Overconfidence and Informational Cascades*

Shachar Kariv†
UC Berkeley

March 21, 2005

Abstract

This paper combines behavioral economics and 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 broadcast 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 rational decisions. This paper shows that the presence of overconfident agents intensifies the free-rider problem of rational agents, since, even if overconfident agents have very limited information, by making it public, they trigger an uninformative everlasting cascade stage, that otherwise need not start. With the help of numerical simulations, this paper shows that having overconfident agents cannot break the poor information flow intrinsic to erroneous uniform behavior.

Journal of Economic Literature Classification Numbers: D82, D83

*I acknowledge helpful discussions of Andrew Caplin, Stefano DellaVigna, 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 of 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 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

[1] For 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.

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. 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.

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 learning[2]. BW conclude that overconfidence creates a large positive externality on the accumulated public information. This paper reverses their conclusion but the paper should not be seen as a test on BW *per se* as the models are different in several aspects. Rather the point of this paper is that under different modeling frameworks, the overconfidence story can be conceptually different. We integrate the results later in the paper.

To focus on the implications of overconfidence for social learning, we exclude ad hoc learning rules or other forms of bounded rational learning which could be relevant to issues of social learning[3]. This allows us to conclude 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?

[2] Smith 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.

[3] Recent work in this area includes Ellison and Fudenberg (1993, 1995), among others.

In the model, rational agents choose their optimal action conditional on the information available to them in a Bayesian way; overconfident agents are Bayesian except for their tendency to overweigh their private information relative to information revealed from the history of actions. Similar definitions of overconfidence are employed in 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 Smith and Sørensen (2000) 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) must. 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, 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 will 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 types 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 standard model of social learning that too much public information can be socially harmful.

A great deal of attention has been paid to overconfidence in the behavioral economics literature. The common meaning of overconfidence in these works is a mistake in self-judgment or favoring private information. The kind of psychological regularity that is characteristic of overconfidence makes sense in a wide range of situations[4]. Because such a large body of evidence of overconfidence exists[5], 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, 2005) and Goeree, Palfrey, Rogers, and McKelvey (2004), find that overconfidence is a significant factor in the dynamics.

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 $i = 1, 2, \cdots, n$. 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.

[4] See Daniel, Hirshleifer and Subrahmanyam (1998), Camerer and Lovallo (1999), Benabou and Tirole (1999, 2000), among others.

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

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

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

where Ω is a random variable defined by

$$\Omega = \sum_{i=1}^{n} \omega_i$$

and ω_i is agent i's private signal about Ω.

We assume that the ω_i's are identically and independently distributed with *c.d.f.* F over a compact support with convex hull $[-\alpha, \alpha]$, and F has no atoms and satisfies symmetry,

$$F(\omega) = 1 - F(-\omega), \ \forall \omega \in [-\alpha, \alpha].$$

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

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. The summation version of Ω makes the model nicely tractable, but some clarifications are in order. 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 overconfidence is distributed independently across agents[6]. 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 a benchmark. The key result gives sufficient condition

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

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*

$$a_i = \begin{cases} 1 & if \quad \omega_i \geq \omega_i^* \\ 0 & if \quad \omega_i < \omega_i^* \end{cases}$$

for some cutoff [7]*.*

The decision problem of agent i is

$$\underset{a_i \in \{0,1\}}{Max}\, a_i \mathbb{E}\left[U(a_i, \omega) \mid \omega_i, \mathcal{I}_i\right]$$

which can be summarized as

$$a_i = 1 \text{ if and only if } \mathbb{E}[\textstyle\sum_{j \leq n} \omega_j \mid \omega_i, \mathcal{I}_i] \geq 0.$$

where $\mathcal{I}_i = \mathcal{I}\left(\{a_j : j < i\}\right)$. Since ω_i and $\mathcal{I}_i$ do not provide any information about the content of successors' signals, we obtain

$$a_i = 1 \text{ if and only if } \mathbb{E}[\textstyle\sum_{j \leq i} \omega_j \mid \omega_i, \mathcal{I}_i] \geq 0,$$

and thus,

$$a_i = 1 \text{ if and only if } \omega_i \geq -\mathbb{E}[\textstyle\sum_{j<i} \omega_j \mid \mathcal{I}_i].$$

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*

$$a_i = \begin{cases} 1 & if \quad \omega_i \geq \hat{\omega}_i \\ 0 & if \quad \omega_i < \hat{\omega}_i \end{cases} \tag{1}$$

where

$$\hat{\omega}_i = -\mathbb{E}[\textstyle\sum_{j<i} \omega_j \mid \mathcal{I}_i] \tag{2}$$

is the optimal history-contingent cutoff.

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

The optimal history-contingent cutoff $\hat{\omega}_i$ is sufficient to characterize agent i's behavior, and thus the cutoff process $\{\hat{\omega}_i\}_{i=1}^n$ characterizes the social behavior. Note 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 $\hat{\omega}_i$ after every possible history of actions $(a_j)_{j<i}$ and $\hat{\omega}_i$ is derived in (2) via Bayes' rule. In Appendix B, we describe agents' behavior formally and discuss the essential elements of the weak perfect Bayesian equilibrium.

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 $\exists i$ such that $\hat{\omega}_j \in C_1$ (C_0) $\forall j \geq i$.*

Analogously, a limit-cascade on action $x = 1$ $(x = 0)$ occurs when the process of cutoffs $\{\hat{\omega}_i\}$ converges almost surely to a random variable $\hat{\omega}_\infty = \lim_{n\to\infty} \hat{\omega}_n$, with $supp(\hat{\omega}_\infty) \subseteq C_1$ $(supp(\hat{\omega}_\infty) \subseteq C_0)$.

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

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

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

Definition 3 (Herd behavior) *Herd behavior occurs when $\exists i$ such that*

$$\lim_{n\to\infty} l_i^n/n = 1.$$

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 $\{\hat{\omega}_i\}$. From the definition of equilibrium (see Appendix), we know that $\mathcal{I}_i \subseteq \mathcal{I}_{i+1} \subseteq \mathcal{I}$, 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 $\{\hat{\omega}_i, \mathcal{I}_i\}$ be an equilibrium. Then $\{\hat{\omega}_i\}$ is a martingale with respect to $\{\mathcal{I}_i\}$ and there exists a random variable $\hat{\omega}_\infty$ such that $\hat{\omega}_n$ converges to $\hat{\omega}_\infty$ almost surely.*

In words, $\{\hat{\omega}_i\}$ 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 $\hat{\omega}_\infty = \lim_{n\to\infty} \hat{\omega}_n$. 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 must[8]. The rate of convergence of the cutoff process and its limiting value depend on what the exact realization of private signals is[9]. 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 $\hat{\omega}_i$ when $a_{i-1} = 1$

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

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

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

[9] We 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.

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

$$F(\omega) \leq \frac{(\omega+1)}{(\omega+2)}$$

for any $\omega \in [-1, 0]$*, 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, like in BHW, 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,

Theorem 1 (No overconfidence) *With only rational agents, (i) an informational cascade need not arise, i.e.,* $\hat{\omega}_i \notin \{C_0, C_1\}$ *for all* i*, and (ii) a limit-cascade, i.e.,* $\hat{\omega}_\infty = \lim_{n\to\infty} \hat{\omega}_i$*, with* $\hat{\omega}_\infty \in \{C_0, C_1\}$*, 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. 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.

Stated in the language of the previous section, the following definition of overconfidence serves as the working definition. Recall that although whether an agent is overconfident is unobservable by others, the fraction of overconfident agents p and the degree of overconfidence $\tilde{\omega}$ are public knowledge.

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*

$$a_i = \begin{cases} 1 & if \quad \omega_i \geq \tilde{\omega}_i \\ 0 & if \quad \omega_i < \tilde{\omega}_i \end{cases} \tag{3}$$

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

As Figure 1 illustrates, if a herd of action $a = 1$ precedes a rational agent i then according to (1) and (2) her cutoff $\hat{\omega}_i$ is close to -1. Hence, in the subset $[\hat{\omega}_i, 0)$, 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 $\tilde{\omega}_i \geq \hat{\omega}_i$, and thus she has a smaller imitation set $[\tilde{\omega}_i, 0) \subseteq [\hat{\omega}_i, 0)$.

[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 $\mathcal{I}_i'$ is generated by p, $\tilde{\omega}(\hat{\omega})$ 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

$$a_i = \begin{cases} 1 & if \quad \omega_i \geq \hat{\omega}_i \\ 0 & if \quad \omega_i < \hat{\omega}_i \end{cases}$$

where the history-contingent cutoff rule is now described by

$$\hat{\omega}_i = -\mathbb{E}\left[\sum_{j<i} \omega_j \mid \mathcal{I}_i'\right] \tag{4}$$

As in the model with no overconfidence, the cutoff $\hat{\omega}_i$ 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 $\tilde{\omega}_i$ is given by the mapping $\hat{\omega} \mapsto \tilde{\omega}(\hat{\omega})$.

Even in the presence of overconfident agents, the process of rational cutoffs $\{\hat{\omega}_i\}$ 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* $\{\hat{\omega}_i, \mathcal{I}_i'\}$ *be an equilibrium. Then* $\{\hat{\omega}_i\}$ *is a martingale with respect to* $\{\mathcal{I}_i'\}$ *and there exists a random variable* $\hat{\omega}_\infty$ *such that* $\hat{\omega}_n$ *converges to* $\hat{\omega}_\infty$ *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 $(\hat{\omega}_i)$ 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 some a-cascade set. For an informational cascade, say on $a = 1$, to start, a discrete jump of $(\hat{\omega}_i)$ into C_1 must occur in a finite time. With overconfidence, the law of motion for $\hat{\omega}_i$ when $a_{i-1} = 1$ is given by

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

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

$$\hat{\omega}_{i-1} - \mathbb{E}^+(\hat{\omega}_{i-1}) - p\Delta^+(\hat{\omega}_{i-1}) \leq -1$$

where $\Delta^+(\hat{\omega}_{i-1}) \equiv \mathbb{E}^+(\tilde{\omega}_{i-1}) - \mathbb{E}^+(\hat{\omega}_{i-1})$ is the rational learning attributable to overconfidence, which is positive and increasing in the degree of overconfidence.

Hence, to have a cascade the trick is to choose a proportion of overconfident agents p and level of overconfidence such that a discrete jump into C_1 takes place. Once the cutoff process enters the cascade set C_1 action $a = 1$ is always taken thereafter. Note that even with a small fraction of overconfident agents p information is lumpy, so during a herd actions can be informative enough to toss all successors into an informational cascade. This is because someone who overweighs private information, reveals more information about the private signal, and, as a result, may stimulate an informational cascade.

The difference between informational cascades and herd behavior obviously matters very little if during a herd public information is accumulated so slowly that it has almost no economic value *ex ante*. However, note that the informational premium of being outside of a cascade is that a deviation that reveals lots of information can never be ruled out. Thus, there could be a significant informational loss to being in a cascade since substantial valuable information is completely suppressed by agents' decisions.

It is easy to see that the value of information lost is unbounded since

$$\mathbb{E}[(\textstyle\sum_{i\leq n}\omega_i)^2 \Big| \textstyle\sum_{i\leq n}\omega_i \geq 0] = n^2\mathbb{E}[\omega_i]^2$$

and $n^2\mathbb{E}[\omega_i]^2 \to \infty$ as $n \to \infty$, but whether rational agents are better off or worse off depend on the informational trade off between the additional information revealed by overconfident actions and the decrease of information revealed by rational agents' actions.

To illustrate the short-run dynamics of the model and to understand the dissimilarities between the case of no overconfidence and overconfidence, consider a finite herd followed by a deviator. With only rational agents, the deviator reveals clear cut information regarding her private signal that dominates the accumulated public information. Thus, her successor will be slightly in favor of joining the deviation. This is referred to by Smith and Sørensen (2000) as the overturning principle.

After adding overconfident agents, in contrast, successors to a deviator will believe that it is likely to be an overconfident act and thus will be less in favor of joining the deviation. The reason is that overconfident agents have a larger no herding subset of private signals and thus, put side by side with a rational agent, an overconfident agent reveals more of her private information by joining a herd but less by avoiding it.

To illustrate, assume that a long finite herd of $a = 1$ precedes some rational agent $i-1$. Then, her cutoff is close to -1. If she receives an extreme contrary signal, $-1 \leq \omega_{i-1} < \hat{\omega}_{i-1}$, she deviates by choosing $a =$

0. With only rational agents, since necessarily $\hat{\omega}_{i-1} > \mathbb{E}^{-}(\hat{\omega}_{i-1})$, having observed the deviation a rational agent i overturns the behavior by setting $\hat{\omega}_i > 0$ (but yet be close to zero). On the other hand, with overconfident agents $\mathbb{E}^{-}(\tilde{\omega}_{i-1}) \geq \mathbb{E}^{-}(\hat{\omega}_{i-1})$ since $\tilde{\omega}_{i-1} > \hat{\omega}_{i-1}$, and thus for some level of overconfidence the behavior will fail to overturn, $\hat{\omega}_i < 0$, since the deviation has less impact on public information.

Summarizing,

Theorem 2 (Overconfidence) *With overconfident agents, (i) an informational cascade arises, i.e., for any fraction of overconfident agents p and mapping of overconfidence $\hat{\omega} \mapsto \tilde{\omega}(\hat{\omega})$ there exists some finite i such that $\hat{\omega}_i \in C_0$ or $\hat{\omega}_i \in C_1$ for all $j \geq i$, and (ii) herd violations fall short of offsetting the accumulated information favoring the contrary action.*

5 Closed Forms

To illustrate the dynamics of the model, we need to further specialize the model by assuming that for each agent i, the signal ω_i is uniformly distributed, so that the history-contingent cutoff rules can be obtained in closed forms. It is particularly useful as it allows carrying out some comparative statics exercises to explore further the qualitative features of the informational trade off caused by overconfidence.

With only rational agents the dynamics of the cutoff rule $\hat{\omega}_i$ is described in a closed form recursively as follows

$$\hat{\omega}_i = \begin{cases} \frac{-1+\hat{\omega}_{i-1}}{2} & \text{if} \quad a_{i-1} = 1 \\ \frac{1+\hat{\omega}_{i-1}}{2} & \text{if} \quad a_{i-1} = 0 \end{cases}$$

where $\hat{\omega}_1 = 0$.

The impossibility of an informational cascade follows immediately since for every i, $-1 < \hat{\omega}_i < 1$. It implies that the agent i's decision always depends in a non-trivial way on her private information in the sense that for some signals she will choose $a_i = 1$ and for other signals she will choose $a_i = 0$. However, notwithstanding the impossibility of informational cascade, the model predicts that herd behavior must arise.

To see that, note that if we assume that the first two agents choose action 0, the third agent's cutoff is $\hat{\omega}_3 = 3/4$; if the first three agents choose 0, the fourth agent's cutoff is $\hat{\omega}_4 = 7/8$; and if the first k agents choose 0, the $(k+1)$ agent's cutoff is $\hat{\omega}_{k+1} = 1 - 2^{-(k-1)}$. Hence, any successive agent who also chooses action 0 reveals less of her private information and makes it more difficult for her predecessor not to choose action 0, but at each date the cutoff lies strictly below C_0.

More precisely, let p_i denote the probability that $a_j = 0$ for all agents $j \geq i$ conditional on the history $a_l = 0$ for all agents $l < i$, and note that $p_i = p_{i+1}\hat{\omega}_i$ since $\hat{\omega}_i = \Pr(a_{i+1} = 0 \,|a_j = 0, \forall j \leq i)$. Then, $p_i > 0$ implies that $p_{i+1} > p_i$ and in fact $p_n \to 1$ as $n \to \infty$. Beginning with the first agent, $\log p_1 = \sum_{i=1}^{n} \log(1 - 2^{-i})$, and as $n \to \infty$, $\log p_1 \approx -\sum_{i=1}^{\infty} \log 2^{-i} > -\infty$ so that $p_1 > 0$ and $p_n \to 1$ as $\hat{\omega}_n \to 1$.

On the other hand, if the fourth agent chooses action 1 after the first three agents choose 0, her decision reveals that her signal lies in the interval $[7/8, 1]$ and the fifth agent's cutoff is $\hat{\omega}_5 = -1/16$. Hence, the longer a cluster of agents acts alike, the larger the asymmetry between the information revealed by imitation and deviation. Notice that a deviator induces her successor to be slightly in favor of joining the deviation.

After adding overconfident agents to the model, simple calculations show that the adjusted cutoff dynamics of rational agents follow the process

$$\hat{\omega}_i = p^{\frac{\hat{\omega}_{i-1} - \hat{\omega}_{i-1}}{2}} + \begin{cases} \frac{-1+\hat{\omega}_{i-1}}{2} & \text{if} \quad a_{i-1} = 1 \\ \frac{1+\hat{\omega}_{i-1}}{2} & \text{if} \quad a_{i-1} = 0 \end{cases}$$

where $\hat{\omega}_1 = 0$ and in the language of the previous section $(\hat{\omega}_{i-1} - \hat{\omega}_{i-1})/2$ is the rational learning of agent i attributable to overconfidence which is positive (negative) if $a_{i-1} = 1$ ($a_{i-1} = 0$).

To begin with, we illustrate, with an extreme example, how the behavior alters because of overconfidence. Suppose that the second agent is so overconfident that she ignores the information revealed from the first's decision and follows her own signal, $\tilde{\omega}_2 = 0$. If the third agent is rational and identifies the nature of the second's decision, whenever the first two decisions coincide, say $a_1 = a_2 = 1$, the third should also choose $a_3 = 1$ regardless of her private signal, $\hat{\omega}_3 = -1$. Hence, the surfeit of information revealed by the second's actions leads to an informational cascade.

Since the model resolves so well in closed form for the uniform signal distribution, it is useful to illustrate how cascades start with the following form of overconfidence.

Example $\hat{\omega} \mapsto \tilde{\omega}(\hat{\omega})$ such that $\tilde{\omega}(\hat{\omega}) = \hat{\omega}^k$ for any $\hat{\omega} \notin \{C_0, C_1\}$ for some $k > 1$ odd and $\tilde{\omega}(\hat{\omega}) = \hat{\omega}$ otherwise.

Note that k is the (constant) degree of overconfidence and that overconfidence is bounded in strength as overconfident actions diminish once the cutoff process jumps into a cascade. Figure 2 illustrates the process of rational cutoffs $\{\hat{\omega}_i\}$ for the first few agents when all choose action 1 with different degrees of overconfidence k (Figure 2a) and different proportions of overconfident agents p (Figure 2b). Note that the adjusted cutoff process easily escapes into C_1, meaning that an informational cascade occurs. In this state of affairs the actions of overconfident agents do not remain informative during a cascade, and thus it is an absorbing state.

[Figure 2 here]

As to the welfare properties of the equilibria, the likelihoods of correct decisions with and without overconfidence can not be found analytically since conditional on the payoff-relevant state, private signals are negatively correlated. However, simulations show certain directional effects, which, to the extent that we can cover finite group sizes, we conjecture that they are robust[10]. Figure 3 summarizes simulations that were carried out with a group size $n = 10$ with different degrees of overconfidence k (Figure 3a) and different proportions of overconfident agents p (Figure 3b).

[Figure 3 here]

Note that the *ex ante* probability that rational agents make a correct decision increases over time but decreases with both increasing proportion of overconfident agents and increasing degree of overconfidence. From a social point of view, not only that having overconfident agents does not make rational agents better off, it also decreases average welfare as overconfident agents are more likely to take a suboptimal action.

6 Conclusion

In this paper we test how robust the theory is to the well-known behavioral phenomenon of individual overconfidence. Overconfident agents overweigh their private information relative to the public information revealed by the decisions of others, and thus, when following a herd, they reveal more of the information available to them. However, we show that overconfidence trades the additional information revealed by overconfident decisions against more information that is being suppressed by perfectly rational decisions.

[10]Numerical simulation are carried out by MatLab. Experiments are repeated until the marginal change in the average of an for additional 10^7 experiments is less than 10^{-5}.

While with only rational agents informational cascades are impossible, they arise after adding overconfidence agents, for the surfeit of information revealed by overconfident actions can lead to an informational cascade. This is because someone who overweighs private information, reveals more information about the private signal, and, as a result, may stimulate an informational cascade. Thus, the presence of overconfident agents intensifies the free-rider problem of rational agents and cannot break the poor information flow intrinsic to uniform behavior or improve decisions accuracy and welfare.

There is a large literature on overconfidence. The most closely related paper is by BW who find that overconfidence can be useful, especially if the group is large enough to benefit from its positive information externality. The reason for the different results is that their model is based on the binary-signal-binary-action model of BHW. In such a binary signal structure, all herds are cascades since once two consecutive decisions coincide no signal can lead to a deviation. Thus, information is revealed from decisions made before a cascade starts, from the two decisions that start the cascade and from overconfident deviations from a cascade. Furthermore, rational agents do not suppress their own information more because of the additional dissemination of information available from overconfident decisions.

We expect, however, that the results of BW to hold also in a continues-signal model with atomic tails, i.e., when the private signal distribution is close enough to BHW's binary private signal, but we know of no sufficient conditions which guarantee that some overconfidence can be useful. Given that BW have reached a very different conclusion, it is natural to ask about the robustness of the results to different information structures. Whether this would lead to sharply different results is unclear, since all the decision rules would have to be changed to reflect the new environment. Obviously, different information structures may lead to different outcomes. This is an important subject for future research.

7 Omitted Proofs

Proof of Proposition 2 By definition (Billingsley 1986), $\{\hat{\omega}_i\}$ is a submartingale with respect to $\{\mathcal{I}_i\}$ if the following four conditions hold: (i) $\mathcal{I}_i \subseteq \mathcal{I}_{i+1}$, (ii) $\hat{\omega}_i$ is measurable $\mathcal{I}_i$, (iii) $\mathbb{E}[|\hat{\omega}_i|] < \infty$, and (iv) with probability 1, $E[\hat{\omega}_{i+1}|\mathcal{I}_i] \geq \hat{\omega}_i$. $\{\hat{\omega}_i\}$ is a supermartingale if $\{-\hat{\omega}_i\}$ is a submartingale, and $\{\hat{\omega}_i\}$ is a martingale if it is both a sub- and supermartingale.

The first condition follows directly from the definition of weak perfect

Bayesian equilibrium. The second conditions follows directly from the definition the cutoff strategy. The third holds because $U(1,\cdot)$ is bounded. To establish the fourth condition, note that using symmetry,

$$\begin{aligned}\mathbb{E}[\omega_{i-1} \quad | \quad \hat{\omega}_{i-1}] &= F(\hat{\omega}_{i-1})\mathbb{E}^{-}(\hat{\omega}_{i-1}) + [1 - F(\hat{\omega}_{i-1})]\mathbb{E}^{+}(\hat{\omega}_{i-1}) \\ &= \int_{-1}^{\hat{\omega}_{i-1}} \omega dF + \int_{\hat{\omega}_{i-1}}^{1} \omega dF \\ &= \int_{-1}^{1} \omega dF \\ &= 0\end{aligned}$$

where $\mathbb{E}^{+}(\xi) \equiv \mathbb{E}\left[\omega | \omega \geq \xi\right]$ and $\mathbb{E}^{-}(\xi) \equiv \mathbb{E}\left[\omega | \omega < \xi\right]$, and thus $\mathbb{E}[\hat{\omega}_i \mid \hat{\omega}_{i-1}] = \hat{\omega}_{i-1}$.

From the martingale convergence theorem, there exists a random variable $\hat{\omega}_\infty$ such that $\hat{\omega}_n \to \hat{\omega}_\infty$ almost surely and $\mathbb{E}[\hat{\omega}_\infty] = \mathbb{E}[\hat{\omega}_i]$ for all i.

Proof of Proposition 3 Using the symmetry of F, when $a_{i-1} = 1$ the law of motion for $\hat{\omega}_i$ is given by

$$\begin{aligned}\hat{\omega}_i &= \hat{\omega}_{i-1} - \mathbb{E}^{+}(\hat{\omega}_{i-1}) \\ &= \hat{\omega}_{i-1} - \tfrac{1}{1-F(\hat{\omega}_{i-1})} \int_{\hat{\omega}_{i-1}}^{1} \omega dF \\ &= \hat{\omega}_{i-1} + \tfrac{1}{1-F(\hat{\omega}_{i-1})} \int_{-1}^{\hat{\omega}_{i-1}} \omega dF \\ &\leq \hat{\omega}_{i-1} - \tfrac{F(\hat{\omega}_{i-1})}{1-F(\hat{\omega}_{i-1})}\end{aligned}$$

and with strict inequality as long as $\hat{\omega}_{i-1} > -1$. The condition that $\hat{\omega}_i > -1$ for all i follows by direct calculations.

8 Appendix

A: Information and payoff structures We analyze a model first proposed by Gale (1996), with a particularly tractable information and payoff structures. Next, we place the model in the domain of the social learning literature.

A general model of social learning comprises a finite set of agents indexed by $i = 1, ..., n$, a finite set of actions $\mathcal{A} \subset \mathbf{R}$, a set of states of nature Ω, and a common payoff function $U(a, \cdot)$.

Uncertainty is represented by a probability measure space $(\Omega, \mathcal{I}, \mathbf{P})$, where Ω is a compact metric space, $\mathcal{I}$ is a σ-field, and $\mathbf{P}$ a probability measure. Each agent i receives an informative private signal $\sigma_i(\omega)$, a function of the state of nature ω, and uses this private information to identify an irreversible payoff-maximizing action.

We assume that $\Omega = \Omega_1 \times \cdots \times \Omega_n$, where Ω_i is an interval $[-\alpha, \alpha]$ and the generic element is $\omega = (\omega_1, ..., \omega_n)$. For each i, the signals are assumed to satisfy $\sigma_i(\omega) = \omega_i, \forall \omega \in \Omega$ where the random variables $\omega_1, ..., \omega_n$ are independently and continuously distributed, that is $\mathbf{P} = \mathbf{P}_1 \times \cdots \times \mathbf{P}_n$. Specifically, $\mathbf{P}_i$ is given by a *c.d.f.* F over the compact support with convex hull $[-\alpha, \alpha]$, such that $\mathbb{E}[\omega] = 0$, F has no atoms and satisfies symmetry, i.e., $F(\omega) = 1 - F(-\omega)\ \forall \omega \in [-\alpha, \alpha]$.

Note that we start with n finite, yet we are interested in approximating the behavior when the group size n is arbitrarily large. In other words, we analyze the limit behavior of a sequence of group sizes indexed by n. Although the information of agent i ω_i about the state of the world $\sum_{i=1}^{n} \omega_i$ is not constant across different sized groups n, the underlying decision problem, the optimal decision rule, and hence our results are independent of n as in the standard social learning models.

B: Weak perfect Bayesian equilibrium We analyzed the weak perfect Bayesian equilibrium of the game. Agent i's choice of action is described by a random variable $X_i(\omega)$ and her information is described by σ_i and a σ-field $\mathcal{I}_i = \mathcal{I}(\{X_j : j < i\})$. Since $\mathcal{I}_i$ represents the agent's publicly available information, it must be the σ-field generated by the random variables $\{X_j : j < i\}$, and since the agent's choice can only depend on the information available to her, X_i must be measurable with respect to $(\sigma_i, \mathcal{I}_i)$. Finally, since X_i is optimal, there cannot be any other $\mathcal{I}_i$-measurable choice function that yields a higher expected utility.

These are the essential elements of the weak perfect Bayesian equilibrium stated below.

Definition 5 *A weak perfect Bayesian equilibrium consists of a sequence of random variables $\{X_i\}$, $\{\sigma_i\}$ and σ-fields $\{\mathcal{I}_i\}$ such that for each $i = 1, ..., n$ (i) $X_i : \Omega \to \mathcal{A}$ is $(\sigma_i, \mathcal{I}_i)$-measurable, (ii) $\mathcal{I}_i = \mathcal{I}(\{X_j : j < i\})$, and (iii) for any $(\sigma_i, \mathcal{I}_i)$-measurable function $x : \Omega \to \mathcal{A}$,*

$$\mathbb{E}[U(x(\omega), \omega)] \leq \mathbb{E}[U(X_i(\omega), \omega)].$$

Note that the definition of equilibrium does not require optimality off the path of play. Since there are no strategic interactions, there is no incentive to make an out-of-equilibrium move in order to signal to successors.

References

[1] Anderson, L. and C. Holt (1997) "Information Cascades in the Laboratory." *American Economic Review,* 87, pp. 847-62.

[2] Benabou, R. and J. Tirole (2002) "Self-Confidence and Personal Motivation." *Quarterly Journal of Economics,* 117, pp. 871-915.

[3] Benabou, R. and J. Tirole (2003) "Intrinsic and Extrinsic Motivation." *Review of Economic Studies*, forthcoming.

[4] Banerjee, A. (1992) "A Simple Model of Herd Behavior." *Quarterly Journal of Economics,* 107, pp. 797-817.

[5] Bernardo, A. and I. Welch (2001) "On the Evolution of Overconfidence and Entrepreneurs." *Journal of Economics & Management Strategy*, 10, pp. 301-330.

[6] Bikhchandani, S., D. Hirshleifer and I. Welch (1992) "A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascade," *Journal of Political Economy* 100, pp. 992-1026.

[7] Bikhchandani, S., D. Hirshleifer and I. Welch (1998) "Learning from the Behavior of Others: Conformity, Fads, and Informational Cascades," *Journal of Economic Perspective*, 12, pp. 151-170.

[8] Billingsley, P. (1986). *Probability and Measure, 2nd edition.* John Wiley and Sons, New York.

[9] Camerer, C. and D. Lovallo (1999) "Overconfidence and Excess Entry: an Experimental Approach." *American Economic Review* 89**,** pp. 306-318.

[10] Çelen, B. and S. Kariv (2004a) "Observational Learning Under Imperfect Information," *Games and Economic Behavior,* 47(1)*,* pp. 72-86.

[11] Çelen, B. and S. Kariv (2004b) "Distinguishing Informational Cascades from Herd Behavior in the Laboratory," *American Economic Review,* 94, pp. 484-497.

[12] Çelen, B. and S. Kariv (2005) "An Experimental Test of Observational Learning under Imperfect Information," *Economic Theory*, 26(3), pp. 677-699.

[13] Chamley, C. and D. Gale (1994) "Information Revelation and Strategic Delay in a Model of Investment." *Econometrica*, 62(5), pp. 1065-85.

[14] Daniel, K., D. Hirshleifer and A. Subrahmanyam (1998) "A Theory of Overconfidence, Self-Attribution, and Security Market Under- and Over-reaction," *The Journal of Finance*, 53, pp. 1839-1886.

[15] DeBondt, W. and R. Thaler (1995) "Financial Decision-Making in Markets and Firms: A Behavioral Perspective," in R., Jarrow, V. Maksimovic and W. Ziemba, eds., *Finance*, Elsevier North Holland, Amsterdam, Handbook in Operation in Research and Management Science, Vol. 9 Chap 13, pp. 385-410.

[16] Ellison, G. and D. Fudenberg (1993) "Rules of Thumb for Social Learning." *Quarterly Journal of Economics*, 101, pp. 612-643.

[17] Ellison, G. and D. Fudenberg (1995) "Word of Mouth Communication and Social Learning." *Journal of Political Economy*, 110, pp. 93-126.

[18] Gale, D. (1996) "What Have We Learned from Social Learning?" *European Economic Review*, 40, pp. 617-628.

[19] Goeree, J., T. Palfrey, b. Rogers, and D. McKelvey (2004) "Self-correcting Information Cascades." Caltech Social Science Working Paper No. 1197.

[20] Gul, F. and R. Lundholm (1995) "Endogenous Timing and the Clustering of Agents' Decisions," *Journal of Political Economy,* 103(5), pp. 1039-66.

[21] Lee, I. H. (1993) "On the Convergence of Informational Cascades," *Journal of Economic Theory,* 61(2), pp. 396-411.

[22] Moscarini, G., M. Ottaviani and L. Smith (1998) "Social Learning in a Changing World," *Economic Theory,* 11, pp. 657-65.

[23] Nöth, M. and M. Weber (2003) "Information Aggregation with Random Ordering: Cascades and Overconfidence," *Economic Journal*, 113, 166-89.

[24] Smith, L. and P. Sørensen (2000) "Pathological Outcomes of Observational Learning," *Econometrica*, 68, pp. 371-398.

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

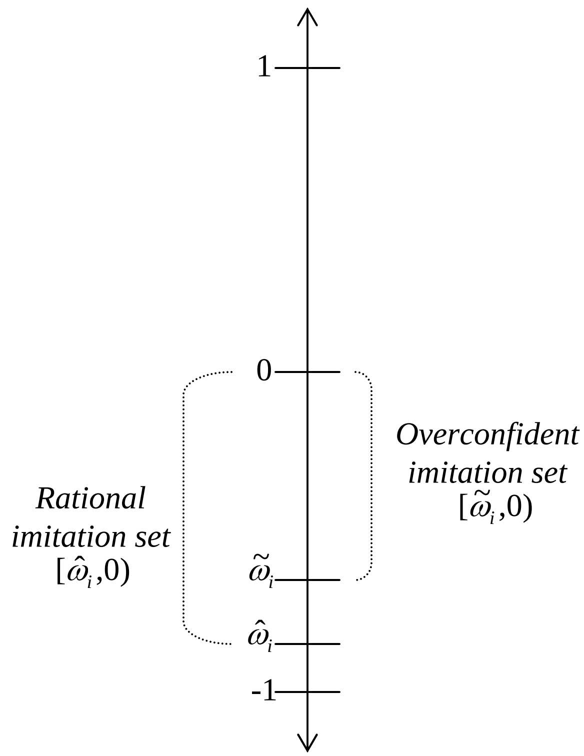

Figure 2a. The sequences of cutoffs when all agents choose action 1 with different degrees of overconfidence

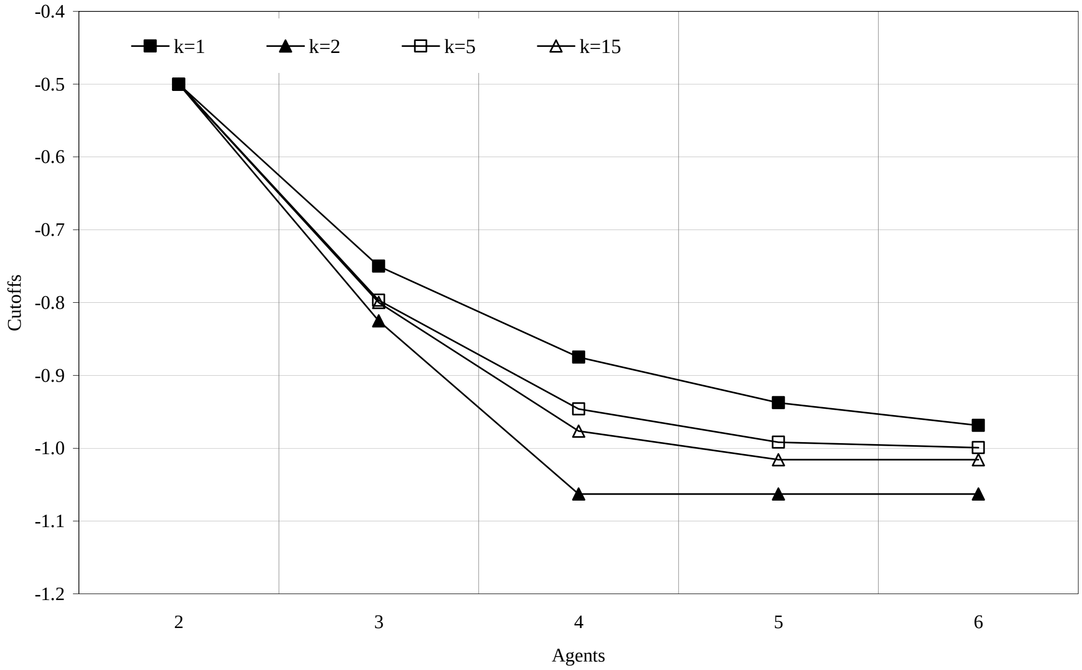

Figure 2b. The sequences of cutoffs when all agents choose action 1 with different proportions of overconfident agents

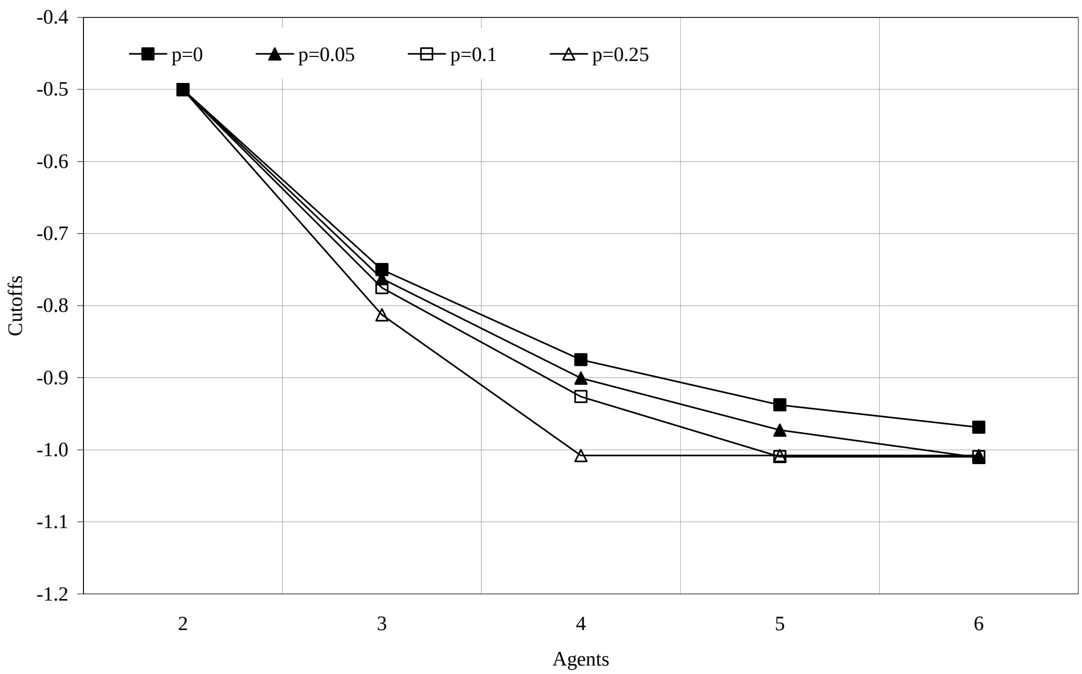

Figure 3a. The probability that a rational agent makes a correct decision with different degrees of overconfidence

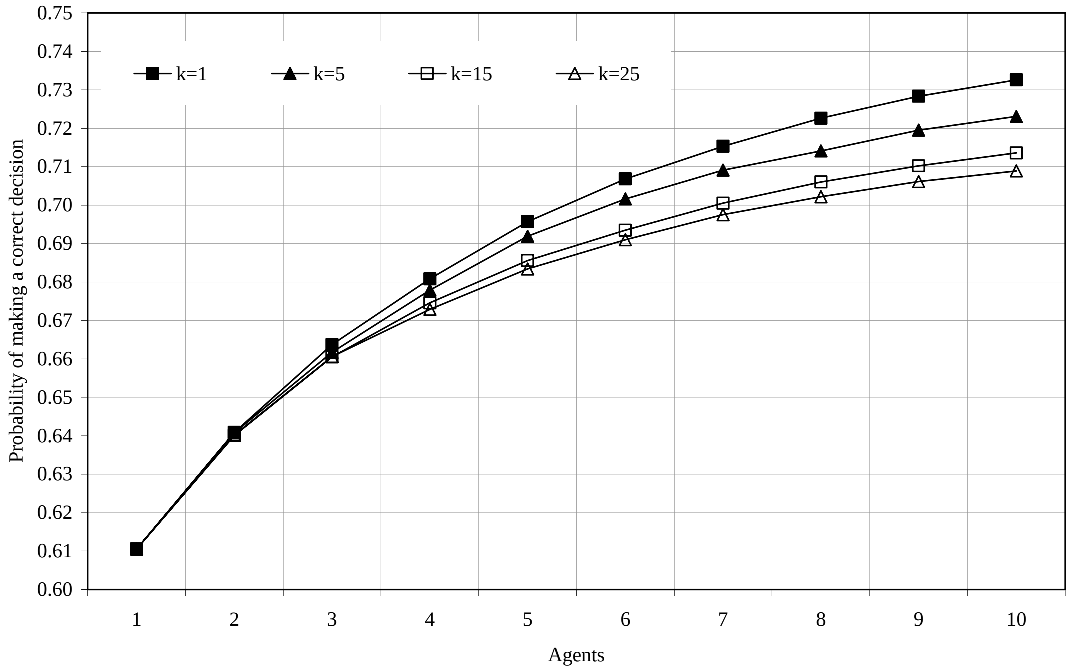

Figure 3b. The probability that a rational agent makes a correct decision with different proportions of overconfident agents