working paper
department
of economics

PATHOLOGICAL OUTCOMES OF OBSERVATIONAL LEARNING

Lones Smith
Peter Sørensen

96-19 June 1996

massachusetts
institute of
technology

50 memorial drive
cambridge, mass. 02139

PATHOLOGICAL OUTCOMES OF OBSERVATIONAL LEARNING

Lones Smith
Peter Sørensen

96-19 June 1996

96-19

*Pathological Outcomes of Observational Learning**

Lones Smith†
Department of Economics, M.I.T.

Peter Sørensen‡
Nuffield College and M.I.T.

June 28, 1996
(original version: July, 1994)

Abstract

This paper systematically analyzes and enriches the observational learning paradigm of Banerjee (1992) and Bikhchandani, Hirshleifer, and Welch (1992). Our contributions fall into three categories.

First, we develop what we consider to be the 'right' analytic framework for informational herding (convergence of actions and convergence of beliefs, using a Markov-martingale process). We demonstrate its power and simplicity in four major ways: (1) We decouple herds and cascades: Cascades might never arise, even though herds must. (2) We show that wrong herds can arise iff the private signals have uniformly bounded strength. (3) We determine when the mean time to start a herd is finite, and show that (absent revealing signals) it is *infinite* when a correct herd is inevitable. (4) We prove that long-run learning is unaffected by background 'noise' from crazy/trembling decisions.

Second, we explore a new and economically compelling model with multiple types, and discover that a 'twin' observational pathology generically appears: *confounded learning.* It may well be impossible to draw any further inference from history even while it continues to accumulate privately-informed decisions!

Third, we show how the martingale property of *likelihood ratios* is neatly linked with the stochastic stability of the learning dynamics. This not only allows us to analyze herding with noise, and convergence to our new confounding outcome, but also shows promise for optimal experimentation.

---

*From its initial presentation at the MIT theory lunch, this paper has evolved and been much improved and polished by feedback from seminars at Yale, IDEI Toulouse, CEPREMAP, IEA Barcelona, Pompeu Fabra, Copenhagen, IIES Stockholm, Stockholm School of Economics, Nuffield College, LSE, CentER, Free University of Brussels, Western Ontario, Brown, Pittsburgh, Berkeley, UC San Diego, UCLA, University College of London, Edinburgh, and Cambridge. We wish also to specifically acknowledge Abhijit Banerjee, Kong-Pin Chen, Glenn Ellison, John Geanakoplos, David Hirshleifer, David Levine, Preston McAfee, Thomas Piketty, Ben Polak, and Jörgen Weibull for useful comments, and Chris Avery for a key role at a very early stage in this project. The usenet group `sci.math.research` has also proved invaluable. All errors remain our responsibility. Smith and Sørensen respectively acknowledge financial support for this work from the National Science Foundation and the Danish Social Sciences Research Council.

†e-mail address: `lones@lones.mit.edu`

‡e-mail address: `SORENSEN@cole.nuff.ox.ac.uk`

Contents

1. INTRODUCTION

Suppose that a countable number of individuals each must make a once-in-a-lifetime binary decision — encumbered solely by uncertainty about the state of the world. Assume that preferences are identical, and that there are no congestion effects or network externalities from acting alike. Then in a world of complete and symmetric information, all would ideally wish to make the same decision.

But life is more complicated than that. Assume instead that the individuals must decide sequentially, all in some preordained order. Suppose that each may condition his decision both on his (endowed) private signal about the state of the world and on all his predecessors' decisions, but *not* their private signals. The above simple framework was independently introduced in Banerjee (1992) and Bikhchandani, Hirshleifer, and Welch (1992) (hereafter, simply BHW). Their perhaps surprising common conclusion was that with positive probability a 'incorrect herd' would arise: Despite the surfeit of available information, after some point, everyone can make the identical less profitable decision.

In this paper, we systematically analyze and enrich the informational herding paradigm, as it so happens to offer both simple lessons and deeper insights into rational learning. Our embellishment upon the herding story is best motivated by means of the following counterfactual. Assume that we are in a potential herd in which one million consecutive individuals have followed suit on some action, but suppose that the very next individual deviates. What then could Mr. one million and two conclude? First, he could decide that his predecessor had a more powerful signal than everyone else. To capture this, we shall generalize the private information beyond discrete signals, and admit the possibility that there is no uniformly most powerful signal. Second, he might opine that the action was irrational or an accident. We shall thus add noise to the herding model. Third, he possibly might decide that different preferences provoked the contrary choice. On this score, we shall consider the herding model with multiple types. Here, we find that herding is not the only possible 'pathological' outcome: We may well converge to a situation where history offers no decisive lessons for anyone, and everyone must forever rely on his private signal!

Here is an overview of the paper, and how we view our contributions.

1. Proper Analytic Framework; Nonrobustness of Past Results; Extensions.

• THE 'RIGHT' THEORY. Our analysis is focused through the two lenses of *convergence of beliefs* (learning) and *convergence of actions* (herding). As such, two stochastic processes constitute the building blocks for our theory: (*i*) the public likelihood ratio is a martingale given the state, and (*ii*) the vector (action taken, likelihood ratio) a Markov chain. We prove the power and simplicity of this informational herding framework in four major ways:

The Markovian aspect of the dynamics allows us to drastically narrow the range of possible long run outcomes, as we need only focus on the ergodic set. This set is wholly unrelated to initial conditions, and depends only on the transition dynamics of the model. By contrast, the martingale property affords us a different glimpse into the long run dynamics, tying them down to the initial conditions in expectation. As it turns out, this allows us to eliminate from consideration the not inconceivable elements of the ergodic set where everyone entertains entirely false beliefs in the long run.

• Bad Herds? Rational learning is fruitful because Bayes-updating of *any* space of private signals from a common prior yields *private beliefs* that are *informative* of the true state — a simple corollary of our *no introspection property.* We focus on the role played by the support of the private beliefs: Incorrect herds can develop exactly when individuals have *bounded* private beliefs, i.e. there do not exist arbitrarily strong private signals. Intuitively, whenever individual private signals are *uniformly* bounded in strength, history has the potential to mislead everyone: Eventually even the most doctrinaire individual dare not quarrel with the conclusion of histories that aggregate enough information.

With *unbounded private beliefs,* learning is complete. Soon enough, a wise enough doctrinaire individual will appear whose contrary action will radically shift public beliefs and overturn any would-be incorrect herd. Yet by our overturning principle, this logic cannot be turned on its head to rule out correct herds.[1] The admission of arbitrarily tenacious individuals is largely a modelling decision. For it not only provides us with a richer model than Banerjee (1992) and BHW, and opens lines of inquiry that make no sense in their framework, but also sheds critical light into the exact failure of incomplete learning: For as we approach this idealized extreme, bad herds become vanishingly implausible.

• Cascades? BHW introduced the colorful terminology of a *cascade* for an infinite train of individuals acting irrespective of the content of their signal. Yet we argue that the label 'cascades literature' is malapropos. For cascades are the exception and not the rule outside BHW's discrete signal world. All but one (rather contrived) example in this paper attests to this fact. With generic signal distributions, individuals always eventually settle on an action (a herd), but no decision is ever a foregone conclusion (a cascade)! With these two notions decoupled, the resulting analysis is much richer — for it is no longer so clear why herds must arise. This occurs, we argue, because *public beliefs* must converge (as per usual), and since belief convergence implies action convergence. This simple insight (*the overturning principle*) into herding is true because contrary actions radically swing beliefs.

• Expected Time To Herd. We characterize when herds arise in *finite expected time.* We find that what matters is not how fast the truth is learned but rather how slowly

[1]The link between unbounded support beliefs and complete learning was implicit in Smith (1991).

error is rooted out: There must be enough contrarians to overturn any temporary herd fast enough. Our surprising discovery is that while learning is complete with unbounded beliefs, the correct herd always requires infinite mean time to start in some state!

• NOISE. That a single individual can 'overturn the herd' is the key analytic insight into herding. So it is only natural to undermine the impact of contrarians by adding noise. Counterintuitively, even with a constant inflow of crazy/trembling individuals, complete learning still obtains. Everyone (sane) eventually learns the true state of the world.

Besides definitively resolving and further exploring informational herding, the above results set the context needed for what are our most striking and innovative findings below.

2. New Herding Economics: Confounding with Multiple Preference Types.

We next relax a critical (if under-appreciated) premise of the original herding results, namely that all individuals have the *same preferences.* Surely this is anything but an apt description of the world. Multiple types offers a parallel reason why the actions of isolated individuals need not greatly matter — but with much richer consequences. Let's fix ideas with a hopefully familiar example. Suppose that on a highway under construction, depending upon how the detours are arranged, those going to Houston should merge either right (in state R) or left (in state L), with the opposite for those headed toward Dallas. If one knows that 70% are headed toward Houston, then absent any strong signal to the contrary, Dallas-bound drivers should take the lane 'less traveled by'. This yields two herding outcomes: 70% left or 70% right. But another rather subtle possibility may arise. For as the chance q that observed history accords state R rises from 0 to 1, the chance $r(q)$ that a Houston driver merges right gradually increases from 0 to 1, and conversely for Dallas drivers. If for some q, cars are equilikely in states R and L to merge right, or $r_R(q) = r_L(q)$, then no inference can be drawn from additional decisions, and all learning stops! Of course, whether such a fixed point exists is far from obvious, and even if so, why need we converge there? The surprising content of Theorem 9 is that for non-degenerate specifications, such a 'confounding' outcome *does exist,* and furthermore, dynamics *will converge* upon it with positive probability — even with arbitrarily strong private signals!

3. Likelihood Ratios as Martingales ⇒ Exponentially Fast/Stable Learning.

The above convergence result and the deduction of 'rational herds' with noise both depend on the rate of belief convergence. With noise, the inflow of rational contrarians may be forever choked off without entering a cascade if learning is exponentially fast. Such rapid belief convergence also implies the *local stochastic stability* of learning needed for confounded learning to arise with multiple types. The common ingredient is a simple link we have found between the martingale character of the likelihood ratio and the exponential stability of learning: This neat result holds if near a fixed point, posterior beliefs aren't (degenerately) equally responsive to prior beliefs for every action taken. We feel this technique has broad applicability beyond the herding paradigm, into optimal experimentation.

Section 2 outlines the benchmark model and some preliminary results. Section 3 provides the associated action and belief convergence results; noise is added in section 4. Our new model with heterogeneous preferences is explored in section 5. A host of needed (or related) results and proofs are appendicized, including a convergence criterion for Markov-martingale processes, and a new local stability result for stochastic difference equations.

2. THE STANDARD MODEL

2.1 Some Notation

An infinite sequence of individuals $n = 1, 2, \ldots$ takes actions in that exogenous order. *Everyone observes the ordered actions of all predecessors.* A background probability space $(\Omega, \mathcal{E}, \nu)$ captures all uncertainty in our model. First, the action payoffs are random: There are $S = 2$ possible *states of the world* (or just *states*), $s = H$ ('high') and $s = L$ ('low'). Formally, Ω is partitioned into events $\Omega^H \cup \Omega^L$, called H and L. Let the common prior belief be $\nu(H) = \nu(L) = 1/2$.[2] Our results extend to any finite number of states, but at significant algebraic cost and dubious conceptual gain (see our 1996 working paper).

Everyone chooses from a finite action set $\langle a_m, m \in \mathcal{M} \rangle$, where $\mathcal{M} = \{1, \ldots, M\}$. One might think of investors deciding whether to 'invest' or 'decline' an investment project of uncertain value. Action a_m pays off $u^s(a_m)$ in state $s \in \{H, L\}$, *the same for all individuals,* and everyone acts so as to maximize his expected payoff. We assume that WLOG no action is weakly dominated, and at least two undominated actions exist. Before selecting an action, an individual observes the entire action *history* profile, loosely denoted h.

Individual n receives a private random signal, $\sigma_n \in \Sigma$, about the state of the world. *Conditional on the state,* $\{\sigma_n\}$ are i.i.d., and drawn according to the probability measure μ^s in state $s \in \{H, L\}$.[3] To ensure that no signal will *perfectly* reveal the state of the world, we shall insist that μ^H and μ^L be mutually absolutely continuous (a.c.).[4] Thus, there exists a positive, finite Radon-Nikodym derivative $g = d\mu^L/d\mu^H : \Sigma \to (0, \infty)$ of μ^L w.r.t. μ^H. And to avoid trivialities, we shall *rule out* $g = 1$ almost surely,[5] so that μ^H and μ^L are not the same measure — i.e. some signals are *informative* about the state.

[2] Common priors is standard (Harsanyi (1967–68)), and a flat prior WLOG (see our working paper).

[3] So $\sigma_n : \Omega \to \Sigma$ is a random variable; $\mu^s \equiv \mu^s_n = \nu^s \circ \sigma_n^{-1}$, where measure ν^s conditions ν on event Ω^s.

[4] See Rudin (1987). Measure μ^L is a.c. w.r.t. μ^H if $\mu^H(S) = 0 \Rightarrow \mu^L(S) = 0\ \forall S \in \mathcal{S}$, the σ-algebra on Σ. By the Radon-Nikodym Theorem, a unique $g \in L^1(\mu^H)$ exists with $\mu^L(S) = \int_S g\, d\mu^H$ for every $S \in \mathcal{S}$.

[5] With μ^H, μ^L mutually a.c., 'almost sure' assertions are well-defined without specifying the measure.

2.2 Private Beliefs

The second source of uncertainty in our model is private information. Given signal $\sigma \in \Sigma$, an individual uses Bayes' rule to arrive at what we shall refer to as his *private belief* $p(\sigma) = 1/(g(\sigma)+1) \in (0,1)$ that the state is H. Conditional on the state, private beliefs are i.i.d. across individuals because signals are. In state $s \in \{H, L\}$, p is distributed with a c.d.f. F^s on $(0,1)$. The distributions F^H and F^L are subtly linked. Since μ^H and μ^L are mutually absolutely continuous, so are the associated distributions F^H and F^L. Thus there exists a Radon-Nikodym derivative $f \equiv dF^H/dF^L$, which reduces to $f^H(p)/f^L(p)$ when each F^s has a density. The next result neatly illustrates the folk wisdom that posterior beliefs are sufficient for one's information, and ultimately drives all learning that occurs.

Lemma 1 *(a) (No Introspection Condition) The derivative $f \equiv dF^H/dF^L$ of private belief c.d.f.s F^H, F^L satisfies ($\bigstar$): $f(p) = p/(1-p)$ almost surely in $p \in (0,1)$. Conversely, if ($\bigstar$) then F^H, F^L arise from updating a common prior with some signal measures μ^H, μ^L.*
(b) The difference $F^L(p) - F^H(p)$ is nondecreasing or nonincreasing as $p \lesseqgtr 1/2$.
(c) (Beliefs are Informative) $F^H(p) < F^L(p)$ except when both terms are 0 or 1.

Proof: If the individual further updates his private belief p by asking of its likelihood in the two states of the world, he must learn nothing more. So, $p = f(p)/[1+f(p)]$, as desired. Conversely, given $f(p) = p/(1-p)$, let σ have distribution F^s in state s, $s \in \{H, L\}$.

Part (a) implies part (b), and is strict when $p \in \operatorname{supp}(F) \setminus \{1/2\}$.[6] Hence (c) follows. □

Thus, the conditional distribution functions F^H and F^L can be taken as the stochastic primitives of the model. Notice that Lemma 1(a) implies that they have a common support, say $\operatorname{supp}(F)$, equal to the range of $p(\cdot)$ on Σ. The structure of $\operatorname{co}(\operatorname{supp}(F)) \equiv [\underline{b}, \bar{b}] \subseteq [0,1]$ plays a major role in the paper. Note that $0 \le \underline{b} < 1/2 < \bar{b} \le 1$ since μ^L and μ^H are distinct, and $F^L(\underline{b}) = 0$ and $F^H(\bar{b}-) = 1$ as there are no perfectly revealing signals. We call the private beliefs *bounded* if $0 < \underline{b} < \bar{b} < 1$, and *unbounded* if $\operatorname{co}(\operatorname{supp}(F)) = [0,1]$.[7]

For clarity, we introduce two leading examples which will be progressively embellished.

• Unbounded Beliefs Example. Let μ^L have probability density $g^L(\sigma) = 2 - 2\sigma$, and μ^H the density $g^H(\sigma) = 2\sigma$ (left panel of figure 1). So $g(\sigma) = g^L(\sigma)/g^H(\sigma) = (1-\sigma)/\sigma$, yielding $p(\sigma) = \sigma$, and thus the simple formulae $F^H(p) \equiv p^2 < F^L(p) \equiv 2p - p^2$.

• Bounded Beliefs Example. Next, let μ^H be Lebesgue (uniform) measure on $[0,1]$, and choose μ^L with Radon-Nikodym derivative $g(\sigma) = \frac{d\mu^L(\sigma)}{d\mu^H(\sigma)} = 3/2 - \sigma$ on $[0,1]$. Then $p(\sigma) = 1/(g(\sigma)+1) = 2/(5-2\sigma)$. Note how p maps $[0,1]$ injectively onto $[\underline{b}, \bar{b}] = [2/5, 2/3]$,

[6] The support of a measure is the smallest closed set of full measure; $\operatorname{co}(A)$ is the convex hull of set A.
[7] To exhaust all possibilities we should also consider supports that are bounded above and not below, etc., but this exercise in generality yields no new insights.

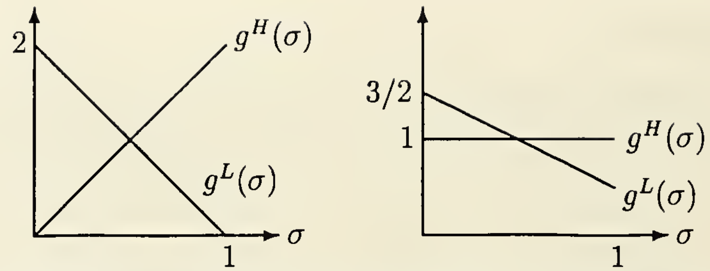

Figure 1: **Signal Densities.** Graphs for the unbounded (left) and bounded (right) beliefs examples. Observe how, for instance, signals near 0 are very strongly in favor of state L, as in Lemma 1.

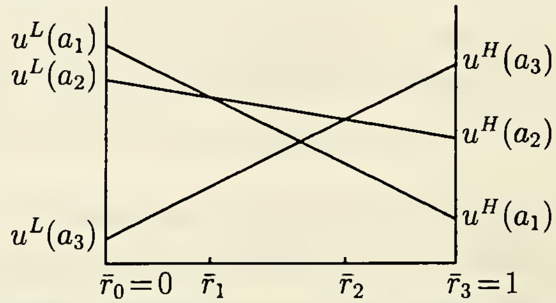

Figure 2: **Payoff Frontier.** The diagram depicts the (expected) payoff of each of three actions as a function of the posterior belief r that the state is H. The individual simply chooses the action yielding the highest payoff. Action a_2 is an *insurance action,* while actions a_1 and a_3 are *extreme actions.*

with $p(\sigma) \leq p \Leftrightarrow 2/(5-2\sigma) \leq p \Leftrightarrow (5p-2)/2p \geq \sigma$. Here, the support of F^H, F^L is bounded, and the distribution of $p \in [2/5, 2/3]$ is $F^H(p) = \mu^H[0, (5p-2)/2p] = (5p-2)/2p$ in state H and in state L:

$$F^L(p) = \int_{p(\sigma)\leq p} g(\sigma)\, d\sigma = \int_0^{\frac{5p-2}{2p}} \frac{3}{2} - \sigma\, d\sigma = \frac{1}{2}(3\sigma - \sigma^2)\Big|_0^{\frac{5p-2}{2p}} = \frac{(5p-2)(p+2)}{8p^2}$$

2.3 Action Choice

Given a posterior belief $r \in [0,1]$ in state H, the expected payoff of choosing action a is $ru^H(a) + (1-r)u^L(a)$. Figure 2 depicts the content of the next (standard) result.

Lemma 2 *The interval* $[0,1]$ *partitions into subintervals, or* action basins, $\tilde{I}_1, \ldots, \tilde{I}_M$ *overlapping at endpoints only, such that action a_m is optimal with posterior belief $r \in \tilde{I}_m$.*

Proof: As noted, the payoff of each action is a linear function of r. But by assumption, action a_m is strictly best for some r; therefore, there must exist a single open subinterval of $[0,1]$ where it strictly dominates all other actions. That this is a partition follows from the fact that there exists at least one optimal action for each posterior $r \in [0,1]$. □

We now WLOG order the actions so that a_m is optimal exactly when the posterior $r \in [\bar{r}_{m-1}, \bar{r}_m] \equiv \tilde{I}_m$, where $0 = \bar{r}_0 < \bar{r}_1 < \cdots < \bar{r}_M = 1$. We employ the tie-breaking rule[8] that individuals take action a_m (versus a_{m+1}) at $r = \bar{r}_m$. The *extreme action* a_M (resp. a_1) is optimal when one is certain that the state is H (resp. L), while *insurance actions* $a_2, \ldots, a_{M-1}$ are respectively optimal as confidence shifts from state L to state H.

• EXAMPLES CONT'D. In our running examples, the two actions may be a_1 = Decline and a_2 = Invest. To Invest yields payoff u in state H and -1 in state L; to Decline from investing yields payoff 0 in both states. Since, by assumption, action a_1 is undominated, we posit $u > 0$. Indifference prevails when $0 = \bar{r}u - (1 - \bar{r})$, so $\bar{r} = 1/(1 + u)$.

2.4 Individual Learning

Individual decision rules map from private beliefs and history to own actions. In a Bayesian equilibrium, everyone knows all decision rules and the common prior,[9] and can compute the ex ante chance $\pi^s(h)$ of any history h in each state s. This yields the public *likelihood ratio* $\ell(h) = \pi^L(h)/\pi^H(h)$ that the state is L *versus* H, and the *public belief* $q(h)$ in state H, i.e.

$$q(h) = \frac{\pi^H(h)}{\pi^H(h) + \pi^L(h)} = 1/(1 + \ell(h))$$

So $q(h)$ is the posterior ensuing from a neutral private belief and history observation h.

A final application of Bayes rule yields the *posterior belief* r (that the state is H) in terms of the public history — or equivalently the likelihood ratio $\ell(h)$ — and the private belief p:

$$r = \frac{p\,\pi^H(h)}{p\,\pi^H(h) + (1-p)\,\pi^L(h)} = \frac{p}{p + (1-p)\,\ell(h)} = \frac{1}{1 + \frac{1-p}{p}\ell(h)} \tag{1}$$

Lemma 3 (Private Belief Thresholds) *Given history h, there are $M + 1$ thresholds $0 = \bar{p}_0(h) \le \bar{p}_1(h) \le \ldots \le \bar{p}_M(h) = 1$, such that a_m is chosen iff the private belief satisfies $p \in (\bar{p}_{m-1}(h), \bar{p}_m(h)]$, where*

$$\frac{\bar{p}_m(h)}{1 - \bar{p}_m(h)} = \frac{\bar{r}_m}{1 - \bar{r}_m}\ell(h) \tag{2}$$

This is true given (i) the tie-break rule, (ii) the RHS of (1) is strictly increasing in p, and (iii) the reformulation of (1) as posterior odds $(1-r)/r$ equal private odds $(1-p)/p$ times the likelihood ratio $\ell(h)$. Note that if $\bar{p}_{m-1}(h) = \bar{p}_m(h)$ at some h, then a_m is never chosen.

Since the likelihood ratio is informationally sufficient for the history, we suppress the explicit dependence $\ell(h)$, and write $\bar{p}_m(\ell)$ instead of $\bar{p}_m(h)$, where $\ell \mapsto \bar{p}_m(\ell)$ is increasing.

• EXAMPLES CONT'D. In both examples, (2) yields $\bar{p}(\ell) = \ell/(u + \ell)$ if $\bar{r} = 1/(1 + u)$.

---

[8] For generic models, this choice does not matter. Even when it does change the probabilistic course of nongeneric models, the tie-breaking rule does not change the statement of any of our theorems.

[9] We assume common knowledge of rationality. Section 4 partially backs away from this.

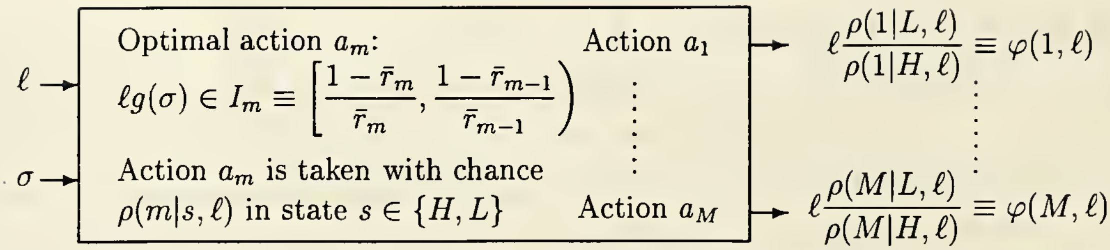

Figure 3: **Individual Black Box.** Everyone bases his decision on both the public likelihood ratio ℓ and his private signal σ, resulting in his action choice a_m and a likelihood ratio to confront successors. One takes action a_m iff one's *posterior likelihood* lies in the interval I_m, where $I_1, \ldots, I_M$ partition $[0, \infty]$

2.5 Corporate Learning as a Markov-Martingale Process

We let ℓ_n and q_n, respectively, be the public likelihood ratio and belief after Mr. n chooses action m_n,[10] with $\ell_0 = 1$ and $q_0 = 1/2$ (null initial history). Signals, and thereby actions, are random, and so $\langle \ell_n \rangle_{n=1}^{\infty}$ and $\langle q_n \rangle_{n=1}^{\infty}$ are stochastic processes, described by

$$\rho(m|s,\ell) = F^s(\bar{p}_m(\ell)) - F^s(\bar{p}_{m-1}(\ell)) \tag{3}$$

$$\varphi(m,\ell) = \ell\rho(m|L,\ell)/\rho(m|H,\ell) \tag{4}$$

Here, $\rho(m|s,\ell)$ is the chance that a (*rational*) individual takes action a_m, given the public likelihood ℓ, and the true state $s \in \{H, L\}$. Faced with ℓ_n, if individual n takes action m_n, we move to $\ell_{n+1} = \varphi\langle m_n, \ell_n\rangle$. Figure 3 schematically summarizes this transition.

Our insights are best expressed by considering $\langle m_n, \ell_n \rangle$ as a time homogeneous *Markov process* on the state space $\mathcal{M} \times [0, \infty)$. Given $\langle m_n, \ell_n \rangle$, (3) and (4) imply that the next state is $\langle m_{n+1}, \varphi(m_{n+1}, \ell_n)\rangle$ with probability $\rho(m_{n+1}|H, \ell_n)$ in state H. In principle, such a two-dimensional process with continuous state variables could be very ill-behaved, and possibly chaotic. But Lemma 5 attests to how well-behaved is the likelihood process.

The next result is quite standard (but for completeness, is proven in the appendix).

Lemma 4 (The Unconditional Martingale) *The public belief $\langle q_n \rangle$ is a martingale, unconditional on the state of the world.*[11]

This martingale describes the forecast of subsequent public beliefs by individuals *in the model,* who do not know the true state of the world: Prior to receiving his signal, individual n's expectation of the public belief that will confront his successor is the current one. But for our purposes, an unconditional martingale does not tell us all we want to know about convergence. For that, we will condition on the state of the world, and that will render

[10] Or equivalently, action a_{m_n}. Throughout the paper, m will denote actions, and n individuals.

[11] We really ought to specify the accompanying sequence of σ-algebras to the stochastic process; however, these will be suppressed because they are simply the ones generated by the process itself.

the public belief $\langle q_n \rangle$ a *sub*martingale in state H (and a *super*martingale in state L), i.e. $E[q_{n+1} \mid H, q_1, \ldots, q_n] \geq q_n$. Essentially, the public beliefs are expected to become weakly more focused on the true state of the world — a result much weaker than we seek. For the *modeller,* a much more useful martingale is one that conditions on knowledge of the state of the world, namely, the likelihood process $\langle \ell_n \rangle = \langle (1 - q_n)/q_n \rangle$.[12]

Lemma 5 (Likelihood Ratios as a Conditional Martingale) *Assume state H.*
(a) The stochastic process of likelihood ratios $\langle \ell_n \rangle$ is a martingale conditional on state H.
(b) The likelihood ratio process $\langle \ell_n \rangle$ converges almost surely to a r.v., $\ell_\infty = \lim_{n\to\infty} \ell_n$, with $\mathrm{supp}(\ell_\infty) = [0, \infty)$. *So fully incorrect learning ($\ell_n \to \infty$) almost surely cannot occur.*

Proof: Given the value of ℓ_n, the conditional expectation of ℓ_{n+1} in state H is

$$E[\ell_{n+1} \mid H, \ell_1, \ldots \ell_n] = \sum_{m \in \mathcal{M}} \rho(m|H, \ell_n)\ \ell_n\ \frac{\rho(m|L, \ell_n)}{\rho(m|H, \ell_n)} \ = \ \ell_n \sum_{m \in \mathcal{M}} \rho(m|L, \ell_n) \ = \ \ell_n$$

Since the likelihood ratios are non-negative random variables, the result follows from the Martingale Convergence Theorem. (See Breiman (1968), Theorem 5.14.) Or, since the stochastic evolution of $\langle \ell_n \rangle$ is mean-preserving, convergence to any dead wrong belief a.s. cannot occur: The odds against the truth are not permitted to explode.[13]

Easley and Kiefer (1988), and others, underscore that unlike statistical learning where information may well accrue at a 'constant rate', complete learning is not at all a foregone conclusion in a economic model of costly experimentation. When information has a price, complete learning is generally deemed too expensive. One might expect that the resulting pathological outcomes to the learning dynamics persist this observational learning setting.

GUIDING QUESTIONS: **1.** Is there a *herd,* or *action convergence*: Does the first coordinate of the process $\langle m_n, \ell_n \rangle$ settle down? If so, on action a_M? And in finite expected time?
2. *Belief convergence* obtains — the second coordinate of the Markov process $\langle m_n, \ell_n \rangle$ converges — since ℓ_∞ exists by Lemma 5. But must a *cascade* arise, where everyone takes action a_m irrespective of signal realizations — $\rho(m|H, \ell_n) = \rho(m|L, \ell_n) = 1$ for some m after some stage n? Or perhaps only a *limit cascade* on a_m will arise, i.e. $\rho(m|H, \ell_n) \to 1$ as $n \to \infty$. And is learning *complete* — or do beliefs converge to the truth, i.e. $\ell_n \to 0$? Otherwise, learning is *incomplete* (beliefs not eventually focused on state H).
3. What is the *link* between action and belief convergence? The rough logic of BHW, valid with a discrete signal space, is: (*i*) cascades must occur, and (*ii*) cascades imply herds, and

[12]See, for instance, Doob (1953), section II.7.
[13]Another proof of this fact uses public beliefs (see, for instance, Bray and Kreps (1987)).

thus (iii) herds occur. The second step is irrefutable: In a cascade on action a_m, Mr. n ignores his signal and takes a_m, revealing no new information, and so $\ell_{n+1} = \ell_n$. So the cascade still obtains at stage $n+1$, as private belief thresholds are unchanged by (2). So $\rho(m|H, \ell_{n+1}) = 1$ too.

We shall prove that (i) is not robust: Cascades generically needn't arise — and *cannot* with unbounded beliefs. This leads to a tougher question: Do limit cascades imply herds?

3. THE MAIN RESULTS

In this section, we first characterize the limit ℓ_∞. We then prove that this convergence of beliefs implies convergence of actions, i.e. limit cascades imply that herds eventually occur. We conclude by discussing the speed at which beliefs converge, and the intertwined issue of whether the mean time to entry into a herd is finite.

3.1 Belief Convergence

We first must understand exactly why individuals might wish to ignore their signals.

Lemma 6 (Action Absorbing Basins) *For each action a_m, there is a possibly empty interval $J_m = \{\ell \,|\, \mathrm{supp}(F) \subseteq [\bar{p}_{m-1}(\ell), \bar{p}_m(\ell)]\} \subset I_m \subset [0, \infty]$, such that if the public likelihood ratio $\ell \in \mathrm{int}\,(J_m)$, then the posterior likelihood ratio $\ell g(\sigma) \in I_m$ almost surely in private signals σ. Action a_m is then taken almost surely, and ℓ is unchanged. Also,*
(a) With bounded private beliefs, $J_1 = [\bar{\ell}, \infty]$ and $J_M = [0, \underline{\ell}]$ for some $0 < \underline{\ell} < \bar{\ell} < \infty$;
(b) With unbounded private beliefs, $J_M = \{0\}$, $J_1 = \{\infty\}$, and all other basins are empty.

The appendicized proof is intuitive: With bounded private beliefs, the posterior odds are only boundedly far from ℓ. So once ℓ is sufficiently near 0 or ∞ or perhaps even in favor of an insurance action, all private signals must lead to the same action. But with unbounded beliefs, every public likelihood $\ell \in (0, \infty)$ will be swamped by some mass of private signals.

Remarks. **1.** The absorbing basins are inversely ordered, or $J_{m_2} \ll J_{m_1}$ iff $m_2 > m_1$. **2.** The lemma asserts that to each 'extreme action' corresponds an action absorbing basin. But $J_m \neq \varnothing$ is also possible for an 'insurance' action a_m when beliefs are bounded. For instance, take $u_H(a_1) = u_L(a_3) = 1$, $u_L(a_1) = u_H(a_3) = 0$, and $u_H(a_2) = u_L(a_2) = 1 - \varepsilon$. Then the insurance action a_2 has an action absorbing basin for small enough $\varepsilon > 0$.
3. By rearranging an expression like (2), one can show that ℓ is in J_m precisely when

$$\bar{r}_{m-1} \leq \frac{\underline{b}}{\underline{b} + (1-\underline{b})\ell} \quad \text{and} \quad \frac{\bar{b}}{\bar{b} + (1-\bar{b})\ell} \leq \bar{r}_m$$

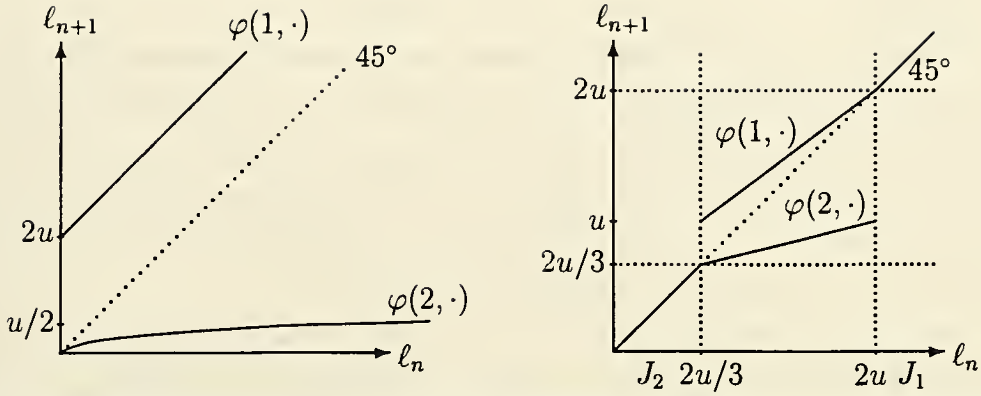

Figure 4: **Continuations and Absorbing Basins.** Continuation functions for the examples: unbounded private beliefs (left), and bounded private beliefs (right). By the martingale property, the expected continuation lies on the diagonal. The stationary points are where both arms hit the diagonal (impossible here), or where one arm is taken with zero chance ($\ell = 0$ in the left panel; $\ell = 2u/3$ in the right).

So, *an action absorbing basin is larger the smaller is the support* $[\underline{b}, \bar{b}]$, *and the larger is the interval* $[\bar{r}_{m-1}, \bar{r}_m]$. Only extreme action absorbing basins exist for large enough $[\underline{b}, \bar{b}]$.

• Unbounded Beliefs Example Cont'd. Private beliefs $p \in (0,1)$ are distributed as $F^H(p) = p^2$ and $F^L(p) = 2p - p^2$. So $\text{supp}(F) = [0,1]$, and private beliefs are unbounded; the basins collapse to the extreme points, $J_1 = \{\infty\}$, $J_M = \{0\}$. With our two actions, we have $\bar{p} = \ell/(u+\ell)$, where $u > 0$. Thus we let $\rho(1|H,\ell) = \ell^2/(u+\ell)^2$, and $\rho(1|L,\ell) = \ell(\ell+2u)/(u+\ell)^2$. We now get $\varphi(1,\ell) = \ell + 2u$ and $\varphi(2,\ell) = u\ell/(u+2\ell)$, shown in figure 4.

• Bounded Beliefs Example Cont'd. Here, $F^H(p) = (5p-2)/2p$ and $F^L(p) = (5p-2)(p+2)/8p^2$ for $p \in [2/5, 2/3]$. With $\bar{p} = \ell/(u+\ell)$, active dynamics occur when $\ell \in (2u/3, 2u)$. For $\ell \le 2u/3$, we have $\rho(1|H,\ell) = \rho(1|L,\ell) = 0$, i.e. action a_2 is taken a.s., and thus its absorbing basin is $J_2 = [0, 2u/3]$. For $\ell \ge 2u$, we similarly find $J_1 = [2u, \infty]$. For $\ell \in (2u/3, 2u)$ we have $\rho(1|H,\ell) = (3\ell - 2u)/2\ell$ and $\rho(1|L,\ell) = (3\ell-2u)(3\ell+2u)/8\ell^2$. By the martingale property, $\varphi(1,\ell) = u/2 + 3\ell/4$ and $\varphi(2,\ell) = u/2 + \ell/4$. (See figure 4.)

We now argue that *limit cascades* must occur: Dynamics must tend to one of the basins.

Theorem 1 (Limit Cascades) *The likelihood ratio process $\ell_n \to \ell_\infty$, and the random variable ℓ_∞ has support $J \equiv J_1 \cup \cdots \cup J_M$, the* absorbed set.

The appendicized proof is (we feel) intuitive. Theorems B.1–2 precisely characterizes ℓ_∞: Any point $\hat{\ell} \in \text{supp}(\ell_\infty)$ must be *stationary* for the Markov process, i.e. either an action doesn't occur ($\rho(m|\hat{\ell}) = 0$) or it teaches us nothing ($\varphi(m,\hat{\ell}) = \hat{\ell}$). If at least two actions are taken in the limit, then the least one is taken with greater chance in state L than state H, simply because beliefs are informative (Lemma 1(c)). This will violate stationarity.

Theorem 2 (Long-run Learning: Complete and Incomplete) *Assume state H.*
(a) *With bounded private beliefs, if $\ell_0 \notin J_M$, then with positive chance $\ell_\infty \in J_1 \cup \cdots \cup J_{M-1}$.*

(*b*) *With unbounded private beliefs,* $\ell_n \to 0$ *almost surely.*

Proof: For (*a*), if $\ell_\infty \in J_1$ with positive chance, we are done. Otherwise, if $\ell_\infty \notin J_1$ a.s., then $\langle \ell_n \rangle$ is uniformly bounded above by $\bar{\ell}$, the infimum of J_1, introduced in Lemma 6. By Lebesgue's Dominated Convergence Theorem, the mean of $\langle \ell_n \rangle$ is preserved in the limit, i.e. $E[\ell_\infty] = \ell_0$, the initial ratio. So, if $\ell_0 > \underline{\ell}$ it cannot be the case that $\text{supp}(\ell_\infty) \subseteq J_M = [0, \underline{\ell}]$. (*b*) Theorem 1 asserts $\ell_\infty \in J$ a.s. By Lemma 6, with unbounded beliefs $J = \{0\} \cup \{+\infty\}$, while Lemma 5 proves $\text{supp}(\ell_\infty) \subseteq [0, \infty)$ if the state is H. □

To underscore how *non*intuitive is this conclusion, notice that this is one situation where strict inequality holds in *Fatou's Lemma,* or $1 = \lim_{n\to\infty} E[\ell_n] > E[\lim_{n\to\infty} \ell_n] = 0$, and so $\langle \ell_n \rangle$ must be unbounded. While the process $\langle \ell_n \rangle$ occasionally gets arbitrarily large, corresponding to arbitrarily long trains of individuals choosing the least optimal action, the longer is the train, the less likely it is to occur. On balance, the Theorem 2 tells us that all such trains must almost surely come to an end. For a classic analogy to the behavior of $\langle \ell_n \rangle$, think of the behavior of a driftless random walk (or Brownian motion) starting at 1 with an absorbing barrier at 0. With probability one, it eventually hits 0 and is absorbed.

Still, this is an arresting result, on two counts.

PUZZLE # 1. Why can't individuals eventually be wholly mistaken about the state of the world? For as noted in section 2, convergence towards totally incorrect beliefs appears self-enforcing. But the martingale convergence theorem ruled out that limit.

PUZZLE # 2. Why complete learning? Why aren't correct herds periodically broken up (just as incorrect herds are)? Couldn't beliefs be ever cycling betwixt confidence in H and L, so that the ergodic distribution assigns weight to non-stationary beliefs? It is here that the martingale convergence theorem succeeds where Markovian arguments fail, and establishes that beliefs must eventually settle down: Limit cycles cannot occur. Note that the analysis of BHW — which did not appeal to martingale methods — only succeeded because their stochastic process necessarily settled down in some (stochastic) finite time.

3.2 Action Convergence

Suppose a *putative herd* has begun: A string of individuals has taken the identical action, but a cascade has not yet begun. If someone then acts in a contrary fashion, his successors have no choice but to concede the strength of his signal, thus sharply revising the public belief. We say that the putative herd has been *overturned* by the unexpected action. More formally, if Mr. n chooses action a_m, then $n+1$ should, *before* he observes his own private signal, find it optimal to follow suit because he knows no more than n, and since it is common knowledge that n rationally chose a_m: So after n's action, the public belief is $q \in \tilde{I}_m \equiv (\bar{r}_{m-1}, \bar{r}_m]$. The next lemma codifies this logic (proof appendicized), for it proves central to an understanding of the entire observational learning paradigm.

Lemma 7 (The Overturning Principle) *For any history, if someone optimally takes action a_m, the updated likelihood ratio $\ell \in I_m$, and an uninformed successor follows suit.*

This result precludes infinitely many contrary actions in a limit cascade, for that would negate belief convergence — eg. for the running examples, the principle correctly predicts $\varphi(1, \ell) \in [u, \infty)$ and $\varphi(2, \ell) \in [0, u)$ (see figure 4). So *belief convergence implies action convergence,* yielding the action counterparts to Theorems 1–2 (details in working paper).

Theorem 3 (Herds) *A herd on some action will almost surely arise in finite time.*
(a) With bounded private beliefs, absent a cascade on the most profitable action a_M (in state H) from the outset, a herd arises on an action other than a_M with positive probability.
(b) With unbounded private beliefs, individuals almost surely settle on the optimal action.

The bounded beliefs analysis proves in generality (and we think, simplicity) the major pathological herding finding in Banerjee (1992) and BHW.[14] Strictly bounded beliefs so happens to be the mainstay for their striking result, as the characterization makes clear.[15]

The major reason to emphasize the above characterization is the continuous transition from incomplete to complete learning as private beliefs tend from bounded to unbounded.

Theorem 4 (Continuity) *Fix the payoffs and prior beliefs. If* co(supp(F^k)) *converges to* $[0, 1]$,[16] *then the chance of an incorrect limit cascade vanishes as $k \to \infty$.*

Proof: Only basins $J_1^k = [\bar{\ell}^k, \infty]$ and J_M^k remain once co(supp(F^k)) is close enough to $[0, 1]$. If π^k is the chance of a correct limit cascade, then $E\ell_\infty \geq (1 - \pi^k)\bar{\ell}^k$. But $E\ell_\infty \leq \ell_0$ by Fatou's Lemma, so that $\pi^k \geq 1 - \ell_0/\bar{\ell}^k$. As $k \to \infty$, $\bar{\ell}^k \to \infty$, and so $\pi^k \to 1$. □

REMARK: FRAGILITY OF LIMIT CASCADES. BHW devote considerable attention to the fragility of herds, pointing out that the release of a small amount of public information can undo a herd. Our results on the nonexistence of cascades only serve to strengthen this insight. Even in the limit, we have shown that generically (i.e. without atoms) the likelihood ratio lies at the edge of the absorbing basin. Consequently, *arbitrarily little* public information will break the limit cascade. By contrast, the limit belief in BHW is bounded away from the edge of the absorbing basin, thus inoculating the model to the release of sufficiently small packets of public information.

[14] And extended to $M > 2$ actions. BHW also handled several states (done in our working paper).

[15] While BHW didn't consider unbounded private signals, their working paper introduces perfectly informative signals under the rubric of 'pseudo-cascades' (*ruled out* by our assumption of mutually a.c. information). Complete learning in that context is clear. For once the public beliefs overwhelm all but the perfectly revealing signals, the very next contrary action reveals the state of the world, and we're done!

[16] We mean the Hausdorff topology: If co(supp(F^k)) $= [a_k, b_k]$, then $a_k \to 0$, and $b_k \to 1$.

3.3 Must Cascades Exist with Bounded Beliefs?

Or rather, we ask "Can cascades exist...?" As $\langle m_n, \ell_n \rangle$ is not a *finite state* Markov chain, Theorem 2 cannot assert finite time convergence (a cascade) — for we may have $\ell_n \to \hat{\ell} \in J$ but always $\ell_n \notin J$. Unlike BHW, cascades need not obtain, though herds must.

In the running bounded beliefs example, a cascade on action a_1 obtains iff $\ell \geq 2u$. In a herd (and so limit cascade) on a_1, with $\ell_n \to \hat{\ell} \in [2u, \infty)$, figure 4 clearly shows that $\langle \ell_n \rangle$ *never enters* J_1, *but only approaches its edge. A cascade never starts.*[17] So learning never ceases: No matter how many individuals have followed suit, a contrarian might still appear. *A cascade on* a_i *can only arise with a nonmonotonic likelihood transition function* $\varphi(i, \cdot)$.

With a *discrete signal distribution,* BHW deduced that cascades must occur, and we can easily see why this is true. For if a herd starts on action a_m, with int $(J_m) \neq \varnothing$, then since $\ell_n < \hat{\ell} \equiv \inf(J_m)$ implies $F^L(\bar{p}_{m-1}(\ell_n)) = F^H(\bar{p}_{m-1}(\ell_n)) = 0$, we have

$$\ell_{n+1} = \varphi(m, \ell_n) \equiv \ell_n \frac{F^L(\bar{p}_m(\ell_n)) - F^L(\bar{p}_{m-1}(\ell_n))}{F^H(\bar{p}_m(\ell_n)) - F^H(\bar{p}_{m-1}(\ell_n))} = \ell_n \frac{F^L(\bar{p}_m(\ell_n))}{F^H(\bar{p}_m(\ell_n))} \tag{5}$$

So $\ell_{n+1}/\ell_n > \inf\{F^L(p)/F^H(p) | p \in \text{int}(\text{supp}(F))\} > 1$ by Lemma A.1, and $\langle \ell_n \rangle$ must 'jump into' J_m in boundedly finite time. This needn't occur with general signal spaces.

One might prematurely conclude that cascades can *only* arise with discrete signal distributions, and thus are nongeneric. A counterexample to this conjecture is in Appendix F.

3.4 Mean Time to Herd

We illustrate the power of our framework by answering an important question: How long is it until a herd starts? To this end, let ℓ_0 lie in the *communicating basin* $\mathcal{C} \subseteq [0, \infty)$, the smallest interval that $\langle \ell_n \rangle$ cannot exit if it starts at ℓ_0. Inside $\mathcal{C}$ is at least one action absorbing basin.[18] Relabel WLOG the actions that can be taken in $\mathcal{C}$ as $1, 2, \ldots, M$. Thus, $\mathcal{C} = I_1 \cup \cdots \cup I_M$ (in reverse order). Call a string of identical actions that eventually comes to an end a *temporary herd*; therefore, a putative herd is either a herd or temporary herd. It suffices that temporary herds are of uniformly bounded expected duration, and that the entry rate into putative herds is boundedly positive after a fixed number of periods.

Lemma 8 (Putative Herds) *There is* $\varepsilon^* > 0$ *and* $n^* < \infty$ *such that a putative herd starts in at most* n^* *steps with chance at least* ε^*.

The appendicized proof rules out $\langle \ell_n \rangle$ remaining long in $\cup_j \{I_j | J_j = \varnothing\}$, delaying the start of a putative herd. Intuitively, at least until it starts, at least two actions can occur, and

[17] This also follows analytically: If $\ell_n < 2u$, then $\ell_{n+1} = \varphi(1, \ell_n) = u/2 + 3\ell/4 < u/2 + 3u/2 = 2u$ too.

[18] It is possible for $\langle \ell_n \rangle$ to jump over action absorbing basin J_i if $\varphi(i, \cdot)$ is locally decreasing nearby.

with boundedly positive chance, $\langle \ell_n \rangle$ marches monotonically toward some I_m with $J_m \neq \varnothing$.

Let e_n be n's exit chance from a putative herd, and $E_n = (1-e_1)\cdots(1-e_n)$ the chance Mr. $1, 2, \ldots, n$ participate in it (since signals are conditionally i.i.d.). Then the chance of a herd is $E_\infty = F_1 = (1-e_1)(1-e_2)\cdots > 0$ exactly when $\sum e_n < \infty$.[19] Let b_n be n's exit chance from a herd conditional on its being temporary. The chance that a given herd at stage n is permanent is $F_n = (1-e_n)(1-e_{n+1})\cdots$, and so by Bayes rule, $b_n = e_n/(1-F_n)$.

Lemma 9 (Temporary Herds) *The expected length of a temporary herd is* $\displaystyle\sum_n \frac{E_n - E_\infty}{1 - E_\infty}$.

Since the acid test $\sum_n (E_n - E_\infty) < \infty$ can fail even when $E_\infty > 0$ by Lemma E.1(b), a good principle is: *How quickly a herd starts depends not on how quickly individuals become convinced of a state but on how slowly they can radically shift beliefs.* This leads us to

Theorem 5 (The Speed of Action Convergence: How Long Until the Herd?)
(a) With atoms at the edge of $\operatorname{supp}(F)$,[20] *herds begin in finite mean time.*
Let $F^H(p) = c(p-\underline{b})^\gamma + O((p-\underline{b})^{\gamma+1})$ *and* $F^L(p) = 1 - d(\bar{b}-p)^\delta + O((\bar{b}-p)^{\delta+1})$ *near* $\underline{b}$ *and* $\bar{b}$.
(b) With bounded beliefs, herds begin in finite mean time iff $\gamma, \delta < 2$.
(c) With unbounded beliefs, herds begin in finite mean time in state H *iff* $\gamma < 2$ *and* $\delta > 2$.

So for bounded beliefs, extreme signals in favor of the true state must have an unbounded density, to ensure a correct herd in finite time. With unbounded beliefs, our proof explains how if there are many extreme signals available in state H, and conversely few in L, or $\gamma < 2$ and $\delta > 2$, then respectively, *temporary* herds on a_M and a_1 end quickly enough that the eventual herd on a_M starts in finite time in state H — but then herds start in infinite time in state L! So, herds cannot start in finite mean time in both states. *With unbounded beliefs, complete learning eventually obtains, but is expected (ex ante) to take forever!*

• Unbounded Beliefs Example Cont'd. With $\gamma = \delta = 2$, the mean time to herd is infinite, since a wrong temporary herd on a_1 in state H ends in *infinite* mean time.

Theorem 6 (Time Discounting) *The expected δ-discounted fraction of time until a herd starts vanishes as $\delta \to 1$, provided each F^s has polynomial weight in the tails.*

Proof: If a herd starts in period n, the discounted fraction of time until then is $1-\delta^n$. The expected discounted time to finishing any temporary herd is $(1-\delta)\sum_1^\infty E_n\delta^n$. The proof of Theorem 5 makes clear that $E_n = O(1/n^a)$ for some $a > 0$ if F^L has polynomial weight in its tail. The result follows from $\liminf_{x\to 1^-}(1-x)\sum_1^\infty x^n/n^a = 0$ for all $0 < a < 1$.[21]

[19] For $\exp(-x) \geq 1 - x$ yields $\exp(-\sum_k e_k) > (1-e_1)(1-e_2)\cdots$, while an induction argument on the number N of terms, followed by $N \to \infty$, proves $(1-e_1)(1-e_2)\cdots \geq 1 - (\sum_k e_k)$ if $e_i \in [0,1)$ for all i.

[20] i.e., $0 < F^H(\underline{b}) < F^H(\bar{b}^-) < 1$. With bounded beliefs, such extreme signals aren't perfectly revealing.

[21] Michael Larsen (U. Pennsylvania) gave us a quick proof: Let $\eta > 0$. As the early terms of $S(x) = \sum_1^\infty x^n/n^a$ don't affect our limit, assume $1/n^a < \eta$; then $(1-x)S(x) < (1-x)\eta(1+x+x^2+\cdots) = \eta$.

4. NOISE

The pivotal role played by the overturning principle is somewhat unsettling. The large weight accorded isolated actions is both economically implausible and theoretically unappealing;[22] therefore, we now add 'noise' to the system: Some percentage of individuals either by design ('craziness') or mistake ('trembling') do not choose wisely. We assume that being noisy is not public information, and that this trait is distributed independently across individuals. Since all actions are expected to occur, none can have drastic effects.[23]

Below, we address the simplest case of craziness noise, and appendicize the discussion of trembling noise. So with chance κ_m, Mr. n chooses action a_m, irrespective of history.[24] We assume a positive fraction $\kappa = 1 - \sum_{m=1}^{M} \kappa_m > 0$ of *rational* individuals, each of whom take action a_m with chance $\rho(m|s, \ell) = F^s(\bar{p}_m(\ell)) - F^s(\bar{p}_{m-1}(\ell))$. So action a_m is now taken with chance $\psi(m|s, \ell)$ yielding likelihood continuation $\varphi(m, \ell)$, where

$$\psi(m|s, \ell) = \kappa_m + \kappa\rho(m|s, \ell) \tag{6}$$

$$\varphi(m, \ell) = \ell\psi(m|L, \ell)/\psi(m|H, \ell) \tag{7}$$

4.1 Convergence of Beliefs

Since $\ell = \sum_{m=1}^{M} \psi(m|H, \ell)\varphi(m, \ell)$, $\langle \ell_n \rangle$ is still a martingale in state H, with almost sure finite limit ℓ_∞. The interval structure of $J_1, \ldots, J_M$ (Lemma 6) is also still valid. So $\rho(m|H, \ell) = \rho(m|L, \ell) = 1$ for some m implies $\rho(m'|H, \ell) = \rho(m'|L, \ell) = 0$ for all $m' \neq m$, and thus *rational* individuals take action a_m almost surely. Contrary actions will perforce be adjudged ex post as noisy, and will simply be ignored. The next result asserts that *Theorems 1 and 2 are robust to noise.* Statistically constant behavior of noisy individuals doesn't affect long run learning by rational individuals, as it can be filtered out.

Theorem 7 (Long Run Learning with Noise) *In the noisy model, $\ell_n \to \ell_\infty$ in state H, where* $\text{supp}(\ell_\infty) \subseteq [0, \infty)$. *With bounded beliefs, $\ell_\infty \in J$ almost surely, and $\ell_\infty \in J_M$ with chance less than* 1 *if $\ell_0 \notin J_M$. With unbounded beliefs, $\ell_\infty = 0$ almost surely.*

Proof: As all ψ are boundedly positive, we need only check stationarity of $\hat{\ell} \in \text{supp}(\ell_\infty)$, or $\varphi(m|H, \hat{\ell}) = \hat{\ell}$.

$$\varphi(m|H, \hat{\ell}) = \hat{\ell}\frac{\kappa_m + \kappa\rho(m|L, \hat{\ell})}{\kappa_m + \kappa\rho(m|H, \hat{\ell})} = \hat{\ell}$$

[22] Note that the herd fragility discussed in BHW refers instead to the release of public information.

[23] One might think that informational free-riders constitute a different form of noise — i.e. where some individuals receive no private signal, and simply free-ride off the public information. But they do not require special treatment: Simply let F^H and F^L each have an atom at 1/2.

[24] Equivalently, every action is commonly misperceived by all successors.

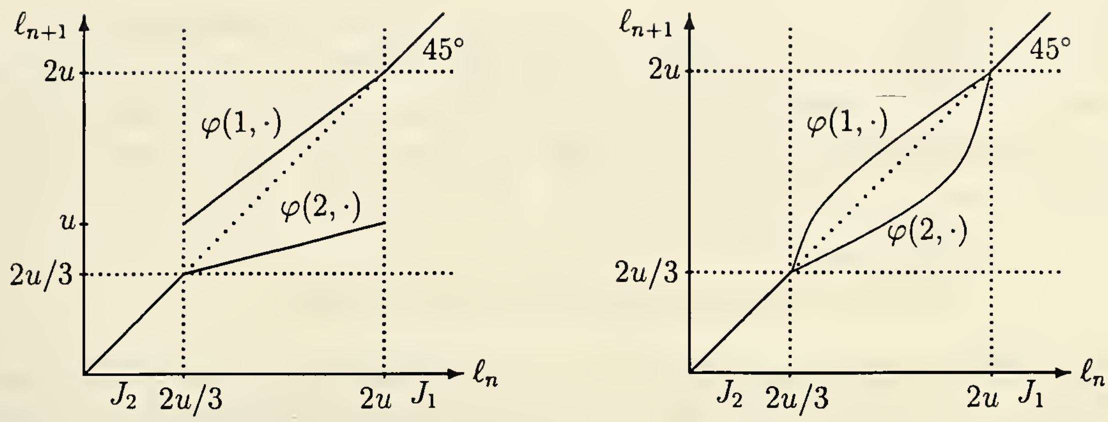

Figure 5: **Continuations.** Here, we juxtapose the two continuation likelihood functions for the BOUNDED BELIEFS EXAMPLE — first for only rational individuals, then with some crazy types. We see in the right graph that the discontinuity vanishes, corresponding to the failure of the overturning principle.

and so $\rho(m|H,\hat{\ell}) = \rho(m|L,\hat{\ell})$. That $\hat{\ell} \in J$ now ensues from the proof of Theorem 1, but is *conceptually easier* here (one case and not two). Finally, mimic the proof of Theorem 2. □

4.2 Convergence of Actions: The Failure of the Overturning Principle

While long-run learning is unaffected by constant background noise, herding is not so resilient: Noisy individuals — like cats — simply do not herd. It is natural to ask about *rational herds*: Do the rational agents herd? The failure of the overturning principle offers a special challenge, as herd violations only minimally impact public beliefs, being deemed irrational acts. This severs our clean implication: belief convergence ⇒ action conformity.

Let us illustrate how Lemma 7 fails with noise. The left panel in figure 5 depicts our running BOUNDED BELIEFS EXAMPLE, where one can see that for ℓ near J_m, $|\varphi(m,\ell)-\ell|$ is small *and* $|\varphi(m',\ell)-\ell|$ is bounded away from 0 for all $m' \neq m$. Observe how each stationary point ℓ^* is only fixed under one continuation, since contrary actions can't occur there: By the overturning principle, $\varphi(m',\ell)-\ell \approx 0$ only for $m'=m$ when ℓ near J_m. Introduction of a small amount of noise effects a remarkable sea change in figure 5, as seen in the right panel of figure 5. For if all actions occur with boundedly positive chance, and ρ and ψ are continuous, $\varphi(m',\ell)-\ell = 0$ for all m' at the fixed point, by Theorem B.1. For instance,

$$\varphi(m',\ell)-\ell = \ell\,\frac{\kappa\left[\rho(m'|L,\ell)-\rho(m'|H,\ell)\right]}{\kappa\rho(m'|H,\ell)+\kappa_{m'}} = \frac{\beta(m',\ell)-\ell}{1+\kappa_{m'}/(\kappa\rho(m'|H,\ell))} \tag{8}$$

where $\beta(m',\ell) \equiv \ell\rho(m'|L,\ell)/\rho(m'|H,\ell)$ is the old noiseless continuation. So, with noise, $|\varphi(m',\ell)-\ell|$ vanishes for all m', and any $\ell \in J_m$: Indeed, $\varphi(m,\ell)=\ell$ since the numerator is 0, and for $m' \neq m$, $\varphi(m',\ell)=\ell$ because the denominator is infinite (as $\rho(m'|H,\ell)=0$).

That $\varphi(\tilde{m}, \hat{\ell}) - \hat{\ell} = 0$ at a fixed point $\hat{\ell}$ allows us to make simple deductions about the rate of convergence for this system. Appendix C develops a theory of stability for stochastic dynamical systems like this one. Given functions φ and ψ that are, respectively, C^1 and continuous at the extremes, Corollary C.1 asserts that

$$\theta \equiv \prod_{m=1}^{M} |\varphi_\ell(m, \hat{\ell})|^{\psi(m|H,\hat{\ell})} = \text{rate that } \langle \ell_n \rangle \text{ converges to fixed point } \hat{\ell}$$

i.e., the frequency-weighted geometric mean of the continuation derivatives. The proof of the next result contains the core essence of our most broadly applicable theoretical finding.

Lemma 10 (Rate of Belief Convergence) *Assume that F^H, F^L have C^1 tails,*[25] *with derivatives f^H and f^L. If private beliefs are bounded and $f^H(\underline{b}), f^L(\bar{b}) > 0$, then $\theta < 1$.*

Proof: Clearly, $\sum_{m=1}^{M} \psi(m|H, \ell) \equiv 1$, while the martingale property yields the identity $\ell \equiv \sum_{m=1}^{M} \psi(m|H, \ell)\varphi(m, \ell)$. If all functions are differentiable, we then have

$$1 = \sum_{m=1}^{M} \psi(m|H, \ell)\varphi_\ell(m, \ell) + \sum_{m=1}^{M} \psi_\ell(m|\ell)\varphi(m, \ell) \tag{9}$$

At a fixed point $\hat{\ell} \in J_m$, the second sum in (9) vanishes, since $\varphi(m', \hat{\ell}) = \hat{\ell}$ for all m', and $\sum_{m'=1}^{M} \psi_\ell(m'|H, \ell) = 0$. So the arithmetic mean of the derivatives $\langle \varphi_\ell(m', \hat{\ell}) \rangle$ is 1. If any are negative, then $\langle \ell_n \rangle$ eventually jumps into J_m (convergence rate 0). If all are non-negative, the *arithmetic mean-geometric mean inequality* neatly proves that convergence occurs at a rate $\theta \leq 1$, with equality iff *all* derivatives are 1. Let $\beta(m', \ell) \equiv \ell\rho(m'|L, \ell)/\rho(m'|H, \ell)$ be the noiseless continuation. Then $\varphi_\ell(m, \hat{\ell}) = (\kappa_m + \kappa\beta_\ell(m, \hat{\ell}))/(\kappa_m + \kappa)$, which is strictly less than 1: $\beta_\ell(m, \hat{\ell}) = 1 + \hat{\ell}[f^H(\underline{b}) - f^L(\underline{b})] < 1$, given Lemma 1(a), bounded beliefs ($\underline{b} > 0$), and informative beliefs ($\underline{b} < 1/2$). So $\theta < 1$. □

Whether a rational herd eventually starts turns on the speed of convergence of the public likelihood ratios $\langle \ell_n \rangle$. Suppose we have a limit cascade $\ell_n \to \hat{\ell} \in J_m$. When can we rule out an infinite subsequence of rational 'herd violators', whose private beliefs counteract the public belief? In light of the (first) Borel-Cantelli Lemma,[26] this occurs with zero chance provided $\sum_{n=1}^{\infty} (1 - \rho(m|H, \ell_n)) < \infty$ for almost surely all $\langle \ell_n \rangle$.[27] With bounded beliefs, this inequality holds if F^H and F^L have C^1 tails, with $f^H(\underline{b}) \neq 0$ and $f^L(\bar{b}) \neq 0$. For then the convergence of $\ell_n \to 0$, and thus $1 - \rho(m|H, \ell_n) \to 0$ is exponentially fast, given the positive tail density. So, Lemma 10 implies that

[25] Namely, each is C^1 in some open neighborhood of $\bar{b}$ and $\underline{b}$.
[26] We mean the non-standard conditional version of the Lemma, e.g. Corollary 5.29 of Breiman (1968).
[27] Since not everyone is rational, this inequality in fact assumes a worst case.

Theorem 8 (Rational Herds) *Assume that private beliefs are bounded, and that F^H and F^L have C^1 tails. Then rational herds must arise.*

Lemma 10 doesn't apply to unbounded beliefs, and the slow rate of convergence may not suffice for the Borel-Cantelli Lemma above. This must remain a (hard) open question.

5. MULTIPLE INDIVIDUAL TYPES

We now venture into new territory and investigate a richer and often more realistic model with heterogeneous preferences, or *types*. A model with multiple but observable types is informationally equivalent to the single preference world. So just as with the noise formulation, we assume that types are private information. Still one can learn from history by comparing the proportions choosing each action with the known type frequencies. This inference intuitively ought to be fruitful, barring nongenericities, like equal frequencies and exactly opposed vNM preferences. A curious new twist is then introduced — *confounded learning*: Dynamics may well converge upon an outcome in which each action is taken with the same probability in all states. This twin pathological learning outcome is by one measure more robust than wrong herds, as it may arise even with unbounded private beliefs. It also turns out to be an all-round fundamentally different economic beast.

5.1 The Model and an Overview

Given are finitely many *types* $t = 1, \ldots, T$, where t determines both vNM preferences over the given action set $a_1, \ldots, a_M$, as well as one's private signal distribution. Let λ^t denote the known proportion of type t, and assume the types are i.i.d. across individuals. Since an one's type is private information, this formulation is identical to noise if all but one type has state-independent preferences. *In contrast to noise, all decisions may depend on (and thus be informative of) private signals.* This radically changes the analysis.

As before, $\langle \ell_n \rangle$ is a convergent martingale in state H. Each type still employs a posterior belief threshold rule when choosing an action; however, they will generally disagree on the desirability of the actions. Let $\rho^t(m|s, \ell)$ be the chance that someone of type t will choose action a_m, given state $s \in \{H, L\}$ and public likelihood ℓ. Next, call the set of ℓ yielding $\rho^t(m|H, \ell) = \rho^t(m|L, \ell) = 1$ the action absorbing basin J^t_m. Similarly, let $J^t = J^t_1 \cup J^t_2 \cup \cdots \cup J^t_M$ for each t, so that $\bar{J} = J^1 \cap J^2 \cap \cdots \cap J^T$ is the overall *absorbed set*: all ℓ where each type finds himself in an absorbing basin. So if just one type has unbounded private signals, then $\bar{J} = \{0, \infty\}$. If $\ell_n \in \bar{J}$, every action is chosen irrespective of private information, and thus provides no information. Conversely, if $\ell_n \notin \bar{J}$, some type's action is not a foregone conclusion. But unlike before, we cannot conclude that $\ell_\infty \in \bar{J}$ almost surely, as there are potentially more limit points.

In a limit cascade, a *type-specific herd* may arise: Everyone of the same type will take the same action. But if types differ in their vNM preferences, unless all should take one and the same action, an overturning principle doesn't work. As with noise, we must apply the speed of convergence reasoning from section 4.2 to conclude that herds must arise.

The dynamics of the likelihood ratio are described in the usual notation by (7) and

$$\psi(m|s,\ell) = \sum_{t=1}^{T} \lambda^t \rho^t(m|s,\ell) \tag{10}$$

If ψ is continuous in ℓ, then any fixed points must satisfy (A-4): $\psi(m|H,\ell) = 0$ or $\varphi(m,\ell) = \ell$. By Theorems 2 and 7, all solutions to this criterion lie in the absorbed set, but with multiple types, this is no longer true. We call its solutions ℓ^* outside $\bar{J}$ *confounding outcomes.* Such may exist since $\varphi(m,\ell^*) = \ell^*$ when actions are taken with equal chance in the two states:

$$\psi(m|L,\ell^*) = \psi(m|H,\ell^*) \quad \forall m \tag{11}$$

Crucially, history is most certainly *not* totally informative at ℓ^*. For if it were, agents would ignore it and their decisions then would generally be informative. Rather history has become *precisely so informative* as to choke off any *additional* inferences. The distinction with a cascade is both compelling and sweet. Both private signals *and* history affect decisions in a confounding outcome, whereas in a limit cascade, history becomes totally decisive, and private signals wholly inconsequential: Yet both are pathological outcomes.

To verify that *confounded learning* can occur, or $\ell_n \to \ell^*$, it suffices to check the generic inequality $\varphi_\ell(1,\ell^*) \neq \varphi_\ell(2,\ell^*)$, *with both terms positive.* For as in the proof of Lemma 10, this will imply the *local stability* of the point ℓ^*. As defined in Appendix C, this means that if $\langle \ell_n \rangle$ starts close enough to ℓ^*, then $\ell_n \to \ell^*$ with positive probability: Since $\langle \ell_n \rangle$ cannot cycle, limit cascades and confounded learning are the only possible 'pathological' (incomplete learning) outcomes in finite-action observational learning models.

5.2 Examples of Confounded Learning

We show that confounding outcomes may exist, and confounded learning can occur.

In the examples, we have $M = 2$ actions and $T = 2$ types. Types differ in their preferences, but not their signal distributions. Type U is 'usual', preferring action a_2 in state H, a_1 in state L, and a_1 for private beliefs below $\bar{p}^U(\ell) = \ell/(u+\ell)$. The preferences

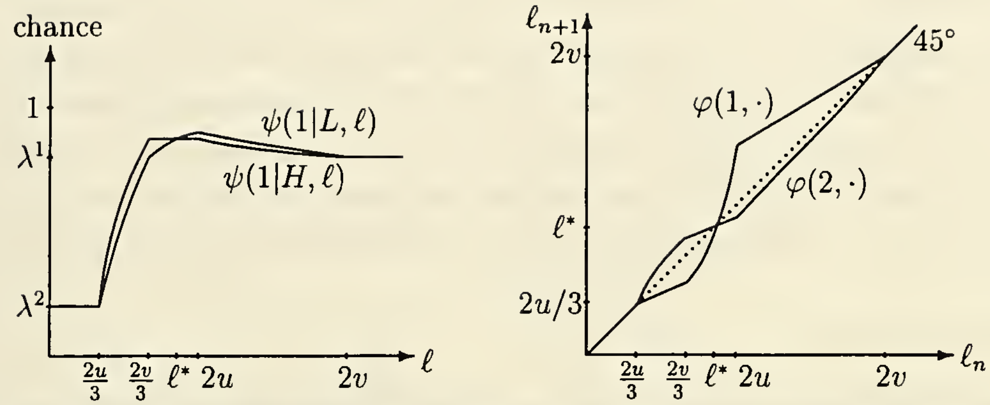

Figure 6: **Confounded Learning.** Based on our Bounded Beliefs Example, with $\lambda^U = 4/5$, $u = v/2$. In the left graph, the curves $\psi(1|H,\ell)$ and $\psi(1|L,\ell)$ cross at the confounding outcome ℓ^*, where no additional decisions are informative. At ℓ^*, 7/8 choose action a_1, and counterintuitively 7/8 lies outside the convex hull of λ^V and λ^U — eg. in the introductory driving example, more than 70% of cars may merge right in a confounding outcome. The right graph depicts continuation likelihood dynamics.

of type V are opposite: action a_1 always pays zero, while a_2 yields payoff v in state H and -1 in state L. He thus chooses a_1 for private beliefs *above* the threshold $\bar{p}^V(\ell) = \ell/(v+\ell)$. WLOG, $v \geq u > 0$, so *vNM preferences are not exactly opposed if* $v > u$.

• Bounded Beliefs Example Cont'd. The transition probabilities for type U are now $\rho^U(1|H,\ell) = (3\ell-2u)/2\ell$ and $\rho^U(1|L,\ell) = (3\ell-2u)(3\ell+2u)/8\ell^2$, where $\ell \in (2u/3, 2u)$. For type V, $\rho^V(1|H,\ell) = (2v-\ell)/2\ell$ and $\rho^V(1|L,\ell) = (2v+\ell)(2v-\ell)/8\ell^2$, where $\ell \in (2v/3, 2v)$. With bounded beliefs, the absorbing basins complicate the dynamics. The two types take action 2 with certainty in the intervals $J_2^U = [0, 2u/3]$ and $J_2^V = [2v, \infty)$, respectively. If these overlap, then dynamics nonessentially differ from those in section 3 because only one of the types ever makes an informative choice for any ℓ, thus precluding confounded learning. The same remark holds if $J_1^U = [2u, \infty)$ and $J_1^V = [0, 2v/3]$ overlap.

For $u \leq v$, no overlap arises if $2v/3 < u$. So consider the dynamics for $\ell \in (2v/3, 2u)$:

$$\psi(1|H,\ell) = \lambda^U \frac{3\ell-2u}{2\ell} + \lambda^V \frac{2v-\ell}{2\ell} \quad \text{and} \quad \psi(1|L,\ell) = \lambda^U \frac{(3\ell-2u)(3\ell+2u)}{8\ell^2} + \lambda^V \frac{(2v-\ell)(2v+\ell)}{8\ell^2}$$

by (10). Figure 6 graphs these functions. We can rewrite (11) for a confounding outcome as

$$\frac{\lambda^U}{\lambda^V} = \frac{(2v-\ell)(3\ell-2v)}{(2u-\ell)(3\ell-2u)} \equiv \xi(\ell)$$

If $u > v$ then ξ maps $(2v/3, 2u)$ onto $(0, \infty)$, and so a confounding outcome exists for any λ^U, λ^V. Next, $\varphi_\ell(1, \ell^*) = (3\lambda^U/4 + \lambda^V/4)/\psi(1|H, \ell^*)$ generically differs from $\varphi_\ell(2, \ell^*) = (\lambda^U/4 + 3\lambda^V/4)/\psi(2|H, \ell^*)$ and both are positive. So confounded learning can occur.

Since the functions φ are increasing, the system either starts in a cascade in $[0, 2u/3]$

or $[2v, \infty)$, or starts and thus is trapped in $[2u/3, \ell^*]$ or $[\ell^*, 2v]$. As all probability mass is eventually concentrated at the endpoints, there can only arise only a wrong herd or confounded learning if $\ell_0 \in [\ell^*, 2v]$, and a correct herd or confounded learning if $\ell_0 \in [2u/3, \ell^*]$. Just as in the proof of Theorem 2, since $\langle \ell_n \rangle$ is a bounded martingale and $E[\ell_\infty] = \ell_0$, both possible outcomes have positive probability in each case. More generally, with only a single confounding outcome, limit cascades must occur with positive chance.

5.3 The Basic Theory

⋆ **Theoretical Robustness.** We have shown by example that confounding outcomes can exist, and confounded learning may arise. We now finish outlining the theory.

Theorem 9 (Confounded Learning) *Assume there are $T \geq 2$ types.*
(a) If the private signal distribution is atomless, then for nondegenerate specifications of preference and type proportions, confounding outcomes exist, and confounded learning obtains with positive chance. No one confounding outcome must occur with bounded beliefs.
(b) Generically, at any confounding outcome only two actions are taken.
(c) Confounded learning is generically robust to the addition of craziness noise.
(d) When confounded learning does not occur, a limit cascade arises, i.e. $\ell_n \to \hat{\ell} \in \bar{J}$, and is almost surely correct with unbounded beliefs.
(e) If $M > 2$ with unbounded beliefs, or if the private signal distributions are discrete, then generically no confounding outcome exists.

Proof: The first point has been addressed by the example, which is not nongeneric.

(*b*) Let us consider the equations that a confounding outcome ℓ^* must solve. First, with bounded beliefs some actions may not occur at all at ℓ^*. Assume that $M_0 \leq M$ actions are taken with positive probability at ℓ^*. Equation (11) reduces in (à la Walras' Law) to $M_0 - 1$ independent equations, since $1 = \sum_{m=1}^{M_0} \psi(m_i|H,\ell) = \sum_{i=1}^{M_0} \psi(m_i|L,\ell) = 1$. Generically, $M_0 - 1$ equation in one unknown ℓ can only be solved when $M_0 = 2$.

(*c*) Equality (11) still obtains if state independent noise is added to both sides. The sign of $\tilde{\varphi}_\ell(i,\ell)$, as it equals

$$\frac{d}{d\ell}\frac{\tilde{\psi}^L(i|\ell)}{\tilde{\psi}^H(i|\ell)} = \frac{\tilde{\psi}^H(i|\ell)\tilde{\psi}^L_\ell(i|\ell) - \tilde{\psi}^L(i|\ell)\tilde{\psi}^H_\ell(i|\ell)}{\tilde{\psi}^H(i|\ell)^2} = \frac{\tilde{\psi}^L_\ell(i|\ell) - \tilde{\psi}^H_\ell(i|\ell)}{\tilde{\psi}^H(i|\ell)} = \frac{\alpha\psi^L_\ell(i|\ell) - \alpha\psi^H_\ell(i|\ell)}{\alpha\psi^H(i|\ell) + (1-\alpha)\kappa_i}$$

when (11) holds. Finally, only for nongeneric noise will $\tilde{\varphi}_\ell(1,\ell) = \tilde{\varphi}_\ell(2,\ell)$.

(*d*) Any $\hat{\ell} \in \ell_\infty$ must satisfy (A-4). This precludes all but limit cascades, where $\psi(m|\hat{\ell}) = 1$ for some action a_m, and confounding outcomes, where $\psi(m|\hat{\ell}) > 0$ for at least two actions a_m. The unbounded beliefs argument is by now standard.

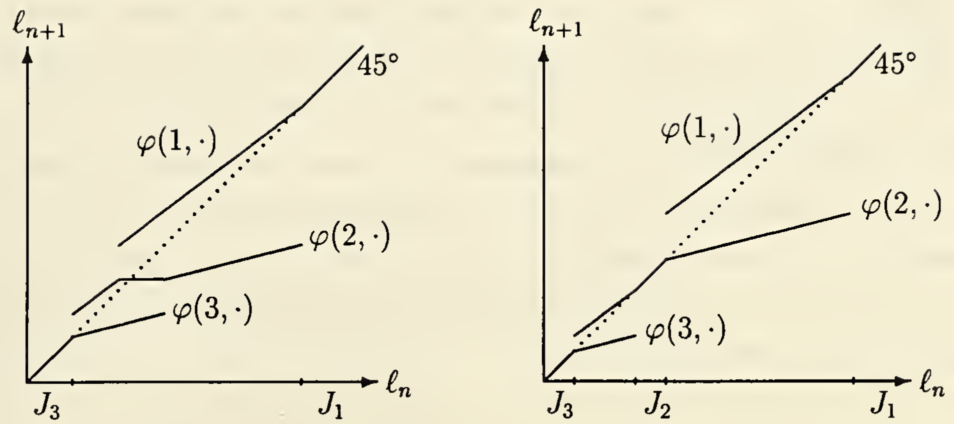

Figure 7: **Continuations and Absorbing Basins.** This is based on the Bounded Beliefs Example, but here with one insurance and two extreme actions. In the first graph, preferences are such that there is no interior basin, but in the second graph there is an interior basin. Observe in the second graph, that with the basin, for any value of ℓ at most two actions are in play. This is the kind of example which can be used to construct models with $M > 2$ and confounded learning.

(*e*) If one of $M > 2$ types has unbounded private beliefs, then for no $\ell \in (0, \infty)$ are only two actions taken with positive chance. So confounding outcomes only appear in degenerate models. Next,

$$\frac{\lambda^U}{\lambda^V} = \frac{\rho^V(1|H,\ell) - \rho^V(1|L,\ell)}{\rho^U(1|L,\ell) - \rho^U(1|H,\ell)}$$

is a reformulation of (11), and so if F^H and F^L are discrete, then the RHS will only assume a countable number of values, and confounding outcomes will generically not exist. □

Remarks. **1.** While only two actions will occur with positive chance at generic confounding outcomes, *generic models with $M > 2$ actions can still have confounding outcomes.* For with bounded beliefs, only two actions may well be taken over a range of ℓ. This will happen when there are absorbing basins for the insurance actions, as in figure 7.

2. Being a fixed point, we can say little about the *uniqueness* of confounding outcomes ℓ^* — except that with discrete distributions, they are not unique when they exist (for an interval around ℓ^* will satisfy (11) because F^H and F^L are locally constant).

3. Our example posited two types with different preference orderings over actions. If only vNM preferences differ, then one can show that no confounding outcome can exist.

⋆ **Economic Importance.** We now return to our introductory example, and shed some light on exactly *when* one should expect to see our new confounding phenomenon.

• The Driving Example Revisited with Unbounded Beliefs. Posit that Houston (type U) drivers should merge right (action a_1) in state H, left (action a_2) in state L, with the reverse true for Dallas (type V) drivers. Going to the wrong city yields zero always. Getting home by the right lane a_1 is preferred, as it has fewer potholes. The payoff vector of the Houston-bound is $(u, 0)$ in state H and $(0, 1)$ in state L; for Dallas drivers, it is $(0, 1)$ and $(v, 0)$.[28] The proportions are $\lambda^U = .7$ and $\lambda^V = .3$.

Claim *With a differentiable signal distribution and unbounded beliefs, a confounding point exists if preferences are more disparate than type frequencies:* $u/v < \lambda^U/\lambda^V < v/u$.

Proof: We show that $\psi(1|H, \ell)$ lies below $\psi(1|L, \ell)$ near $\ell = 0$ and above it near $\ell = \infty$, and therefore by continuity the transition probability curves must at some ℓ coincide. Differentiating

$$\psi(1|s, \ell) = \lambda^U F^s(\ell/(u + \ell)) + \lambda^V[1 - F^s(\ell/(v + \ell))]$$

near $\ell = 0$ yields $\psi_\ell(1|s, 0) = f^s(0)[\lambda^U/u - \lambda^V/v]$. Since $f^L(0) > f^H(0)$ by Lemma A.1, $\psi_\ell(1|H, 0) > \psi_\ell(1|L, 0)$ iff $\lambda^U/\lambda^V > u/v$. The reverse inequality near $\ell = \infty$ is similar. □

The intuition for the existence of confounding points is simple, and common to bounded beliefs. Clearly, if nearly all drivers are Houston-bound, then it is uniformly true that more will merge right in state H than in state L. But otherwise, if preference differences dominate, then what matters near extreme beliefs is how many contrarians *of either type* exist. By Lemma 1, there are many more (infinitely more with unbounded beliefs) doctrinaire contrarians when they are right than when they are wrong. Thus, more will merge right (resp. left) for public beliefs near $p = \underline{b}$ (resp. $p = \bar{b}$) in state H than state L (conversely).

6. CONCLUSION

This paper has explored and expanded upon the so-called herding literature. We hope our analysis underscores a rich theory that ensues from attention to likelihood ratios and their conditional martingale property in theoretical learning models.

⋆ **Related Literature.** We think it noteworthy that Milgrom's (1979) convergence theorem for competitive bidding also turns on the bounded-unbounded signal knife-edge.

We must underscore that our complete learning results are *in no way* related to Lee (1993), where a rich continuous action space effectively allows for a one-to-one map of signals ↔ actions. As roughly foreseen by Banerjee (1992), this precludes herding in much the same way as statistical decision problems do. A key touchstone of herding is the coarse inference of predecessors' signals as they pass through a lumpy filter (like action choices). [29] *there is no loss of generality in our model with two signal values compared to a model with many signal values.*

[28] By a payoff renormalization, this is equivalent to our standard payoff structure.

[29] To quote from Lee (on page 397): *From the standpoint of information revelation, a sparse action*

⋆ **Lessons for the Experimentation Literature.** Our companion paper Smith and Sørensen (1996a), or SS1, presents a different slant on this field, drawing a formal parallel between bad herds and incomplete learning in optimal experimentation: observational learning $\leftrightarrow$ experimentation, as actions $\leftrightarrow$ signals, as belief thresholds $\leftrightarrow$ actions. SS1 also makes it clear that even if a patient social planner (unable to observe signals) could control the action choices, he too would succumb to bad herds, given bounded private beliefs.

In light of SS1, that more than two actions generically cannot be taken at a confounding outcome speaks to a general property of optimal experimentation models. For instance, SS1 cites published examples of incomplete learning, and all have binary signals.[30] And in fact, given exogenous signal functions and just two states, only nongeneric models yield beliefs for which more than two signals optimally arise with equal chance in either state.

Our stochastic stability condition also may be a useful tool for optimal experimentation problems. Since informational herding is formally equivalent to single person learning, we suspect that by focusing on the the stochastic dynamic process of likelihood ratios rather than the intricacies of dynamic optimization, our stability program offers significant hope. For example, in the popular two-state, continuous-action learning models, it suffices to verify that near a fixed point, the continuation posterior does not always have the same slope as a function of the prior belief after each observed signal. This condition clearly has nothing to do with the particulars of the actual dynamic optimization. Indeed, this paper has afforded this simple insight precisely because the optimization problem is so trivial.

⋆ **Where Do We Go Now?** Smith and Sørensen (1996b) relaxes the key assumption that one can perfectly observe the ordered action history. While this is yet another (more plausible) reason for why isolated contrary actions might have little effect, we are motivated by deeper concerns. For absent martingales, the resulting analysis is radically different that standard rational learning theory, and forces one to think (à la Blackwell) abstractly about information, and to delve deeply into the theory of urn processes.

A. ON BAYESIAN UPDATING OF DIVERSE SIGNALS

The set-up (measures μ^s and distribution functions F^s, $s = H, L$) is taken from §2.1-2.2. The characterization of the Radon-Nikodym derivative of F^H, F^L in Lemma 1 implies

---

set provides little means to convey the private information. Consequently the infinite sequence of private signals adds little to the updating of the posterior distribution and the whole sequence of individuals may end up choosing the wrong action.

[30] An additional one that has come to our attention is Kihlstrom, Mirman, and Postlewaite (1984).

Lemma A.1 (Extreme Beliefs are Informative)
(a) $F^H(p) \leq pF^L(p)/(1-p)$ *holds for all* $p \in (0,1)$*, and is strict when* $F^L(p) > 0$.
(b) The ratio F^H/F^L *is weakly increasing, and strictly so on* $(\underline{b}, \bar{b}]$.

Proof: Since $f = dF^H/dF^L$ is a strictly increasing function, we have

$$F^H(p) = \int_{r \leq p} f(r)\, dF^L(r) < f(p) \int_{r \leq p} dF^L(r) = pF^L(p)/(1-p) \tag{A-2}$$

for any p with $F^L(p) > 0$. Thus, whenever $F^L(p) > F^L(q) > 0$, we have

$$F^H(p) - F^H(q) = \int_q^p f(r)\, dF^L(r) > [F^L(p) - F^L(q)] f(q) > [F^L(p) - F^L(q)] F^H(q)/F^L(q)$$

where we have used (A-2). It immediately follows that $F^H(p)/F^L(p) > F^H(q)/F^L(q)$. □

B. FIXED POINTS OF MARKOV-MARTINGALE SYSTEMS

The general framework that we introduce here includes, but is not confined to, the evolution of the likelihood ratio $\langle \ell_n \rangle$ over time viewed as a stochastic difference equation.[31]

Given is a finite set $\mathcal{M}$, and Borel measurable functions $\varphi(\cdot\,, \cdot) : \mathcal{M} \times \mathbb{R}_+ \to \mathbb{R}_+$, and $\psi(\cdot\,|\,\cdot) : \mathcal{M} \times \mathbb{R}_+ \to [0,1]$ satisfying:

- $\psi(\cdot\,|\,\ell)$ is also a probability measure on $\mathcal{M}$ for all $\ell \in \mathbb{R}_+$, or $\sum_{m \in \mathcal{M}} \psi(m|\ell) = 1$.
- ϕ and ψ jointly satisfy the following 'martingale property' for all $\ell \in \mathbb{R}_+$:

$$\sum_{m \in \mathcal{M}} \psi(m|\ell)\ \varphi(m, \ell) = \ell \tag{A-3}$$

For our application, $\psi(m|\ell)$ is the chance that the next agent takes action a_m when faced with likelihood ℓ, and $\varphi(m, \ell)$ is the resulting continuation likelihood ratio.

Next, for any B in the Borel σ-algebra $\mathcal{B}$ on $\mathbb{R}_+ = [0, \infty)$, define a transition probability $P : \mathbb{R}_+ \times \mathcal{B} \to [0,1]$:

$$P(\ell, B) = \sum_{m \,|\, \varphi(m,\ell) \in B} \psi(m|\ell)$$

Let $\langle \ell_n \rangle_{n=1}^{\infty}$ be a Markov process[32] with transition from $\ell_n \mapsto \ell_{n+1}$ governed by P, and

[31] Arthur, Ermoliev, and Kaniovski (1986) consider a stochastic system with a seemingly similar structure — namely, a 'generalized urn scheme'. But their approach, differs fundamentally from ours insofar as here it is of importance not only how many times a given action has occurred, but exactly *when* it occurred.

[32] Technically, $\ell_n : \Omega^H \to \mathbb{R}_+$ is a measurable Markov process on $(\Omega^H, \mathcal{E}^H, \nu^H)$, where $\mathcal{E}^H$ and $\mathcal{E}^L$ correspond to the restriction of sigma field $\mathcal{E}$ to Ω^H and Ω^L, respectively. Despite a countable state space, standard convergence results for discrete Markov chains have no bite, as states are in general transitory.

$E\ell_1 < \infty$. Then $\langle \ell_n \rangle$ is a martingale, true to the above casual label of (A-3):

$$E[\ell_{n+1}|\ell_1, \ldots, \ell_n] = E[\ell_{n+1}|\ell_n] = \int_{\mathbb{R}_+} tP(\ell_n, dt) = \sum_{m \in \mathcal{M}} \psi(m|\ell_n)\varphi(m, \ell_n) = \ell_n$$

By the Martingale Convergence Theorem, we have $\ell_n \to \ell_\infty = \lim_{n\to\infty} \ell_n \geq 0$ a.s.

Theorem B.1 (Stationarity) *If $\ell \mapsto \varphi(m, \ell)$ and $\ell \mapsto \psi(m|\ell)$ are continuous for all $m \in \mathcal{M}$, and $\ell_n \to \ell_\infty$ almost surely, then for all $\ell \in \operatorname{supp}(\ell_\infty)$, stationarity $P(\ell, \{\ell\}) = 1$ obtains, i.e.*

$$\psi(m|\ell) = 0 \quad or \quad \varphi(m, \ell) = \ell \qquad for\ all\ m \in \mathcal{M} \tag{A-4}$$

Indeed, $\langle \ell_n \rangle$ must also converge weakly (in distribution) to ℓ_∞. Since $\langle \ell_n, m_n \rangle$ is also a Markov chain, its limiting distribution is intuitively invariant for the transition P, as in Futia (1982). The a.s. convergence then implies that the invariant limit must be pointwise invariant. While Theorem B.1 admits a proof along these lines, the continuity assumptions are subtly hard-wired into the final stage of Futia's proof of this fact. As we wish to do away with continuity, we establish an even stronger result. That (A-4) is violated for m exactly when neither $\psi(m|\ell)$ nor $\varphi(m, \ell) - \ell$ is zero suggests

Theorem B.2 (Generalized Stationarity) *Assume that the open interval $I \subseteq \mathbb{R}_+$ has the property*

$$\exists \varepsilon > 0\ \forall \ell \in I\ \exists m \in \mathcal{M}: \ \psi(m|\ell) > \varepsilon, \ |\varphi(m, \ell) - \ell| > \varepsilon \tag{$\star$}$$

Then I cannot contain any point from the support of the limit, ℓ_∞.

Proof: Let I be an open interval satisfying ($\star$) for $\varepsilon > 0$, and suppose for a contradiction that there exists $\bar{\ell} \in I \cap \operatorname{supp}(\ell_\infty)$. Let $J = (\bar{\ell} - \varepsilon/2, \bar{\ell} + \varepsilon/2) \cap I$. By ($\star$), for all $\ell \in J$, there exists $m \in \mathcal{M}$ with $\psi(m|\ell) > \varepsilon$ and $\varphi(m, \ell) \notin J$. Since $\bar{\ell} \in \operatorname{supp}(\ell_\infty)$, $\ell_n \in J$ eventually with positive probability. But whenever $\ell_n \in J$, $\ell_{n+1} \notin J$ with chance at least ε. That is, the conditional chance that the process stays in J in the next period is at most $1 - \varepsilon$. So the process $\langle \ell_n \rangle$ almost surely eventually exits J. This contradicts the claim that with positive chance $\langle \ell_n \rangle$ is eventually in J. Hence, $\bar{\ell}$ cannot exist. □

Corollary *Assume that $\bar{\ell} \in \operatorname{supp}(\ell_\infty)$. Then for each $m \in \mathcal{M}$, either $\ell \mapsto \varphi(m, \ell)$ or $\ell \mapsto \psi(m|\ell)$ is discontinuous at $\bar{\ell}$, or the stationarity condition* (A-4) *obtains.*

Proof: If there is an m such that $\bar{\ell}$ does not satisfy (A-4) and both $\ell \mapsto \varphi(m, \ell)$ and $\ell \mapsto \psi(m|\ell)$ are continuous, then there is an open interval I around $\bar{\ell}$ in which $\psi(m|\ell)$ and $\varphi(m, \ell) - \ell$ are both bounded away from 0. This implies that ($\star$) obtains, and so Theorem B.2 yields an immediate contradiction. □

Finally, it is obvious that the corollary implies Theorem B.1.

C. STABLE STOCHASTIC DIFFERENCE EQUATIONS

In this appendix, we first develop a *global stability* criterion for *linear* stochastic difference equations. We then use it to derive a result on *local stability* of a *nonlinear* systems.[33] There is a small extant literature that treats models like ours at an abstract level. Bellman (1954) is cited as the first work in this area, Furstenberg (1963) is a classic article, and Kifer (1986) is a modern textbook. Ellison and Fudenberg (1995) have independently obtained related 'concrete' results, yet not as flexible — state-dependent transition chances is critical for our needs, and extrapolation to many dimensions an added benefit of our different approach.

Further, we add an analysis of rates of convergence. There is a small literature which treats the problems we discuss here, but we have not been able to exactly recognize our results. Overall, the literature is aimed at determining an *exact rate* $\bar{\theta}$ of convergence of $\langle \ell_n \rangle$ to a fixed point, such that $\bar{\theta}^{-n}\ell_n \to \zeta$ for some $\zeta \neq 0$. We, on the other hand, are happy to settle for an upper bound $\bar{\theta}$ such that $\theta^{-n}\ell_n \to 0$ for all $\theta > \bar{\theta}$. In the one-dimensional linear case we treat in Lemma C.1, we do obtain the exact rate, but, as we discuss later on, we do not determine the exact rate in higher dimensions.

⋆ **Linear Stochastic Difference Equations.** Fix $a, b \in \mathbb{R}$, $p \in [0,1]$, and let $\ell_0 \in \mathbb{R}$. Let

$$\ell_n = \begin{cases} a\ell_{n-1} & \text{with probability } p \\ b\ell_{n-1} & \text{with probability } 1-p \end{cases} \tag{A-5}$$

define a Markov process $\langle \ell_n \rangle$. This can be recast as: $\ell_n \equiv a^{y_n} b^{1-y_n} \ell_{n-1}$, where $\langle y_n \rangle$ is a sequence of stochastic indicator functions defined by $y_n = 1$ when $\sigma_n \leq p$, and $y_n = 0$ otherwise, where $\langle \sigma_n \rangle$ is a sequence of i.i.d. uniform-$[0,1]$ random variables.

Lemma C.1 (Stability of Linear Homogeneous Systems) *Define* $\bar{\theta} = |a|^p |b|^{1-p}$.
(*a*) *Almost surely,* $\theta^{-n}\ell_n \to 0$ *for all* $\theta > \bar{\theta}$. *In particular,* $\ell_n \to 0$ *almost surely if* $\bar{\theta} < 1$.
(*b*) *If* $\bar{\theta} < 1$ *and* $\mathcal{N}_0$ *is any open ball around 0, then there is a positive probability that* $\ell_n \in \mathcal{N}_0$ *for all* n, *provided* $\ell_0 \in \mathcal{N}_0$.

Proof: (*a*) Let $Y_n \equiv \sum_{k=1}^n y_k$. Then $|\ell_n| = \left(|a|^{\frac{Y_n}{n}} |b|^{\frac{n-Y_n}{n}}\right)^n |\ell_0|$. Since $Y_n/n \to p$ a.s. by the Strong Law of Large Numbers, the result follows from $|a|^{\frac{Y_n}{n}} |b|^{\frac{n-Y_n}{n}} \to \bar{\theta}$ a.s.

(*b*) If $ab\ell_0 = 0$, then $\langle \ell_n \rangle$ with positive chance jumps to 0 at once, and stays there. Let $ab\ell_0 \neq 0$. Since $\ell_n \to 0$ a.s., all but finitely many terms lie inside any open ball $\mathcal{N}_0$ around

[33] We are coining terms here. We call a fixed point $\bar{\ell}$ of a stochastic difference equation *locally stable* if $\Pr(\lim_{n\to\infty} \ell_n = \bar{\ell}) > 0$ whenever $\ell_0 \in \mathcal{N}_{\bar{\ell}}$, a small enough neighborhood about $\bar{\ell}$. If $\Pr(\lim_{n\to\infty} \ell_n = \bar{\ell}) > 0$ for all ℓ_0, then $\bar{\ell}$ is *globally stable*.

0 a.s., or $\Pr\left(\bigcup_{k\in\mathbb{N}}\bigcap_{n\geq k}\{\omega\in\Omega^H|\ell_n\in\mathcal{N}_0\}\right)=1$. So $\Pr\left(\{\omega\in\Omega^H|\forall n\geq k,\ell_n\in\mathcal{N}_0\}\right)>0$ for some k. So with positive chance, $\langle\ell_n\rangle$ stays inside $\mathcal{N}_0$ starting at that ℓ_k. WLOG $k=1$ since dynamics are time invariant. With scale invariant dynamics, any $\ell_0\in\mathcal{N}_0$ will do. □

REMARK 1. First, time invariance allows us to conclude that $\ell_n\in\mathcal{N}_0$ for all n with positive probability if $\ell_0=\bar{\ell}_M$. Second, the equations are linear, so that if $||\ell_0||<||\bar{\ell}_M||$ then $\ell_n\in||\ell_0||/||\ell_M||\mathcal{N}_0$ for all n. So, we have now a smaller ball around zero, such that if the system started in the small ball, it would remain in the large ball. But finally notice that from any point of the original large ball, one can reach the inner ball in only a finite number of steps, which occurs with positive probability too.

REMARK 2. Observe that it is not the *arithmetic mean* of the coefficients $pa+(1-p)b$, but their *geometric mean* that determines the behavior of the linear system. If we reformulate the criterion by first taking logarithms, as in $p\log(|a|)+(1-p)\log(|b|)<0$, then this is reminiscent of stability results from the theory of differential equations. It is common for the logarithm to enter when translating from difference to differential equations.

REMARK 3. It is straightforward to generalize Lemma C.1 to the case of more than two continuations, i.e. where y_n has arbitrary finite support. The analysis for multidimensional ℓ_n is also of importance, but unfortunately in that case only one half of the lemma goes through. Indeed, let $\ell_n\in\mathbb{R}^n$ and assume

$$\ell_n=\begin{cases}A\ell_{n-1} & \text{if } y_n=1\\ B\ell_{n-1} & \text{if } y_n=0\end{cases}$$

where A and B are given real $n\times n$ matrices. Let $||A||$ and $||B||$ denote the operator norms of the matrices.[34] Then the following half of Lemma C.1 goes through, with nearly unchanged proof (using $||AB||\leq||A||\,||B||$): If $\theta>\bar{\theta}=||A||^p||B||^{1-p}$, then $\theta^{-n}\ell_n\overset{\text{a.s.}}{\to}0$, i.e. ℓ_n converges a.s. to zero at rate $\bar{\theta}$. As this part of Lemma C.1 is the only result applied in the sequel, our local stability assertions will also go through in multidimensional models.

Call $\bar{\theta}$ a *convergence rate* of $\ell_n\to\ell^*$ if $\theta^{-n}\ell_n\to0$ for all $\theta>\bar{\theta}$. This definition is not very tight, for if $\langle\ell_n\rangle$ converges to ℓ^* at the rate θ', then it also does so at any rate $\theta''\in[\theta',1]$. Perhaps we ought to narrow down the convergence rate to the infimum rate; however, that is impractical, for in multidimensional settings we do not have the converse part of Lemma C.1. In general, it is possible to find convergence rates smaller than $\bar{\theta}$. Consider for instance the case where A is the projection onto a linear subspace, and B is the projection onto its orthogonal complement, then $\bar{\theta}=1$, but 0 is a.s. reached in finite time. We prefer to maintain the possibility of calling $\bar{\theta}$ a convergence rate, even if it is not

[34] That is, $||A||=\sup_{|x|=1}|Ax|$.

the tightest such. See Kifer (1986) for more precise results for linear systems.

⋆ **Nonlinear Stochastic Difference Equations.** Consider the revised process defined by:

$$\ell_n = \begin{cases} \varphi(1, \ell_{n-1}) & \text{with probability } \psi(1|\ell_{n-1}) \\ \varphi(2, \ell_{n-1}) & \text{with probability } \psi(2|\ell_{n-1}) \end{cases} \tag{A-6}$$

where transitions are independent: $\ell_n = \varphi(1, \ell_{n-1})$ iff $\sigma_n \leq \psi(1|\ell_{n-1})$. We care about the fixed points $\hat{\ell}$ of (A-6):

$$\varphi(1, \hat{\ell}) = \hat{\ell} \quad \text{and} \quad \varphi(2, \hat{\ell}) = \hat{\ell} \tag{A-7}$$

Theorem C.1 (Local Stability of Nonlinear Systems) *Assume that at a fixed point $\hat{\ell}$ of* (A-7), *each $\psi(i|\cdot) > 0$ is continuous and $\varphi(i, \cdot)$ is Lipschitz,*[35] *with Lipschitz constant L_i. If the stability criterion*

$$\bar{\theta} = L_1^{\psi(1|\hat{\ell})} L_2^{1-\psi(1|\hat{\ell})} < 1 \tag{A-8}$$

obtains, then for some open ball around $\hat{\ell}$, if ℓ_0 lies inside it, $\langle \ell_n \rangle$ will with positive chance forever remain inside it, while $\ell_n \to \hat{\ell}$. Also, whenever $\ell_n \to \hat{\ell}$, it converges at the rate $\bar{\theta}$.

Proof: First, we majorize (A-6) locally around $\hat{\ell}$ by a linear stochastic difference equation of the form (A-5), and then argue that Lemma C.1 applies to our original non-linear system.

WLOG $L_1 \leq L_2$. As $\psi(1, \cdot) > 0$ is continuous, (A-8) obtains — and thus $L_1 < 1$ holds — in a neighborhood $\mathcal{N}(\hat{\ell})$ of $\hat{\ell}$. Choose $\mathcal{N}(\hat{\ell})$ small enough so that for some $p \in [0, 1]$,

$$L_1^p L_2^{1-p} < 1, \; \psi(1|\ell) > p, \text{ and } ||\varphi(i, \ell) - \varphi(i, \hat{\ell})|| \leq L_i ||\ell - \hat{\ell}||, \text{ for } i = 1, 2$$

for all $\ell \in \mathcal{N}(\hat{\ell})$. Fix $\ell_0 \in \mathcal{N}(\hat{\ell})$. Define a new stochastic process $\langle \tilde{\ell}_n \rangle$ with $\tilde{\ell}_0 = \ell_0 \in \mathcal{N}(\hat{\ell})$ given, and

$$\tilde{\ell}_n - \hat{\ell} = \begin{cases} L_1(\tilde{\ell}_{n-1} - \hat{\ell}) & \text{if } y_n = 1 \\ L_2(\tilde{\ell}_{n-1} - \hat{\ell}) & \text{if } y_n = 0 \end{cases}$$

where $\langle y_n \rangle$ is our earlier i.i.d. indicator sequence. Lemma C.1 then asserts $\tilde{\ell}_n \to \hat{\ell}$ a.s., and $\tilde{\ell}_n \in \mathcal{N}(\hat{\ell})$ for all n with positive chance given $\tilde{\ell}_0 \in \mathcal{N}(\hat{\ell})$. In any realization of $\langle \sigma_n \rangle$ yielding $\tilde{\ell}_n \in \mathcal{N}(\hat{\ell})$ for all n and $\tilde{\ell}_n \to \hat{\ell}$, the linear process $\langle \tilde{\ell}_n \rangle$ majorizes the non-linear one $\langle \ell_n \rangle$: As $\psi(1|\ell) > p$ when $\ell \in \mathcal{N}(\hat{\ell})$, we have $y_n = 1 \Rightarrow \ell_n = \varphi(1, \ell_{n-1})$, and the Lipschitz property yields the majorization $||\tilde{\ell}_n - \hat{\ell}|| \geq ||\ell_n - \hat{\ell}||$ for all n (in this realization). So $\ell_n \to \hat{\ell}$ for any such $\langle \sigma_n \rangle$ realization. In summary, $\ell_n \to \hat{\ell}$ with positive probability.

Finally, the rate of convergence is $\bar{\theta}$, since for any $\theta > \bar{\theta}$, a small enough neighborhood

[35] The function $f : \mathbb{R}^m \to \mathbb{R}^k$ is *Lipschitz* at the point $\hat{x}$ with *Lipschitz constant* $L \geq 0$ if there exists a *neighborhood* $\mathcal{N}(\hat{x})$ such that $\forall x \in \mathcal{N}(\hat{x}) : \; ||f(x) - f(\hat{x})|| \leq L||x - \hat{x}||$. If f is continuously differentiable at $\hat{x}$, then f is Lipschitz with any constant $L > ||Df(\hat{x})||$.

$\mathcal{N}_p(\hat{\ell})$ exists for which $L_1^p L_2^{1-p} < \theta$. Whenever $\ell_n \to \hat{\ell}$, $\langle \ell_n \rangle$ eventually stays in $\mathcal{N}_p(\hat{\ell})$, wherein it is dominated by $\langle \tilde{\ell}_n \rangle$, which converges at rate $L_1^p L_2^{1-p}$ by Lemma C.1. □

Corollary C.1 (C¹-Local Stability) *If each $\varphi(i,\cdot)$ is also continuously differentiable, then Theorem C.1 is true with* $\bar{\theta} = |\varphi_\ell(1,\hat{\ell})|^{\psi(1|\hat{\ell})}|\varphi_\ell(2,\hat{\ell})|^{(1-\psi(1|\hat{\ell}))} < 1$.

Finally, let us briefly explain precisely how Corollary C.1 carries through in several dimensions, i.e. when ℓ is a vector (for several states). In that case $\varphi(m,\ell)$ is a vector function, and we care about its matrix derivative $D_\ell\varphi(m,\hat{\ell})$ at the stationary point. We then must assume that the operator norm $||D_\ell\varphi(m,\hat{\ell})||$ (which is the same as the largest eigenvalue) mis less than 1. Then the proof goes through, largely as before. We must also be more careful with the dominance argument. Rather than choosing a constant L_1 larger than $|\varphi_\ell(1,\ell)|$, we have to choose a matrix A with the same eigenspaces as $D_\ell\varphi(m,\hat{\ell})$, and with all numerically larger eigenvalues.

D. MORE STATES AND ACTIONS

We can handle any finite number S of states. Given pairwise mutually absolutely continuous measures μ^s for each state, we fix one reference state, and use it to define $S-1$ likelihood ratios, each a convergent conditional martingale. But the optimal decision rules would become notationally cumbersome to write down. Rather than the simple partitioning of $[0,1]$ into closed subintervals, we would now have a unit simplex in $\mathbb{R}^{S-1}$ sliced into closed convex polytopes. We leave it to the reader to ponder the optimal notation.[36] In Theorem B.2 (and its proof) we need to refer to the open intervals I and J as open balls.

If a single action a is optimal in two states of the world, which will arise if there are fewer actions than states, it will be impossible to statistically distinguish between these two states in the limit. So, even with unbounded beliefs, we cannot possibly get complete learning. But while we do not get full learning, in the terminology of Aghion, Bolton, Harris, and Jullien (1991), we get *adequate learning:* the limit beliefs are such that the correct action is chosen optimally.

In the same vein, with more than two states, the long-run ties of BHW may occur, whereby more than one action is optimal in a given state. In that case, when the true state of the world has two or more optimal actions, and there are unbounded beliefs, full learning will obtain, but we will not necessarily observe that all individuals take one

[36]The exact formulation of what constitutes full-support beliefs, which is outlined in Smith and Sørensen (1996b), is also slightly non-trivial.

particular action. In short, the overturning lemma will fail among the actions that are tied in the long run. But again, since the individuals will eventually get the optimal payoff, the learning is adequate.

The analysis also goes through virtually unchanged with a denumerable action space. Rather than a finite partition of $[0,1]$ in Lemma 2, we get a countable partition, and thus a countable set of posterior belief thresholds $\bar{r}$.[37] In this way, Lemma 3 will yield the threshold functions $\bar{p}$ just as before. The martingale properties of the model are preserved.

The convergence result Theorem 2 does not depend on the action space being denumerable. In the bounded beliefs proof, a technical complication arises, as our choice of the least m such that $\bar{p}_m(\ell) > \underline{b}$ was well-defined because there were only finitely many actions. Otherwise, we could instead just pick m so that $\bar{p}_m$ is close enough to $\underline{b}$ such that all the "bounded away" assertions hold. Similarly, in the proof of unbounded beliefs case, we could substitute a minimum action threshold $\bar{p}_1$ by one that is arbitrarily close to 0.

Complications are more insidious when it comes to Theorem 3. With $M = \infty$, both results still obtain without any qualifications provided a unique action is optimal for posteriors sufficiently close to 0 and 1, for then the overturning principle is still valid near the extreme actions. But otherwise, we must change our tune. For instance, with unbounded beliefs, there may exist an infinite sequence of distinct optimal 'insurance' action choices made such that the likelihood ratio nonetheless converges. This obviously requires that the optimality intervals $[\bar{r}_{m-1}, \bar{r}_m]$ shrink to a point, which robs the overturning argument of its strength. Yet this is not a serious non-robustness critique, because the payoff functions of the actions taken by individuals must then converge!

Under noise, the only subtlety that arises is with the trembling formulation, where we shall insist upon a finite support of the tremble from any ℓ.

E. SOME ASYMPTOTICS FOR DIFFERENCE EQUATIONS

Lemma E.1 *Let $x_i \in (0,1)$ for all i, and define $X_n = (1-x_1)\cdots(1-x_n)$.*
(a) Let $X_\infty = 0$ and $x_n = a/n + b_n$, with $\langle b_n \rangle$ summable. Then $X_n = O(n^{-a})$, and so $\sum_n X_n = \infty$ if $a \le 1$, $\sum_n X_n < \infty$ if $a > 1$, with $\sum_n X_n \ll \infty$ if $\sum |b_n| \ll \infty$.[38]
(b) Let $X_\infty > 0$. Then $\sum_n (X_n - X_\infty) < \sum n x_n$, and $\sum_n n x_n = \infty \Rightarrow \sum_n (X_n - X_\infty) = \infty$.
(c) If $X_\infty > 0$ and $x_i < \bar{x} < 1$ for all i, then $\sum_n n x_n \ll \infty \Rightarrow X_\infty \gg 0$.

[37]This may mean that we cannot necessarily well order the order the belief thresholds, nor as a result the actions.

[38]Say $x \gg c$ (eg. $x \gg 0$ is x boundedly positive) of a function or random variable x if there exists some $c_0 > c$ with $x(\cdot) \ge c_0$ always. Likewise $x \ll \infty$ (or x boundedly finite) if $x(\cdot) \le c_0$ for some $c_0 < \infty$.

Proof of (a):[39] The binomial theorem yields $d_n \equiv \log(1-a/n-b_n) - \log(1-1/n)^a = \log(1-a/n-b_n) - \log(1-a/n+c_n)$, for $c_n = O(1/n^2)$. By the mean value theorem, $|d_n| < |b_n+c_n|/\min(1-a/n-b_n, 1-a/n+c_n) < 2|b_n+c_n|$ for large enough n (since then $1-a/n-b_n > 1/2$ and $(1-1/n)^a > 1/2$). So $\sum |b_n| < \infty$ iff $\sum |d_n| < \infty$, and $X_n = n^{-a}\exp(d_1+\cdots+d_n) \equiv D_n n^{-a}$, where $0 < \underline{D} < D_n < \bar{D} < \infty$. So $X_n = O(1/n^a)$. The boundedly finite parts follow with such an exact bound.

Proof of (b): We've shown before that $(1-x_1)(1-x_2)\cdots > 1-(x_1+x_2+\cdots)$ if all x_i lie in $(0,1)$. This means that $\sum_{k>n} x_k > 1 - X_\infty/X_n$. Summing both sides over n and appealing to $\sum_n \sum_{k>n} x_k = \sum n x_n$ yields $\sum n x_n > \sum(1-X_\infty/X_n) > \sum(X_n - X_\infty)$.

Conversely, the identity $\sum n(X_{n-1}-X_n) = \sum(X_n - X_\infty)$ follows by rearranging terms of partial sums to $n=N$, and taking limits as $N \to \infty$. Our sought for result follows from

$$\sum n x_n = \sum n(X_{n-1}-X_n)/X_{n-1} < \sum n(X_{n-1}-X_n)/X_\infty = \sum(X_n - X_\infty)/X_\infty$$

Proof of (c): Since $\bar{x} < 1$, there exists $\alpha(\bar{x}) > 1$ with $1 - x \geq \exp(-\alpha(\bar{x})x)$ for all $x \in [0,\bar{x}]$.[40] So $X_\infty = (1-x_1)(1-x_2)\cdots \geq \exp(-\alpha(\bar{x})\sum_k x_k) \geq \exp(-\alpha(\bar{x})\sum_k k x_k)$. □

Lemma E.2 (Vanishing Dynamic Systems) *Let g be nondecreasing, with $g(0)=0$.*[41]
(a) If $\dot{x} = -g(x)$ and $y_{n+1} - y_n = -g(y_n)$ with $x(0) \geq y_0 > 0$, then $x(n) > y_n$.
(b) If $g' < 1$ exists, then $(1-g'(z))\dot{z} = -g(z)$ and $y_0 \geq z(0) - g(z(0)) \Rightarrow y_n \geq z(n) - g(z(n))$.

Proof of (a): By induction, assume $x(n) > y_n$. If $x(n+1) \geq y_n > y_{n+1}$, we're done. So assume $x(n+1) < y_n$. With x decreasing, there exists $t' \in (n, n+1)$ with $x(t') = y_n$. Then

$$x(n+1) = x(t') - \int_{t'}^{n+1} g(x(t))dt \geq x(t') - g(x(t'))(n+1-t') > x(t') - g(x(t')) = y_n - g(y_n) = y_{n+1}$$

Proof of (b): If not, there is a least n with $z(n) - g(z(n)) \leq y_n$ and $z(n+1) - g(z(n+1)) > y_{n+1}$. Since $\dot{z} < 0$ and $z - g(z)$ is increasing in z, we have $z(t) - g(z(t)) > y_{n+1} = y_n - g(y_n)$ for all $t \in [n, n+1]$, whereupon $z(t) > y_n$ — once again because $z - g(z)$ is increasing in z. As g is also monotone, $(1 - g'(z(t))) \cdot z(t) = -g(z(t)) \leq -g(y_n)$ for $t \leq n+1$. Then

$$z(n+1) - g(z(n+1)) = z(n) - g(z(n)) + \int_n^{n+1} \frac{d}{dt}(z(t) - g(z(t)))\,dt \leq y_n - g(y_n) = y_{n+1}$$

[39] We are very grateful to Herman Rubin of Purdue University (Statistics) for the proofs of (a) and (b).
[40] If $\xi(x) = 1 - x - e^{-\alpha x}$, then $\xi(0) = 0$ and $\xi(1) < 0$, and on $[0,\infty)$, $\xi(x) \gtrless 0$ for $x \gtrless \bar{x}(\alpha) \in (0,1)$, as $\xi'' < 0$. Since $\xi'(\bar{x}(\alpha)) < 0$, it follows that $\bar{x}'(\alpha) > 0$. Finally, $\alpha(\bar{x})$ is the inverse function to $\bar{x}(\alpha)$.
[41] We are very grateful to Anthony Quas of Cambridge University (Statistics) for result (b) and its proof.

F. OMITTED PROOFS AND EXAMPLES

⋆ **Fact: Fair Priors is WLOG.** As alluded in §2, unfair priors (i.e. states H and L not 50/50) is equivalent to a simple payoff renormalization. For Lemma 2 is still valid, as it refers only to the posterior beliefs, while the key defining indifference relation

$$\bar{r}_m u^H(a_m) + (1-\bar{r}_m)u^L(a_m) = \bar{r}_m u^H(a_{m+1}) + (1-\bar{r}_m)u^L(a_{m+1}) \tag{A-9}$$

implies that

$$\text{posterior odds } = \frac{1-\bar{r}_m}{\bar{r}_m} = -\frac{u^H(a_m) - u^H(a_{m+1})}{u^L(a_m) - u^L(a_{m+1})}$$

Since priors merely multiply the posterior odds by a common constant, the thresholds $\bar{r}_0, \ldots, \bar{r}_M$ are all unchanged if we merely multiply all payoffs in state H by the same constant. Likewise, constants added to payoffs in any state do not affect any $\bar{r}_i$.

⋆ **The Unconditional Martingale: Proof of Lemma 4.** Given the action history $\{m_1, \ldots, m_{n-1}\}$, the conditional expectation of the next public belief is (by the Markovian assumption),

$$\begin{aligned} E[q_{n+1} \mid q_n] &= q_n \sum_{m \in \mathcal{M}} \rho(m|H, q_n) \frac{1}{1 + \ell_n \frac{\rho(m|L,q_n)}{\rho(m|H,q_n)}} + (1-q_n) \sum_{m \in \mathcal{M}} \rho(m|L, q_n) \frac{1}{1 + \ell_n \frac{\rho(m|L,q_n)}{\rho(m|H,q_n)}} \\ &= q_n \sum_{m \in \mathcal{M}} \rho(m|H, q_n) \frac{q_n \rho(m|H, q_n) + (1-q_n)\rho(m|L, q_n)}{q_n \rho(m|H, q_n) + (1-q_n)\rho(m|L, q_n)} = q_n \end{aligned}$$

⋆ **Absorbing Basins: Proof of Lemma 6.** Since $\bar{p}_m(\ell)$ is increasing in m by Lemma 3, $[\bar{p}_{m-1}(\ell), \bar{p}_m(\ell)]$ is an interval for all ℓ. Then J_m is the closed interval of all ℓ that fulfill

$$\bar{p}_{m-1}(\ell) \leq \underline{b} \quad \text{and} \quad \bar{p}_m(\ell) \geq \bar{b} \tag{A-10}$$

Then interior disjointness is obvious. Next, if $\operatorname{int}(J_m) \neq \varnothing$ then $F^H(\bar{p}_{m-1}(\ell)) = 0$ and $F^H(\bar{p}_m(\ell)) = 1$ for all $\ell \in \operatorname{int}(J_m)$. The individual will choose action a_m a.s., and so no updating occurs; therefore, the continuation value is a.s. ℓ, as required.

With bounded beliefs, one of the inequalities in (A-10) holds for some ℓ, but no ℓ might simultaneously satisfy both. As Lemma 3 yields $\bar{p}_0(\ell) \equiv 0$ and $\bar{p}_M(\ell) \equiv 1$ for all ℓ, we must have $J_M = [0, \underline{\ell}]$ and $J_1 = [\bar{\ell}, \infty]$, where $\bar{p}_{M-1}(\underline{\ell}) \equiv \underline{b}$ and $\bar{p}_1(\bar{\ell}) \equiv \bar{b}$ define $0 < \underline{\ell} < \bar{\ell} < \infty$.

Finally, let $m_2 > m_1$, with $\ell_1 \in J_{m_1}$ and $\ell_2 \in J_{m_2}$. Then

$$\bar{p}_{m_2-1}(\ell_1) \geq \bar{p}_{m_1}(\ell_1) \geq \bar{b} > \underline{b} \geq \bar{p}_{m_2}(\ell_2) \geq \bar{p}_{m_2-1}(\ell_2)$$

and so $\ell_2 < \ell_1$ because $\bar{p}_{m_2-1}$ is strictly increasing in ℓ.

With unbounded beliefs, $\underline{b} = 0$ and $\bar{b} = 1$. Hence, $\bar{p}_{m-1} = 0$ and $\bar{p}_m = 1$ for $\ell \in J_m$ by (A-10). By Lemma 3, this only happens for $m = 1$ and $\ell = \infty$, or $m = M$ and $\ell = 0$.

⋆ **Limit Cascades: Proof of Theorem 1.** We first proceed here under the simplifying assumption that ρ and φ are continuous in ℓ. By Theorem B.1, stationarity at the point $\hat{\ell}$ yields $\rho(m|\hat{\ell}) = 0$ or $\varphi(m, \hat{\ell}) = \hat{\ell}$. Assume $\hat{\ell}$ meets this criterion, and consider the *smallest* m such that $\rho(m|\hat{\ell}) > 0$, so $F^H(\bar{p}_{m-1}(\hat{\ell})) = F^L(\bar{p}_{m-1}(\hat{\ell})) = 0$. Then $\varphi(m, \hat{\ell}) = \hat{\ell}$ implies $F^H(\bar{p}_m(\hat{\ell})) = F^L(\bar{p}_m(\hat{\ell})) > 0$. Since $F^H \succ_{FSD} F^L$ by Lemma 1(c), this equality is only possible if $F^H(\bar{p}_m(\hat{\ell})) = F^L(\bar{p}_m(\hat{\ell})) = 1$. Thus, $\hat{\ell} \in J_m$, as required.

Next abandon continuity. Suppose by way of contradiction that there exist a point $\hat{\ell} \in \mathrm{supp}(\ell_\infty)$ with $\hat{\ell} \notin J$. Assume WLOG the state is H. Then for some m we have $0 < F^H(\bar{p}_m(\hat{\ell})-) < 1$, so that individuals will strictly prefer to choose action a_m for some private beliefs and a_{m+1} for others. Consequently, $\bar{p}_m(\hat{\ell}) > \underline{b}$, and since $\bar{p}_0(\hat{\ell}) = 0 \leq \underline{b}$, the least such m satisfying $\bar{p}_m(\hat{\ell}) > \underline{b}$ is well-defined. So we may assume $F^H(\bar{p}_{m-1}(\hat{\ell})-) = 0$.

Next, $F^H(\bar{p}_m(\ell)) > 0$ in a neighborhood of $\hat{\ell}$. There are two possibilities:

CASE 1. $F^H(\bar{p}_m(\hat{\ell})) > F^H(\bar{p}_{m-1}(\hat{\ell}))$.

Here, there will be a neighborhood around $\hat{\ell}$ where $F^H(\bar{p}_m(\ell)) - F^H(\bar{p}_{m-1}(\ell)) > \varepsilon$ for some $\varepsilon > 0$. From (3), $\psi(m|\ell) = \rho(m|H, \ell)$ is bounded away from 0 in this neighborhood, while (4) reduces to $\varphi(m, \ell) = \ell F^L(\bar{p}_m(\ell))/F^H(\bar{p}_m(\ell))$, which is also bounded away from $\hat{\ell}$ for ℓ near $\hat{\ell}$. Indeed, $\bar{p}_m(\hat{\ell})$ is in the interior of co(supp(F)), and so Lemma 1 guarantees us that $F^L(\bar{p}_m(\ell))$ exceeds and is bounded away from $F^H(\bar{p}_m(\ell))$ for ℓ near $\hat{\ell}$ (recall that $\bar{p}_m$ is continuous). By Theorem B.2, $\hat{\ell} \in \mathrm{supp}(\ell_\infty)$ therefore cannot occur.

CASE 2. $F^H(\bar{p}_m(\hat{\ell})) = F^H(\bar{p}_{m-1}(\hat{\ell}))$.

This can only occur if F^H has an atom at $\bar{p}_{m-1}(\hat{\ell}) = \underline{b}$, and places no weight on $(\underline{b}, \bar{p}_m(\hat{\ell})]$. It follows from $F^H(\bar{p}_{m-1}(\hat{\ell})-) = 0$ and $\bar{p}_{m-2} < \bar{p}_{m-1}$, that $F^H(\bar{p}_{m-2}(\ell)) = 0$ for all ℓ in a neighborhood of $\hat{\ell}$. Therefore, $\psi(m-1|\ell)$ and $\varphi(m-1, \ell) - \ell$ are bounded away from 0 on an interval $[\hat{\ell}, \hat{\ell} + \eta)$, for some $\eta > 0$. On the other hand, the choice of m ensures that $\psi(m|\ell)$ and $\varphi(m, \ell) - \ell$ are boundedly positive on an interval $(\hat{\ell} - \eta', \hat{\ell}]$, for some $\eta' > 0$. So, once again Theorem B.2 (observe the order of the quantifiers!) proves that $\hat{\ell} \notin \mathrm{supp}(\ell_\infty)$.

⋆ **Overturning Principle: Proof of Lemma 7.** If n optimally chooses a_m, his signal σ_n must satisfy

$$\frac{1 - \bar{r}_{m-1}}{\bar{r}_{m-1}} > \ell(h) g(\sigma_n) \geq \frac{1 - \bar{r}_m}{\bar{r}_m} \tag{A-11}$$

Let $\Sigma(h)$ denote the set of all signals σ_n that satisfy (A-11). Then individual n chooses action a_m with probability $\int_{\Sigma(h)} d\mu^H$ (resp. $\int_{\Sigma(h)} g\, d\mu^H$) in state H (resp. state L). This yields the continuation

$$\ell(h, a_m) = \ell(h)\frac{\int_{\Sigma(h)} g\, d\mu^H}{\int_{\Sigma(h)} d\mu^H}$$

Now just cross-multiply, and use inequality (A-11) to bound the right hand integral.

⋆ **Herds: Proof of Theorem 3.** With bounded private beliefs, we need only combine Theorem 2 and Lemma 7. We shall prove that when the limit value $\hat{\ell}$ is in J_m, then a herd on a_m must arise in finite time. Notice that $\ell \in J_m$ implies $\bar{p}_m(\ell) \geq \bar{b}$. Consequently, (2) yields

$$\ell \geq \frac{1-\bar{r}_m}{\bar{r}_m}\frac{\bar{b}}{1-\bar{b}} > \frac{1-\bar{r}_m}{\bar{r}_m}$$

where the strict inequality follows from $\bar{b} > 1/2$. Similarly, $\ell < (1-\bar{r}_{m-1})/\bar{r}_{m-1}$ for all $\ell \in J_m$, and so the closed interval J_m lies in the interior of $[(1-\bar{r}_m)/\bar{r}_m, (1-\bar{r}_{m-1})/\bar{r}_{m-1})$. Therefore, whenever $\ell_n \to \hat{\ell} \in J_m$, we have $\ell_n \in [(1-\bar{r}_m)/\bar{r}_m, (1-\bar{r}_{m-1})/\bar{r}_{m-1})$ for $n > N$ and N big enough. By the overturning principle, only action a_m is taken after period N.

Now posit unbounded beliefs. Assume WLOG that the state is H, with a_M optimal. Theorem 2 asserts that $\ell_n \to 0$ a.s., and so ℓ_n is eventually in $[0, (1-\bar{r}_{M-1})/\bar{r}_{M-1})$. But by Lemma 7, whenever any other action than a_M is taken, we exit that neighborhood of 0.

⋆ **Calculating the Chance of a Correct Herd.** In the bounded beliefs example, given $\ell_0 = 1$, let a correct herd happen with chance π in state H. Theorem 2 tells us that a limit cascade arises a.s., and by the the reasoning in §3 about this example, $\text{supp}(\ell_\infty) \subseteq \{2u, 2u/3\}$. (Such a tight prediction is clearly impossible with cascades.) Lebesgue's Dominated Convergence assures us that $E[\ell_\infty \mid H] = \ell_1 = 1$ because $|\ell_n| \leq 2u$. The identity $1 = \pi(2u/3) + (1-\pi)(2u)$ then implicitly defines π whenever $2u/3 < 1 < 2u$.

⋆ **Cascades with Smooth Signals.** For a discrete jump into an absorbing basin, for instance $[2u, \infty)$ in figure 4, simple graphical reasoning tells us we need the left derivative $\varphi_{\ell-}(1, 2u) < 0$. Since $\varphi(1,\ell) = \ell F^L(\bar{p}(\ell))/F^H(\bar{p}(\ell))$, the private beliefs odds $F^L(\bar{p}(\ell))/F^H(\bar{p}(\ell))$, which is decreasing by Lemma A.1, must be more than unit-elastic in ℓ. The trick is to choose smooth but 'nearly' discrete private signal distributions. Assume that μ^H is Lebesgue measure on $[0,1]$, and that μ^L satisfies $d\mu^L/d\mu^H = g$, where $g' > 0$, $g(0) > 0$, and $g(1) < \infty$. The belief $p(\sigma) = 1/(1+g(\sigma))$ in state H is decreasing in σ, and has support $[1/(1+g(1)), 1/(1+g(0))]$. Given the inverse $\sigma(p) = g^{-1}((1-p)/p)$, the belief distributions are $F^H(p) = \int_{\sigma(p)}^1 d\sigma$ and $F^L(p) = \int_{\sigma(p)}^1 g(\sigma)d\sigma$. Now action a_2 requires that $\ell g(\sigma) \leq u$, and since $g(\sigma) \leq g(0)$, we have $\inf(J_1) = u/g(0)$. We thus only need $\varphi_\ell(1, u/g(0)) = 1 - g(0)(1-g(0))/g'(0)$. It suffices to choose $g'(0)$ small for fixed

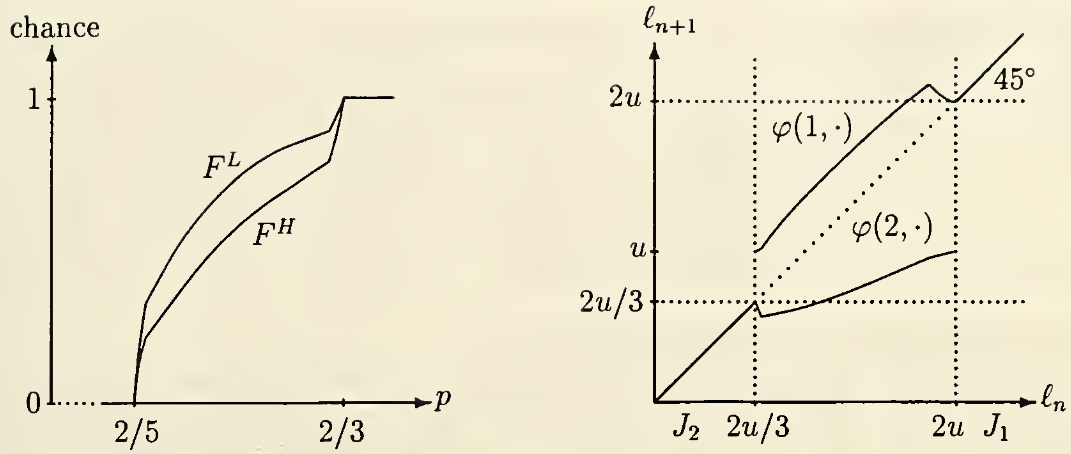

Figure 8: **Nonmonotonicity.** The left graph exhibits F^H and F^L that work in the CASCADES WITH SMOOTH SIGNALS EXAMPLE. 'nearly' having atoms at the edge of their supports, they mimic the effect in BHW: With lumpy information, actions are boundedly informative, and so a single decision can toss all successors into a cascade. The right graph shows how the corresponding continuation functions for the likelihood ratio are no longer monotonic.

$g(0) \in (0, 1)$. Figure 8 gives such an example.[42]

⋆ **Putative Herds: Proof of Lemma 8.** Let $\ell_n \in I_m \subsetneq \mathcal{C}$, with WLOG $m > 1$, and assume we are not yet in a putative herd. Absent a cascade ($\ell_n \in J_m \subset I_m$), at least two actions are possible. Then either

a) $J_m \neq \varnothing$ and action a_m is taken with chance $\gg 0$; or

b) an action a_j ($j < m$) occurs with chance $\gg 0$, whereupon $\ell_{n+1} \in I_j$, by Lemma 7; or

c) actions a_j ($j > m$) are sufficiently more likely than a_j ($j < m$) that when a_m is taken (chance $\gg 0$), $\ell_{n+1}/\ell_n \gg 1$; so in boundedly many steps, $\langle \ell_n \rangle$ enters I_j, some $j < m$.

For if not (a) or (b), then a_m, a_{m+1} is taken with boundedly positive chance. So there exists $b^* \in (\underline{b}, \bar{b})$ and $\varepsilon > 0$, a_m occurs only for private beliefs $p < b^*$, and a_{m-1} iff $p < \underline{b}+\varepsilon$. *After* a_m,

$$\frac{\ell_{n+1}}{\ell_n} \geq \frac{F^L(b^*) - F^L(\underline{b}+\varepsilon)}{F^H(b^*) - F^H(\underline{b}+\varepsilon)} \uparrow \frac{F^L(b^*)}{F^H(b^*)} > 1 \quad \text{as } \varepsilon \downarrow 0$$

where the monotonicity and inequality follow from Lemma A.1(a) and (b).

Now, a putative herd starts at once with chance $\gg 0$ in case (a), while (b) and (c) inductively help establish Lemma 8: There is $\varepsilon_m^* > 0$ and $k_m > 0$ such that ℓ transits from I_m into I_j (some $j < m$) in k_m steps with chance ε_m^*. Then set $k^* = k_1^* + \cdots + k_M^*$, and $\varepsilon^* = \varepsilon_1^* \cdots \varepsilon_M^*$ (by the conditional independence of the private signals).

⋆ **Temporary Herds: Proof of Lemma 9.** By the alternate formula for the mean of a positive r.v., the mean time to exiting from a temporary herd — that is, *conditional*

[42]Namely, $g(\sigma) = (2\sigma + 5)/10$ for $\sigma \in [0, 1/4]$, $g(\sigma) = (18\sigma + 1)/10$ for $\sigma \in [1/4, 3/4]$, and $g(\sigma) = (2\sigma + 13)/10$ for $\sigma \in [3/4, 1]$.

on eventually exiting is, the sum of chances of exiting after every period $n = 1, 2, \ldots$, or

$$\begin{aligned}\sum_{n=1}^{\infty}(1-b_1)\cdots(1-b_n) &= \sum_{n=1}^{\infty}\prod_{k=1}^{n}\frac{(1-F_k)-e_k}{1-F_k} = \sum_{n=1}^{\infty}\prod_{k=1}^{n}\frac{(1-e_k)-F_k}{1-F_k} \\ &= \sum_{n=1}^{\infty}\prod_{k=1}^{n}\frac{(1-e_k)(1-F_{k+1})}{1-F_k} = \sum_{n=1}^{\infty}\frac{E_n(1-F_{n+1})}{1-F_1} = \sum_{n=1}^{\infty}\frac{E_n-F_1}{1-F_1}\end{aligned}$$

with the last equality is true because $E_n F_{n+1} = F_1$.

⋆ **Mean Time To Herd: Proof of Theorem 5.** Consider a temporary herd on a_M in state H: As $\ell_n \in I_M$ and $E[\ell_\infty|H, \ell_n] = \ell_n$, $\langle \ell_n \rangle$ converges J_M with positive chance, and so $E_\infty > 0$. An upper bound to $\sum ne_n$ implies a *uniform upper* bound to the length of the temporary herds, and a *uniform lower* bound to E_∞, by Lemma E.1(b) and (c).

• Part (a): Extreme Signals are Atoms. We have bounded beliefs, and as above we wish to check that $\sum_n ne_n \ll \infty$. As seen in subsection 3.3, a putative herd ends in boundedly many, say at most N, steps. Then $e_n = 0$ for $n > N$, and $\sum_n ne_n < N^2$.

• Nonatomic Tails: Preliminary. We study how fast $\langle e_n \rangle$ vanishes in a putative herd on (say) a_M. When a_M is taken repeatedly, $\langle \ell_n \rangle$ obeys the recurrence

$$\ell_{n+1} - \ell_n = -\ell_n \frac{F^L(\bar{p}_{M-1}(\ell_n)) - F^H(\bar{p}_{M-1}(\ell_n))}{1 - F^H(\bar{p}_{M-1}(\ell_n))} \equiv -\eta(\ell_n) \tag{A-12}$$

Let preferences result in the private belief threshold $\bar{p}_{M-1}(\ell) = \ell/(u+\ell)$ for some $u > 0$ (Lemma 3). If $\underline{\ell}$ solves $\underline{b} = \bar{p}(\underline{\ell})$, then Lemma 1(b) implies that $\eta(\underline{\ell}) = 0$, *and* that η is increasing near $\underline{\ell}$ (as each of its factors is). Hence, we may apply Lemma E.2 to the differential equation $\dot{z} = -\eta(z)$.

• Fact 1. The proof will proceed with weaker assumptions than in the theorem, in all but one case: unbounded beliefs and $\delta = 2$. So $F^H(p) = c(p-\underline{b})^\gamma + o((p-\underline{b})^\gamma)$ and $F^L(p) = 1 - d(\bar{b}-p)^\delta + o((\bar{b}-p)^\delta)$. Now $f^H(p) = c\gamma(p-\underline{b})^{\gamma-1} + o((p-\underline{b})^{\gamma-1})$. Then with bounded beliefs ($\underline{b} > 0$), Lemma 1(a) yields $f^L(p) = c\gamma(1-\underline{b})/\underline{b}(p-\underline{b})^{\gamma-1} + o((p-\underline{b})^{\gamma-1})$, which integrates to $F^L(p) = c(1-\underline{b})/\underline{b}(p-\underline{b})^\gamma + o((p-\underline{b})^\gamma)$; similarly, $f^L(p) = c\gamma p^{\gamma-2} + o(p^{\gamma-2})$ and $F^L(p) = c\gamma/(\gamma-1)p^{\gamma-1} + o(p^{\gamma-1})$ with unbounded beliefs.[43]

• Fact 2. Consider the differential equation $\dot{x} = -j(x-\underline{x})^\gamma$. For $\gamma < 1$, the solution reaches $\underline{x}$ in finite time. For $\gamma = 1$, we get $x = \underline{x} + \exp(-jt+d)$, d arbitrary. For $\gamma > 1$, we get $x_t = \underline{x} + [(\gamma-1)(jt+h)]^{1/(1-\gamma)}$, $h > 0$ arbitrary.

• Part (b): Bounded Beliefs. One key approximation that we use here and below

[43]Observe that both integrals are valid since necessarily only $\gamma > 0$ (resp. $\gamma > 1$) renders f^L integrable near $\underline{b}$ with bounded (resp. unbounded) beliefs.

is $\bar{p}_{M-1}(\ell) - \underline{b} = \bar{p}_{M-1}(\ell) - \bar{p}(\underline{\ell})[u/(u+\underline{\ell})^2](\ell - \underline{\ell}) + O((\ell - \underline{\ell})^2)$. Recalling our formula in (A-12), Fact 1 and some tedious algebra produces $\eta(\ell) = \alpha(\ell - \underline{\ell})^\gamma + o((\ell - \underline{\ell})^\gamma)$, where $\alpha = c[u - \underline{\ell}][u/(u+\underline{\ell})^2]^\gamma$. There exist constants $k < 1 < K$ (with $K - k$ arbitrarily small when $|\ell - \underline{\ell}|$ is) with $k\alpha(\ell - \underline{\ell})^\gamma < \eta(\ell) < K\alpha(\ell - \underline{\ell})^\gamma$ for ℓ near $\underline{\ell}$.

Assume $\gamma < 2$. Then $\langle \ell_n \rangle$ is bounded above by the solution to the above differential equation of Fact 2, by Lemma E.2(a). Since $e_n = F^H(\bar{p}_{M-1}(\ell_n)) = c[u/(u+\underline{\ell})^2](\ell_n - \underline{\ell})^\gamma + o((\ell_n - \underline{\ell})^\gamma)$, we have (*i*) for $\gamma < 1$, $e_n = 0$ after finitely many steps; (*ii*) for $\gamma = 1$, $\langle e_n \rangle$ converging exponentially to 0; and (*iii*) for $\gamma \in (1,2)$, $e_n = O(n^{-\gamma/(\gamma-1)})$, so $\langle n e_n \rangle$ is summable. In all cases, there is an open neighborhood $\mathcal{N}$ containing the basin J_M, with $\sum n e_n \ll \infty$ when $\langle \ell_n \rangle$ is in $\mathcal{N}$. But outside of $\mathcal{N}$, we see from (A-12) that $\langle \ell_n \rangle$ is falling in our putative herd by boundedly positive decrements, or $\ell_{n+1} < \ell_n - \varepsilon$ for some $\varepsilon > 0$, for then $\eta(\ell) > \varepsilon$. It thus reaches $\mathcal{N}$ in at most $(1 - \bar{r}_M)/(\varepsilon \bar{r}_M)$ steps. Thus, $\sum n e_n \ll \infty$.

With $\gamma \geq 2$, η is differentiable at $\underline{\ell}$ with $\eta'(\underline{\ell}) = 0$. Applying Lemma E.2(b), we can see that $\ell_n = O(n^{-1/(\gamma-1)})$, and that $e_n = O(n^{-\gamma/(\gamma-1)})$. So $\sum_n e_n = \infty$.

• Part (*c*): Unbounded Beliefs. Mimicking the fist part of the bounded beliefs analysis yields $\eta(\ell) = \alpha\ell^\gamma + o(\ell^\gamma)$, where $\alpha = cu^{1-\gamma}\gamma/(\gamma-1)$. Notice that $\eta'(0) = 0$.

We now have two separate cases in state H: We want $\sum n e_n \ll \infty$ for a putative herd on a_M, and $\sum E_n \ll \infty$ for one on a_1 (which can't be a herd, and so $E_\infty = 0$). In the first case, $e_n = O(n^{-\gamma/(\gamma-1)})$, and thus $\sum n e_n \ll \infty$ iff $\gamma < 2$, just as with bounded beliefs.

For the second case, Lemma E.1(*a*) provides a simple test for $\sum E_n \ll \infty$ using $e_n = 1 - F^H(\bar{p}_1(\ell_n))$ alone. A putative herd on a_1 in state H is analytically equivalent to one on a_M in state L (with the derived criterion on γ then applied to δ). So, we let $\langle \ell_n \rangle$ evolve as before, but consider $e_n = F^L(\bar{p}_{M-1}(\ell_n))$. Now Lemma E.2(a) and (b) yield $[(\gamma-1)(K\alpha n + h)]^{-1/(\gamma-1)} \leq \ell_n \leq [(\gamma-1)(k\alpha n + h)]^{-1/(\gamma-1)}$, with $k < 1 < K$, both arbitrarily close to one as ℓ_n is close to 0. So, $\alpha/[(K\gamma - 1)(\alpha n + h)] \leq e_n \leq \alpha/[(\gamma-1)(k\alpha n + h)]$. For $\gamma < 2$, we can find $a < 1$ with $e_n \leq a/n$ eventually. Then $\sum_n E_n = \infty$. Likewise for $\gamma > 2$, $a > 1$ exists with $e_n \geq a/n$ eventually, and so $\sum_n E_n \ll \infty$.

The case $\gamma = 2$ requires special attention — as already noted, we go back to the assumptions of the theorem in this case.[44] Then it can be proved exactly as before that $F^L(p) = 2cp^2 + O(p^3)$. We have $\eta(\ell) = \alpha\ell^2 + O(\ell^3)$, and it can be directly verified that $\eta(\ell)/(1 - \eta'(\ell)) = \alpha\ell^2 + O(\ell^3)$. So, we consider $dz/dt = -\alpha z^2 + bz^3$ with $\alpha > 0$. Integration yields $z(t) = 1/[\alpha t + c - (b/\alpha)\log(-b + \alpha/z(t))]$ for $z(t) < \alpha/b$, i.e. z close enough to zero. Using our prior knowledge that $z(t) = O(1/t)$ together with $\log(1+x) = x + O(x^2)$, we get $z(t) = 1/[\alpha t + c - (b/\alpha)[\log(\alpha/z(t)) - b/(\alpha t) + O(1/t^2)]] = 1/(\alpha t + c') + O(\log(t)/t^2)$.

[44] If we had used the more general assumption we could not conclude below that $e_n - 1/n$ is summable. We think a slightly larger error term likely also suffices, but we shall refrain from pursuing this.

Lemma E.2(b) then yields asymptotically $\ell_n = 1/(\alpha n + h) + O(\log(n)/n^2)$, and then we get $e_n = \alpha/(\alpha n + h) + O(\log(n)/n^2)$. Then $e_n - 1/n$ is summable, and we conclude from Lemma E.1(a) that $\sum_n E_n = \infty$. □

⋆ **On the Trembling Form of Noise.** A qualitatively different form of noise instead posits that one may 'tremble', à la Selten (1975): Individuals randomly take a suboptimal action with some idiosyncratic chance. Someone planning to take action a_m will instead opt for a_j with probability $\kappa_m^j(\ell)$ for $j \neq m$, possibly dependent on ℓ. For simplicity, we insist that all $\kappa_m^j(\ell)$ be bounded away from 0. So the chance of deviating from an optimal action a_m is $\kappa_m(\ell) = \sum_{j \neq m} \kappa_m^j(\ell)$. Even with constant values of κ_m^j, this form of noise is history-dependent, and must be bounded above: $\kappa_m(\ell) < 1/2M$ for all m. The dynamics in state H are now described by (7) and

$$\psi(m|s,\ell) = [1 - \kappa_m(\ell)]\rho(m|s,\ell) + \textstyle\sum_{j \neq m} \kappa_j^m(\ell)\rho(j|s,\ell) \tag{A-13}$$

Observe crucially that *craziness is a special case of trembling,* where $\kappa_j^m(\ell)$ is invariant across j and ℓ: Regardless of plans, one accidentally takes action a_m with fixed chance κ_m.

Proof of Theorem 7 for Trembling Individuals: We now show that Theorem 7 is also valid for trembling noise. Indeed, all actions must be taken with positive probability, and so $\psi(m|\ell)$ is indeed bounded away from 0 by (A-13). We wish to argue once more that $\varphi(m,\ell) - \ell = 0$ is satisfied under exactly the same conditions as in the proofs of Theorems 2 and 2. Let $\ell \neq 0$ be a stationary point, and assume by way of contradiction that more than one action is taken with positive probability. Then $0 < F^H(\bar{p}_m(\ell)) < 1$ for some m. For any such m, we can use (7) to rewrite $\varphi(m,\ell) = \ell$ as follows:

$$[1 - \kappa_m(\ell)]\,[\rho(m|L,\ell) - \rho(m|H,\ell)] = \sum_{k \neq m} \kappa_k^m(\ell)\,[\rho(k|H,\ell) - \rho(k|L,\ell)] \tag{A-14}$$

Here, the sum on the right hand side may have negative and positive terms, but notice that

$$\begin{aligned} &\sum_{k=1}^{m} [\rho(k|H,\ell) - \rho(k|L,\ell)] \\ &= \sum_{k=1}^{m} \left[F^H(\bar{p}_k) - F^H(\bar{p}_{k-1}) - F^L(\bar{p}_k) + F^L(\bar{p}_{k-1})\right] = F^H(\bar{p}_m) - F^L(\bar{p}_m) \end{aligned}$$

Recall that the function $F^L - F^H$ is first increasing, then decreasing, and that the threshold $\bar{p}_m$ is increasing in m. Thus, the negative terms in the sum can at most sum to (minus) the number $\bar{F}(\ell) \equiv \max_{m=1,\ldots,M}\{F^L(\bar{p}_m) - F^H(\bar{p}_m)\}$. Since $\sum_{k\neq m}[\rho(k|H,\ell) - \rho(k|L,\ell)] = \rho(m|H,\ell) - \rho(m|L,\ell)$, and $\kappa_k^m(\ell) \leq \kappa_k(\ell) \leq 1/2M$ by assumption, the left hand side of (A-14) obeys the inequality

$$[1-\kappa_m(\ell)]\,[\rho(m|L,\ell) - \rho(m|H,\ell)] \leq \frac{1}{2M}[\rho(m|L,\ell) - \rho(m|H,\ell)] + \frac{1}{2M}\bar{F}(\ell)$$

and so

$$\left[1 - \frac{1}{M}\right][\rho(m|L,\ell) - \rho(m|H,\ell)] \leq \frac{1}{2M}\bar{F}(\ell)$$

As this holds for all m, we may sum over $m = 1, \ldots, \tilde{m}$, and discover that

$$\left[F^L(\bar{p}_{\tilde{m}}) - F^H(\bar{p}_{\tilde{m}})\right] \leq \tilde{m}\frac{1}{2M-2}\bar{F}(\ell) < \frac{M}{2M-2}\bar{F}(\ell)$$

which is impossible by definition of $\bar{F}(\ell)$ (and $M \geq 2$). Hence, the equations $\varphi(m,\ell) = \ell$ could only be solved by an ℓ for which only one action is optimal. The proof of Theorem 2 obtains once again, while $\varphi(1,\ell) - \ell$ is bounded away from 0 on the interval I of Theorem B.2, and so Theorem 2 goes through just as before also.

References

AGHION, P., P. BOLTON, C. HARRIS, and B. JULLIEN (1991): "Optimal Learning by Experimentation," *Review of Economic Studies*, 58(196), 621–654.

ARTHUR, W. B., Y. M. ERMOLIEV, and Y. M. KANIOVSKI (1986): "Strong Laws For a Class of Path-Dependent Stochastic Processes with Applications," in *Stochastic Optimization: Proceedings of the International Conference, Kiev, 1984*, ed. by I. Arkin, A. Shiraev, and R. Wets. Springer-Verlag, New York.

BANERJEE, A. V. (1992): "A Simple Model of Herd Behavior," *Quarterly Journal of Economics*, 107, 797–817.

BELLMAN, R. (1954): "Limit Theorem For Non-commutative Operations. I," *Duke Mathematical Journal*, 21, 491–500.

BIKHCHANDANI, S., D. HIRSHLEIFER, and I. WELCH (1992): "A Theory of Fads, Fashion, Custom, and Cultural Change as Information Cascades," *Journal of Political Economy*, 100, 992–1026.

BRAY, M., and D. KREPS (1987): "Rational Learning and Rational Expectations," in *Kenneth Arrow and the Ascent of Economic Theory*, ed. by G. Feiwel, pp. 597–625. Macmillan, London.

BREIMAN, L. (1968): *Probability*. Addison-Wesley, Reading, Mass.

DOOB, J. L. (1953): *Stochastic Processes*. Wiley, New York.

EASLEY, D., and N. KIEFER (1988): "Controlling a Stochastic Process with Unknown Parameters," *Econometrica*, 56, 1045–1064.

ELLISON, G., and D. FUDENBERG (1995): "Word of Mouth Communication and Social Learning," *Quarterly Journal of Economics*, 109, 93–125.

FURSTENBERG, H. (1963): "Noncommuting Random Products," *Transactions of the American Mathematics Society*, 108, 377–428.

FUTIA, C. A. (1982): "Invariant Distributions and the Limiting Behavior of Markovian Economic Models," *Econometrica*, 50, 377–408.

HARSANYI, J. C. (1967–68): "Games with Incomplete Information Played by Bayesian Players," *Management Science*, 14, 159–182, 320–334, 486–502.

KIFER, Y. (1986): *Ergodic Theory Of Random Transformations*. Birkhauser, Boston.

KIHLSTROM, R., L. MIRMAN, and A. POSTLEWAITE (1984): "Experimental Consumption and the 'Rothschild Effect'," in *Bayesian Models in Economic Theory*, ed. by Boyer, and Kihlstrom. Elsevier, New York.

LEE, I. H. (1993): "On the Convergence of Informational Cascades," *Journal of Economic Theory*, 61, 395–411.

MILGROM, P. (1979): "A Convergence Theorem for Competitive Bidding," *Econometrica*, 47, 670–688.

RUDIN, W. (1987): *Real and Complex Analysis*. McGraw-Hill, New York, 3rd edn.

SELTEN, R. (1975): "Reexamination of the Perfectness Concept for Equilibrium Points in Extensive Games," *International Journal of Game Theory*, 4, 25–55.

SMITH, L. (1991): "Error Persistence, and Experiential versus Observational Learning," Foerder Series working paper.

SMITH, L., and P. SØRENSEN (1996a): "Informational Herding as Experimentation Déjà Vu," MIT Working Paper.

——— (1996b): "Rational Social Learning Through Random Sampling," MIT mimeo.