Available online at www.sciencedirect.com

ScienceDirect

Journal of Economic Theory 175 (2018) 713–729

JOURNAL OF Economic Theory

www.elsevier.com/locate/jet

Notes

Herding with costly information ☆

S. Nageeb Ali

Pennsylvania State University, United States

Received 24 May 2014; final version received 6 February 2018; accepted 19 February 2018

Abstract

This paper incorporates costly information into a model of observational learning. Individuals would like to avoid the cost of buying information and free-ride on the public history. The paper characterizes when learning is nevertheless complete. Necessary and sufficient conditions for complete learning follow from an elementary principle: a player purchases information only if it can influence her action. With a "coarse" action space, learning is complete if and only if for every cost $c > 0$, a positive measure of types can acquire, at cost less than c, an experiment that can overturn the public history. With a "rich" action space, learning is complete if and only if for every cost $c > 0$, a positive measure of types can acquire any informative signal at cost weakly less than c. The results are applied to financial markets to evaluate when markets are informationally efficient despite information being costly.

JEL classification: D82; D83

Keywords: Social learning; Herding; Information acquisition; Responsiveness

☆ This is a revision of the paper, "Social Learning with Endogenous Information". I thank Chris Chambers, Ben Golub, Johannes Horner, Vijay Krishna, Ignacio Monzon, Manuel Mueller-Frank, Omar Nayeem, Paul Niehaus, Henrique Roscoe de Oliveira, Mallesh Pai, Andres Santos, Ron Siegel, Joel Sobel, Tom Wiseman, and especially Aislinn Bohren and Navin Kartik. This paper has benefited significantly from the careful reading and suggestions of an Associate Editor and referee. Tetsuya Hoshino, Erik Lillethun, and Garima Singal provided excellent proofreading. This work is financially supported by NSF grant SES-1127643.

E-mail address: nageeb@psu.edu.

https://doi.org/10.1016/j.jet.2018.02.009
0022-0531/

1. Introduction

Why do individuals imitate each other? The observational learning literature, initiated by Banerjee (1992) and Bikhchandani et al. (1992), offers the following perspective: when Alice sees many others before her choosing the same action, she infers that sufficiently many of them have private information that favors that action. That inference can induce her to follow suit even if her own information indicates otherwise. By joining the herd, Alice's action obscures her information from future players, and thus, induces informational inefficiency.

Understanding how this motive for imitation influences long-run behavior is the main theme of the herding literature. Are herds guaranteed to form, and if so, can they persist indefinitely on incorrect actions? In a seminal paper, Smith and Sørensen (2000) show that herding is inevitable but herds persist indefinitely on incorrect actions if and only if information is of bounded persuasiveness.

These insights are developed in a setting where all individuals obtain information for free. But individuals often have to devote time and resources to acquire information. Once players find it costly to acquire information, there is a new motive for herding, namely that a player can use the "wisdom of the crowd" to avoid incurring the cost of information acquisition. For example, when choosing among health insurance plans, individuals typically find it costly to learn about the various characteristics of these plans, and one might instead choose the same plans that one learns that one's co-workers have chosen (Sorensen, 2006). Analogously, in the adoption of agricultural technology, individuals may herd on the technological choices of others rather than learn about the efficacy of various technological interventions oneself (Conley and Udry, 2001). Information costs strengthen the motive to herd.

The goal of this paper is to ask the main question of the herding literature once information is costly—when do herds persist only on correct actions and never on incorrect actions? This paper answers this question by combining the most basic principle of information demand—*information is valuable only if it can change one's action*[1]—with the martingale techniques of Smith and Sørensen (2000) to characterize long-run learning. The lesson provided by the main result, Theorem 1, is simple:

> *Herds persist only on correct actions if and only if for every interior public belief, players acquire information that can overturn it with positive probability.*

This result applies in both discrete and continuous action spaces, with information costs being modeled non-parametrically, and allowing for players to be heterogeneous in their information acquisition costs. Below, I describe more of the framework, the obstacles that are encountered, and the intuition for this result.

Individuals sequentially choose from a menu of options while being uncertain about the realization of a payoff-relevant state of the world. Prior to making that choice, each individual observes the full history of prior choices, and then chooses whether to conduct an experiment (or test) whose stochastic result reveals some information about the underlying state of the world. The cost of that experiment may depend both on its characteristics (e.g., its informativeness) and on characteristics particular to that individual. This heterogeneity captures differences in both

[1] This fundamental property is implicit in Blackwell's comparison of information structures (Blackwell, 1951, 1953) and is described by Arrow (1971) who attributes it to Marschak (1959). This property is also used by Schlee (1990) and Grant et al. (1998) as a point of contrast for theories in which information has intrinsic value.

absolute and comparative advantage in information acquisition across players. For example, Bob may find it more costly than Alice to acquire any kind of information. Alternatively, Alice may find it less costly than Carol to evaluate the long-run cost of a health insurance plan, but finds it more costly to evaluate the quality of covered medical services.

Because Alice observes the full history of prior actions before choosing which test to run, the choice of information acquisition is *endogenous*: if all her predecessors choose the same option, she has less reason to acquire information (since the chance that she'll overturn their consensus is low), but if instead, those before her have not settled on a single choice, information may be valuable. Hence, Alice's choice of tests varies across histories.

In addition to being endogenous, Alice's choice of information acquisition must reflect her strategic reasoning (and her reasoning about others' reasoning). She recognizes that her predecessors faced the same choice as she does, and some of them may have been tempted to not acquire any information at all. She cannot directly observe who acquired information nor what they learned, but forms inferences based on the prior history.

The heterogeneity of players and the endogenous nature of information acquisition make it infeasible to tractably characterize short-run and medium-run behavior without parametric restrictions on the decision problem and information structure. This paper sidesteps these issues by using asymptotic methods to characterize long-run behavior. While beliefs about the true state ebb and flow as individuals acquire information and act, the Martingale Convergence Theorem guarantees that learning eventually stops. Does it stop before individuals have settled on the right action?

Theorem 1 summarized above offers the necessary and sufficient condition for learning to settle only on the right action, and I explain the role of that condition. For learning to converge to the truth (with probability 1), at any belief short of certainty, information must be acquired with positive probability. If at any such interior belief, there is strictly positive probability that a player is willing to acquire information that can overturn it, then that interior belief is unstable: either the prior consensus is overturned or it isn't, and in either case, the outcome is informative to subsequent players. Thus, a sufficient condition for actions to settle only on the right action is that there is some form of information that can both overturn the prior consensus and is sufficiently inexpensive that a positive mass of players acquires it.

This condition is also necessary. If at some interior public belief, the only information that can overturn it is too costly for all players, then no one is willing to acquire any information. At that point, learning stops at an interior belief, which implies that with positive probability, players may have converged to taking the wrong action.

This necessary and sufficient condition bears different implications for different decision problems. With a "coarse" action space (e.g. finite), information can overturn only if it can "swing" a player's action whereas with a sufficiently rich action space, all that the information has to do is induce a player to "tweak" her action. This paper develops the language of *responsiveness* to describe this distinction between "coarse" and "rich" action spaces, and shows that the distinction is more than that of finite versus continuous action spaces.

A decision problem is *responsive* if any change in a player's belief changes her optimal action. The quintessential example is a "prediction problem" proposed by Lee (1993)[2] in which each player chooses a in $[0, 1]$, and obtains a payoff of $-(a - \omega)^2$ in state ω. In this problem,

[2] This prediction problem has also been studied by others within the social learning literature. A partial and incomplete list is Vives (1997), Burguet and Vives (2000), Eyster and Rabin (2010, 2014), and Guarino and Jehiel (2013).

the optimal action shifts monotonically with a player's beliefs about the states. By contrast, if a player's optimal action does not always shift with her beliefs, that decision problem is *unresponsive*. A particular kind of unresponsive decision problem is where the extreme action that she chooses when she is absolutely certain of the state remains her optimal action when she is sufficiently confident; such a decision problem is *unresponsive at certainty*. Every finite action model is unresponsive at certainty because every weakly undominated action is optimal for a range of beliefs. However, a continuous action choice model may also be unresponsive at certainty; see Section 3.2 for examples that include risky investments and lumpy actions that have fixed costs.

Theorem 1 characterizes long-run behavior. For responsive decision problems, learning obtains whenever some informative experiment is conducted with positive probability at every interior belief. By contrast, if the decision problem is unresponsive at certainty, it must be that at each interior public belief, players (with positive probability) purchase signals that can overturn that belief. A sufficient condition is that the affordable signal induces unbounded likelihood ratios, but this is not necessary; it is necessary and sufficient that for every degree of confidence, there is an affordable signal that can generate that degree of confidence. Accordingly, stronger conditions are needed to preclude bad herds in this case.

The paper applies these results to financial markets, particularly to revisit the question of Grossman and Stiglitz (1980): when information is costly, can markets be informationally efficient? I consider a trading model in which noisy and speculative traders deal with a market-maker in a fixed sequential order, and each trader faces the choice of buying or selling a security, or remaining inactive. Theorem 3 shows that the flexibility of the price mechanism renders the decision problem responsive even though each trader chooses from finitely many options.

Related literature Burguet and Vives (2000) are the first to investigate the role of costly information in social learning. Theirs is a model of social learning wherein at each time period, each of a continuum of short-lived players faces the decision problem of Lee (1993), which is responsive. Each player chooses the precision with which to observe the state of the world and takes an action. Because learning would be trivially resolved by observing the choices of a continuum of players, the model assumes that all past play is observed with noise. Players face the same information acquisition costs, which have a parametric structure: each player chooses the precision of a normally distributed signal where higher precision is more costly. Their analysis and mine are complementary in that they study how assumptions on the marginal cost of information evaluated at 0 influence the prospects for complete learning when all players are identical. By contrast, I focus on a setting where the costs of information acquisition vary across players. Moreover, the results here pertain to both responsive and unresponsive decision problems in the sequential choice of individual players, all of whom perfectly observe the past history. While every experiment in their environment necessarily induces unbounded likelihood ratios, my results show that this may be unnecessary for complete learning in this responsive decision problem. Instead, Theorem 1 shows that learning is complete so long as the experiment is even mildly informative, with induced likelihood ratios that are bounded away from 0 and 1.

Hendricks et al. (2012) study a sequential search problem in which each player can choose a costly perfect signal about her value for the product, and observes only the fraction of people who have purchased the product (and not the previous history). Because players do not observe the completely ordered history, their analysis does not afford an elementary martingale treatment, and hence, their analysis and setting aren't comparable to that here.

Mueller-Frank and Pai (2016) study a costly search process in which the payoff from each action is distributed i.i.d. and when a player "searches an action", he observes the value of that action perfectly without learning that of any other. The action space here is finite, and therefore, the decision problem is unresponsive at certainty. Information takes a particular form in their paper: it is perfect about one action and completely noisy about others. Complete learning obtains if and only if some type in the support can obtain perfect noise-free information for free, and otherwise, learning is incomplete. Apart from the differences in the framework and approach, the results of my paper show that information need not be perfect for complete learning to obtain.

Perhaps surprisingly, none of these prior papers nest the standard herding model and this paper is the first to do so. The motivation here is to offer a more direct contrast of observational learning with and without information acquisition costs. Accordingly, rather than adopt a particular parametric structure, the analysis pursues a more elementary approach that yields a simpler and more general treatment of long-run behavior across both responsive and unresponsive decision problems. Thus, the scope and form of analysis here both significantly differ from the papers mentioned above and showcase how the techniques that Smith and Sørensen (2000) develop for observational learning without information costs are useful even when information is costly.

2. Model

Each of an infinite sequence of players $i = 1, 2, 3, \ldots$ chooses an action a_i from a menu $\mathcal{A}$ that contains at least two actions and is a compact subset of $\Re$. The payoff from these actions depends on the state of the world ω, which is high ($\omega = 1$) or low ($\omega = 0$). Choosing action a in state ω generates a payoff of $u(a, \omega)$, which is continuous in a for each ω. No action is weakly dominated: if $u(a, 0) > u(a', 0)$, then $u(a, 1) < u(a', 1)$. Therefore, there is no loss of generality in setting $u(a, 0)$ to be strictly decreasing in a and $u(a, 1)$ to be strictly increasing in a. The lowest and highest actions in $\mathcal{A}$ are $\underline{a}$ and $\overline{a}$ respectively, and these extreme actions maximize $u(\cdot, 0)$ and $u(\cdot, 1)$ respectively. The players' common prior attributes probability $\pi \in (0, 1)$ to the high state.

After player i observes actions chosen by predecessors, namely $(a_j)_{j=0,\ldots,i-1}$, she chooses whether to acquire information about the state, and if she does so, exactly what kind of information to acquire. Acquiring information is modeled as the choice of an experiment (or test) from a set of experiments $\mathcal{X}$. An experiment X in $\mathcal{X}$ generates realizations (or outcomes) in $[0, 1]$ governed by the cumulative distribution function $F_X(\cdot, \omega)$ in state ω, and these realizations are drawn independently of the realizations observed by previous players.[3] For every experiment X, the distributions $F_X(\cdot, 0)$ and $F_X(\cdot, 1)$ are mutually absolutely continuous, and have support $Supp(X)$. Following convention, the realization of an experiment is normalized to be the posterior belief that the state is high that would be induced with a neutral prior, and thus, realizations are drawn from a subset of $[0, 1]$.[4] Each experiment X is at least partially informative; in other words, there exists $p \in [0, 1]$ such that $F_X(p, 0) \neq F_X(p, 1)$. The convex hull of $Supp(X)$ is denoted $\left[\underline{p}(X), \overline{p}(X)\right]$.

If the player chooses not to acquire any information, then she incurs no costs from information acquisition. If she chooses to acquire information, her cost depends both on the experiment she chooses and a privately observed cost-parameter that describes her strengths and weaknesses

[3] In other words, even if two players conduct the same experiment—or observe the same signal process—the realizations that they observe are i.i.d. conditional on the state.

[4] In other words, if F_X is differentiable at p, and $p \in Supp(X)$, $p = \frac{f_X(p,1)}{f_X(p,1)+f_X(p,0)}$. If an individual has prior π, the signal realization p generates a posterior likelihood ratio of $\frac{p\pi}{(1-p)(1-\pi)}$.

in information acquisition. That privately observed cost parameter is denoted by θ and drawn according to a distribution ρ whose support is Θ, a compact subset of a metric space endowed with the Borel σ-algebra. Thus, a player earns a payoff of $u(a, \omega)$ in state ω from her action choice of a if she acquired no information at all, and $u(a, \omega) - c(X, \theta)$ if she chooses to run the experiment X when her cost parameter is θ, where $c(X, \theta)$ is non-negative.[5]

Continuity and compactness conditions are imposed on this cost function. Because each experiment X is summarized by a pair of distribution functions $(F_X(\cdot, 0), F_X(\cdot, 1))$, convergence in distributions (or equivalently, weak convergence) defines the appropriate notion of continuity for $\mathcal{X}$.[6] I assume that the set of experiments $\mathcal{X}$ is compact in this topology, and that $c(X, \theta)$ is continuous in $X \times \Theta$.[7] Because c is continuous on the compact set $\mathcal{X} \times \Theta$, the Heine-Cantor Theorem implies that c is uniformly continuous.

Each player observes actions of all predecessors but not their information. The public history before player $t > 1$ chooses to acquire information is $h^t \equiv (a_i)_{i=1,\dots,t-1}$. After observing the public history and her type, a player chooses an experiment X to conduct and after observing its realization, chooses an action a. I study Perfect Bayesian Equilibria (henceforth PBE). In a PBE, σ, and history h^t, let $\mu_t(h^t) \equiv Pr(\omega = 1|h^t)$ summarize the *public belief* after history h^t. Consider the set $\mathcal{H}$ of infinite length histories, and for a history h_∞ in $\mathcal{H}$, let h^t_∞ be its truncation to actions in periods $1, \dots, t-1$. For $\omega \in \{0, 1\}$, let $\mathcal{H}_\omega$ denote the set of histories in $\mathcal{H}$ such that $\lim_{t\to\infty} \mu^t(h^t_\infty) = \omega$. Learning is *complete* if for each $\omega \in \{0, 1\}$, $Pr\,(h_\infty \in \mathcal{H}_\omega \mid \omega) = 1$, and otherwise, learning is *incomplete*.

3. When is learning complete?

3.1. The affordability of information

This section introduces terminology to categorize costs of information acquisition. Imagine that a player has budgeted $k \geq 0$ to spend on acquiring information. Fixing an experiment X, and depending on the realized cost parameter, some players may be able to afford X on the budget of k while others may not. The ex ante probability that a player can afford X without fully expending a budget of k is $G(X, k) \equiv \rho\,(\{\theta \in \Theta : c(X, \theta) \leq k\})$, where θ is the cost-parameter that describes a player's costs of information acquisition, and ρ is the probability distribution that governs θ. Experiment X is *affordable* if $G(X, k) > 0$ for every $k > 0$: in other words, for every strictly positive budget, there is always a strictly positive mass of types that can afford experiment X with that budget.[8] The opposite of affordability—namely that X is *unaffordable*—implies that

[5] While the model assumes a one-shot process of information acquisition, it can also serve as a reduced-form for a sequential information acquisition environment (e.g. Wald, 1947; Moscarini and Smith, 2001) in which each individual can choose dynamically how much information to acquire. Suppose that each player can acquire multiple signals sequentially, conditioning the acquisition of a signal on the realizations of signals she has already acquired. A reduced-form version of this dynamic setting corresponds to $\mathcal{X}$ being every feasible strategy and $c(X, \theta)$ being the expected cost for a player whose cost parameter is θ. I thank Thomas Wiseman for making this suggestion.

[6] A sequence of experiments $\{X_n\}_{n=1,2,\dots}$ converges in distribution to X if for every ω and p at which $F_X(p, \omega)$ is continuous, $F_{X_n}(p) \to F_X(p)$. This continuity notion is metrized using the Levy-Prokhorov metric (Billingsley, 1995).

[7] In other words, when a sequence of experiments $\{X_n\}_{n=1,2,\dots}$ converges in distribution to X, and taking a sequence $\{\theta_n\}_{n=1,2,\dots}$, then $c(X, \theta) = \lim_{n\to\infty} c(X_n, \theta_n)$.

[8] Continuity of $c(X, \theta)$ in θ ensures that experiment X is affordable if and only if the support of the distribution of costs of conducting X includes 0.

there is some cost $\varepsilon > 0$ such that almost-surely, a player cannot afford to conduct experiment X for a budget less than ε.

More broadly, *information is affordable* if the set of experiments, $\mathcal{X}$ contains an affordable experiment, and otherwise, *information is unaffordable*. The following result establishes that if information is unaffordable, there must be fixed or lumpy costs in information acquisition.

Lemma 1. *If information is unaffordable, there exists $\varepsilon > 0$ such that $c(X, \theta) > \varepsilon$ for every X, θ.*

The argument for Lemma 1 relies on c being continuous on a compact set (and hence, uniformly continuous) and the measure over types, ρ, being full-support on Θ.

The above discussion pertains to all experiments. But in certain settings—e.g., a finite-action space—a player only values those experiments that can change her mind, swinging her optimal action from one to another. To this end, say that *overturning information is affordable* if for every $p_*, p^* \in (0, 1)$, there exists an affordable experiment X such that $\underline{p}(X) \leq p_*$ and an affordable experiment Y such that $\overline{p}(Y) \geq p^*$.[9] The affordability of overturning information implies that for any required level of confidence, there are affordable experiments that can (with positive chance) induce at least that level of confidence.

I call it "overturning" because of its role in overturning beliefs: if it is satisfied, then regardless of the budget for information acquisition, there is a positive probability of types who are acquiring information that can "overturn" the public history. A sufficient, but not necessary, condition is that there exists an affordable experiment whose likelihood ratios are unbounded.[10]

Contrasting to this definition, say that *overturning information is unaffordable* if there exists $\varepsilon > 0$, $p_* > 0$, and $p^* < 1$ such that $c(X, \theta) < \varepsilon$ implies that $Supp(X) \subset [p_*, p^*]$. In other words, whenever a player is restricted to a sufficiently small budget, then all the experiments that can be afforded on this budget assume likelihood ratios that are a subset of $\left[\frac{p_*}{1-p_*}, \frac{p^*}{1-p^*}\right]$, and are therefore bounded from both directions.

A careful reader may see that overturning information being unaffordable is stronger than the negation of it being affordable, since the latter might bound likelihood ratios from only one direction. Imposing bounds from both directions simplifies the exposition of Theorem 1 without changing its qualitative message. This issue is revisited after stating Theorem 1, and the corresponding result using merely the negation of overturning information is described.

The cost structure of information accounts for two dimensions of interest: (i) each player chooses what kind of information to observe, and it may be that more precise information is more costly, and (ii) players are heterogeneous in the private costs incurred in learning the true state. Prior analyses of social learning have focused exclusively on either of these channels, to the exclusion of the other, and have not studied their interaction.[11]

---

9 For an experiment X, recall that the convex hull of the support of X is $[\underline{p}(X), \overline{p}(X)]$.

10 For example, suppose that $\mathcal{X}$ is a set of experiments $\{X_3, X_4, \ldots\}$ such that $Supp(X_n) = \{\frac{1}{n}, \frac{1}{3}, \frac{2}{3}, \frac{n-1}{n}\}$, where the ex ante probability that X_n generates a realization outside of $\{\frac{1}{3}, \frac{2}{3}\}$ is at most $\frac{1}{n}$, and $c(X_n, \theta) = \frac{1}{n}$. Then, there is no experiment that induces unbounded likelihood ratios, but nevertheless, overturning information is affordable.

11 Burguet and Vives (2000) model homogeneous players and account for how the marginal cost of information influences social learning in a continuous-action prediction problem. Mueller-Frank and Pai (2016) model heterogeneous players all of whom access a noise-free experiment but vary in their costs of accessing that experiment, and describe when learning is complete.

3.2. The responsiveness of a decision problem

The other determinant of long-run learning is the nature of the decision problem: clearly, there is an important difference between the continuous-action prediction problem posed by Lee (1993) (and studied by others) and the finite-action decision problem of Bikhchandani et al. (1992). This paper revisits this distinction because it affects not only the dynamics of social learning (as pointed out in previous work) but also players' value for information.

For a belief that places probability μ on the state being high (i.e., $\omega = 1$), let $a^*(\mu)$ denote the maximizer of the expected payoff $V(a, \mu) \equiv \mu u(a, 1) + (1 - \mu)u(a, 0)$. If there are multiple optimal actions at a belief, select the lowest one for convenience (this tie-breaking rule does not affect the analysis). Because $V(a, \mu)$ satisfies increasing-differences in its arguments, the maximizing action is non-decreasing with a player's belief that the state is high.

Lemma 2. *The optimal action, $a^*(\mu)$, is non-decreasing in the belief, μ, that the state is high.*

Since $\mathcal{A}$ contains no weakly dominated actions, $\underline{a}$ is uniquely optimal when $\mu = 0$ and $\overline{a}$ is uniquely optimal when $\mu = 1$.

Let us pose Lee's continuous-action prediction problem in this language. The payoff for each player is $u(a, \omega) = -(a - \omega)^2$, the choice of actions, $\mathcal{A}$, is [0, 1], and the optimal action at belief μ, $a^*(\mu)$, is simply μ. The conventional logic for why social learning obtains in this problem is that the optimal action reveals a player's belief to subsequent players perfectly. Responsiveness, defined below, generalizes this property.

Say that the decision problem $(\mathcal{A}, u)$ is *responsive* if $a^*(\mu) \neq a^*(\mu')$ whenever $\mu \neq \mu'$, and otherwise, it is *unresponsive*. The prediction problem is, of course, responsive since a player's belief always changes with her action. The quintessential example of unresponsiveness is the finite-action decision problem typically studied in the herding literature, where each (weakly undominated) action is optimal for a range of beliefs.

While a finite action space generates an unresponsive decision problem, decision problems may be unresponsive even when each player chooses from a continuum of weakly undominated actions. For example, take the prediction problem studied before but truncate the action space to $\left[\frac{1}{4}, 1\right]$. Now the optimal action rule is $a^*(\mu) = \max\left\{\frac{1}{4}, \mu\right\}$, and thus, the player's optimal action is locally unresponsive whenever she attributes probability less than $\frac{1}{4}$ to the state being high. Her action, at these beliefs, is (locally) unresponsive to changes in belief. Information is still valuable for such a player: whenever μ is in $\left(0, \frac{1}{4}\right)$, a player places strictly positive probability on $\omega = 1$, and if she knew that were the true state, she would choose a different action.

This truncation may appear to be artificial. Examples 1 and 2 show how this issue surfaces organically in investment problems with continuous investment choices.

Example 1. Suppose, as in Chari and Kehoe (2004), that players choose how much to invest in a risky project. The menu of actions is $\mathcal{A} = [0, 1]$ and investing a in the risky project generates a payoff $h(a)$ if the technology is good ($\omega = 1$) and 0 otherwise; however much is left is invested in a safe asset, which offers 1:1 returns corresponding to the investment. The payoff function then is $u(a, \omega) = h(a)\mathbb{1}_{\omega=1} + 1 - a$, in which h is strictly concave, $1 < h'(0) < \infty$, and $h(0) = 0$. Observe that $a^*(\mu) = 0$ if $\mu < \frac{1}{h'(0)}$, and therefore, the decision problem is unresponsive.[12]

[12] A tractable example of such an investment function is $h(a) = \log(2a + 1)$, in which case $a^*(\mu) = \max\{\mu - \frac{1}{2}, 0\}$.

Example 2. Suppose that players choose, sequentially, how much to invest in a risky asset where each individual has wealth 1 and a strictly concave utility function, which for simplicity is $\log(\cdot)$. Investing a in the risky asset yields a payoff of ka where $k > 1$ if the asset is good ($\omega = 1$), and 0 otherwise. The optimal allocation is $a^*(\mu) = \max\left\{\frac{k\mu-1}{k-1}, 0\right\}$. Since a player chooses to invest nothing in the risky asset if $\mu < \frac{1}{k}$, the decision problem is unresponsive.

In both of these examples, and the truncated prediction problem, responsiveness fails at extreme actions: for some state ω, a player chooses the exact same action when she attributes high probability to ω as she would if she were completely certain. Theorem 1 identifies that failing responsiveness in this particular way is critical for herding. Say that a decision problem is *unresponsive at certainty* if an extreme action is optimal at an interior belief: $a^*(\mu) \in \{\underline{a}, \overline{a}\}$ for μ in $(0, 1)$. In other words, an extreme action is chosen for an interval of beliefs, and not only for a degenerate belief. Apart from the examples above, a decision problem with finitely many actions is unresponsive at certainty because any action that is optimal at certainty is also optimal near-certainty.

3.3. Main result

Each player would like to avoid incurring the cost of buying information and is inclined to free-ride on the wisdom of others. Thus, she purchases information only if it may overturn the public history. This calculus couples the responsiveness of the decision problem with affordability conditions to generate necessary and sufficient conditions for complete learning.

Theorem 1. *Fix a prior* $\pi \in (0, 1)$.

1. *If* $(\mathcal{A}, u)$ *is responsive, learning is complete if and only if information is affordable.*
2. *If* $(\mathcal{A}, u)$ *is unresponsive at certainty, learning is complete if overturning information is affordable, and is incomplete if overturning information is unaffordable.*
3. *If* $(\mathcal{A}, u)$ *is unresponsive, there exists an open set of priors,* Π*, and a set of experiments* $\mathcal{X}$ *such that if* $\pi \in \Pi$*, learning is incomplete even if information is affordable.*

When the decision problem is responsive, all that is necessary and sufficient is that some experiment—even one that is only mildly informative—be affordable. By contrast, if the decision problem is unresponsive at certainty, a more stringent condition is necessary and sufficient: the affordable experiment has to overturn the public history. Thus, stronger conditions are required for complete learning in decision problems that are unresponsive at certainty.

Whenever learning is incomplete, it is also inadequate in the sense of Aghion et al. (1991): because beliefs place positive probability on two distinct states at which the optimal action is also distinct, the action that is chosen when players are near certainty are not the same that would have been chosen if all uncertainty were resolved. In such cases, players still value information that can overturn the public history, but any such experiment is too costly for any player to purchase it.

The intuition for the first two parts of this result is described below, and the full proof is in the next subsection. Suppose throughout the sketch below that the public belief μ places probability $1 - \varepsilon$ on a state, where ε is small. A player's maximum willingness to pay for information at such a history is small given that she is almost certain about the state.

Let us first begin with the case of responsive decision problems. In this case, at any interior belief, a player finds *all* information to be valuable because anything she learns influences her subsequent choice. Affordability of information guarantees that there exists an experiment such that for a strictly positive measure of types, the cost of information is sufficiently low that they are willing to acquire it. Because their actions reveal this information to subsequent players, public beliefs continue to ebb and flow until certainty is reached.[13] By contrast, if information is unaffordable, then conducting any experiment involves a strictly positive cost, regardless of θ, that is bounded from below. Accordingly, once the public belief is sufficiently close to certainty, then no player is willing to conduct any experiment regardless of its signal-to-noise ratio. Learning ceases short of certainty, which implies that with positive chance, players may be adopting the wrong action.

Now consider the case of a decision problem that is unresponsive at certainty. When a player is near-certainty, information is valuable if and only if it can swing one's action; in other words, the information has to overturn the public belief. When overturning information is affordable, one can be assured that regardless of the degree of public confidence, there exist affordable experiments that can overturn it. Therefore, information is acquired with strictly positive probability. Because that information influences actions with positive probability, learning never ceases at an interior belief. By contrast, if overturning information is unaffordable, then there exists a degree of near-certainty such that any experiment that can overturn that degree of confidence is simply too costly for anyone to be willing to acquire it. At that point, learning ceases because even though information has value, that value is outweighed by its cost.

The above result references the unaffordability of overturning information, but returning to the discussion in Section 3.1 about this being stronger than the negation of affordable overturning information, one may be left wondering about the exact necessary and sufficient condition. Say that a decision problem is unresponsive at $\omega \in \{0, 1\}$ if $a^*(\omega) = a^*(\mu)$ for some $\mu \in (0, 1)$. Similarly, say that information that overturns ω is affordable if for every ε, there exists an affordable experiment that has realizations that would reduce the posterior belief that the state is ω to below ε. Now, the same argument as that of Theorem 1 establishes that if $(\mathcal{A}, u)$ is unresponsive at ω, then learning is complete at $\Omega\backslash\{\omega\}$ if and only if information that overturns ω is affordable.

Benchmarks In evaluating Theorem 1, a natural comparison is to the behavior if information were completely costless. Ali (2018) offers this benchmark. Responsive decision problems inherit the key characteristic of Lee (1993): observational learning reduces to a pure statistical problem so that whenever players have access to any information, learning is guaranteed. By contrast, decision problems that are unresponsive at certainty inherit the properties of finite-action games, and analogous to Smith and Sørensen (2000), learning is complete if and only if the set of likelihood ratios induced by experiments is unbounded. Whenever learning is incomplete, it is also inadequate since limit beliefs place positive probability on both states, and the optimal action differs across these states.

---

[13] Theorem 11 differs from and complements the results of Burguet and Vives (2000). Apart from the difference in setting—they study intergenerational learning where at each time period, each of a continuum of one-period lived agents choose actions, and future generations observe only a noisy signal of the average action—there are two important qualitative differences. First, their analysis involves only information that can have unbounded likelihood ratios whereas the analysis here confirms that learning is assured even if likelihood ratios are bounded. Second, their analysis emphasizes the marginal cost of information at 0, which has no counterpart in the analysis here, and the analysis here emphasizes the heterogeneity of information acquisition costs across players, which has no counterpart in their analysis.

To evaluate the role that free-riding plays once information is costly, another natural benchmark is to the behavior of a centralized social planner who considers when to acquire information. Consider a social planner who chooses actions at $t = 1, 2, 3, \ldots$ and with costs of acquiring information in period t being $c(\cdot, \theta_t)$, where θ_t is drawn i.i.d. according to the measure ρ. She discounts payoffs from period t by δ^{t-1} where $\delta < 1$, but does not observe payoffs until the end of the game.[14] When is her long-run learning complete?

Theorem 2. *For a centralized Social Planner who is not perfectly patient ($\delta < 1$), learning is complete either if $(\mathcal{A}, u)$ is responsive and information is affordable or if $(\mathcal{A}, u)$ is unresponsive at certainty and overturning information is affordable. If information is unaffordable, learning is incomplete.*

The sufficient conditions for decentralized learning to be complete are also sufficient for centralized learning to be complete, and information being unaffordable is sufficient to guarantee that centralized learning is incomplete. The main difference between the decentralized and centralized solutions is when the decision problem is certainly unresponsive, overturning information is unaffordable, but information is affordable. In such cases—for example, if bounded information were free—the centralized Social Planner would learn completely while the decentralized solution still features incomplete learning.[15]

3.4. Proof of Theorem 1

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 $l_0^t(h^t) = 1/l_1^t(h^t)$. Treat $\langle l_i^t(\cdot)\rangle_{t=1}^{\infty}$ as a stochastic process, and 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 affordable, and let X be an affordable experiment. Observe that when the public belief is $\mu \in (0, 1)$, the value of experiment X for a responsive individual is

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

whereas the value of not acquiring information is

[14] I fix δ rather than studying the limits as $\delta \to 1$ because in that latter case, the Planner would be willing to undertake many experiments, analogous to bandit environments with perfectly patient players (Easley and Kiefer, 1988; Aghion et al., 1991; Ali, 2011).

[15] I have compared behavior to that of a centralized Social Planner who observes the entire sequence of signals. As highlighted by Smith et al. (2017), a more subtle comparison is to a "teams-problem" whereby the Social Planner is designing the optimal rule for a team to follow that cannot observe each other's information. With costless information, they show that the optimal solution rewards contrarianism. With costly information acquisition, the Social Planner has to reward the right information cost types to engage in information acquisition. It is a question for future study to characterize when such a teams-problem features complete learning, and whether its solution can be implemented via a simple scheme.

$$\bar{V}(\mu) \equiv \mu \int_0^1 u(a^*(\mu), 1) dF_X(p, 1) + (1-\mu) \int_0^1 u(a^*(\mu), 0) dF_X(p, 0). \tag{1}$$

Because X is informative, $B(\mu, p)$ differs from μ with strictly positive probability, and because u is responsive, $a^*(B(\mu, p)) \neq a^*(\mu)$. By revealed preference, $V(X, \mu)$ is strictly higher than $\bar{V}(\mu)$: the decision maker can guarantee herself the payoff of $\bar{V}(\mu)$ after choosing experiment X by choosing $a^*(\mu)$ for every realization of p. Because she deviates from this plan with strictly positive probability, $V(X, \mu) - \bar{V}(\mu) > 0$. Therefore, there exists $\varepsilon > 0$ such that $V(X, \mu) - \bar{V}(\mu) > \varepsilon$. Since X is affordable, $G(X, \varepsilon) > 0$, and therefore, a strictly positive measure prefers experiment X to not acquiring any information at all.

This observation guarantees that learning is complete, i.e., 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 $\varepsilon < l$. The argument in the previous paragraph implies that at this public belief, a strictly positive measure of types obtains information and takes actions that perfectly reveal their signal realizations. The Strong Law of Large Numbers implies then that the probability that the public likelihood ratio remains perpetually in $(l - \varepsilon, l + \varepsilon)$ is 0, and therefore, $l > 0$ cannot be in the support of l_i^∞. Therefore, the support of l_i^∞ is $\{0\}$.

Suppose instead that information is unaffordable. Define $\gamma \equiv \max_{\omega \in \Omega} |u(\underline{a}, \omega) - u(\overline{a}, \omega)|$ as the difference in payoffs between taking the best and worst possible actions. This term, γ, is used to bound the value of information. Observe that by definition of $a^*(\mu)$,

$$\begin{aligned} &\max\{\mu u(\bar{a}, 1) + (1-\mu) u(\bar{a}, 0), \mu u(\underline{a}, 1) + (1-\mu) u(\underline{a}, 0)\} \\ &\quad \leq \mu u(a^*(\mu), 1) + (1-\mu) u(a^*(\mu), 0), \end{aligned}$$

and therefore, for every μ,

$$\begin{aligned} &[\mu u(\overline{a}, 1) + (1-\mu) u(\underline{a}, 0)] - [\mu u(a^*(\mu), 1) + (1-\mu) u(a^*(\mu), 0)] \\ &\quad \leq \min\{(1-\mu)\gamma, \mu\gamma\}. \end{aligned}$$

The LHS describes the value from obtaining a perfect signal, and the RHS offers a bound on that value. Select $\varepsilon < \gamma$ such that according to Lemma 1, $c(X, \theta) > \varepsilon$ for every $X \in \mathcal{X}$. Observe that if $(1-\mu)\gamma < \varepsilon$ or $\mu\gamma < \varepsilon$, then no player is willing to incur a cost of ε for even a perfectly revealing experiment. Define $\mu^* = 1 - \frac{\varepsilon}{\gamma}$ and $\mu_* = \frac{\varepsilon}{\gamma}$. Therefore, if the public belief ever crosses outside of $[\mu_*, \mu^*]$, no player acquires any information. 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 yields that $E[l_i^\infty] = \lim_{t \to \infty} E[l_i^t] = l_i^0$, where the second equality follows from $\langle l_i^t(\cdot) \rangle_{t=1}^\infty$ being a martingale. Because l_i^0 is strictly positive, we have reached 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 overturning information is affordable, and towards a contradiction that for $\mu \in (0, 1)$, $\frac{\mu}{1-\mu}$ or its reciprocal is in the support of l_i^∞. Let X be an affordable experiment such that 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 $\int_{\tilde{P}} dF_X(p, \omega) > 0$ for every $\omega \in \Omega$. The value of experiment X for a responsive individual is $V(X, \mu)$ as characterized in (1). Mirroring the argument of Case

1, it follows from revealed preference that $V(X,\mu) - \bar{V}(\mu) > 0$, and therefore, there exists $\varepsilon > 0$ such that $V(X,\mu) - \bar{V}(\mu) > \varepsilon$. Because X is affordable, $G(X,\varepsilon) > 0$, and thus, a strictly positive measure of types prefers experiment X to not acquiring information. Therefore, a strictly positive measure of types are conducting experiments that influence their actions. 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 overturning information is unaffordable. Consider ε, p_*, p^* such that $c(X,\theta) < \varepsilon$ implies $Supp(X)$ is a subset of $[p_*, p^*]$. 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 the decision problem is unresponsive at certainty, either $\underline{\mu} > 0$ or $\overline{\mu} < 1$ or both. Define

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

$$\mu^+ = \max\left\{1 - \frac{\varepsilon}{\gamma}, \frac{\overline{\mu}(1-p_*)}{\overline{\mu} + p_* - 2\overline{\mu}p_*}\right\}$$

Observe that for a public belief $\mu \in [0, \mu_+] \cup [\mu^+, 1]$, the most that a player would pay for a signal that perfectly reveals the state is ε. So if a player acquires information at all, she acquires a bounded experiment X in which $Supp(X)$ is a subset of $[p_*, p^*]$. However, beliefs are sufficiently concentrated around a state i that for all $p \in [p_*, p^*]$, $a^*(B(\mu, p)) = a^*(\delta_i)$, where δ_i is the belief that puts probability 1 on state i. Since her actions are unaffected by the realizations of X, she has no gain from conducting experiment X, and subsequent players gain no information from observing her actions. Therefore, once the public belief enters $[0, \mu_+] \cup [\mu^+, 1]$, learning must cease.

Towards clarifying the impact on learning, suppose that $\omega = 1$ and $\mu_+ > 0$. If $\pi \le \mu_+$, then we are done. 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 learning would cease above that threshold. 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. Because $a^*(\mu)$ is non-decreasing (as shown in Lemma 2), being non-responsive implies that there exists $\overline{\mu}$ and $\underline{\mu} < \overline{\mu}$ such that $a^*(\mu)$ is constant for every $\mu \in [\underline{\mu}, \overline{\mu}]$. Suppose that the prior π is in (μ, μ') and that the set of experiments $\mathcal{X}$ contains only the experiment X such that with a prior π, the posterior beliefs induced by X all lie within $[\underline{\mu}, \overline{\mu}]$. Suppose that the cost parameter, θ, is distributed uniformly from $[0, 1]$ and that $c(X,\theta) = \theta$ for every θ. Then for every $k > 0$, $G(X,k) = k > 0$ and therefore, X is affordable. Nevertheless, no player acquires information because no realization of experiment X changes her action from $a^*(\pi)$. Therefore, learning is incomplete even though information is affordable. □

4. Are markets informationally efficient?

Once information is costly, do markets aggregate information? This is the classic question posed by Grossman and Stiglitz (1980). The analysis below revisits this issue using a model of sequential trading (Glosten and Milgrom, 1985; Avery and Zemsky, 1998) but with costly information.

Consider a single security whose true value, $V(\omega)$, depends on the realization of $\omega \in \{0,1\}$, and satisfies $0 \le V(0) < V(1) < \infty$. Players share a common prior that $\omega = 1$ with probability $\pi \in (0,1)$. Of the pool of traders, some are *noise traders*, and the rest are *speculative traders*. In a period, the trader is a speculative trader with probability $q \in (0,1)$, independently of the type of other traders. Traders arrive in a fixed sequential order, and each faces a one-time trading decision: she can buy or sell a unit of the security, or remain inactive. Conditional on a trader being a noisy trader, she chooses each of these actions with probability $\frac{1}{3}$. A speculative trader has a payoff of $V(\omega) - \overline{P}$ from buying a security at price $\overline{P}$, a payoff of $\underline{P} - V(\omega)$ from selling a security at price $\underline{P}$, and a payoff of 0 from being inactive. These prices are set by a risk-neutral market-maker who has no private information and is subject to competition. After each history h^t, the market-maker posts a bid price $\underline{P}(h^t)$ at which he is willing to buy a security and an ask price $\overline{P}(h^t)$ at which he is willing to sell a security. The market-maker's zero profit condition implies that $\overline{P}(h^t) \ge \mu(h^t)V(1) + (1-\mu(h^t))V(0) \ge \underline{P}(h^t)$, where $\mu(h^t)$ is the public belief that $\omega = 1$ after history h^t.

Each speculative trader can conduct an experiment X that offers information about ω, and incurs cost $c(X,\theta)$, where θ is the information cost type of that speculative trader. Individuals observe the full history of past trading behavior and prices. As in Section 2, consider a set of $\mathcal{H}$ of infinite length histories, and for such a history, let h^t_∞ be its truncation to periods $1,\dots,t-1$.

One may investigate two roles for prices. The first is the degree to which asymmetric information influences the bid-ask spread, and the second is the degree to which information is aggregated. Formally, *the bid-ask spread vanishes* if $Pr\left(lim_{t\to\infty}(\overline{P}(h^t_\infty) - \underline{P}(h^t_\infty)) = 0 \mid \omega\right) = 1$ for every state ω; in other words, with probability 1, the gap between bid and ask prices vanishes. Analogously, *prices converge to value* if $Pr(lim_{t\to\infty} \max\{\overline{P}(h^t_\infty) - V(\omega), \underline{P}(h^t_\infty) - V(\omega))\} = 0 \mid \omega) = 1$ for every state ω. In other words, with probability 1, market prices converge to the true value of the asset. The following result characterizes conditions for each.

Theorem 3. *Prices converge to value if and only if information is affordable, but the bid-ask spread vanishes both when information is affordable and when it is unaffordable.*

Here is the intuition. With prices being set by a competitive market-maker, the appropriate version of responsiveness is that of "responsiveness in equilibrium": when the public belief is μ (which feeds into the prices set by the market-maker), and a trader's private belief is ν, does a trader's optimal action $a^*(\nu,\mu)$ change as her belief ν varies in the interval $[\mu-\varepsilon, \mu+\varepsilon]$? The answer is yes because she wishes to sell the security whenever $\nu < \mu$ and buy it whenever $\nu > \mu$. Accordingly, a speculative trader values any informative experiment, which leads to complete learning if information is affordable. On the other hand, when information is unaffordable, then prices do not converge to value even though the bid-ask spread vanishes. Thus, in conjunction with Lemma 1, markets are informationally inefficient when the process of information acquisition is "lumpy" or has fixed costs.

These results offer a herding analogue to the failures of information aggregation conjectured by Grossman and Stiglitz (1980): once the price becomes too informative, individuals may stop acquiring information if there are "fixed" or lumpy costs in information acquisition, and that force generates a long-run wedge between the price of an asset and its fundamental value. Absent such lumpiness, informational efficiency is guaranteed. The argument of Theorem 3 illustrates the importance of flexible prices in dynamically adjusting the tradeoffs (based on public beliefs) between buying and selling securities. Were prices fixed and not adjusting with beliefs, traders would face a decision problem that is unresponsive at certainty, and in that case, overturning information would need to be affordable. The flexibility of prices guarantees that learning is complete even if the only affordable information is only minimally informative.

Appendix A. Omitted proofs

Proof of Lemma 1. Suppose that the conclusion is false. Then there exists a sequence $(X_n, \theta_n)_{n=1,2,\ldots}$ such that $c(X_n, \theta_n) \leq \frac{1}{n}$. Because $\mathcal{X} \times \Theta$ is compact, the sequence has a convergent subsequence (Aliprantis and Border, 2006, Theorem 2.31). Let (X^*, θ^*) be a limit point of that subsequence. As a limit point, $c(X^*, \theta^*) \leq \frac{1}{n}$ for each n, and therefore, $c(X^*, \theta^*) = 0$.

Fix $\varepsilon > 0$. Because c is continuous, it follows that there exists $\delta > 0$ such that for every θ in $N_\delta(\theta^*)$, $c(X^*, \theta) \leq \varepsilon$. Observe that

$$G(X^*, \varepsilon) = \rho(\{\theta \in \Theta : c(X^*, \theta) \leq \varepsilon\}) \geq \rho(N_\delta(\theta^*)) > 0$$

where the equality is the definition of G, the weak inequality follows from $N_\delta(\theta^*) \subseteq \{\theta \in \Theta : c(X^*, \theta) \leq \varepsilon\}$, and the strict inequality follows from θ^* being in the support of ρ (because ρ is a full-support distribution on Θ). So $G(X^*, \varepsilon) > 0$ for every $\varepsilon > 0$ and therefore, X^* is affordable. Taking the contrapositive establishes the result. □

Proof of Lemma 2. Observe that $V_\mu(a, \mu) = u(a, 1) - u(a, 0)$. Since $u(a, 1)$ is increasing in a and $u(a, 0)$ is decreasing in a, it follows that for $a'' > a'$, $V_\mu(a'', \mu) > V_\mu(a', \mu)$. Therefore V satisfies increasing-differences in (a, μ), and hence, the maximizer $a^*(\mu)$ is non-decreasing in μ. □

Proof of Theorem 2. First, consider the case where (resp. overturning) information is affordable, and $(\mathcal{A}, u)$ is responsive (resp. unresponsive at certainty). Suppose that $\omega = 0$ (the argument is identical for $\omega = 1$) and suppose towards a contradiction that learning were incomplete. So there exists $\mu \in (0, 1)$ such that $\frac{\mu}{1-\mu}$ is in the support of l_0^∞. The one-shot deviation principle establishes a contradiction. Consider an affordable informative (resp. overturning) experiment X. We have already established in the proof of Theorem 1 that there exists $\varepsilon > 0$ such that $V(X, \mu) - \bar{V}(\mu) > \varepsilon$. Because $G(X, \varepsilon) > 0$, it follows that for a strictly positive measure of cost-parameters, the Social Planner benefits from a one-shot deviation where she conducts experiment X rather than not acquiring information, which is a contradiction.

Now suppose that information is unaffordable. Let $\varepsilon < \gamma$ be such that $c(X, \theta) > \varepsilon$. Consider beliefs $\mu^* = 1 - \frac{(1-\delta)\varepsilon}{\gamma}$ and $\mu_* = \frac{(1-\delta)\varepsilon}{\gamma}$. Notice that for beliefs $\mu \in [\mu^*, 1] \cup [0, \mu_*]$, the Social Planner would not be willing to pay ε once for a fully revealing signal. Therefore, if learning does not stop beforehand, it stops once beliefs reach this region. □

Proof of Theorem 3. Fix an equilibrium of the game. Each trader's best response at a history h^t depends only on the public belief $\mu(h^t)$. Let $\overline{P}(\mu)$ and $\underline{P}(\mu)$ denote the ask and bid prices for a history that generates public belief μ.

As before, $\langle l_i^t \rangle$ is a martingale conditioning on $\omega = i$, and converges almost-surely to a random variable l_i^∞. Suppose that $\mu \in (0, 1)$, $\frac{\mu}{1-\mu}$ or its reciprocal is in the support of l_i^∞. For μ to be a limit belief, it must be that there is no more trading done on the basis of information, and therefore, $\overline{P}(\mu) = \underline{P}(\mu) = \mu V(1) + (1 - \mu) V(0)$. Therefore, the bid-ask spread vanishes.

If a speculative trader does not acquire information, then it is optimal at those prices for her to remain inactive (or at least, she is indifferent between trading and not). Suppose that information is affordable. Let X be an informative experiment, which implies that with strictly positive probability, X generates realizations p such that $B(\mu, p) \neq \mu$. But if $B(\mu, p) < \mu$, then it is strictly optimal for her to sell the security, and if $B(\mu, p) > \mu$, then it is strictly optimal for her to buy the security. Therefore, it follows, as in the proof of Theorem 1 that the value of acquiring signal X strictly exceeds that of not acquiring any information. Since X is affordable, a strictly positive measure of types acquires information and trades speculatively. But then setting $\overline{P}(\mu) = \underline{P}(\mu) = \mu V(1) + (1-\mu)V(0)$ is not an equilibrium price for the market-maker.

Suppose instead that information is unaffordable, and let $\varepsilon > 0$ be a cost such that any informative experiment costs more than ε. If the public belief is μ, an upper-bound on how much a speculative trader values perfect information is

$$\begin{aligned}&\mu\,[V(1) - (\mu V(1) + (1-\mu)V(0))] + (1-\mu)\,[\mu V(1) + (1-\mu)V(0) - V(0)]\\ &\quad = 2\mu(1-\mu)(V(1) - V(0)).\end{aligned}$$

Therefore, there exists μ_* and μ^* such that if $\mu < \mu_*$ or $\mu > \mu^*$, the value of perfect information is smaller than ε. Consider towards a contradiction an infinite history h_∞ such that in state ω, $lim_{t\to\infty} \max\{\overline{P}(h^t_\infty) - V(\omega), \underline{P}(h^t_\infty) - V(\omega))\} = 0$. Such a history would require players to acquire information and trade speculatively for μ arbitrarily close to 0 or 1. But once μ escapes $[\mu_*, \mu^*]$, no speculative trader has any incentive to conduct an experiment, leading to a contradiction. □

References

Aghion, P., Bolton, P., Harris, C., Jullien, B., 1991. Optimal learning by experimentation. Rev. Econ. Stud. 58, 621–654.

Ali, S.N., 2011. Learning self-control. Q. J. Econ. 126, 857–893.

Ali, S.N., 2018. On the role of responsiveness in rational herds. Econ. Lett. 163, 79–82.

Aliprantis, C.D., Border, K., 2006. Infinite Dimensional Analysis: a Hitchhiker's Guide, 3rd edition. Springer, New York, N.Y..

Arrow, K.J., 1971. The value of and demand for information. Decis. Organ. 2, 131–139.

Avery, C., Zemsky, P., 1998. Multi-dimensional uncertainty and herd behavior in financial markets. Am. Econ. Rev. 88, 724–748.

Banerjee, A., 1992. A simple model of herd behavior. Q. J. Econ. 107, 797–817.

Bikhchandani, S., Hirshleifer, D., Welch, I., 1992. A theory of fads, fashion, custom, and cultural changes as informational cascades. J. Polit. Econ. 100, 992–1026.

Billingsley, P., 1995. Probability and Measure, third edition. Wiley Series in Probability and Mathematical Statistics. Wiley, New York.

Blackwell, D., 1951. Comparison of experiments. In: Neyman, J. (Ed.), Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability. University of California Press, Berkeley, CA, pp. 93–102.

Blackwell, D., 1953. Equivalent comparisons of experiments. Ann. Math. Stat. 24, 265–272.

Burguet, R., Vives, X., 2000. Social learning and costly information acquisition. Econ. Theory 15, 185–205.

Chari, V.V., Kehoe, P.J., 2004. Financial crises as herds: overturning the critiques. J. Econ. Theory 119, 128–150.

Conley, T., Udry, C., 2001. Social learning through networks: the adoption of new agricultural technologies in Ghana. Am. J. Agric. Econ. 83, 668–673.

Easley, D., Kiefer, N.M., 1988. Controlling a stochastic process with unknown parameters. Econometrica 56, 1045–1064.

Eyster, E., Rabin, M., 2010. Naive herding in rich-information settings. Am. Econ. J. Microecon. 2, 221–243.

Eyster, E., Rabin, M., 2014. Extensive imitation is irrational and harmful. Q. J. Econ. 129, 1861–1898.

Glosten, L.R., Milgrom, P.R., 1985. Bid, ask and transaction prices in a specialist market with heterogeneously informed traders. J. Financ. Econ. 14, 71–100.

Grant, S., Kajii, A., Polak, B., 1998. Intrinsic preference for information. J. Econ. Theory 83, 233–259.

Grossman, S.J., Stiglitz, J.E., 1980. On the impossibility of informationally efficient markets. Am. Econ. Rev. 70, 393–408.

Guarino, A., Jehiel, P., 2013. Social learning with coarse inference. Am. Econ. J. Microecon. 5, 147–174.

Hendricks, K., Sorensen, A., Wiseman, T., 2012. Observational learning and demand for search goods. Am. Econ. J. Microecon. 4, 1–31.

Lee, I.H., 1993. On the convergence of informational cascades. J. Econ. Theory 61, 395–411.

Marschak, J., 1959. Remarks on the economics of information. In: Contributions to Scientific Research in Management. University of California, pp. 79–98.

Moscarini, G., Smith, L., 2001. The optimal level of experimentation. Econometrica 69, 1629–1644.

Mueller-Frank, M., Pai, M.M., 2016. Social learning with costly search. Am. Econ. J. Microecon. 8, 83–109.

Schlee, E., 1990. The value of information in anticipated utility theory. J. Risk Uncertain. 3, 83–92.

Smith, L., Sørensen, P., 2000. Pathological Outcomes of Observational Learning. Econometrica 68, 371–398.

Smith, L., Sørensen, P., Tian, J., 2017. Informational Herding, Optimal Experimentation, and Contrarianism. University of Wisconsin.

Sorensen, A.T., 2006. Social learning and health plan choice. Rand J. Econ. 37, 929–945.

Vives, X., 1997. Learning from others: a welfare analysis. Games Econ. Behav. 20, 177–200.

Wald, A., 1947. Foundations of a general theory of sequential decision functions. Econometrica, 279–313.