

Yale University

EliScholar – A Digital Platform for Scholarly Publishing at Yale

Cowles Foundation Discussion Papers

Cowles Foundation

11-1-2009

Biased Social Learning

Helios Herrera

Johannes Hörner

Follow this and additional works at: <https://elischolar.library.yale.edu/cowles-discussion-paper-series>



Part of the [Economics Commons](#)

Recommended Citation

Herrera, Helios and Hörner, Johannes, "Biased Social Learning" (2009). *Cowles Foundation Discussion Papers*. 2062.

<https://elischolar.library.yale.edu/cowles-discussion-paper-series/2062>

This Discussion Paper is brought to you for free and open access by the Cowles Foundation at EliScholar – A Digital Platform for Scholarly Publishing at Yale. It has been accepted for inclusion in Cowles Foundation Discussion Papers by an authorized administrator of EliScholar – A Digital Platform for Scholarly Publishing at Yale. For more information, please contact elischolar@yale.edu.

BIASED SOCIAL LEARNING

By

Helios Herrera and Johannes Hörner

November 2009

COWLES FOUNDATION DISCUSSION PAPER NO. 1738



**COWLES FOUNDATION FOR RESEARCH IN ECONOMICS
YALE UNIVERSITY
Box 208281
New Haven, Connecticut 06520-8281**

<http://cowles.econ.yale.edu/>

BIASED SOCIAL LEARNING

HELIOS HERRERA AND JOHANNES HÖRNER

ABSTRACT. This paper examines social learning when only one of the two types of decisions is observable. Because agents arrive randomly over time, and only those who invest are observed, later agents face a more complicated inference problem than in the standard model, as the absence of investment might reflect either a choice not to invest, or a lack of arrivals. We show that, as in the standard model, learning is complete if and only if signals are unbounded. If signals are bounded, cascades may occur, and whether they are more or less likely than in the standard model depends on a property of the signal distribution. If the hazard ratio of the distributions increases in the signal, it is more likely that no one invests in the standard model than in this one, and welfare is higher. Conclusions are reversed if the hazard ratio is decreasing. The monotonicity of the hazard ratio is the condition that guarantees the presence or absence of informational cascades in the standard herding model.

Keywords: Informational herds, Cascades, Selection bias.

JEL codes: D82, D83

1. INTRODUCTION

The process of learning in social contexts confronts the same difficulties as any other statistical analysis. The data available to an individual may be subject to selection bias. Consider the leading example used by Bikchandani, Hirshleifer and Welch [2] (henceforth BHW), for instance. Upon learning that a paper has been previously rejected, a referee at a second journal tilts toward rejection. But what if, as is usually the case, he did not learn about this rejection? Surely, he would nevertheless wonder about the paper's journey onto his desk, and speculate about

Date: First Version: December 1th 2007, This Version: November 2, 2009.

Herrera: SIPA, Columbia University. Email: hhh2108@columbia.edu. Hörner: Yale University. Email: johannes.horner@yale.edu.

rejections the paper might have gone through. While publications are by definition observable, rejections are not.

To wit, there are far more significant instances of such bias in academia. How easy is it to publish a paper that finds inconclusive empirical evidence? In medical and social sciences, studies whose findings are statistically insignificant get filed away, biasing the published papers toward positive results.¹

The difficulty in interpreting the absence of negatives is encountered everywhere. Is no one waiting in this line because cabs come by all the time, or because this isn't actually a cab line? Do the low figures of tax evasion reflect the success of deterrent policies, or the success of tax evaders? Why do 90% of mutual funds truthfully claim to have performance in the first quartile of their peers? (The other three quarters of funds have closed. See Elton, Gruber and Blake [7].)

This paper develops a model of biased social learning and revisits the findings of the literature. In this model, individuals arrive randomly over time. As in Smith and Sørensen [13], each agent has some private, conditionally independent information about the relevance of taking some decision immediately upon arrival—say, making an investment. As often in the social learning literature, we assume that the payoff from investing depends on the state of the world, but not on what earlier or later individuals decide. Therefore, values are common, and externalities are purely informational.² As is standard as well, signal distributions satisfy the strict monotone likelihood ratio property (MLRP).

What sets this model apart from standard models is the following informational assumption. While the decision to invest (but not the payoff from investing) is observable to all future individuals, the failure to do so, and in fact, the mere arrival of individuals (who do not invest), remains hidden. Individuals arriving later will observe “positives” (if and when earlier individuals invested), but not “negatives” (if and when earlier individuals chose to abstain).

¹This phenomenon is known as the “file-drawer problem,” or the “publication bias.” See Scargle [11]. As a result, prominent medical journals no longer publish results of drug research sponsored by pharmaceutical companies unless that research was registered in a public database from the start. Some of them also encourage publication of study protocols in their journal.

²We shall also discuss at length a version in which there is only one investment opportunity, in which case there is an obvious payoff externality.

Therefore, every individual faces a complex problem of statistical inference: given the observed history, and the randomness in the arrival of individuals, how likely is it that some individuals had the opportunity to invest, but chose not to? And if so, what were their private signals? Note that, in this problem, time plays a key role, as it becomes increasingly more likely, as time passes by, that some individuals must have had the opportunity to invest.

In this context, we ask whether biased social learning exacerbates or mitigates herding. Could it be the case that some investment opportunities, or lucrative projects, remain unexploited because agents considering making it suspect that others must have thought of it, or even tried it before them? How many entrepreneurs, or scientists, stumbling across a new idea, chose not to follow through this idea because of the rational belief that they were unlikely to be the first to think of it?

Our first main result shows that, qualitatively, the absence of negatives does not alter the conditions under which cascades can, or cannot occur. If the informativeness of signals is bounded, wrong herds can occur (that is, they will occur for some prior and payoff parameters). On the other hand, if signals are unbounded, learning is necessarily complete; whether the state of the world is such that investment is profitable or not, agents will eventually learn it.

On the other hand, our second main result shows how, quantitatively, the absence of negatives affects the probability of a wrong herd. Consider the case of bounded signals (so that cascades may occur). What is the probability that no agent ever invests, while agents should, in the case of biased learning, relative to this probability in the benchmark model of BHW, in which all decisions, to invest or not, are observed? As it turns out, the comparison of these probabilities hinges upon a simple statistical property of the signal distribution. If signals satisfy the increasing hazard ratio property (IHRP), that is, if the ratio of the hazard rates increases in the signal, then the probability of no one ever investing is lower under biased learning, *independently* of the state of the world. Conversely, if the hazard ratio is decreasing, then this probability is lower in the benchmark model. While already used in the statistical literature, IHRP appears to be new in the literature on social learning. Yet as we show, it also plays a key role in the BHW model. Namely, IHRP is the necessary and sufficient condition under for the absence of informational cascades, namely provided the decision of the first individual depends on his signal, the decision of all later individuals will do as well. That is, it ensures that the posterior public

belief necessarily stays in the learning region provided that the prior lies in it.³ While biased learning always leads to higher investment (relative to the benchmark model) under IHRP, it nevertheless leads to lower welfare, at least in the version of our model in which there is only one investment opportunity.⁴

In practice, biased learning comes in at least two kinds. It can be, as we have assumed so far, that the absence of negatives indicates that no agent had an opportunity to invest, rather than a choice not to invest. Alternatively, it may be that only successful investments are observed, while failed investments, no investments, and no opportunity to invest are all observationally indistinguishable. In this second variant, if investments may also succeed, albeit at a lower rate, in the unfavorable state, the resulting inferential problem remains non-trivial. In a last section, we show that the qualitative findings (i.e., our first result) extend to this environment as well.

The first models of sequential decisions and observational learning by Banerjee [1] and Bickchandani, Hirshleifer and Welch [2], and their subsequent generalization by Smith and Sørensen [13] all assume that all actions are observed by later individuals. Namely, agents could observe the precise sequence of decisions made by all the predecessors. Later work, notably Çelen and Kariv [4], Callander and Hörner [3] and Smith and Sørensen [12] relaxes this assumption and considers the case in which either a subsample of the sequence, or a statistic thereof is observable. As these authors show, the asymptotic properties of social learning may radically change. For instance, Çelen and Kariv [4] show that, when agents only observe the action of their immediate predecessor, beliefs do not converge. Therefore, complete learning never occurs, as beliefs and actions end up cycling. Hence, limiting the information available to agents may alter the qualitative properties of learning in general, although it turns out not to do so in the case of biased learning. Callander and Hörner [3] show that if agents can only observe the fraction of agents having taken each action, rather than the entire sequence, then it might be optimal to take the action that was adopted by the minority of predecessors. A similar observational assumption is made in Hendricks, Sorensen and Wiseman [8].

Chari and Kehoe [6] develop an observational learning model with a similar investment bias, in the sense that more information is revealed after observing an investment than after observing

³More precisely, IHRP guarantees that this is the case after “good news,” that is, after an observed investment decision. There is a corresponding property for the case of “bad news.”

⁴This version allows us to focus on the history of no observed investment.

a decision not to investment. Each investment amount is an observable continuous variable, so the investor's private signal can be fully inferred (when this investment amount is positive). As in usual models, in case of a non-investment is observed then only a truncation on the investor's private signal can be inferred. In a sense, their model adds information in a biased fashion to the standard model (investment decisions become more informative), while we suppress information in a biased fashion, by assuming that decisions not to invest are not observable.

Our model is also related to models of endogenous timing such as the elegant paper of Chamley and Gale [5]. In their model as in ours, whether an agent has an opportunity to invest or not is a random variable. In their model, there is a finite number of agents who are all present from the start, and may choose to wait before investing, if they wish to. Inefficiently low investment occurs because agents might decide to wait, in the hope that others will act first, and thereby reveal valuable information. In our model, agents arrive at random times, and the total number of agents having the opportunity to invest is almost surely countably infinite, so that the collective information of the agents reveals the state. In both their model and ours, agents must be careful in interpreting an observed absence of investment. In Chamley and Gale [5], this absence might also reflect strategic delay, rather than bad news, while in ours, it might simply be because no agent happened to face this choice. As we show, in the version of our model in which the game stops after the first investment opportunity, agents have no incentive to engage in strategic delay.

Section 2 introduces the set-up. Section 3 develops the analysis and parametric examples. The qualitative results are stated in Section 4. Section 5 focuses on the quantitative results regarding the probability that no one ever invests. Section 6 discusses the variant in which only successful investments are observed. Most proofs are relegated to an appendix.

2. SET-UP

2.1. Information. Imagine a situation in which there are two states of the world. The state of the world is denoted $\theta \in \{0,1\}$. We refer to state 0 as the bad state, and to state 1 as the good state. The *ex ante* probability of the good state is denoted

$$p_0 := \mathbb{P}[\theta = 1] \in (0, 1).$$

There is a countable infinity of agents (or individuals, or players). Each agent receives a private signal (or type) $x \in X := [0, 1]$. Private signals are conditionally independent across

agents. The conditional distribution of this signal is identical across agents. Conditional on state θ , the distribution (c.d.f) is denoted F_θ , and assumed twice differentiable on $(0, 1)$, with density f_θ that is strictly positive on $(0, 1)$.

Private signals provide valuable information about the state to the agents. The distributions are assumed throughout to satisfy the strict monotone likelihood ratio property (MLRP henceforth), and this assumption will be implicit in all formal statements. That is, defining the likelihood ratio

$$l(x) := \frac{f_1(x)}{f_0(x)},$$

on the interval $(0, 1)$, we assume that l is strictly increasing. This guarantees that higher values of the signal lead to higher posterior probabilities that the state is good, for all priors. See Milgrom [9]. In some instances, without loss of generality, we will set $l(1/2)$ to 1, so that the signal $x = 1/2$ leaves any given prior probability belief unchanged.

2.2. Actions and Payoffs. Each agent i faces a binary choice. He may either invest or not. The decision to invest is denoted I , while the decision not to is denoted N . An action, therefore, is an element $a_i \in \{N, I\}$. Investing is costly: the action I entails a cost $c \in (0, 1)$. The return from investment is random, and depends on the state of the world. We normalize its expectation to 0 in the bad state, and to 1 in the good state. The payoff from not investing is set to 0, so that, under complete information, an agent would invest if and only if the state were good. To summarize, the payoff of an agent is given by

$$\begin{aligned} u(N) &:= 0, \\ u(I) &:= -c + \begin{cases} 0 & \text{if } \theta = 0, \\ 1 & \text{if } \theta = 1. \end{cases} \end{aligned}$$

Note that it is optimal for an agent to invest if and only if he assigns a probability of at least c to the good state. Throughout, we refer to this probability as the agent's belief.

2.3. Timing and Histories. We are now ready to describe the extensive-form game. Time is continuous and the horizon is infinite. There is Poisson arrival process defined over dates $t \in \mathbb{R}_+$, with associated random point process $\{T_i\}_{i \geq 0}$, with $T_0 := 0$, and $T_i \leq T_{i+1}$ for all $i \geq 0$. The intensity of the Poisson process, λ , is constant and independent of the state of the world.

The random variable T_i determines agent i 's *arrival time*. That is, agent i must take an action j at the date of the realization t_i of T_i . Because agents cannot delay their decision, any discounting is irrelevant and ignored.

Arrival times are not observed, and neither are decisions not to invest. Private signals are not observed either. Further, agents do not know their index, i.e. agent i does not know that, by definition of his index, $i - 1$ agents had the opportunity to invest before him.⁵ However, decisions to invest are observed (of course, the corresponding arrival time is then inferred). A (complete) history up to date t , then, specifies the state of the world, the infinite sequence of private signals of players, the date t and the sequence $\{(t_i, a_i)\}_i$, with $t_i \leq t$, for all i , of arrival times and actions taken by the corresponding agent. This sequence is (almost surely) finite. Agents, however, only observe a subset of these arrival times. The relevant history, then, is the public history $h_t := (t, \{(t_i, I)\}_i)$, which is the subset of the complete history that includes all times at which an agent decided to invest, as well as the current date t . Note that the public history does not include the identity i of the agents that decided to invest, so that it is not possible to infer from such a history how many agents actually had the opportunity to invest up to time t . Denote the set of public histories up to time t by H_t , and let $H := \cup_{t \geq 0} H_t$ denote the set of all histories. Set $H_0 := (0, \{\emptyset\})$. A history $h_t = (t, \{\emptyset\})$ indicates that no investment has taken place up to date t .

2.4. Strategies and Equilibrium. A strategy for agent i specifies, for each possible signal x , arrival time t_i and public history h_t , $t = t_i$, an action choice (possibly mixed). That is, a behavior strategy for agent i is a measurable mapping

$$\sigma_i : X \times H \rightarrow \Delta\{N, I\},$$

where $\Delta\{N, I\}$ denotes the set of lotteries over $\{N, I\}$. Since a player only knows his own type, this is a game of incomplete information, and one should therefore specify each player's assessment over the complete history, for each possible (x_i, h_t) , $t = t_i$. Given that agents cannot delay their decision, and given the payoff specification, it is clear that the only relevant probability that affects their choice is their belief p_i over the state of the world, given (x, h_t) . We build this directly in the definition of an equilibrium.

⁵Their belief about their rank is the improper uniform prior, so that they are *a priori* equally likely to be anywhere in the sequence.

A *Perfect Bayesian Equilibrium* consists of a strategy profile $\sigma := \{\sigma_i\}_i$ and a profile of beliefs $p_i : X \times H \rightarrow [0, 1]$, all i , such that (i) each player's strategy is a best-reply at every information set, and (ii) the beliefs p_i are consistent with Bayes' rule at every information set that is reached with positive probability, given σ .

Note that, given MLRP, the optimal strategy of an agent must be a pure strategy, and more precisely a cut-off strategy. That is, he should invest if and only if his signal is high enough. We focus on symmetric equilibria. This implies that an equilibrium will be uniquely determined by a measurable function x_t that specifies the cut-off type above which an agent i invests at time t . Along with the prior belief p_0 , this determines, in particular, the *public belief* $p_t := \mathbb{P}[\theta = 1|h_t]$ (henceforth, belief) about the state given the observed history h_t . Note also that p_t is a summary statistic for h_t . In the sequel, it we shall also use the likelihood ratio of p_t , denoted L_t , and defined by

$$L_t := \frac{p_t}{1 - p_t} = \frac{\mathbb{P}[\theta = 1|h_t]}{\mathbb{P}[\theta = 0|h_t]}.$$

Because L is strictly increasing in p , we shall sometimes, with an abuse of terminology, refer to this ratio as the public belief as well.

3. ANALYSIS

3.1. Threshold Signal and Public Belief. As mentioned, an equilibrium can be summarized by two functions of the public history h_t : the belief p_t that the state is good, and the cut-off x_t such that an agent invests at date t if and only if his signal exceeds x_t .

Bayes' rule provides one relationship between x_t and the belief p_t . Namely, p_t determines x_t since, given p_t , an agent of type x_t must be indifferent between investing or not, at least when x_t is in $(0, 1)$:⁶

$$x_t := x \text{ solves } \mathbb{E}[\theta|x, h_t] = \mathbb{P}[\theta = 1|x, h_t] = c.$$

Using Bayes' rule, this means that the threshold x_t solves

$$\frac{f_1(x) p_t}{f_1(x) p_t + f_0(x) (1 - p_t)} = c.$$

⁶If, given p_t , it is optimal to invest independently of the signal, set x_t equal to 0. Similarly, set it to 1 if it is always optimal not to invest.

Using the likelihood ratio, the relationship between the threshold x and the public belief L takes the simple product form, for any given c ,

$$(1) \quad l(x) L = \frac{c}{1-c},$$

which highlights the inverse relationship between the public belief and the signal cut-off: if agents are more optimistic, that is, if the public belief L is higher, a lower threshold signal x is required for an investment opportunity to be deemed profitable.

There is a second relationship between x_t and p_t . Namely, the threshold x_t determines the evolution of the belief p_t , along with the initial value p_0 . The evolution of p_t over some small interval of time $(t, t + dt]$, $dt > 0$, depends on whether an investment occurs or not during that interval. Accordingly, we divide the analysis into two cases.

3.2. Evolution without Investment. Consider first the case in which there has been no investment up to time t . Note that, in this case, p_t is a continuous function, because the distribution over arrival times is continuous. It follows that x_t is continuous as well. Let

$$G_{\theta,t} := \mathbb{P}[h_t = (t, \{\emptyset\}) | \theta]$$

be the probability of this event, conditional on state θ . Since investments arrive at rate $\lambda F_\theta(x_t)$, this probability is given by

$$G_{\theta,t} = e^{-\lambda \int_0^t (1 - F_\theta(x_s)) ds}.$$

To see this, note that, from time t to time $t + dt$, this probability evolves as follows (neglecting terms of order dt^2 or higher):

$$G_{\theta,t+dt} = G_{\theta,t} \cdot (F_\theta(x_t) \cdot \lambda dt + 1 \cdot (1 - \lambda dt)).$$

Indeed, the probability that no one invests up to time $t + dt$ is the probability that no one invests up to time t , multiplied by the probability that no one invests in the time interval $(t, t + dt]$. This latter probability is the sum of two terms. With probability $1 - \lambda dt$, no agent arrives during this time interval. With probability $F_\theta(x_t) \cdot \lambda dt$, some agent arrives in the time interval, but his signal is below the threshold. Because x_t is continuous, it follows that $G_{\theta,t}$ is differentiable and solves

$$\frac{G'_{\theta,t}}{G_{\theta,t}} = -\lambda (1 - F_\theta(x_t)),$$

along with $G_{\theta,0} = 1$. This integrates out to the formula above.

After a no-investment history $h_t = (t, \{\emptyset\})$, the public belief evolves according to

$$p_t = \mathbb{P}[\theta = 1 | h_t] = \frac{\mathbb{P}[h_t | \theta = 1] \cdot \mathbb{P}[\theta = 1]}{\mathbb{P}[h_t]} = \frac{G_{1,t} p_0}{\mathbb{P}[h_t]}.$$

Since $G_{1,t}$ is a function of $(x_s)_{s \leq t}$, this provides a second relationship between the belief p_t and the threshold x_t .

We now combine the two relationships. As pointed out, the threshold x_t solves

$$\frac{f_1(x) p_t}{f_1(x) p_t + f_0(x) (1 - p_t)} = \frac{f_1(x) G_{1,t} p_0}{f_1(x) G_{1,t} p_0 + f_0(x) G_{0,t} (1 - p_0)} = c.$$

Using the formula for $G_{\theta,t}$, and setting

$$\gamma := \frac{1 - p_0}{p_0} \frac{c}{1 - c},$$

it follows that x_t solves

$$(2) \quad l(x) = \gamma \frac{G_{0,t}}{G_{1,t}} = \gamma e^{\lambda \int_0^t (F_0(x_s) - F_1(x_s)) ds}.$$

Since the right-hand side is differentiable in t , and l is differentiable, the function x_t must be differentiable as well. By the implicit function theorem, the function x_t solves

$$\frac{l'(x_t) x_t'}{l(x_t)} = \lambda (F_0(x_t) - F_1(x_t)),$$

with initial condition $x_0 = l^{-1}(\gamma)$. Integrating, we obtain the following implicit characterization of the threshold x_t :

$$t = g(x_t) := \frac{1}{\lambda} \int_{x_0}^{x_t} \frac{l'(x)}{l(x)} (F_0(x) - F_1(x))^{-1} dx.$$

While the right-hand side admits no closed-form solution in general, we shall provide a few examples in which we can solve for x_t explicitly. Note that this also gives us a characterization of p_t , or equivalently, L_t , since

$$L_t = \frac{c}{1 - c} \frac{1}{l(x_t)}.$$

While we have focused so far on the history in which there is no investment at all, observe that the same analysis applies to the evolution of the belief for arbitrary histories, over any interval of time over which no investment takes place, provided the initial condition is accordingly modified. If s is the last date at which an investment is observed, and the belief immediately after this investment is p_s , we simply replace p_0 by p_s in the definition of γ .

Finally, observe that a simple change of variables yields

$$\int_0^t (1 - F_\theta(x_s)) ds = \int_{x_0}^{x_t} (1 - F_\theta(x)) g'(x) dx = \frac{1}{\lambda} \int_{x_0}^{x_t} \frac{l'(x)(1 - F_\theta(x))}{l(x)(F_0(x) - F_1(x))} dx.$$

This gives us the following formula for the probability of no investment in state θ , namely

$$G_{\theta,t} = e^{-\int_{x_0}^{x_t} \frac{l'(x)(1 - F_\theta(x))}{l(x)(F_0(x) - F_1(x))} dx}.$$

3.3. Evolution after an Investment. If an investment occurs at date t , the evolution of the belief is discontinuous at date t . The belief p_t jumps up, since agents become suddenly more optimistic (by MLRP). Simultaneously, the threshold x_t jumps down. More precisely, let (L_t, x_t) denote the belief and threshold immediately before the investment, and (L_t^+, x_t^+) these values immediately after the investment. By Bayes' rule, the public belief jumps up to

$$L_t^+ = \frac{1 - F_1(x_t)}{1 - F_0(x_t)} L_t > L_t.$$

Therefore, the threshold x_t jumps down to the solution to

$$l(x_t^+) L_t^+ = \frac{c}{1 - c}.$$

Taken together, the formulas derived in the last two subsections allow us to solve recursively for the threshold x_t after any arbitrary history h_t .

3.4. Examples. We provide here a pair of parametric examples. Consider first the case in which the distributions are given by, for all $x \in [0, 1]$,

$$F_1(x) = \frac{e^{ax} - 1}{e^a - 1}, \text{ and } F_0(x) = \frac{e^a - e^{a(1-x)}}{e^a - 1},$$

for some $a > 0$. Note that the range of the likelihood ratio is $l(x) \in [e^{-a}, e^a]$. The parameter a is a measure of the informativeness of the private signals, as a larger a implies a larger range of possible likelihood ratio values.

For the initial condition $\gamma = 1$ (which obtains, for instance, for $p_0 = 1/2$ and $c = 1/2$), the evolution of the belief and of the threshold up to the first investment are given by

$$x_t = \frac{1}{a} \ln \left(\frac{e^{\frac{\lambda t + a}{2}} + 1}{e^{\frac{\lambda t - a}{2}} + 1} \right), \quad p_t = \frac{e^a}{e^a + \left(\frac{e^{\frac{\lambda t + a}{2}} + 1}{e^{\frac{\lambda t - a}{2}} + 1} \right)^2}.$$

The details are in the appendix. As discussed, if an investment takes place at some date, both the public belief and the threshold jump. The evolution of the cut-off and the public belief is shown below for the case in which $\lambda = a = 1$, and an investment takes place at date $t = 5$.

The *ex ante* probability of no investment ever in the good state is given by

$$G_{1,t} = \left(\frac{e^{-\frac{a}{2}} + e^{-\frac{\lambda t}{2}}}{e^{-\frac{a}{2}} + 1} \right)^2.$$

Since the limit of this probability as $t \rightarrow \infty$ is bounded away from zero, there is a positive probability that no one ever invests, although the state is good. This probability decreases as the informativeness of the signals a increases.




FIGURE 1. Cut-off (left) and belief (right) over time with an investment at $t = 5$.

As a second example, take the power distributions given by, for all $x \in [0, 1]$,

$$F_1 = x^2, \quad F_0 = 1 - (1 - x)^2.$$

The probability that no one ever invests while the state is good is given by

$$G_{1,t} = \frac{1 - x_t}{x_t} e^{\frac{1}{2x_t} - 1},$$

which tends to 0 as $t \rightarrow \infty$, because, along such a history, the cut-off x_t tends to one. Here as well, details can be found in the appendix. This means that, in this example, almost surely, investments will eventually take place when the state is good. Since this is true for the first investment, it is also true for later investments, so that the total number of investments is unbounded.

As is easy to check, the likelihood ratio of the signal distribution is a bounded function in the first example, while it is not in the second. We shall prove in the next section that this distinction explains the different asymptotic properties of these two examples.

4. ASYMPTOTIC PROPERTIES

Although the learning process is biased, its asymptotic properties are the same as in traditional models of social learning. As is standard, we define private signals to be *unbounded* if $\lim_{x \rightarrow 0} l(x) = 0$ and $\lim_{x \rightarrow 1} l(x) = \infty$. This means that extreme signals are arbitrarily informative. Signals are *bounded* if the first limit is strictly positive, and the second is finite. Note that this does not partition the set of all distributions (for instance, it could be that $f_1(0)$ is equal to zero, but $f_0(1)$ is not).

From equation (1), it follows that an agent with the highest possible private signal, signal $x = 1$, will be indifferent between investing or not (if ever) if $l(1)L_t = c/(1 - c)$, where L_t is the likelihood ratio of the public belief resulting from the public history up to date t . Let \underline{L} denote the highest likelihood ratio for which, given that the public history leads to this likelihood ratio, it is optimal for such an agent not to invest (with the convention that $\underline{L} = 0$ if he always does). It follows that, if the signals are unbounded, $\underline{L} = 0$, while with bounded signals,

$$\underline{L} := \frac{c}{1 - c} \frac{1}{l(1)} < \frac{c}{1 - c}.$$

Similarly, define \bar{L} to be the lowest likelihood ratio for which an agent with the lowest possible signal, signal $x = 0$, finds it optimal to invest. With unbounded signals, $\bar{L} = +\infty$, while with bounded signals,

$$\bar{L} := \frac{c}{1 - c} \frac{1}{l(0)} > \frac{c}{1 - c}.$$

In terms of beliefs, this means that, defining

$$\underline{p} := \frac{\underline{L}}{1 + \underline{L}}, \text{ and } \bar{p} := \frac{\bar{L}}{1 + \bar{L}},$$

the probabilities \underline{p}, \bar{p} are in $(0, 1)$ if signals are bounded, and in $\{0, 1\}$ if they are unbounded. If signals are unbounded, then, independently of the history up to t , an agent arriving at date t will follow his signal if this signal is extreme enough. While it is not hard to see how this implies complete learning if the state is good, this is only slightly subtler if the state is bad: although later agents do not observe the informative actions of the agents with sufficiently low

signals, they will infer as much from their absence over the long-run. We see here the key role of two assumptions: the arrival rate is common knowledge (so that the absence of negatives can be correctly interpreted), and there is only one action that is hidden (so that agents can infer it from its absence). When signals are bounded, cascades can happen, just as in the traditional model, and for the same reason: histories leading to a belief above \bar{p} , or to the belief \underline{p} have positive probability under either state.⁷ In both cases, eventually, almost all agents take the same action. This discussion is summarized in the following set of results.

Proposition 1. *Beliefs converge a.s., with limit*

$$p_\infty := \lim_t p_t \in \{\underline{p}\} \cup [\bar{p}, 1].$$

The threshold x_t a.s. converges to either 0 or 1, with limit

$$x_\infty := \lim_t x_t \in \{0, 1\}.$$

Proof. See Appendix. □

This does not yet say that beliefs converge to the correct value if signals are unbounded, but simply that they converge to either 0 or 1. Turning to investments, we have the following.

Proposition 2. *If signals are unbounded, the probability that there is no investment ever in the good state is zero. It is positive if they are bounded.*

Proof. See Appendix. □

If the number of investments is finite, then, after the last investment, the public belief would decrease to the lower bound \underline{p} . On the other hand, if there was an infinite number of investments, the belief cannot converge to the lower bound \underline{p} , because, after each investment, the public belief must exceed the cost c . Hence, in that case, beliefs must converge to (no less than) the upper bound \bar{p} . This is summed up in the following lemma.

Lemma 3. *The total number of investments is finite (resp., infinite) if and only if the belief converges to \underline{p} (resp., greater than or equal to \bar{p}).*

⁷The asymmetry is due to the fact that, in case of an investment, the posterior belief might jump above \bar{p} , but in case of no investment, the belief trajectory is continuous.

As expected, complete learning occurs if and only if signals are unbounded.

Proposition 4. *The belief converges to the correct value almost surely if and only if signals are unbounded.*

Proof. See Appendix. □

Together with Proposition 1, this means that, if signals are unbounded, the number of investments is a.s. infinite in the good state, and finite in the bad state.

5. ONE-INVESTMENT GAME

In both the traditional, BHW model, and in this model, learning is complete when signals are unbounded. Therefore, we turn our attention to the case of bounded signals. To compare the likelihood of cascades under biased learning, relative to traditional learning, one should compare, in particular, the probability that there is a total of exactly k investments over the infinite horizon under both scenarios, for all integers k . While it might be possible to do so, we shall make a simpler comparison, by focusing on the probability that at least one investment is ever made. We may, and will interpret this as a game in which there is a single investment opportunity; obviously, this interpretation does not affect the agents' behavior, since agents cannot choose the timing of their decision. In fact, assuming that agents cannot wait before taking their action is not restrictive. If other agents do not delay either, an agent cannot gain by waiting: either nothing happens, or the game ends. This result does not require agents to discount future payoffs. More formally,

Proposition 5. *If all other agents follow the cut-off strategy x_t described in Section 3, delaying investment is not optimal, for any discount rate.⁸*

Proof. See Appendix. □

Observe that this version of the game admits independent economic interpretations, in which the winner takes all. It is of no use to discover a product that has been already patented, or to prove a result that has been already published.

⁸While it seems plausible that acting immediately is also the unique equilibrium strategy in this game, addressing this question would require defining the continuous-time game which is beyond the scope of this paper.

5.1. **The BHW Model.** The BHW model differs from ours in two respects: all actions are observed, and arrivals are not random. But given that actions are observed, whether arrivals are random or not is irrelevant to the decisions of the agents, and we might, for concreteness, keep on viewing arrivals as random. The public history up to date t is summarized by the sequence of individuals who arrived up to this date, and what their decisions were. Because we are assuming here that a single investment ends the game, this further reduces to the number n of agents who arrived and chose not to invest. Between each arrival, the threshold that characterizes the optimal strategy is now constant. Let x_k denote this cut-off when there have been k decisions not to invest so far.

Therefore, the probability that no agent invests, among the first n agents, conditional on the state θ , is given by

$$B_{\theta,n} := \prod_{k=0}^{n-1} F_{\theta}(x_k).$$

The thresholds x_k can then be solved recursively. Clearly, $x_0 = l^{-1}(\gamma)$, and, from Bayes' rule, x_n , $n \geq 1$, solves

$$l(x_n) \frac{B_{1,n}}{B_{0,n}} = \gamma,$$

since the right-hand side is the likelihood ratio of an agent with private signal x_n , given the public history. Thus, the thresholds solve the first-order difference equation

$$l(x_{n+1}) = \frac{F_0(x_n)}{F_1(x_n)} l(x_n).$$

Our objective is then to compare the limit of this probability,

$$B_{\theta,\infty} := \prod_{k=0}^{\infty} F_{\theta}(x_k),$$

with the analogous probability under biased learning derived in subsection (3.2), namely,

$$G_{\theta,\infty} := e^{-\int_{x_0}^1 \frac{1-F_{\theta}(x)}{F_0(x)-F_1(x)} \frac{l'(x)}{l(x)} dx}.$$

While those two expressions bear little in common, we shall see that the comparison hinges upon a simple statistical property that we define next.

5.2. The Increasing Hazard Ratio Property (IHRP). By the strict monotone likelihood ratio property, the mapping $x \mapsto (c/(1-c))/l(x)$ is strictly decreasing, mapping $(0, 1)$ onto (\underline{L}, \bar{L}) . Define the *hazard ratio* at the signal x as the ratio of the hazard rates conditional on the good and the bad state, that is,

$$H(x) := \frac{1 - F_0(x)}{1 - F_1(x)} l(x).$$

The (strict) *increasing hazard ratio property* (IHRP) holds if this mapping is strictly increasing. This property has been introduced in the statistical literature by Kalashnikov and Rachev (1986). Both parametric examples of Section 3.4. satisfy IHRP.

The relevance of IHRP for the BHW model is almost immediate. Because the posterior likelihood ratio after an investment is given by

$$\frac{1 - F_1(x)}{1 - F_0(x)} L,$$

where x is the cut-off signal and L is the prior likelihood ratio, and because the cut-off signal solves $L = (c/(1-c))/l(x)$, it follows that this posterior likelihood ratio is equal to $(c/(1-c))/H(x)$. To put it differently, IHRP states that this posterior likelihood ratio is decreasing in the cut-off signal, or alternatively, since the cut-off signal is decreasing in the prior likelihood ratio L (under MLRP), that the posterior likelihood ratio L^+ is increasing in the prior likelihood ratio L (after an investment).

To be a little more formal, let $L^+ : \mathbb{R}^+ \rightarrow \mathbb{R}^+$ denote this function mapping the prior likelihood ratio into the posterior likelihood ratio, and let us say that the signals satisfy *updating monotonicity after good news* (UMG) if L^+ is strictly increasing. Given our discussion, the following is immediate.

Proposition 6. *Under MLRP, IHRP is necessary and sufficient for UMG.*

A corresponding property can be defined for the event in which an agent does not invest. Define the function K by $K(x) = F_0(x)l(x)/F_1(x)$ for all $x \in (0, 1)$. The (strict) *increasing failure ratio property* (IFRP) holds if this mapping is strictly increasing. Both parametric examples of Section 3.4. satisfy IFRP. Observe that the posterior likelihood ratio, conditional on such an event, is given by

$$\frac{(c/(1-c))}{K(x)},$$

where x is a decreasing function of the prior likelihood ratio. It follows that this posterior likelihood ratio, L^- , is an increasing function of the prior likelihood ratio L . That is, let us say that the signals satisfy *updating monotonicity after bad news* (UMB) if L^- is strictly increasing. Then it immediately follows that, under MLRP, IFRP is necessary and sufficient for UMB.

As it turns out, these concepts play an important role for the existence of cascades. If the public likelihood ratio L lies in (\underline{L}, \bar{L}) , as defined previously, an agent's strategy will depend on his private signal: he will invest if and only if his signal is above some cut-off in $(0, 1)$. (Recall that we are focusing on the case of bounded signals, so that (\underline{L}, \bar{L}) is a proper subset of $(0, +\infty)$, that is (\underline{p}, \bar{p}) is a proper subset of $(0, 1)$.) On the other hand, if this ratio ever exits the interval (\underline{L}, \bar{L}) a cascade starts and observational learning stops: all agents take the same action independently of their signal.

Observe now that, by definition of \bar{L} , and of the mapping L^+ , \bar{L} is a fixed-point of this mapping as: $\frac{1-F_1(0)}{1-F_0(0)}\bar{L} = \bar{L}$. Similarly, \underline{L} is a fixed-point of the mapping L^- . Therefore, updating monotonicity after good news guarantees that, if the prior likelihood ratio start below \bar{L} , the posterior remains below it. Similarly, UMB guarantees that the posterior likelihood ratio is above \underline{L} if the prior is.

As a result, along with MLRP, IHRP (and IFRP) is a sufficient condition guaranteeing that, provided that the first agent uses a strategy that depends on his private signal, then all later agents will: if the game does not start with a cascade, a cascade will never occur. Conversely, it is a necessary condition, in the sense that, if either condition is violated, it is possible to find parameters p_0 and c for which a cascade occurs with positive probability.

That IHRP (and IFRP) are the necessary and sufficient conditions for (the absence of) cascades appears to be new. The early literature (for instance, BHW) established that cascades occur when signals are discrete, while later contributions (in particular, Smith and Sørensen, 2001) showed that this result does not necessarily hold with continuous signals. Indeed, it is easy to see that IHRP and IFRP are necessarily violated in the case of discrete signals.⁹ On the other

⁹The definition above assumes a continuous density. In the discrete case, the function H must be defined by

$$H(k) \quad : \quad = \frac{\sum_{j \geq k} \mathbb{P}_1[x_j]}{\sum_{j \geq k} \mathbb{P}_0[x_j]} L = \frac{\sum_{j \geq k} \mathbb{P}_1[x_j]}{\sum_{j \geq k} \mathbb{P}_0[x_j]} \frac{\mathbb{P}_1[x_k]}{\mathbb{P}_0[x_k]},$$

which is decreasing at every boundary point $L = \frac{\mathbb{P}_1[x_k]}{\mathbb{P}_0[x_k]}$.

hand, unbounded signals (more precisely, $f_1(0) = 0$) is stronger than (local) IHRP, in the sense that it implies $\bar{L} = \infty$, so that clearly L^+ has to be increasing for high enough likelihood ratios.

To get some intuition for these results, recall that $H(x)$ is the ratio of the hazard rate $h_1(x)/h_0(x)$, where

$$h_\theta(x) := \frac{f_\theta(x)}{1 - F_\theta(x)}$$


is the hazard rate in state θ , namely the density at the cut-off x , conditional on an investment. If the public belief increases, the threshold x decreases, and therefore, according to IHRP, so does $H(x)$. That is, as the public belief increases, and conditional on an investment, it becomes less likely that the investor was least optimistic (among agents who would invest), in the good state relative to the bad state. To see why this is consistent with updating monotonicity, suppose that it were the case that, to the contrary, the least optimistic investor was conditionally more likely in the good state than in the bad state. Because of MLRP, an investment would still be good news, but, since higher signals are better news, this would mitigate the good news. The likelihood ratio might not increase much for higher values of L , and this might violate the monotonicity of the mapping L^+ .

While it is not clear *a priori* whether it is more reasonable to assume that the hazard ratio is increasing or decreasing, note that, in case it is decreasing, the failure of updating monotonicity implies that a single decision to invest will trigger an investment cascade. That is, the dichotomy is rather extreme: if the first agent uses a strategy that depends on his signal, all later agents will do so if the hazard ratio is increasing (assuming IFRP), while a single investment will trigger a cascade if the hazard ratio is decreasing.¹⁰ This can be most easily understood with the help of the following figure.

It is worth mentioning that, of course, IHRP neither implies, nor is implied by MLRP. While MLRP states that higher signals are better news, IHRP states that this remains true, conditional on truncations i.e. substituting $f_\theta(x)$ with $h_\theta(x)$. Therefore, these two stochastic orders are related, as IHRP implies MLRP if in addition $h_\theta(x)$ decreases in the state θ . The relationship is explored at the end of the appendix. Also, because Smith and Sørensen [13] show that monotonicity of the mapping L^+ is implied by the assumption that the private belief log-likelihood

¹⁰In fact, a decreasing hazard ratio is inconsistent with both bounded signals and MLRP. But it is consistent with MLRP and signals that are not unbounded, as the example $F_1(x) = x$, $F_0(x) = x(3 - x^2)/2$ illustrates.

ratio is log-concave, while IHRP is necessary and sufficient, it follows indirectly that the log-concavity assumption is stronger than IHRP (and arguably more complicated to verify). On the implications of log-concavity, see also Smith and Sørensen [14].



5.3. The Result. We are finally ready to compare the probabilities of no investment ever in both models, conditional on a given state. One might suspect that this comparison depends on the state, but it turns out that this is not the case, because the ratio of these conditional probabilities is the same in both models, as the next lemma establishes.

Lemma 7. *It holds that*

$$\frac{G_{1,\infty}}{G_{0,\infty}} = \frac{B_{1,\infty}}{B_{0,\infty}}.$$

Proof. By definition,

$$G_{1,\infty} = e^{-\int_{x_0}^1 \left(1 + \frac{1-F_0(x)}{F_0(x)-F_1(x)}\right) \frac{l'(x)}{l(x)} dx} = \frac{\gamma}{l(1)} G_{0,\infty},$$

and similarly, by definition,

$$B_{1,\infty} = \frac{\prod_{k=0}^{\infty} F_1(x_k)}{\prod_{k=0}^{\infty} F_0(x_k)} B_{0,\infty} = \frac{\gamma}{l(1)} B_{0,\infty}.$$

The result follows. \square

The main result of this section establishes that the probability that no investment ever takes place is higher in the BHW model than in the biased learning model, if IHRP holds. That is, under IHRP, biased learning leads to higher investment, independently of the state. If the hazard ratio is constant, both models lead to the same amount of investment, and biased learning leads to lower investment if the hazard ratio is decreasing.

Proposition 8. *Assume IHRP. Conditional on either state, the probability of investment is always larger in the hidden action model:*

$$G_{\theta,\infty} < B_{\theta,\infty}.$$

This inequality is reversed if the hazard ratio is decreasing.

Proof. See Appendix. \square

Unfortunately, it is difficult to provide a simple intuition for this result. The proof hinges on the inequality

$$F_{\theta}(x_k) \geq e^{-\int_{x_k}^{x_{k+1}} \frac{1-F_{\theta}(x)}{F_0(x)-F_1(x)} \frac{l'(x)}{l(x)} dx},$$

which holds under IHRP because it gives us precisely the appropriate lower bound on the integrand appearing on the right-hand side. This inequality can be interpreted as follows. Start from the same (initial public belief and hence) initial threshold x_k in both models, define x_{k+1} as the new threshold after bad news in the BHW model and define Δt_k as the time it takes for the threshold to go from x_k up to x_{k+1} in the hidden action model. Whichever the state, the probability that there is one investment in the biased learning model during the time span Δt_k is always larger than the probability that the next agent that arrives invests in the BHW model. That is, the comparison is between what the next agents does in the BHW model and what none, one, or many agents do if they arrive in the time span Δt_k in the hidden action model.

5.4. Welfare Comparison. Since the probabilities of cascades are not the same in both models, it is natural to wonder which one aggregates information better. In the good state, it is optimal for someone to invest, while in the bad state, it is optimal for all to abstain. Therefore, we define welfare by the expectation of the utility of this eventual outcome (either someone eventually invests, or no one does). Since players do not internalize the informational externalities, there is no reason to expect *a priori* that having more information is necessarily better. Indeed, whether this is the case depends here again on the hazard ratio, as the next result establishes.

Proposition 9. *Under IHRP, the welfare in the hidden action model is lower than in the benchmark model. It is higher if the hazard ratio is decreasing.*

Proof. The welfare in the hidden action model is, by definition,

$$\begin{aligned}
 W(G) &:= \mathbb{E}_\theta [(1 - G_\theta) u(I) + (G_\theta) u(N)] \\
 &= p_0 (1 - G_1) (1 - c) + (1 - p_0) (1 - G_0) (-c) \\
 &= p_0 (1 - c) \left((1 - G_1) - \gamma \left(1 - \frac{l(1)}{\gamma} G_1 \right) \right) \\
 &= p_0 (1 - c) (1 - \gamma + (l(1) - 1) G_1),
 \end{aligned}$$

and likewise for $W(B)$. Since $l(1) > 1$, as 1 is the highest signal, this expression is increasing in G_1 , and therefore

$$W(G) < W(B),$$

whenever the IHRP is satisfied. The inequality is obviously reversed if the hazard ratio is decreasing. \square

Therefore, even though the model with biased learning performs better in the good state, as it always leads to a higher probability of investment, it achieves a lower welfare than the BHW model, at least under IHRP.

5.5. Many Investments. So far, the comparison between the model with biased learning and the benchmark BHW model has been performed in the game with one investment. Comparing the two models more generally is difficult. For example, computing the probability of exactly two investments involves a summation (or, in the biased learning model, an integration) over the times at which one agent first invested. Even in the simple example of an exponential distribution, that

satisfies IHRP (see Section 3.4.), closed-form solutions appear elusive. Nevertheless, numerical computations can be performed in this case, which we briefly present here.

Consider first the model with biased learning. Let $G_\theta^0(x_0, x_t)$ denote the probability of no investment in state θ during the time interval required for the threshold to go from an initial value x_0 to the value x_t , and let $G_\theta^n(x_0, 1)$ denote the probability of exactly n eventual investments in state θ (i.e., over the time interval required for the threshold to go from the initial value x_0 to 1.) This probability must satisfy the recursion

$$G_\theta^{n+1}(x_0, 1) = \int_{x_0}^1 G_\theta^0(x_0, x_t) \lambda (1 - F_\theta(x_t)) G_\theta^n(x_t^+, 1) dx_t,$$

as the probability that there are n investments overall can be decomposed as the sum of the probabilities, as t varies, that the first investment occurs exactly at time t , and that exactly $n - 1$ investments overall are made afterwards (given the resulting new initial value).

In the exponential example, considering the state $\theta = 1$ for instance, we obtain

$$G_1^0(x_0, x_t) = e^{-\lambda \int_0^t (1-F_1) dt} = e^{-\int_{x_0}^{x_t} \frac{1-F_1}{t} dx} = \left(\frac{e^{ax_0} - 1}{e^{ax_t} - 1} \right)^2,$$

and the recursion becomes

$$G_1^{n+1}(x_0, 1) = \int_{x_0}^1 \left(\frac{e^{ax_0} - 1}{e^{ax} - 1} \right)^2 \left(2a \frac{e^{ax}}{e^{ax} - 1} \right) G_1^n(x/2, 1) dx.$$

This integration can only be performed explicitly in the case $n = 1$.

A similar decomposition can be used in the BHW model: the probability of k overall investments is the sum, over the index n of the first agent to invest, of the probability that the first agent to invest is the n -th agent, times the probability that there are exactly $k - 1$ agents investing in the game (where, for the latter probability, we use as initial belief the public belief resulting from a first investment by the n -th agent.)

These recursions allow numerical computation of these probabilities. Figure 2 below depicts the (log plot of) the probabilities of k eventual investments in both models, in the exponential example with parameter $a = 1/4$ (the same pattern arises for all values of a). These probabilities cross exactly once. That is, there exists an integer k^* such that, for $k < k^*$, the probability of k eventual investments is higher in the BHW model, while the opposite is true for $k \geq k^*$. It can be numerically verified that the probability distribution of k or less eventual investments of

the BHW model first-order stochastically dominates the one of the biased learning model. This suggests that, at least in this example, restricting attention to the one-investment game does not seem too misleading.




FIGURE 2. Log plot of the probabilities of k eventual investments for $a = 1/4$ (dots: BHW benchmark; squares: biased learning model).

6. PAYOFF INFORMATION: ONE SUCCESSFUL INVESTMENT GAME

As discussed in the introduction, there are many applications in which not only the opportunity to invest is hidden, but also decisions to invest that resulted in a failure. To discuss such a variation, it is necessary to introduce another dimension to the model, namely, we must specify the probabilities with which, conditional on a given state, a decision to invest is successful. To simplify, we assume here that the game ends with the first successful investment. This adds another layer to the inferential problem. The absence of an observed success has now three possible explanations: either no agent had the opportunity to invest, or agents with such an opportunity chose not to, given their private signal, or they chose to, but failed.

Define P_θ as the probability that an investment succeeds, conditional on the state θ . Assume that

$$0 < P_0 < P_1 < 1,$$

so that success is more likely in the good state. Successes are conditionally independent across agents who invest.

The definitions of histories and strategies, as well as the derivations are quite similar to those in the baseline model with biased learning, and we shall omit the details. Let $h_t = (t, \{\emptyset\})$

denote a history without successful investment up to date t , and

$$\bar{G}_{\theta,t} := \mathbb{P}[h_t|\theta]$$

the probability of the history h_t , conditional on state θ . The optimal strategy is still a cut-off strategy, with cut-off \bar{x}_t . Note that the function $\bar{G}_{\theta,\cdot}$ satisfies $\bar{G}_{\theta,0t} = 1$ and, for all t (ignoring terms of order dt^2 or higher)

$$\bar{G}_{\theta,t+dt} = \bar{G}_{\theta,t} ((F_\theta(\bar{x}_t) + (1 - P_\theta)(1 - F_\theta(\bar{x}_t))) \cdot (\lambda dt) + 1 \cdot (1 - \lambda dt)),$$

for $dt > 0$, with the obvious interpretation. Taking limit and solving the resulting differential equation gives

$$\bar{G}_{\theta,t} = e^{-\lambda P_\theta \int_0^t (1 - F_\theta(x_s)) ds}.$$

After the history h_t , the cut-off \bar{x}_t must solve

$$(3) \quad l(\bar{x}_t) = \gamma \frac{\bar{G}_{0,t}}{\bar{G}_{1,t}} = \gamma e^{\lambda(P_1 \int_0^t (1 - F_1(x_s)) ds - P_0 \int_0^t (1 - F_0(x_s)) ds)}.$$

Differentiating with respect to t and integrating by separation of variables yields equation

$$t = \bar{g}(\bar{x}_t) := \frac{1}{\lambda} \int_{x_0}^{\bar{x}_t} \frac{l'(x)}{l(x)} (P_1 (1 - F_1(x)) - P_0 (1 - F_0(x)))^{-1} dx,$$

which implicitly defines \bar{x}_t , along with the initial condition $x_0 = l^{-1}(\gamma)$.

It turns out that this model has similar properties than the baseline model, because the cut-offs \bar{x} and x share the same rate of convergence, as the next lemma shows.

Lemma 10. *There exists $\alpha, \beta > 0$ such that, for all t ,*

$$x_t < \bar{x}_{\alpha t} < x_{\beta t}.$$

Therefore, the hidden failure model has the same qualitative behavior as the hidden action model.

Therefore, the probability of an eventual (and successful) investment is one if the state is good, in both models, if signals are unbounded, and it is less than one otherwise. The comparison between cut-offs is in general ambiguous. Nevertheless, we have the following result.

Proposition 11. *If p_0 is low enough, and signals are sufficiently bounded ($l(1) < (1 - P_0)/(1 - P_1)$), the cut-off \bar{x} is below the cut-off x , i.e., for all t ,*

$$x_t \geq \bar{x}_t.$$

If p_0 is high enough and signals are unbounded, the inequality is reversed.

Therefore, when the prior is low enough and signals are sufficiently bounded, and for any time t , the probability that an investment occurs by time t is higher when failures are not observable. Intuitively, assume we start from the same initial prior and observe “no news” in both models. In the baseline model, this means simply that some pessimistic agents might have arrived and chose not to invest, while in the hidden failure model, it might also be that some optimistic agent might have invested and failed. Hence the more pronounced caution by agents, and the lower threshold. Of course, this intuition is incomplete, since inferences do depend on the threshold function itself. The increased caution of earlier agents mitigates the bad news that a no-investment history conveys, making later agents relatively less cautious. This is why additional conditions are required to ensure that these thresholds can be globally ranked.

7. CONCLUSIONS

At least as far as convergence is concerned, we have shown that relaxing the assumptions that all types of decisions are observable does not change significantly the asymptotic learning properties of the model. Beliefs always converge and will converge to the true value for sure if and only if private beliefs are unbounded. So, even if one action is hidden and can only be inferred, the market aggregates the information correctly anyway when beliefs are unbounded.

With bounded signals, whichever model delivers a higher probability of investment depends on a property of the signal distribution. If the hazard ratio is increasing in the signal, then investment is more likely in the model with biased learning, and welfare is lower. These conclusions are reversed if the hazard ratio is decreasing. In the standard model, the property of increasing hazard ratio precisely captures the condition ensuring that cascades cannot occur.

Those two results each raise their own set of questions. How little information does it take to obtain complete learning with unbounded signals? Does the increasing hazard ratio property always ensure that more information is better in social learning? It is particularly intriguing that this property turns out to play a key role both in the comparison between both models, and within the BHW model itself.

8. APPENDIX

8.1. Proofs.

Proof of Proposition 1. Since the public belief p_t is a bounded martingale, the martingale convergence theorem implies almost sure convergence to a value $p_\infty \in [0, 1]$.

Beliefs cannot converge to any interior value

$$p_\infty \in (\underline{p}, \bar{p}),$$

namely, beliefs cannot settle within any $\varepsilon > 0$ of any interior value $p_\infty \in (\underline{p}, \bar{p})$, as beliefs would jump discretely after any investment, or would decrease to \underline{p} otherwise. The threshold signal x_t must converge because it is a continuous function of the belief. It must converge to its boundaries $x_\infty \in \{0, 1\}$, as the limit threshold cannot be any interior value $x_\infty \in (0, 1)$: otherwise, a later investment would almost surely occur, and the belief and the threshold would jump. The result follows. \square

Proof of Proposition 2. In a history without any investment, the belief would converge to zero and the threshold to one. We now show that with unbounded beliefs this is a zero probability history, and conversely:

$$l(1) = +\infty \iff \lim_t G_{1,t} = 0.$$

Observe that the probability of no investment ever under state 1 converges, as it is an increasing and bounded function. Since

$$F_0(x_t) - F_1(x_t) \leq 1 - F_1(x_t),$$

it follows that from (3) that

$$G_{1,t} = e^{-\int_0^t \lambda(1-F_1(x_s))ds} \leq \frac{c}{1-c} \frac{1-p_0}{p_0} \frac{1}{l(x_t)},$$

so

$$\lim_t G_{1,t} = 0 \quad \text{if } l(1) = +\infty.$$

Conversely, if beliefs are bounded we have that $l(1) < \infty$. Then

$$h(x) := \frac{l(x)}{l'(x)} \lambda(F_0(x) - F_1(x))$$

converges to 0, and, so for all x sufficiently close to 1,

$$x'_t = h(x) > -h'(1)(1 - x_t) - M(1 - x_t)^2,$$

with $h'(1) < 0$ as $h(x) > 0$, for some $M > 0$. This implies that, for all t sufficiently large,

$$1 - x_t \leq \frac{-h'(1)}{M + C_1 e^{-h'(1)t}},$$

for some constant C_1 . Since

$$1 - F_1(x_s) \leq f_1(1)(1 - x_s) + C_2(1 - x_s)^2,$$

and

$$\int^t \frac{-h'(1) ds}{M + C_1 e^{-h'(1)s}} = \ln(C_1 + M e^{-h'(1)t}) / M < \ln(C_1 + M) / M,$$

it follows that $G_{1,t}$ is bounded below, so that

$$l(1) < +\infty \Rightarrow \lim_t G_{1,t} > 0.$$

□

Proof of Proposition 4. Assume signals are unbounded. By Proposition 1, the public belief must either converge to 0 or 1. Assume that the state is good. The belief cannot converge to 0, because that would imply that the number of investments is finite, contradicting Proposition 2 (recall that, after an investment, the public belief exceeds c).

If the state is bad, then the public belief converges to zero. By the martingale property, if the belief converges to 1 in the good state it must converge to zero in the bad state, namely

$$p_t = \mathbb{E}_t[p_\infty] = \mathbb{P}_\infty[\theta = 1]p_t + \mathbb{P}_\infty[\theta = 0](1 - p_t), \text{ and so } \mathbb{P}_\infty[\theta = 1] = 1 \Rightarrow \mathbb{P}_\infty[\theta = 0] = 0.$$

Assume that signals are bounded. In the good state, beliefs can converge to the wrong value \underline{p} because there is a positive probability of a no investment history. In the bad state, beliefs converge to (no less than) \bar{p} with positive probability. Indeed, if they did not, then it would be the case that $\bar{p} = 1$ (by definition of \bar{p}), and this is impossible with bounded beliefs. □

Proof of Proposition 5. Let $\tau(x, t)$ denote the stopping time of a player arriving at instant t with signal x , and let $F(s; t, x)$, $s \geq t$, denote the corresponding c.d.f. Fix an equilibrium and suppose for the sake of contradiction that $0 < F(s; t, x) < 1$ for some x, t and $s > t$, for some finite s . Let q_τ denote the private belief of this agent at time $\tau \geq t$, given his signal x and the equilibrium

strategies (conditional on the event E_τ that no one invested up to τ). Observe that q_τ is non-increasing, and constant over some interval of time $[t', t'')$ if and only if $F(t''; s, x) = F(t'; s, x)$ for all $s \leq t'$ and $x \in [0, 1]$. Assume that it is strictly profitable to invest at time t with signal x . Then, because the payoff from investing is strictly increasing in q_t , $F(s; t, x) < 1$ is only possible if $q_s = q_t$, and player i assigns probability one to no one investing before (or at the same) time than he does. In particular, any other player arriving in the interval of time (t, s) must invest with probability zero in that interval of time. Consider the event that some player arrives in this interval of time with a signal $x \geq x_t$, i.e. a player whose payoff from investing immediately is strictly positive. (x_t here depends obviously upon the equilibrium strategies.) This event has strictly positive probability, and thus, given the equilibrium strategies, there exists a player arriving at some time $t' \in (t, s)$ whose probability of investing first (after s) is strictly less than 1. This player would profitably gain from investing immediately at time t' . Assume now that it is strictly unprofitable to invest at time t with signal x . Plainly it remains unprofitable to invest at any later time, and so it cannot be that $0 < F(s; t, x)$ for some finite s . Therefore, a player is unwilling to delay unless $x = x_t$, i.e. he is indifferent between investing immediately or never. This event has zero probability, and so does not affect the analysis, in particular the determination of x_t . \square

Proof of Proposition 6. IHRP means that the inverse of H , denoted m , i.e.

$$m(x) := \frac{f_0(x)/(1 - F_0(x))}{f_1(x)/(1 - F_1(x))},$$

is strictly decreasing, so that, for $m := m(x_k)$,

$$\frac{1 - F_1(x)}{1 - F_0(x)} \leq ml(x),$$

for all $x \in [x_k, x_{k+1}]$, with equality for $x = x_k$. That is, we have

$$\frac{F_0(x) - F_1(x)}{1 - F_0(x)} \leq ml(x) - 1,$$

which implies that the right-hand side is positive, as the left-hand side is.

Given Lemma 5, it is enough to show $B_{\theta, \infty} > G_{\theta, \infty}$ for $\theta = 0$. We have

$$\prod_{k=0}^{\infty} F_0(x_k) > e^{-\int_{x_0}^1 \frac{1-F_0(x)}{F_0(x)-F_1(x)} \frac{l'(x)}{l(x)} dx} = \prod_{k=0}^{\infty} e^{-\int_{x_k}^{x_{k+1}} \frac{1-F_0(x)}{F_0(x)-F_1(x)} \frac{l'(x)}{l(x)} dx},$$

so it suffices to show that, for all k ,

$$\ln F_0(x_k) + \int_{x_k}^{x_{k+1}} \frac{l'(x)}{l(x)} \frac{1 - F_0(x)}{F_0(x) - F_1(x)} dx \geq 0.$$

We have

$$\begin{aligned} & \int_{x_k}^{x_{k+1}} \frac{l'(x)}{l(x)} \frac{1 - F_0(x)}{F_0(x) - F_1(x)} dx \geq \int_{x_k}^{x_{k+1}} \frac{l'(x)}{l(x)(ml(x) - 1)} dx \\ & = \ln \frac{ml(x_{k+1}) - 1}{l(x_{k+1})} - \ln \frac{ml(x_k) - 1}{l(x_k)} = \ln \frac{ml(x_k) - \frac{F_1(x_k)}{F_0(x_k)}}{ml(x_k) - 1}, \end{aligned}$$

using that $l(x_{k+1}) = \frac{F_0(x_k)}{F_1(x_k)}l(x_k)$. Therefore, the inequality

$$\ln F_0(x_k) + \int_{x_k}^{x_{k+1}} \frac{l'(x)}{l(x)} \frac{1 - F_0(x)}{F_0(x) - F_1(x)} dx \geq 0,$$

would be implied by

$$\ln F_0(x_k) + \ln \frac{ml(x_k) - \frac{F_1(x_k)}{F_0(x_k)}}{ml(x_k) - 1} \geq 0,$$

Rearranging, this is equivalent to

$$\frac{1 - F_1(x_k)}{1 - F_0(x_k)} \geq ml(x_k),$$

but this is the case, since in fact both sides are equal by definition of m . By immediate inspection, this chain of inequalities is tight if the hazard ratio is constant, and reversed if it is decreasing. \square

Proof of Lemma 10. For any t choose a time t' such that $x_t = \bar{x}_{t'}$. Because

$$\begin{aligned} t &= \int_{x_0}^{x_t} \frac{l'}{l} (F_0 - F_1)^{-1} dx \text{ and} \\ t' &= \int_{x_0}^{\bar{x}_{t'}} \frac{l'}{l} (P_1(1 - F_1) - P_0(1 - F_0))^{-1} dx, \end{aligned}$$

the ratio of the integrands, by l'Hospital rule, is given by

$$\begin{aligned} \lim_{x \rightarrow 1} \frac{F_0 - F_1}{P_1(1 - F_1) - P_0(1 - F_0)} &= \lim_{x \rightarrow 1} \frac{f_0 - f_1}{P_0 f_0 - P_1 f_1} = \frac{1 - l(1)}{P_0 - P_1 l(1)} \\ &= \begin{cases} 1/P_1 & \text{if } l(1) = \infty, \\ \frac{1-l(1)}{P_0 - P_1 l(1)} & \text{if } l(1) < \infty. \end{cases} \end{aligned}$$

The result follows. \square

Proof of Proposition 11. We need to show that

$$\int_{x_0}^{x_t} \frac{l'}{l} \left((F_0 - F_1)^{-1} - (P_1(1 - F_1) - P_0(1 - F_0))^{-1} \right) dx > 0.$$

A sufficient condition for $t < t'$ is

$$P_1(1 - F_1) - P_0(1 - F_0) > (F_0 - F_1),$$

or

$$\frac{1 - P_0}{1 - P_1} > \frac{1 - F_1}{1 - F_0}.$$

The left-hand side is constant and larger than one. As for the right-hand side, note that, because

$$\frac{f_0}{1 - F_0} > \frac{f_1}{1 - F_1},$$

it has a positive derivative, and so it is increasing from 1 (at $x = 0$) to $l(1)$ as

$$\lim_{x \rightarrow 1} \frac{1 - F_1}{1 - F_0} = l(1).$$

Hence, a sufficient condition for lower investment is

$$\frac{1 - P_0}{1 - P_1} > l(1).$$

which requires signals to be bounded: $l(1) < \infty$.

Conversely, if $l(1) = \infty$, then there exists a $\xi \in (0, 1]$ such that

$$\frac{1 - F_1(\xi)}{1 - F_0(\xi)} = \frac{1 - P_0}{1 - P_1},$$

and we then have

$$x > \xi \implies \frac{1 - F_1(x)}{1 - F_0(x)} > \frac{1 - P_0}{1 - P_1}.$$

Therefore, for $x_0 \geq \xi$,

$$\int_{x_0}^{x_t} \frac{l'}{l} \left((F_0 - F_1)^{-1} - (P_1(1 - F_1) - P_0(1 - F_0))^{-1} \right) dx < 0,$$

and the result follows. \square

8.2. **Exponential Example.** Take the following c.d.f.

$$F_1 = \frac{e^{ax} - 1}{e^a - 1}, \quad F_0 = \frac{e^a - e^{a(1-x)}}{e^a - 1}.$$

The threshold can be expressed as

$$\begin{aligned} \lambda t &= \int_{x_0}^x 2a \int \frac{e^a - 1}{1 + e^a - e^{as} - e^{a(1-s)}} ds = \left[2 \ln \left(\frac{e^{as} - 1}{e^a - e^{as}} \right) \right]_{x_0}^x \quad \text{or} \\ \lambda t &= 2 \ln \left(\frac{e^{ax} - 1}{e^a - e^{ax}} \right) - A_0 \quad \implies \quad x_t = \frac{1}{a} \ln \left(\frac{e^{\frac{\lambda t + A_0}{2} + a} + 1}{e^{\frac{\lambda t + A_0}{2}} + 1} \right). \end{aligned}$$

With a uniform prior, we have

$$\left(x_0 = \frac{1}{2} \quad \implies \quad A_0 = -a \right) \iff x_t = \frac{1}{a} \ln \left(\frac{e^{\frac{\lambda t + a}{2}} + 1}{e^{\frac{\lambda t - a}{2}} + 1} \right).$$

The likelihood ratio and the public belief are, respectively,

$$\begin{aligned} l(x_t) &= e^{a(2x_t - 1)} = e^{-a} \left(\frac{e^{\frac{\lambda t + a}{2}} + 1}{e^{\frac{\lambda t - a}{2}} + 1} \right)^2 \rightarrow e^a, \\ L_t &= e^a \left(\frac{e^{\frac{\lambda t - a}{2}} + 1}{e^{\frac{\lambda t + a}{2}} + 1} \right)^2 \rightarrow e^{-a} \quad \implies \quad p_t = \frac{L_t}{1 + L_t} \rightarrow \frac{1}{1 + e^a}. \end{aligned}$$

Because

$$e^{ax_t} = \frac{e^{\frac{\lambda t + a}{2}} + 1}{e^{\frac{\lambda t - a}{2}} + 1}, \quad \text{and} \quad \int_0^t \frac{1}{e^{\frac{\lambda t - a}{2}} + 1} dt = t - \frac{2}{\lambda} \ln \left(\frac{e^{\frac{\lambda t - a}{2}} + 1}{e^{-\frac{a}{2}} + 1} \right),$$

the probability of no investment in the good state is

$$\begin{aligned} G_{1,t} &= e^{-\int_0^t \lambda(1-F_1)dt} = e^{-\lambda \int_0^t (1 - \frac{e^{ax} - 1}{e^a - 1}) dt} = e^{-\lambda \int_0^t \frac{1}{e^{\frac{\lambda t - a}{2}} + 1} dt} \\ &= e^{-\lambda t} \left(\frac{e^{\frac{\lambda t - a}{2}} + 1}{e^{-\frac{a}{2}} + 1} \right)^2 = \left(\frac{e^{-\frac{a}{2}} + e^{-\frac{\lambda t}{2}}}{e^{-\frac{a}{2}} + 1} \right)^2 \rightarrow \left(\frac{e^{-\frac{a}{2}}}{e^{-\frac{a}{2}} + 1} \right)^2. \end{aligned}$$

Given the symmetry

$$F_0(x; a) = F_1(x; -a),$$

the probability of no investment in the bad state is

$$G_{0,t} = \left(\frac{e^{\frac{a}{2}} + e^{-\frac{\lambda t}{2}}}{e^{\frac{a}{2}} + 1} \right)^2 \rightarrow \left(\frac{e^{\frac{a}{2}}}{e^{\frac{a}{2}} + 1} \right)^2.$$

After an investment, the threshold $x_{\tau+}$ decreases discontinuously (e.g. take $a = 1$)

$$\begin{aligned} l(x_{\tau+}) &= \frac{1 - F_0(x_\tau)}{1 - F_1(x_\tau)} l(x_\tau) = (e^{-x_\tau}) l(x_\tau) \\ e^{2x_{\tau+}-1} &= e^{x_\tau-1} \implies x_{\tau+} = \frac{x_\tau}{2}, \end{aligned}$$

and the belief increases discontinuously to

$$p_{\tau+} = \frac{(e^{x_\tau}) p_\tau}{(1 - p_\tau) + (e^{x_\tau}) p_\tau} > p_\tau.$$

8.3. Polynomial Example. Take the following c.d.f.

$$F_1 = x^2, \quad F_0 = 1 - (1 - x)^2.$$

In this example beliefs are unbounded, namely: $l(x) \in [0, \infty]$. For $\lambda = 1$, $c = 1/2$ ($\implies x_0 = 1/2$), the threshold is determined implicitly by

$$t = \int_{0.5}^{x_t} \left(\frac{1}{2x^2(1-x)^2} \right) dx = \frac{1}{2} \left(\frac{1}{1-x_t} - \frac{1}{x_t} \right) + \ln \left(\frac{x_t}{1-x_t} \right).$$

The change of variables

$$t = \frac{1}{2} \left(\frac{1}{1-x} - \frac{1}{x} \right) + \ln \left(\frac{x}{1-x} \right) := g(x)$$

gives

$$\int_0^t (1 - F_1(x_s)) ds = \int_{0.5}^{x_t} (1 - F_1(x)) g'(x) dx = \ln \frac{x_t}{1-x_t} - \frac{1}{2x_t} + 1,$$

so the probability that nobody invests in the good state is

$$G_{1,t} = e^{-\lambda \int_0^t (1 - F_1(x_s)) ds} = \frac{1 - x_t}{x_t} e^{\frac{1}{2x_t} - 1} \rightarrow 0,$$

and the probability that nobody invests in the bad state is

$$G_{0,t} = e^{\frac{1}{2x_t} - 1} \rightarrow e^{-\frac{1}{2}} \simeq 61.$$

8.4. Hazard Function and Hazard Ratio. For θ belonging to any ordered state space, define the hazard function h_θ , the likelihood ratio $l_{\theta,\theta'}$ and hazard ratio $m_{\theta,\theta'}^{-1}$ as

$$\begin{aligned} h_\theta & : = \frac{f_\theta(x)}{1 - F_\theta(x)}, \\ \theta > \theta' & : \quad l_{\theta,\theta'} := \frac{f_\theta(x)}{f_{\theta'}(x)}, \quad m_{\theta,\theta'}^{-1} := \frac{h_\theta(x)}{h_{\theta'}(x)}. \end{aligned}$$

Note that the properties of monotone likelihood ratio property (MLRP), the decreasing hazard function property (DHFP) and the increasing hazard ratio property (IHRP) can be stated as

$$\begin{aligned} \text{MLRP} & \iff \frac{\partial^2}{\partial\theta\partial x} \ln(f_\theta(x)) > 0, \\ \text{DHFP} & \iff \frac{\partial^2}{\partial\theta\partial x} \ln(1 - F_\theta(x)) > 0, \\ \text{IHRP} & \iff \frac{\partial^2}{\partial\theta\partial x} \ln\left(\frac{f_\theta(x)}{1 - F_\theta(x)}\right) > 0. \end{aligned}$$

Lemma 12.

$$\begin{aligned} \text{IHRP} + \text{DHFP} & \implies \text{MLRP}, \\ \text{MLRP} & \implies \text{DHFP}. \end{aligned}$$

Proof. The first implication is trivial as IHRP is by definition the first inequality and DHFP by definition the second inequality:

$$\frac{\partial^2}{\partial\theta\partial x} \ln(f_\theta(x)) > \frac{\partial^2}{\partial\theta\partial x} \ln(1 - F_\theta(x)) > 0.$$

For the second implication we need to show that

$$\theta > \theta' \quad \frac{f_\theta(x)}{\int_x^1 f_\theta(z) dz} < \frac{f_{\theta'}(x)}{\int_x^1 f_{\theta'}(z) dz},$$

which implies

$$\begin{aligned} \int_x^1 (f_{\theta'}(x) f_\theta(z) - f_\theta(x) f_{\theta'}(z)) dz & > 0, \\ \int_x^1 f_{\theta'}(x) f_{\theta'}(z) (l_{\theta,\theta'}(z) - l_{\theta,\theta'}(x)) dz & > 0, \end{aligned}$$

where the latter step is implied by MLRP. □

REFERENCES

- [1] Banerjee, A. (1992). “A Simple Model of Herd Behavior,” *Quarterly Journal of Economics*, **107**, 797–817.
- [2] Bikhchandani, S., Hirshleifer D. and I. Welch (1992). “A Theory of Fads, Fashion, Custom and Cultural Change as Informational Cascades,” *Journal of Political Economy*, **100**, 992–1026.
- [3] Callander S. and J. Hörner (2009). “The Wisdom of the Minority,” *Journal of Economic Theory*, **144**, 1421–39.
- [4] Çelen B. and S. Kariv (2004). “Observational Learning under Imperfect Information,” *Games and Economic Behavior*, **47**, 72–86.
- [5] Chamley C. and D. Gale (1994). “Information Revelation and Strategic Delay in a Model of Investment,” *Econometrica*, **62**, 1065–1085.
- [6] Chari, V. V. and P. J. Kehoe (2004). “Financial crisis as herds: Overturning the critiques,” *Journal of Economic Theory*, **119**, 128–150.
- [7] Elton, E., Gruber, M. and C. Blake (1996). “Survivorship Bias and Mutual Fund Performance,” *Review of Financial Studies*, **4**, 1097–1120.
- [8] Hendricks, K., Sorensen, A. and T. Wiseman (2009). “Observational Learning and Demand for Search Goods,” mimeo, University of Texas at Austin.
- [9] Milgrom, P. R. (1981). “Good News and Bad News: Representation Theorems and Applications,” *The Bell Journal of Economics*, **12**, 380–391.
- [10] Kalashnikov, V. V. and S. T. Rachev (1986). “Characterization of queuing models and their stability,” in *Probability Theory and Mathematical Statistics*, ed. Yu K. Prohorov et al., **2**, 37–53, VNU Science Press, Amsterdam.
- [11] Scargle, J. D. (2000). “Publication Bias: The “File-Drawer” Problem in Scientific Inference,” *Journal of Scientific Exploration*, **14**, 91–106.
- [12] Smith, L. and P. Sørensen (1997). “Rational Social Learning with Random Sampling,” mimeo, University of Michigan.
- [13] Smith, L. and P. Sørensen (2001). “Pathological Outcomes of Observational Learning,” *Econometrica*, **68**, 371–398.
- [14] Smith, L. and P. Sørensen (2008). “Informational Herding and Optimal Experimentation,” mimeo, University of Michigan.