history

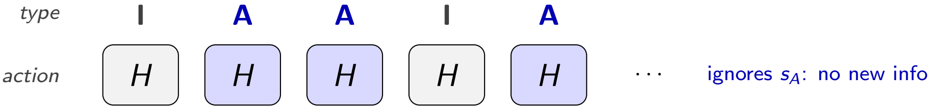

(b) A-type echoes a shared s_A

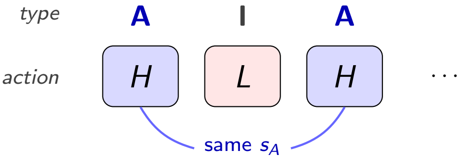

A-type herds before I-type and s_A is never revealed.

Regime 1: A Lower Bound on W_I

Proposition (lower bounds for W_I)

Lower bound $\Phi(\lambda, q_I)$ exists for W_I under the signal tie and random tie breaking rule and $\Phi(\lambda, q_I)$ is increasing in λ.

Comparison with BHW.

- *Signal tie breaking:* $\Phi > W^{\text{sig}}_{\text{BHW}}(q_I)$ for *every* $\lambda > 0$. The echo-herding effect dominates.
- *Random tie-breaking, narrow adoption (λ small):* $W_I < W^{\text{rand}}_{\text{BHW}}(q_I)$. The echo-discount effect dominates: a small share of A-types delays the I-cascade and gives later I-types more opportunities to use their own signals (Wu, 2015).
- *Random tie breaking, wide adoption ($\lambda > \frac{1}{2}$):* $\Phi > W^{\text{rand}}_{\text{BHW}}(q_I)$. The echo-herding effect dominates.

Regime 2 (Precise AI, $q_A > q_I$): Reveal the AI signal

Suppose $a_1 = H$. Effective precision $q_1 = \lambda q_A + (1-\lambda) q_I > q_I$, so every later I-type plays H regardless of her own signal.

(a) s_A agrees with $a_1 = H$

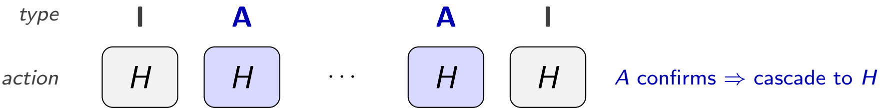

(b) s_A disagrees with $a_1 = H$

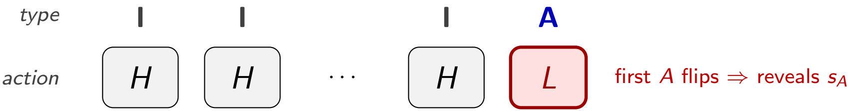

Departure from BHW. An A-type can disagree after an arbitrarily long run of Hs.

Regime 2: Strong Region ($q_A > \bar{q}(q_I)$)

When $\bar{q}(q_I) := q_I^2/[q_I^2 + (1 - q_I)^2]$ is exceeded, the revealed s_A alone is strong enough to push the public belief outside the I-separating interval. The cascade locks onto the direction of s_A.

Corollary

$$\boxed{W_I(q_A, q_I, \lambda) \;=\; 1 - q_A} \qquad \text{independent of } \lambda \text{ and } q_I.$$

- The incorrect-cascade probability equals the AI's own error rate.
- λ governs only the *speed* of revelation, not the eventual outcome.
- The wisdom of the crowd disappears.

Regime 2: Intermediate Region ($q_I < q_A \leq \bar{q}(q_I)$)

Suppose $\theta = H$. After s_A is revealed, public belief sits inside the I-separating interval, at a position that depends on whether $s_A = h$ or $s_A = \ell$.

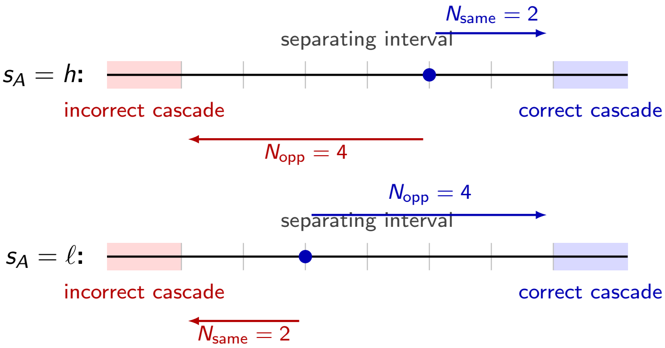

$N_{\text{same}}, N_{\text{opp}}$ = net number of subsequent actions in the same (opposite) direction as s_A needed to reach each cascade.

Regime 2: Discussion

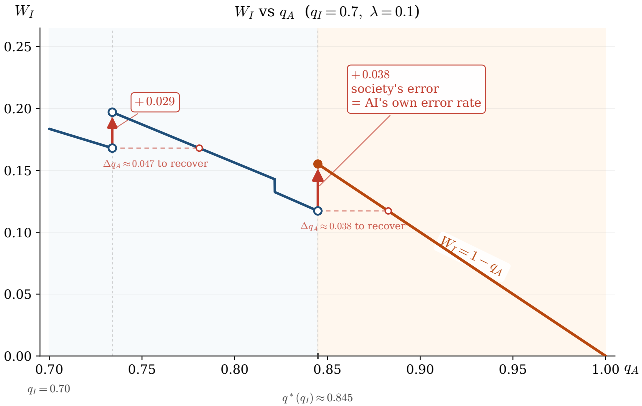

$q_I = 0.7$, $\lambda = 0.1$.

A more precise AI can deliver a higher incorrect-cascade probability.

Conclusion

Take-aways.

- **Weak AI:** AI signal is never disclosed and strictly raises incorrect-cascade probability when $\lambda > \frac{1}{2}$
- **Strong AI:** the incorrect-cascade probability equals the AI's own error rate.
- **Intermediate:** A more precise AI need not improve social learning.

Future directions.

- Allow *imperfect correlation* across AI users
- Endogenize AI adoption (λ as a choice).
- Multiple competing algorithms; platform / disclosure design.

Thank you. yiqun.wang.23@ucl.ac.uk

Preview of Findings

This paper builds on the BHW model (Bikhchandani, Hirshleifer, and Welch, 1992) by allowing some agents to use a common AI signal.

Two regimes by precision ordering.

- **Regime 1, AI less precise than independent signals:** the AI signal is *never* publicly revealed. Incorrect cascades become strictly more likely than in BHW if more than half of agents are using AI.
- **Regime 2, AI more precise than independent signals:** the AI signal is eventually revealed. Two sub-cases:
  - *Strong:* the incorrect-cascade probability equals the AI's own error rate, independent of AI adoption and independent signal precision.
  - *Intermediate:* A more precise AI can deliver a *higher* incorrect-cascade probability.

$\Rightarrow$ Regulating AI by precision alone is unlikely to be effective.

Back to overview

Related Literature

Social learning. Bikhchandani, Hirshleifer, and Welch (1992); Banerjee (1992); Smith and Sørensen (2000).

Heterogeneous-precision. Wu (2015): low-precision but *independent* laymen can *delay* premature cascades and improve learning.

Common-source models. Eyster and Rabin (2014); DeMarzo, Vayanos, and Zwiebel (2003).

Algorithmic monoculture. Kleinberg and Raghavan (2021); Bommasani et al. (2022).

Back to overview