Social Learning with Endogenous Information

S. Nageeb Ali[*]

May 23, 2014

Abstract

I introduce costly information acquisition to the observational learning framework. Each individual chooses from a set of actions, is uncertain about the state of the world, and wishes to take higher actions in the higher state. Individuals make choices in an exogenous order, observing all actions chosen by predecessors. The departure from standard herding models is that information is no longer free: instead, individuals choose from a compact set of signals and vary in their costs of acquiring a signal. I find conditions for complete learning in both discrete and continuous action spaces with and without costly information acquisition.

[*]University of California, San Diego (email: snali@ucsd.edu). I thank Jim Andreoni, Aislinn Bohren, Chris Chambers, Johannes Horner, Navin Kartik, Mark Machina, Manuel Mueller-Frank, Mark Muendler, Omar Nayeem, Paul Niehaus, Mallesh Pai, Andres Santos, Joel Sobel, Joel Watson, and Tom Wiseman. Erik Lillethun and John Rehbeck provided excellent research assistance. This work is financially supported by NSF grant SES-1127643.

1 Introduction

Motivating Question: The rational herding literature, initiated by Banerjee (1992) and Bikhchandani, Hirshleifer, and Welch (1992), offers an understanding of why individuals imitate others, and how each of a large group of individuals may make the wrong choice even if their collective information reveals all uncertainty. It is now well-understood that two critical frictions generate inefficiencies:

a) *Beliefs are unobservable to successors*: When choosing from a coarse set of actions, an individual's action may only partly reveal her posterior belief to successors. Were the information or beliefs of each player publicly observable or inferable by her successors, learning is necessarily complete.

b) *Private beliefs are bounded*: In a coarse action space, individuals herd with strictly positive probability on the wrong action if all signals induce bounded private beliefs, but such possibilities are precluded if signals may induce unbounded likelihood ratios (Smith and Sørensen 2000).

These two principles organize our understanding of social learning but emerge from frameworks in which information is acquired at no cost. Yet, rational herding is sometimes relevant because information is costly: each individual has a motive to not acquire information herself but to instead free ride on the wisdom of others. In this vein, Banerjee (1992) suggests the importance of understanding costly information acquisition in this setting:

> "The most serious departure of our model from reality is probably our assumption that signals to the agents are essentially free; a more realistic analysis would combine the question of incentives for obtaining these signals with the kinds of considerations we discuss...."

This paper explores the general implications of costly information acquisition for sequential decision making, and examines when learning is complete despite the cost of acquiring information.

Summary of Model and Results: I study a model with a binary state space where each player prefers higher actions in the higher state and lower actions in the lower state. Each player observes actions taken by predecessors and then chooses whether to acquire a signal from a compact set of signals, each of which is noisy. Each player has a type drawn from a rich type space that reflects the costs she faces in acquiring signals. Costs of information acquisition need not be ordered insofar as type θ may face a lower cost of acquiring signal s than type θ' but a higher cost of acquiring signal s'. The only assumption is that a player's cost of acquiring signals respects continuity in both the signal and the player's type.

Prior to describing the results that I obtain with costly information acquisition, it is useful to re-visit conditions needed for complete learning when information is exogenously provided. Say that a utility function is *responsive* if a player chooses different actions at different beliefs, and is *unresponsive at certainty* if for at least one state ω, the action that she chooses when she is certain of ω is also that which she also chooses whenever she is "almost-certain" that the state is ω. Unresponsiveness at certainty describes not only the standard observational learning environment (Banerjee 1992; Bikhchandani et al. 1992; Smith and Sørensen 2000) in which each undominated action is chosen over an interval of beliefs, but also the continuous investment model of Chari and Kehoe (2004) in which individuals pool on not investing whenever they are sufficiently pessimistic about the value of investing. By contrast, the paragon of responsiveness is the continuous action quadratic loss environment—studied by Lee (1993) and others—in which each player optimally chooses an action that matches her belief, and so even in a neighborhood of certainty, a player's action is different from that at absolute certainty.

I first prove a benchmark result that shows that responsiveness is the pivotal condition that characterizes whether unbounded signals are necessary for complete learning when information acquisition is costless. When the utility function is unresponsive at certainty, unbounded signals are needed to guarantee complete learning, thereby unifying the results from Smith and Sørensen (2000) and Chari and Kehoe (2004). By contrast, if the utility function is responsive, then even bounded signals suffice for learning. Although the distinction between coarse and rich action spaces is familiar, this benchmark result is not present in the prior literature, and formalizes the key distinction between coarse and rich action spaces. Responsiveness plays an intuitive role, corresponding exactly to the *inference problem* fundamental to observational learning: can a player's beliefs, and hence information, be perfectly inferred from her action? With perfect inference, observing a player's action is as good as observing her information, but otherwise, learning may be incomplete (if signals are bounded). When other features of the strategic environment—e.g. market prices (Glosten and Milgrom 1985; Avery and Zemsky 1998) or communication (Shiller 2000; Çelen, Kariv, and Schotter 2010)—ensure that signal realizations are publicly observable or inferable, responsiveness ceases to matter and complete learning can obtain even if the utility function is unresponsive at certainty.

My main result shows that *responsiveness* characterizes conditions for complete learning even when information is acquired at a cost. If players have responsive utility functions, then complete learning is guaranteed so long as some type can obtain an informative signal for free; were the cost of all informative signals to have a strictly positive uniform lower bound, then learning is incomplete. On the other hand, if players have utility functions that are unresponsive at certainty, learning is guaranteed if that free signal is also *unbounded*; were all unbounded signals to have some minimal cost, learning is incomplete.

Responsiveness is pivotal with and without costly information acquisition but for different reasons. Without costly information acquisition, the key issue is inference, as illustrated by the benchmark result, and responsiveness ceases to matter when all information is public. By contrast, when information is costly, the sufficient conditions for complete and incomplete learning are independent of whether past signal realizations are observable, and so responsiveness must play a different role when information is costly.[1] Responsive matters because it determines what kind of information is valuable.

To illustrate this mechanism, suppose that each player has a finite set of actions, and so the utility

[1]The observability of other peoples' information of course influences the speed of learning, or the probability with which inefficient herding materializes.

function is unresponsive at certainty. An individual values information acquisition only if signal realizations can influence her action. When she puts $1 - \varepsilon$ probability on a state, the cost of information acquisition that she is willing to bear is $O(\varepsilon)$. If unbounded signals have a strictly positive minimal cost, then the only signals that she is willing to purchase have bounded likelihood ratios. Information acquisition stops once these signals can no longer overturn the public belief. By contrast, a responsive decision maker values information at all interior beliefs, and so she is willing to acquire even signals that induce bounded beliefs. Because the value of information is vanishing as beliefs converge to the truth, information must also be available at arbitrarily low cost, but it need no longer be able to induce unbounded beliefs.

This result has implications for market settings in which prices are believed to overturn herds. Suppose that there were a competitive market maker, as in Glosten and Milgrom (1985) and Avery and Zemsky (1998), who sets prices over time so that all players can infer the information that would lead a player to buy or sell a unit of stock. Such an institution can aggregate information and induce complete learning when information is costless. However, when information is costly, matters are quite different: if the action space is finite (as in their model), then once the public history of actions and prices becomes very informative, all subsequent individuals lose their incentive to acquire information. The logic, intuitively, is a sequential manifestation of Grossman and Stiglitz (1980): when the history of market prices and actions conveys a lot of information, each subsequent investor may stop acquiring information, and thereby generate choices that lead to incomplete learning. When information is costly, even a competitive market maker may not engender complete learning even if he reveals the information of participants.

Related Literature: This note builds on prior papers that illuminate the role of costly information acquisition in social learning in several environments.

Burguet and Vives (2000) and Chamley (2004) study models of social learning in which players choose from a continuous action space with quadratic loss payoffs (as in Lee 1993), choose the precision of a normally distributed signal at a cost that is increasing in its precision, and are identical in their information acquisition costs. They find that complete learning obtains if the marginal cost of precision for a completely noisy signal is zero. This is a setting in which payoffs are responsive, players share the same costs of information acquisition, the state space is uncountable, and the information environment is restricted to normal signals (which means that all signals take unbounded likelihood ratios).

Hendricks, Sorensen, and Wiseman (2012) study a sequential search problem in which each player faces a choice to purchase or not purchase a product, freely obtains some private information about her value for the product, can choose to acquire at some cost a perfect signal about their value for the product, and observe only the fraction of people who have purchased the product. They prove that when the product's quality is low, learning is complete, but when it is high, learning may be incomplete. Surprisingly, they are able to derive a closed-form for the probability of an incorrect herd, which they use to generate comparative statics and compare to data from an online music market. Because players do not observe the completely ordered history, the sequence of public beliefs in their setting does not conform to a martingale, and hence, their analysis and setting is very different from that here.

Mueller-Frank and Pai (2013) also study a framework with costly search in which the payoff from each of a finite set of actions is distributed i.i.d. and when an individual "searches an action," he observes the value of that action perfectly without learning that of any other. Search is costly, and types are ordered in the cost of search. Since information is perfect, the only source of noise is that of the cost of acquiring information. They characterize conditions for complete learning: it obtains if one type can acquire *perfect information* at no cost at all, and thus, a positive measure of types obtain perfect information at cost no greater than ε.

These environments have substantially differed from each other, as have their analyses and approach, and none of these settings nests the canonical observational learning framework. Despite their conceptual similarities, the focus of Mueller-Frank and Pai (2013) on some type obtaining *perfect* signals at no cost in a *discrete* action space contrasts with the focus of Burguet and Vives (2000) and Chamley (2004) on *noisy* signals with a *continuous* action space. This contrast makes it difficult to see what might be needed for complete learning when information acquisition costs are considered in the canonical framework, and how these requirements vary with the action space. Motivated by the prior literature, this note focuses on the economic forces at the core of this problem without drawing upon a particular parametric structure for information. Building directly on the canonical observational social learning framework facilitates a more direct comparison with the standard setting in which information is free.

To tackle this problem, I take an elementary approach and focus on a simple logic: people are willing to buy information only if it can change their action with positive probability, and people's willingness to pay for information vanishes as their beliefs converge to certainty. Coupled with a martingale approach, this argument permits a combined treatment of perfect and imperfect signals, discrete and continuous action spaces, and allows for rich heterogeneity in the cost of information (e.g., types need not be ordered in their costs of acquiring information). Through the lens of responsiveness, the analysis facilitates a comparison to the case in which information is freely observed, and illustrates why the distinction between signal realizations being privately and publicly observed may be critical when information is freely available but is less so when information is costly. The framework also permits straightforward extensions to settings in which each player can dynamically choose how much information to acquire (as in Wald 1947) and settings in which players have heterogeneous priors and preferences.

2 Examples

Binary Actions: Each of a sequence of players $t = 1, 2, 3, \ldots$ chooses an action a from $\{0, 1\}$, and observes all prior actions. Payoffs depend on the realization of $\omega \in \{0, 1\}$, and each state is equally likely: a player obtains a payoff of 1 if her action matches the state of the world and 0 otherwise.

There are two sources of information. The first signal (or experiment), s_1, has discrete support $\left\{\frac{1}{4}, \frac{3}{4}\right\}$ and a conditional distribution $Pr(s_1 = p \mid \omega) = p\omega + (1-p)(1-\omega)$. The other signal, s_2, has continuous support $[0,1]$, and conditional density $f(p,\omega) = 2p\omega + 2(1-p)(1-\omega)$. Note that if a signal's realization is p, that is also the induced posterior belief that $\omega = 1$.

Were information acquisition costless, and individuals limited to signal s_1, a cascade inevitably forms, possibly on the wrong action. The impediment is that an individual's action may be insensitive to his information, and thus, those signal realizations are lost to subsequent players. By contrast, were individuals to obtain signal s_2 for free, complete learning is guaranteed.

Suppose that information acquisition is costly, and that each player i has a type θ_i drawn uniformly from $[0,1]$ that determines her cost of acquiring information. The cost of signal s_1 is θ_i whereas that of acquiring signal s_2 is $2\theta_i + c$ for $c \geq 0$. Thus, lower types find it less costly to obtain better information, and type 0 always obtains signal s_1 for free.

Suppose that $c > 0$: individuals may acquire information when uncertain about the state, but stop doing so once they are sufficiently certain. Specifically, suppose that the public belief that $\omega = 1$ strictly exceeds $\max\{3/4, 1-c\}$. Exceeding $3/4$ guarantees that signal s_1 has no value to all types: there is no reason to buy information that cannot influence one's choice. The second condition guarantees that signal s_2 has less value than its cost: were s_2 to perfectly reveal the state, its net value is still strictly less than c. Thus, learning stops once this cascade set is reached.

Now suppose that $c = 0$: for every interior public belief μ, type 0 has a strictly positive gain from acquiring signal s_2, and by continuity, so do all nearby types. This acquisition and accumulation of a signal that induces unbounded beliefs from a strictly positive measure of types guarantees that eventual beliefs almost-surely escape μ's neighborhood.

The distinction between whether an "unbounded" signal has a strictly positive minimum cost for all types, or is free for some type drives whether learning is complete. Note that in this example, it ceases to matter for the prospects of complete learning whether prior signal realizations are observed. Even if prior individuals' information / signal realizations are publicly observed, learning is incomplete whenever unbounded signals are costly for all types, and complete if it is free for some type. Incomplete learning is no longer driven by the logic that a coarse action space "filters" information.

Continuous Actions: Now suppose that the action space is $[0,1]$ and that a player's payoff from taking action a in state ω is $-(a-\omega)^2$. When information is exogenously generated, an individual equates her action to her posterior belief and thereby reveals her signal realization to those who follow. Yet, as we argued above, "filtering" is not relevant once information is costly.

Nevertheless, complete learning is guaranteed in the continuous action space even if $c > 0$. The difference is that signal s_1 is valuable at every interior belief. Thus, an interval of types neighboring $\theta = 0$ will acquire signal s_1, forcing social learning to settle on only the truth. The main result of this paper generalizes this distinction. I show that both with and without costly information acquisition, the necessary and sufficient conditions for complete learning can be connected to the *responsiveness* of individual utility functions to information.

3 Model

Actions and Payoffs: Each of an infinite sequence of players $t = 1, 2, 3, \ldots$ makes a single choice from $\mathcal{A}$, a compact subset of $\Re$, in which $|\mathcal{A}| \geq 2$. The payoff from action a in state of the world $\omega \in \Omega \equiv \{0, 1\}$ is $u(a, \omega)$, which is continuous in a for each ω. Suppose that no action is weakly dominated, and thus, if $u(a, 0) > u(a', 0)$, then $u(a, 1) < u(a', 1)$. Thus, we can order the actions from best to worst in state 0, and this order coincides with that from worst to best in state 1. Accordingly, there is no loss of generality in assuming that $u(a, 0)$ is strictly decreasing in a and $u(a, 1)$ is strictly increasing in a. Since $\mathcal{A}$ is compact, I denote by $\underline{a}$ and $\overline{a}$ the lowest and highest actions in $\mathcal{A}$. I bound the gap in payoffs between the best and worst actions in any state by $\gamma \equiv \max_{\omega \in \Omega} |u(\underline{a}, \omega) - u(\overline{a}, \omega)|$.

Beliefs and Information Acquisition: Each individual is uncertain about the state of the world. The common prior attributes probability $\pi \in (0, 1)$ to $\omega = 1$. Each individual can choose to acquire information, which involves running an experiment or observing a signal process. This choice is modeled as the selection of a random variable s from a compact set S, endowed with the weak topology. Every random variable s generates realizations described by a pair of cumulative distribution functions $(F_s(\cdot, 0), F_s(\cdot, 1))$ that is i.i.d. conditional on the state. The distributions $F_s(\cdot, 0)$ and $F_s(\cdot, 1)$ are mutually absolutely continuous—so no signal realization perfectly reveals or eliminates a state—and have common support $\Lambda(s)$. I normalize the realizations of signals by the posterior beliefs that they induce with a neutral prior, so that $\Lambda(s)$ is the range of private posterior beliefs that s may induce when an individual begins with a neutral prior.[2] S contains a *completely uninformative* signal process s_0 for which $\Lambda(s_0) = \{\pi\}$. Every other signal process s is *informative*, i.e., there exists $p \in [0, 1]$ such that $F_S(p, 0) \neq F_S(p, 1)$.[3] The closed convex hull of $\Lambda(s)$ is $\left[\underline{p}(s), \overline{p}(s)\right]$. I distinguish signals by how their realizations influence beliefs.

Definition 1. A signal process s is *bounded* if $0 < \underline{p}(s) \leq \overline{p}(s) < 1$ and *unbounded* if $\underline{p}(s) = 0$ and $\overline{p}(s) = 1$.

For a bounded signal process, there exists some $\underline{p} > 0$ such that after observing a realization of that signal process, almost-surely, the observer would ascribe at least $\underline{p}$ probability to each state if she begins with a neutral prior. By contrast, an unbounded signal process can push observers arbitrarily

[2] In other words, if F_s is differentiable at p, and $p \in \Lambda(s)$, $p = \frac{f_s(p,1)}{f_s(p,1)+f_s(p,0)}$. When an individual has prior π, the signal realization p generates a posterior likelihood ratio of $\frac{p\pi}{(1-p)(1-\pi)}$.

[3] By construction, each informative signal satisfies the monotone likelihood ratio property, and so $F_s(\cdot, 1)$ first order stochastically dominates $F_s(\cdot, 0)$.

close to putting probability 1 on a single state. However, unbounded signal processes differ from perfect information insofar as every signal process has a common support across states.[4]

A player's cost of observing a signal depends on the signal and her *information cost type*, θ, drawn from Θ, a compact metric space endowed with the Borel $\sigma-$algebra. Information cost types are drawn i.i.d. by a Borel measure ρ with support Θ_ρ. The costs of acquiring signal s for a player of type θ is $c(s,\theta) \geq 0$. As a normalization, let $c(s_0,\theta) = 0$ for all θ in Θ, and assume that c is continuous in $S \times \Theta$.

Timing and Observation: Each player observes actions of all predecessors but not their information. The public history before player $t > 1$ chooses to acquire information is $h^t \equiv (a_i)_{i=1,\ldots,t-1}$. After observing the public history and her type, a player chooses a signal s to observe and after observing its realization, chooses an action a. She earns a payoff of $u(a,\omega) - c(s,\theta)$ in state ω. I study Perfect Bayesian equilibria in which indifference is broken by the player choosing the lowest among all her optimal actions. In a PBE, σ, and history h^t, let $\mu_t(h^t) = Pr(\omega = 1|h^t)$ summarize the *public belief* after history h^t. Consider a set $\mathcal{H}$ of infinite length histories, and for such a history h_∞, let h^t_∞ be its truncation to actions in periods $1,\ldots,t-1$. For $\omega \in \{0,1\}$, let $\mathcal{H}_\omega$ denote the set of histories in $\mathcal{H}$ such that $\lim_{t\to\infty} \mu^t(h^t_\infty) = \omega$.

Definition 2. Learning is *complete* if for each $\omega \in \{0,1\}$, $Pr\left(h_\infty \in \mathcal{H}_\omega \mid \omega\right) = 1$, and otherwise, learning is *incomplete*.

4 When is Learning Complete?

For each belief μ, let $a^*(\mu)$ denote the (lowest) action that maximizes $\mu u(a,1) + (1-\mu)u(a,0)$.

Definition 3. The pair $(\mathcal{A},u)$ is **responsive** if $a^*(\mu) \neq a^*(\nu)$ whenever $\mu \neq \nu$, and otherwise, it is **unresponsive**. $(\mathcal{A},u)$ is **unresponsive at certainty** if there exists $\varepsilon > 0$ such that at least one of the following holds: (i) $a^*(\varepsilon) = \underline{a}$, (ii) $a^*(1-\varepsilon) = \overline{a}$.

Choice is responsive if any incremental change in an individual's belief induces her to tweak her action. Choice is unresponsive if the optimal action is constant over a set of beliefs, and unresponsive at certainty if that set is in a neighborhood of absolute certainty of some state ω.

The standard social learning framework with a coarse set of actions typifies unresponsiveness at certainty since any action that is optimal at certainty of a state remains optimal in a neighborhood of certainty. The paragon of responsiveness is quadratic loss utility with a rich action set described in Section 2. Notice that the individual is responsive because her payoffs are locally flat at her optimal action at each state; instead, if the action set were restricted to $\mathcal{A}' = [\varepsilon, 1]$ for some $\varepsilon > 0$, then the

[4] A common example of an unbounded noisy signal process is the normal signal $N(\omega,\sigma)$ in which $\sigma > 0$, but nevertheless, arbitrarily high realizations push beliefs arbitrarily in favor of $\omega = 1$ and arbitrarily low realizations push beliefs arbitrarily in favor of $\omega = 0$.

decision maker takes action ε whenever she believes that $\omega = 1$ with probability no greater than ε, and is therefore unresponsive at certainty. Such an environment—with continuous actions that are unresponsive at certainty—is studied in the investment game of Chari and Kehoe (2004).[5]

4.1 Benchmark Result: Costless Information

Suppose that information is exogenous: each player observes independent realizations from an informative signal s for free. To simplify analysis and exposition, I focus on continuous signals for which $F_s(\cdot, \omega)$ has a strict positive density on its support $\Lambda(s) = [\underline{p}(s), \overline{p}(s)]$.

Theorem 1. *The following describes conditions for complete learning.*

a) *If $(\mathcal{A}, u)$ is responsive, then learning is complete for every $\pi \in (0, 1)$.*
b) *If $(\mathcal{A}, u)$ is unresponsive at certainty, then for every prior $\pi \in (0, 1)$, learning is incomplete if and only if s is bounded.*
c) *If $(\mathcal{A}, u)$ is unresponsive, then there exists an open set of prior-signal combinations (π, s) such that learning is incomplete.*

Responsiveness plays an intuitive role: do actions fully reveal information? Actions are fully revealing when $(\mathcal{A}, u)$ is responsive, and this perpetual accumulation of information inexorably binds beliefs to the truth (a.s.). Otherwise, information is lost because an individual's action is a coarse signal of his beliefs. When this coarseness manifests at extreme beliefs, extreme signals are needed to shift beliefs and actions from the herd. Were each signal realization publicly observed, or perfectly inferable from market prices (Avery and Zemsky 1998) or communication (Shiller 2000; Çelen et al. 2010), responsiveness ceases to matter. Thus, the only role of responsiveness is to ensure perfect *inference*.

Theorem 1 connects prior results on herding, building directly on Lee (1993) and Smith and Sørensen (2000). While the distinction between coarse and fine action spaces is one that is often alluded to in the herding literature, the exact property that describes when learning is complete has not been characterized in past work (to my knowledge). Lee (1993) studies conditions on the action set that are necessary and sufficient for complete learning when the payoff function is quadratic loss and the signal has finite support. His analysis emphases the importance of "connectedness" in his setting. Theorem 1 extends his insight to all action spaces and utility functions, and articulates why his insight does not apply to continuous action environments, as in Chari and Kehoe (2004), that are unresponsive at certainty.[6]

[5] Their particular setting has $\mathcal{A} = [0, 1]$ and $u(a, \omega) = h(a)\omega + 1 - a$, in which h is strictly concave, $1 < h'(1) < h'(0) < \infty$, and $h(0) = 0$. Observe that the optimal action is 0 whenever the player believes that the probability that $\omega = 1$ is no more than $\frac{1}{h'(0)}$.

[6] Also closely related, Arieli and Mueller-Frank (2013) show in a more general framework than that here that $(\mathcal{A}, u)$ is responsive for all but a meager set of continuous utility functions so long as $\mathcal{A}$ contains no isolated points. My analysis complements their work insofar as they study the implications of responsiveness in portfolio optimization and repeated interaction in networks, and not observational learning with and without costly information acquisition.

4.2 Main Result: Costly Information

Now suppose information is costly. For any fixed budget on information acquisition, an individual may prefer to purchase the signal that has the greatest informativeness (Blackwell 1951, 1953) or accuracy (Lehmann 1988; Persico 2000). However, the results below on learning do not require specifying which signal is chosen by which type; instead, I can rely upon the ability of all types to choose her most preferred signal in equilibrium. I define conditions on the costs of information acquisition:

1. The cost function, c, satisfies **costly information** if there exists $\varepsilon > 0$ such that $c(s, \theta) \geq \varepsilon$ for every informative signal s and type θ in Θ_ρ.
2. The cost function, c, satisfies **virtually free information** if there exists a type θ in Θ_ρ such that $c(s, \theta) = 0$ for at least one informative signal s.
3. The cost function, c, satisfies **costly unbounded information** if there exist $\varepsilon > 0$ and cutoffs $0 < \underline{p} < \overline{p} < 1$ such that for any signal s and type θ in Θ_ρ for which $c(s, \theta) \leq \varepsilon$ then $\Lambda(s) \subset [\underline{p}, \overline{p}]$.
4. The cost function, c, satisfies **virtually free unbounded information** if there exists a type θ in Θ_ρ such that $c(s, \theta) = 0$ for at least one unbounded signal.

The virtually free conditions do not imply that all types can obtain the relevant signals for free, but that some type can do so, and hence, information may be free for only a measure 0 set of types. By continuity, a strictly positive fraction obtain such information at cost no greater than ε; when ε is small, this fraction of players who obtain the information cheaply may also be small. With virtually free unbounded information, the unbounded signal that is cheap may be extremely noisy (e.g., a normal signal centered at ω with high variance), but its tails can sway beliefs arbitrarily. I also observe that because $c(s, \theta)$ is continuous, information can be costly only if s_0 is an isolated point in S.

Theorem 2. *The following describes conditions on c that guarantee complete and incomplete learning in responsive and unresponsive environments.*

a) *Suppose $(\mathcal{A}, u)$ is responsive. Then for every $\pi \in (0, 1)$, learning is complete if information is virtually free, and incomplete if information is costly.*
b) *Suppose $(\mathcal{A}, u)$ is unresponsive at certainty. Then for every prior $\pi \in (0, 1)$, learning is complete if unbounded information is virtually free, and incomplete if unbounded information is costly.*
c) *Suppose $(\mathcal{A}, u)$ is unresponsive. Then there exists an open set of priors such that learning is incomplete even if information is virtually free.*

Suppose that for all but a $\rho-$measure 0 set of types, $c(s, \theta) > 0$ for every informative signal s. Then the above characterization is invariant to the observability of past signal realizations.

The similarity of Theorems 1 and 2 may indicate that the mechanism by which responsiveness facilitates learning is identical, namely that of *inference*. But the inability to observe previous signals is not the impediment to complete learning: responsiveness continues to matter even if all prior signal realizations are publicly observed, so long as information cannot freely accumulate.

The principle that guides learning is that only information that can influence behavior is worth buying. If the environment is unresponsive at certainty, then when a player places $1-\varepsilon$ probability on a state, her maximum willingness to pay for information is $O(\varepsilon)$. Yet, if unbounded signals have a minimal cost that exceeds this "budget," unbounded signals are never purchased. However, bounded signals would not influence one's action once beliefs are sufficiently extreme, and are therefore worthless. By contrast, if an unbounded signal is free for any type, then sufficiently close types have a strict incentive to acquire it. Their actions reflect such information, and can overturn any incorrect herd.[7]

When the environment is responsive, milder conditions on the cost of information guarantee learning. At an interior belief, a player always values information even if it induces bounded likelihood ratios. The persistent injection of information and its revelation through smooth actions guarantee complete learning. It is only if information is costly that learning ends once sufficient confidence is reached.

As mentioned earlier, these results have implications on how market and financial interactions influence herding behavior. When information is costly, whether a player's investment and trading decisions are responsive, as in Lee (1993), or unresponsive at certainty, as in Chari and Kehoe (2004), influence whether complete learning emerges, regardless of whether there is a competitive market maker or external financial institution that aggregates information through prices.

Remark 1. *There is a gap between virtually free unbounded information and costly unbounded information insofar as they omit the setting in which unbounded private signals may be acquired by type* $\underline{\theta}$ *at arbitrarily low cost but not for free. Whether learning obtains in such cases depends upon the specific parametric structure for costly information acquisition, and details of the action and state space.*[8] *The challenge of deriving general results is that the value of information acquisition need not be concave (Radner and Stiglitz 1984; Chade and Schlee 2002; Keppo, Moscarini, and Smith 2008), and so a player's willingness to buy "a small bit of information" depends upon how that small bit is modeled.*[9] *Thus, I conjecture that different settings generate different conclusions for when inexpensive but not free signals engender complete learning. For the same reason, I do not analyze the speed of learning or the probability with which incorrect herds form since these computations hinge on the specific parametric*

[7] The discussion so far pertains to the convergence of beliefs, and not actions. Suppose that learning is incomplete, and the prior is such that players do not begin by herding on any action at the outset. Theorem 2 establishes that beliefs in the "learning phase" remain interior, and therefore, individuals may herd on the wrong action with strictly positive probability. However, learning may never cease at any finite history: there may always be a vanishing small probability of someone rationally purchasing a signal that could overturn the herd.

[8] Burguet and Vives (2000) and Chamley (2004) study this question in the context of normally distributed signals in a continuous action and state environment.

[9] As in Keppo, Moscarini, and Smith (2008), consider an example with normal signals in which the cost of a signal centered at ω and with variance $1/\gamma$ is $c(\gamma,\theta)$. Based on their analysis, it appears that it would be insufficient for complete learning for $c'(0,\theta)=0$ for a strictly positive measure of types since the marginal value of information is 0 at $\gamma=0$. However, if if $c'(0,\theta)=c''(0,\theta)=0$ for a strictly positive measure of types, then there is a strictly positive level of information acquisition at every μ in $\left(\frac{1}{2},1\right)$, thereby guaranteeing complete learning in the limit.

structure of information acquisition.

Remark 2. *Another important benchmark is that of "externalities": how much inefficiency is caused by players not internalizing the benefits that others would have from their information? This benchmark corresponds to the same decision maker choosing actions at $t = 1, 2, 3, \ldots$ and with costs of acquiring information in period t drawn i.i.d. according to the measure ρ. When would she choose to completely learn the state of the world? It is straightforward to establish that if she is perfectly patient, then she would learn under a wide range of cost conditions, exactly analogous to the complete learning results that emerge in bandit environments with perfectly patient players (Aghion, Bolton, Harris, and Jullien 1991; Ali 2011). Thus, if unbounded information is costly, the population's learning is incomplete relative to a social planner who weighs every agent equally.*[10]

4.3 Proof of Theorem 2

Let $B(\mu, p)$ be the Bayesian posterior when a player has a prior μ and observes a signal realization p. For each state ω, consider the likelihood ratio with respect to the other state: $l_1^t(h^t) = \frac{1-\mu^t(h^t)}{\mu^t(h^t)}$ and $l_0^t(h^t) = 1/l_1^t(h^t)$. I treat $\langle l_i^t(\cdot)\rangle_{t=1}^{\infty}$ as a stochastic process, and it is straightforward to see that conditioning on $\omega = i$, it is a non-negative martingale. The Martingale Convergence Theorem ensures that it converges almost-surely to a random variable l_i^{∞} whose support is in $[0, \infty)$.

Case 1: $(\mathcal{A}, u)$ is Responsive: Suppose that information is virtually free, and let (s, θ) be an informative signal-type pair for which $c(s, \theta) = 0$. Observe that when the public belief is $\mu \in (0, 1)$, the value of signal s for a responsive individual is

$$V(s, \mu) \equiv \mu \int_0^1 u(a^*(B(\mu, p)), 1)dF_s(p, 1) + (1 - \mu) \int_0^1 u(a^*(B(\mu, p)), 0)dF_s(p, 0)$$

whereas the value of not acquiring information is

$$V(s_0, \mu) \equiv \mu \int_0^1 u(a^*(\mu), 1)dF_s(p, 1) + (1 - \mu) \int_0^1 u(a^*(\mu), 0)dF_s(p, 0).$$

Since s is informative, $B(\mu, p)$ differs from μ with strictly positive probability, and because u is responsive, $a^*(B(\mu, p)) \neq a^*(\mu)$. I argue by revealed preference that $V(s, \mu)$ is strictly higher than $V(s_0, \mu)$: the decision maker can upon choosing signal s guarantee herself the payoff of $V(s_0, \mu)$ by choosing $a^*(\mu)$ for every realization of p. Because he deviates from this plan with strictly positive probability implies that he is strictly better off. Since $c(s, \theta) = 0$, type θ has a strict incentive to acquire signal s at the public

[10] If the social planner is not perfectly patient, comparisons are more subtle. If unbounded information is costly, she would not completely learn if $(\mathcal{A}, u)$ is unresponsive at certainty since the cost of information acquisition at some point outweighs the discounted continuation value from learning. If unbounded information is virtually free, then she may completely learn if she can obtain unbounded information at low cost sufficiently often (with respect to her patience).

belief μ, and by continuity of $c(\cdot,\cdot)$, there exists δ such that for all $\theta' \in N_\delta(\theta)$, a player of type θ' strictly prefers signal s to signal s_0.

I use this to now prove that learning is complete, i.e., that conditional on $\omega = i$, the support of l_i^∞ is $\{0\}$. Suppose towards a contradiction that $l > 0$ is in the support of l_i^∞. Consider any $\tilde{l} \in (l-\varepsilon, l+\varepsilon)$, in which $\varepsilon < l$. By above, when the public belief is $\tilde{l}$, a strictly positive measure of types obtain information and take actions that perfectly reveal their signal realizations. The Strong Law of Large Numbers implies that the probability that the public likelihood ratio remains perpetually in $(l-\varepsilon, l+\varepsilon)$ is 0, which implies that the support of l_i^∞ is $\{0\}$.

Suppose instead that information is costly, and let $\varepsilon < \gamma$ be such that if $c(s,\theta) < \varepsilon$, then s is uninformative. Consider beliefs $\mu^* = 1 - \frac{\varepsilon}{\gamma}$ and $\mu_* = \frac{\varepsilon}{\gamma}$. If the public belief ever crosses outside of $[\mu_*, \mu^*]$, no player is willing to budget even ε for even a perfect signal given that the maximal gain from learning the true state is γ. These are cascade regions and if learning does not stop earlier, it stops when μ^t enters $[0,\mu_*] \cup [\mu^*, 1]$. I argue that these cascade regions generate incomplete learning. Suppose that $\omega = i$, and towards a contradiction that learning is complete. Then, l_i^t cannot exceed $\frac{\gamma-\varepsilon}{\varepsilon}$ since players stop acquiring information once it does so. The Bounded Convergence Theorem then ensures that $E[l_i^\infty] = \lim_{t\to\infty} E[l_i^t]$, which because $\langle l_i^t(\cdot)\rangle_{t=1}^\infty$ is a martingale, coincides with $l_i^0 > 0$. This yields a contradiction to the claim that $Pr(l_i^\infty = 0 \mid \omega = i) = 1$.

Case 2: $(\mathcal{A}, u)$ **is Unresponsive at Certainty:** Suppose that unbounded signals are virtually free. Let (s,θ) be an unbounded signal-type pair for which $c(s,\theta) = 0$. Mirroring the argument of Case 1, the value of signal s for a responsive individual is

$$V(s,\mu) \equiv \mu \int_0^1 u(a^*(B(\mu,p)),1)dF_s(p,1) + (1-\mu)\int_0^1 u(a^*(B(\mu,p)),0)dF_s(p,0)$$

Observe that so long as the public belief μ is in $(0,1)$, there exists a set of signal realizations $\tilde{P}$ such that for every $p \in \tilde{P}$, $a^*(B(\mu,p)) \neq a^*(\mu)$ and $F_s(\tilde{P},\omega) > 0$ for every $\omega \in \Omega$. Therefore, it follows that by revealed preference, $V(s,\mu) > V(s_0,\mu)$, and so type θ strictly prefers to acquire signal s than to acquire no information. Therefore, each type in a sufficiently small neighborhood of type θ chooses to acquire an informative signal that influences her action with positive probability. Since the posterior beliefs after these actions depart from μ with strictly positive probability, it follows that when $\omega = i$, $\langle l_i^t\rangle$ converges almost-surely to the random variable l_i^∞ with support $\{0\}$.

Suppose that unbounded signals are costly. Consider $\varepsilon, \underline{p}, \overline{p}$ such that $c(s,\theta) < \varepsilon$ implies $\Lambda(s)$ is a subset of $[\underline{p}, \overline{p}]$. Let $\underline{\mu}$ and $\overline{\mu}$ be the highest and lowest beliefs such that $a^*(\underline{\mu}) = \underline{a}$ and $a^*(\overline{\mu}) = \overline{a}$. Since behavior is unresponsive at certainty, either $\underline{\mu} > 0$ or $\overline{\mu} < 1$ or both. Define

$$\mu_+ = \min\left\{\frac{\varepsilon}{\gamma}, \frac{\underline{\mu}(1-\overline{p})}{\underline{\mu}+\overline{p}-2\underline{\mu}\overline{p}}\right\},$$
$$\mu^+ = \max\left\{1-\frac{\varepsilon}{\gamma}, \frac{\overline{\mu}(1-\underline{p})}{\overline{\mu}+\underline{p}-2\overline{\mu}\underline{p}}\right\}$$

Observe that for a public belief $\mu \in [0, \mu_+] \cup [\mu^+, 1]$, the most that a player would pay for a perfect signal is ε. So if a player acquires information at all, she acquires a bounded signal s in which $\Lambda(s)$ is a subset of $[\underline{p}, \overline{p}]$. However, beliefs are sufficiently concentrated around a state i that even if she were to receive the most negative signal about that state, she still picks the optimal action associated to $\omega = i$. Since her actions are unaffected by this information, she has no incentive to acquire such signals either. Therefore, once the public belief enters $[0, \mu_+] \cup [\mu^+, 1]$, no player has any incentive to acquire information.

Towards clarifying the impact on learning, suppose that $\omega = 1$ and $\mu_+ > 0$. If $\pi \leq \mu_+$, then we are done since no player ever acquires information. Otherwise, if learning has not stopped before, it does so once $\langle l_1^t \rangle$ exceeds $\frac{1-\mu_+}{\mu_+}$. Suppose towards a contradiction that learning is complete at $\omega = 1$. It must be that $l_1^t < \frac{1-\mu_+}{\mu_+}$ for every t since otherwise players would stop acquiring information. The Bounded Convergence Theorem implies that $E[l_1^\infty] = \lim_{t\to\infty} E[l_1^t] = l_1^0 > 0$, which contradicts $Pr(l_1^\infty = 0 \mid \omega = 1) = 1$. An analogous argument applies if $\omega = 0$ and $\mu^+ < 1$.

Case 3: $(\mathcal{A}, u)$ **is Unresponsive:** Suppose that $(\mathcal{A}, u)$ is unresponsive: trivially, consider the setting described in Theorem 1, letting the signal described therein be the only informative signal, available for free to a null set of types, and at price $\varepsilon > 0$ to all others.

Note that identical arguments apply if past signal realizations are observed and that for all but a $\rho-$measure 0 set of types, $c(s, \theta) > 0$ for every informative signal s.

5 Extensions

5.1 Multiple States

Theorem 2 generalizes to a richer finite state space with more than two states and a general (multi-dimensional) action space for the case in which $(\mathcal{A}, u)$ is responsive. When $(\mathcal{A}, u)$ is unresponsive at certainty, the challenge with multiple states is that without greater structure, players cannot order actions or the set of signal realizations that induce those actions. This issue does not arise if $\mathcal{A}$ is a compact subset of $\Re$ and standard monotonicity assumptions are imposed on payoffs and information. Suppose that $u(a, \omega)$ satisfies the single-crossing property so that higher actions are preferred in higher states, and that the realizations of informative signals can be ordered according to the monotone likelihood ratio property so that higher realizations are relatively more likely in higher states. In such cases, higher signal realizations induce higher actions (Athey 2002). With this structure, Theorem 2's characterization of learning when $(\mathcal{A}, u)$ is unresponsive at certainty extends to finite state spaces with more than two states and a one dimensional action space.

5.2 Sequential Information Acquisition

I model information acquisition as a single-stage process but can extend results to a sequential information acquisition environment (e.g. Wald 1947) in which each individual can choose how long to acquire information.[11] Suppose that each player can acquire multiple signals sequentially, conditioning the acquisition of a signal on the realizations of signals she has already acquired. Suppose that $(\mathcal{A}, u)$ is unresponsive at certainty, and that bounded information is virtually free but unbounded information is costly. One may wish to limit players to a uniform upper bound on the number of signals each can acquire, since otherwise a player who finds bounded information to be free can acquire that signal a countless number of times so as to have unbounded beliefs at no cost. With this restriction on the strategy space, every strategy in which only bounded signals are consulted induce bounded private beliefs. Once the public belief is sufficiently extreme, no player will find it sequentially rational to acquire an unbounded signal at costs of $\varepsilon > 0$ or higher, and beliefs cannot be overturned by any feasible strategy that acquires only bounded signals. Thus, Theorem 2 applies when we restrict information acquisition strategies to a uniform upper bound on the number of signals acquired.

5.3 Heterogeneous Priors and Preferences

I have focused on a setting in which all players share the same payoffs. The results extend seamlessly to including dominant strategy or "crazy" types whose preferred action is independent of ω. A more interesting extension is that in which players have "monotone" private preferences (every player prefers weakly higher actions in the higher state and weakly lower actions in the lower state), as in Goeree, Palfrey, and Rogers (2006) and Wiseman (2008). A diversity of preferences, or beliefs, in this case can lead to complete learning even if each type of player is unresponsive at certainty.

Let me model this issue via the lens of heterogeneous beliefs.[12] Suppose that in addition to player i's cost type, she has a prior belief $\pi_i \in [0,1]$ that the state of the world is $\omega = 1$, and a player's belief is independent of her cost type. Suppose π_i is distributed independently according to an atomless cdf H that has a strictly positive density on its support $[\underline{\pi}, \overline{\pi}]$, and this is common knowledge. Consider a case in which social learning fails when players share common priors and preferences: each player has the same payoff function $u(a, \omega)$ that is is unresponsive at certainty with $\underline{\mu} > 0$ and $\overline{\mu} < 1$. Finally, let unbounded information be costly but bounded information be virtually free.

[11] I thank Thomas Wiseman for suggesting this extension.

[12] There is an equivalent model in which players share common priors and their payoffs for each action are augmented by individual private shocks.

Theorem 3. *Learning is complete if* $[\underline{\pi}, \overline{\pi}] = [0,1]$ *and incomplete if* $\underline{\pi} > 0$ *and* $\overline{\pi} < 1$. .

A rich diversity of beliefs (or preferences) of this form facilitates learning even if $(\mathcal{A}, u)$ is unresponsive at certainty because at every public belief, there is always a player who values any informative signal, even if it is bounded. As these individuals acquire and inject information, the public belief converges almost-surely to almost all types being close to certainty of the state.

I study heterogeneous priors and preferences when players share the same order on actions, since such heterogeneity changes the results of the prior section. Dispensing with this form of monotonicity generates the possibility for confounded learning (Smith and Sørensen 2000) even with costless information, and adding costs to information only cements this possibility.

5.4 Costly Observability of History

I assume that the public history is freely observable while acquiring private information is costly. In some settings, individuals may find it costly to observe the choices of all predecessors, and may instead choose to observe a subset.[13] While a complete analysis of the tradeoff between costly observations of the public history and acquiring independent information is beyond my scope here, particular examples highlight how introducing partial observability of history can foster social learning. Suppose that the majority of players face costs as modeled in this paper and observe the public history for free, but that a fraction find it too costly to observe any aspect of the public history. If bounded information is virtually free, and the common prior is not extreme, a positive measure of these individuals would acquire independent information about the state so as to improve their decision making. In doing so, they inject information into the public history, and thereby foster learning by others. Thus, learning may be complete even if unbounded information is costly and preferences are unresponsive at certainty.[14]

A Appendix

Proof of Theorem 1. For a measurable set of actions A, let

$$P(A, \mu) \equiv \{p \in \Lambda(s) : a^*(B(\mu, p)) \in A\},$$
$$\alpha(A, \mu, \omega) \equiv \int_{P(A,\mu)} dF_s(p, \omega).$$

An action a is in the support, $\overline{\mathcal{A}}(\mu)$, if for every $\varepsilon > 0$, $\alpha((a-\varepsilon, a+\varepsilon), \mu, \omega) > 0$ for every ω. Let $p^*(a, \mu)$ and $p_*(a, \mu)$ be the *sup* and *inf* of $P(\{a\}, \mu)$ respectively. Since $F_s(\cdot, \omega)$ is continuously

[13] Such a setting is studied by Çelen and Kariv (2005), Kultti and Miettinen (2006), Smith and Sorensen (2008), Monzón and Rapp (2009), Acemoglu, Dahleh, Lobel, and Ozdaglar (2011), and Çelen and Hyndman (2012).

[14] An analogue of this with costless information is established in Acemoglu, Dahleh, Lobel, and Ozdaglar (2011) and Bohren (2013).

differentiable, it follows that for every measurable subset A, $\alpha(A,\mu,\omega)$ is continuous in μ. Let $\beta(a,\mu)$ be the updated public belief when action $a \in \overline{\mathcal{A}}(\mu)$ is chosen at public belief μ; for every action a such that $p^*(a,\mu) \neq \underline{p}(s)$, $\beta(a,\mu)$ is continuous in μ. The cascade set of beliefs is

$$\mathcal{C}(s) \equiv \bigcup_{a\in\mathcal{A}} \{\mu \in [0,1] : P(\{a\},\mu) = \Lambda(s)\}.$$

The lemma below relates the public belief to actions that it can induce.

Lemma 1. *A public belief $\mu \in \mathcal{C}(s)$ if and only if $\alpha(A,\mu,0) = \alpha(A,\mu,1)$ for every measurable A.*

Proof. If $\mu \in \mathcal{C}(s)$, it trivially follows that there exists an action a such that $F_s(p^*(a,\mu),\omega) - F_s(p_*(a,\mu),\omega) = 1$ for each ω. Suppose that $\mu \notin \mathcal{C}(s)$. Then there exists an action $\tilde{a}$ such that $F_s(p^*(\tilde{a},\mu),\omega) \in (0,1)$, and consider the set of actions $[\underline{\theta},\tilde{a}]$: by Lemma A.1 of Smith and Sørensen (2000), it follows that $F_s(p^*(\tilde{a},\mu),1) < F_s(p^*(\tilde{a},\mu),0)$. □

For each state ω, consider the likelihood ratio with respect to the other state: $l_1^t(h^t) = \frac{1-\mu^t(h^t)}{\mu^t(h^t)}$ and $l_0^t(h^t) = 1/l_1^t(h^t)$. I treat $\langle l_i^t(\cdot)\rangle_{t=1}^{\infty}$ as a stochastic process, and it is straightforward to see that it is a non-negative martingale conditioning on $\omega = i$. The Martingale Convergence Theorem ensures that it converges almost-surely to a random variable l_i^{∞} whose support is in $[0,\infty)$.

Lemma 2. *Conditional on $\omega = i$, the likelihood ratio l is in the support of l_i^{∞} implies that $\frac{1}{1+l}$ is a subset of $\mathcal{C}(s)$ if $i = 1$, and $\frac{l}{1+l}$ is a subset of $\mathcal{C}(s)$ if $i = 0$.*

Proof. Suppose towards a contradiction that the support of l_1^{∞} includes l such that $\mu = \frac{1}{1+l}$ is not in $\mathcal{C}(s)$. Consider action $\tilde{a}$ such that $F_s(p^*(\tilde{a},\mu),\omega) \in (0,1)$, and $\beta(\tilde{a},\mu) < \mu$; such an action must exist by Lemma 1 and the law of iterated expectations. By monotonicity, for each $a \in [\underline{a},\tilde{a}] \cap \overline{\mathcal{A}}(\mu)$, $|\beta(a,\mu) - \mu| \geq |\beta(\tilde{a},\mu) - \mu|$. Let $\tilde{\alpha} = \frac{\alpha([\underline{a},\tilde{a}],\mu,1)}{2}$. Since $\alpha(\cdot,\mu,\omega)$ and $\beta(\cdot,\mu)$ are continuous in μ, it follows that there exists $\varepsilon > 0$ such that for every $\mu' \in (\mu-\varepsilon,\mu+\varepsilon)$, the updated belief is in $(\mu-\varepsilon,\mu+\varepsilon)$ with probability at most $1-\tilde{\alpha}$, yielding a contradiction. An analogous argument applies for l_0^{∞}. □

Now suppose $(\mathcal{A},u)$ is responsive. Then $\mathcal{C}(s) = \{0,1\}$, and since the martingale convergence theorem ensures that l_i^{∞} has support in $[0,\infty)$, Lemma 2 implies that $Pr(l_i^{\infty} = 0 \mid \omega = i) = 1$.

Suppose that $(\mathcal{A},u)$ is unresponsive at certainty. For an unbounded signal s, $\mathcal{C}(s) = \{0,1\}$, and so as above, $Pr(l_i^{\infty} = 0 \mid \omega = i) = 1$. Now suppose that s is bounded. Let $\underline{\mu}$ and $\overline{\mu}$ be the highest and lowest beliefs respectively such that $a^*(\underline{\mu}) = \underline{a}$ and $a^*(\overline{\mu}) = \overline{a}$. We consider the following cases below.

1. Suppose that $0 < \underline{\mu} < \overline{\mu} < 1$. Define

$$l_* \equiv \left(\frac{1-\underline{\mu}}{\underline{\mu}}\right)\left(\frac{1-\underline{p}(s)}{\underline{p}(s)}\right), l^* \equiv \left(\frac{1-\overline{\mu}}{\overline{\mu}}\right)\left(\frac{1-\overline{p}(s)}{\overline{p}(s)}\right).$$

It follows that once l_1^t enters $[0, l^*] \cup [l_*, \infty]$, all subsequent individuals choose the same action regardless of their signal realization. Learning is incomplete in both $\omega = 0, 1$.

2. Suppose that $\underline{\mu} > 0$ but $\overline{\mu} = 1$. To show that learning is incomplete with strictly positive probability, it suffices to establish that there exists l such that $Pr(l_1^\infty > l \mid \omega = 1) > 0$. Suppose otherwise. Then, $E[l_1^\infty \mid \omega = 1] = 0$. However, it must also be that for every t, $Pr(l_1^t < l_*) = 1$ since otherwise, there is positive probability that the public likelihood ratio converges to a positive number. Since l_1^t is dominated by l_*, we can apply the Bounded Convergence Theorem to establish that $E[l_1^\infty \mid \omega = 1] = \lim_{t\to\infty} E[l_1^t \mid \omega = 1]$, which equals $l_1^0 > 0$ since $\langle l_1^t \rangle$ is a martingale, yielding a contradiction.
3. Suppose that $\underline{\mu} = 0$ but $\overline{\mu} < 1$. The argument is analogous to that above by considering the stochastic process $\langle l_0^t \rangle$.

Finally, suppose that $(\mathcal{A}, u)$ is unresponsive. Consider an action a and a range of beliefs $[\underline{\mu}, \overline{\mu})$ such that for every $\mu \in [\underline{\mu}, \overline{\mu})$, $a^*(\mu) = a$. Consider any combination of prior-signal combination (π, s) such that $\pi \in [\underline{\mu}, \overline{\mu})$, and $B(\pi, \overline{p}(s)) < \overline{\mu}$ and $B(\pi, \underline{p}(s)) > \underline{\mu}$. For such combinations, every individual chooses action a regardless of her signal realization. □

Proof of Theorem 3. Normalize the public belief $\mu^t(h^t)$ as if the prior were neutral, and let $B(\mu, p, \pi_i)$ denote the posterior belief of player i when the public belief is μ, she observes signal realization p, and her prior belief is π_i. First, suppose that $\underline{\pi} = 0$ and $\overline{\pi} = 1$, and a bounded signal s such that $c(s, \theta) = 0$, and $\Lambda(s) \subset [\underline{p}, \overline{p}]$ including the end-points. At a public belief μ, consider a prior π_i that satisfies

$$\frac{\overline{\mu}}{1-\overline{\mu}} \left(\frac{\mu}{1-\mu}\right)^{-1} \left(\frac{\overline{p}}{1-\overline{p}}\right)^{-1} < \frac{\pi_i}{1-\pi_i} < \frac{\overline{\mu}}{1-\overline{\mu}} \left(\frac{\mu}{1-\mu}\right)^{-1} \left(\frac{\underline{p}}{1-\underline{p}}\right)^{-1}.$$

Observe that a player with prior π_i has a strictly positive value for the signal s since she chooses different actions in every small neighborhood of $\overline{p}$ than she does in every small neighborhood of $\underline{p}$. Since there is a strictly positive measure of players with such priors, there exists for every $\varepsilon > 0$, a strictly positive measure of types (θ_i, π_i) that strictly prefer to obtain the bounded signal s at cost no less than ε to no information at all. Since the public belief departs from μ with strictly positive probability, it follows that l_i^∞ has support $\{0\}$. By contrast if $\underline{\pi} > 0$ and $\overline{\pi} < 1$, this case of Theorem 2 is easily extended by re-defining μ_+ and μ^+ to account for the range of possible priors. □

References

Acemoglu, D., M. A. Dahleh, I. Lobel, and A. Ozdaglar (2011): "Bayesian learning in social networks," *The Review of Economic Studies*, 78, 1201–1236.

Aghion, P., P. Bolton, C. Harris, and B. Jullien (1991): "Optimal learning by experimentation," *The review of economic studies*, 58, 621–654.

Ali, S. N. (2011): "Learning self-control," *The Quarterly Journal of Economics*, 126, 857–893.

Arieli, I. and M. Mueller-Frank (2013): "Inferring Beliefs From Actions," Department of Economics, University of Oxford.

Athey, S. (2002): "Monotone comparative statics under uncertainty," *The Quarterly Journal of Economics*, 117, 187–223.

Avery, C. and P. Zemsky (1998): "Multi-Dimensional Uncertainty and Herd Behavior in Financial Markets," *American Economic Review*, 88, 724–48.

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

Bikhchandani, S., D. Hirshleifer, and I. Welch (1992): "A theory of fads, fashion, custom, and cultural changes as informational cascades," *Journal of Political Economy*, 100, 992–1026.

Blackwell, D. (1951): "Comparison of Experiments," in *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability*, ed. by J. Neyman, Berkeley, CA: University of California Press, 93–102.

——— (1953): "Equivalent comparisons of experiments," *The Annals of Mathematical Statistics*, 24, 265–272.

Bohren, J. A. (2013): "Informational Herding with Model Misspecification," University of Pennsylvania.

Burguet, R. and X. Vives (2000): "Social learning and costly information acquisition," *Economic Theory*, 15, 185–205.

Çelen, B. and K. Hyndman (2012): "Social learning through endogenous information acquisition: an experiment," *Management Science*, 58, 1525–1548.

Çelen, B. and S. Kariv (2005): "An experimental test of observational learning under imperfect information," *Economic Theory*, 26, 677–699.

Çelen, B., S. Kariv, and A. Schotter (2010): "An experimental test of advice and social learning," *Management Science*, 56, 1687–1701.

Chade, H. and E. Schlee (2002): "Another look at the Radner–Stiglitz nonconcavity in the value of information," *Journal of Economic Theory*, 107, 421–452.

Chamley, C. P. (2004): *Rational Herds*, Cambridge, UK: Cambridge University Press.

Chari, V. V. and P. J. Kehoe (2004): "Financial crises as herds: overturning the critiques," *Journal of Economic Theory*, 119, 128–150.

Glosten, L. R. and P. R. Milgrom (1985): "Bid, Ask and Transaction Prices in a Specialist Market with Heterogeneously Informed Traders," *Journal of Financial Economics*, 14, 71–100.

Goeree, J. K., T. R. Palfrey, and B. W. Rogers (2006): "Social learning with private and common values," *Economic Theory*, 28, 245–264.

Grossman, S. J. and J. E. Stiglitz (1980): "On the impossibility of informationally efficient markets," *American Economic Review*, 70, 393–408.

Hendricks, K., A. Sorensen, and T. Wiseman (2012): "Observational learning and demand for search goods," *American Economic Journal: Microeconomics*, 4, 1–31.

Keppo, J., G. Moscarini, and L. Smith (2008): "The demand for information: More heat than light," *Journal of Economic Theory*, 138, 21–50.

Kultti, K. and P. Miettinen (2006): "Herding with costly information," *International Game Theory Review*, 8, 21–31.

Lee, I. H. (1993): "On the convergence of informational cascades," *Journal of Economic Theory*, 61, 395–411.

Lehmann, E. (1988): "Comparing location experiments," *The Annals of Statistics*, 16, 521–533.

Monzón, I. and M. Rapp (2009): "Observational Learning with Position Uncertainty," Tech. rep., working paper.

Mueller-Frank, M. and M. M. Pai (2013): "Social Search," Department of Economics, University of Pennsylvania.

Persico, N. (2000): "Information acquisition in auctions," *Econometrica*, 68, 135–148.

Radner, R. and J. Stiglitz (1984): "A Nonconcavity in the Value of Information," *Bayesian models in economic theory*, 5, 33–52.

Shiller, R. J. (2000): "Conversation, Information, and Herd Behavior," *American Economic Review*, 85, 181–185.

Smith, L. and P. Sørensen (2000): "Pathological Outcomes of Observational Learning," *Econometrica*, 68, 371–398.

Smith, L. and P. Sorensen (2008): "Rational social learning by random sampling," *Available at SSRN 1138095*.

Wald, A. (1947): "Foundations of a general theory of sequential decision functions," *Econometrica*, 279–313.

Wiseman, T. (2008): "Disagreement leads to complete learning: Sequential choice with continuous types," *Economics Letters*, 100, 53–55.