Economics Letters 163 (2018) 79–82

Contents lists available at ScienceDirect

Economics Letters

journal homepage: www.elsevier.com/locate/ecolet

On the role of responsiveness in rational herds[☆]

S. Nageeb Ali

Pennsylvania State University, 303 Kern Bldg, University Park, PA 16802, United States

HIGHLIGHTS

- This paper unifies results in social learning with coarse and rich action spaces.
- This paper develops the language of responsiveness.
- Learning is complete in responsive decision problems.
- It is complete in unresponsive decision problems if and only if information is unbounded.
- A decision problem can be unresponsive even with a continuous action space.

ARTICLE INFO

Article history:
Received 3 September 2017
Received in revised form 12 November 2017
Accepted 21 November 2017
Available online 2 December 2017

JEL classification:
D83
D7

Keywords:
Herding
Responsiveness
Observational learning
Social learning

ABSTRACT

Models of rational herding typically involve a finite action space. An intuition for herding is that this coarseness of the action space relative to the space of potential beliefs is responsible for herding, and were the action space sufficiently rich, learning would be complete. That intuition is false: simple examples illustrate that learning may be incomplete even if the action space is isomorphic to the space of beliefs. What then distinguishes models with "coarse" versus "rich" action spaces? This paper develops the language of *responsiveness* to formalize this distinction. Responsiveness assesses the sensitivity of optimal actions with respect to their rationalizing beliefs. If the optimal action always changes with beliefs, then complete learning is guaranteed regardless of the information structure. By contrast, if the action that is optimal at certainty remains optimal near-certainty, then complete learning is guaranteed if and only if information can induce unbounded likelihood ratios. The lens of responsiveness unifies results across coarse and rich action spaces.

1. Introduction

The observational learning literature, initiated by Banerjee (1992) and Bikhchandani et al. (1992), studies a population of players sequentially facing the same decision problem, with each player observing the full history of prior choices. The literature offers an informational rationale for imitation: after observing several individuals making the same choice, an individual is inclined to imitate them because their consensus is informative. This imitation can generate herds on the wrong action and potentially lead to incomplete learning.

An intuition commonly expressed for herding is that the coarseness of the action set obscures a player's information. Because others do not directly observe a player's information about the state of the world, they can infer it only through her choice of action. If the action set is finite, that filtering is potentially imperfect. By contrast, Lee (1993) shows that in a continuous-action environment, that filtering problem disappears. Each player's beliefs are perfectly revealed through her actions, and so there is no loss from observing only players' actions and not their information. These contrasting results from Bikhchandani et al. (1992) and Lee (1993) may lead one to believe that there is an important distinction between finite and continuous action spaces.

Focusing on that distinction may be incorrect. One can construct examples in which the set of weakly undominated actions is continuous – and hence the action space is "rich"– and nevertheless herds ensue. Example 2 on p. 5 is one such example, and more examples appear in Huck and Oechssler (1998), Chari and Kehoe (2004), and Ali (2017). These examples raise the question: *if the distinction is not between "coarseness" and "richness," what is the appropriate categorization?*

☆ This paper develops results that were originally in the working paper "Social Learning with Endogenous Information," which has now been divided into this paper and a companion paper on "Herding with Costly Information." This note has benefited from conversations with Erik Eyster, Ben Golub, Joel Sobel, and especially Aislinn Bohren and Navin Kartik, and from the careful reading and suggestions of an anonymous referee. This work is financially supported by NSF grant SES-1127643.

E-mail address: nageeb@psu.edu.

https://doi.org/10.1016/j.econlet.2017.11.029
0165-1765/

This note proposes a new categorization through the lens of *responsiveness*. Responsiveness assesses the sensitivity of optimal actions to their rationalizing beliefs, and thus, specifies a joint test for the state-dependent utility function and the set of feasible actions. If the optimal action always shifts with a player's beliefs, then the decision problem is *responsive* and otherwise, it is *unresponsive*. An unresponsive decision problem is *certainty-unresponsive* if an action that is optimal when a player is completely certain remains optimal whenever she is near-certain.

The prototypical example of the responsive decision problem is the quadratic-loss prediction problem studied by Lee (1993) and others.[1] The quintessential example of a certainty-unresponsive decision problem is any decision problem with a finite action space; however, a decision problem with a continuous action space can also be certainty-unresponsive.

The main result of this paper establishes the following:

Learning is complete in a certainty-unresponsive decision problem if and only if the information structure induces unbounded likelihood ratios. Learning is complete in a responsive decision problem regardless of the information structure.

In other words, all decision problems that are certainty-unresponsive inherit the behavior of models with finite action spaces (identified by Smith and Sørensen, 2000), and responsive decision problems inherit the behavior typically attributed to continuous action spaces. When a decision problem is certainty-unresponsive, once players are nearly certain of a state, their actions cease to reveal their belief. Accordingly, beliefs remain fixed, unless it is the case that there exists a signal realization that can overturn that near-certainty. Thus, an information structure with unbounded likelihood ratios is needed (and suffices) to ensure that learning does not stop until complete certainty of the state is attained.

When the likelihood ratios are bounded, players eventually cease to respond to their information even before reaching complete certainty. In this case, learning is incomplete. One may wonder in this case whether learning is still "adequate" (Aghion et al., 1991): do players learn all that is relevant for their payoffs? Because limit beliefs place strictly positive probability on both states of the world, and the optimal action in each state is distinct, players' learning stops at a place where they would still value learning more. Learning is thus inadequate.

By contrast, when the decision problem is responsive, every player reveals her information through her optimally chosen action. Because each player can perfectly infer the full history of prior signal realizations, observational learning reduces to pure statistical learning. By the Strong Law of Large Numbers, learning is necessarily complete.

The goal of this paper is to develop and show that the lens of responsiveness clarifies and unifies results on coarse and rich action spaces in observational learning. Related ideas appear in several other papers. 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." Arieli and Mueller-Frank (2017) show more generally that decision problems are responsive for all but a meager set of continuous utility functions so long as the set of feasible actions contains no isolated points. Arieli (2017) studies the role of responsiveness in a model of observational learning across generations in which each generation has payoff-interdependencies. My work complements these prior results in that I focus on the standard observational learning framework to revisit the distinction between coarse and rich action spaces. Apart from developing the notion of certainty-unresponsive, the contribution here is to unify results on coarse and rich action spaces in the language of responsiveness. A companion paper (Ali, 2017) highlights the role of responsiveness when information is costly: learning is inhibited by unresponsiveness because any experiment that cannot "swing" a player's action has no value, and so information has to be sufficiently persuasive to have value. By contrast, in responsive decision problems, a player values all kinds of information because she always has a motive to "tweak" her action.[2]

2. Model

2.1. Setup

Decision problem: Each of an infinite sequence of players $t = 1, 2, 3, \ldots,$ chooses from $\mathcal{A}$, the set of feasible actions, which is a compact subset of $\Re$ with at least two distinct elements. The payoff of each action depends on the state of the world ω, which is either "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. Without loss of generality, I assume that no action is weakly dominated, and therefore, 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. The lowest and highest actions in $\mathcal{A}$ are denoted by $\underline{a}$ and $\overline{a}$ respectively, and these are optimal in $\omega = 0$ and $\omega = 1$ respectively. A *decision problem* is a pair $(\mathcal{A}, u)$.

Beliefs and information: Players share a common prior that ascribes probability $\pi \in (0, 1)$ to the state being high. Each player obtains information about the state: given the state ω, an individual obtains a signal realization s_i in $[0, 1]$ that is governed by a conditional c.d.f. $F(\cdot, \omega)$, independently of the signal realizations of other players. No signal realization perfectly reveals the state: the distributions $F(\cdot, 0)$ and $F(\cdot, 1)$ are mutually absolutely continuous, and have common support Σ. As is conventional, I normalize the realization to be the posterior probability that would be ascribed to the high state were the prior neutral.[3]

The information structure is at least moderately informative: there exists $p \in [0, 1]$ such that $F(p, 0) \neq F(p, 1)$. The closure of the convex hull of Σ is denoted $\left[\underline{p}, \overline{p}\right]$. An information structure induces *bounded beliefs* if $0 < \underline{p} \leq \overline{p} < 1$ and induces *unbounded beliefs* if $\underline{p} = 0$ and $\overline{p} = 1$. To simplify exposition, I assume that $F(\cdot, \omega)$ is continuously differentiable for each ω.

Histories and equilibrium: Each player observes actions of all predecessors but not their information. The public history observed by player j is $h^j \equiv (a_i)_{i=1,\ldots,j-1}$. I study Perfect Bayesian Equilibria (henceforth PBE). For a PBE, and equilibrium history h^j, let $\mu(h^j) = Pr(\omega = 1|h^j)$ summarize the *public belief* after history h^j. Consider a set $\mathcal{H}$ of infinite length (equilibrium) histories, and for such a history h_∞, let h^j_∞ be its truncation to actions in periods $1, \ldots, j-1$. For $\omega \in \{0, 1\}$, let $\mathcal{H}_\omega$ denote the set of histories in $\mathcal{H}$ such that $\lim_{j\to\infty} \mu(h^j_\infty) = \omega$. 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*.

1 A partial list is Vives (1993), Eyster and Rabin (2010, 2014), and Guarino and Jehiel (2013).

2 Other analyses of social learning study variations of the standard framework that ensure that "collective behavior" is responsive even if the action space is discrete. Avery and Zemsky (1998) highlight how prices that adjust with the public history can ensure that individuals' pricing choices are responsive to their information. Goeree et al. (2006) augment payoffs with private taste shocks for either action and show that such shocks can induce a strictly positive fraction of types to be sensitive to their information. Analogously, Eyster et al. (2014) study how congestion costs may ensure that some agent behaves in a "responsive" way infinitely-often. Eyster and Rabin (2014) study the quadratic-loss prediction problem across general observational structures, and impose a condition that ensures that the record of actions reveals information.

3 In other words, if F is differentiable at p, and $p \in \Sigma$, $p = \frac{f(p,1)}{f(p,1)+f(p,0)}$.

2.2. Responsiveness of the decision problem

For a belief that places probability μ on the state being high, let $a^*(\mu)$ denote a maximizer of $\mu u(a,1)+(1-\mu)u(a,0)$.[4] This optimal action rule $a^*(\mu)$ is non-decreasing in μ because $u(a,0)$ is decreasing in a and $u(a,1)$ is increasing in a. Because $\mathcal{A}$ contains no weakly dominated actions, $\underline{a}$ is uniquely optimal when $\mu=0$, $\overline{a}$ is uniquely optimal when $\mu=1$, and for each $a\in\mathcal{A}$, there exists μ such that $a=a^*(\mu)$.

Definition 1. A decision problem $(\mathcal{A},u)$ is *responsive* if for every μ,ν in $[0,1]$ where $\mu\neq\nu$, $a^*(\mu)\neq a^*(\nu)$, and otherwise, it is *unresponsive*.

Being responsive is a demanding condition for a decision problem: any change in belief induces a player to "tweak" her action. Formally, it demands that $a^*:[0,1]\to\mathcal{A}$ be injective, and so a rational player's belief can be inferred from her action.[5] A necessary condition for a $(\mathcal{A},u)$ to be responsive is that $\mathcal{A}$ is isomorphic to $[0,1]$, but that is not sufficient, as Example 2 below indicates. Prior to describing the examples, it is worth highlighting a particular failure of responsiveness, namely when it fails near-certainty. This is described below as "certainty-unresponsive."

Definition 2. A decision problem $(\mathcal{A},u)$ is *certainty-unresponsive* if an extreme action, $\underline{a}$ or $\overline{a}$, is optimal at an interior belief: $a^*(\mu)\in\{\underline{a},\overline{a}\}$ for $\mu\in(0,1)$.

A particular example of a certainty-unresponsive decision problem is when $\mathcal{A}$ is finite: each weakly undominated action is optimal over a range of beliefs, including those that are optimal when players are completely certain of the state. But Example 2 and examples in Chari and Kehoe (2004) and Ali (2017) indicate that even continuous action investment problems may be certainty-unresponsive.

I illustrate notions of responsiveness and unresponsiveness using the following examples.

Example 1. Consider the prediction problem studied by Lee (1993): suppose $\mathcal{A}=[0,1]$ and $u(a,\omega)=-(a-\omega)^2$. The optimal action rule is $a^*(\mu)=\mu$, and so $(\mathcal{A},u)$ is a responsive decision problem.

The prediction problem described above is the paragon of responsiveness. However, a slight adaptation of this action space results in the decision problem becoming certainty-unresponsive.

Example 2. Truncate $\mathcal{A}$ defined in Example 1 to $\mathcal{A}'=[\varepsilon,1]$ for some $\varepsilon\in(0,1)$. Then $a^*(\mu)=\min\{\varepsilon,\mu\}$, and thus, the optimal action chosen when the player is certain that $\omega=0$ remains the same when the player ascribes probability at least $1-\varepsilon$ to $\omega=0$. Therefore, the decision problem $(\mathcal{A}',u)$ is certainty-unresponsive.

The action spaces in Examples 1 and 2 are both isomorphic to the space of beliefs, $[0,1]$, and to each other. Nevertheless, the decision problems exhibit different behaviors across beliefs, and this difference – as we shall see in Theorem 1 – results in different implications for herding.

[4] If there are multiple optimal actions at a belief μ, select the lowest one.

[5] In this vein, this work connects to the literature on eliciting beliefs from actions (Savage, 1971; Karni, 2009), where the objective is to design decision problems in which actions reveal beliefs.

3. The main result

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

(a) *If $(\mathcal{A},u)$ is responsive, then learning is complete.*
(b) *If $(\mathcal{A},u)$ is certainty-unresponsive, then learning is complete if the information structure induces unbounded beliefs, and is incomplete if the information structure induces bounded beliefs.*[6]
(c) *If $(\mathcal{A},u)$ is unresponsive, then there exists an open and dense set of prior-signal combinations such that learning is incomplete.*

Responsiveness plays an intuitive role: do actions fully reveal information? When $(\mathcal{A},u)$ is responsive, actions are fully revealing and this perpetual accumulation of information inexorably concentrates public belief to the truth (a.s.), as in Lee (1993). Otherwise, if the decision problem is unresponsive, information is lost because an individual's action is a coarse signal of his beliefs (even if his action space is a continuum). When this coarseness manifests at extreme beliefs – as it does in certainty-unresponsive decision problems – extreme signal realizations are needed to sway actions from an inefficient herd, exactly as in the finite-action case studied by Smith and Sørensen (2000).

Proof of Theorem 1. Let $B(\mu,p)$ denote the posterior probability that the state is high when the prior is μ and the signal realization is p. For a measurable set of actions A, let

$$P(A,\mu)\equiv\{p\in\Sigma : a^*(B(\mu,p))\in A\}$$
$$\text{and}\quad \alpha(A,\mu,\omega)\equiv\int_{P(A,\mu)} dF(p,\omega).$$

I define the "support" of actions:

$$\overline{\mathcal{A}}(\mu)=\{a\in\mathcal{A} : \alpha((a-\varepsilon,a+\varepsilon)\cap\mathcal{A},\mu,\omega)>0 \text{ for every } \varepsilon>0 \text{ and } \omega\in\{0,1\}\}.$$

Let $p^*(a,\mu)$ and $p_*(a,\mu)$ be the sup and inf of $P(\{a\},\mu)$ respectively. Since $F(\cdot,\omega)$ is continuously 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 $p^*(a,\mu)\neq\underline{p}$ or $p_*(a,\mu)\neq\overline{p}$, $\beta(a,\mu)$ is continuous in μ. Define the cascade set of beliefs to be $\mathcal{C}\equiv\bigcup_{a\in\mathcal{A}}\{\mu\in[0,1] : P(\{a\},\mu)=\Sigma\}$.

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

Proof. If $\mu\in\mathcal{C}$, then the definition of $\mathcal{C}$ implies that there exists an action a such that $P(\{a\},\mu)=\Sigma$. Therefore, for each ω, $\alpha(A,\mu,\omega)=1$ if $a\in A$, and $\alpha(A,\mu,\omega)=0$ if $a\notin A$. Suppose that $\mu\notin\mathcal{C}$. Then there exists an action $\tilde{a}$ such that $F(p^*(\tilde{a},\mu),\omega)\in(0,1)$, and consider the set of actions $[\underline{a},\tilde{a}]$: by Lemma A.1 of Smith and Sørensen (2000), it follows that $F(p^*(\tilde{a},\mu),1)<F(p^*(\tilde{a},\mu),0)$. □

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 $l_0^t(h^t)=1/l_1^t(h^t)$. Treat $\langle l_i^t(\cdot)\rangle_{t=1}^{\infty}$ as a stochastic process, and note 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 $[0,\infty)$.

[6] This result is not the tightest possible because the environment, as in Smith and Sørensen (2000), assumes that private beliefs are bounded away from *both* 0 and 1 or from *neither*. A tighter result is possible when one allows for information structures that are bounded away from one but not the other. Say that a decision problem is *certainty-unresponsive* at $\omega\in\{0,1\}$ if $a^*(\mu)=a^*(\omega)$ for $\mu\in(0,1)$. Also, say that private beliefs are unbounded at $\omega\in\{0,1\}$ if $\omega\in[\underline{p},\overline{p}]$. Then the argument of Theorem 1(b) establishes the following tighter claim: *Suppose that the decision problem is certainty-unresponsive at ω. Then learning is complete at $\omega'\neq\omega$ if and only if private beliefs are unbounded at ω'.*

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 $\mathcal{C}$ if $i = 1$, and $\frac{l}{1+l}$ is a subset of $\mathcal{C}$ 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 $\mathcal{C}$. Consider action $\tilde{a}$ such that $F(p^*(\tilde{a}, \mu), \omega) \in (0, 1)$, 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 $\mathcal{C} = \{0, 1\}$, 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 certainty-unresponsive. If the information structure induces unbounded beliefs, $\mathcal{C} = \{0, 1\}$, and so as above, $Pr(l_i^\infty = 0 \mid \omega = i) = 1$. Now suppose that the information structure induces bounded beliefs. Let $\underline{\mu}$ be the highest belief such that $a^*(\underline{\mu}) = \underline{a}$; analogously, let $\overline{\mu}$ be the lowest belief such that $a^*(\overline{\mu}) = \overline{a}$. We consider the following cases below.

1. Suppose that $0 < \underline{\mu} < \overline{\mu} < 1$. Define
$$l_* \equiv \left(\frac{1-\underline{\mu}}{\underline{\mu}}\right)\left(\frac{1-\underline{p}}{\underline{p}}\right), l^* \equiv \left(\frac{1-\overline{\mu}}{\overline{\mu}}\right)\left(\frac{1-\overline{p}}{\overline{p}}\right).$$
It follows that once l_1^t enters $[0, l^*] \cup [l_*, \infty]$, all subsequent players 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. The argument is analogous for $\underline{\mu} = 0$ but $\overline{\mu} < 1$ 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 $B(\pi, \overline{p}) < \overline{\mu}$ and $B(\pi, \underline{p}) > \underline{\mu}$. For such combinations, every player chooses action a regardless of her signal realization. □

References

Aghion, P., Bolton, P., Harris, C., Jullien, B., 1991. Optimal learning by experimentation. Rev. Econ. Stud. 58 (4), 621–654.

Ali, S.N., 2017. Herding with Costly Information, Working Paper.

Arieli, I., 2017. Payoff externalities and social learning. Games Econ. Behav. 392–410.

Arieli, I., Mueller-Frank, M., 2017. Inferring beliefs from actions. Games Econ. Behav. 102, 455–461.

Avery, C., Zemsky, P., 1998. Multi-dimensional uncertainty and herd behavior in financial markets. Am. Econ. Rev. 88 (4), 724–748.

Banerjee, A., 1992. A simple model of herd behavior. Q. J. Econ. 107 (3), 797–817.

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

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

Eyster, E., Galeotti, A., Kartik, N., Rabin, M., 2014. Congested observational learning. Games Econ. Behav. 87, 519–538.

Eyster, E., Rabin, M., 2010. Naive herding in rich-information settings. American Economic Journal: Microeconomics 2, 221–243.

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

Goeree, J.K., Palfrey, T.R., Rogers, B.W., 2006. Social learning with private and common values. Econ. Theory 28, 245–264.

Guarino, A., Jehiel, P., 2013. Social learning with coarse inference. American Economic Journal: Microeconomics 5, 147–174.

Huck, S., Oechssler, J., 1998. Informational cascades with continuous action spaces. Economics Letters 60, 163–166.

Karni, E., 2009. A mechanism for eliciting probabilities. Econometrica 77, 603–606.

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

Savage, L.J., 1971. Elicitation of personal probabilities and expectations. J. Am. Stat. Assoc. 66, 783–801.

Smith, L., Sørensen, P., 2000. Pathological outcomes of observational learning. Econometrica 68, 371–398.

Vives, X., 1993. How fast do rational agents learn? Rev. Econ. Stud. 60, 329–347.