Rational Social Learning

by

Peter Sorensen

Master of Science in Mathematics and Economics
University of Copenhagen, 1992

Submitted to the Department of Economics
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy in Economics

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

May 1996

Author .....
....
Department of Economics
March 19, 1996

Certified by.
....
Lones A. Smith
Associate Professor
Thesis Supervisor

Certified by.
....
Abhijit Banerjee
Associate Professor
Thesis Supervisor

Accepted by ....
Richard S. Eckaus
Chairman, Departmental Committee on Graduate Students

Rational Social Learning

by

Peter Sorensen

Submitted to the Department of Economics
on March 19, 1996, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy in Economics

Abstract

Chapter 1 systematically analyzes and enriches the observational learning paradigm of Banerjee (1992) and Bikhchandani, Hirshleifer, and Welch (1992). We find that herding is quite robust. BHW's kindred notion of *cascades* is not as resilient: While beliefs converge upon a limit where only one action is taken with probability one (a *limit cascade*), this need not occur in finite time. The potential for bad herds is also not without caveat: Absent a uniform bound on the strength of the individuals' private signals, only correct herds arise. Finally, in a world with multiple preference types, a *confounded learning* outcome might arise, where the lesson of history is forever mixed, and private signals always decisive.

In chapter 2 we augment the informational herding model with a social planner maximizing welfare. We find that even when the herding externality is internalized in this fashion, incorrect herds and incomplete learning still obtain. We show how the optimal plan can be implemented by transfers. Along the way we prove that the observational learning models are but special cases of a standard single person experimentation model with myopia. We then re-interpret the incorrect herding outcome as a familiar failure of complete learning in an optimal experimentation problem.

Chapter 3 abandons two standard assumptions in observational learning (or herding) models: (i) that the *entire ordered* action history is observed, and (ii) that it is *actions* which are observed. Rather, we propose the following simple paradigm for social learning: Individuals learn by observing signals of some of their predecessors' posterior beliefs. Since posterior beliefs are informative of the state of the world, any informative statistic thereof is likewise, enabling social learning to be a fruitful enterprise. We discuss the model properties that guarantee complete learning in the long run. One main conclusion is that the process of learning is not severely affected by the assumption that the history is less than perfectly observed. However, with learning based on plain conversation a confounding outcome arises, as coarse signals are not sufficiently informative about posteriors.

Thesis Supervisor: Lones A. Smith
Title: Associate Professor

Thesis Supervisor: Abhijit Banerjee
Title: Associate Professor

Contents

Acknowledgments

I am particularly grateful to my advisor Lones Smith, who has co-authored all three chapters of this thesis.

I thank Christopher Avery, Abhijit Banerjee, Tilman Börgers, Glenn Ellison, John Geanakoplos, David Hirshleifer, Bruno Jullien, David Myatt, Marco Ottaviani, Thomas Piketty, Ben Polack, Robert Shimer, Xavier Vives, Christopher Wallace, Jörgen Weibull, Jennifer Wu, and Peter Zemsky for detailed comments on all or parts of this thesis. Several seminar audiences, especially the MIT theory lunch, have given many valuable suggestions.

I gratefully acknowledge financial support for my graduate studies and thesis research from the Danish Social Sciences Research Council, the Danish Fulbright Commission, the Danish Research Academy, and the American Scandinavian Foundation.

Introduction

Over the last few years a lot of research has been carried out to shed light on a paradigm known as informational herding. The recent interest was sparked independently by Banerjee (1992) and Bikhchandani, Hirshleifer, and Welch (1992). This thesis takes part in the endeavor to rigorously investigate models of rational social learning.

The context is simple: An infinite sequence of individuals must decide on an action choice from a finite menu. Everyone has identical preferences and menus, and each may condition his decision both on his (endowed) private signal about the state of the world and on all his predecessors' decisions; however, observation of others' private signals is impossible. If these private signals have bounded power, it is known that a 'herd' eventually arises, but is not always correct — namely, after some point, all make the identical choice.

In chapter 1 this story is explored with the purpose to question the robustness of the herding predictions. Informational herding is found to be quite robust. The notion of *cascades* is not as resilient: While beliefs converge upon a limit where only one action is taken with probability one (a *limit cascade*), this need not occur in finite time. The potential for bad herds is also not without caveat: Absent a uniform bound on the strength of the individuals' private signals, only correct herds arise. Finally, in a world with multiple preference types, a *confounded learning* outcome might arise, where the lesson of history is forever mixed, and private signals always decisive.

The herding outcome is a striking example of an informational externality: While all individuals collectively know enough to fully determine the state of the world, it is aggregated rather poorly. For in a herd, most individuals almost surely take an action which reveals almost none of their information. Late-comers ideally prefer that their predecessors had better signalled their information with more revealing actions, but early individuals clearly have no incentive to do so.

The purpose of chapter 2 is to investigate the herding externality with more forward-looking behavior. We can neatly address this issue by interpreting the herding model as an individual experimentation model with myopia. When a social planner tries to maximize an average (discounted) of the individuals' welfare, is that the same experimentation model arises without myopia. As is well known, a rational experimenter will forego some current payoff in order to secure knowledge relevant for future payoffs and decisions. Still, we find that even when a planner internalizes the herding externality in this fashion, incorrect herds and incomplete learning still obtain, albeit with chances vanishing as the discount factor converges to one. The social planner's optimum outcome can be decentralized as a constrained social optimum using monetary transfers between agents.

One crucial and arguably untenable assumption in the herding literature is the perfect observability of the entire action history. How important is this modelling assumption? Chapter 3 moves one step beyond the constraints of the simplest observational regimes. Inspired by Banerjee and Fudenberg (1995) we consider learning from *samples*: Individual n only observes a random, ordered or unordered sample of k of his predecessors' actions. Full learning still obtains when private beliefs are unbounded.

In the herding model, the vehicle for social learning is action observation. But Banerjee (1992) and Lee (1993) pointed out that when the action space and payoff functions are continuous, this may well permit a perfect inference of predecessors' signals. This reduces social learning to pure *statistical learning,* yielding fast complete learning. An action choice is one natural coarse informative statistic of a posterior belief, but it need not be the only one. In general, any discrete informative signal of an individual's posterior belief is informative. This imperfect perception model is explored in chapter 3. Combined with the observation of unordered samples it turns out to wreck the complete learning results.

In sum, this thesis attempts to further in several directions the general understanding of the informational herding paradigm. While chapter 2 shows the original herding phenomena of Banerjee (1992) and Bikhchandani, Hirshleifer, and Welch (1992) to be single-person experimentation pathologies rehashed, in chapter 3 we explore their viability in models never anticipated in the experimentation framework.

Chapter 1

Pathological Outcomes of Observational Learning

1.1. INTRODUCTION

Suppose that a countable number of individuals each must make a once-in-a-lifetime binary decision[1] — 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. After some individual, everyone would make the identical less profitable decision.

This is a compellingly robust 'pathological' outcome: Individuals do not eventually learn the true state of the world despite the surfeit of available information. So let us

[1] An economic example might be the decision of which profession to enter.

seriously analyze the model whence it arises. For we believe that learning from others' actions is economically important, and as such, merits the theoretical scrutiny long afforded the single-person learning results, and the rational expectations learning literature.

In this paper, we systematically analyze and enrich the observational learning paradigm. Our analysis shall be focused through two lenses. BHW found in their model that a herd almost surely arises — which we shall call *convergence of actions*: i.e. everyone eventually settles on a common decision. It will be a correct or incorrect herd as the action is ex post optimal or suboptimal. To illustrate the cognate phenomenon of *convergence of beliefs,* BHW introduced the colorful terminology of a *cascade;* this refers to any infinite train of individuals who decide to act irrespective of the content of their signal.

Our embellishment upon the above herding story is best motivated by means of the following counterfactual. Assume that we are in a herd of the Banerjee-BHW sort 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 immediate predecessor had a more powerful signal than everyone else. To capture this, we shall generalize the signal space beyond discrete signals, and admit the possibility that there is no uniformly most powerful signal. Second, he might opine that the action was an act of irrationality or was an accident. We shall thus explore the addition of various forms of noise to the herding model. Third, he possibly might decide that the action was prompted by different preferences. On this note, 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!

1. The Role of the Signal Space. Our preamble highlights a simple property of Bayesian updating from a common prior (the *no introspection condition*). We then focus on the role played by the support of private signals: Incorrect herding can arise exactly when individuals have ***bounded private beliefs,*** i.e. there do not exist arbitrarily strong private signals (private likelihood ratios are bounded away from zero and infinity). This result may not be altogether obvious. Inspired by recent related work by Banerjee and Fudenberg (1995), one might conjecture that herding would not obtain when individuals have sufficiently but boundedly strong private signals. For in that case, a finite history might never provide a sufficiently strong signal to mislead everyone. But this intuition ignores the fact that the strength of history increases with the strength of private signals, so that eventually even the most doctrinaire individual dare not quarrel with its conclusion.

With *unbounded private beliefs,* incorrect herds cannot occur. Eventually, a sufficiently (wisely) doctrinaire individual will appear on the scene whose contrary action will radically shift public beliefs and overturn the would-be herd. This we formalize as *the overturning principle*: convergence of beliefs implies convergence of actions. Next, we show that this logic cannot be turned on its head to rule out correct herds: Complete learning must obtain.[2] The proof relies on the overturning principle to show that belief convergence implies action convergence. For both cases, we ask whether convergence of actions occurs in *finite expected time.* We answer negatively, showing that what is relevant is not how fast individuals learn the truth but rather how quickly they realize error. This in turn depends on the asymptotic size of tail of the signal distributions.

The admission of individuals who are arbitrarily tenacious in their beliefs is largely a modelling decision, and therein lies its ultimate justification. For it provides us with a richer model than possible in Banerjee (1992) and BHW, allowing us to consider natural economic modifications that hinder informational transmission, and ask if convergence still (almost surely) occurs. This natural line of inquiry made no sense in their framework.

Absent the discrete signals of BHW, their main result that cascades eventually occur is false. Every example in this paper attests to this fact. With generic signal distributions, a cascade need never arise even though a herd must! With these two notions decoupled, the resulting analysis is simultaneously more nontrivial, and richer. We find that convergence of beliefs always occurs, but convergence of actions is less certain.

2. Noise. That a single individual can 'overturn the herd' is the key insight into the existence of herds. The natural introduction of noise into the model prevents this from happening. Counterintuitively, even with a constant inflow of mostly 'crazy' individuals, we still find that complete learning result without noise still obtains. Everyone (sane)

[2]The link between unbounded support beliefs and complete learning, not as cleanly expressed, was introduced in Smith (1991).

eventually learns the true state of the world. The deduction of convergence of actions is now only possible by considering the speed of belief convergence.

3. 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. The existence of multiple types offers a parallel reason why the actions of isolated individuals need not greatly matter, but it has 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, corresponding to 70% choosing left or right, respectively. But there is another rather subtle possibility that may arise. For as the probability q that observed history accords state R rises from 0 to 1, the chance $r(q)$ that a Houston (resp. Dallas) driver merges right gradually increases (resp. decreases) from 0 to 1. If perchance for some q, cars ought to merge right with the same probability in both states R and L, 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 1.7 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!

4. Likelihood Ratios as Martingales and the Stochastic Stability of Learning. Focusing on the appropriate stochastic processes turns out to have been a crucial step in our analysis. The basic building block for all our theory is that, conditional on the state, the public likelihood ratio is a martingale *and* the vector of (action taken, likelihood ratio) a time homogeneous Markov chain. The Markovian aspect of the dynamics allows us (just as it did Futia (1982)) 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 — which is unavailable in Futia (1982) — 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. Much more unexpected however is a simple link we discover between the martingale property of the likelihood ratio and the *stability* of the associated stochastic difference equations. This provides the necessary ingredient for the deduction of exponentially fast belief convergence of actions under noise, and with multiple types near the confounding outcome.

On the Herding Literature. As Smith and Sørensen (1996b) explore, that individuals observe lumpy signals (namely their actions) of predecessors' beliefs is the only real touchstone of the herding literature. This characterization rules out the continuous action space model of Lee (1993); with a continuous payoff function, this allows perfectly inference of the underlying privat signals![3] Indeed, herding pathologies were absent from the rational expectations literature for this very reason,[4] since price can adjust continuously.[5]

Section 1.2 outlines the basic model and provides some preliminary results. We characterize belief and action convergence in section 1.3, and again with noise in section 1.4. The model with heterogeneous preferences is explored in 1.5. exercise in optimal experimentation. An appendix, among other things, derives some new convergence criteria for Markov-martingale stochastic processes, as well as some simple results on the local stability of stochastic difference equations.

1.2. THE STANDARD MODEL

1.2.1 Some Notation

We first introduce a background probability space $(\Omega, \mathcal{E}, \nu)$. This space underlies all random processes in the model, and is assumed to be common knowledge.

[3] In contrast, Banerjee's (1992) model had a continuous action space but a discontinuous payoff function.
[4] For a good take on this field, see Bray and Kreps (1987).
[5] This assumes that individuals can continuously adjust their trading quantities. If not, Avery and Zemsky (1992) have shown that a temporary herd may arise.

An infinite sequence of individuals $n = 1, 2, \dots$ sequentially takes actions in that exogenous order. *Individuals observe the actions of all predecessors.* There is uncertainty about the payoffs from these actions. Thus, there are $S = 2$ possible *states of the world* (or more simply, *states*), $s = H$ ('high') and $s = L$ ('low'). Formally, this means that the background state space Ω is partitioned into two events Ω^H and Ω^L, called H and L. Let the common prior belief be that $\nu(H) = \nu(L) = 1/2$.[6] The results of this paper extend to any finite number of states, but at significant notational cost (see Appendix 1.D).

Each individual can choose from a finite set of actions $\langle a_m, m \in \mathcal{M}\rangle$, where $\mathcal{M} = \{1, \dots, M\}$. One might think of investors deciding whether to 'invest' or 'decline' an investment project of uncertain value. Action a_m has a payoff $u^s(a_m)$ in state $s \in \{H, L\}$, *the same for all individuals,* and everyone wishes to take the action that maximizes his expected payoff. We assume WLOG that no action is weakly dominated (by the set of other actions), and to avoid trivialities we insist that at least two such undominated actions exist. Before deciding upon an action, the individual can observe the entire action history. We shall loosely denote the action profile of *any finite number of* individuals as h. How the individual uses that history is considered in the next subsection.

Individual n receives a private random signal, $\sigma_n \in \Sigma$, about the state of the world. *Conditional on the state,* the signals are assumed to be i.i.d. across individuals. It is common knowledge that they are distributed according to the probability measure μ^s in state $s \in \{H, L\}$.[7] 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.).[8] 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,[9] so that μ^H and μ^L are not the same measure; this will ensure that some signals are *informative* about the state.

[6] That individuals have common priors is a standard modelling assumption, see e.g. Harsanyi (1967–68). Also, a flat prior over states is truly WLOG, for it will turn out that more general priors will be formally equivalent to a renormalization of the payoffs, as seen in section 1.2.3 below.

[7] Formally, we mean that $\sigma_n : \Omega \to \Sigma$ is a stochastic variable, and $\mu^s \equiv \mu^s_n = \nu^s \circ \sigma_n^{-1}$, where ν^s is the measure ν conditioned on the event Ω^s, $s \in \{H, L\}$.

[8] 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 $S \in \mathcal{S}$.

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

1.2.2 Private Beliefs

Given signal $\sigma \in \Sigma$, the 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 both F^H and F^L have densities.

Lemma 1.1 (No Introspection Condition) *The derivative f satisfies $f(p) = p/(1-p)$ for almost all $p \in (0,1)$. Conversely, any private belief distributions F^H, F^L satisfying $dF^H/dF^L(p) = p/(1-p)$ for all $p \in (0,1)$ arise from some signal measures μ^H and μ^L.*

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. Consequently, F^H grows faster than F^L exactly when $p \in \text{supp}(F)$ satisfies $p > 1/2$.

Finally, given $f(p) = p/(1-p)$, let σ have distribution F^s in state s, $s \in \{H, L\}$. □

All results to follow are ultimately driven by the conditional distribution functions F^H and F^L, which by Lemma 1.1, can be taken as the stochastic primitives of the model. Notice that Lemma 1.1 implies that F^H and F^L have a common support,[10] say $\text{supp}(F)$, which almost surely coincides with the range of $p(\cdot)$ on Σ. The structure of $\text{supp}(F)$ will play a major role in the paper. In fact, $\text{co}(\text{supp}(F)) \equiv [\underline{b}, \bar{b}] \subseteq [0,1]$ with $0 \leq \underline{b} < 1/2 < \bar{b} \leq 1$.[11] The strict inequalities follow directly from the assumption that μ^L and μ^H are distinct. Also, since there are no perfectly revealing signals, $F^L(\underline{b}) = 0$ and $F^H(\bar{b}-) = 1$. We shall call the private beliefs *bounded* if $0 < \underline{b} < \bar{b} < 1$; if $\text{co}(\text{supp}(F)) = [0,1]$, private beliefs are *unbounded.* To exhaust all possibilities we should also consider supports that are bounded above and not below, or vice versa, but this exercise in generality yields no new insights.

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

---

[10] As usual, the support of a measure on the Borel-algebra is the smallest closed set of measure 1.

[11] Here, $\text{co}(A)$ denotes the convex hull of the set A.

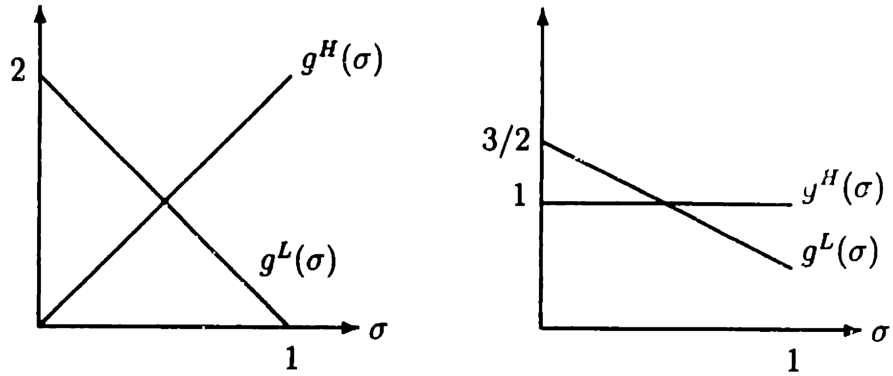

Figure 1-1: **Signal Densities.** Graphs for the unbounded (left) and bounded (right) beliefs examples.

UNBOUNDED BELIEFS EXAMPLE. Let μ^L have probability density $g^L(\sigma) = 2-2\sigma$, and μ^H the density $g^H(\sigma) = 2\sigma$. Hence, $g(\sigma) = g^L(\sigma)/g^H(\sigma) = (1-\sigma)/\sigma$, yielding $p(\sigma) = \sigma$, and thus $F^H(p) = p^2$ and $F^L(p) = 2p - p^2$. The left panel of figure 1-1 depicts the above densities g^H and g^L. Observe how (*i*) signals near 0 very strongly favor of state L, and more generally, (*ii*) the important (yet trivial) principle that signals in favor of a state are more likely to occur in that state. Ultimately, this drives the full learning results we strive for.

BOUNDED BELIEFS EXAMPLE. Next, let μ^H be Lebesgue (uniform) measure on $[0,1]$, and choose μ^L so that the Radon-Nikodym derivative is $g(\sigma) = \frac{d\mu^L(\sigma)}{d\mu^H(\sigma)} = 3/2 - \sigma$ on $[0,1]$. Then $p(\sigma) = 1/(g(\sigma)+1) = 2/(5-2\sigma)$. Note how p maps $[0,1]$ injectively onto $[\underline{b}, \bar{b}] = [2/5, 2/3]$, with $p(\sigma) \le p \Leftrightarrow 2/(5-2\sigma) \le p \Leftrightarrow (5p-2)/2p \ge \sigma$. Here, the support of F^H and F^L is bounded, and the distribution of $p \in [2/5, 2/3]$ is described in state H by $F^H(p) = \mu^H[0, (5p-2)/2p] = (5p-2)/2p$, and in state L by

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

1.2.3 Action Choice

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

Lemma 1.2 *The interval* $(0,1)$ *partitions into relatively closed subintervals* $\bar{I}_1, \ldots, \bar{I}_M$ *overlapping at endpoints only, such that action* a_m *is optimal when the posterior* $r \in \bar{I}_m$.

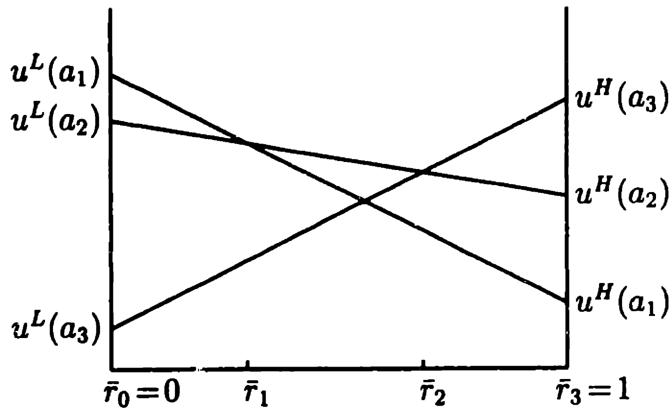

Figure 1-2: **Payoff Frontier.** The diagram depicts the (expected) payoff of each of three actions as a function of the posterior belief r that the state is H. The individual simply chooses the action yielding the highest payoff. Action a_2 is an *insurance action*, while actions a_1 and a_3 are *extreme actions.*

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 strictly order the actions so that a_m is optimal exactly when the posterior $r \in [\bar{r}_{m-1}, \bar{r}_m] \equiv \bar{I}_m$, where $0 = \bar{r}_0 < \bar{r}_1 < \cdots < \bar{r}_M = 1$. Let's further introduce the tie-breaking rule[12] that individuals take action a_m (versus a_{m+1}) at $r = \bar{r}_m$. Note that action a_M (resp. a_1) is optimal when one is certain that the state is H (resp. L). With more actions than states, we can think of a_1 and a_M as 'extreme' actions, and $a_2 \ldots, a_{M-1}$ as 'insurance' actions. Of course, the latter may well be chosen, since decisions are often taken without the luxury of perfectly focused beliefs.

We can now see how unfair priors (i.e. not 50/50 across states H and L) are equivalent to a simple payoff renormalization, as asserted earlier. For the characterization in Lemma 1.2 is still valid, as it refers only to the posterior beliefs, while the key defining indifference relation

$$\bar{r}_m u^H(a_m) + (1-\bar{r}_m)u^L(a_m) = \bar{r}_m u^H(a_{m+1}) + (1-\bar{r}_m)u^L(a_{m+1}) \tag{1.1}$$

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

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 unfair priors merely serve to 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. By the same token, a constant may be added to payoffs in any state without changing optimal thresholds.

EXAMPLES CONT'D. In our running examples, we shall usually consider the case where there are given two possible actions, a_1 = Invest and a_2 = Decline. 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 $u > 0$. From the indifference equation (1.1), we get $0 = \bar{r}u - (1-\bar{r})$ so $\bar{r} = 1/(1+u)$.

1.2.4 Individual Learning

We now consider how an individual's optimal decision rule incorporates the observed action history and his own private belief. In so doing, we could proceed inductively, and first derive the first individual's decision rule as a function of his private belief; next, we could describe how the second individual bases his decision on the private belief *and* on the first individual's action, and so on. Instead, we shall collapse this reasoning process, and simply say that an individual knows the decision rules of all predecessors, and acts accordingly.[13] He can use the common prior to calculate the ex ante (that is, as of time-0) probability of any action profile h in either of the two states. We shall denote these probabilities by $\pi^H(h)$ and $\pi^L(h)$, and let the resulting *likelihood ratio* that the state is L versus H (that is, *not H versus L*) be $\ell(h) = \pi^L(h)/\pi^H(h)$. Similarly, let $q(h)$ be the *public belief* that the state is H, i.e.

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

[13] Note that we have assumed common knowledge of rationality. Section 1.4 backs away from this, only assuming it among all rational individuals.

Think of $q(h)$ as the belief an individual facing the history h would entertain if he had a purely neutral private belief. Given the one-to-one relationship between q and $\ell = (1-q)/q$, we may also loosely refer to the likelihood ratio as the public belief.

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:

$$r = \frac{p\,\pi^H(h)}{p\,\pi^H(h) + (1-p)\,\pi^L(h)} = \frac{p}{p + (1-p)\,\ell(h)} = \frac{1}{1 + \frac{1-p}{p}\ell(h)} \tag{1.2}$$

Lemma 1.3 (Private Belief Thresholds) *After history h is observed, there exist threshold values $0 = \bar{p}_0(h) \leq \bar{p}_1(h) \leq \ldots \leq \bar{p}_M(h) = 1$, such that a_m is chosen exactly when the private belief satisfies $p \in (\bar{p}_{m-1}(h), \bar{p}_m(h)]$, where for all $m = 0, \ldots, M-1$,*

$$\frac{\bar{p}_m(h)}{1 - \bar{p}_m(h)} = \frac{\bar{r}_m}{1 - \bar{r}_m}\ell(h) \tag{1.3}$$

The proof is simple. The thresholds simply come from a well-known reformulation of (1.2) as posterior odds $(1-r)/r$ equal the private odds $(1-p)/p$ times the likelihood ratio $\ell(h)$. The strict inequalities are consequences of the tie-breaking rule, and the fact that (1.2) is strictly increasing in p. If $\bar{p}_{m-1}(h) = \bar{p}_m(h)$ at some history h, then action a_m is not taken.

Later, when referring to (1.3) and elsewhere, we shall suppress the explicit dependence of $\ell(h)$ on h, and write $\bar{p}_m(\ell)$ instead of $\bar{p}_m(h)$. This is only a mild abuse of notation because the likelihood ratio is a sufficient statistic for the history, with respect to the decision to be made by the individual. Written as such, $\ell \mapsto \bar{p}_m(\ell)$ is an increasing function.

Examples Cont'd. in both examples, (1.3) yields $\bar{p}(\ell) = \ell/(u+\ell)$ if $\bar{r} = 1/(1+u)$.

1.2.5 Corporate Learning as a Markov-Martingale Process

We shall denote the public likelihood ratio and belief after individual n acts as ℓ_n and q_n, respectively.[14] Since the first individual has not observed any history, we shall normalize $\ell_0 = 1$. As signals, and thereby actions, are random, both $\langle \ell_n \rangle_{n=1}^{\infty}$ and $\langle q_n \rangle_{n=1}^{\infty}$

[14]Throughout the paper, m subscripts will denote actions, and n subscripts individuals.

are stochastic processes, and it is important to understand their dynamics. Define

$$\rho(m|s,\ell) = F^s(\bar{p}_m(\ell)) - F^s(\bar{p}_{m-1}(\ell)) \tag{1.4}$$

$$\varphi(m,\ell) = \ell\rho(m|L,\ell)/\rho(m|H,\ell) \tag{1.5}$$

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 a_m, we move to $\ell_{n+1} = \varphi(m, \ell_n)$. Figure 1-3 schematically summarizes this transition.

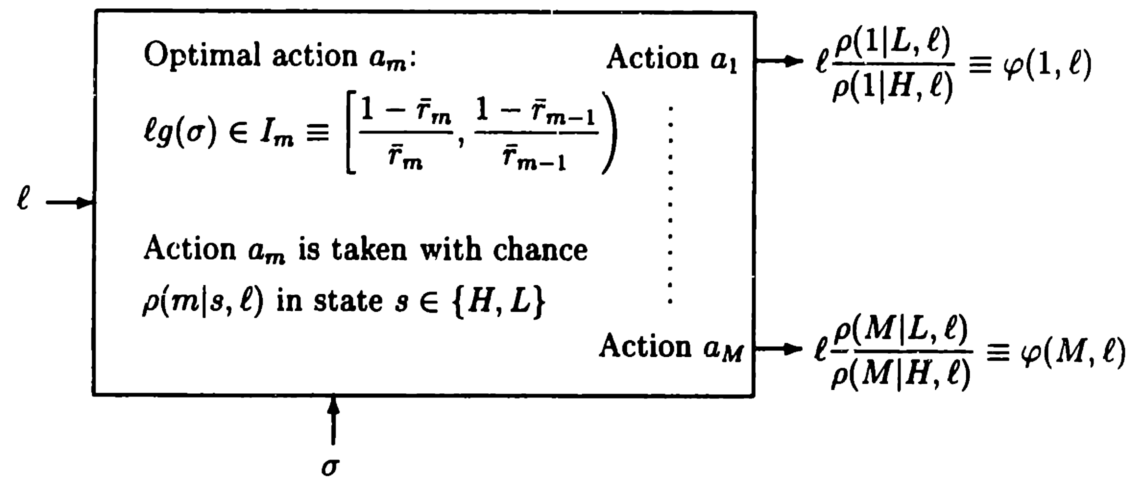

Figure 1-3: **Individual Black Box.** An individual bases his decision on both ℓ and his private signal σ, resulting in his action choice a_m and a likelihood ratio to confront successors.

Many of our insights derive from the simple fact that our dynamics constitute a time homogeneous *Markov process* in (m_n, ℓ_n), where m_n is the index of the action choice of individual n, and ℓ_n the public likelihood ratio resulting from his decision. Thus, the state space is $\mathcal{M} \times [0, \infty)$. Given (m_n, ℓ_n), (1.4) and (1.5) imply that the next state is $(m_{n+1}, \varphi(m_{n+1}, \ell_n))$ with probability $\rho(m_{n+1}|H, \ell_n)$ in state H. We wish to characterize the ergodic properties of this Markov process: In the long run, what actions occur, and what beliefs are held? Note that in principle, such a two-dimensional process could be very ill-behaved, for continuous state variables can embed chaotic processes, or even worse. Fortunately, Lemma 1.5 will attest to how well behaved is the likelihood ratio.

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

Lemma 1.4 (The Unconditional Martingale) *The public belief* $\langle q_n \rangle$ *is a martingale,* *unconditional on the state of the world.*[15]

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 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.

Given that $\langle q_n \rangle$ is expected to increase in state H, it is comforting and maybe surprising that $\langle \ell_n \rangle = \langle (1 - q_n)/q_n \rangle$ remains constant in expectation. This result is, of course, well-known from statistics.[16] Yet perhaps because of the singular focus on experimentation models in which incomplete learning was a distinct possibility, this is (to the best of our knowledge) its first application to the economics learning literature. At any rate, likelihood ratios will be the only workhorse in our paper, and we intend to showcase their power.

Lemma 1.5 (The Conditional Martingale) *Conditional on the state of the world H (resp. state L), the likelihood ratio $\langle \ell_n \rangle$ (resp. $\langle 1/\ell_n \rangle$) is a martingale.*

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

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

Lemma 1.6 *In state H (resp. state L), the likelihood ratio $\langle \ell_n \rangle$ (resp. $\langle 1/\ell_n \rangle$) converges almost surely to a limiting stochastic variable, $\ell_\infty = \lim_{n \to \infty} \ell_n$, with $\operatorname{supp}(\ell_\infty) = [0, \infty)$. Thus, fully incorrect learning, or $\ell_n \to \infty$ in state H, cannot occur with positive probability.*

Proof: Since the likelihood ratios are non-negative random variables, the result follows from the Martingale Convergence Theorem. (See Breiman (1968), Theorem 5.14.) □

[15] 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.

[16] See, for instance, Doob (1953), section II.7. The idea is simply that the likelihood ratio is the Radon-Nikodym derivative of the state L distribution against the state H distribution of the public history.

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.[17]

Easley and Kiefer (1988), among others, underscore that unlike in the statistics literature, complete learning is not at all a foregone conclusion in a costly experimentation model. For in a stationary statistical learning model, information keeps accruing at a 'constant rate', whereas experimentation has its price in economic settings, and complete learning is generally deemed too expensive. The resulting pathological outcomes to the learning dynamics persist in models of observational learning.

Herding, Cascades, and Complete Learning.

In the sequel, our goals are two-fold. For definiteness, let the state be H.

1. Is there *convergence of actions*? In the terminology of the introduction, do *herds* arise? With a finite action space, this is convergence of the first (action) coordinate projection of the Markov process (m_n, ℓ_n): Does m_n settle down — and if so, at $m_n = M$? And does this occur in finite expected time?

2. As it happens, we must first investigate *convergence of beliefs,* and characterize the limiting stochastic variable ℓ_∞. Must a *cascade* arise, so that $\rho(m|H, \ell_n) = \rho(m|L, \ell_n) = 1$ for some m after some stage n? In other words, will individuals eventually take action a_m irrespective of signal realizations? Or will only a *limit cascade* on some action a_m arise, i.e. where $\rho(m|H, \ell_n) \to 1$ as $n \to \infty$. Limit cascades thus refer to convergence of the second (belief) coordinate of the Markov process (m_n, ℓ_n). Finally, is learning *complete* — or do beliefs converge to the truth, i.e. $\ell_n \to 0$? Otherwise, if posterior beliefs are not eventually arbitrarily focused on the true state of the world, then we shall say that learning is *incomplete.*

3. What is the *link* between action and belief convergence? Observe the logic of BHW, valid with their discrete signal space: (i) cascades must occur, and (ii) cascades imply herds, and thus herds occur. The second step is irrefutable: If a cascade on action a_m arises at stage n, then by Lemma 1.3, individual n must take action a_m; therefore, his action reveals no private information, and so $\ell_{n+1} = \ell_n$. The cascade on action a_m still obtains

[17]Another proof of this fact uses public beliefs (see, for instance, Bray and Kreps (1987)).

at stage $n+1$, as private belief thresholds are unchanged by (1.3). So $\rho(m|H, \ell_{n+1}) = 1$ too.

But, as we shall soon see, (i) is generally false: Generically, cascades need not arise with bounded beliefs, and cannot arise with unbounded beliefs, as $(0, \infty) \subseteq (\bar{p}_{m-1}(\ell_{n+1}), \bar{p}_m(\ell_{n+1})]$ is impossible. This leads to the more involved question: Will limit cascades imply herds?

1.3. 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.

1.3.1 Belief Convergence

Lemma 1.6 assures us that $\ell_\infty = \lim_{n\to\infty} \ell_n$ almost surely exists. Theorem A-1.2 then yields a precise characterization of ℓ_∞: Any point $\hat{\ell} \in \text{supp}(\ell_\infty)$ must be stationary for the Markov process, i.e. either $\rho(m|\hat{\ell}) = 0$ or $\varphi(m, \hat{\ell}) = \hat{\ell}$. We now exhibit a salient reason for stationarity, where all learning stops, and individuals act irrespective of their signals.

Lemma 1.7 (Action Absorbing Basins) *For each action a_m, there is a possibly empty interval $J_m = \{\ell \,|\, \text{supp}(F) \subseteq [\bar{p}_{m-1}(\ell), \bar{p}_m(\ell)]\} \subset [0, \infty]$, such that when $\ell \in \text{int}\,(J_m)$ the individual takes action a_m with probability one, and the likelihood ratio is unchanged. Also,*
(i) With bounded private beliefs, $J_1 = [\bar{\ell}, \infty]$ and $J_M = [0, \underline{\ell}]$ for some $0 < \underline{\ell} < \bar{\ell} < \infty$;
(ii) 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.

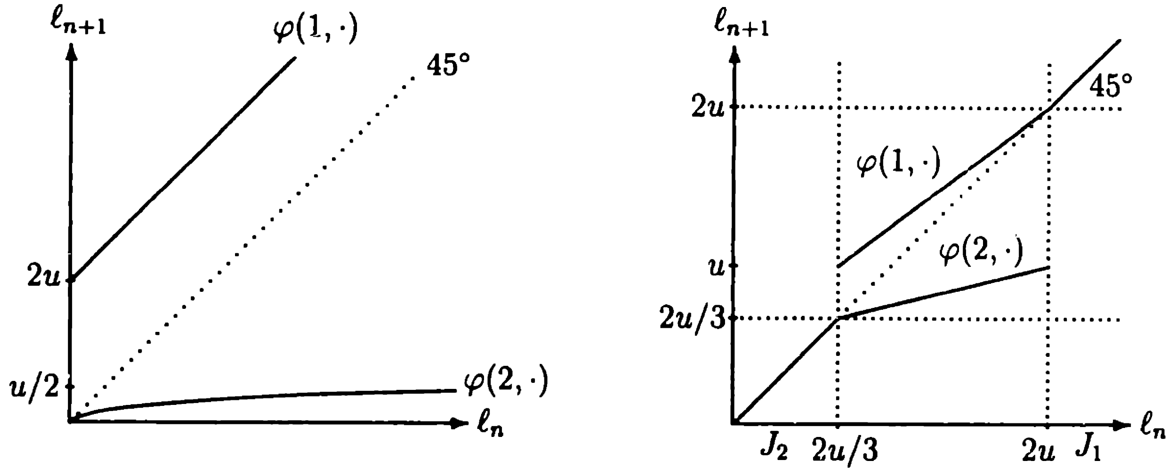

Figure 1-4: **Continuations and Absorbing Basins.** Continuation functions for the examples: unbounded private beliefs (left), and bounded private beliefs (right). By the martingale property, the expected continuation lies on the diagonal. The stationary points are where both arms hit the diagonal (impossible here), or where one arm is taken with zero chance ($\ell = 0$ in the left panel; $\ell = 2u/3$ in the right).

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 (1.3), one can show that ℓ is in J_m precisely when

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

Thus, *an action absorbing basin is larger the smaller is the support* $[\underline{b}, \bar{b}]$, *and the larger is the interval* $[\bar{r}_{m-1}, \bar{r}_m]$. So for fixed $\varepsilon > 0$ above, the absorbing basin eventually vanishes as the bounded support $[\underline{b}, \bar{b}]$ tends to $(0, \infty)$. More generally, for fixed preferences, action absorbing basins only obtain on the extreme actions for large enough $[\underline{b}, \bar{b}]$.

UNBOUNDED BELIEFS EXAMPLE CONT'D. We have $F^H(p) = p^2$ and $F^L(p) = 2p - p^2$ for $p \in (0,1)$. Since $\text{supp}(F) = [0,1]$, the private beliefs are unbounded, and 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 1-4.

BOUNDED BELIEFS EXAMPLE CONT'D. We have $F^H(p) = (5p-2)/2p$ 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 $\ell \leq 2u/3$, 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 $\ell \geq 2u$, 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 1-4.)

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

Theorem 1.1 (Limit Cascades and Incomplete Learning) *Assume the private beliefs are bounded, and let $\ell_n \to \ell_\infty$. Then $\ell_\infty \in J_1 \cup \cdots \cup J_M$ almost surely. Moreover, in state H, if $\ell_0 \notin J_M$, then with positive chance $\ell_\infty \in J_1 \cup \cdots \cup J_{M-1}$.*

Proof: A complete proof of the first claim is in Appendix 1.E. We proceed here under the simplifying assumption that ρ and φ are continuous in ℓ. By Theorem A-1.1, 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 $F^H(\bar{p}_m(\hat{\ell})) = F^L(\bar{p}_m(\hat{\ell})) = 0$. 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 A-1.1, this equality is only possible if $F^H(\bar{p}_m(\hat{\ell})) = F^L(\bar{p}_m(\hat{\ell})) = 1$. Thus, $\hat{\ell} \in J_m$.

Now for the second statement, we just have to rule out $\ell_\infty \in J_M$ almost surely. If $\ell_\infty \in J_1$ with positive probability, we are done. Otherwise, if $\ell_\infty \notin J_1$ almost surely, then $\langle \ell_n \rangle$ is uniformly bounded above by $\bar{\ell}$, the infimum of J_1, introduced in Lemma 1.7. By Lebesgue's Dominated Convergence Theorem, the mean of $\langle \ell_n \rangle$ is preserved in the limit, i.e. $E[\ell_\infty] = \ell_0$. So, if $\ell_0 > \underline{\ell}$ it cannot be the case that $\text{supp}(\ell_\infty) \subseteq J_M = [0, \underline{\ell}]$. □

Next we present the counterpart to Theorem 1.1 that was not considered in Banerjee (1992) and BHW. Strictly bounded beliefs turns out to have been the mainstay for their striking pathological herding results.

Theorem 1.2 (Complete Learning) *If the private beliefs are unbounded then almost surely $\ell_n \to 0$ in state H, and $\ell_n \to \infty$ in state L.*

Proof: As usual, assume WLOG the state is H. By the same proof as for Theorem 1.1, $\ell_\infty \in J_1 \cup \cdots \cup J_M$ a.s.. By Lemma 1.7, with unbounded beliefs $J_1 \cup \cdots \cup J_M = \{0\} \cup \{+\infty\}$. Finally, Lemma 1.6 proves $\text{supp}(\ell_\infty) \subseteq [0, \infty)$ and we can conclude that $\ell_\infty = 0$ a.s.. □

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 1.2 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 1.2, 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.

1.3.2 Action Convergence

We shall now argue that limit cascades imply herds. Suppose that a long string of individuals has taken the identical action, but that we are not yet in a cascade. If someone then takes a contrary action, his successors have no choice but to concede the strength of his signal, and sharply revise the public likelihood. We say that the would-be herd has been *overturned* by the unexpected action. We shall more rigorously formulate it as the *overturning principle,* as it proves central to an understanding of the observational learning paradigm: Assume individual n chooses action a_m. Then individual $n+1$ should, *before* he observes his own private signal, find it optimal to choose action a_m because he knows no more than n, and since it is common knowledge that n rationally chose a_m. So, after individual n's action, the public belief satisfies $q(h, a_m) \in (\bar{r}_{m-1}, \bar{r}_m]$, which is the content of the next lemma.

Lemma 1.8 (The Overturning Principle) *For any history h, if an individual optimally takes action a_m, then the updated likelihood ratio must satisfy*

$$\ell(h, a_m) \in \left[\frac{1-\bar{r}_m}{\bar{r}_m}, \frac{1-\bar{r}_{m-1}}{\bar{r}_{m-1}}\right) \tag{1.6}$$

Proof: Since individual n optimally chose action m, his signal σ_n must have satisfied

$$\frac{1-\bar{r}_{m-1}}{\bar{r}_{m-1}} > \ell(h)g(\sigma_n) \geq \frac{1-\bar{r}_m}{\bar{r}_m} \tag{1.7}$$

Let $\Sigma(h)$ denote the set of all signals σ_n that satisfy (1.7). 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). Thus,

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

Now just cross-multiply, and use inequality (1.7) to bound the right hand integral. □

EXAMPLES. In our running examples, with the two actions, we have $\bar{r} = 1/(1+u)$, so the overturning principle predicts that $\varphi(1,\ell) \in [u, \infty)$ and $\varphi(2,\ell) \in [0, u)$ must hold. This is indeed the case in our examples so far, and in figure 1-4.

The overturning principle will imply that contrary actions cannot forever arise in a limit cascade: Or, convergence in beliefs implies convergence in actions.

Theorem 1.3 (Herds) *Assume that private beliefs are bounded. Then a herd on some action will almost surely arise in finite time. Absent a cascade on the most profitable action a_M from the outset, a herd arises on an action other than a_M with positive probability.*

Proof: We need only combine Theorem 1.1 and Lemma 1.8. 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, (1.3) 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. □

So the bottom line is that all individuals eventually stop paying significant heed to their private signals. The result is sometimes a 'correct' and sometimes an 'incorrect' herd. This is essentially the major pathological learning result obtained by Banerjee (1992) and BHW, albeit extended to $M > 2$ actions.[18] Our new result completes the characterization.

Theorem 1.4 (Correct Herds) *If the private beliefs are unbounded, then almost surely individuals eventually all take the optimal action.*

Proof: Assume WLOG that the state is H, so that a_M is optimal. Theorem 1.2 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 1.8, whenever any other action than action a_M is taken, we exit that neighborhood of 0. □

1.3.3 Continuity From Bounded To Unbounded Beliefs

We now show that transition between incomplete learning with bounded private beliefs, and complete learning with unbounded private beliefs, is continuous.

Theorem 1.5 (Continuity) *Fix the payoffs and the prior beliefs. If* co(supp(F^k)) *converges to* $[0,1]$,[19] *then the chance of an incorrect limit cascade vanishes as* $k \to \infty$.

Proof: With fixed preferences, only the two extreme action 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 $\pi^k \to 1$. □

1.3.4 Must Cascades Exist with Bounded Beliefs?

Because $\langle m_n, \ell_n \rangle$ is not a *finite state* Markov chain, Theorem 1.1 does not imply convergence in finite time — for we could conceivably have $\ell_n \to \hat{\ell} \in J_1 \cup \cdots \cup J_M$ but

[18]The analysis of BHW also handled more states. We address this generalization in appendix 1.D.
[19]We mean the Hausdorff topology: If co(supp(F^k)) $= [a_k, b_k]$, then $a_k \to 0$, and $b_k \to 1$.

$\ell_n \notin J_1 \cup \cdots \cup J_M$ for all n. We now show that cascades (convergence in finite time) needn't obtain, even though herds must. In light of BHW, this may well come as a surprise.

In the running BOUNDED BELIEFS EXAMPLE, a cascade on action a_1 obtains iff $\ell \geq 2u$. Posit a limit cascade on a_1, with $\ell_n \to \hat{\ell} \in [2u, \infty)$, so that a herd on a_1 obtains, by Theorem 1.3. Then figure 1-4 clearly shows that *despite a herd on a_1, ℓ_n never enters J_1, but only approaches its edge. A cascade never starts.*[20] So learning never ceases: No matter how long a string of individuals have followed suit, one can never rule out a contrarian.

BOUNDED BELIEFS EXAMPLE CONT'D. Given $\ell_0 = 1$, we can calculate the exact probability π of a correct limit cascade in state H. Theorem 1.1 tells us that a limit cascade arises a.s., and by the above reasoning about this example, only $\ell_\infty \in \{2u, 2u/3\}$ is possible. (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 $|\ell_n| \leq 2u$; therefore, the identity $1 = \pi(2u/3) + (1-\pi)(2u)$ implicitly defines π whenever $2u/3 < 1 < 2u$.

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 $\operatorname{int}(J_m) \neq \varnothing$, then[21]

$$\ell_{n+1} = \varphi(m, \ell_n) \equiv \ell_n \frac{F^L(\bar{p}_m(\ell_n)) - F^L(\bar{p}_{m-1}(\ell_n))}{F^H(\bar{p}_m(\ell_n)) - F^H(\bar{p}_{m-1}(\ell_n))} = \ell_n \frac{F^L(\bar{p}_m(\ell_n))}{F^H(\bar{p}_m(\ell_n))} \tag{1.8}$$

This implies $\ell_{n+1}/\ell_n > \inf\{F^L(p)/F^H(p) | p \in \operatorname{int}(\operatorname{supp}(F))\} > 1$ by Lemma A-1.1, and so $\langle \ell_n \rangle$ must 'jump into' J_m in boundedly finite time.

One might prematurely conclude that cascades can *only* arise with discrete distributions, and thus are nongeneric phenomena. The next example disproves this conjecture.

CASCADES WITH SMOOTH SIGNALS. For a discrete jump into an absorbing basin, for instance $[2u, \infty)$ in figure 1-4, 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-1.1, 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 $g' > 0$,

[20] 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.
[21] Note that $F^L(\bar{p}_{m-1}(\ell_n)) = F^H(\bar{p}_{m-1}(\ell_n)) = 0$ since $\ell_n < \hat{\ell} \equiv \inf(J_m)$.

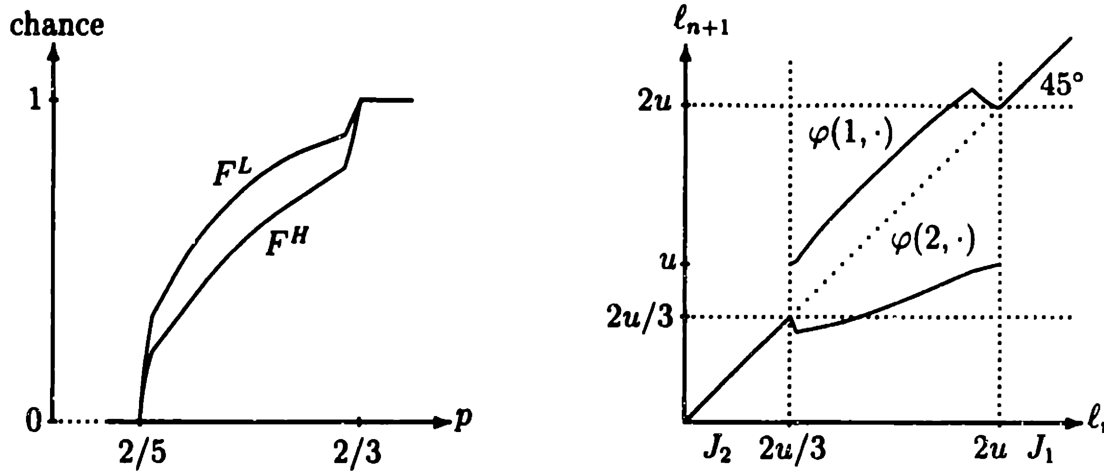

Figure 1-5: **Nonmonotonicity.** The left graph exhibits F^H and F^L that work in the CASCADES WITH SMOOTH-SIGNALS EXAMPLE. As they 'nearly' have atoms at the edge of their supports, they can mimic the effect in BHW: With lumpy information, actions are boundedly informative, and so a single decision can toss all successors into a cascade. The right graph shows how the corresponding continuation functions for the likelihood ratio are no longer monotonic.

$g(0) > 0$, 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 $F^L(p) = \int_{\sigma(p)}^1 g(\sigma)d\sigma$. Since $\inf(J_1) = u/g(0)$,[22] we 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 $g(0) \in (0,1)$. Figure 1-5 gives such an example.[23]

1.3.5 How Fast Do Beliefs Converge?

We have seen that the likelihood ratio converges almost surely. We now turn to the speed of this convergence — a topic we are interested in not just for its own sake, but also rather for what it will tell us about action convergence. This motivation suggests that we should investigate the convergence rate of something more tangible than the likelihood ratios. Let $e(m, \ell_n)$ be the *exit rate* from a current would-be herd on action a_m given likelihood ℓ_n. This is the chance that a contrary action is taken. When it is understood, suppress m and abbreviate $e_n \equiv e(m, \ell_n)$. Let $E_n \equiv (1-e_1)\cdots(1-e_n)$ be

[22]Action a_2 requires that $\ell g(\sigma) \le u$, and since $g(\sigma) \le g(0)$, we have $\inf(J_1) = u/g(0)$.

[23]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 $\sigma \in [3/4,1]$.

the chance that the would-be herd does not end by individual n. Clearly, convergence requires that a (permanent) herd obtains with positive probability, which occurs exactly when $E_\infty \equiv \lim_{n\to\infty} E_n > 0$. It is well-known asymptotic result that this holds exactly when the $\langle e_n \rangle$ are summable: $\sum_n e_n < \infty$.[24]

Before venturing further, say that a deterministic sequence $\langle x_n \rangle$ converges to $\bar{x}$ at *rate* $\theta \in [0,1]$ if $\hat{\theta}^{-n}|x_n - \bar{x}| \to 0 \; \forall \hat{\theta} > \theta$. Namely, $x_n \to \bar{x}$ 'at least as fast as' $\theta^n \to 0$. So convergence at rate θ clearly implies convergence at any rate $\theta' \in (\theta, 1]$ too. Convergence at any rate $\theta < 1$ is known as *exponential* (or exponentially fast). So convergence at rate 0 is a well-defined notion, while convergence at rate 1 is possible but perhaps not exponential. Note that the exponential convergence of $\langle e_n \rangle$ to 0 is sufficient but not necessary for $\sum_n e_n < \infty$.

To determine the convergence rate of $\langle e_n \rangle$, we need further assumptions. Say that *extreme signals are rare* if $\liminf_{p \searrow \underline{b}} F^H(p)/[p-\underline{b}] = 0$ or $\liminf_{p \nearrow \bar{b}} [1-F^L(p)]/[\bar{b}-p] = 0$, where $\mathrm{co}(\mathrm{supp}(F)) = [\underline{b}, \bar{b}]$. For instance, if there exists *continuous* densities f^H and f^L, then extreme signals are rare iff $f^H(\underline{b}) = 0$ or $f^L(\bar{b}) = 0$. Lemma 1.1 implies that extreme signals *must* be rare with unbounded beliefs.

Lemma 1.9 (Belief Convergence) *With bounded beliefs, if extreme signals are not rare, $\langle e_n \rangle$ converges to $\underline{b}$ or $\bar{b}$ at rate $\theta \in [0,1)$. With unbounded beliefs, $\langle e_n \rangle$ converges to 0 in state H at no rate $\theta < 1$.*

Proof: PART I: BOUNDED BELIEFS. Consider a limit cascade realization $\ell_n \to \hat{\ell} \in J_m \neq \varnothing$ corresponding to a herd on action a_m. Either a cascade occurs and $\ell_n \in \mathrm{int}(J_m)$ in finite time, and so $\theta = 0$, or $\hat{\ell}$ is on the rim of the interval J_m, say $\hat{\ell} = \inf\{\ell : \ell \in J_m\}$. In that case, once the herd has begun, $\ell_{n+1} = \varphi(m, \ell_n)$ and $\langle \ell_n \rangle$ satisfies (1.8). Thus, for $\ell_n < \hat{\ell}$,[25]

$$0 < \varphi(m, \hat{\ell}) - \varphi(m, \ell) = (\hat{\ell} - \ell) - \ell \left[\frac{F^L(\bar{p}_m(\ell))}{F^H(\bar{p}_m(\ell))} - 1 \right] \tag{1.9}$$

Appealing to Lemma A-1.2, it follows from the non-rare extreme signals assumption that there exist an $\varepsilon > 0$ so $F^L(\bar{p}_m(\ell))/F^H(\bar{p}_m(\ell)) - 1 > \varepsilon(\bar{p}_m(\hat{\ell}) - \bar{p}_m(\ell))$, once $\bar{p}_m(\ell)$ is

[24]Indeed, $e^{-x} \geq 1 - x$ shows that $E_\infty > 0 \Rightarrow \sum_n e_n < \infty$. The reverse direction is somewhat trickier.

[25]We have exploited the fact that by the definition of J_m, we have $F^H(\bar{p}_m(\hat{\ell})) = F^L(\bar{p}_m(\hat{\ell})) = 1$.

sufficiently close to $\bar{p}_m(\hat{\ell})$. On the other hand, $\bar{p}_m$ is a differentiable function of ℓ, with a bounded derivative near $\hat{\ell}$. Thus, there exist $\delta > 0$ such that $\ell(F^L(\bar{p}_m(\ell))/F^H(\bar{p}_m(\ell))-1) > \delta(\hat{\ell}-\ell)$ for ℓ close enough to $\hat{\ell}$; so $\varphi(m,\cdot)$ is Lipschitz with constant $\theta = 1-\delta$ near $\hat{\ell}$. This implies that $\langle \ell_n \rangle$ converges at rate θ. A simple argument (omitted in this version) then shows that $\langle e_n \rangle$ also vanishes exponentially fast.

PART II: UNBOUNDED BELIEFS. In state H, $\ell_n \to 0$ and eventually everyone chooses action a_M, and so

$$\ell_{n+1} = \varphi(M, \ell_n) = \ell_n \frac{1 - F^L(\bar{p}_{M-1}(\ell_n))}{1 - F^H(\bar{p}_{M-1}(\ell_n))}$$

For now we shall assume that *extreme signals are not very rare at* 0: i.e. there exists $K > 0$ so that for all $p \in (0,1)$, $F^L(p)/p < K$, and similarly, *extreme signals are very rare at* 1 if there exist $K > 0$ so that for all $p \in (0,1)$, $(1 - F^H(p))/(1-p) < K$. By Lemma A-1.2, with unbounded beliefs and very rare extreme signals, there exist $\varepsilon > 0$ such that $(1-F^L(p))/(1-F^H(p)) > 1-\varepsilon p$, once p is close enough to zero. Since $\bar{p}_{M-1}(\ell) = \ell/(u+\ell)$ for some $u > 0$, we can conclude that $\langle \ell_n \rangle$ eventually satisfies $\ell_{n+1} > \ell_n - \delta\ell_n^2$ for some $\delta > 0$. Fix n so large that $\ell_{n+1} > \ell_n - \delta\ell_n^2$ and $1 - 2\delta\ell_n > 0$, and take then N so large that $\ell_n \geq 1/((N+n)\log(N+n))$. Then we get $\ell_{n+1} \geq 1/((N+n+1)\log(N+n+1))$. Thus $\sum_n 1/(n\log(n)) = \infty$. (proof to be finished) □

1.3.6 How Fast do Actions Converge? Mean Time To Herding

The power of the Markovian analysis now comes to the fore, in answer to the natural question: 'When does the herd start?'[26] As the dynamics generically are not absorbed in our state space $\mathcal{M} \times [0,\infty)$, this is not an off-the-shelf calculation. Rather, we wish to know when the discrete component $\{m_n\}$ settles down, and we believe our approach is new. Given the Markovian state (m,ℓ), let $\tau(m,\ell)$ be the random waiting time until a herd begins (i.e. the last individual whose action differs from his predecessor). Theorems 1.3 and 1.4 imply that $\tau(m,\ell) < \infty$ a.s., but we also know that $\tau(m,\ell)$ has no uniform upper bound. So is the *mean waiting time* $\bar{\tau}(m,\ell) = E[\tau(m,\ell)|(m,\ell)]$ *until a herd begins* finite?

[26]For this subsection, we wish to acknowledge generous mathematical assistance from Peter Matthews (University of Maryland Baltimore County), Bertram Walsh (Rutgers University), and others on the usenet group **sci.math.research**.

To proceed, consider an *immediate herd* on action a_m, i.e. where a_m is taken from now on. Recursively define the sequence of likelihood ratios in such an immediate herd, via $\varphi^0(m,\ell) = \ell$, and $\varphi^{n+1}(m,\ell) = \varphi(m,\varphi^n(m,\ell))$. Then the *probability that a herd now obtains on action a_m in state (m,ℓ)* is clearly $\alpha(m,\ell) = \prod_{n=0}^{\infty}\psi(m|\varphi^n(m,\ell))$. With unbounded beliefs, because a herd must arise a_M, we have $\alpha(m,\ell) = 0$ for all $m \neq M$.

UNBOUNDED BELIEFS EXAMPLE CONT'D. We can explicitly calculate $\alpha(2,\ell)$. To find $\ell_n = \varphi^n(2,\ell)$, we solve the difference equation $\ell_{n+1} = \varphi(2,\ell_n)$. But since $\varphi(2,\ell) = u\ell/(u+2\ell)$, the substitution $w = (u+\ell)/\ell$ reduces the recursion to $w_{n+1} = 2 + w_n$, or $w_n = w_0 + 2n$. Substituting this into $\psi(2|\ell) = 1 - \ell^2/(u+\ell)^2 = 1 - 1/w^2$ yields $\psi(2|\varphi^n(2,\ell)) = 1 - 1/(w_o + 2n)^2$. Finally, we discover[27]

$$\alpha(2,\ell) = \prod_{n=0}^{\infty}\psi(2|\varphi^n(2,\ell)) = \frac{[\Gamma(w/2)]^2}{\Gamma((w+1)/2)\Gamma((w-1)/2)} = \frac{[\Gamma((u+\ell)/2\ell)]^2}{\Gamma((u+2\ell)/2\ell)\Gamma((u/2\ell)}$$

where Γ is the Gamma function. In particular, notice that with the neutral belief $\ell = u$ (eg. $u = 1$ and fair prior $\ell_0 = 1$), we get the pretty expression $\alpha(2,u) = 2/\pi$, where $\pi = 3.14159\ldots$. So the chance of a correct herd from the outset can be quite substantial.

We shall now characterize $\bar{\tau}$ in terms of α. Define the functional operator V from the space of nonnegative functions into itself:

$$(V \circ f)(m,\ell) = 1 - \alpha(m,\ell) + \sum_{k=1}^{M}\psi(k|\ell)f(k,\varphi(k,\ell)) \tag{1.10}$$

Lemma 1.10 *The function $\bar{\tau}$ is the smallest (possibly infinite) nonnegative fixed point of V.*

Proof: First, $\bar{\tau}$ actually is a fixed point of V, because

$$\begin{aligned}\bar{\tau}(m,\ell) = E[\tau(m,\ell)] &= \sum_{n=1}^{\infty}\Pr(\tau(m,\ell)\geq n) = \Pr(\tau(m,\ell)\geq 1) + \sum_{n=2}^{\infty}\Pr(\tau(m,\ell)\geq n) \qquad (1.11)\\ &= 1 - \Pr(\tau(m,\ell)=0) + \sum_{n=2}^{\infty}\Pr(\tau(m,\ell)\geq n|\tau(m,\ell)\geq 1)\Pr(\tau(m,\ell)\geq 1)\\ &= 1 - \alpha(m,\ell) + E[\tau(m,\ell)|\tau(m,\ell)\geq 1] = (V\circ\bar{\tau})(m,\ell)\end{aligned}$$

[27]See section 12-13 of Whittaker and Watson (1963) for reductions of infinite products of terms like $[(n+(w_0+1)/2)(n+(w_0-1)/2)]/(n+w_0/2)^2$, as we have.

Next, we show (à la negative dynamic programming) that $\bar{\tau}$ is the smallest fixed point of V. Being a sum (1.11), we could in principle calculate $\bar{\tau}$ by iterating (1.10) starting with $f_0 = 0$. Since V is monotonic ($f_1 \leq f_2 \Rightarrow V \circ f_1 \leq V \circ f_2$), the resulting weakly increasing sequence of functions $f_n = V^n \circ f_0$ is pointwise convergent in $[0, \infty]$. So for any other f solving $V \circ f = f \geq f_0$, the monotonicity of V implies $f \geq f_n$ for all n. Thus $f \geq \bar{\tau}$. □

AN ILLUSTRATION. We now apply the functional equation to solve analytically for the mean time in a $2 \times 2 \times 2$ example (two states, two signals, two actions) from BHW.

- signals s_1 and s_2, resulting in $F^L(1-p) = p = 1 - F^H(1-p)$, some $p \in (1/2, 1)$
- two actions a_1 and a_2, with a_2 preferred iff the posterior belief exceeds $1/2$

The prior belief is $q = 1/2$, $\ell_0 = 1$. With BHW's special tie-break rule that indifferent individuals choose the action opposite their predecessor, updating ℓ_n is simple, for the action reveals the signal For instance, when $\ell_n = 1$, signal s_2 leads to action a_2, and thus $\ell_{n+1} = (1-p)/p$; in that case, one further signal s_2 results in a_2 and thus $\ell_{n+1} = [(1-p)/p]^2 \in J_2$, while a_1 leads back to the neutral $\ell_{n+1} = 1$. So,

- α is fully described by $\alpha(2, [(1-p)/p]^2) = \alpha(1, [p/(1-p)]^2) = 1$, $\alpha(2, (1-p)/p) = p\alpha(2, [(1-p)/p]^2) = p$, $\alpha(1, p/(1-p)) = 1-p$, $\alpha(1,1) = (1-p)^2$, $\alpha(2,1) = p^2$.
- Clearly, for ℓ in the absorbing basins, $\bar{\tau}(2, [(1-p)/p]^2) = \bar{\tau}(1, [p/(1-p)]^2) = 0$. Lemma 1.10 also yields four equations in four unknowns: $\bar{\tau}(2, (1-p)/p) = 1 - p + (1-p)\bar{\tau}(1,1)$, $\bar{\tau}(1, p/(1-p)) = p + p\bar{\tau}(2,1)$, $\bar{\tau}(1,1) = 1 - (1-p)^2 + p\bar{\tau}(2, (1-p)/p) + (1-p)\bar{\tau}(1, p/(1-p))$, and $\bar{\tau}(2,1) = 1 - p^2 + p\bar{\tau}(2, (1-p)/p) + (1-p)\bar{\tau}(1, p/(1-p))$.

These four equations have a straightforward, if messy, algebraic solution. When p approaches 1, we approach unbounded beliefs, and the probability of immediately being in a herd, $(1-p)^2 + p^2$ converges to one. Similarly, the mean time to herding converges to zero, as, in the limit, $\bar{\tau}(2, (1-p)/p)), \bar{\tau}(2,1) \to 0$ and $\bar{\tau}(1, p/(1-p)), \bar{\tau}(1,1) \to 1$.

We wish to apply the sharp characterization of $\bar{\tau}$ to deduce another knife-edge result.

Conjecture (Mean Time to Herd) *The expected time $\bar{\tau}(m, \ell)$ until a herd starts is boundedly finite or always infinite exactly as the private beliefs are bounded or unbounded.*

Support for this conjecture comes from the following analyses of our examples.

BOUNDED BELIEFS EXAMPLE CONT'D. Partition the range of possible likelihood ratios into $K_2 = [2u/3, u)$ and $K_1 = [u, 2u]$, where we take actions a_2 and a_1, respectively. As there is a positive probability of eventually herding on a_1, $\alpha(1,\ell) > 0$ in a region $\mathcal{N}$ of $2u$. Since $\alpha(1,\ell)$ is increasing in ℓ, it achieves its minimum over K_2 at $\ell = u$. From any $\ell \geq u$, there is positive chance of reaching $\mathcal{N}$ in a finite number of steps, so $\alpha(1,\ell) > \alpha(1,u) > 0$ for all $\ell > u$ (and so $\mathcal{N} = K_1$). Similarly, $\alpha(2,\ell) > \alpha(2,u) > 0$ for all $\ell < u$. Thus, with chance at least $\varepsilon = \min\langle\alpha(1,u), \alpha(2,u)\rangle$ a herd now obtains for any ℓ. But from any $\ell \in K_2$, there is a chance of at least $\varepsilon\rho(1,\ell)$ of starting a herd on a_1 — unconditional on whether a herd now obtains. This chance vanishes as ℓ approaches $2u/3$. Still we have

CLAIM: Conditional on not yet being in a herd, there is a boundedly positive chance of starting a herd immediately.

Proof: Our major concern is when ℓ is near $2u/3$ or $2u$ and no herd obtains. Suppose that action a_2 is taken repeatedly. Then ℓ_n approaches $2u/3$ exponentially fast, by Lemma 1.9; therefore, $e_n \equiv \rho(1,\ell_n)$ must vanish exponentially fast, or $e_n = \theta^n$. Since we are not yet in a herd by assumption, the conditional chance of taking a_1 from ℓ_n is of the form $e_n / \sum_{\nu=n}^{\infty} e_\nu = 1/\sum_{\nu=0}^{\infty} \theta^\nu > 0$.

Just as with the formula for the mean arrival time of a Poisson event, the claim implies that the expected time to jump into the herd is finite.

UNBOUNDED BELIEFS EXAMPLE CONT'D. As all terms in the sum (1.11) are non-negative, we shall simply find a subset of terms whose sum is infinite. WLOG, start the iteration from the state $(1,\ell_0)$ with $0 < \ell_0 < \infty$. Eventually a herd will start on action a_2, but $\tau(1,\ell_0) \geq N$ is true if, for instance, action a_1 is chosen the next N consecutive times. In that case, $\varphi(1,\ell) = \ell + 2u$ implies $\ell_n = \ell_0 + 2nu$. Together with $\psi(1|\ell) = \ell^2/(u+\ell)^2$, we see that $\tau(1,\ell_0) \geq N$ has chance at least

$$\prod_{n=0}^{N-1} \psi(1|\ell_n) = \prod_{n=0}^{N-1} \psi(1|\ell_0 + 2n) = \left(\prod_{n=0}^{N-1} \frac{\ell_0 + 2n}{u + \ell_0 + 2n}\right)^2$$

If $\ell_0 = 3$ and $u = 1$, then this reduces to $2^{2-4N}\binom{2N}{N}^2$. This is asymptotically $1/N$, and therefore this sequence of probabilities is not summable, so that $\bar{\tau}(1, \ell_0) = \infty$.

Conversely, we fully expect that the exponential convergence of $\langle \ell_n \rangle$ with bounded beliefs will yield a finite mean time to herd. It is possible that the distinction between exponential and non-exponential convergence provides the accurate criterion for finite/infinite mean time to herd. We are currently investigating this.

1.3.7 Related Literature

As noted, BHW and Banerjee first noticed the existence of herds with bounded beliefs, namely Theorem 1.3. We have shown that BHW's superficially stronger result than our Theorem 1.1, asserting that cascades and not just limit cascades occur, is not robust to general signal spaces.[28]

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.

We wish to make perfectly clear that our complete learning results are *in no way* related to recent work by Lee (1993). Lee's assumption that individuals have a continuous action space effectively allows for a one-to-one correspondence between signals and actions. Individuals can thus perfectly infer their predecessors' signals — whence the social learning problem is reduced to a single person statistical decision problem![29] It is essential that in a herding model the inference of predecessors' signals be garbled as it passes through an

[28]While BHW did not discuss the case of unbounded private signals, in their working paper, they introduce perfectly informative signals under the rubric of 'pseudo-cascades'. (We rule them out by our assumption of mutually a.c. information.) The deduction of complete learning in that context is *much* simpler, both conceptually and mathematically. For as soon as 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!

[29]To be fair, Banerjee makes a similar point in a passing observation.

action choice. Even for unbounded beliefs, this puts complete learning in jeopardy: As n grows, individual actions become an ever more garbled statistic for the true signal.[30] Perhaps inspired by this logic, Lee disposes of the idea that an unbounded signal support can be helpful, with his declaration (page 398) that *there is no loss of generality in our model with two signal values compared to a model with many signal values.* Of course, we have shown that the unbounded private beliefs assumption guarantees that complete learning *does* take place. Hopefully, this underscores how unexpected are our results!

1.4. NOISE

We now turn to the economic robustness of the existing theory, by striking at its central underpinnings. The pivotal role played by the overturning principle is rather unsettling. It does not seem 'reasonable' that such large weight be accorded a single individual's action; therefore, we now make the reasonable assertion that there is 'noise' in the system: A small percentage of individuals either deliberately (that is, they are 'crazy'), or by accident (i.e. they 'tremble') do not choose their optimal action. Without 'common knowledge of rationality', no action can have drastic effects, simply because the 'unexpected' is really expected to happen every now and then.

Two competing theses then seem plausible at this point:

HYPOTHESIS #1. The statistically constant nature of noisy individuals does not jeopardize the learning by the rational individuals in the long run, as it can be filtered out: If $\ell_n \to 0$ absent noise, then this is still true with the noise.

HYPOTHESIS #2. The learning will be incomplete, as the contributions of arbitrarily informed individuals will be lost to the background noise. The resulting model will then be tantamount to one with bounded beliefs.

In fact, we show that the first intuition is correct. Provided there is not too much noise, the belief convergence results of Theorems 1.1 and 1.2 still hold — but it is a harder

[30]In fact, Lee himself affirms this distinction (on page 397): *From the standpoint of information revelation, a sparse action 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.*

to resolve the issue of herds. And if the noise is paramount, then the learning may well break down.

1.4.1 Two Forms of Noise

Just to be clear, whether an individual is noisy is not public information; further, this trait is distributed independently across individuals. We shall be making some comparisons to the noise-free model treated till now. So, from the section 1.2 we recall the definition (1.4) that *rational* individuals take action m with chance $\rho(m|s,\ell) = F^s(\bar{p}_m(\ell)) - F^s(\bar{p}_{m-1}(\ell))$.

CRAZINESS. We first posit the existence of crazy individuals in the model. Assume that with probability κ_m, individual n will choose action a_m, irrespective of history. To avoid trivialities, we assume a positive fraction $\kappa = 1 - \sum_{m=1}^{M} \kappa_m > 0$ of 'sane' individuals. Equivalently, public history is observed with a little noise (common to all successors): For each individual, there is then a chance κ_m that action a_m is observed by all successors, independently of which action was actually chosen. In the language of Appendix 1.B, the dynamics of the likelihood ratio in state H are now described by

$$\psi(m|s,\ell) = \kappa_m + \kappa\rho(m|s,\ell) \tag{1.12}$$

$$\varphi(m,\ell) = \ell\psi(m|L,\ell)/\psi(m|H,\ell) \tag{1.13}$$

TREMBLING. In the second manifestation of noise, everyone is rational, but some may 'tremble', in the sense of Selten (1975). In particular, individuals randomly take a suboptimal action with some idiosyncratic chance. Given ℓ, an individual planning to take action a_m will instead opt for a_j with probability $\kappa_m^j(\ell)$ for $j \neq m$. Then we have $\kappa_m(\ell) = \sum_{j\neq m} \kappa_m^j(\ell)$ as the total chance of deviating from action a_m, given a_m is optimal. To rule out too much noise, assume that all $\kappa_m(\ell) < 1/2M$ for all m. On the other hand, to avoid completely trivializing the noise, we further insist that all $\kappa_m^j(\ell)$ be bounded away from 0. In state H, the dynamics are now given by (1.13) and

$$\psi(m|s,\ell) = [1-\kappa_m(\ell)]\rho(m|s,\ell) + \sum_{j\neq m} \kappa_j^m(\ell)\rho(j|s,\ell) \tag{1.14}$$

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.[31]

1.4.2 Convergence of Beliefs

With noise, $\langle \ell_n \rangle$ is still a martingale in state H, since $\ell = \sum_{m=1}^{M} \psi(m|H,\ell)\varphi(m,\ell)$. Also, the interval structure of the action absorbing basins $J_1,\ldots,J_M$ in $[0,\infty)$ is still valid. So when $\rho(m|H,\ell) = \rho(m|L,\ell) = 1$ for some m, then $\rho(m'|H,\ell) = \rho(m'|L,\ell) = 0$ for all $m' \neq m$, and thus *rational* individuals take action a_m a.s. Consequently, contrary actions will be adjudged ex post to have been noisy, and will simply be ignored.

Theorem 1.6 (Limit Cascades with Noise) *Augment the standard model by either type of noise, and assume the state is H. Then $\ell_n \to \ell_\infty$ for some random variable ℓ_∞ with* $\text{supp}(\ell_\infty) \subseteq [0,\infty)$. *With bounded beliefs, $\ell_\infty \in J_1 \cup \cdots \cup J_M$ almost surely, and $\ell_\infty \in J_M$ with probability less than 1 unless $\ell_0 \in J_M$. With unbounded beliefs, $\ell_\infty = 0$ almost surely.*

Proof: Posit craziness noise. The Martingale Convergence Theorem assures us that ℓ_∞ exists, and is almost surely finite. Also

$$\varphi(m,\ell) - \ell = \ell\,\kappa \frac{\rho(m|L,\ell) - \rho(m|H,\ell)}{\psi(m|H,\ell)}$$

and because (1.12) yields $\psi(m|H,\ell) > \varepsilon$ whenever $\rho(m|H,\ell) > \varepsilon/\kappa$, we have $\varphi(m,\ell) - \ell = 0$ is satisfied for $\ell \in \text{supp}(\ell_\infty)$ under exactly the same circumstances as in the proofs of Theorems 1.1 and 1.2: namely, $\ell = 0$ or $\rho(m|L,\ell) = \rho(m|H,\ell)$. Those proofs now go

[31] We could imagine a third form of noise whereby a fraction of individuals receive no private signal, and therefore simply free-ride off the public information. These are analogous to the 'noise traders' that richly populate the financial literature. But they require no special treatment here, as they are subsumed in the standard model outlined in section 1.2. For if μ^H and μ^L have a common atom accorded the same probability under each measure, then F^H and F^L will each have an atom at 1/2. Since a noise trader is precisely someone who has the private belief equal to the common prior, namely 1/2, all results from section 1.3 now carry over.

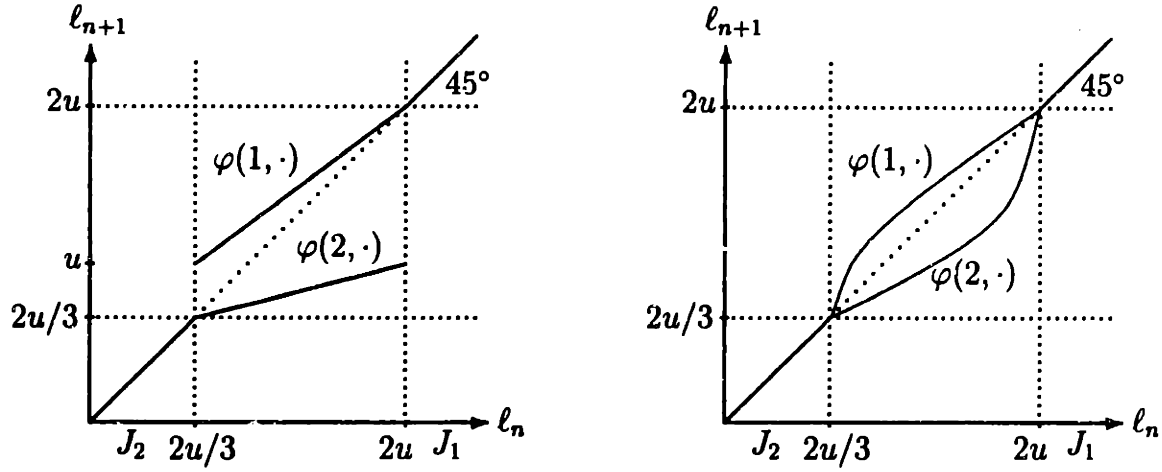

Figure 1-6: **Continuations.** Here, we juxtapose the two continuation likelihood functions for the BOUNDED BELIEFS EXAMPLE — first for only rational individuals, then with some crazy types. We see in the right graph that the discontinuity vanishes, corresponding to the failure of the overturning property.

through.

Appendix 1.E contains the more involved argument for trembling agents. □

1.4.3 Convergence of Actions: The Failure of the Overturning Principle

Whether learning is complete is surprisingly unaffected by a small amount of constant background noise. But analogs of the herding results need modification, for noisy individuals — like cats — do not herd. We shall define a *rational herd* as a herd restricted to non-noisy individuals, but even then, the failure of the overturning principle offers a special challenge: Herd violations only minimally impact public beliefs, as they are deemed irrational acts. This severs our clean implication: belief convergence ⇒ action conformity.

Let us illustrate how Lemma 1.8 fails with noise. The gist of the overturning principle is that unexpected actions can dramatically impact beliefs: For ℓ close to J_m, $|\varphi(m,\ell)-\ell|$ is small, *but* $|\varphi(m',\ell)-\ell|$ is bounded away from 0 for all $m' \neq m$. The left graph in figure 1-6 depicts this property for our running BOUNDED BELIEFS EXAMPLE. Observe how each stationary point of the process is only fixed under one of the continuations. But this is fine, because contrary action are not expected to occur at that point. Thus, the overturning principle implies that for only one action a'_m is $\varphi(m',\ell)-\ell=0$ near a fixed point ℓ.

Introduction of a small amount of noise effects a remarkable sea change in figure 1-6.

For then all actions occur with positive probability $\psi(m'|H,\ell) > 0$, and so Theorem A-1.2 requires that $\varphi(m',\ell) - \ell$ be arbitrarily close to 0 near a stationary point for all m'. If ρ and ψ are also continuous, then the corollary to Theorem A-1.2 says that $\varphi(m',\ell) - \ell = 0$ for all m' at the stationary point. For instance, with craziness, we generally have

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

where $\beta(m',\ell) = \ell\rho(m'|L,\ell)/\rho(m'|H,\ell)$ denotes the old continuation function without noise. So, the addition of noise pushes $|\varphi(m',\ell) - \ell|$ closer to 0 for all m': For any $\ell \in J_m$, we have $\varphi(m,\ell) = \ell$ because the numerator is 0, and $\varphi(m',\ell) = \ell$ because the denominator is infinite ($\rho(m'|H,\ell) = 0$). In the right graph in figure 1-6 we have loosely depicted what the continuations will look like with noise added to our running bounded beliefs example.

That $\varphi(\tilde{m},\check{\ell}) - \check{\ell} = 0$ at a fixed point $\check{\ell}$ allows us to make simple deductions about the rate of convergence for this system. Appendix 1.C develops a theory of stability for stochastic dynamical systems like this one. Given functions φ that are C^1 outside the absorbing basin, and ψ that are continuous,[32] Theorem asserts that $\ell_n \to \check{\ell}$ at the rate

$$\theta = \prod_{m=1}^{M} |\varphi'(m,\check{\ell})|^{\psi(m|H,\check{\ell})}$$

i.e., the geometric average of the continuation derivatives, weighted by their frequencies.

Lemma 1.11 (Rate of Belief Convergence) *Assume F^H and F^L are C^1 with derivatives f^H and f^L. With bounded beliefs, if $f^H(\underline{b}) \neq 0$ and $f^L(\bar{b}) \neq 0$, $\theta < 1$. With unbounded beliefs, if $f^L(0) < \infty$ and $f^H(1) < \infty$, $\theta = 1$, and $\sum_{n=1}^{\infty} \ell_n = \infty$.*

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

$$1 = \sum_{m=1}^{M} \psi(m|H,\ell)\varphi'(m,\ell) + \sum_{m=1}^{M} \psi'(m|\ell)\varphi(m,\ell) \tag{1.15}$$

[32] The appendix actually assumes less, namely that we are dealing with Lipschitz functions. This means that the analysis we present here could be translated to non-C^1 functions.

and $\sum_{m=1}^{M} \psi'(m|\ell) = 0$. At a stationary point $\check{\ell}$, the second sum in (1.15) vanishes. Indeed, $\psi(m|H,\ell) > 0$ for all m with noise, and thus $\varphi(m,\check{\ell}) = \check{\ell}$ for all m too, and the terms $\psi'(m|\ell)$ sum to 0. In other words, the arithmetic average of the derivatives $\varphi'(m,\check{\ell})$ is 1. If any of them is negative, $\langle \ell_n \rangle$ eventually jumps into the basin, and just as with the no noise case, the convergence rate is then 0. If all of the derivatives are non-negative, the arithmetic mean-geometric mean inequality proves that $\theta \leq 1$. So, in general, convergence occurs at a rate $\theta \leq 1$, with equality iff *all* the derivatives are one. If the beliefs are bounded, and the limit of $\langle \ell_n \rangle$ is in J_m, it is not hard to prove that $\varphi'(m,\check{\ell}) = (\kappa_m + \kappa\beta'(m,\check{\ell}))/(\kappa_m + \kappa)$. This is strictly less than one because the derivative of $\beta(m,\ell) = \ell\rho(m|L,\ell)/\rho(m|H,\ell)$ is (see Lemma 1.9 — the assumptions on $f^H(\underline{b})$ and $f^L(\bar{b})$ mean that extreme signals are rare). On the contrary, when the beliefs are unbounded, all of the $\varphi(\tilde{m},\cdot)$ will have a derivative of 1 because β does. □

We now address the question of rational herds. Will one eventually start? This turns on the speed of convergence of the public likelihood. For a herd arises iff an infinite string of individuals has private beliefs too weak to counteract the public belief. Suppose we have a limit cascade toward some $\ell \in J_m$. Then a herd on a_m eventually starts so long as there is not an infinite string of rational 'herd violators'. In light of the (first) Borel-Cantelli Lemma,[33] this occurs with zero chance provided $\sum_{n=1}^{\infty} (1 - \rho(m|H,\ell_n)) < \infty$ almost surely. With bounded beliefs, this inequality holds under the assumptions that F^H and F^L are C^1 with $f^H(\underline{b}) \neq 0$ and $f^L(\bar{b}) \neq 0$, because the convergence of $\langle \ell_n \rangle$ is at an exponential rate. So, Lemma 1.11 implies that ***with bounded beliefs herds must arise, even with noise.*** With unbounded beliefs, we have seen in Lemma 1.11 that the rate of convergence is very slow (when extreme signals are very rare), and we do not know whether the Borel-Cantelli inequality occurs almost surely or even with positive probability if the beliefs are unbounded. Therefore we do not know whether herds will occur with unbounded beliefs.

[33]We actually mean the non-standard conditional version of the Lemma, e.g. Corollary 5.29 of Breiman (1968).

1.5. MULTIPLE INDIVIDUAL TYPES

Having characterized the theory with identical preferences, we now venture into new territory and investigate the richer and more realistic model with heterogeneity. As with the noise formulation, we assume that an individual's type is his private information. Yet everyone will be able to extract information from history by comparing the proportion of individuals choosing each action with the known frequencies of preference types. So long as all individuals know the type frequencies and these are not equal, this inference intuitively ought to be fruitful.[34] Here, an interesting new twist is introduced: The learning dynamics may well converge upon an outcome in which each action is taken with the same probability in all states. We shall argue that this 'twin pathology', which we dub *confounded learning,* arises even with unbounded private beliefs — that is, even when an incorrect herd cannot arise. So, even with unbounded beliefs, the learning need not be complete.

1.5.1 The Model and an Overview

We modify the standard herding model and suppose there are finitely many *types* $t = 1, \ldots, T$. An individual's type determines his vNM preferences over the given set of actions $a_1, \ldots, a_M$, as well as the distribution whence his private signal is drawn. Let λ^t denote the known proportion of individuals of type t, and assume the types are i.i.d. across individuals. Since an individual's type is private information, this formulation reduces to that of noise if all but one type has state-independent preferences. *In general, unlike noise, the decisions of everyone will potentially depend on (and thus be informative of) their private signals; this fundamentally changes the nature of the analysis and results: for instance, complete learning requires no restrictions on the type proportions.*

As before, we can easily prove that $\langle \ell_n \rangle$ is a convergent martingale in state H. Each type will still employ a posterior belief threshold rule when choosing an action; however, the types 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

[34]With nongeneric specifications, such as equal frequencies and exactly opposed vNM preferences, this logic fails.

likelihood ℓ, and call the set of all ℓ yielding $\rho^t(m|H,\ell) = \rho^t(m|L,\ell) = 1$ the action absorbing basin J^t_m. Similarly, define $J^t = J^t_1 \cup J^t_2 \cup \cdots \cup J^t_M$ for all types t, so that $J = J^1 \cap J^2 \cap \cdots \cap J^T$ are the *absorbed outcomes*: all ℓ where each type finds himself in an absorbing basin. If $\ell_n \in J$, for no matter which action is observed, it must have been chosen irrespective of private information, and thus there is nothing to be learned from it. Conversely, if $\ell_n \notin J$, actions must convey information to some type. Notice also that if just one type has unbounded private signals, then $J = \{0\} \cup \{\infty\}$. But unlike Theorems 1.1 and 1.2, we cannot conclude that $\ell_\infty \in J$ almost surely, as there are potentially more limit points.

When a limit cascade arises, we can also show that a *type-specific herd* may arise: By this, we now mean that everyone of the same type will take the same action. Yet, when some types' vNM preferences are not identical, the overturning principle will (sometimes) fail here just as it did with noise: Unexpected actions need not radically affect beliefs, because the successors will also entertain the hypothesis that the individual was simply of a different type. Only if the cascade is on a joint action basin which leads all preference types to take one and the same action will the overturning principle hold to prove that a herd must arise. Otherwise, we can apply the speed of convergence reasoning from section 1.4.3 to conclude that herds must arise (at least with bounded beliefs).

The dynamics of the likelihood ratio is described in the usual notation by (1.13) and

$$\psi(m|s,\ell) = \sum_{t=1}^{T} \lambda^t \rho^t(m|s,\ell) \tag{1.16}$$

If all ψ functions are continuous in ℓ, then the potential limit outcomes must satisfy (A-4): in state H, $\psi(m|H,\ell) = 0$ or $\varphi(m,\ell) = \ell$. Theorems 1.1, 1.2, and 1.6 essentially proved that all solutions to this criterion are found in the absorbing basins. With multiple types, this is no longer true. We shall refer to solutions ℓ^* of criterion (A-4) which are not in J as *confounding outcomes,* and say that *confounded learning* obtains if $\ell_n \to \ell^*$. At a confounding outcome, criterion (A-4) is met because both actions are taken with equal probability in the two states of the world, or

$$\psi(m|L,\ell^*) = \psi(m|H,\ell^*) \quad \forall m \tag{1.17}$$

and so $\varphi(m,\ell^*) = \ell^*$.

Observe that at ℓ^*, it is *not* true that history is totally uninformative, for otherwise individuals would ignore it and their decisions would be informative of the state of the world. Rather it is the case that it is precisely so informative as to render impossible any *additional* inferences from it. The distinction with a cascade is rather sweet. Individuals' private signals in a confounding outcome are totally decisive for their actions, whereas in a limit cascade, it is history that is decisive, and individuals' private signals are (in the limit) wholly inconsequential: Yet both are pathological outcomes to the learning: Observational learning stops short of a focused belief on the true state of the world.

Have we catalogued all possible pathological outcomes to the learning? As noted, the Markovian nature of the stochastic process does not help us much, because the state space is not finite. There might perchance also exist a stable cycle, or even worse. But the dual character of our stochastic process is truly helpful, for the martingale structure of ℓ_n rules out all outcomes in which ℓ_n does not converge.

1.5.2 Examples of Confounded Learning

We shall now show that by example that confounding outcomes may exist, and confounded learning can occur.

In the examples, for simplicity we have only $M = 2$ actions, and $T = 2$ types: $\{U, V\}$. The types only differ in their preferences, not in their private signal distributions. Type U has our usual preferences, preferring action a_2 over a_1 in state H and vice versa in state L. Also, he chooses action a_1, when his private belief is below $\bar{p}^U(\ell) = \ell/(u+\ell)$. Type V has the opposite preferences: action a_1 guarantees zero payoff, while a_2 yields payoff v in state H and -1 in state L. He thus chooses a_1 exactly when his private belief is *above* the threshold $\bar{p}^V(\ell) = \ell/(v+\ell)$. Preference parameters are $v \geq u > 0$, WLOG.

Unbounded Beliefs Example Cont'd. We recall from previous calculations that

$\rho^U(1|H,\ell) = \ell^2/(u+\ell)^2$ and $\rho^U(1|L,\ell) = \ell(\ell+2u)/(u+\ell)^2$. By analogy, $\rho^V(1|H,\ell) = v(2\ell+v)/(v+\ell)^2$ and $\rho^V(1|L,\ell) = v^2/(v+\ell)^2$. Thus, we get from equation (1.16) that

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

The confounding outcomes are those $\ell^* \in (0,\infty)$ solving (1.17), or

$$\frac{\lambda^U}{\lambda^V} = \frac{(u+\ell)^2 v}{(v+\ell)^2 u} \equiv \eta(\ell)$$

One can easily show that $\eta(\ell)$ is a weakly increasing function, spanning $[u/v, v/u]$. Since this range always contains 1, for fixed $v \neq u$ there will be solutions to equation (1.16) provided the proportions of the two types are sufficiently close to 1/2. Moreover, one can show that $\varphi_\ell(1,\ell^*)$ and $\varphi_\ell(2,\ell^*)$ are both positive and not equal, which implies the *local stability* of the point ℓ^* (just as in the proof of Lemma 1.11). As defined in Appendix 1.C, this means that if ℓ_n starts close enough to ℓ^*, then $\ell_n \to \ell^*$ with positive probability.

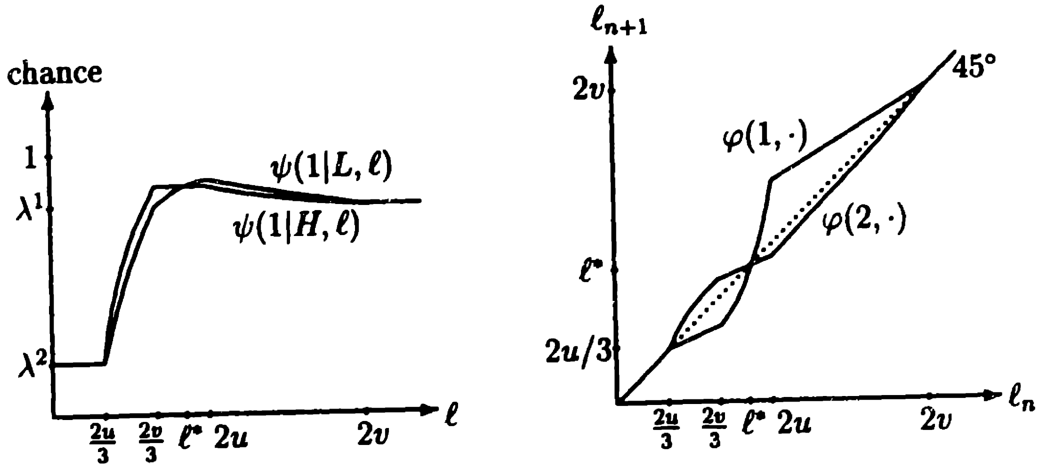

Figure 1-7: **Confounded Learning.** This is based on our Bounded Beliefs Example, with $\lambda^U = 4/5$, $u = v/2$. In the left graph, the curves $\psi(1|H,\ell)$ and $\psi(1|L,\ell)$ cross at the confounding outcome ℓ^*, where no additional decisions are informative. At ℓ^*, 7/8 choose action a_1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

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 1.3 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 $u \leq v$, there are no overlaps when $2v/3 < u$. 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 (1.16). Figure 1-7 graphs these functions. We can rewrite (1.17) for a confounding outcome as

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

If $u > v$ then η maps $(2v/3, 2u)$ onto $(0, \infty)$, and so a confounding outcome exists for any λ^U, λ^V. Next, $\varphi_\ell(1, \ell^*), \varphi_\ell(2, \ell^*) > 0$ and $\varphi_\ell(1, \ell^*) \neq \varphi_\ell(2, \ell^*)$ generically.[35] As in the proof of Lemma 1.11, the confounding outcome is locally stable: 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 the interval $[2u/3, \ell^*]$ or $[\ell^*, 2v]$. As all probability mass is eventually concentrated at the endpoints, only an incorrect herd or confounded learning can arise if $\ell_0 \in [\ell^*, 2v]$, while only a correct herd or confounded learning can arise if $\ell_0 \in [2u/3, \ell^*]$. And just as in the proof of Theorem 1.1, since $\langle \ell_n \rangle$ is a bounded martingale, $E[\ell_\infty] = \ell_0$ implies that in each case, both possible outcomes have positive probability. More generally, with only a single confounding outcome, limit cascades must occur with positive chance.

By the same token, in the UNBOUNDED BELIEFS EXAMPLE, if $\ell_0 < \ell^*$, Fatou's lemma implies $E[\ell_\infty] \leq \ell_0$, and thus complete learning occurs with positive probability.

[35]Namely, $\varphi_\ell(1, \ell^*) = (3\lambda^U/4 + \lambda^V/4)/\psi(1|H, \ell^*)$ and $\varphi_\ell(2, \ell^*) = (\lambda^U/4 + 3\lambda^V/4)/\psi(2|H, \ell^*)$.

1.5.3 The Basic Theory

We have shown by example that confounding outcomes can exist, and confounded learning may arise. We now address robustness and the inter-relationship with limit cascades.

Theorem 1.7 (Confounded Learning) *Assume there are $T \geq 2$ types.*

(1) *Let the private signal distributions be atomless. Then for nondegenerate specifications of preferences and type proportions, confounding outcomes exist, and confounded learning obtains with positive probability. With bounded private beliefs, no single confounding outcome can attract all mass.*

(2) *Generically, at any confounding outcome only two actions are taken.*

(3) *When confounded learning does not occur, a limit cascade arises, i.e. $\ell_n \to \hat{\ell} \in J$, and is almost surely correct with unbounded beliefs.*

(4) *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 examples, which are not nongeneric.

(2) 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 $\hat{M} \leq M$ actions are taken with positive probability at ℓ^*. Equation (1.17) reduces in (à la Walras' Law) to $\hat{M} - 1$ independent equations, since $1 = \sum_{m=1}^{\hat{M}} \psi(m|H, \ell) = \sum_{m=1}^{\hat{M}} \psi(m|L, \ell) = 1$. Generically, $\hat{M} - 1$ equation in one unknown ℓ can only be solved when $\hat{M} = 2$.

(3) As $\langle \ell_n \rangle$ is a convergent martingale, eventually it must settle down to some $\hat{\ell} \in \ell_\infty$ which must satisfy (A-4). This precludes all but limit cascades, where $\psi(m|\hat{\ell}) > 0$ for exactly one action a_m, and a confounding outcome, where $\psi(m|\hat{\ell}) > 0$ for two actions a_m. With unbounded beliefs, $\langle \ell_n \rangle$ cannot explode to infinity (in state H) as before, and thus the only possible limit cascade is $\hat{\ell} = 0$.

(4) Let $M > 2$. If one type has unbounded private beliefs, then for no $\ell \in (0, \infty)$ are only two actions taken with positive probability. By (3), confounding outcomes only

appear in degenerate models. Next, observe that (1.17) implies

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

If F^H and F^L are discrete, then the right hand side 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 probability at a confounding outcome, this does not mean that generic models with $M > 2$ actions do not have confounding outcomes. For with with bounded beliefs, it may occur that only two actions are taken for some range of ℓ. This will happen when there are absorbing basins for the insurance actions, as in figure 1-8.

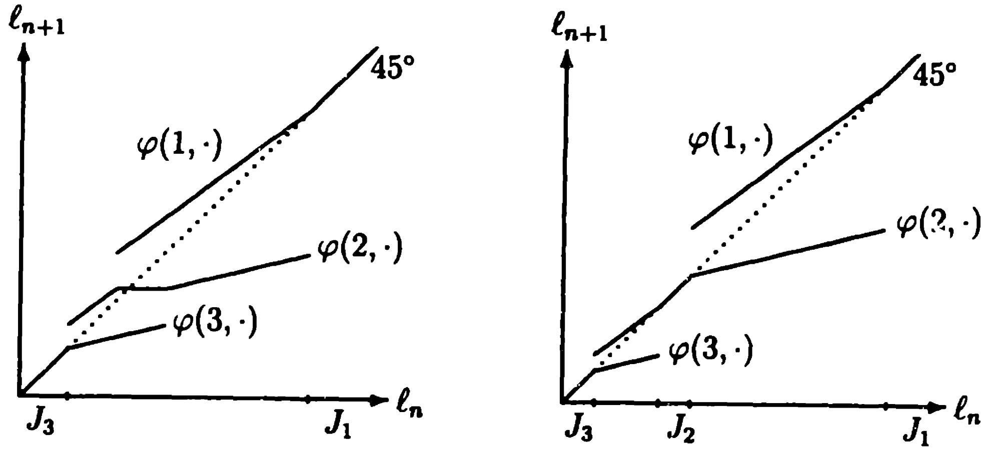

Figure 1-8: **Continuations and Absorbing Basins.** This is based on the BOUNDED BELIEFS EXAMPLE, but here with one insurance and two extreme actions. In the first graph, preferences are such that there is no interior basin, but in the second graph there is an interior basin. Observe in the second graph, 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.

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

3. In a limit cascade, if the likelihood ratio converges to a basin where all types take the same action, then the overturning principle holds as before, and so herds must arise in finite time. For all other limit cascades, type-specific herds still arise in finite time

with bounded beliefs, just as with noise. With unbounded beliefs, this is likewise an open question.

We now return to the example from the introduction, and see if we can say for sure that a confounding outcome exists with unbounded beliefs.

THE DRIVING EXAMPLE REVISITED. 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 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 $(v, 0)$.[36] The proportions are $\lambda^U = .7$ and $\lambda^V = .3$.

CLAIM. *With a differentiable signal distribution, and unbounded beliefs, a confounding point exists if $u/v < \lambda^U/\lambda^V < v/u$.*

Proof: We simply show that $\psi(1|H,\ell)$ lies below $\psi(1|L,\ell)$ near $\ell = 0$ and above it a near $\ell = \infty$, and so the curves must cross. Differentiation of

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

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-1.1, $\psi_\ell(1|H,0) > \psi_\ell(1|L,0)$ exactly when $\lambda^U/\lambda^V > u/v$. Similar methods reveal the reverse inequality near $\ell = \infty$. □

Thus, provided preferences are sufficiently disperse, we have $u/v < 7/3$ and $7/3 < v/u$, and thus a confounding outcome must exist.

1.6. SUMMARY AND AN EXTENSION

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.

Some of our results have turned on the knife-edge between bounded and unbounded

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

beliefs. We wish to point out that Milgrom's ((1979)) convergence theorem for competitive bidding relied on an analogous distinction. Just as Pesendorfer and Swinkels (1995) recently replaced unbounded private beliefs by an assumption of a large number of agents each with bounded beliefs, so too in our setting, it might possible envision a herding model with unbounded beliefs as tantamount to one with bounded beliefs and with each agent replaced by the simultaneous entry of a Poisson-sized cohort.

In a separate work in progress, Smith and Sørensen (1996b), we relax (among other things) the key implicit assumption that individuals can perfectly observe the order of all predecessors' moves. While this is yet another (more plausible) reason for why isolated contrary actions might have very little effect, we are motivated by deeper concerns. For martingales simply do not obtain there, and the resulting analysis is strikingly different than standard rational learning theory! Still, in that setting, the assumption of private beliefs having full-support appears pivotal.

1.A. CONSEQUENCES OF BAYES UPDATING

The setup is taken from section 1.2.1. First, since (*i*) μ^H and μ^L are mutually absolutely continuous, and (*ii*) $p(\sigma)$ is in the support of F^s exactly when σ is in the support of μ^s, $s = H, L$, both F^H and F^L have the same support. We let $\text{co}(\text{supp}(F)) \equiv [\underline{b}, \bar{b}] \subseteq [0, 1]$, and the fair priors assumption implies that $0 \leq \underline{b} < 1/2 < \bar{b} \leq 1$.

The characterization of the Radon-Nikodym derivative of F^H and F^L in Lemma 1.1 implies

Lemma A-1.1 (Distributional Properties)

1. *The difference $F^L(p) - F^H(p)$ is weakly increasing for $p \leq 1/2$ and weakly decreasing for $p \geq 1/2$. Both properties are strict for all $p \in \text{supp}(F) \setminus \{1/2\}$.*
2. *$F^L(p) > F^H(p)$ holds except when $F^L(p) = F^H(p) = 0$ or $F^L(p) = F^H(p) = 1$.*
3. *$F^H(p) \leq pF^L(p)/(1-p)$ holds for all $p \in (0,1)$, and is strict when $F^L(p) > 0$.*
4. *The ratio F^H/F^L is weakly increasing, and strictly so on $(\underline{b}, \bar{b}]$.*

Proof: Since by Lemma 1.1, F^H grows faster than F^L exactly when $p \in \text{supp}(F)$ satisfies $p > 1/2$, the first two points follow.

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

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

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-1). It immediately follows that $F^H(p)/F^L(p) > F^H(q)/F^L(q)$. □

The final result requires the assumptions of rare and very rare extreme signals, as defined in section 1.3.

Lemma A-1.2 *Assume that extreme signals are not rare at $\bar{b}$. Then there exist $\varepsilon > 0$ such that for all $p \in (\bar{b} - \varepsilon, \bar{b})$, $F^L(p)/F^H(p) - 1 > \varepsilon(\bar{b} - p)$. Similarly, if extreme signals are not rare at $\underline{b}$, then there exist $\varepsilon > 0$ such that for all $p \in (\underline{b}, \underline{b} + \varepsilon)$, $1 - (1 - F^L(p))/(1 - F^H(p)) > \varepsilon(p - \underline{b})$. If extreme signals are very rare at 1, there there exist $K > 0$ so $F^H(p)/F^L(p) > 1 - K(1-p)$ for p near 1. Similarly, if extreme signals are very rare at 0, then there exist $K > 0$ so $(1 - F^L(p))/(1 - F^H(p)) < 1 - Kp$ for p near 0.*

Proof: Notice that for $p < \bar{b}$, $F^L(p)/F^H(p) - 1 > \varepsilon(\bar{b} - p)$ is equivalent to $[F^L(p) - F^H(p)]/[\bar{b} - p] > \varepsilon F^H(p)$. If the ε were not to exist as claimed, then $\liminf_{p \nearrow \bar{b}} [F^L(p) - F^H(p)]/[\bar{b} - p] = 0$. Since $F^H(p) \to 1 = F^H(\bar{b})$ as $p \nearrow \bar{b}$, extreme signals would be rare at $\bar{b}$. The other half with rare beliefs follows from analogous arguments.

Next, assume extreme signals are very rare at 0, and let L be such that $F^L(p)/p < L$. With unbounded beliefs, $F^H(p)/p$ stays near 0 as p is near zero. Consequently, there exist a $K > 0$ such that for p near 0, $(F^L(p) - F^H(p))/(p + pF^H(p)) < K$. Simple algebra proves $(1 - F^L(p))/(1 - F^H(p)) < 1 - Kp$. When extreme signals are very rare at 1, the proof is similar.

□

1.B. DISCRETE DYNAMICAL 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.[37]

The context is as follows. Given is a finite set $\mathcal{M}$, and measurable functions $\varphi(\cdot, \cdot) : \mathcal{M} \times \mathbf{R}_+ \to \mathbf{R}_+$, and $\psi(\cdot|\cdot) : \mathcal{M} \times \mathbf{R}_+ \to [0,1]$ meeting two restrictions. First, $\psi(\cdot|\ell)$ must be a probability measure for all $\ell \in \mathbf{R}_+$, or

$$\sum_{m \in \mathcal{M}} \psi(m|\ell) = 1.$$

Second, the following 'martingale property' must hold for all $\ell \in \mathbf{R}_+$:

$$\sum_{m \in \mathcal{M}} \psi(m|\ell)\, \varphi(m, \ell) = \ell \qquad \text{(A-2)}$$

Finally, equip $\mathbf{R}_+ = [0, \infty)$ with the Borel σ-algebra $\mathcal{B}$, and define a transition probability $P : \mathbf{R}_+ \times \mathcal{B} \to [0,1]$ as follows:

$$P(\ell, B) = \sum_{m \,|\, \varphi(m,\ell) \in B} \psi(m|\ell) \qquad \text{(A-3)}$$

for any $B \in \mathcal{B}$. For our application, $\psi(m|\ell)$ is the chance that the next agent takes action a_m

[37] Arthur, Ermoliev, and Kaniovski (1986) consider a stochastic system with a seemingly similar structure — namely, a 'generalized urn scheme'. Their approach, however, 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. We do make one application of their method in Smith and Sørensen (1996b).

when faced with likelihood ℓ, and $\varphi(m,\ell)$ is the resulting continuation likelihood ratio.

Suppose for definiteness that we are given a (measurable) Markov stochastic process $\langle \ell_n \rangle_{n=1}^{\infty}$ on $(\Omega^H, \mathcal{E}^H, \nu^H)$, where for each n, $\ell_n : \Omega^H \to \mathbb{R}_+$.[38] Transition from ℓ_n to ℓ_{n+1} is described by the transition probability P. We assume that $E\ell_1 < \infty$; in applications we shall always assume that ℓ_1 is identically 1, so this is not restrictive.[39] Denote by $\mathcal{F}_n$ the σ-field in $(\Omega^H, \mathcal{E}^H)$ generated by $(\ell_1, \ldots, \ell_n)$. Clearly, ℓ_n is $\mathcal{F}_n$-measurable, and it follows from (A-2) that $\langle \ell_n, \mathcal{F}_n \rangle$ is actually a martingale,[40] thus justifying our earlier casual description of property (A-2). Indeed,

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

Since $\langle \ell_n \rangle$ is a martingale on $\mathbb{R}_+$, the Martingale Convergence Theorem asserts that it converges almost surely in $\mathbb{R}_+$, say to $\ell_\infty = \lim_{n\to\infty} \ell_n$. We now characterize the limit.

Theorem A-1.1 (Stationarity) *Assume that for all $m \in \mathcal{M}$, the two functions $\ell \mapsto \varphi(m,\ell)$ and $\ell \mapsto \psi(m|\ell)$ are continuous. Suppose that $\ell_n \to \ell_\infty$ almost surely. Then for all $m \in \mathcal{M}$ and for all $\ell \in \operatorname{supp}(\ell_\infty)$, stationarity obtains, i.e.*

$$\psi(m|\ell) = 0 \quad or \quad \varphi(m,\ell) = \ell \tag{A-4}$$

The proof is deferred. That implication (A-4) is truly a stationarity condition is best seen — by means of (A-3) — in its alternative formulation $P(\ell, \{\ell\}) = 1$.

The intuition behind Theorem A-1.1 is rather simple. Since ℓ_n converges almost surely to ℓ_∞, it also converges weakly (in distribution) to ℓ_∞. As the process (ℓ_n, m_n) is also a Markov chain, it is intuitive that the limiting distribution is invariant for the transition P, as described in Futia (1982). The a.s. convergence in turn implies that the invariant limit must be pointwise invariant. In fact, we can prove Theorem A-1.1 along these lines, but the continuity assumptions are subtly hard-wired into the final stage of the argument proving that the limiting distribution is invariant. As we wish to do away with continuity, we establish an even stronger result. Motivated by the

[38] $\mathcal{E}^H$ and $\mathcal{E}^L$ correspond to the the restricted sigma fields of $\mathcal{E}$ on Ω^H and Ω^L, respectively.

[39] Notice that the system has a discrete transition function; therefore, if ℓ_1 has a discrete distribution the process will be a discrete (in fact, countably infinite) Markov chain. One might think that it would be possible to apply standard results about the convergence of discrete Markov chains, but in fact such results are not useful here. While the state space is certainly countable, all states (which will soon be interpreted as likelihood functions) are in general transitory, and so standard results are useless.

[40] No ambiguity arises if we simply say that $\langle \ell_n \rangle$ is a martingale.

fact that (A-4) is violated for m exactly when neither $\psi(m|\ell)$ nor $\varphi(m,\ell)-\ell$ is zero, we have

Theorem A-1.2 (Generalized Stationarity) *Assume that the open interval $I \subseteq \mathbb{R}_+$ has the property*

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

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

Proof: Let I be an arbitrary open interval satisfying $(\star)$ for $\varepsilon > 0$, and suppose by way of 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 $m \in \mathcal{M}$ such that $\psi(m|\ell) > \varepsilon$, and $\varphi(m,\ell) \notin J$. Because $\bar{\ell} \in \operatorname{supp}(\ell_\infty)$, there is positive probability that $\ell_n \in J$ eventually. But whenever $\ell_n \in J$, there is a probability of at least ε that $\ell_{n+1} \notin J$. That is, the conditional probability that the process stays in J in the next period is at most $1-\varepsilon$. Recursively, the chance that $\langle \ell_n \rangle$ forever remains in J after stage n is at most $\prod_{k=n}^{\infty}(1-\varepsilon) = 0$. This contradicts the claim that with positive probability the process $\langle \ell_n \rangle$ 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 A-1.2 yields an immediate contradiction. □

Finally, it is obvious that the corollary implies Theorem A-1.1.

1.C. STABILITY OF A STOCHASTIC DIFFERENCE EQUATION

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.[41] Further, we add an analysis of rates of convergence.

[41] We are coining terms here. We call a fixed point $\bar{y}$ of a stochastic difference equation *locally stable* if $\Pr(\lim_{n\to\infty} y_n = \bar{y}) > 0$ whenever $y_0 \in \mathcal{N}_{\bar{y}}$, a small enough neighborhood about $\bar{y}$. If $\Pr(\lim_{n\to\infty} y_n = \bar{y}) > 0$ for all y_0, then $\bar{y}$ is *globally stable*.

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 A-1.3, we do obtain the exact rate, but, as we discuss later on, we do not determine the exact rate in higher dimensions. The cited literature is thus forced to make stronger assumptions than we do. Bellman (1954) is cited as the first work in this area, Furstenberg (1963) is a classic article, and Kifer (1986) is a modern textbook. Results related to ours have recently been derived by Ellison and Fudenberg (1995), yet their method is different from ours and we believe they could not reach the same generality (several dimensions).

Consider linear stochastic difference equations of the following form. An i.i.d. stochastic process $\langle y_n \rangle$ on $\mathbb{R}$ is given, such that $\Pr(y_n = 1) = p = 1 - \Pr(y_n = 0)$. Define the auxilliary stochastic process $\langle y_n \rangle$ on $\mathbb{R}$ as follows: ℓ_0 is given, and

$$\ell_n = \begin{cases} a\ell_{n-1} & \text{if } y_n = 1 \\ b\ell_{n-1} & \text{if } y_n = 0 \end{cases} \tag{A-5}$$

where a and b are fixed real constants. Define $\bar{\theta} = |a|^p|b|^{1-p}$. The first lemma is a straightforward generalization of the standard stability criterion for linear difference equations:

Lemma A-1.3 (Global Stability) *Almost surely, $\theta^{-n}\ell_n \to 0$ for all $\theta > \bar{\theta}$, and $\theta^{-n}|\ell_n| \to \infty$ for all $\theta \in (0, \bar{\theta})$.*

In particular, $\ell_n \to 0$ almost surely if $\bar{\theta} < 1$, and $|\ell_n| \to \infty$ almost surely if $\bar{\theta} > 1$.

Proof: Define the stochastic variable $Y_n \equiv \sum_{k=1}^{n} y_k$ so that $\ell_n = a^{Y_n}b^{n-Y_n}\ell_0$, and thus $|\ell_n| = \left(|a|^{\frac{Y_n}{n}}|b|^{\frac{n-Y_n}{n}}\right)^n |\ell_0|$. As the Strong Law of Large Numbers yields $Y_n/n \to p$ almost surely, the result follows immediately from $|a|^{\frac{Y_n}{n}}|b|^{\frac{n-Y_n}{n}} \to \bar{\theta}$ almost surely. □

REMARK. 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.

It is straightforward to generalize Lemma A-1.3 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.[42] Then the following half of Lemma A-1.3 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 A-1.3 is the only result applied in the sequel, our local stability assertions will also go through in multidimensional models.

It may be noted that our definition of the convergence rate is not very tight. By our definition, if $\langle \ell_n \rangle$ converges to ℓ^* at the rate θ', then it also converges 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 A-1.3. 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 the tightest such. See Kifer (1986) for more precise results for linear systems.

Next we provide a different stability result, helpful for the later characterization of non-linear systems. We consider the one-dimensional system (A-5) once again.

Lemma A-1.4 *If $\bar{\theta} < 1$ and $\mathcal{N}_0$ is any open ball around 0, then there is a positive probability that $y_n \in \mathcal{N}_0$ for all n, provided $y_0 \in \mathcal{N}_0$.*

Proof: First, if $ab\ell_0 = 0$, then there is a positive probability of jumping to 0 immediately, and the system will stay there. So, assume that $ab\ell_0 \neq 0$. Lemma A-1.3 asserts that $\ell_n \to 0$ almost surely. So for any open ball $\mathcal{N}_0$ around 0, we have $\Pr\left(\bigcup_{M\in\mathbf{N}} \bigcap_{n\geq M}\{\omega \in \Omega^H : \ell_n \in \mathcal{N}_0\}\right) = 1$. There must then be some $M \in \mathbf{N}$ such that $\Pr\left(\{\omega \in \Omega^H : \forall n \geq M, \ell_n \in \mathcal{N}_0\}\right) > 0$. Also, for given ℓ_0, supp(ℓ_M) is finite, and $\ell_0 \neq 0$ implies $\ell_M \neq 0$; therefore, there exists $\bar{\ell}_M \in \mathcal{N}_0 \setminus \{0\}$ such that $\ell_n \in \mathcal{N}_0$ for all $n \geq M$ if $\ell_M = \bar{\ell}_M$. Finally, it is simple to see that the dynamical

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

system is time and scale invariant, so we may conclude that given any $\ell_0 \in \mathcal{N}_0$ there is a positive probability of staying there. □

We next use these results to investigate the local stability of non-linear stochastic dynamical systems. The system is now described by

$$\ell_n = \begin{cases} \varphi(1, \ell_{n-1}) & \text{if } y_n = 1 \\ \varphi(2, \ell_{n-1}) & \text{if } y_n = 0 \end{cases} \tag{A-6}$$

where $\varphi(1, \cdot)$ and $\varphi(2, \cdot)$ are given functions, and ℓ_0 is given. Moreover, we no longer need the i.i.d. assumption on the stochastic process $\langle y_n \rangle$. Rather, we only require that there exist a sequence of i.i.d. stochastic variables $\langle \sigma_n \rangle$, each uniformly distributed on $[0, 1]$, such that $y_n = 1$ exactly when $\sigma_n \le \psi(1|\ell_{n-1})$ and $y_n = 0$ otherwise, so $\Pr(y_n = 1) = \psi(1|\ell_{n-1}) = 1 - \Pr(y_n = 0)$. Thus the distribution of y_n is allowed to depend on ℓ_{n-1} in a simple way. We are concerned with stationary points $\hat{\ell}$ of (A-6), namely the solutions to

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

We shall need a little bit of terminology: we say that the function $f : \mathbf{R}^m \to \mathbf{R}^n$ is *Lipschitz* with Lipschitz constant $L \ge 0$ at the point $\hat{x}$ if there exists a neighborhood $\mathcal{N}(\hat{x})$ such that $\forall x \in \mathcal{N}(\hat{x}) : \; ||f(x) - f(\hat{x})|| \le L||x - \hat{x}||$. Note immediately, that if f is continuously differentiable at $\hat{x}$, then f is Lipschitz with any constant $L > ||Df(\hat{x})||$.

Proposition A-1.1 (Local Stability) *Fix a stationary point $\hat{\ell}$ given by (A-7). Assume that at $\hat{\ell}$, $\psi(1|\cdot)$ is continuous, and $\varphi(1, \cdot)$ and $\varphi(2, \cdot)$ are both Lipschitz, with respective Lipschitz constants L_1 and L_2. Suppose also that $\hat{\ell}$ satisfies the first order stability criterion*

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

Then there is an open ball around $\hat{\ell}$, such that starting from any point ℓ_0 in this ball there is a positive probability that ℓ_n always stays in the ball while $\ell_n \to \hat{\ell}$. Further, whenever $\ell_n \to \hat{\ell}$ it converges at the rate $\bar{\theta}$.

Proof: We proceed as follows. First, we majorize the nonlinear dynamical system around $\hat{\ell}$ by a linear stochastic difference equation of the form just treated (that satisfies the conclusions of Lemma A-1.4). Next we argue that the conclusion of Lemma A-1.4 must apply to the original non-linear dynamical system.

By the continuity assumption on $\psi(1,\cdot)$, the inequality (A-8) obtains in a neighborhood of $\hat{\ell}$. Also, we may assume WLOG that $L_1 \leq L_2$, so in particular $L_1 < 1$. Pick a small enough open ball $\mathcal{N}(\hat{\ell})$ around $\hat{\ell}$, and pick $p \in [0,1]$, such that for all $\ell \in \mathcal{N}(\hat{\ell})$:

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

Fix $\ell_0 \in \mathcal{N}(\hat{\ell})$. Introduce a new stochastic process $\langle \tilde{y}_n \rangle$ defined by $\tilde{y}_n = 1$ when $\sigma_n \leq p$, and $\tilde{y}_n = 0$ otherwise. Use this to define a new stochastic process $\langle \tilde{\ell}_n \rangle$ by $\tilde{\ell}_0 = \ell_0$, and

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

Observe that $\langle \tilde{y}_n \rangle$ are independent because $\langle \sigma_n \rangle$ are. Thus Lemma A-1.3 is valid, and tells us that $\tilde{\ell}_n \to \hat{\ell}$ a.s., while Lemma A-1.4 asserts that there is a positive probability that $\tilde{\ell}_n \in \mathcal{N}(\hat{\ell})$ for all n, if only $\ell_0 \in \mathcal{N}(\hat{\ell})$. We consider such a realization of $\langle \sigma_n \rangle$ whereby $\tilde{\ell}_n \in \mathcal{N}(\hat{\ell})$ for all n and $\tilde{\ell}_n \to \hat{\ell}$. We now prove that in such a realization the linear process $\langle \tilde{\ell}_n \rangle$ majorizes the non-linear system $\langle \ell_n \rangle$. Because $\psi(1|\ell) > p$ when $\ell \in \mathcal{N}(\hat{\ell})$, we have $\tilde{y}_n = 1 \Rightarrow y_n = 1$. Then the Lipschitz property yields the desired majorization for all n (in this realization), namely $||\tilde{\ell}_n - \hat{\ell}|| \geq ||\ell_n - \hat{\ell}||$. This is proved by induction in n. Indeed, when $\sigma_n \leq \psi(1|\ell_{n-1})$ we get the majorization

$$||\ell_n - \hat{\ell}|| = ||\varphi(1,\ell_{n-1}) - \varphi(1,\hat{\ell})|| \leq L_1 ||\ell_{n-1} - \hat{\ell}|| \leq ||\tilde{\ell}_n - \hat{\ell}||$$

A similar calculation holds when $\sigma_n > \psi(1|\ell_{n-1})$. So, for any such realization of $\langle \sigma_n \rangle$, $\ell_n \to \hat{\ell}$. We thus conclude that $\ell_n \to \hat{\ell}$ with positive probability.

Finally, consider the rate of convergence. For any given $\theta > \bar{\theta}$ we could have picked the neighborhood $\mathcal{N}(\hat{\ell})$ so small that $L_1^p L_2^{1-p} < \theta$. Whenever ℓ_n converges to $\hat{\ell}$ the process must eventually be in $\mathcal{N}(\hat{\ell})$. In here, we have proved that the process is dominated by the linear system, and Lemma A-1.3 proves that the linear system converges at rate $L_1^p L_2^{1-p}$. □

Note that choosing the smallest possible Lipschitz constants also yields the minimum $\bar{\theta}$. It is straightforward to see that Proposition A-1.1 holds when the process has more than two continuations in each period. It is similarly easy to see how it goes through in several dimensions, i.e. when ℓ is a vector. A peasant's version of Theorem A-1.1 is as follows.

Corollary (C^1-Local Stability) *Fix a stationary point $\hat{\ell}$ of* (A-7). *Assume that at $\hat{\ell}$, $\varphi(1,\cdot)$ and $\varphi(2,\cdot)$ are both continuously differentiable, while $\psi(1|\cdot)$ is continuous. Suppose also that $\hat{\ell}$ satisfies the first order stability criterion*

$$\bar{\theta} = |\varphi_\ell(1,\hat{\ell})|^{\psi(1|\hat{\ell})}|\varphi_\ell(2,\hat{\ell})|^{(1-\psi(1|\hat{\ell}))} < 1 \tag{A-9}$$

Then starting from any point ℓ_0 in some open ball around $\hat{\ell}$, there is a positive probability that ℓ_n always stays in the ball while $\ell_n \to \hat{\ell}$. Further, whenever $\ell_n \to \hat{\ell}$ it converges at the rate $\bar{\theta}$.

We should note that the convergence criterion (A-9) is exact in the one-dimensional case. Thus it follows from a simple application of Lemma A-1.3 that if $\bar{\theta} > 1$ then there is probability zero that the system converges to $\hat{\ell}$.

Finally, let us explain precisely how Corollary carries through in several dimensions, i.e. when ℓ is a vector. In that case $\varphi(m,\ell)$ is a vector function, and of interest is its matrix derivative at the stationary point, $D_\ell\varphi(m,\hat{\ell})$. The geometric average to be 1 is the average of operator norms, $||D_\ell\varphi(m,\hat{\ell})||$ (which is the same as the largest eigenvalue). Then the proof goes through, largely as before. It is only necessary to be a little 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.

1.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 $\mathbf{R}^{S-1}$ sliced into closed convex polytopes. We leave it to the reader to ponder the optimal notation.[43] In Theorem A-1.2 (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

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

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 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 1.2, we get a countable partition, and thus a countable set of posterior belief thresholds $\bar{r}$.[44] In this way, Lemma 1.3 will yield the threshold functions $\bar{p}$ just as before. The martingale properties of the model are preserved.

The convergence results Theorems 1.1 and 1.2 do not depend on the action space being denumerable. In the proof of Theorem 1.1, 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 Theorem 1.2, 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 Theorems 1.3 and 1.4. 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 Theorem 1.4, 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.

---

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

1.E. OMITTED PROOFS

Proof of Lemma 1.4. Given the action history $\{m_1, \ldots, m_{n-1}\}$, the conditional expectation of the next public belief is (by the Markovian assumption),

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

Proof of Lemma 1.7. Since $\bar{p}_m(\ell)$ is increasing in m by Lemma 1.3, $[\bar{p}_{m-1}(\ell), \bar{p}_m(\ell)]$ 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 disjointness is obvious. Next, if $\operatorname{int}(J_m) \neq \emptyset$ 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, it is clear that we can always ensure one of the inequalities in (A-10) for some ℓ, but simultaneously attaining the two may well be impossible. As Lemma 1.3 yields $\bar{p}_0(\ell) \equiv 0$ and $\bar{p}_M(\ell) \equiv 1$ for all ℓ, it follows that we must have $J_M = [0, \underline{\ell}]$ and $J_1 = [\bar{\ell}, \infty]$, where $0 < \underline{\ell} < \bar{\ell} < \infty$ satisfy $\bar{p}_{M-1}(\underline{\ell}) = \underline{b}$ and $\bar{p}_1(\bar{\ell}) = \bar{b}$.

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.

If the beliefs are unbounded, one has $\underline{b} = 0$ and $\bar{b} = 1$. By (A-10), for ℓ to be in J_m we must have $\bar{p}_{m-1} = 0$ and $\bar{p}_m = 1$. From the definition of the thresholds in Lemma 1.3 this can only happen for $m = 1$ with $\ell = \infty$, or for $m = M$ with $\ell = 0$.

Proof of Theorem 1.1. Suppose by way of contradiction that there exist a point $\hat{\ell} \in \operatorname{supp}(\ell_\infty)$ with $\hat{\ell} \notin J_1 \cup \cdots \cup J_M$. 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 (1.4), $\psi(m|\ell) = \rho(m|H,\ell)$ is bounded away from 0 in this neighborhood, while (1.5) 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 A-1.1 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 A-1.2, $\hat{\ell} \in \text{supp}(\ell_\infty)$ 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 bounded away from 0 on an interval $(\hat{\ell} - \eta', \hat{\ell}]$, for some $\eta' > 0$. So, once again Theorem A-1.2 (observe the order of the quantifiers!) proves that $\hat{\ell} \notin \text{supp}(\ell_\infty)$.

Proof of Theorem 1.6, Case 2: Trembling Individuals. All actions are taken with positive probability, and so $\psi(m|\ell)$ is indeed bounded away from 0 by (1.14). 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 1.1 and 1.2. 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 (1.13) to rewrite $\varphi(m,\ell) = \ell$ as follows:

$$[1 - \kappa_m(\ell)]\,[\rho(m|L,\ell) - \rho(m|H,\ell)] = \sum_{k\neq m} \kappa_k^m(\ell)\,[\rho(k|H,\ell) - \rho(k|L,\ell)] \tag{A-11}$$

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-11) 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 1.1 obtains once again, while $\varphi(1,\ell) - \ell$ is bounded away from 0 on the interval I of Theorem A-1.2, and so Theorem 1.2 goes through just as before also.

Chapter 2

Mitigating the Informational Herding Externality

2.1. INTRODUCTION

The last few years has seen a flood of research on a paradigm known as informational herding. We ourselves have actively participated in this research herd (Smith and Sørensen (1996a), or simply SS) that was sparked independently by Banerjee (1992) and BHW: Bikhchandani, Hirshleifer, and Welch (1992). The context is seductively simple: An infinite sequence of individuals must decide on an action choice from a finite menu. Everyone has identical preferences and menus, and each may condition his decision both on his (endowed) private signal about the state of the world and on all his predecessors' decisions; however, observation of their private signals is verboten.

If these private signals have bounded power, it is known that a 'herd' eventually arises, but is not always correct — namely, after some point, all make the identical choice, possibly unwise. This simple pathological outcome has understandably attracted much fanfare.

In SS we explored, embellished, and fully fleshed out this story, and found that herding is quite robust. BHW's kindred notion of *cascades* is not as resilient: While beliefs converge upon a limit where only one action is taken with probability one (a *limit cascade*), this need not occur in finite time. The potential for bad herds is also not without caveat: Absent a uniform bound on the strength of the individuals' private signals, only correct herds arise.

Finally, in a world with multiple preference types, a *confounded learning* outcome might arise, where the lesson of history is forever mixed, and private signals always decisive.

Addressing the Herding Externality. Robustness aside, the herding outcome is a striking example of an informational externality: While all individuals collectively know enough to fully determine the state of the world, it is aggregated rather poorly. For in a herd, most individuals almost surely take an action which reveals almost none of their information. Late-comers ideally prefer that their predecessors had better signalled their information with more revealing actions, but early individuals clearly have no incentive to do so.

The main purpose of our paper is to investigate the herding externality with more forward-looking behavior. We set up and solve the *social planner's problem* when the individual's welfare is weighted together. Equivalently, this is an observational learning model where individuals care about successors. The social planner will forego some current payoff in order to secure knowledge relevant for future payoffs and decisions, just as altruistic individuals are willing to sacrifice some personal gain for posterity. Such an outcome can be decentralized as a constrained social optimum using monetary transfers between agents. We find that even when individuals or a planner internalizes the herding externality in this fashion, incorrect herds and incomplete learning still obtain, but with chances vanishing as the discount factor converges to one.

A Possible Link, and a Puzzle. Is informational herding fundamentally a *new* phenomenon? We have been piqued by its similarity to the familiar failure of complete learning in an optimal experimentation problem. One classic example is Rothschild's (1974) analysis of the *two-armed bandit:* An infinite-lived impatient monopolist optimally experiments with two possible prices each period, with purchase chances for each price being fixed, unknown draws from $[0, 1]$. Rothschild showed that the monopolist would (i) eventually settle down on one of the prices, and (ii) with positive probability select the less profitable price.

Aiming for more than a casual analogy between this outcome and herding, we must recast the observational learning paradigm as a single person optimization problem. This suggests considering the *forgetful experimenter,* who each period receives a new informative signal, takes an optimal action, and then promptly forgets his signal; the next period, he can reflect only on his action choice. Alas, this is neither a very satisfying model of rationality, nor can it be the basis for an optimal experimentation problem. How can an experimenter not observe the private signals, and yet take informative actions?

The Second Puzzle, and a Resolution. An interesting sequel to Rothschild's work was McLennan (1984), who permitted the monopolist the flexibility to charge one of a continuum of prices, but assumed only two possible linear purchase chance 'demand curves'. When the demand curves crossed, he found that the monopolist may settle down on the suboptimal uninformative price.

Rothschild's and McLennan's models give examples of *potentially confounding actions,* as introduced in EK: Easley and Kiefer (1988). In brief, these are actions that are optimal for *unfocused* beliefs for which they are invariants (i.e. taking the action leaves the beliefs unchanged). Of particular significance is the proof in EK (on page 1059) that with finite state and action spaces, there will *generically* not exist any potentially confounding actions, and thus complete learning must arise.[1] Rothschild and McLennan might be seen as separate anticipations of EK's general insight. Rothschild escapes it by means of a continuous state space, whereas McLennan resorts to a continuous action space. Yet there appears no escape for the herding paradigm, where both flavors of incomplete learning, limit cascades and confounded learning, generically arise with two actions and two states.

Our resolution of these puzzles respects the quintessence of the herding paradigm that predecessors' signals are hidden from view. In a nutshell, we imagine that individuals do not act for themselves, but rather furnish optimal history-contingent 'decision rules' for agent machines that automatically map any realized private signal into an action choice. As such, this experimentation problem is the exact same problem that was solved by the social planner.

Overview. Section 2.2 outlines the basic herding model. Section 2.3 characterizes the planner's problem. In section 2.4, we re-interpret the herding paradigm as a case of optimal experimentation.

---

[1] For instance, payoff assignments in a one-armed bandit — where the safe arm is a potentially confounding action — are not generic in $\mathbb{R}^2$.

2.2. OBSERVATIONAL LEARNING MODEL

In this section we set up a very general observational learning model, that generalizes SS, and thus BHW and Banerjee (1992), and also happens to cover Lee (1993). This generality will facilitate the ensuing mapping into the experimentation literature.

Information. An infinite sequence of individuals $n = 1, 2, \ldots$ takes actions in that exogenous order. There is uncertainty about the payoffs from these actions. The elements of the parameter space $(\Theta, \mathcal{F})$ are referred to as *states of the world.* There is a given common prior belief, the probability measure λ over Θ.

Individual n receives a private random signal, $\sigma_n \in \Sigma$, about the state of the world. As demonstrated in SS (Lemma 1), we may assume WLOG that the private signal received by an individual is actually his *private belief,* i.e. we let σ be the measure over Θ which results from Bayesian updating given the prior λ and observation of the private signal. So, we let Σ be the space of all probability measures over $(\Theta, \mathcal{F})$, and $\mathcal{G}$ the associated sigma-algebra. Conditional on the state, the signals are assumed to be i.i.d. across individuals. It is common knowledge that they are distributed according to the probability measure μ^θ in state $\theta \in \Theta$. To ensure that no signal will *perfectly* reveal the state of the world, we shall insist that all μ^θ be mutually absolutely continuous (a.c.), for $\theta \in \Theta$.[2]

Bayesian Decision-Making. Everyone chooses from a fixed action set A, equipped with the sigma-algebra $\mathcal{A}$. Action a earns a nonstochastic payoff $u(a, \theta)$ in state $\theta \in \Theta$, the same for all individuals, where $u : A \times \Theta \mapsto \mathbb{R}$ is measurable. It is common knowledge that everyone is rational, i.e. seeks to maximize their expected payoff. Before deciding upon an action, everyone first observes his private signal and the entire action history. Individuals use the observed action history h to gain information about the state of the world.

An individual's optimal decision rule uses the observed action history and his own private belief. As in SS, we simply assume that an individual can compute the decision rules of all predecessors, and can use the common prior to calculate the ex ante (time-0) probability distribution over action profiles h in either state. Knowing these probabilities,

[2]See Rudin (1987). Measure μ^1 is a.c. w.r.t. μ^2 if $\mu^2(S) = 0 \Rightarrow \mu^1(S) = 0 \; \forall S \in \mathcal{S}$, the sigma-algebra on Σ. By the Radon-Nikodym Theorem, a unique $g \in L^1(\mu^2)$ exists with $\mu^1(S) = \int_S g \, d\mu^2$ for every $S \in \mathcal{S}$. With μ^H, μ^L mutually a.c., 'almost sure' assertions are well-defined without specifying the measure.

Bayes' rule implies a unique *public belief* $\pi = \pi(h) \in \Sigma$ for any history h. A final application of Bayes' rule also given the private signal σ yields the *posterior belief* $\rho \in \Sigma$.

Given the posterior belief $\rho \in \Sigma$, individual n picks the action $a \in A$ which maximizes the expected payoff $\bar{u}_a(\rho) = \int_\Theta u(a,\theta)d\rho(\theta)$. We assume that an optimal action $a = a(\rho)$ always exists.[3] The solution defines an *optimal decision rule* $x : \Sigma \to \mathcal{P}(A)$ ($\mathcal{P}(A)$ denotes the set of probability measures over $(A, \mathcal{A})$). That is, x is an element of the space X of maps from Σ to $\mathcal{P}(A)$. A rule x produces an implied distribution over actions $\nu = x(\sigma)$ simultaneously for all private beliefs σ. Which x is optimal depends on π.

Observational Learning as a Stochastic Process. Given the optimal decision rule x, the probability distribution of signals σ — and thus actions a — neatly depends on the state θ. In fact, the density is $\psi(a|\theta, x) \equiv \int x(\sigma)(a)d\mu^\theta(\sigma)$, and unconditional on the state, it is $\psi(a|\pi, x) \equiv \int_\Theta \psi(a|\theta, x)\pi(d\theta)$. This in turn implies a distribution over next period public beliefs π_{n+1}. Thus, the public beliefs $\langle \pi_n \rangle$ follow a discrete-time Markov process with state-dependent transition probabilities. By the law of iterated expectations, this process is also a martingale, unconditional on the state $\theta \in \Theta$.

2.3. PATIENCE

2.3.1 Special Assumptions

Our analytical results below will require a number of restrictions on the above general herding model. For instance, if $(\Theta, \mathcal{F})$ is a continuum subset of $\mathbf{R}$ equipped with the Borel sigma-algebra, then the space of mappings X is rather unwieldy: the space of measurable mappings from the measures Σ over Θ to the measures $\mathcal{P}(A)$ over A.

So just as in SS, we assume that $\Theta = \{H, L\}$ (or is more generally finite). High and low states are equilikely ex ante, and so have prior $\lambda(L) = \lambda(H) = 1/2$. Private belief σ expresses the chance of state H, and thus Σ is the interval $[0,1]$, and $\mathcal{G}$ the Borel sigma-algebra over $[0,1]$. Let supp(μ) be the (common) support[4] of each probability measure μ^θ.

[3] Absent a unique solution, we must take an arbitrary measurable selection from the solution correspondence.

[4] As usual, the support of a measure on the Borel-algebra is the smallest closed set of measure 1.

If $\text{supp}(\mu) \subseteq (0,1)$, then we call the private beliefs *bounded,* while if $\text{co}(\text{supp}(\mu)) = [0,1]$, they are *unbounded.* In that case, arbitrarily strong private signals can occur.

Lee (1993) showed that with these restrictions, a continuous action space can easily allow for simple statistical learning: With the common knowledge of individual n's public belief, observation of his action perfectly reveals his private signal. We thus make the standard lumpiness assumption that $A = \{a_1, \ldots, a_M\}$ is finite. We assume that no action is dominated by the others, and this implies the simple interval structure that action a_m is optimal exactly when the posterior ρ is in some sub-interval of $[0,1]$. We order the actions such that a_m is optimal for posteriors $\rho \in [\bar{r}_{m-1}, \bar{r}_m]$, where $0 = \bar{r}_0 < \bar{r}_1 < \ldots < \bar{r}_M = 1$.

2.3.2 The Social Planner

The Planner's Problem. Consider now an informationally constrained social planner SP trying to maximize the expected discounted average welfare of the individuals in the herding model: If the realized payoff sequence realized is $\langle u_n \rangle$, then the SP's payoff (value) is $v(\pi, \delta) = (1-\delta)\sum_{n=1}^{\infty} \delta^{n-1} u_n$. We also assume that SP neither knows the state nor can observe the individuals' private signals, but can both observe and punish/reward any actions taken. For now, simply assume that SP has the power to instruct each individual to use any rule x.

Detailed analysis of SP's problem is in the appendix. SP respects the ordering of actions, i.e. higher signals map into higher actions (see Lemma A-2.2):

Lemma 2.1 (Interval Structure) *For any π, the optimal rule $x \in X$ is described by thresholds $0 = \bar{x}_0 \leq \bar{x}_1 \leq \ldots \leq \bar{x}_M = 1$ such that action a_m is taken when $\sigma \in (\bar{x}_{m-1}, \bar{x}_m)$, and when $\sigma = \bar{x}_m$ there is a randomization between a_m and a_{m+1}.*

It is immediate that SP can achieve at least the welfare that comes out of the equilibrium of the herding model, for one option is to implement the equilibrium rules. The interesting question is how different is SP's outcome from the equilibrium? To answer this we describe SP's long run behavior in detail. First, as is standard, and proved in SS,

Lemma 2.2 *The belief process $\langle \pi_n \rangle$ is a martingale unconditional on the state, which converges a.s. to some limiting stochastic variable π_∞. It is concentrated on $(0,1]$ in*

Figure 2-1: **Typical value function.** Stylized Graph of $v(\pi,\delta)$, $\delta \geq 0$

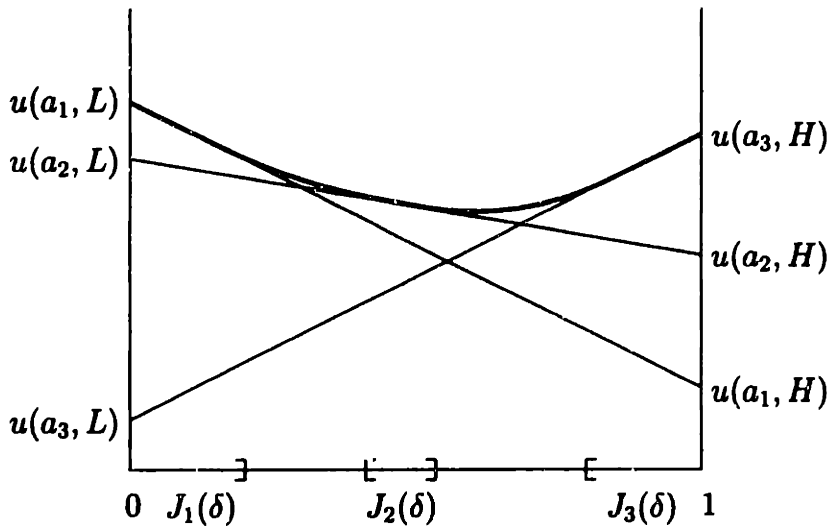

state H.

So beliefs must settle down, and SP is never dead wrong about the state. The next result states that the limiting belief π_∞ precludes further learning.

Proposition 2.1 (Absorbing Basins) *For each $a_m \in A$, a possibly empty interval exists $J_m(\delta) \subset [0,1]$, such that when $\pi \in J_m(\delta)$, SP optimally chooses z which a.s. induces a_m.*

- *For all $\delta \in [0,1)$, the limit belief π_∞ is concentrated on the basins $J_1(\delta) \cup \cdots \cup J_M(\delta)$.*
- *With unbounded private beliefs, $J_1(\delta) = \{0\}$ and $J_M(\delta) = \{1\}$; all other $J_m(\delta)$ are empty.*
- *If the private beliefs are bounded, then $J_1(\delta) = [0, \underline{\pi}(\delta)]$ and $J_M(\delta) = [\bar{\pi}(\delta), 1]$, where $0 < \underline{\pi}(\delta) < \bar{\pi}(\delta) < 1$. The larger is δ, the smaller are all basins. For large enough δ, all basins disappear except for J_1 and J_M, while $\lim_{\delta\to 1} J_1(\delta) = \{0\}$ and $\lim_{\delta\to 1} J_M(\delta) = \{1\}$.*

This characterization of the stationary points of the stochastic process of beliefs $\langle \pi_n \rangle$ directly generalizes the analysis for $\delta = 0$ in SS. See figure 2-1 for an illustration of how the basins are determined from the shape of the optimal value function. All but the limit belief result are established in the appendix. To see why that one is true — that a *limit cascade* must occur — observe that for any belief $\hat{\pi}$ not in any basin, at least two signals in A are realized with positive probability. Moreover, the highest such signal is more likely in state H, and the lowest more likely in state L. It follows that next period's belief will differ from $\hat{\pi}$ with positive probability. By the results in appendix B of SS, $\hat{\pi}$ cannot be in the support of π_∞.

Proposition 2.2 (Long Run Learning)

- *For unbounded private beliefs, π_∞ is concentrated on the truth for any $\delta \in [0, 1)$.*
- *With bounded private beliefs, learning is incomplete for any $\delta \in [0, 1)$: Unless $\pi_0 \in J_M(\delta)$, there is positive probability in state H that π_∞ is not in $J_M(\delta)$.*
- *The chance of incomplete learning with bounded private beliefs vanishes as $\delta \uparrow 1$.*

Proof: Given the absorbing basin characterization of Proposition 2.1, the unbounded beliefs result is a corollary of Lemma 2.2, while the incomplete learning conclusion for bounded beliefs follows just as in Theorem 1 of SS. We now extend that proof to establish the limiting result for $\delta \uparrow 1$. First, Proposition 2.1 assures us that for δ close enough to 1, π_∞ places all weight in $J_1(\delta)$ and $J_M(\delta)$. The likelihood ratio $\ell_n \equiv (1 - \pi_n)/\pi_n$ is a martingale conditional on state H. Since all private beliefs σ have likelihood ratio $(1-\sigma)/\sigma$ bounded above by some $\bar{\ell} < \infty$, the sequence $\langle \ell_n \rangle$ is bounded above by $\bar{\ell}(1 - \underline{\pi}(\delta))/\underline{\pi}(\delta)$, and the mean of ℓ_∞ must equal its prior mean $(1 - \pi_0)/\pi_0$. Since $\lim_{\delta \to 1} \underline{\pi}(\delta) = 0$, the weight that π_∞ places on $J_1(\delta)$ must vanish as $\delta \to 1$. □

Observe how incomplete learning to some extent plagues even an extremely patient SP.

We are now positioned to reformulate the learning results at the level of actions. Fix $\delta \in [0, 1)$ and consider a belief π which is very close to but not inside the basin $J_m(\delta)$. Conditional on taking some action $a \neq a_m$, there is a large loss of expected current payoff. To make up for this there must be a large gain in next period value, and since the value function is continuous it must be that the updated belief conditional on observing a is far from π. By proposition 2.2 we then have the following immediate corollary.

Proposition 2.3 (Herds)

- *For unbounded private beliefs, a herd eventually starts on the correct action.*
- *With bounded private beliefs, a herd on some action eventually starts. Unless $\pi_0 \in J_M(\delta)$, a herd arises on an action other than a_M with positive chance in state H for any $\delta \in [0, 1)$.*
- *The chance of an incorrect herd with bounded private beliefs vanishes as $\delta \uparrow 1$.*

It is no surprise that SP ends up with full learning with unbounded beliefs, for even selfish individuals will. More interesting is that SP optimally incurs the risk of an ever-lasting incorrect herd. Herding is truly a robust property of the observational learning paradigm.

Optimal Taxes. How does SP steer the choices away from the myopic equilibrium solution? He taxes or subsidizes the actions according to the following simple scheme. Given the current public belief π, if the individual takes observable action a he receives the (possibly negative) transfer $\tau(a|\pi)$. Faced with such incentives, it is easy to see that individuals still optimally choose private belief threshold rules.

Now, let the belief be π, and let x^{*} be the optimal rule for SP. How are the taxes determined? When the individual observes the private belief σ, it is mapped into the posterior $\rho(\pi,\sigma) = \pi\sigma/[\pi\sigma + (1-\pi)(1-\sigma)]$. The selfish herder's threshold $\bar{x}_m$ is then determined by the indifference equation $\bar{u}_{a_m}(\rho(\pi,\bar{x}_m)) + \tau(a_m|\pi) = \bar{u}_{a_{m+1}}(\rho(\pi,\bar{x}_m)) + \tau(a_{m+1}|\pi)$. So, only the difference $\tau(a_{m-1}|\pi) - \tau(a_m|\pi)$ matters for how the individuals trade-off between the two actions, and the SP can ensure that the threshold belief is optimally chosen, $\bar{x}_m = \bar{x}^{*}_m$, by suitably adjusting this difference.

Since all face the same incentives if a constant is added to all M taxes, we can also insist that SP achieve *expected budget balance* each period: i.e. the expected contribution from everyone is zero, or $0 = \sum_{m=1}^{M} \psi(a_m|\pi,\bar{x})\tau(a_m|\pi)$. This uniquely determines the transfers.

It would be nice to provide some robust properties of these taxes, but these are by no means obvious. All we can say is that individuals are rewarded for experimenting, i.e. making non-myopic choices, thereby providing information to later individuals. What these high-information choices are is hard to tell.

2.4. MODEL COMPARISON

We now return to the general herding model of section 2.2. A first step in recasting that general model of observational learning as a single person experimentation problem is to respect the individuals' selfishness. Thus, we must study an impatient experimenter with discount factor 0 (no 'active experimentation').

But to avoid a forgetful experimenter, we must regard the observational learning story

OBSERVATIONAL LEARNING MODEL	IMPATIENT EXPERIMENTER MODEL
States $\theta \in \Theta$	States $\theta \in \Theta$
Belief after n individuals π_n	Belief after n observations π_n
Optimal decision rule $x \in X$	Optimal action $x \in X$
Private signal of individual n, σ_n	Randomness in the nth experiment
Action taken by individual, $a \in A$	Observable signal $a \in A$
Density over actions $\psi(a\|\theta, x)$	Density over observables $\psi(a\|\theta, x)$
Payoffs (private information)	Payoffs (unobserved)

Table 2.1: **Embedding.** This table displays how our single-type observational learning model fits into the impatient single person experimentation model.

from a new perspective. Consider individual n, who uses both the public belief π_n and his private signal σ_n in forming and acting upon his posterior beliefs ρ_n. We may separate these two steps using the conditional independence of π_n and σ_n. Mr. n can be regarded as: (i) observing π_n, but *not* his private signal; (ii) optimally determining the rule $x \in X$, and submitting it to an agent 'choice' machine; and (iii) letting that machine observe his private signal and take his action $a \in A$ for him. The ultimate payoff $u(a, \theta)$ is unobserved, lest that provide an additional signal of the state of the world. If private beliefs σ have distribution μ^θ in state θ, the impatient experimenter will choose the same optimal decision rule x described in section 2.2, resulting in action $a \in A$ with chance $\psi(a|\theta, x)$.

We now precisely describe the single-person experimentation model. The parameter space is Θ. In period n, the experimenter EX chooses an action (the rule) $x \in X$. Given x, a random observable statistic $a \in A$ is realized with probability $\psi(a|\theta, x)$ in state θ. Finally, EX updates beliefs using this information alone.[5] Table 2.1 summarizes the embedding. EX's problem is formally the same as that of SP in the previous section.

Notice how this mapping addresses both lead puzzles. First, the experimenter never knows the private beliefs σ, and thus does not forget them. Second, the pathological learning outcomes are entirely consistent with EK's generic complete learning result for

[5]This experimentation model does not strictly fit into the EK mold, where the instantaneous reward in period n depends only on the action and the observed signal, but (unlike here) not on the parameter θ in Θ. This is the structure of Aghion, Bolton, Harris, and Jullien (1991) (ABHJ), where payoffs are not necessarily observed. If we wish, we may simply posit that the experimenter has fair insurance, and simply earns his expected payoff each period rather than his random realized payoff. Then his behaviour will be exactly the same, yet he will not learn anything from the payoff.

models with finite action and state spaces. Simply put, *actions do not map to actions but to signals when one rewrites the observational learning model as an experimentation model.* The true action space for EX is the infinite space X.

SS considered two major modifications of the herding model. One was to add i.i.d. noise to the individual decision problem. Noise is easily incorporated here by adding an exogenous chance of a noisy signal (random action). SS also allowed for T different types of preferences, with individuals randomly drawn from one or the other type population. Multiple types is addressed by simply imagining that EX chooses a T-vector of optimal decision rules from X^T with (only) the choice machine observing the task and private belief, and choosing the action a as before.

Finally, let us interpret our characterization of SP's optimal policy as well-known experimentation pathologies. Proposition 2.1 is an expression of EK's Theorem 5 that the limiting belief π_∞ precludes further learning. The incomplete learning result of Proposition 2.2 does not fall under the rubric of EK's Theorem 9, where it is shown that if the optimal value function v is *strictly convex* in beliefs π, learning is complete for δ close enough to 1. For here, EX optimally behaves myopically for very extreme beliefs: $v(\pi) = u_{a_1}(\pi)$ for π near 0, and $v(\pi) = u_{a_M}(\pi)$ for π near 1, both *affine* functions. This points to the source of the incomplete learning: lumpy signals rather than impatience.

2.5. CONCLUSION

Very recent models of rational social learning, like Banerjee and Fudenberg (1995) and Smith and Sørensen (1996b), have relaxed the assumption that the action history is perfectly observed. This renders the public belief process no longer a martingale — a fact that we are exploring in a work in progress. As a result, a mapping of these models into the experimentation literature, where the martingale property perforce obtains, is impossible.

Techniques from Smith and Sørensen (1996b) can be used to broaden the mapping beyond the simple realm of action observation to models with posterior belief misperception: That is, only an imperfect signal (drawn from a finite set) of one's posterior belief is observed. Provided the entire history of such signals is observed, everything goes through.

2.A. APPENDIX

We now rigorously establish Proposition 2.1 in section 2.3. Our analysis here is inspired of that in ABHJ and in sections 9.1 and 9.2 of Stokey and Lucas (1989).

A strategy s_n for period n is a map from Σ to X. It prescribes the rule $x_n \in X$ which must be used, given belief π_n. The social planner chooses a strategy profile $s = (s_1, s_2, \ldots)$, which in turn determines the stochastic evolution of the model — i.e. a distribution over the sequences of realized actions, signals, payoffs, and beliefs.

The value function $v(\cdot, \delta) : \Sigma \mapsto \mathbf{R}$ for the social planning problem with discount factor δ is $v(\pi, \delta) = \sup_s E[(1-\delta)\sum_{n=1}^{\infty} \delta^{t-1} u_n | \pi]$ where the expectation is taken using the distribution of processes implied by s. Now, define the Bellman operator T_δ by

$$T_\delta v(\pi) = \sup_{x \in X} \left\{ \sum_{a_m \in A} \psi(a_m|\pi, x) \left[(1-\delta)\bar{u}_m(q(\pi, x, a_m)) + \delta v(q(\pi, x, a_m))\right] \right\} \tag{A-1}$$

where $q(\pi, x, a)$ is the Bayes-updated belief obtained from π when a is observed and rule x is applied. Recall that $\bar{u}_m(\pi) = \pi u(a_m, H) + (1-\pi)u(a_m, L)$ denotes the expected payoff from a_m at belief π. We used the fact that $\bar{u}_m$ is affine to reach the expression for the current value.

Note that for $v \geq v'$ we have $T_\delta v \geq T_\delta v'$. As is standard, T_δ is a contraction, and $v(\cdot, \delta)$ is its unique fixed point in the space of bounded, continuous, weakly convex functions. Next define the function v_0 by $v_0(\pi) = \max_m \bar{u}_m(\pi)$, that is the expected utility an individual would obtain by choosing the myopically optimal action.

Lemma A-2.1 *The iterated $T_\delta^n v_0$ is a pointwise increasing sequence of weakly convex functions which converge to $v(\cdot, \delta)$. For $\delta_1 \geq \delta_2$, $v(\pi, \delta_1) \geq v(\pi, \delta_2)$ for all π.*

Proof: Consider maximization of $\sum_{a_m \in A} \psi(a_m|\pi, x)\left[(1-\delta)\bar{u}_m(q(\pi, x, a_m)) + \delta v(q(\pi, x, a_m))\right]$ over x for given π. One possible policy is to choose x such that the myopically optimal signal a_m occurs with probability one. Then $q(\pi, x, x(\sigma)) = \pi$ a.s. and the obtained value is $v_0(\pi)$. Optimizing over all $x \in X$, we see that $T_\delta v_0(\pi) \geq v_0(\pi)$ for all π. By induction, it follows that $T^n v_0 \geq T^{n-1} v_0$, so we get a pointwise increasing sequence converging to the fixed point $v(\cdot, \delta)$.

Now consider any convex function v with $v \geq v_0$ pointwise. For any chosen x, we have $\sum_{a_m \in A} \psi(a_m|\pi, x)\bar{u}_m(q(\pi, x, a_m)) \leq \sum_{a_m \in A} v(q(\pi, x, a_m))$. Now, consider $\delta_1 \geq \delta_2$. Since there is more weight on the larger component of (A-1), we have $T_{\delta_1} v_0 \geq T_{\delta_2} v_0$. Simple induction verifies that $T_{\delta_1}^n v_0 \geq T_{\delta_2}^n v_0$, since one possible policy under δ_1 is to choose the x optimal under δ_2. □

Lemma A-2.2 (Interval Structure) *For any π, one optimal rule $x \in X$ is described by thresholds $0 = \bar{x}_0 \leq \bar{x}_1 \leq \ldots \leq \bar{x}_M = 1$ such that action a_m is taken when $\sigma \in (\bar{x}_{m-1}, \bar{x}_m)$, and when $\sigma = \bar{x}_m$ there is a randomization between a_m and a_{m+1}.*

Proof: We prove via the Bellman equation that any rule x which does not satisfy this simple interval structure can be improved upon. For $m_1 < m_2$ define Σ_1 resp. Σ_2 to be the subsets of signals in Σ which are mapped with positive probability into a_{m_1} resp. a_{m_2}, and assume that they are *not* ordered with $\Sigma_1 \leq \Sigma_2$. If $q(\pi, x, a_{m_1}) > q(\pi, x, a_{m_2})$, the rule is improved by remapping signals leading to a_{m_1} into a_{m_2}, and vice versa, by our ordering of actions.

Next, assume that $q(\pi, x, a_{m_1}) \leq q(\pi, x, a_{m_2})$. For any $\bar{x} \in (0,1)$ let $\tilde{\Sigma}_1(\bar{x}) \equiv (\Sigma_1 \cup \Sigma_2) \cap [0, \bar{x}]$ and $\tilde{\Sigma}_2(\bar{x}) \equiv (\Sigma_1 \cup \Sigma_2) \cap [\bar{x}, 1]$. Consider then the modified rule $\tilde{x}$ which equals x, except that a_{m_i} is taken for signals in $\tilde{\Sigma}_i(\bar{x})$, and where $\bar{x}$ is calibrated such that $\psi(a_{m_i}|\pi, x) = \psi(a_{m_i}|\pi, \tilde{x})$ (it may be necessary for $\tilde{x}$ to randomize over the two actions at signal $\bar{x}$ to accomplish that). Since signals more in favor of state H are mapped into a_{m_1} under $\tilde{x}$, we find $q(\pi, \tilde{x}, a_{m_1}) \leq q(\pi, x, a_{m_1})$, and similarly $q(\pi, \tilde{x}, a_{m_2}) \geq q(\pi, x, a_{m_2})$. Thus, the use of $\tilde{x}$ instead of x implies a mean preserving spread of the updated belief, and since the value function is weakly convex, this weakly improves the value in the Bellman equation. □

Proposition 2.1 is now established in claims.

Claim A-2.1 (Basins) *For each $a_m \in A$, a possibly empty interval $J_m(\delta)$ exists, s.t. when $\pi \in J_m(\delta)$, SP optimally chooses x such that a_m occurs a.s. (learning stops). For any $\delta \in [0,1)$, $0 \in J_1(\delta)$ and $1 \in J_M(\delta)$.*

Proof: For the first half, we really need only prove that $J_m(\delta)$ must be an interval. If $\pi \in J_m(\delta)$, then a_m is the optimal choice, and the value is $v(\pi, \delta) = v_0(\pi) = \bar{u}_m(\pi)$. Conversely, if $v(\pi, \delta) = v_0(\pi) = \bar{u}_m(\pi)$ then $\pi \in J_m(\delta)$ and a_m is the optimal choice. As $\bar{u}_m(\pi)$ is an affine function of π, and $v(\cdot, \delta)$ is weakly convex, $J_m(\delta)$ must be an interval. Second, for $\pi_n = 0$, a_1 is optimal in the short run, and no matter which rule is applied, the updated $\pi_{n+1} = \pi$ a.s. Similarly when $\pi_n = 1$. □

Claim A-2.2 (Monotonicity) *The larger is δ, the smaller are all basins: $\forall a_m \in A$, $\forall \delta_1, \delta_2 \in [0,1)$, if $\delta_1 \geq \delta_2$ then $J_m(\delta_1) \subseteq J_m(\delta_2)$.*

Proof: As seen in Lemma A-2.1, $v(\pi, \delta_1) \geq v(\pi, \delta_2) \geq \bar{u}_m(\pi)$ for all π, when $\delta_1 \geq \delta_2$. For $\pi \in J_m(\delta_1)$, we know $v(\pi, \delta_1) = \bar{u}_m(\pi)$ and thus $v(\pi, \delta_2) = \bar{u}_m(\pi)$. The optimal value can thus be obtained by inducing a_m a.s., so that $\pi \in J_m(\delta_2)$. □

Claim A-2.3 (Unbounded Beliefs) *With unbounded private beliefs, $J_1(\delta) = \{0\}$ and $J_M(\delta) = \{1\}$; all other $J_m(\delta)$ are empty.*

Proof: SS establish that all $J_m(0)$ are empty, except for $J_1(0) = \{0\}$ and $J_M(0) = \{1\}$. Now apply Claim A-2.1 and Claim A-2.2. □

Claim A-2.4 *If the private beliefs are bounded, then $J_1(\delta) = [0, \underline{\pi}(\delta)]$ and $J_M(\delta) = [\bar{\pi}(\delta), 1]$, where $0 < \underline{\pi}(\delta) < \bar{\pi}(\delta) < 1$.*

Proof: We prove that for sufficiently low beliefs it is optimal to choose a rule x such that a_1 occurs with probability one; the argument for large beliefs is very similar. Since action a_1 is optimal at belief $\pi = 0$, and is not weakly dominated, there must be some positive length interval $I = [0, \tilde{\pi}]$ on which $\bar{u}_1(\pi) = v_0(\pi)$, i.e. a_1 is the optimal choice for beliefs in I. Moreover, by the affinity of each $\bar{u}_m$, $\exists \underline{u} > 0$ such that on the interval $[0, \tilde{\pi}/2]$ $\bar{u}_1(\pi) > \bar{u}_m(\pi) + \underline{u}$ for all $m \neq 1$.

No observation $a \in A$ can ever yield a stronger signal than any $\sigma \in \text{supp}(\mu) \subseteq [\underline{\sigma}, \bar{\sigma}] \subset (0,1)$. So any initial belief π is updated to at most $\bar{q}(\pi) = \pi\bar{\sigma}/[\pi\bar{\sigma} + (1-\pi)(1-\bar{\sigma})]$. For π sufficiently small, $\bar{q}(\pi) \in [0, \tilde{\pi}/2]$ and $\bar{q}(\pi) - \pi$ is arbitrarily small. By the weak convexity of v, $v(\bar{q}(\pi), \delta) - v(\pi, \delta)$ is then arbitrarily small — in particular, less than $\underline{u}(1-\delta)/\delta$ for small enough π. It follows directly from the Bellman equation $T_\delta(v) = v$ corresponding to (A-1) that it is suboptimal to risk any other outcome than a_1 for such small beliefs. □

Claim A-2.5 (Limiting Patience) *For large enough δ, all basins disappear except for $J_1(\delta)$ and $J_M(\delta)$, while $\lim_{\delta\to 1} J_1(\delta) = \{0\}$ and $\lim_{\delta \to 1} J_M(\delta) = \{1\}$.*

Proof: Select an $x^* \in (1/2, 1)$ such that $1 > \mu^H([x^*, 1]) > \mu^L([x^*, 1]) > 0$. Updating π with observation of the event that σ is in $[x^*, 1]$ yields the belief $q(\pi) = \pi\mu^H([x^*,1])/[\pi\mu^H([x^*,1]) + (1-\pi)\mu^L([x^*,1])]$. Over any compact subinterval $I \subset (0,1)$, $q(\pi) - \pi \geq \varepsilon$ for some $\varepsilon > 0$.

Start with any given $\delta \in [0,1)$ and m with $1 < m < M$. Write $J_m(\delta) = [\pi_1, \pi_2]$ where $0 < \pi_1 \leq \pi_2 < 1$, and pick I to contain $J_m(\delta)$ and take ε as described. We will consider the possible choice $\bar{x}_{m-1} = 0$, $\bar{x}_m = x^*$, $\bar{x}_{m+1} = 1$ (see Lemma A-2.2). q maps the interval $[\pi_2 - \varepsilon/2, \pi_2]$ into (but not onto) $[\pi_2 + \varepsilon/2, 1]$. Choose $\ddot{u} > 0$ so large that $\bar{u}_m(\pi) < \bar{u}_{m+1}(\pi) + \ddot{u}$ for all $\pi \in [0,1]$, and $\bar{v} > 0$ so small that $v(\pi, \delta) > \bar{u}_m(\pi) + \bar{v}$ for all $\pi \in [\pi_2 + \varepsilon/2, 1]$. As the value functions increase with δ by Lemma A-2.1, we have $v(\pi, \delta') > \bar{u}_m(\pi) + \bar{v}$ for all $\delta' > \delta$. If $\delta' > \delta$ is so large that $(1-\delta')\ddot{u} < \delta'\bar{v}$, it follows from the Bellman equation that our suggested policy x beats inducing a_m a.s. when $\pi \in [\pi_2 - \varepsilon/2, \pi_2]$. By iterating this procedure a finite number of times, each time cutting away a chunk of size $\varepsilon/2$ of the interval $J_m(\delta)$, we see that the interval must have vanished for large enough δ.

If $m = 1$ or $m = M$, this procedure can still be applied repeatedly, to show that $J_m(\delta) \cap I$ vanishes for large enough δ for any closed $I \subset (0,1)$. □

We now provide the necessary analysis for Proposition 2.3. In light of the interval structure found in Lemma A-2.2, SP's problem can be reduced to that of determining the probabilities $\psi(a_m|\pi)$ with which each action should be chosen (i.e., just choosing how large a chunk of the signal space should map into each given action). Thus, the choice set is the compact M-simplex. The objective function in the Bellman equation corresponding to A-1 is continuous in this choice vector and in π, and it follows from the Theorem of the Maximum (e.g. Theorem I.B.3 of Hildenbrand (1974)) that the non-empty correspondence of optimal rules is upper hemi-continuous in π.

Lemma A-2.3 (Overturning) *Given $\delta \in [0,1)$, and given $a_m \in A$ with $J_m(\delta) \neq \emptyset$, there exists a $\varepsilon > 0$ and an open interval K with $J_m(\delta) \subset K$, such that $\forall \pi \in K$ and $\forall a \neq a_m$, $||q(\pi, x, a) - \pi|| > \varepsilon$ when x is the optimal rule.*

Proof: For $\pi \in J_m(\delta)$ we know that the only optimal rule is to a.s. induce a_m. As the correspondence of optimal rules is u.h.c. in π, it must be true that near $J_m(\delta)$ it is optimal that only the most extreme elements of supp(μ) induce actions other than a_m. In the case of bounded beliefs, π is far from 0 and 1, so it is immediate that the public belief must move far when observing actions other than a_m.

With unbounded beliefs, an extra argument is needed. Consider WLOG π near 0. We have $q(\pi, x, a_1) \leq \pi$, i.e. the updated belief is close to the current, so $v(q(\pi, x, a_1), \delta) - v(\pi, \delta)$ is close to zero. Conditional on choosing an action other than a_1, there is a large immediate loss which must be made up for by gained future value (according to the Bellman equation). Since v is continuous, it follows that $q(\pi, x, a_m)$ must be far from π. □

Chapter 3

Rational Social Learning with Random Sampling

3.1. INTRODUCTION

Consider the following canonical model: An infinite sequence of exogenously ordered individuals must each make a once-in-a-lifetime binary decision, with uncertain common payoff. Decisions are made on the basis of private information and the knowledge of what all predecessors have done. In this setting, Banerjee (1992) and Bikhchandani, Hirshleifer, and Welch (1992) independently found that a bad herding outcome may arise: There is a positive chance that everyone eventually settles upon the less profitable decision. Smith and Sørensen (1996a) (hereafter SS) have recently shown, among other things, that this conclusion obtains precisely when the signals possessed by the individuals have bounded informativeness. Otherwise, barring other innovations in the model, complete learning obtains.

One crucial and arguably untenable assumption in the whole 'herding' literature is the *perfect* observability of history, and in particular, knowledge of the exact order of all decisions. A modelling assumption to be sure, but how important is it? In fact, SS exploited this assumption to deduce the *overturning principle*: A single individual's choice to violate a would-be herd can overturn the weight of ten trillion predecessors' recorded actions. An arguably less objectionable observational structure was adopted earlier in Smith (1991), who posited that entrants could only observe *aggregates,* namely, the number of individuals choosing each option, and not the order in which these actions were taken.[1] This paper moves one step beyond the constraints of these observational regimes. Inspired by Banerjee and Fudenberg (1995) (hereafter BF) we first consider learning from *samples*: Individual n observes a random, ordered or unordered sample of k (itself potentially random) of his predecessors' actions, possibly with some weighting (e.g. favoring more recent individuals). This might reflect word-of-mouth learning, where individuals profit from the action choice of randomly chosen people before deciding to act.

What is Rational Social Learning? Let's parse this. Since Bray and Kreps (1987), *rational* learning in any context has come to mean learning that is fully Bayesian, and exploits all available information. *Social* learning is that which occurs in a world of sequential decision-making by privately informed individuals. The imprimatur of *rational social learning* is that everyone learns fully from his predecessors' *posterior beliefs.* Each Bayes posterior is eventually woven into the informational fabric to be sampled by posterity. Observations of predecessors' signals or payoffs is not a part of social learning. Yet one can argue that such time intensive activities do occur: People enjoy sharing from their experiences. True enough, but such personal information does not add to the 'stock' of social knowledge unless it is passed on again, this time as hearsay embedded in the posterior beliefs. Social learning describes the latter form of information acquisition. As in BF, payoff observations are better modelled as part and parcel of the private signals obtained by the individuals. For a payoff observation is an ideal informative signal, being a function of the state of the world alone, and not the action history.

In the herding model, the vehicle for social learning was action observation. But Banerjee (1992) and Lee (1993) have pointed out that when the action space and payoff functions are continuous, this may well permit a perfect inference of predecessors' signals. This transparently reduces to pure *statistical learning,* and thus yields quick complete

[1]That paper embellished the analysis by assuming individuals played 'two-armed bandits' after making their observation. But at the aggregate level, individual 'error persistence' (settling forever on the wrong action) is just as easily captured with the assumption that individuals live but once and sometimes err in their one-shot decisions. Also, it was not clear to what extent the complete learning result in Smith (1991) was simply an artifact of the continuum of players assumption.

learning in the long run. What has therefore additionally generally set social learning apart from its statistical counterpart is that posterior beliefs are garbled through a *coarse* or *discrete* filter.

An action choice is certainly one natural coarse informative statistic of a posterior belief, but it need not be the only one. In general, one can learn from any discrete informative signal of an individual's posterior belief. Such a more encompassing framework might well describe *plain conversation,* for instance: Verbal discourse may allow one to make an informed estimate of the posterior of another, but also has the potential to totally mislead. Shiller (1995), for one, underscores that this transparent information transmission mechanism is ignored at one's peril in the social learning context.[2] Since posterior beliefs are informative of the true state of the world, we are essentially capturing the necessary and sufficient conditions for social learning to be a fruitful enterprise.

A Foundation for Rational Social Learning. There is unfortunately a major theoretical roadblock to further progress. Once one abandons the assumption that the *order* of one's historical observations is known, the model can no longer be resolved by martingale analysis — for beliefs no longer constitute a filtration. To resolve the difficulties, we adopt the basic principle of BF. By sheer imitation of a randomly drawn action from his sample, an individual can always guarantee himself the same expected payoff as his average sampled predecessor. Under general assumptions on the sampling mechanism in use, this principle serves to prove that the expected welfare of the average sampled population keeps increasing, and that again can be used to prove when complete learning arises in the long run.

The continuum of agents assumption in BF makes it a continuous and deterministic version of a stochastic discrete agent model such as ours. But since learning *is* about uncertainty, we find our discrete agent approach much more natural. The modelling simplicity of the BF is manifest:[3] Theirs is a continuous time dynamical system that admits

---

[2]He writes: *This flow of conversation serves to exchange a wide variety of information to be held in common by the group.*

[3]The precursor 'rules of thumb' paper Ellison and Fudenberg (1995) is not a 'rational' learning model; our work is therefore not easily comparable to theirs.

This same comment applies to a superficially similar model of social learning due to Arthur and Lane (1994). There, individuals observe some predecessors' actions and random payoffs, but only attempt to resolution by means of an easily-interpreted Lyapounov function. Yet, a comparison of our results to theirs proves that their deterministic approach often gives a false description of the underlying stochastics. We are able to provide richer scenarios for incomplete learning.

The 'Vives Effect'. Consider the reason for the infamous bad herding outcome: Over time, individuals tend to rely less on their private information, and thus add less to the 'informational pool' of history. This phenomenon is described (on page 331) in Vives (1993), writing on the speed of learning in a social learning model.

> ... information revelation through prices can be a victim of its own success ...
> An exogenous change which raises the informativeness of prices ... for a given responsiveness of agents to private information, is partially counterbalanced by the agents giving less weight to their private signals.

Vives' insight applies forwards and backwards: If information transmission becomes easier, the new information inflow slows down, while if learning from others is rendered more arduous, the information inflow accommodates by speeding up. In other words, social learning tends to be 'self-correcting', and on balance, the net effect of more efficient learning or better informed early individuals is ambiguous. A very similar phenomenon is also present in Banerjee (1993) on the spread of 'rumors'. When the information transmission mechanism is endogenous, speeding up the spread of rumors does not affect the informativeness of receiving the rumor, for the faster it spreads, the sooner it should be received to be trusted. Banerjee in fact derives a perfect neutrality result.[4]

As testimony to the Vives Effect, it is not necessarily an improvement for all individuals that everyone observe more about their predecessors. For instance, individual 100 may be better off living in a world where everybody randomly sample 49 rather than 50 predecessors. For the actions of the first 99 individuals will tend to be more correlated in the second scenario than in the first.

The Vives Effect serves to confound most attempts at comparing social learning models. One cannot simply say "Information transmission is better in model X, and thus individuals

---

infer the action payoff distributions, and ignore the informativeness of the actual action choices made.

[4]But this neutrality result followed more simply because the speed that a rumor spreads in Banerjee's model in fact provides the only normalization of time, and thus ought have no other real effects.

are better off." This is an important issue which is well worth our elaboration. To design good frameworks for social learning, a planner must take into account that the population may be better off with less information transmission.

Ergodics: Complete and Incomplete Learning. Just as in SS, whenever the informativeness of private signals is bounded, there cannot be 'complete learning' in the sense of beliefs converging to a point mass on the truth. Once history becomes sufficiently persuasive, everyone subsequently acts irrespective of their private information, and therefore social informational accumulation stops. If this occurs in finite time, then with positive probability a significant proportion of individuals will then be taking the wrong action.

The assumption of unboundedly informative private signals is thus an essential ingredient for complete learning. But in contrast to SS, this may not be *sufficient,* depending on the nature of the posterior garbling process. For instance, when posterior garbling has the howsoever slight potential to radically mislead, as with plain conversation rather than action observation, we show that learning generally grinds to a halt short of complete learning at a *confounding point.*

Results from the statistical literature on urn schemes prove handy to characterize the action distribution under incomplete learning. Remarkably, in the most basic regime where each individual observes the action of one randomly drawn predecessor, the incomplete learning does not lead to herd behavior in the population as a whole. While eventually each individual optimally imitates the observed predecessor's action, action diversity persists in the population with probability one. So, there can be herding at the level of individuals without our observing the herd behavior in the composition of decisions in the population. While that may be important to keep in mind when applying herd models to real world observations, it is a knife-edge result. If individuals observe more than one predecessor, the imitation aspect will imply that almost all the population makes identical decisions in the long run.

Overview. In section 3.2, we set up the model, and describe exactly why sampling posterior beliefs is informative. Section 3.3 shows exactly how action observation regimes improve welfare over time. In sections 3.4 and 3.5, we study the model where a random sample of predecessors' actions is observed. Section 3.7 displays the confounding outcomes that arise with plain conversation. Section 3.6 discusses some difficulties with extending our analysis, but adds a wealth of nice results for a special version of the model.

3.2. THE MODEL

3.2.1 Signals and Bayesian Updating

The model structure is assumed to be common knowledge. We first introduce a background probability space $(\Omega, \mathcal{E}, \nu)$ that underlies all random processes in the model. The uncertainty is captured by two *states of the world* (or more simply, states), namely $\theta = H$ ('high') and $\theta = L$ ('low'). Formally, this means that the background state space Ω is partitioned into two events Ω^H and Ω^L, called H and L. WLOG, let the common prior belief be that $\nu(H) = \nu(L) = 1/2$. That individuals have common priors is a standard modelling assumption, see e.g. Harsanyi (1967–68), and crucial to what we have to say.

SIGNALS. An infinite sequence of exogenously ordered individuals $n = 1, 2, \ldots$ sequentially takes actions. Each individual n receives a private signal $\sigma_n \in \Sigma$ about the state of the world. Conditional on the state, the signals are assumed to be i.i.d. across individuals. It is common knowledge that the signal is distributed according to the probability measure μ^θ in states $\theta = H, L$. To ensure that no signal will perfectly reveal the state, we shall insist that μ^H and μ^L be mutually absolutely continuous. Consequently, there exists a positive and finite Radon-Nikodym derivative $g = d\mu^L/d\mu^H : \Sigma \to (0, \infty)$ of μ^L w.r.t. μ^H. To avoid trivialities, we rule out $g = 1$ almost surely, so that μ^H and μ^L are not the same measure; hence, some signals are informative about the state.

PRIVATE BELIEFS. Using Bayes' rule, the individual arrives 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 states $\theta = H, L$, the private belief p has distribution F^θ on $(0, 1)$. It will follow from Lemma 3.2 that F^H and F^L have a common support, denoted supp(F). By construction, co(supp$(F)) \equiv [a, b] \subseteq [0, 1]$

Figure 3-1: **Signals.** The first picture is an example graph of the signal densities in the two states. The second picture is the corresponding graph of the posterior cumulative distribution functions.

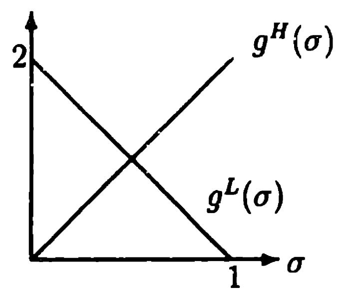

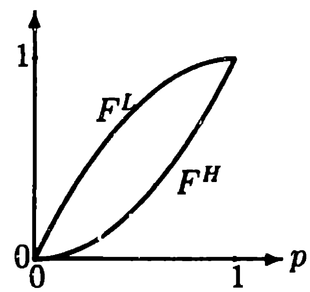

with $0 \leq a < 1/2 < b \leq 1$.[5] The strict inequalities follow directly from the assumption that μ^L and μ^H are distinct. We shall call the private beliefs *bounded* if $0 < a < b < 1$; if $\text{co}(\text{supp}(F)) = [0,1]$, private beliefs are *unbounded.* To exhaust all possibilities we should also consider supports that are bounded above and not below, and conversely, but this exercise yields no new insights.

EXAMPLE. In one specific example that we consider, μ^H and μ^L are described by densities on $(0,1)$, μ^L has the density $g^L(\sigma) = 2-2\sigma$ while μ^H has the density $g^H(\sigma) = 2\sigma$. Then we have $g(\sigma) = g^L(\sigma)/g^H(\sigma) = (1-\sigma)/\sigma$. From this specification it easily follows that $p(\sigma) = \sigma$, and thus $F^H(p) = p^2$ and $F^L(p) = 2p - p^2$. We see that the beliefs are unbounded, i.e. the support of F^H and F^L is all of $[0,1]$. Figure 3-1 shows the densities g^H and g^L, and we see how for instance signals near 0 are very strongly in favor of state L. Figure 3-1 also illustrates the very important (trivial) principle that signals in favor of state L are more likely to occur in state L than in state H. Ultimately, this is what drives the full learning results we strive for; there is an underlying tendency for the truth to be revealed in the signals. This is formally expressed in Lemma 3.2.

SIGNAL QUALITY. The above example may appear arbitrary, but there is method to our madness. Imagine that individuals are told one of two possible statistically true statements "with chance q, the state is high/low", where the *signal quality* q is distributed over $(0,1)$ according to the measure γ. Formally, the signal takes either the value σ_H or σ_L, where $P(\sigma = \sigma_H|H) = q = 1 - P(\sigma = \sigma_H|L)$.[6] After being told that the state

[5]Here, $\text{co}(A)$ denotes the convex hull of the set A.

[6]This is just one possible way of specifying signal quality. The reader may have different notions. We

is high (resp. low), the individual revised his belief to $q/[q + (1-q)] = q$ (resp. $1-q$), since his prior was flat. Thus, in state H, the generalized density of individuals with private belief p is $f^H(p) = p[d\gamma(p) + d\gamma(1-p)]$, while similarly in the low state it is $f^L(p) = (1-p)[d\gamma(p) + d\gamma(1-p)]$. This story of binary signals (i.e. being told 'high' or 'low') with variable certitude is less general than an arbitrary state-dependent signal with support Σ, since it implies the symmetry property of Lemma 3.1.

In the special case where p is uniform over $(0,1)$, we get the earlier densities $f^H(p) = 2p$ and yielding distributions $F^H(p) = p^2$ and $F^L(p) = 2p - p^2$. This simple symmetric density is the basis for many simulations we have performed.

Lemma 3.1 (Symmetry) *Under the signal quality structure, the distributions of private beliefs satisfy* $F^H(p) = 1 - F^L(1-p)$ *for all* $p \in (0,1)$.

Proof: From $f^H(p) = p[d\gamma(p) + d\gamma(1-p)]$ and $f^L(p) = (1-p)[d\gamma(p) + d\gamma(1-p)]$ we see that $f^H(p) = f^L(1-p)$. Consequently,

$$F^H(p) = \int_0^p f^H(r)\,dr = \int_{1-p}^1 f^L(r)\,dr = 1 - F^L(1-p)$$

□

General Signal Distributions. Even without symmetry, there is a one-to-one relationship between the distribution functions F^H and F^L, owing to the fact that posterior sampling is just as informative as sampling private signals. To formalize this, note that the assumption of mutual absolute continuity of μ^H and μ^L implies the same of the associated measures F^H and F^L. Thus, there exists a Radon Nikodym derivative $f = dF^H/dF^L$, and if both F^H and F^L have densities, we have $f(p) = f^H(p)/f^L(p)$. SS establish the following crucial but simple result.

Lemma 3.2 (No Introspection Condition) *Sampling private belief is just as informative as sampling signals, or* $dF^H/dF^L(p) = p/(1-p)$ *for all* $p \in (0,1)$.[7]

should point out that with our specification, the signal qualities q and $1-q$ are equally useful for the individual, and when the signal quality is below 1/2 the signals have 'the opposite meaning', in that σ_H is a signal in favor of state L and vice versa.

[7]Technically, the result holds only a.s. wrt. the measure F^L on $(0,1)$, but since we only care about what is true almost surely, we can just modify f such that it holds everywhere.

The lemma is intuitive:[8] No one ought to be able to make a further inference from the posterior he has by further inquiring as to its likelihood in the two states of the world. This lemma exactly quantifies the simple principle we noted from figure 3-1: Signals in favor of a state occur more frequently in that state than in other states.

SS establishes one useful consequence of Lemma 3.2. Knowing that p is in the interval $[0, r]$ is more in favor of state L than knowing that $p = r$ (unless r is below the support for p). Conversely for $[r, 1]$. Thus,

Lemma 3.3 *The inequalities $(1-p)F^H(p) \leq pF^L(p)$ and $(1-p)(1-F^H(p)) \geq p(1-F^L(p))$ obtain for all $p \in [0,1]$. Moreover, the first inequality is strict when $p > a$, and the second is strict wher $p < b$, where $[a, b]$ is the convex hull of the support of F^H.*

3.2.2 Action Choice

Individuals in the model care about their posterior beliefs to the extent that it enables them to make wise choices. Everyone can take one of M possible actions, $a_1, a_2, \ldots, a_M$, whose payoffs are state dependent. For convenience, we shall proceed under the assumption that there are only two actions. This is not crucial to our results, except in subsection 3.6.2.[9] Action a_m has a (common to all individuals) vN-M payoff $u^\theta(a_m)$ in state $\theta = H, L$, and we order the actions such that $u^H(a_1) \leq \ldots \leq u^H(a_M)$. In particular, with 2 actions we assume w.l.o.g. that $u^H(a_1) = u^L(a_1) = 0$, $u^H(a_2) = u$, $u^L(a_2) = -1$. These payoffs can be interpreted this way: Action a_1 is to decline from a given investment opportunity, while a_2 is to invest.

The objective of the individual is to take the action that maximizes his expected payoff. Using the optimal rule, the payoff becomes a *convex* function of the posterior belief. SS shows that with M possible actions, one can partition the belief space into a finite union of intervals $[0, \bar{r}_1), [\bar{r}_1, \bar{r}_2), \ldots, [\bar{r}_{M-1}, 1]$ such that the individual takes action a_k exactly when his posterior lies in $[\bar{r}_{k-1}, \bar{r}_k)$. This interval structure only depends on the payoffs.

[8] A proof is easy: The individual could use his sample of the realized private belief p to form a belief about the state of the world using Bayes' rule. By the law of iterated expectations, this gives him back his original belief. So, $p = f(p)/[1 + f(p)]$, as desired.

[9] In Appendix 3.A we generalize to more than two actions.

3.2.3 The Dynamic Behavior of Beliefs

We have shown that observing the private beliefs of predecessors is just as informative as observing their signals. Likewise, sampling posterior beliefs or actions allows one to make imperfect inferences about the underlying signals whence they were formed. The action of individual n is fixed after his turn. As such, there is an invariant probability distribution over actions that he will produce forever afterwards. Individuals learn from history by drawing some sample of predecessors' reports. Below, we will aim to admit as large a class of sampling rules as possible.

For any given sampling mechanism, the observation by individual n of any sample depends stochastically on the true state of the world and the realized history of length n. As the model is common knowledge, he can calculate the probabilities of making that observation in either state. Having drawn his information about the past, individual n can then form the *social belief* q_n in state H. This would be individual n's posterior belief in state H had he a purely neutral private belief.

As private signals and thereby posterior belief observations are random, $\langle q_n \rangle_{n=1}^{\infty}$ is a stochastic process. Individual n may form his posterior belief r_n from the social belief q_n and the private belief p_n,

$$r_n = \frac{p_n q_n}{p_n q_n + (1 - p_n)(1 - q_n)} \tag{1}$$

using Bayes' rule. We are interested in the asymptotic behavior of social beliefs $\langle q_n \rangle$. The most important question to answer is whether there is complete learning in the long run: Are the social beliefs eventually correct? If so, then by corollary, almost all individuals will attain the maximum payoff in the long run. For when social beliefs are very close to being correct, only misleading posterior signals and very extreme incorrect private signals will lead one to act suboptimally, which will occur very seldomly.

REMARK. From the individuals' standpoint, historical information may not be refined over time. Unlike the information structure in SS, one does not make observations that include one's predecessors'; therefore, the entire martingale apparatus of SS is inapplicable here. But from the modeller's point of view, which we sometimes adopt, the fully specified history clearly generates a filtration, or a sequence of progressively-refined σ-algebras.

3.3. IMPROVING WELFARE

3.3.1 A Key Result for Rational Social Learning

What follows is the key to much of our analysis. Individuals can use their observed sample of predecessors to obtain an expected welfare which is higher than that of the average sampled predecessor. Our version of this result is essentially a stochastic analogue to the deterministic Lemma 1 in Banerjee and Fudenberg (1995), granted model differences. The driving principle is very much the same: given an observed sample of actions, an individual could randomly select one of the reported actions and copy that action himself; given the freedom of choice, the individual can only do even better than this.

Now, fix a sampling mechanism and focus on some individual $n > 1$. A sample s is a finite collection of actions. There are only finitely many distinguishable samples, $s_1, \ldots, s_{K(n)}$ that can be drawn by n. According to the evolution of the model, we and the individual can calculate the probability of each sample in each state. So, let $P^\theta(s_k)$ be the probability of sample s_k in state θ. Also, let ρ_k be the probability that a uniformly drawn individual from s_k was someone who reported action a_2.[10] Then define the probability that an average sampled predecessor took a_2 to be

$$R^\theta = \sum_{k=1}^{K} \rho_k P^\theta(s_k) \tag{2}$$

The sampling by individual n, composed with uniform resampling of one by him, implies a distribution over the individual population $1, \ldots, n-1$. For each individual, there is a certain probability that he took action a_2 in each state. R^θ expresses the expected proportion of the predecessors, weighted by the sampling probabilities, who took action a_2. So, $(uR^H - R^L)/2$ is the expected welfare of the average sampled population.

If individual n observes the sample s_k he forms the social belief $q_k = P_k^H/[P^H(s_k) + P^L(s_k))$. The social and private beliefs are the basis for the individual's posterior r as

[10]Precisely, if there are j out of $\hat{j}$ individuals in s_k who reported a_2, then $\rho_k = j/\hat{j}$. If the size of the sample is zero, we let ρ_k be the probability with which an individual with posterior 1/2 will choose a_2. So, in that case, ρ_k is either one or zero. This is a matter of definition which may influence the interpretation of the average sample payoff when samples of size zero can occur.

in (1). Individual n takes action a_2 exactly when $ru - (1-r) \geq 0$, i.e. exactly when the private belief p is above the threshold $(1-q)/[uq + (1-q)]$. As we know the distribution functions F^θ for the private belief, we can calculate the probability that individual n takes a_2 in each state,

$$Q^\theta = \sum_{k=1}^{K} P^\theta(s_k)[1 - F^\theta(\frac{P^L(s_k)}{uP^H(s_k) + P^L(s_k)})] \tag{3}$$

We are now in a position to formulate

Theorem (Principle of Improved Welfare) *Assume that there is perfect action observation. The expected welfare of individual $n > 1$, namely $(uQ^H - Q^L)/2$, is (weakly) larger than that of his average sampled predecessor, $(uR^H - R^L)/2$.*

Proof: We have to prove that $uR^H - R^L \leq uQ^H - Q^L$. Using (2) and (3), write[11]

$$\begin{aligned} & uQ_n^H - Q_n^L - uR_n^H + R_n^L \\ &= \sum_{k=1}^{K} \rho_k \left\{ P_k^L F^L\left(\frac{P_k^L}{uP_k^H + P_k^L}\right) - uP_k^H F^H\left(\frac{P_k^L}{uP_k^H + P_k^L}\right)\right\} \\ &- \sum_{k=1}^{K}(1-\rho_k)\left\{P_k^L\left[1 - F^L\left(\frac{P_k^L}{uP_k^H + P_k^L}\right)\right] - uP_k^H\left[1 - F^H\left(\frac{P_k^L}{uP_k^H + P_k^L}\right)\right]\right\} \end{aligned} \tag{4}$$

Lemma 3.3 yields the inequalities

$$P_k^L F^L\left(\frac{P_k^L}{uP_k^H + P_k^L}\right) \geq uP_k^H F^H\left(\frac{P_k^L}{uP_k^H + P_k^L}\right)$$

and

$$P_k^L\left[1 - F^L\left(\frac{P_k^L}{uP_k^H + P_k^L}\right)\right] \leq uP_k^H\left[1 - F^H\left(\frac{P_k^L}{uP_k^H + P_k^L}\right)\right]$$

□

Notice that we formulated this principle at great generality. At this point, we have made no important assumptions about the sampling mechanism. On the other hand, if the sampling mechanism is rather strange, it may not be so useful to know that the individual did better than his average sampled predecessor. More on that below.

[11] In the proof we write P_k^H (resp. P_k^L) as shorthand notation for $P^H(s_k)$ (resp. $P^L(s_k)$).

3.3.2 Recursive Sampling

We now focus on a restricted class of sampling mechanisms built from a simple principle. Assume that the mechanism is such that each individual samples k of his predecessors' reports. k itself may be random, but it has support on some finite set $\{0, \ldots, J\}$. It is not possible to identify the order of the individuals sampled. With two actions, the individuals can identify at most K different samples (generally, with any finite number of actions there will only be finitely many distinguishable samples).

Recursive Sampling. Any given sampling mechanism induces a measure over the individuals in $(1, ..., n)$ when individual $n + 1$ samples: this is the (normalized to one) frequency with which the individuals are sampled. We say that the sampling mechanism is *recursive* when this induced measure over $(1, ..., n)$ consists of some weight π_n on n, and some weight $1 - \pi_n$ on the previous distribution over $(1, ..., n - 1)$.

Geometric Weighting. Many recursive sampling mechanisms can be constructed, but the following class of strikes us as the most natural one, due to its stationary nature. With geometric weighting, the relative frequency with which individual ν is sampled by n is $\rho^{n-\nu}$ where $\rho \geq 0$. If $\rho = 0$, we interpret this to mean that only the immediate predecessor is sampled—clearly, then the sample size must be one. If $\rho < 1$, there is relatively high weight on recent predecessors (what happened long ago is discounted). If $\rho = 1$ we have *proportional sampling,* all predecessors occur with even frequency. If $\rho > 1$, recent individuals are undersampled (only slowly are they added to the stock of public knowledge).

Corollary (Increasing Average Welfare) *Under Recursive Sampling, the expected welfare of the average sampled population never decreases.*

Proof: The above Principle of Improved Welfare says that the expected welfare of individual n is larger than the average sampled welfare among individuals $1, \ldots, n - 1$. The average sampled expected welfare of $1, \ldots, n$ is a weighted average between that of individual n and the average welfare among $1, \ldots, n - 1$. □

Beyond Recursive Sampling? So far we have treated certain sample mechanisms in detail. But what about those mechanisms which do not satisfy the recursive sampling requirement? Our proof of increasing average welfare does not immediately extend to such mechanisms. A key argument was that since each individual is better off than his average sampled predecessor, the average sample welfare keeps increasing. Now, assume that each individual observes the action of its two most immediate predecessors. The following sequence of welfares for the first four individuals is consistent with each individual beating his average sample: $(0, 3, 2, 2.8, \ldots)$. However, individual 4's average sampled welfare is 2.5 while that of individual 5 will be only 2.4.

Notice, that this was not a closed-form example respecting the special assumptions of our model, and that the example did not say anything about the long run. It was merely an example to point out how one particular logic of our proof could fail. While the above Corollary plays a crucial role in our analysis of whether full learning can occur, it is obviously not ruled out that the below learning results can extend to a far wider class of sampling mechanisms. In fact, we are able to derive results for far more sampling mechanisms in section 3.6 below, but then we have to restrict our model in other ways. Which sampling mechanisms admit full learning in the general model, and which don't, remains an open question for now.

3.4. LONG RUN LEARNING

The average sampled welfare keeps increasing, but is bounded above, so it must converge to some limit. The most important question to answer is then whether there is full learning or not: will the share of a_2-takers converge to 1 in state H and to 0 in state L? If so, sampling a_2 will become infinitely informative in the long run. To nail down the long-run distributions of actions in the two states of the world, we need to analyze the potential limits. We say that learning is *complete* iff $\lim_{n\to\infty} R_n^H = \lim_{n\to\infty} Q_n^H = 1$ and $\lim_{n\to\infty} R_n^L = \lim_{n\to\infty} Q_n^L = 0$. The result is

Proposition 3.1 (Convergence in Distribution) *Assume Recursive Sampling.*

1. If $\sum_{n=1}^{\infty} \pi_n = \infty$, if the private beliefs are unbounded, and if no samples are of size zero, then learning is complete.

2. If $\sum_{n=1}^{\infty} \pi_n = \infty$ and there are bounded private beliefs then learning is incomplete.

3. If $\sum_{n=1}^{\infty} \pi_n < \infty$, then learning is incomplete in any case.

4. If there exists $\varepsilon > 0$ such that for any one individual there is chance at least ε that the sample size is zero, then learning is incomplete.

REMARK 1. We argue in section 3.7 that the full learning result fails with misperceptions.

Proof: As noted, the sequence $U_n \equiv uR_n^H - R_n^L$ converges monotonically. From the initial conditions it starts strictly above -1 so the limit is in $(-1, u]$. With individual number n the expected increase is $U_{n+1} - U_n = \pi_n[uQ_n^H - Q_n^L - uR_n^H + R_n^L]$.

CASE 1. Here we will prove that $U_n \to u$, and thus $R_n^H \to 1$, $R_n^L \to 0$. If $\sum_{n=1}^{\infty} \pi_n = +\infty$, it must be true that $\liminf_{n\to\infty}[uQ_n^H - Q_n^L - uR_n^H + R_n^L] = 0$. We shall need to look at this difference in detail. Recall how we rewrote it in (4). We observed then that the difference is positive since it is the sum of a number of positive terms. In order for this sum to be very close to zero, it is necessary that all terms are very close to zero. Applying Lemma 3.3 we see that this is only possible if for each sample k it has very small probability of occurring in at least one of the two states of the world (otherwise the sample is very informative and improves welfare too much). Now, we had (2) which said that $R^\theta = \sum_{k=1}^{K} \rho_k P_k^\theta$ where ρ_k is the chance of resampling a_2 when drawing a random individual from s_k. We claim that non-pure samples must occur with very low probability in each state of the world, if n is also large. The reason is simple. If there were a good probability of sampling a mixed sample s_k of size at most $K-1$, then there would be a good probability of sampling any of the samples $s_1, \ldots, s_K$ in that state, for one could simply draw s_k a large number of times and then draw randomly from this, recreating any of the other samples with good probability. So, all of the mixed samples must have very low probability. When sampling, the probabilities of sampling $s_1, \ldots, s_K$ have to add up to one, so some pure samples must have large probability. The only way this can be reconciled, is if R^H and R^L are both very close to one or zero. Also, from the fact that all the inequalities of Lemma 3.3 need to collapse, it is necessary that R^H is approximately $1 - R^L$. Since the welfare $uR^H - R^L$ was always increasing, it is impossible to get very close to $R^H = 0 = 1 - R^L$, so we must be close to $R^H = 1 = 1 - R^L$. Finally, since now $\limsup_{n\to\infty} R_n^H - R_n^L = 1$, and since $uR_n^H - R_n^L$ is increasing, we can conclude with

Lyapounov that R_n^H and R_n^L must converge.

CASE 2. With bounded beliefs, the argument is basically the same. It is important to notice that the whole system grinds to a halt as soon as all available samples are sufficiently informative to overturn all private beliefs. That is the reason why none of the limits could be the full learning one.

CASE 3. When $\sum_{n=1}^{\infty} \pi_n < \infty$, we can choose a number N so large that $\sum_{n=N}^{\infty} \pi_n < 1$. It follows that for each successor of N, there is probability at least $1 - \sum_{n=N}^{\infty} \pi_n > 0$ of drawing the entire sample among individuals $1, \ldots, N$. Since at most N private signals have been aggregated to this point, any such sample is of bounded informativeness. Since any $n > N$ receives samples of bounded informativeness with positive probability, there is not complete learning.

CASE 4. In the long run, at least a fraction ε of the individuals observe nothing, and they must rely entirely on their private signal. Therefore, there is in the long run a bounded away from zero fraction of decliners in the pool, and a bounded away from zero fraction of investors, too. With the upper bound on sample sizes, there will eventually be a minimum positive chance of observing all investors, and a minimum positive chance of sampling all decliners. In conclusion, there is a bound to how informative samples can be. □

For the geometrically weighted mechanisms, $\sum_1^{\infty} \pi_n < \infty$ iff $\rho > 1$. Thus, proportional sampling, and mechanisms that put higher weight on the recent past permit full learning, while geometric weighting in favor of the remote past rules out complete learning. When $\sum_1^{\infty} \pi_n < \infty$, the past gets over-sampled in the long run, barring the complete aggregation of available information. The problem is that the first few individuals are observed with non-vanishing probability in the long run. The below Corollary in Section 3.6 can be seen as a rather precise converse. Generalizing beyond recursive sampling mechanisms, it shows that there is full learning when the weight given to any given individual converges to zero in the long run.

Once $\sum_{n=1}^{\infty} \pi_n = \infty$, the distinction that full learning occurs iff there are unbounded private beliefs is the same as in SS. The ergodic intuition is similar. Some interior distribution over actions can be stable iff there are bounded beliefs.

Our result stands in contrast to results of Banerjee and Fudenberg (1995). In their sections 3 and 4 they operate with proportional sampling. They find that with sample sizes larger than 1, no interior distribution can persist (their Lemma 2). Therefore, they get full learning even with bounded beliefs (their Theorem 1). They do so by showing that mixed populations cannot persist in the long run, as they are not stationary states. We consider their analysis faulty for the following reason. They posit a continuum agent model for analytical convenience. This essentially means that they take the average across stochastic outcomes of our model already as the model is progressing (small n). For instance, if we find with bounded beliefs that there is some probability $\pi^\theta \in (0,1)$ that (almost) all individuals are taking a_1, and some complementary probability $1-\pi^\theta$ that all individuals take a_2, they will find that there is a good probability of drawing mixed samples from the population (e.g. there is probability $2\pi^\theta(1-\pi^\theta)$ of drawing (a_1, a_2)). However that is an illusion; with probability very close to one only "pure" samples will be drawn. So, while we agree that only pure populations can survive in the long run (when samples are larger than one), we have proved that with bounded private beliefs the populations do with positive probability "go wrong" in the sense that the sub-optimal action is being taken by all.

Herding With Bounded Beliefs? One important question yet to be addressed is whether there is probability one that almost all individuals eventually take identical actions when the beliefs are bounded. Action a_2 is more likely to be sampled in state H than in state L. Still, we know that there is some significant remaining probability of sampling a_1 in state H. Is it the case that in each stochastic run through our model the population gets focused on one or the other action, or is it just that in state H all populations are essentially mixed but with a strong overweight of a_2?

To address this question we go back to work on random urns by Freedman (1965). Assume proportional sampling. When the sample size is always one, the long-run outcome is mixed populations. But this is only a knife-edge result, for as soon as the average sample size is larger than one, only separated populations can survive. The intuition is rather clear. As long as only one predecessor is seen, it can never be observed that the population is mixed, so all samples can be very informative (leading to imitation), thus replicating the mixed population.[12] However, as soon as larger samples can be observed, the mixed samples are going to be less informative, and they will allow the individuals to follow their private information, which in turn tends to purify the population (in state H, private signals in favor of a_2 occur more frequently than in state L).

General Sampling Mechanisms. When we relaxed the assumption of recursive sampling mechanisms, we saw before that our proof of Corollary fell apart. Therefore we have difficulties with the proof of Proposition 3.1, and we cannot prove complete learning under general sampling mechanisms. On the other hand, we can be quite general with the incomplete learning result resulting from bounded beliefs. The idea behind incomplete learning is that as soon as the beliefs based on samples get sufficiently focused, then individuals with bounded beliefs follow their sample blindly, not using their own information. Once this happens, no more information flows into the system. Since this effect kicks in before the beliefs are completely focused, it can happen in state H as well as in state L that almost all individuals take action a_1. So, there is incomplete learning.

Rigorously, we can offer the following. Assume that each individual draws a sample of at most size K, i.e. the sample size is uniformly bounded. The sample may be ordered or unordered, that is not important for the following result.

Proposition 3.2 (Bounded Beliefs) *If a sampling mechanism has uniformly bounded sample size, and private beliefs have bounded support, then learning cannot be complete.*

Proof: Consider for a moment the possibility that individual one would take action a_2 with probability one. Then no samples would ever become informative, and all individuals would take action a_2. Thus, in this special case, complete learning would be ruled out. We now continue under the assumption that individual 1 takes action a_1 with positive probability.

We proceed by contradiction. Assume that there is adequate learning, i.e. that p_n converges to one, where p_n is the probability with which individual n takes the correct action. Consider the sample s^* which has all individuals taking action a_1, the action which is optimal in state L. Since the sample size is bounded above by K, the probability of

[12]This is also why Banerjee and Fudenberg (1995) get this case right. They too find, that with sample size one mixed populations can survive in the long run.

observing s^* must converge to one in state L and to zero in state H. So, the belief based on observation of s^* must converge to zero. So, there is some number N so large that after N individuals, any successor who draws s^* will ignore the private signal and take action a_1. Now, in state H there is positive probability that the first N individuals all take action a_1 (thanks to our technical assumption). Whenever that occurs, a herd on action a_1 follows, for each successor will observe s^* and take action a_1. But that is a contradiction to the claim that the probability of observing s^* converges to zero in state H. □

3.5. ALMOST SURE CONVERGENCE

Stronger than convergence in distribution is almost sure convergence, to which we now turn.

Convergence of Aggregates? We have just established a weak form of convergence, and it is natural to ask how much stronger a result is possible. For instance, can the convergence be *uniform*? The answer to this question will depend on the exact form of the sampling mechanism. But for proportional sampling, we offer the following negative result: *Even with unbounded beliefs and action observation, there is a positive chance that an infinite subsequence of individuals chooses the contrary action.* It is easy to see why. For with random sampling of one (or more) predecessors, each individual will almost surely be sampled by infinitely many successors,[13] and the early influential individuals have a positive chance of doing anything. Since history become arbitrarily informative, anyone sampling such an individual will eventually choose to follow their lead and choose the same suboptimal action. So, even though almost all individuals eventually take the right action, there is also a large probability (perhaps 1), that a subsequence of individuals take the suboptimal action.

Convergence of Proportions. One can think of this section as the probabilistic counterpart to the preceding section, which concerned convergence in mean of the proportions. If one focuses on the report proportions, the model now falls prey to techniques

[13]Fix N. For each $n > N$, individual n samples N with probability $1/n$. Since $\sum_{n=N+1}^{\infty} 1/n = \infty$, it follows from the Borel-Cantelli lemma, that individual N will be sampled infinitely often.

developed in Arthur, Ermoliev, and Kaniovski (1986) (henceforth AEK). They investigate *generalized urn schemes of the Polya kind,* studying the evolution of pools consisting of balls of one of a finite number of colors. In our context, a ball will correspond to a decision-making individual, and the color of the ball is the chosen action: 'invest' and 'decline'. AEK provide general conditions under which one can describe the limit distribution of the pool. This corresponds to our desire to describe which actions are chosen in the long run.

In fact, AEK do have a result (Theorem 3.1) that might seem to prove a.s. convergence for our model, but we are reluctant to apply their result, since there is one condition we cannot readily check. The application is fine when there is positive probability of sample sizes 2 or larger. However, there is a difficulty when the sample size is one. Let us hint at the difference this way: figure 3-2 below graphs the expected next ball as a function of the current pool. With a sample size of one, and given there is herding, individuals will with large probability follow their sampled individual, so in the long run this curve lies along the diagonal. However, with larger samples, when the population is unfocused there is a good chance to observe mixed samples, and therefore there is a tendency to move far away from status quo (individuals apply their own information), so the curve will only intersect the diagonal near 0 and 1, the only possible long-run outcomes.

So, the AEK theorem proves that there is a.s. convergence of the action proportion in our model when the samples are larger than one. We now turn to focus on the *fixed sample size one* case. As we cannot apply the AEK theorem, we will instead rewrite our model the way they suggest and modify their proof to cover the current model. Even the rewriting itself sheds light on the dynamics in our context.

Fix the state $\theta = H$. Thus far we have focused on $R_n^H = P_n^H(a_2)$, the expected chance of sampling an investor (i.e. someone who took action a_2). We now turn to the *realized* chances of sampling an investor. Let x_n be the realized chance of sampling one investor out of the first $n-1$. Let's write the dynamics as described in terms of x. If $x \in [0,1]$ is the proportion of investors in the pool, then individual n himself becomes an investor with chance

$$\xi_n^H(x) = x\left[1 - F^H\left(\frac{R_n^L}{uR_n^H + R_n^L}\right)\right] + (1-x)\left[1 - F^H\left(\frac{1-R_n^L}{u(1-R_n^H)+(1-R_n^L)}\right)\right]$$

Figure 3-2: **Expected Motion.** This graph shows the function $\xi_n^H(x)$ for fixed n. As n goes to infinity, ξ_n^H converges towards the diagonal.

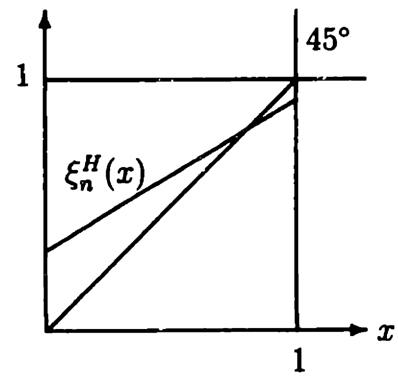

in state H, refer equation (2). Let β_n be the following stochastic indicator variable: $\beta_n = 1$ if individual n invests (chance $\xi_n^H(x_n)$), and otherwise $\beta_n = 0$. Then the evolution of x_n under recursive sampling is described by $x_{n+1} = (1 - \pi_n)x_n + \pi_n\beta_n$. If we define $\eta_n(x_n) = \beta_n - \xi_n^H(x_n)$, this becomes

$$x_{n+1} - x_n = \pi_n \left[\xi_n^H(x_n) - x_n + \eta_n(x_n)\right] \tag{5}$$

Since $E[\eta_n(x_n)|x_n] = 0$, the drift term is $\pi_n(\xi_n^H(x) - x)$. In fact, the equation for the evolution of R_n^H

$$R_{n+1}^H - R_n^H = \pi_n(\xi_n^H(R_n^H) - R_n^H)$$

is essentially a deterministic analogue of (5). As such, the convergence of R_n^H helps to prove the almost surely convergence of the random variable x_n. Note that $\xi_n(x)$ is linear in x, and that $\xi_n(x) - x$ vanishes as $n \to \infty$. In fact, $x \mapsto \xi_n(x)$ is a contraction for all finite n, but only weakly so in the limit; therefore we must modify the AEK approach. Figure 3-2 graphs the expected motion.

Most importantly, AEK point out a simple fact: Because $E[\eta_n(x_n)|x_n] = 0$, the cumulative perturbations variable $\mu_n = \sum_{k=1}^n \pi_k \eta_k(x_k)$ is a martingale with respect to the σ-algebra generated by $(x_1, \ldots, x_n)$. And $\eta_n(x_n)$ and $\eta_m(x_m)$ are uncorrelated for $m \neq n$ and for any given x_n and x_m; therefore, the variance of μ_n is $\sum_{k=1}^n \pi_k^2 [\mathrm{Var}(\eta_k(x_k))] \leq \sum_{k=1}^\infty \pi_k^2$. If we assume that $\sum_{k=1}^\infty \pi_k^2 < \infty$, then the Martingale Convergence Theorem for variables of bounded variance[14] proves that μ_n almost surely converges to a random limit μ_∞. More importantly, the tail perturbations $\sum_{k=n}^{\infty} \pi_k \eta_k(x_k)$ converge to zero. This essentially means that the drift of the system will determine its evolution. This yields

Proposition 3.3 (Action Convergence) *Assume sample size one, and recursive sampling with $\sum_n^\infty \pi_n = \infty$ and $\sum_n^\infty \pi_n^2 < \infty$. Then $R_n^H - x_n$ converges a.s. If the private belief are unbounded then $x_n \to 1$ a.s.*

Proof: We prove that whenever the accumulated perturbations converge (which happens almost surely), $R_n^H - x_n$ must converge as well. Fix a realization with convergent accumulated perturbations, and study the sequences of real numbers, (x_n) and (η_n).

Simple inspection of the above equations yields the inequalities $|x_{n+1} - x_n| \leq \pi_n$, $|R_{n+1} - R_n| \leq \pi_n$, $|\xi(x_n) - \xi(R_n^H)| \leq |x_n - R_n^H|$, $(x_{n+1} - R_{n+1}^H)^2 \leq (x_n - R_n)^2 + \pi_n^2 + 4\pi_n\eta_n$.

Since $\sum_{n=1}^\infty \pi_n < \infty$, we have $\pi_n \to 0$ as $n \to \infty$. Since the step sizes of x_n and R_n^H thus converge to zero, we conclude that one of these three alternatives is true: (i) there exists N such that $x_n > R_n^H$ for all $n > N$, or (ii) there exists N such that $x_n < R_n^H$ for all $n > N$, or (iii) there exists an infinite subsequence n_k such that $R_{n_k}^H - x_{n_k} \to 0$ as $n_k \to \infty$.

Consider alternative (i). After N we have that $|\xi(x_n) - \xi(R_n^H)| \leq |x_n - R_n^H|$ can be written as $\xi(x_n) - \xi(R_n^H) \leq x_n - R_n^H$ since ξ is increasing. Thus, $x_{n+1} - R_{n+1}^H \leq x_n - R_n^H + \pi_n\eta_n$. Let $\bar{\omega} = \liminf_n x_n - R_n^H$, we will prove that the sequence $x_n - R_n^H$ converges to $\bar{\omega}$. Let any $\varepsilon > 0$ be given. Choose $N' > N$ such that for all $n, m > N'$ we have $|\mu_n - \mu_m| < \varepsilon/2$ (by Cauchy convergence). Then choose some $N'' > N'$ such that $x_{N''} - R_{N''}^H < \bar{\omega} + \varepsilon/2$. We can iterate the above inequality to give for all $k > 0$,

$$x_{N''+k} - R_{N''+k}^H \leq x_{N''} - R_{N''}^H + \sum_{n=N''}^{N''+k-1} \pi_n\eta_n = x_{N''} - R_{N''}^H + \mu_{N''} - \mu_{N''+k} \leq \bar{\omega} + \varepsilon$$

In the case of alternative (ii), a similar proof yields that $x_n - R_n^H$ converges to its lim sup.

Next, consider alternative (iii). We will prove that $x_n - R_n^H$ converges to zero. Let any $\varepsilon > 0$ be given. Choose $N' > N$ such that for all $n, m > N'$ we have $|\mu_n - \mu_m| < \varepsilon/3$ and such that $\sum_{k=N'}^\infty \pi_k^2 < \varepsilon/3$. Then pick some $N'' > N'$ such that $(x_{N''} - R_{N''}^H)^2 < \varepsilon/3$. We

[14] See for instance Theorem 5.14 in Breiman (1968).

iterate one of the above inequalities to give for all $k > 0$,

$$(x_{N''+k} - R^H_{N''+k})^2 \leq (x_{N''} - R^H_{N''})^2 + \sum_{n=N''}^{N''+k-1} \pi_n^2 + \sum_{n=N''}^{N''+k-1} \pi_n \eta_n \leq \varepsilon$$

Finally, if the beliefs are unbounded, we know from Proposition 3.1 that $R^H_n \to 1$, so x_n must converge a.s. Since R^H_n is the state H mean of x_n, and since $x_n \leq 1$, it also follows that the limit must be 1 a.s. □

3.6. SCOPE

3.6.1 The Basic Postulate of Rational Social Learning

Building on the fact that observation of one predecessor is enough to generate full learning with unbounded private beliefs, we could hope to make inter-model comparisons. In particular, it sounds intuitively right that if all individuals sample more information than one predecessor, then there must also be full learning. The idea is to use the following induction: if all individuals until n had higher welfare with the bigger samples, then individual n can take his bigger sample, throw away all but one of the observations, and get an observation of one which is better off than in the samplesize-one model. However, there are some difficulties with this argument. Consider it in more detail. The reason why n is better off with the one individual left in his large sample must be that his posterior beliefs are shifted by a mean preserving spread (MPS). Hopefully, this would follow directly from the fact that his predecessors posterior belief was shifted by a MPS. So, the key principle driving the proof would have to be this:

Conjecture (Basic Postulate of Social Learning) *Individuals prefer to sample the action of more informed individuals, in the sense of Blackwell.*

However, we can give a counterexample to this conjecture. Assume two actions, and assume that $\bar{r} = 2/3$ is the threshold posterior belief between them. Assume that the observee first has this distribution of his posterior: q takes the values [.3 .7] with probabilities [.5 .5]. Then the action reveals the belief, and the observer will get the same distribution of belief based on the action observation. Now change the distribution by a MPS to [.3 .6 .8] with probability [.5 .25 .25] Then the observer's belief is distributed [.4 .8] with probability [.75 .25]. But this was not a MPS.

The Vives Effect. As pointed out, a major complicating factor in comparing models with different sampling mechanisms is that observing a larger sample need not improve welfare. It is simply not true that the ex ante expected utility of individual n must be at least as large in a model with a sample size of $k+1$ as in the k-sample size model. Computer simulations reveal that with $F^H(p) = (5p-2)/(2p)$ and $F^L(p) = (5p-2)(p+2)/(8p^2)$, individuals following number 84 expect to be better off if everyone samples 49 instead of 50 predecessors.[15] The reason why it can be less informative to draw a sample of 50 actions than a sample of 49 actions, when each individual has been better off in the 50-model so far, must be that the 50-model actions have a higher degree of correlation. This is also what we would suspect: In the 50-model, individuals have used their own signal to a lesser extent, and therefore they have taken more correlated actions.[16]

The Vives Effect explains the phenomenon: If individual n has a more informative sample in the $k+1$ sample model, then his posterior is more reflective of his sample and less of his private information. This effect potentially undermines the quality of the pool whence individual $n+1$ draws his signal. Notice that this implies a self-correcting property of the informativeness of the samples. If the individuals following 84 onwards receive a more informative sample in the 49-model, then they will reveal less of their private information. So, the individuals in the 50-model are not likely to fall far behind, and it is a possibility that later individuals will again prefer the 50-model over the 49-model.

3.6.2 Symmetric Two Action Case

We will now concentrate on a very special case in which the above basic postulate is true. Several nice results can be proved through the postulate. While we showed above that the postulate fails under a small modification of this model, the obtained results are still valid conjectures for our general model. We have no counterexamples to any of

[15]The initial conditions are such that all predecessors' (unordered) actions are observed up until there are more than 49 (resp. 50) predecessors in the pool. Note, that the computer simulations only ran to about 200 individuals, and the ordering of individuals' welfare may be reversed again for later individuals.

[16]If the reader so desires, the Vives effect may be labeled as a congestion effect.

Propositions 3.4, 3.5 or the Full Learning Corollary.

SYMMETRIC TWO ACTION SETUP. Assume that there are two available actions, and that the threshold belief between them is 1/2. As before, the prior belief is $q_0 = 1/2$ (this is now a crucial assumption). Also, the signal distribution is of the symmetric kind mentioned in the model setup. So, the private posterior distribution satisfies $F^H(p) = 1 - F^L(1-p)$ according to Lemma 3.1. To underscore the symmetry, assume that an individual with the threshold posterior belief 1/2 mixes over the two actions with equal weight.

It is easy to see that this model evolves in a symmetric fashion. Most importantly, the individual's posterior belief distribution in H is the mirror image of the distribution in L. Recall that $P_n^\theta(s_k)$ is the chance with which sample s_k is drawn in state θ.

Lemma 3.4 (Symmetric Evolution) *Assume the symmetric two action setup. For each individual n and each possible sample s_k we have $P_n^H(s_k) = 1 - P_n^L(s_k)$.*

So, if individual n draws sample s_k, he forms the belief $P_n^H(s_k)$. That happens with probability $P_n^H(s_k)$ in state H and $1 - P_n^H(s_k)$ in state L. We can now state

Lemma 3.5 (Basic Postulate of Social Learning) *Consider the symmetric two action setup. Individuals prefer to sample the action of more informed individuals, in the sense of Blackwell.*

Proof: The mean of an individual's posterior is necessarily equal to his prior, 1/2. The symmetry implies that there is probability 1/2 that the belief is at or below 1/2, leading to action a_1. A MPS is known to improve the individual's ex ante expected welfare. Since there is ex ante probability 1/2 of each action, the improvement in welfare must mean that the probability of taking action a_1 in state L is increased, just as the probability of a_2 in state H is increased. Now, sampling this individual's action is more informative: if action a_1 is observed there is a stronger belief in state L, if action a_2 is observed there is a stronger belief in state H. Thus, the (observation-based) social belief distribution is changed by a MPS. Since the private signal is independent of the history, the ultimate posterior is also changed by a MPS (Theorem 12 in Blackwell (1951)). □

We shall now consider a broad class of sampling mechanisms. The principle is to compare most mechanisms to one where full learning takes place, but very slowly. So, we start by defining our basic mechanism.

SAMPLING THE DISTANT PAST. We assume here that individual n observes the action of *one uniformly drawn* predecessor among the $\kappa(n)$ first individuals. The function $\kappa : \mathbf{N} \mapsto \mathbf{N}$ is assumed to have the property $\kappa(n) \leq n$. Further, we shall restrict attention to weakly increasing functions κ. Obviously, for complete learning to occur, it is necessary that $\kappa(n) \to \infty$ when $n \to \infty$. Otherwise the samples will always be drawn among the same finite number of individuals, and a finite number of private signals and derived statistics could never suffice to reveal the true state of the world.

It could be expected that it were also necessary for κ to converge towards infinity at a sufficiently fast rate, in order that late individuals take sufficiently informed actions. But it turns out that this is irrelevant for this full learning result. To compare with the incomplete learning result we reached in Proposition 3.1 for mechanisms oversampling the past, notice one important difference. In the mechanisms we considered back then, the problem was that early individuals were forever sampled with non-vanishing probability. However, in the current model a given individual is sampled with chance at most $1/\kappa(n)$, and that converges to zero as n tends to infinity.

Proposition 3.4 (Distant Past) *Assume Symmetry. Assume the sampling mechanism is to sample the distant past, that κ is weakly increasing, and that it satisfies $\kappa(n) \to \infty$ when $n \to \infty$. $P_n^H(s_2)$ converges to 1 with unbounded beliefs, and to a limit less than one with bounded beliefs.*

Proof: We proceed basically as in the proof of proposition 3.1. One crucial argument is that the sequence of $P_n^H(s_2)$ is weakly increasing, converging to a limit P. This follows form the above basic postulate and the fact that κ is weakly increasing. A simple induction then proves that later individuals are better off. Since $\kappa(n)$ converges to infinity with n, we observe that the sequence $P_{\kappa(n)}^H(s_2)$ converges to the same limit P. Just as before, we can then conclude that the limiting P must necessarily be such that individuals are herding in the limit. With unbounded beliefs it must therefore be $P = 1$, while with bounded beliefs the $P_n^H(s_2)$ sequence cannot possibly reach to 1. □

INFORMATIVE SAMPLING MECHANISMS. We shall say that a given sampling mechanism Σ is in the class of informative sampling mechanisms if the following is true. There exists a weakly increasing function $\kappa : \mathbf{N} \mapsto \mathbf{N}$ with $\lim_{n\to\infty} \kappa(n) = \infty$ such that if individual n draws samples according to Σ and picks one individual in his sample according to the uniform distribution, then the resulting distribution over individuals first order dominates that of uniformly sampling one of the first $\kappa(n)$ individuals.

Informative sampling mechanisms greatly generalize the recursive sampling shemes. Let us illustrate the definition with some examples. If Σ is to observe all predecessors, then it can be compared to sampling just one random predecessor. If Σ is to observe the j most immediate predecessors, then this can be compared to observing just one random predecessor; since the weight is concentrated on the j most recent individuals it first order dominates the mechanism with weight on all predecessors. If Σ is to draw a sample of random size $j \geq 1$ from among the $\kappa(n)$ predecessors, then it can be compared to sampling just one predecessor among the $\kappa(n)$. If Σ is to draw from some weighted distribution over the j most immediate predecessors, then it can be compared to sampling uniformly among the $n - j$ first predecessors.

The following proposition builds on the basic postulate and shows how we can compare welfare under informative mechanisms to welfare when sampling the distant past.

Proposition 3.5 (Large Samples) *Assume Symmetry. Unconditional on the true state of the world, the ex ante expected welfare of individual n is weakly larger under the informative mechanism Σ than under any mechanism of sampling the distant past to which Σ can be compared.*

Proof: The proof proceeds by induction. Unconditional on the true state of the world ($\theta = H, L$), denote by $G_n(r)$ the distribution of individual n's ultimate posterior belief when everyone draws a sample according to Σ, and by $H_n(r)$ the corresponding distribution when all draw from a dominated mechanism of sampling the distant past. The induction claim number n is that for all $j = 1, \ldots, n$ the distribution G_j is weakly more informative than H_j.

For $n = 1, 2$, the observation structure must be identical in the two models, and the posteriors for the two first individuals are distributed exactly the same. Assume the induction hypothesis obtains up to individual $(n-1)$, for some $n > 2$.

Consider the effect on individual n of moving from the mechanism Σ world to the dominated world. We proceed in three stages. First, he is clearly worse off if just he discards all but one randomly drawn observation, since more information always helps. Second, incorporate the effect of all predecessors living in the dominated world also. In this case, the distribution G_ν of the posterior belief of the individual ν left in his sample is more informative than H_ν in the dominated model, by induction assumption. Third, in the dominated model of sampling the distant past, later individuals are better off, so sampling uniformly from among the $\kappa(n)$ first individuals in the world would give an individual who's even worse off. By the above basic postulate, our considered individual n is better off in an ex ante sense under the Σ mechanism, for he could afford to throw away information and still make a more informative observation. But since this is true *for all convex value functions* of his posterior belief, the opposite direction of Blackwell's result implies that the posterior of n himself in the Σ world is a MPS (and thus more valuable) than his posterior in the dominated world. This proves the induction hypothesis. □

Corollary (Full Learning) *Assume Symmetry, and unbounded private beliefs. Under any informative sampling mechanism Σ there is full learning.*

Proof: Immediate from Propositions 3.4 and 3.5. □

3.7. IMPERFECT PERCEPTIONS

We now wish to relax a key assumption underlying the herding paradigm — namely, that it is an individual's *action* that is the informative statistic observed by successors who sample him. Actions are a coarse informative of the posteriors belief as a_k is taken iff the posterior belief of the individual is in $[\bar{r}_{k-1}, \bar{r}_k)$. Action observation only serves to pool nearby posterior beliefs, and does not otherwise add 'noise': There is simply no chance that a high posterior results in a low action.

More generally, imagine that there are finitely many ordered *posterior signals* or *reports* $\mathcal{S} = \{s^1, s^2, \ldots, s^M\}$, with higher reports more informative of higher posteriors, satisfying the *monotone likelihood ratio principle* (MLRP). To be exact, there are M *imperfect perception functions* $\psi_m : [0,1] \to [0,1]$ such that $1 = \sum_m \psi_m(r)$ for all r. $\psi_m(r)$ is the conditional probability that signal s_m occurs when r is the posterior. The MLRP condition is that $\psi_{m_1}(r_1)\psi_{m_2}(r_2) \geq \psi_{m_1}(r_2)\psi_{m_2}(r_1)$ whenever $m_1 > m_2$ and $r_1 > r_2$. Action observation is the special case where ψ_m is the indicator function on $[\bar{r}_{m-1}, \bar{r}_m)$, in which case all inequalities are of the form $0 \geq 0$.

For a simple story of this general case, think of a door-to-door pollster sampling the opinion of several individuals. At the end of each interview, he must register the individual as belonging to one of M ordered categories, from 'strongly disagrees' to 'strongly agrees'. Clearly, in such a situation, there is some latitude for the type of misperception modelled here. Inasmuch as this is a stylized feature of verbal communication, we wish to interpret misperceptions as a reduced-form model of plain conversation.

Here is another story generating the same formal model. Assume that the individuals' vN-M preferences are not exactly the same. Rather, the preferences are drawn randomly from a continuum of types before the action is taken. We assume that individuals have an identical ordering of the actions, such that in state H $a_1 \prec \ldots \prec a_M$ (this for the purpose of the MLRP).[17] Yet, it is randomly determined for each individual which beliefs map into which actions. Then, given observation of action a_m it is not given which interval the posterior must have lied in, and therefore we get the same equation as with imperfect perceptions.

Once again we concentrate on the case where only two signals can be conveyed. An individual then samples randomly one of his predecessors, and he gets to make an observation of the simple form: if the posterior of the sampled individual was p, then there is chance $\psi(p)$ that the predecessor is perceived as signal s^I (an investor), while there is chance $1 - \psi(p)$ that the predecessor is perceived as s^D (a decliner). The function $\psi : [0,1] \to [0,1]$ is yet to be specified. In general, we assume that ψ is weakly increasing, so that if the sampled individual really had a higher posterior in favor of state H, then there is also a higher probability that he is perceived as an investor; that is the natural MLRP assumption to make in this context.

[17]If we relax the assumption that actions are ordered the same way, we are likely to see confounding outcomes similar to those introduced in Smith and Sørensen (1996a). However, as shall be displayed below, confounding outcomes arise even without the conflicting preferences.

Let $\phi^\theta(q)$ be the probability (in state θ) that someone who samples a predecessor who himself faced the social belief q of state H, will receive an encouraging report s^I. Thus, for $\theta = H, L$,

$$\phi^\theta(q) = \int_0^1 \psi\left(\frac{pq}{pq + (1-p)(1-q)}\right) dF^\theta(p) \tag{6}$$

We define P_n^θ to be the chance in state θ that individual $n+1$ draws report s^I. We start from $P_1^\theta = \phi^\theta(q_0)$.

Individual $n > 1$ does not observe the actual sampling and posterior history, but only the report of the one sampled individual. If he observes s, he forms the social belief $q_n(s)$ in state H, where $q_n(s^D) = (1-P_n^H)/(2-P_n^H-P_n^L)$ or $q_n(s^I) = P_n^H/(P_n^H+P_n^L)$. The social and private beliefs are the basis for the individual's posterior as in (1). The dynamics are described as follows. In state H, individual $n+1$ observes s^I with chance P_n^H, resulting in a chance $\phi^H(q_n(s^I))$ that the report from individual n will be s^I.

We have seen above that there is complete learning in the action observation case when the private signal space is unbounded. We would now expect that if ψ were a continuous function with $\psi(0) = 0$ and $\psi(1) = 1$, then the result would go through. Alas, this is wrong. With unbounded beliefs, the learning may actually stop short of complete learning, ending at something that looks mostly like confounded learning, as described in SS.

For the rest of this section, assume that all samples are of size one, drawn uniformly among all predecessors (proportional sampling). How can learning grind to a halt? It must be the case that P^H and P^L solve the two stationarity equations (compare with equation (4))

$$P^H\phi^H\left(\frac{P^H}{P^H+P^L}\right) + \left(1-P^H\right)\phi^H\left(\frac{1-P^H}{2-P^H-P^L}\right) = P^H$$
$$P^L\phi^L\left(\frac{P^H}{P^H+P^L}\right) + \left(1-P^L\right)\phi^L\left(\frac{1-P^H}{2-P^H-P^L}\right) = P^L$$

These stationarity equations are not only satisfied at $P^H = 1$ and $P^L = 0$, as was the case with action observation and unbounded beliefs, but also possibly at some values $P^H < 1$, $P^L > 0$. Moreover, the interior solution can be a stable solution, with the system converging to it regardless of where it starts (except at the other stationary point $P^H = 1$

and $P^L = 0$).

EXAMPLE. Take our standard unbounded symmetric private belief example, where $f^H(p) = 2p$ and $f^L(p) = 2p - p$, for $p \in (0,1)$. We will consider the class of piecewise linear misperception functions ψ that are defined through the parameter $a \in [0, 1/2]$ by

$$\psi(p) = \begin{cases} 0 & \text{for } p \leq a \\ (p-a)/(1-2a) & \text{for } a < p < 1-a \\ 1 & \text{for } p \geq 1-a \end{cases}$$

For $a = 1/2$ this becomes the indicator function that corresponds to action observation, so we know that there is complete learning. ψ is symmetric in the sense that $\psi(p) = 1 - \psi(1-p)$. Also, notice that ψ is increasing and continuous at 0 and 1, and that for $a > 0$ there are intervals of perfect information transmission, as ψ is flat near 0 and 1. Yet, for $a < 1/2$ the general problem arises that there is always some doubt, upon observation of an investor, whether it might have been a sample of a decliner that was misperceived. This, of course, is what can hinder the complete learning from taking place. Assume as always that there are fair priors, so the system starts at $P_1^H = P_1^L = 1/2$. Then it is an easy exercise to prove that the system evolves symmetrically: for all n, $P_n^H = 1 - P_n^L$. Therefore the dynamics are fully described in one variable P_n^H. The system of difference equations boils down to one difference equation, and for that equation it is easily proved that $P_{n+1}^H > P_n^H$ iff

$$P_n^H \phi^H(P_n^H) + \left(1 - P_n^H\right) \phi^H(1 - P_n^H) > P_n^H \tag{7}$$

Computer generated graphs for the above choice of ψ reveal that when $P^H = 0$, the left hand side is zero, expressing that this is a stationary point; for P^H above zero, the left hand side lies below 1, and thus the dynamics tend to bring P_n^H away from 0. Then, there is a crossing point $P^* > 0$, at which the left hand side equals 1, and the system is at rest. Above P^*, the left hand side is above 1, implying that P_n^H will move downwards. The conclusion to be drawn about such dynamics, is that P^* is the stable long run outcome of the model, to which the system will converge. So, there is less than complete learning.

The larger is a, the smaller is the confounded learning point P^*, and since we know there is complete learning for $a = 1/2$, it will have to vanish when a approaches $1/2$. Thus far we are unable to conclude if there was a region of parameter values close to $a = 1/2$ with complete learning or not.

The above example was constructed to investigate when confounded learning can permit full learning to go through. The robust result is this.

Proposition 3.6 *Assume the symmetric two action setup, and that the misperception function is symmetric, i.e.* $\psi(q) = 1 - \psi(1-q)$. *Assume that the misperception function* ψ *is* n *times differentiable at 1, and that one of the derivatives* $\psi^{(\tilde{n})}(1)$ *is not zero,* $\tilde{n} < n$. *Then there exists a confounded learning point near* 1, *and there will not be full learning.*

Proof: The aim is to prove that inequality (7) is violated in a neighborhood of $P^H = 1$. That follows from $\tilde{n}$ times differentiation of the function $h(P) \equiv P\phi^H(P) + (1-P)\phi^H(1-P) - P$ near $P = 1$. For simplicity, now, assume that $\tilde{n} = 1$, and differentiate once. We see that $h'(P) = \phi^H(P) + Pd\phi^H(P)/dP - \phi^H(1-P) - (1-P)d\phi^H(1-P)/dP - 1$ which evaluated at $P = 1$ boils down to $h'(1) = d\phi^H(1)/dP$. We need to prove that $h'(1) < 0$.

$$\frac{d\phi^H(P)}{dP} = \int_0^1 \psi'\left(\frac{pP}{pP+(1-p)(1-P)}\right)\frac{p(1-p)}{(pP+(1-p)(1-P))^2}\,dF^H(p)$$

so $d\phi^H(1)/dP = \psi'(1) < 0$. □

3.A. APPENDIX: MORE ACTIONS

We prove that it is rather straightforward to generalize the results for the sample size one perfect action observation model to more than two actions. Assume that there is any finite number M of actions. We rule out actions that are weakly dominated by the set of the other actions, so any action will be optimal for some beliefs p.

First, The Principle of Improved Welfare holds. Rather than recreating the expression (4) of its proof, we turn to the following direct proof. One possible strategy for individual n is to always imitate one randomly chosen strategy of his sample. That will give n the same expected welfare as his average sample. When employing the optimal strategy, n's welfare can only be better than that.

As the welfare of n is better than that of his average sample, Corollary is generalized without difficulties. Our main task in this appendix is to prove that Proposition 3.1 can be transferred. When $\sum_1^\infty \pi_n < \infty$, the same proof as before shows that there cannot be complete learning. Also, if the beliefs are bounded, the proof goes largely as before. So, we focus on the case where the private beliefs are unbounded, and $\sum_1^\infty \pi_n = \infty$. As before, with these assumptions, the gap between the welfare of individual n and his predecessors must close to 0 in the long run. This, again, implies that mixed samples are very unlikely after a long while. So, only pure samples are possible, and they need to be very informative, so it must be that some subset of actions occur with high probability in state H, while a complementary subset of actions occur in state L. As an implication, pure samples are extremely informative. Therefore, actions a_1 and a_M will be taken by almost all individuals. And, in conclusion, a_1 must occur almost always in state L, while a_M must occur almost always in state H.

The change from 2 to M actions does not alter the conclusions about almost sure convergence of proportions. When sample sizes larger than 2 can occur, the only long-run stable populations have to be pure. I.e., only one action is taken by almost all individuals. With sample size one, we face the same difficulties as before. The proof that we get non-mixed populations with unbounded beliefs can be constructed as a modification of the proof in the text. However, with bounded beliefs and sample size one, mixed populations can survive in the long run.

Bibliography

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. Spinger-Verlag, New York.

ARTHUR, W. B., and D. LANE (1994): "Information Contagion," in *Increasing Returns and Path Dependence in the Economy*, ed. by W. B. Arthur, pp. 69–97. The University of Michigan Press, Ann Arbor.

AVERY, C., and P. ZEMSKY (1992): "Herding in Financial Markets," PhD Thesis, Chris Avery, Stanford University.

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

——— (1993): "The Economics of Rumours," *Review of Economic Studies*, 60, 309–327.

BANERJEE, A. V., and D. FUDENBERG (1995): "Word-of-Mouth Learning," MIT Mimeo.

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.

BLACKWELL, D. (1951): "Comparison of Experiments," in *Proceedings, Second Berkeley Symposium on Mathematical Statistics and Probability*, pp. 93–102. University of California Press, Berkeley.

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.

FREEDMAN, D. A. (1965): "Bernard Friedman's Urn," *Annals of Mathematical Statistics*, 36, 956–970.

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.

HILDENBRAND, W. (1974): *Core and Equilibria of a Large Economy*. Princeton University Press, Princeton.

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

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

McLennan, A. (1984): "Price Dispersion and Incomplete Learning in the Long Run," *Journal of Economic Dynamics and Control*, 7, 331–347.

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

Pesendorfer, W., and J. Swinkels (1995): "The Loser's Curse and Information Aggregation in Common Value Auctions," Northwestern mimeo.

Rothschild, M. (1974): "A Two-Armed Bandit Theory of Market Pricing," *Journal of Economic Theory*, 9, 185–202.

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.

Shiller, R. J. (1995): "Conversation, Information, and Herd Behavior," Cowles Foundation Discussion Paper no. 1092.

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

Smith, L., and P. Sørensen (1996a): "Pathological Outcomes of Observational Learning," Chapter 1 of this thesis.

——— (1996b): "A Theory of Rational Social Learning," Chapter 3 of this thesis.

Stokey, N. L., and R. E. Lucas (1989): *Recursive Methods in Economic Dynamics.* Harvard University Press, Cambridge, Mass.

Vives, X. (1993): "How Fast do Rational Agents Learn?," *Review of Economic Studies*, 60, 329–347.

Whittaker, E. T., and G. N. Watson (1963): *A Course Of Modern Analysis.* Cambridge University Press, England, 4th edn.