

Overconfidence and Informational Cascades

Shachar Kariv[†]

University of California, Berkeley

July 15, 2004

Abstract

This paper combines behavioral economics and social learning. We test how robust the theory is to the well-known behavioral phenomenon of individual overconfidence (mistaken private information perception, specifically overvaluing). In the context of social learning, overconfident agents overweigh their private information relative to the public information revealed by the decisions of others. Therefore, when following a herd, they broaden more of the information available to them. However, overconfidence trades the additional information revealed by overconfident decisions against more information that is being suppressed by perfectly rational decisions. Thus, the presence of overconfident individuals intensifies the free-rider problem of rational individuals. With the help of numerical simulations, we show that, from a social perspective, the presence of overconfident agents cannot improve decisions accuracy or break the poor information flow intrinsic to erroneous uniform behavior.

I acknowledge helpful discussions of Andrew Caplin, Kfir Eliaz, Douglas Gale, David Hirshleifer, Botond Koszegi, Yaw Nyarko, Ben Polak, Matthew Rabin and Andrew Schotter. This paper has benefited from suggestions by the participants of seminars at several universities.

[†]Department of Economics, University of California, Berkeley, Evans Hall # 3880, Berkeley, CA 94720-3880 (E-mail: kariv@berkeley.edu; URL: www.socrates.berkeley.edu/~kariv).

1 Introduction

The standard model of social learning, first studied by Banerjee (1992) and Bikhchandani, Hirshleifer and Welch (1992) (BHW), and extended by Smith and Sørensen (2000)¹, comprises a set of agents, a finite set of actions, a set of states of nature, and a common payoff function which depends on the agent's own action and on the state of nature. Each agent receives a private signal, a function of the state of nature, and uses this private information to identify a payoff-maximizing action. Thus, each agent's action reveals some information about her private signal, so an agent can generally improve her decision by observing what others do before choosing her own action.

The models of social learning assume a simple sequential structure, in which the order of play is fixed and exogenous. Agents make decisions in some predetermined order, and when an agent's turn comes, she can observe the actions of each of her predecessors. Each agent has received a private signal and her choice of action may reveal this. Therefore, by observing the decisions of her predecessors, any agent may be able to infer their private information and learn from them. In social settings, when agents can observe one another's actions, it is rational for them to learn from each other.

The main results of the social learning literature, are that, despite the asymmetry of information, eventually every agent imitates her predecessor, even though she would have chosen a different action on the basis of her own information alone. In this sense, agents ignore their own information and join the herd. Furthermore, once an agent decides to join the herd, her own information is suppressed, so despite the available information, herds often adopt a suboptimal action. This failure of information aggregation is explained by two facts. First, an agent's action is a coarse signal of her private information and, secondly, after some point, agents suppress their private information and join the herd, so very little information is ever revealed.

This is an important result that helps us understand the basis for uniformity of social behavior. At the same time, the standard model of social learning has several special features that deserve further examination. A central assumption of nearly all social learning models is rational behavior. The agent is rationally comparing her information with that of a large (in the limit, unboundedly large) number of other agents. However, given the experimental evidence showing that individuals' behavior is often odds with Bayesian updating, it is not obvious how social learning phenomena can be captured plausibly in a model based on perfect individual rationality.

Motivated by the large body of evidence from psychological surveys and experiments of the importance of overconfidence, this paper tests how robust the theory is to the well-known behavioral phenomenon of individual overconfidence. In the context of social learning, overconfident agents overweigh their private information relative to the public information revealed by the decisions of others.

¹For surveys see, Gale (1996) and Bikhchandani, Hirshleifer and Welch (1998). Among others, Lee (1993), Chamley and Gale (1994), Gul and Lundholm (1995), Moscarini, Ottaviani and Smith (1998), and Çelen and Kariv (2004a) provide further extensions of the theory.

That is, in Bayesian terms, agents weigh their own information too heavily and give too little weight to the public information. The idea of overconfidence amounts to the judgment bias of base rate fallacy or base rate neglect in the sense that more overconfident agents have more severe base rate fallacies.

A great deal of attention has been paid to overconfidence in the behavioral economics literature. The kind of psychological regularity that is characteristic of overconfidence makes sense in a wide range of situations². The common meaning of overconfidence in these works is a mistake in self-judgment or favoring private information. In fact, because such a large body of evidence of overconfidence has been found in various circumstances³, DeBondt and Thaler (1995) argue that “perhaps the most robust finding in the psychology of judgment is that people are overconfident.”

Laboratory experiments provide a clean test of social learning models by minimizing potentially confounding effects. In a seminal paper, Anderson and Holt (1997) investigate the model of BHW experimentally and demonstrate that cascade behavior can be replicated in the laboratory. Following their pioneering work, a number of experimental papers that have analyzed aspects of social learning, e.g., Nöth and Weber (2001), Çelen and Kariv (2004b) and Goeree, Palfrey, Rogers, and McKelvey (2004), find that overconfidence is a significant factor in the dynamics.

To focus on the implications of overconfidence for social learning, we exclude ad hoc learning rules or other forms of bounded rational learning that could also be relevant to issues of social learning so that any biases result exclusively from overconfidence. Thus, our attention is restricted to Bayesian behavior except for overwriting of private information. This suggests that overconfidence may mitigate the tendency toward herd behavior and raises some important questions: Is imitation resulting from rational choice more likely to occur in the presence of overconfident behavior? Can rational agents learn enough to make good decisions by observing the choices of only a small number of overconfident agents? Is there a higher chance that a group achieves a desirable outcome because of the additional dissemination of information?

The idea of studying social learning with overconfident agents is not new, having been explored by Bernardo and Welch (2001) (BW) who study the relationship between overconfidence and entrepreneurship in the binary-signal model of BHW. In our model, unlike BW, there is a continuous-signal space which allows for a richer pattern of social learning⁴. BW conclude that overconfidence creates a large positive externality on the accumulated public information. This paper appears to reverse their conclusion. But the paper should not be seen as a test on BW per se as the models are different in few aspects.

²See Daniel, Hirshleifer and Subrahmanyam (1998), Camerer and Lovo (1999), Benabou and Tirole (1999, 2000), among others.

³Kent, Hirshleifer and Subrahmanyam (1998) study overconfidence in the context of financial markets and provide (Section I) a summary of the psychological evidence of overconfidence.

⁴Smith and Sørensen (2000) extend the basic model of BHW to allow for richer information structures and provide a more general and precise analysis of the convergence of actions and beliefs.

Rather the point of this paper is that under different modeling frameworks, the overconfidence story can be conceptually different. We provide headway in integrating the results later in the paper.

In the model, there are rational agents who choose their optimal action condition on the information available to them in a Bayesian way, and overconfident agents who are Bayesians except for their tendency to overweigh their private information relative to the information revealed from the history of actions. Similar definitions of overconfidence are also employed in some recent works in economics and finance including BW. We describe the rational agents' optimal strategies recursively; they in turn, characterize the dynamics of learning and actions.

With only rational agents, we replicate the results of the literature and use them as a benchmark. The learning process has the martingale property which allows us to establish convergence of beliefs and actions. An informational cascade (convergence of beliefs in finite time) need not arise but herd behavior (convergence of actions in finite time) must⁵. Thus, agents forever take into account their private signals in a non-trivial way. Adding overconfident agents has no effect on convergence, as the rational agents' learning process does not lose its martingale property. The reason is that overconfident agents use dominated strategies that convey more of their private information.

However, overconfidence trades the additional information revealed by overconfident decisions against more information that is being suppressed by perfectly rational decisions. In fact, by making more of their private information public, the overconfident agents trigger a cascade stage, that otherwise need not start. Hence, the presence of overconfident agents intensifies the free-rider problem of rational agents, and because of the surfeit of information revealed to the group from overconfident actions, in some finite time, the learning process tips into an informational cascade.

An important concept discussed in Smith and Sørensen (2000) is the overturning principle. This asserts that even if many agents have acted alike, it is possible that, because of a rational deviation, the information revealed from the history up until that point cancels out. The reason is that it is optimal to follow a rational deviator because she processes all previously revealed information. In contrast, when the overturning party is overconfident, the same intuition implies that she may not be followed. With an overconfident deviator, her successor should not be so willing to follow in her shoes as she is not a proper Bayesian.

While the convergence properties of the model are quite general, other properties have only been established for particular type of overconfidence. With the help of numerical simulations, we explore further the qualitative features of the informational tradeoff caused by overconfidence. We show that from a social perspective, the presence of overconfident agents cannot break the poor information flow intrinsic to erroneous uniform behavior or improve decisions accuracy and welfare. This result is fully in line with the message of the stan-

⁵ Informational cascades occur when, after some finite time, all agents ignore their private information when choosing an action, while herd behavior occurs when, after some finite time, all agents choose the same action, not necessarily ignoring their private information.

standard model of social learning that too much public information can be socially harmful.

The rest of the paper is organized as follows. The next section outlines the model. Section 3 describes the case of only rational agents as a benchmark. Section 4 adds overconfident agents. Section 5 provides closed forms and comparative statics with a uniform signal distribution, and section 6 concludes. Proofs are gathered in Section 7.

2 The Model

To illustrate the basic workings of overconfidence, it is necessary to consider a special information and payoff structure. The specific model which we analyze builds on Gale (1996). There is a finite number of agents indexed by $i = 1, 2, \dots, n$. Each agent i makes a once-in-a-lifetime decision indicated by $a_i \in \{0, 1\}$. Decisions are made sequentially in an exogenously determined order. All decisions are announced publicly and therefore known to all successors.

The preferences of the agents are assumed to be identical and represented by

$$u(a_i) = \begin{cases} \frac{1}{2} & \text{if } a_i = 1 \\ 0 & \text{if } a_i = 0 \end{cases}$$

where θ_i is a random variable defined by

$$\theta_i = \prod_{j=1}^n \theta_j$$

and θ_i is agent i 's private signal about θ . We assume that the θ_i 's are identically and independently distributed with c.d.f. F over a compact support with convex hull $[-1, 1]$, F has no atoms and satisfies symmetry,

$$F(\theta) = 1 - F(-\theta), \quad \theta \in [-1, 1].$$

Without loss of generality, we assume that $\theta = 1$.

The summation version of θ makes the model nicely tractable, but some clarifications are in order. This setup provides an example of a pure information externality. Each agent's payoff depends on his own action and on the state of nature. It does not depend directly on the actions of other agents. In Appendix A, we summarize this development and explore the relation to the general information structure studied by Smith and Sørensen (2000).

There are two types of agents in the model: rational who choose their optimal action condition on the information available to them in a Bayesian way, and overconfident who are Bayesians except for their tendency to overweigh their private information relative to the information revealed from the history of actions.

We assume that a fraction p of agents are overconfident, and that whether an agent's behavior is overconfident is unobservable by others and that overcon-

fidence is distributed independently across agents⁶. The traits of overconfidence will be stated more precisely after providing the definitions of some key concepts. Finally, it is assumed that the information structure, individual types and all decision rules are common knowledge.

3 The Case of No Overconfidence

We next replicate the results of the literature with only rational agents ($p = 0$) and use them as a benchmark. The key result gives a sufficient condition on the signal distribution guaranteeing the impossibility of informational cascades.

We first provide a definition that will be useful in characterizing the optimal strategy.

Definition 1 agent i follows a cutoff strategy if her decision rule is defined by

$$a_i = \begin{cases} 1 & \text{if } \theta_i \geq \theta_i^c \\ 0 & \text{if } \theta_i < \theta_i^c \end{cases}$$

for some cutoff θ_i^c ⁷.

The decision problem of agent i is

$$\text{Max}_{a_i \in \{0,1\}} a_i \mathbf{E}[U(a_i, \theta) \mid \theta_i, I_i]$$

where $I_i = I(\{a_j : j < i\})$ which can be summarized as

$$a_i = 1 \text{ if and only if } \mathbf{E}^P[\theta_{j=n} \mid \theta_i, I_i] \geq 0.$$

Since θ_i and I_i do not provide any information about the content of successors' signals, we obtain

$$a_i = 1 \text{ if and only if } \mathbf{E}^P[\theta_{j=i} \mid \theta_i, I_i] \geq 0,$$

and thus,

$$a_i = 1 \text{ if and only if } \theta_i \geq -\mathbf{E}^P[\theta_{j < i} \mid I_i].$$

It readily follows that the optimal decision takes the form of a cutoff strategy, which we state in the next proposition.

Proposition 1 For any agent i , the optimal strategy is the cutoff strategy

$$a_i = \begin{cases} 1 & \text{if } \theta_i \geq \hat{\theta}_i \\ 0 & \text{if } \theta_i < \hat{\theta}_i \end{cases} \quad (1)$$

where

$$\hat{\theta}_i = -\mathbf{E}^P[\theta_{j < i} \mid I_i] \quad (2)$$

is the optimal history-contingent cutoff.

⁶For computational simplicity, BW assume, in contrast, that the type of each agent is public knowledge. But their main message that overconfident agents can provide socially useful information is unchanged when types are unknown.

⁷Notice that the tie-breaking assumption is such that $a_n = 1$ when $\theta_i = \theta_i^c$, but these are probability zero events.

The optimal history-contingent cutoff \hat{c}_i is sufficient to characterize agent i 's behavior, and thus the cutoff process $\{\hat{c}_i\}_{i=1}^n$ characterizes the social behavior. In Appendix B, we describe agents' behavior formally and discuss the essential elements of the weak perfect Bayesian equilibrium. It follows directly that a cutoff equilibrium, i.e., an equilibrium in which all follow a cutoff strategy is a weak perfect Bayesian equilibrium since a_i given by (1) is a best response to \hat{c}_i after every possible history of actions $(a_j)_{j < i}$ and \hat{c}_i is derived in (2) via Bayes' rule.

Now, we are ready to define informational cascades and herd behavior, two notions introduced by Banerjee (1992) and BHW to address the same phenomenon but Smith and Sørensen (2000) emphasize the difference between them.

For any agent i and each action $a_i \in \{0, 1\}$, we call the set of cutoffs C_{a_i} such that $C_1 = (-\infty, -1]$ and $C_0 = [1, \infty)$ the a_i -cascade set. Note that an agent who engages who sets her cutoff at C_1 or C_0 takes either action 1 or 0, no matter what her private signal is. In contrast, an agent who reports a cutoff in the interval $(-1, 1)$, indicating that there are some signals that can lead her to choose action 1, some to choose 0.

Definition 2 (informational cascades) An informational cascade on action $a = 1$ ($a = 0$) occurs when \hat{c}_i such that $\hat{c}_j \in C_1$ (C_0) $\forall j < i$.

Analogously, a limit-cascade on action $x = 1$ ($x = 0$) occurs when the process of cutoffs $\{\hat{c}_i\}$ converges almost surely to a random variable $\hat{c} = \lim_n \hat{c}_n$, with $\text{supp}(\hat{c}) \subset C_1$ ($\text{supp}(\hat{c}) \subset C_0$).

Further, we call a finite sequence of agents who act alike a finite herd and, we let

$$l_i^j = \#\{a_j = a_i, i < j \leq n\}$$

denote the length of a finite herd following agent i in a group of size n .

Definition 3 (Herd behavior) Herd behavior occurs when \hat{c}_i such that

$$\lim_n l_i^n / n = 1.$$

Hence, a cascade implies herd behavior but herding is not necessarily the result of an informational cascade. When acting in a herd, agents choose the same action, but they could have acted differently from one another if the realization of their private signals had been different. In an informational cascade, an agent considers it optimal to follow the behavior of her predecessors without regard to her private information.

The first step in the analysis is to establish convergence of the cutoff process $\{\hat{c}_i\}$. From the definition of equilibrium (see Appendix), we know that $l_i \leq l_{i+1} \leq l$, i.e., agents' public information is non-decreasing over time. Hence, the equilibrium payoffs must be non-decreasing over time and, since it is bounded, must converge. We use this result to establish convergence of cutoffs.

Proposition 2 Let $\{\hat{c}_i, l_i\}$ be an equilibrium. Then $\{\hat{c}_i\}$ is a martingale with respect to $\{l_i\}$ and there exists a random variable \hat{c} such that \hat{c}_n converges to \hat{c} almost surely.

In words, $\{\hat{\pi}_i\}$ has the martingale property (following from the martingale property of conditional expectations) so by the Martingale Convergence Theorem, it converges almost surely to a random variable $\hat{\pi} = \lim_n \hat{\pi}_n$. Hence, it is stochastically stable in the neighborhood of the fixed points, -1 and 1 , meaning that there is a limit-cascade. However, since convergence of the cutoff process implies convergence of actions, behavior can not overturn forever. In other words, behavior settles down in some finite time and is consistent with the limit learning.

Hence, with no overconfidence, we agree with Smith and Sørensen (2000) that a cascade need not arise but a limit-cascade and herd behavior must⁸. However, the rate of convergence of the cutoff process and its limiting value depend on what the exact realization of private signals is⁹. Moreover, what the Martingale Convergence Theorem does not imply is converges in finite time.

As for the rate of convergence, consider the law of motion for $\hat{\pi}_i$ when $a_{i-1} = 1$

$$\hat{\pi}_i = \hat{\pi}_{i-1} - \mathbf{E}^+(\hat{\pi}_{i-1}).$$

By symmetry $\mathbf{E}^+(\hat{\pi}_{i-1}) > 0$ and thus $\hat{\pi}_i < \hat{\pi}_{i-1}$ and the inequality is strict as long as $\hat{\pi}_{i-1} > -1$. Notice that the relation $\hat{\pi}_i = \mathbf{E}^+(\hat{\pi}_{i-1})$ is continuous $[-1, 0]$ and $\mathbf{E}^+(-1) = -1$. Furthermore, $\mathbf{E}^+(\hat{\pi}_i) < \hat{\pi}_i$ and the inequality is strict as long as $\hat{\pi}_i < 1$ which implies that $\hat{\pi}_i > -1$. In the next proposition we state a simple sufficient condition on the signal distribution such that a cascade never starts.

Proposition 3 An action 1-cascade is impossible, i.e., $\hat{\pi}_i \neq C_1$ for all i , if

$$F\left(\frac{C_1 + 1}{C_1 + 2}\right) < \frac{C_1 + 1}{C_1 + 2}$$

for any $C_1 \in [-1, 0]$, and an analogous argument applies to action 0-cascade.

In words, there is a constriction on the mass of probability near the edge of the signal support as otherwise a single decision can guide all subsequent agents to a cascade. Nevertheless, this constriction leaves a very broad class of distributions for which an informational cascade with only rational agents is impossible.

Our conjecture is that the impossibility of cascades holds for the model even more generally for a very wide range of signal distributions. However, like Smith and Sørensen (2000) we know of no simple necessary condition that guarantees the impossibility of cascades. Note, however, that with the atomic tails in the discrete-signal setup of BHW informational cascade arises.

Summarizing,

⁸Smith and Sørensen (2000) introduce the Martingale Convergence Principle to show that social learning eventually leads to convergence of actions.

⁹We provide a characterization for the case of bounded beliefs (private beliefs are said to be bounded when there is no private signal that can perfectly reveal the state of the world and to be unbounded otherwise). With unbounded beliefs, Smith and Sørensen (2000) show that learning leads to correct decisions, and with bounded beliefs, we agree with them that what is learned can be incorrect.

Theorem 1 (No overconfidence) With only rational agents, (i) an informational cascade need not arise, i.e., $\hat{c}_i \notin \{C_0, C_1\}$ for all i , and (ii) a limit-cascade, i.e., $\hat{c} = \lim_n \hat{c}_i$, with $\hat{c} \in \{C_0, C_1\}$, implies that a herd on the corresponding action almost surely arises.

4 The Case of Overconfidence

Now, we add overconfident agents ($p > 0$). The main result is that overconfidence can cause agents to act on their private information more strongly than rationally, thereby exacerbating the possibility of informational cascades.

Stated in the language of the previous section, the following definition of overconfidence serves as the working definition.

Definition 4 (Overconfidence) Agent i is said to be overconfident if for any history of actions $(a_j)_{j < i}$ her decision takes the form of the cutoff strategy

$$a_i = \begin{cases} 1 & \text{if } \hat{c}_i \geq \tilde{c}_i \\ 0 & \text{if } \hat{c}_i < \tilde{c}_i \end{cases} \quad (3)$$

where \tilde{c}_i is determined by the commonly known mapping $\hat{c} \mapsto \tilde{c}(\hat{c})$ such that (i) $\tilde{c}(\hat{c}) \in [0, \hat{c})$ for any $\hat{c} \in [0, 1)$ and $\tilde{c}(\hat{c}) = \hat{c}$ for any $\hat{c} \in C_0$, (ii) $\tilde{c}(\hat{c})$ is symmetric, i.e., $\tilde{c}(-\hat{c}) = -\tilde{c}(\hat{c})$, and (iii) \hat{c} is the rational history-contingent cutoff rule given p and $\hat{c} \mapsto \tilde{c}(\hat{c})$.

Overconfident agents' perception about the precision of their private information is overoptimistic, and thus overweigh it relative to the public information revealed by the behavior of others. Put differently, in Bayesian terms, overconfident agents weigh their own information too heavily and give too little weight to public information. Note that although whether an agent is overconfident is unobservable by others, the fraction of overconfident agents p and the degree of overconfidence \tilde{c} are public knowledge.

As Figure 1 illustrates, if a herd of action $a = 1$ precedes a rational agent i then according to (1) and (2) her cutoff \hat{c}_i is close to -1 . Hence, in the subset $[\hat{c}_i, 0)$, which we call an imitation set, private signals are ignored in making a decision and agent i imitates her predecessors' action. Similarly, put side by side with the rational agent, according to (3) the cutoff of an overconfident agent i is $\tilde{c}_i < \hat{c}_i$, and thus she has a smaller imitation set $[\tilde{c}_i, 0) \subset [\hat{c}_i, 0)$.

[Figure 1 here]

Note that overconfidence is symmetric and invariant across histories that reveal the same information. As such, overconfidence does not create a bias toward any action. Furthermore, to avoid trivialities, like BW we assume that overconfidence is bounded in strength, in the sense that it ceases once the history of actions provides a sufficiently decisive guide regarding to what the state is. Thus, informational cascades are unbreakable even in the presence of overconfident agents.

The behavioral postulate is that overconfidence is the agents' only motive other than maximization. Noticeably, there is a trade off between these two motives since they imply different behavior. Therefore, overconfidence is viewed as being irrational or erroneous. However, the tension between overconfidence and maximization is partially reconcilable as overconfidence falls once enough public information is revealed.

With common knowledge of the distribution of types and all decision rules, agent i 's publicly available information I_i^0 is generated by $p, \tilde{v}(\cdot)$ and the history of actions $(a_j)_{j < i}$. Hence, after adding overconfident agents, the optimal decision rule of rational agent i can be summarized as

$$a_i = \begin{cases} 1 & \text{if } v_i \geq \hat{v}_i \\ 0 & \text{if } v_i < \hat{v}_i \end{cases}$$

where the history-contingent cutoff rule is now described by

$$\hat{v}_i = -\mathbf{E}^{h_P} \left[\sum_{j < i} v_j \mid I_i^0 \right] \quad (4)$$

As in the model with no overconfidence, the cutoff \hat{v}_i inherits all the information that a rational agent i learns from the history of actions, and therefore is sufficient to characterize her behavior. Furthermore, Put side by side with a rational agent, an overconfident agent i 's cutoff \tilde{v}_i is given by the mapping $\tilde{v}_i = \tilde{v}(\hat{v}_i)$.

Even in the presence of overconfident agents, the process of rational cutoffs $\{\hat{v}_i\}$ does not lose the martingale property. This has an important implication: rational agents' cutoffs and thus actions are convergent.

Proposition 4 Augment the model by adding overconfident agents and let $\{\hat{v}_i, I_i^0\}$ be an equilibrium. Then $\{\hat{v}_i\}$ is a martingale with respect to $\{I_i^0\}$ and there exists a random variable \hat{v}^* such that \hat{v}_n converges to \hat{v}^* almost surely.

After adding overconfident agents to the model, the proof goes through the same as with only rational agents. Since with overconfident agents the cutoff process of the rational agents (\hat{v}_i) is still a martingale, the previous convergence result is unaffected. The reason behind it is that overconfident agents are in fact equivalent to rational agents with dominant strategies that are at least as informative.

Next, we show that informational cascades start with overconfident agents as the cutoff process easily escapes in some finite time into a-cascade set. For an informational cascade, say on $a = 1$, to start, a discrete jump of (\hat{v}_i) into C_1 must occur in a finite time. With overconfidence, the law of motion for \hat{v}_i when $a_{i-1} = 1$ is given by

$$\hat{v}_i = \hat{v}_{i-1} - (1-p)\mathbf{E}^+(\hat{v}_{i-1}) - p\mathbf{E}^+(\tilde{v}_{i-1})$$

as each rational agent assigns a chance p that she is preceded by an overconfident agent. Thus, it readily follows that to have a cascade on investment we simply need that

$$\hat{v}_{i-1} - \mathbf{E}^+(\hat{v}_{i-1}) - p \mathbf{E}^+(\tilde{v}_{i-1}) > 1$$

Figure 1. Rational and overconfident imitation sets after a finite herd of action 1




Figure 2a. The sequences of cutoffs when all agents choose action 1 with different degrees of overconfidence




Figure 2b. The sequences of cutoffs when all agents choose action 1 with different proportions of overconfident agents




Figure 3a. The probability that a rational agent makes a correct decision with different degrees of overconfidence




Figure 3b. The probability that a rational agent makes a correct decision with different proportions of overconfident agents

