# *Pathological Outcomes of Observational Learning*

Lones Smith
Massachusetts Institute of Technology

Peter Sørensen
Nuffield College and M.I.T.

1996

(original version July, 1994)

# 1 OVERVIEW

- Individuals sequentially choose an action based on private information, and observation of all predecessors' *actions* $\Longrightarrow$ not simple statistical learning
- pure informational externality; no economic externalities
- Banerjee (1992); BHW (1992)
- Two spins on their pathological learning outcome:
  1. *Belief Convergence*, or *Cascades*: Public history eventually becomes so informative that individuals disregard their private information $\Rightarrow$ public beliefs enter an absorbing state, possibly wrong one
  2. *Action Convergence*, or *Herds*: Eventually, all individuals will take the same action, possibly wrong one

- Generalization of the herding model
  1. General private *signal space:* With continuous signals, herds generically may exist without cascades
  2. Unbounded private signal strength: $\exists$ *complete learning* in belief and action space $\Rightarrow$ only a correct herd obtains, and herding pathology disappears!
  3. Addition of a little *noise*: This does away with the 'overturning principle' (that one single individual's contrary action has drastic effects)
  4. *Multiple preference types*: New pathology *confounded learning* arises, even if private signals have unbounded strength
  5. *Link to experimentation literature*: herding is an example of optimal *single agent* learning model

## 2 THE STANDARD MODEL

- Infinite sequence of individuals $1, 2, \ldots$ who act sequentially, in an exogenous order
- Two underlying states of the world, $H$ and $L$ (assume $H$)
- Private *conditionally i.i.d.* signals $\sigma_n$ (with no perfectly revealing signals) & $g(\sigma_n)$ = private $L/H$ odds
- Actions $a_1, \ldots, a_M$ with state dependent payoffs
- Individuals have identical preferences over outcomes
- They observe the full action history, and make an inference about other individuals' signals, updating their own posterior
- The observed history of the first $n-1$ actions leads to a *public belief* $q_n$ that state is $H$, and a *likelihood ratio* $\ell_n = (1-q_n)/q_n$

Figure 1: **The Individual Decision Problem: Frontier of Expected Payoffs and Posterior Thresholds.**

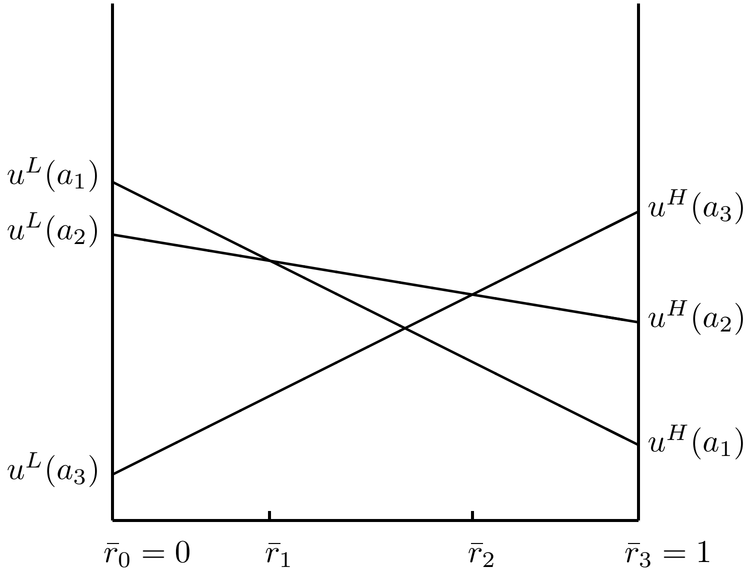

**Private Belief Distributions**

- if $H, L$ are WLOG ex ante equilikely, then individual $n$ has the interim *private belief* $p \equiv p(\sigma_n) = 1/(g(\sigma_n) + 1)$ that the state is $H$
- dist'n of private beliefs $p = p(\sigma)$ is $F^H$ or $F^L$

  *Q: What is the likelihood of L/H given my private beliefs?*

⋆ *No Introspection Condition*:

Any two c.d.f.'s can be rationalized iff $dF^L/dF^H = (1-p)/p$

eg. $F^H(p) = p^2$ and $F^L(p) = 2p - p^2$

$\Rightarrow F^H$ and $F^L$ have the same support, with $\text{co}(\text{supp}(F)) = [\underline{b}, \bar{b}]$ (*'Romeo and Juliet' effect*)

$\Rightarrow F^H \succ_{FSD} F^L$; note: $F^H(p) = F^L(p) \Leftrightarrow F^H(p) \in \{0, 1\}$

## Acting upon Private Beliefs

- given $\ell$ & $p$, posterior belief is $r = p/(p + \ell(1-p))$, by Bayes rule

  $\Rightarrow$ choose $a_m \Leftrightarrow p \in [\bar{p}_{m-1}(\ell), \bar{p}_m(\ell))$; private belief *thresholds* satisfy $\bar{p}'_m(\ell) > 0$ and $0 \equiv \bar{p}_0(\ell) \leq \bar{p}_1(\ell) \leq \ldots \leq \bar{p}_M(\ell) \equiv 1$

  $\Rightarrow$ one takes action $a_m$ with chance $\rho(m|s,\ell) \equiv F^s(\bar{p}_m(\ell)) - F^s(\bar{p}_{m-1}(\ell))$ in state $s \in \{H, L\}$

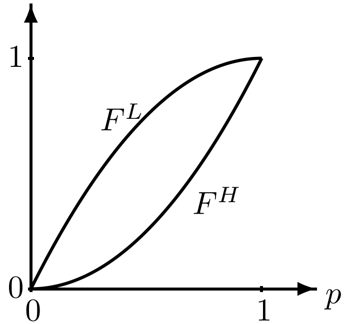

Figure 2: **Individual Black Box.** Individual $n$ bases his action decision $m_n$ on the public history ($\leftrightarrow$ likelihood ratio $\ell_n$) and on his private signal $\sigma_n$, implying a new continuation $\ell_{n+1}$.

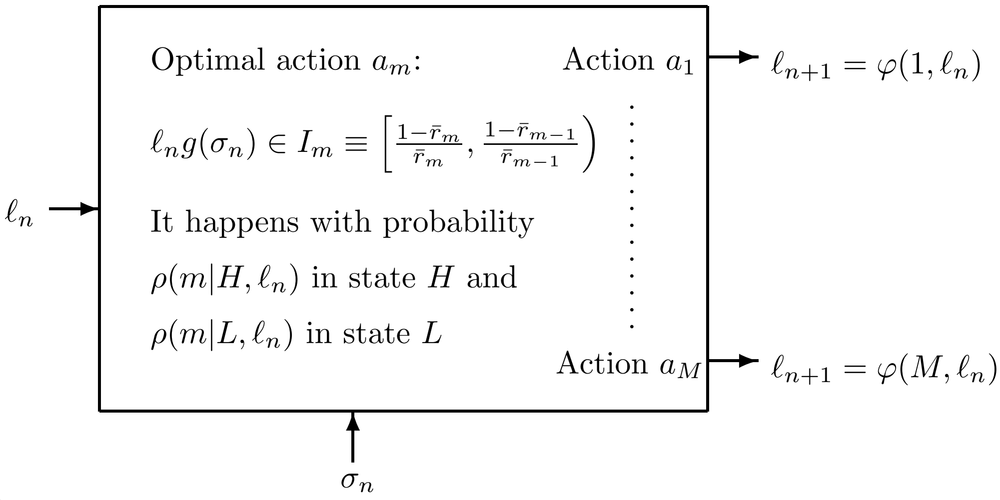

## Corporate Learning as a Martingale Process

- Through the individuals' private signals, their actions $\langle m_n \rangle$ are random, and so $\langle q_n \rangle$ and $\langle \ell_n \rangle$ are stochastic processes
- Individual $n$ takes action $a_{m_n}$ with chance $\rho(m_n|H, \ell_n)$ in state $H$

$$\Rightarrow \ell_{n+1} = \varphi(m_n, \ell_n) \equiv \ell_n \frac{\rho(m_n|L, \ell_n)}{\rho(m_n|H, \ell_n)} \qquad \text{(Bayes' Rule)}$$

- We focus on odds $\langle \ell_n \rangle$ rather than beliefs $\langle q_n \rangle$. Why?

  Because $\langle \ell_n \rangle$ is a *martingale conditional on state* $H$:

  $E[\ell_{n+1} \mid H, \ell_1, \ldots \ell_n] = \sum_m \rho(m|H, \ell_n)\ell_n \frac{\rho(m|L,\ell_n)}{\rho(m|H,\ell_n)} = \ell_n$
- Since $\ell_n \geq 0$ always, MCT applies

  $\Longrightarrow$ conditional on state $H$, $\langle \ell_n \rangle$ converges (a.s.) to the <u>random variable</u> limit $\ell_\infty = \lim_{n\to\infty} \ell_n$ with (<u>finite</u>) values in $[0, \infty)$.

## Corporate Learning as a Markov Process

- $(m_n, \ell_n)$ is a *Markov process* on $\{1, 2, \ldots, M\} \times [0, \infty)$

  $(m_n, \ell_n) \mapsto (m_{n+1}, \varphi(m_{n+1}, \ell_n))$ with chance $\rho(m_{n+1}|H, \ell_n)$

**Theorem B-1 (Stationarity)** *If $\rho$ and $\varphi$ are continuous in $\ell$, then any $\hat{\ell} \in \operatorname{supp}(\ell_\infty)$ satisfies $\forall m: \ \rho(m|H, \hat{\ell}) = 0 \vee \varphi(m, \hat{\ell}) = \hat{\ell}$*

- Intuition: At any $\hat{\ell} \in \operatorname{supp}(\ell_\infty)$, no further information can be gleaned from any action observation
- Special case: *Action absorbing basin* for action $a_m$ is $J_m = \{\ell \,|\, \rho(m|H, \ell) = 1\}$ (hence, $J_m = \{\ell \,|\, \rho(m|L, \ell) = 1\}$)

⋆ $\hat{\ell} = \infty$ is stationary, so can fully incorrect learning occur? No! MCT rules out $\ell_n \to \infty$:

**Basic Concepts**

- Private beliefs are

  1. *bounded* if the private signal has a bounded likelihood range; $g(\sigma)$ and $1/g(\sigma)$ are bounded above
  2. *unbounded* if the convex hull of the range of $g$ is $[0, \infty)$

- With bounded beliefs, there *must* exist action absorbing basins for the two extreme actions, $J_1$ and $J_M$, and there *may* exist absorbing basins for insurance actions
- With unbounded beliefs, action absorbing basins only exist for extreme actions: $J_1 = \{\infty\}$, $J_M = \{0\}$, with $J_2, \ldots, J_{M-1} = \varnothing$
- A *cascade* on action $a_m$ as of individual $n$ means that $\ell_n \in J_m$
- A *herd* on action $a_m$ as of individual $n$ means that all individuals $n, n+1, \ldots$ choose $a_m$ (logically weaker than cascade)

Figure 3: **Continuations & Absorbing Basins.** Bounded support beliefs $g(\sigma) = 1/2 + \sigma$ on $[0, 1]$; one insurance & 2 extreme actions. [Martingale property $\Rightarrow$ E(continuation likelihood) is on diagonal.]

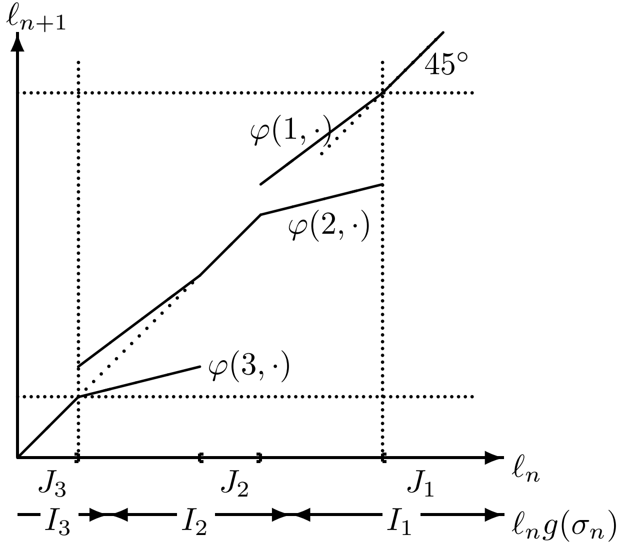

Figure 4: **Continuations & Absorbing Basins, Revisited.** Bounded support beliefs $g(\sigma) = 1/2 + \sigma$ on $[0, 1]$; no insurance actions (because preferences are different).

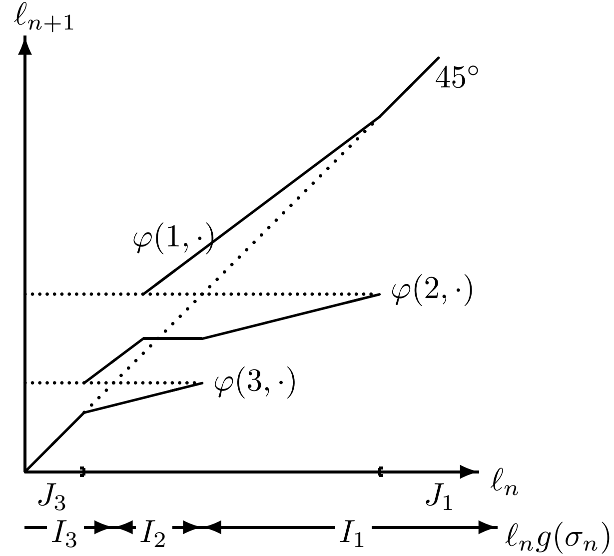

# 4 MAIN RESULTS

## Convergence of Beliefs

**Theorem 1 (Limit Cascades)** *With bounded beliefs,*
(1) $\ell_\infty \in J_1 \cup \cdots \cup J_M$ *almost surely*
(2) $\ell_0 \notin J_M \Longrightarrow \ell_\infty \in J_M$ *a.s. is impossible (state $H$)*

**Theorem 2 (Complete Learning)** *With unbounded beliefs, $\ell_n \to 0$ in state $H$, and $\ell_n \to \infty$ in state $L$.*

---

## Convergence of Actions

**Theorem 3 (Herds)** *With bounded beliefs, a herd on some action will almost surely arise in finite time. Unless there is a cascade on the most profitable action $a_M$ from the very outset, a herd can arise on an action other than $a_M$.*

**Theorem 4 (Correct Herds)** *With unbounded beliefs, eventually everyone takes the optimal action (almost surely).*

**Why Limit Cascades?**

- $\langle \ell_n \rangle$ is a martingale $\Longrightarrow \ell_\infty \equiv \lim_{n\to\infty} \ell_n$ exists, by MCT
- $\hat{\ell} \in \text{supp}(\ell_\infty)$

  $\Longrightarrow \rho(m|H,\hat{\ell}) = 0$ or $\rho(m|H,\hat{\ell}) = \rho(m|L,\hat{\ell})$, by stationarity

  $\Longrightarrow$ any $m$ with $\rho(m|H,\hat{\ell}) > 0$ satisfies $\rho(m|H,\hat{\ell}) = 1$, since beliefs are shifted towards state $H$ if state $H$ is true

**Why Incorrect Limit Cascades?**

- in state $H$, must rule out $\ell_\infty \in J_M$ almost surely
- If $\ell_\infty \in J_1$ with positive probability, we are done; else, $\ell_n \leq \inf J_1 < \infty$.

  $\Longrightarrow E[\ell_\infty] = \lim_{n\to\infty} E[\ell_n] = \ell_0$ by Lebesgue's Dominated Convergence Theorem
- so $\ell_0 \notin J_M = [0, \underline{\ell}]$ implies $\text{supp}(\ell_\infty) \subseteq J_M = [0, \underline{\ell}]$ is impossible

**Why Complete Learning?**

- With unbounded support, limit cascades can only arise on extreme actions $a_1$ and $a_M$ (as $J_2, \ldots, J_{M-1} = \varnothing$)

- $\rho(m|H, \hat{\ell}) \in \{0, 1\} \iff (m, \hat{\ell}) = (1, 0)$ or $(m, \hat{\ell}) = (M, \infty)$

  and martingale property of $\langle \ell_n \rangle \Rightarrow Pr(\ell_\infty = \infty) = 0$ in state $H$

**Why Herds?**

- idea: convergence in beliefs $\Longrightarrow$ convergence in actions
- Indeed, we only have limit cascades and not cascades

$\star$ *The Overturning Principle*

If agent $n$ optimally chooses action $a_m$, then, before observing his private signal, agent $n+1$ would optimally choose $a_m$ too

$\Rightarrow$ one contrary action will completely overturn the public belief ($\ell_{n+1}$ jumps far from $\ell_n$)

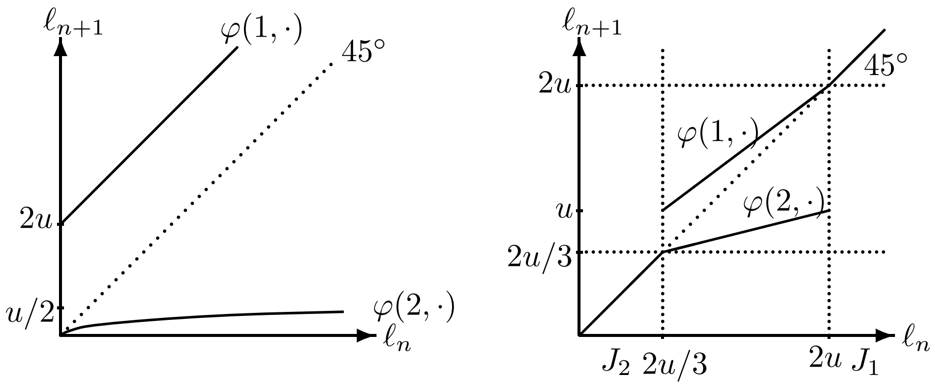


Figure 5: **Continuations.** Binary action examples with unbounded private beliefs (left), and bounded private beliefs (right)

- illustrates the Overturning Principle, and
- shows that a cascade need not arise with bounded beliefs, and
- hints why complete learning arises in unbounded case and not in the bounded case.

**Fast Learning in Belief Space**

- If $\exists$ cont's density $f^H$ of $F^H$ (and thus $f^L$ of $F^L$), then *extreme signals are rare* iff $f^H(\underline{b}) = 0$ or $f^L(\bar{b}) = 0$.
- $\ell_n$ *converges to* $\hat{\ell}$ *at rate* $\bar{\theta} \in [0,1]$ if $|\ell_n - \hat{\ell}| = O(\theta^n)$

**Lemma 9 (Exponential Convergence)** *Assume bounded beliefs and that extreme signals are not rare. In any limit cascade, if* $\hat{\ell} = \lim_{n\to\infty} \ell_n$ *then* $\ell_n$ *converges to* $\hat{\ell}$ *at some rate* $\theta < 1$.

*Proof Idea*: In a limit cascade and herd on action $a_1$, with $\ell_n \uparrow \hat{\ell} = \inf(J_1)$, $n$ chooses action $a_1 \Leftrightarrow n$'s posterior $< \bar{r}_1$ $\Leftrightarrow p(\sigma_n) < \bar{p}_1(\ell_n)$. Thus, with smooth private belief distributions,

$$\ell_{n+1} = \varphi(1, \ell_n) = \ell_n \frac{F^L(\bar{p}_1(\ell_n))}{F^H(\bar{p}_1(\ell_n))} \qquad \text{(Bayes' Rule)}$$

$$\Longrightarrow [\varphi_\ell(1, \hat{\ell}) = \theta < 1 \Leftrightarrow f^L(\bar{p}_1(\hat{\ell})) < f^H(\bar{p}_1(\hat{\ell}))]$$

$$\Longrightarrow \hat{\ell} - \ell_{n+1} = \hat{\ell} - \varphi(1, \ell_n) \doteq \varphi_\ell(1, \hat{\ell})(\hat{\ell} - \ell_n) = \theta(\hat{\ell} - \ell_n)$$

**Fast Learning in Action Space**

- Bounded beliefs: If learning is exponentially fast, then a herd arises in finite expected time, as every abortive herd ends fast:

- $e_n$ = *exit chance* from temporary herd vanishes exponentially fast, so *conditional* exit rates are boundedly positive

⋆ *The key to fast action convergence is how slowly error is discovered by contrarians.*

- Unbounded beliefs: extreme signals in favour of truth are *rare* if $F^L(p) = O(p^\alpha)$ and $1 - F^H(1-p) = O(p^\alpha)$, $\alpha \geq 1$, small $p$

⋆ Case 1: if extreme signals are rare, then ∃ (correct) herd in infinite mean time (the truth is learned, but it takes forever)

- classic example: $F^L(p) = 2p - p^2, F^H(p) = p^2$

⋆ Case 2: if extreme signals are not rare, so $F^L(p) = O(p^\alpha)$ and $1 - F^H(1-p) = O(p^\alpha)$, $\alpha < 1$, then mean time to herd $< \infty$

# 5 NOISE

- Introduce small amount of i.i.d. noise: eg. crazy/misperceived types, or trembling individuals
- this yields new transition chance $\psi(m|s,\ell)$, where

- Trembling: fraction $\kappa_j^m$ should take $a_j$ but take $a_m$
  $\psi(m|H,\ell) = [1-\kappa_m(\ell)]\rho(m|H,\ell) + \sum_{j\neq m} \kappa_j^m(\ell)\rho(j|H,\ell)$
- Craziness (special case): fraction $\kappa_m$ always takes action $a_m$
  $\psi(m|H,\ell) = \kappa_m + (1 - \sum_{j=1}^M \kappa_j)\rho(m|H,\ell)$

**Theorem 6 (Convergence in Beliefs)** *Let $\ell_n \to \ell_\infty$. With bounded beliefs,*

(1) $\ell_\infty \in J_1 \cup \cdots \cup J_M$ *almost surely;*

(2) $\ell_0 \notin J_M \Longrightarrow \ell_\infty \in J_M$ *a.s. is impossible (state H)*

*With unbounded beliefs, $\ell_\infty = 0$ almost surely (state H).*

**Why Complete Learning with Unbounded Beliefs?**

All $\psi$ are bounded away from zero, so we must investigate stationarity: $\varphi(m|H,\hat{\ell}) = \hat{\ell}$

$$\hat{\ell}\frac{\kappa_m + (1 - \sum_{m=1}^{M} \kappa_m)\rho(m|L,\hat{\ell})}{\kappa_m + (1 - \sum_{m=1}^{M} \kappa_m)\rho(m|H,\hat{\ell})} = \hat{\ell}$$

$\Longrightarrow \rho(m|H,\hat{\ell}) = \rho(m|L,\hat{\ell})$, which as before implies that they are zero or one

Figure 6: **Continuations.** The same basic two-action model, first without and then with craziness.

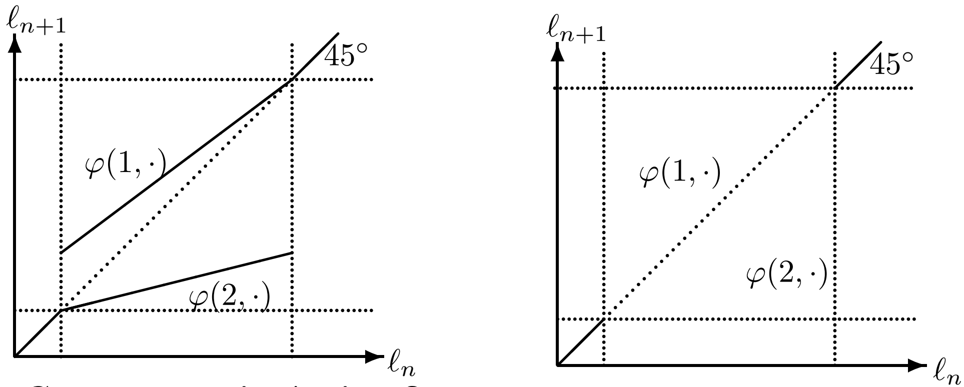


**Convergence in Actions?**

- With noise, the overturning principle fails, so noise-less proof fails: 'Contrary' actions have little impact on public beliefs (discounted as likely irrational actions)

⋆ *With bounded beliefs and non-rare extreme signals, 'rational herds' still arise (a.s.)*

- (first) Borel-Cantelli Lemma $\Longrightarrow$ an infinite string of rational 'herd violators' a.s. can't occur if $\sum_{n=1}^{\infty}(1-\rho(m|H,\ell_n)) < \infty$
- martingale property $\ell \equiv \sum_{m=1}^{M} \psi(m|H,\ell)\varphi(m,\ell)$ & AM-GM $\Rightarrow$

$$\begin{aligned} 1 &= \sum_{m=1}^{M} \psi(m|H,\hat{\ell})\varphi'(m,\hat{\ell}) + \sum_{m=1}^{M} \psi'(m|\hat{\ell})\varphi(m,\hat{\ell}) \\ &= \sum_{m=1}^{M} \psi(m|H,\hat{\ell})\varphi'(m,\hat{\ell}) > \prod_{m=1}^{M} |\varphi'(m,\hat{\ell})|^{\psi(m|H,\hat{\ell})} \equiv \theta \end{aligned}$$

  at a stationary point $\hat{\ell}$, where $\varphi(m,\hat{\ell}) = \hat{\ell}$ for all $m$

- appendix: $\theta < 1$ is the criterion for exponential stability of a stochastic difference equation, i.e. $|\ell_n - \hat{\ell}| \approx \theta^n$ if $\ell_n \to \hat{\ell}$

# 6 MULTIPLE INDIVIDUAL TYPES

- Assume $T$ types of individuals, spread i.i.d. in sequence, with state-dependent preferences (noise = special case)
- new transition probability: $\psi(m|H,\ell) = \sum_{t=1}^{T} \lambda^t \rho^t(m|H,\ell)$
- history is informative with distinct type frequencies $\lambda^1, \ldots, \lambda^T$
- At a *confounded learning point* $\ell^*$, no inference can be drawn from $\ell^*$ as each action occurs with equal chance in states $H$, $L$
  $\Rightarrow \psi(m|H,\ell^*) = \psi(m|L,\ell^*)$, so $\ell^*$ is a stationary point of $\langle \ell_n \rangle$

Figure 7: **Confounded Learning Point.** At $\ell^*$, no inference can be drawn about the true state of the world with two types $A$ and $B$. (Here, $a_2 \succ_A a_1$ and $a_1 \succ_B a_2$ in state $H$; the reverse in state $L$.)

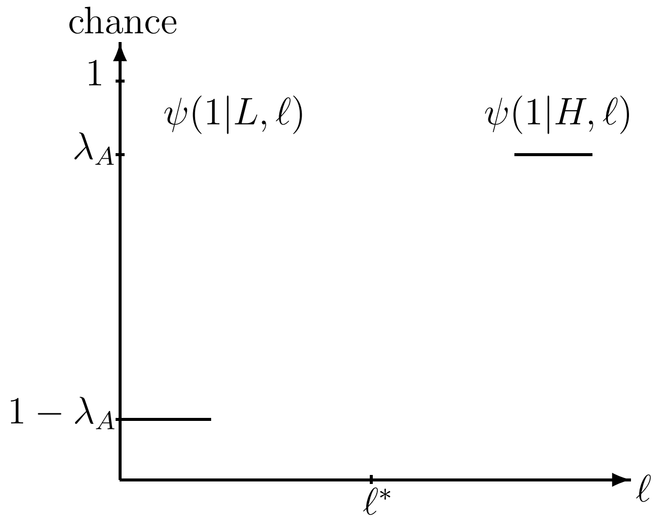

Figure 8: **Confounded Learning Point.** A fixpoint argument suggests the existence of a confounded learning point

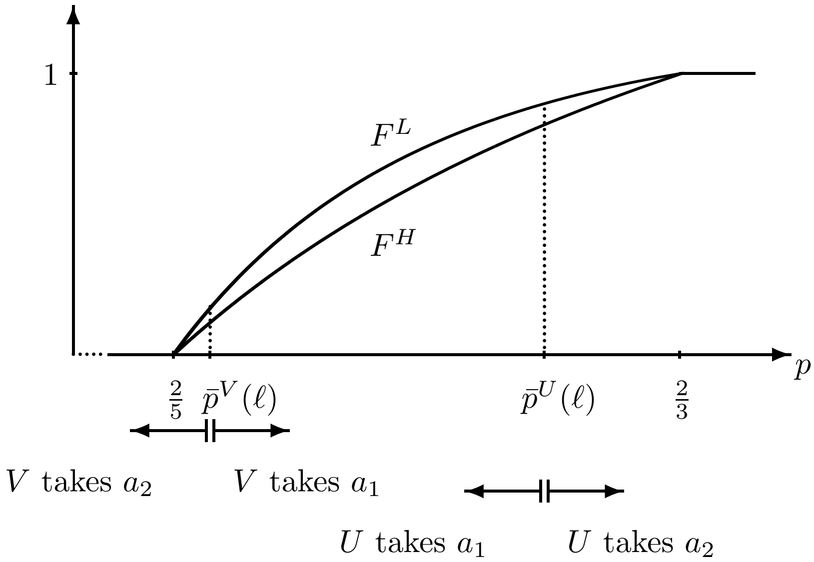

- Still, does *confounded learning* occur, i.e. $\ell_n \to \ell^*$ occur?
  Yes! Just use local stability criterion ($\star$).
- Even with unbounded beliefs, complete learning need no longer obtain: learning may die out, with $\ell_\infty$ unfocused!
- Private signals become totally decisive for individual actions, whereas in a cascade, private signals are ignored

# 7 LINK TO EXPERIMENTATION LITERATURE

- We can map the pathological outcomes of social learning into the standard outcomes of single person experimentation
- Incorrect herd $\leftrightarrow$ settle on suboptimal action, the learning process stops short of revealing the true state (eg. Rothschild (1974) and the two-armed bandit problem)
- Confounded learning $\leftrightarrow$ an outcome where statistics are still generated, but they are identically distributed in the two states

\- Similar to the learning problem in McLennan (1984)

A monopolist faces one of two possible demand curves; consumers arrive one per period, and buy with chances $q = a - bp$ or $q = A - Bp$

- Easley-Kiefer (1988) calls such actions *potentially confounding*, i.e. optimal for unfocused beliefs for *any* experiment realization

- But EK show that this generically doesn't exist for finite state and action spaces!

$\Rightarrow$ so how do we get herding and confounded learning?

$\Rightarrow$ Must write the herding model as a single person experimentation problem

**How to replace everyone with a single experimenter**

- new state space: $\Theta = \{H, L\}$
- new action space: the compact set of $n$ private belief thresholds $X = \{x \in [0,1]^M | 0 \leq x_1 \leq \ldots \leq x_M = 1\}$ *(NOT finite)*
- discount factor $= 0$
- new random expt outcome, or observable signal: old action chosen in herding model from $\{1, 2, \ldots, M\}$.
- Given the action $x$ chosen, the probability that signal $m$ occurs is $\rho(m|s, x) = F^s(x_m) - F^s(x_{m-1})$ in state $s$ without noise, and more generally $\psi(m|s, x)$ with noise.
- to simulate two types, let experimenter choose two sets of thresholds, and not observe which one determines the observed signal