JENA ECONOMIC RESEARCH PAPERS

2008 – 094

Fragility of Information Cascades: An Experimental Study using Elicited Beliefs

by

Frédéric Koessler
Anthony Ziegelmeyer
Juergen Bracht
Eyal Winter

www.jenecon.de

ISSN 1864-7057

The JENA ECONOMIC RESEARCH PAPERS is a joint publication of the Friedrich Schiller University and the Max Planck Institute of Economics, Jena, Germany. For editorial correspondence please contact markus.pasche@uni-jena.de.

Impressum:

Friedrich Schiller University Jena
Carl-Zeiss-Str. 3
D-07743 Jena
www.uni-jena.de

Max Planck Institute of Economics
Kahlaische Str. 10
D-07745 Jena
www.econ.mpg.de

Jena Economic Research Papers 2008 - 094

Fragility of Information Cascades: An Experimental Study using Elicited Beliefs *

Frédéric KOESSLER† Anthony ZIEGELMEYER‡ Juergen BRACHT§ Eyal WINTER¶

November 12, 2008

Abstract

This paper examines the occurrence and fragility of information cascades in laboratory experiments. One group of low informed subjects make predictions in sequence. In a matched pairs design, another set of high informed subjects observe the decisions of the first group and make predictions. According to the theory of information cascades (Bikhchandani, Hirshleifer, and Welch, 1992), if initial decisions coincide, an information cascade should occur: it is rational for subsequent players with low quality information to follow the observed pattern regardless of their private information. However, an information cascade should be fragile: it is always rational for subsequent players with high quality information to follow their private information. In line with existing experiments on information cascades, we find some evidence that low informed subjects follow the herd when it is rational, and this herding behavior occurs more frequently if there is a pronounced imbalance. The main finding of this paper is that *information cascades are not fragile*. We find strong evidence that highly informed subjects follow the herd regardless of their private information. In accordance with those observations we show, by explicitly eliciting subjects' beliefs about the state, that beliefs are not constant in the number of previous decisions that coincide, whether or not an information cascade already occurred. Subjects' behavior can be understood with a statistical model that allows for the possibility of errors in earlier decisions.

KEYWORDS: Information cascades, elicited beliefs, experimental economics

JEL CLASSIFICATION: C72, C92, D82.

*This paper is a revised version of Chapter V of Ziegelmeyer's dissertation (Ziegelmeyer, 2001) which resulted from a collaboration with Frédéric Koessler. It circulated earlier under the title "Behaviors and Beliefs in Information Cascades". The authors have benefited greatly from comments and suggestions from Mohammed Abdellaoui, Nick Feltovick, Dorothea Kübler, Bradley Ruffle, Georg Weizsäcker, participants at F.U.R. X in Torino, 5ièmes Journées d'Economie Expérimentale in Strasbourg, and Economics Dept. Seminar at Ben Gurion University. We are indebted for the financial support provided by the Department of Economics at the University of Strasbourg.

†Paris School of Economics and CNRS. *E-mail:* `koessler@pse.ens.fr`

‡Max Planck Institute for Research into Economic Systems, Strategic Interaction Group, Jena, Germany. *E-mail:* `ziegelmeyer@mpiew-jena.mpg.de`

§University of Aberdeen Business School.

¶Center for Rationality and Interactive Decision Theory, Hebrew University of Jerusalem.

1 Introduction

In recent years a great deal of attention has been focused on situations in which the existence of informational externalities leads to a loss of social welfare. In these situations, the attempt by agents to take advantage of the information of others can lead to the failure to exploit their own information in a socially optimal way. This likely consequence of agents relying on whatever information they have obtained via observation of others' action is what has been called "rational herding" or "information cascades".[1] An information cascade can be defined as a choice sequence in which some agents act as if they ignore their private information and follow the choices made earlier in the sequence by other agents. Rational herding or information cascades have been proposed as explanations of a variety of phenomena, such as fads, fashions, booms and crashes.[2]

As emphasized by Bikhchandani et al. (1992), *the convergence of behavior in rational herding is idiosyncratic and fragile*. More precisely, once an information cascade has started the further choices are uninformative, i.e., the informativeness of this informational cascade does not rise with the number of similar choices. Hence, the social cost of information cascades is that the benefit of diverse information sources is lost but incorrect decisions can be rapidly reversed, e.g., by the arrival of a little extra information. These two consequences of rational herd behavior are the two sides of the same coin, and they constitute the departure point of our experimental study.

We extend existing experimental designs on information cascades in two important respects.[3] First, we rely on a matched pairs design where one group of low informed subjects make predictions in sequence and another set of high informed subjects observe the decisions of the first group and make predictions too. Though the Bayesian Nash equilibrium is characterized by a high occurrence of information cascades in the low informed subjects sequence, rational herd behavior predicts that highly informed subjects should go against the established pattern of identical choices when endowed with a contradicting signal. Second, in one of the two experiments reported here we directly elicit subjects' beliefs using a "quadratic proper scoring rule" which provides subjects with an incentive to report their beliefs truthfully. The quadratic proper scoring rule is an incentive compatible mechanism to elicit beliefs provided that respondents are risk neutral and do not distort probabilities.[4]

[1] The sequential rational herding literature dates back to Banerjee (1992) and Bikhchandani et al. (1992) who first developed simple models of informational cascades.

[2] For a survey, see Bikhchandani, Hirshleifer, and Welch (1998) or Chamley (2003).

[3] Anderson and Holt (1997) produced the first results showing that rational herding occurs in the laboratory. These experimental results have been replicated by, among others, Huck and Oechssler (2000), Hung and Plott (2001), Nöth and Weber (2003), Celen and Kariv (2004) and Kübler and Weizsäcker (2004, 2005).

[4] Sonnemans and Offerman (2001) experimentally investigate whether risk attitudes and probability weighting pose empirical problems for the elicitation of beliefs with a quadratic proper scoring rule. They find that in practice these factors do not affect subjects reported beliefs in an undesired way, i.e., subjects' stated beliefs are not biased by their risk attitudes or probability weighting when they are rewarded with the help of a quadratic proper scoring rule. Others have elicited beliefs in experimental studies. See, e.g., Offerman, Sonnemans, and Schram (1996), Huck and Weizsäcker (2002) and Nyarko and Schotter (2002) for direct evidence on subjects' expectations about others' behavior. More closely related to our study, McKelvey and Page (1990) report on an experiment of a common knowledge inference process where a lottery version of the quadratic proper scoring rule provides incentives for individuals to reveal their current posterior probabilities of an event.

Our results show that, on average, simple informed subjects rationally follow the herd in about 70% of the cases; these results are in accordance with previous experimental studies. However, though information cascades occur in the laboratory, a closer look at subjects' behavior reveals that they treat their private signals as irrelevant and follow the trend of their predecessors *only* when a strong majority of identical choices is observed. In this respect, whereas only about 2/3 of the choices are mimetic after two identical observed choices, almost all subjects follow the herd after more than four identical choices.[5] Our most important and original result is that high informed subjects' behavior entirely contradicts information cascade theory: slightly less than 35% of observed cascades are broken by more informed subjects. Hence, our experiment provides strong evidence that, while information cascades occur, they are not fragile.

The departure from the theory is supported by our analysis of the large data set on elicited beliefs. We observe that subjects' elicited beliefs are strongly consistent with their own behavior, and are not constant in the number of previous decisions that coincide, whether or not an information cascade theoretically occurred several periods before. As in Kübler and Weizsäcker (2004) we try to explain those deviations from the equilibrium predictions with a statistical error-rate model by introducing trembles in players' choices, with different error rates on different levels of reasoning about others' behavior.

The remainder of the paper is structured as follows. Section 2 describes the game used in the experiment, as well as the equilibrium prediction. Section 3 discusses the experimental design. In Section 4 we give an exhaustive analysis of our data and we conclude in Section 5.

2 A Model with Low and High Informed Agents

In this section we present a simple model of exogenous cascades based on the specific parametric model of Bikhchandani et al. (1992). First we consider a low informed setup in which the unique equilibrium outcome is characterized by a high occurrence of information cascades. Next we introduce more informed agents to underline the fragility of rational mimetism.

2.1 Low Informed Setup

Every agent in a sequence of nine agents has to choose one of two possible actions, A or B, where A stands for "predicting urn A" and B for "predicting urn B". Agents have a common prior belief on a payoff relevant state space $\{\alpha, \beta\}$, where α stands for "urn A has been selected", β stands for "urn B has been selected", and $\Pr(\alpha) = 1 - \Pr(\beta) = 0.55$.[6] Urn A contains two a balls and one b ball whereas urn B contains two b balls and one a ball. Guessing the right urn, i.e., choosing action A in state α or choosing action B in state β, yields 10 whereas guessing the wrong urn yields -5.

[5]This result was also obtained in a more recent experiment by Kübler and Weizsäcker (2005).

[6]Contrary to Bikhchandani et al.'s specific model and Anderson and Holt (1997) experiment, our initial events are not equally likely. This difference has the advantage that tie-breaking rules are useless, so the equilibrium outcome is unique.

Before choosing an action, each agent i both observes a single draw from the selected urn, which constitutes his private signal $t_i \in \{a, b\}$, and the public history of action decisions of all preceding individuals $1, 2, \ldots, i-1$. Conditionally to the realization of a state of Nature, the agents' signals are i.i.d. and the conditional probabilities are given by $\Pr(t_i = a \mid \alpha) = \Pr(t_i = b \mid \beta) = 2/3$ and $\Pr(t_i = a \mid \beta) = \Pr(t_i = b \mid \alpha) = 1/3$.

For $i \geq 2$, let $\{A, B\}^{i-1}$ be the space of all possible period i histories of actions chosen by the $i-1$ predecessors of agent i. Denote by h^{i-1} an element of $\{A, B\}^{i-1}$, i.e., h^{i-1} is a sequence of actions up to agent $i-1$. Let $\mu_i : \{a, b\} \times \{A, B\}^{i-1} \to [0, 1]$ be player i's belief (conditional probability given past observed actions and his private signal) that the state of Nature is α. Agent i's belief is given by

$$\mu_i(t_i, h^{i-1}) = \frac{\Pr(t_i \mid h^{i-1}, \alpha)\ \Pr(\alpha \mid h^{i-1})}{\Pr(t_i \mid h^{i-1})},$$

where probabilities are computed with respect to players' strategies and the prior. Given a history h^{i-1}, a signal t_i and a belief $\mu_i(t_i, h^{i-1})$, player i's expected utility is given by $15\mu_i(t_i, h^{i-1}) - 5$ (respectively $15(1-\mu_i(t_i, h^{i-1}))-5$) if he chooses action A (respectively action B). Hence, predicting urn A is relevant for player i if he believes α with probability greater than $1/2$.

In equilibrium players rationally update their beliefs by observing their signal and previously taken actions, and they act rationally given these beliefs. Agent 1's optimal action is to predict in accordance with his signal. If agent 2 observes an A prediction then he predicts A too even if his private signal is b. As the same argument applies for all the rest of the sequence, it is here that an information cascade results. On the contrary, if agent 2 observes a B prediction then he predicts in accordance with his signal. If player 3 observes two B predictions then he follows his predecessors' choices even if he is endowed with an a signal. This implies that the rest of the sequence joins the herd. Once a cascade has started the further choices are uninformative. In other words, after an A prediction not canceled out by previous choices, whatever their positions in the sequence the beliefs of two followers are identical when endowed with the same private signal. Similarly, after two B predictions not canceled out by previous choices, whatever their positions in the sequence, the beliefs of two followers are identical when endowed with the same private signal.

The only history which does not lead to an information cascade is $BABABA\ldots$ Table 1 reports the probability of having no information cascade after any even number of agents lower than eight. There is a less than 5% probability that the fifth agent's decision in the sequence depends on his signal.

Number of agents	Probability of no cascade
2	0.222
4	0.049
6	0.011
8	0.002

Table 1: Probability of no cascade after any even number of agents lower than eight.

2.2 High Informed Setup

In this setup we discriminate between the agents by assuming that one and only one of them receives a more informative signal about the state. We will use the subscript j in order to refer to this agent and we denote his signal $t_j \in \{a_S, b_S\}$. Agent j's signal has a higher "accuracy": we assume that $\Pr(t_j = a_S \mid \alpha) = \Pr(t_j = b_S \mid \beta) = 4/5$ and $\Pr(t_j = a_S \mid \beta) = \Pr(t_j = b_S \mid \alpha) = 1/5$.

Whatever agent j's position in the sequence, he has to decide in accordance with his signal. Indeed, the high-precision agent's signal is twice as informative as any other agent's signal. So if agent j observes an A cascade and gets a b_S signal then his signal overweighs the a signal and the priors ($\mu_j(b_S, A) = 0.38 = \mu_i(b, \emptyset)$), and agent j has to decide in accordance with his signal. If agent j observes a B cascade and gets an a_S signal then his signal and the two b signals just cancel out. This leaves agent j with a belief which equals the prior one ($\mu_j(a_S, BB) = \Pr(\alpha)$). Hence the rational action is to choose A. Of course, if agent j either observes a $BABA\ldots$ sequence or if agent j's signal is in accordance with what he has observed then he has to follow his signal too. The interesting result is that whatever the type of information cascade that has started a more informed agent j, contrary to any less informed agent i, has to break the cascade.

3 Experimental Procedures

Two experiments were run on a computer network[7] at the BETA laboratory of experimental economics (LEES) in winter 2000 using 96 undergraduate students from the University of Strasbourg. No subjects had any training in game theory or economics of information. Each experiment was made up of three sessions which took between $1\frac{1}{2}$ and $2\frac{1}{4}$ hours. Sixteen subjects participated in each session (plus one subject that was used as a monitor).[8] Subjects were randomly assigned to a computer terminal, which was physically isolated from other terminals. Communication, other than through the decisions made, was not allowed. Subjects were instructed about the rules of the game and the use of the computer program through written instructions, which were read aloud by the monitor. A short questionnaire and one dry run followed.[9] Subjects, on average, earned approximately 126 French francs for their participation (including a show-up fee of 15 French francs), which was paid to them in cash at the end of the session.[10]

In the first experiment subjects played either the basic cascade game or the more informed cascade game described in Section 2, fifteen times in the same group. In each round, we implemented this setup in the following way. We built two "lines" of subjects, a "low line" and a "high line". The low line was constituted by nine subjects whereas the high line was only made of seven subjects, and subjects in the high line were only in positions from three to nine. At the beginning of the round,

[7]Based on an application developed by Boun My (2000) designed for Visual Basic.

[8]At least twenty two subjects were invited for each session to be able to select subjects and make sure all of the participating subjects had understood the game.

[9]The dry run was added in order to give some experience to the subjects about the computer program. Subjects did not take decisions in this dry run.

[10]\$1 was approximately 7.5 French francs at the time of the experiment.

a random choice was made between state α and state β, and the probability of choosing state α was 55%. Subjects were then chosen in a random order to observe a single draw from a selected urn. Balls tagged a or b were put in urns labeled A and B and drawn on the computer screen to represent subjects' signals. In the low line the signal's strength, which indicates the probability that this signal is correct, was equal to 2/3 whereas in the high line it was equal to 4/5. Thus, in the low line on each subject's computer screen appeared a ball drawn from an urn containing 3 balls, two "correct" balls and an "incorrect" one. In the high line on each subject's computer screen appeared a ball drawn from an urn containing 5 balls, four "correct" balls and an "incorrect" one. This information structure was common knowledge as being part of the instructions which were read aloud by the monitor. In particular, each subject knew on which type of private information quality (signal strength either of 2/3 or 4/5) was based each previous choice. Whatever his line, each subject observed the previous choices of the subjects *in the low line*. Finally, subjects were asked to make a public prediction about the identity of the selected state. Each subject received 10 French francs for a correct guess and −5 French francs otherwise. At the end of each round uncertainty about the true state was resolved to allow for controlled learning. During a session a subject always belonged either to the low line or to the high line. Table 2 summarizes the progress of a typical round.

Period	Subject's position in the low line	Subject's position in the high line	Observed history
1	1		$\emptyset$
2	2		1st decision in the low line
3	3	1	1st to 2nd decision in the low line
4	4	2	1st to 3rd decision in the low line
$\vdots$	$\vdots$	$\vdots$	$\vdots$
9	9	7	1st to 8th decision in the low line

Table 2: Typical round.

Our innovative design has nice features. First, it allows us to collect a lot of data concerning the potential situations where a cascade should be broken.[11] Second, one can investigate whether subjects' behavior, both the low and high informed types, rely on the position in the decision queue. Third, as low and high informed subjects observe the same history, a highly controlled comparison between low informed and high informed subjects' behavior can be made.

In the second experiment we replicated Experiment 1 but we also elicited subjects' beliefs about the randomly chosen state. For the sake of comparison of subjects' behavior between the two experiments, the same random events, i.e., urns used and private signals, were maintained to run one session in each experiment.

[11] In all, there were 107 potential situations where a cascade break should have been theoretically observed.

Eliciting Beliefs

At each period of the second experiment subjects in the low line were asked to key in a probability vector which represents their beliefs that state α or state β was randomly chosen at the beginning of the round. Subjects in the high line reported their beliefs until the period of their guess. For example, if in a given round a subject in the high line held position 5 then he only reported 5 beliefs. The first elicitation of a subject's belief was made just after he received his signal and before choosing between action A and action B. Otherwise, subjects' beliefs were elicited at the beginning of each period, i.e., after having observed the previous periods' choices. Of course, in period 1, except for the subject who received a signal, subjects' beliefs should reflect the priors. This procedure of beliefs' elicitation allowed us to collect a 9×9 matrix of beliefs for the low line and 7 vectors of beliefs for the high line, whose length goes from 3 to 9, for each round.

Subjects' assessments were rewarded on the basis of a quadratic scoring rule function. Thus, subject i reported his beliefs in period n by keying in a vector $\mu_{in} = (\mu_{in}^{\alpha}, \mu_{in}^{\beta})$ indicating his belief about the probability that the state randomly chosen at the beginning of the round is α or β.[12] In period n, the payoff to subject i when state α was randomly chosen and μ_{in} is the reported belief vector of subject i is given by

$$\Pi_{\alpha} = 0.25 - \frac{1}{8}\left((1-\mu_{in}^{\alpha})^2 + \left(\mu_{in}^{\beta}\right)^2\right) = 0.25\left(1 - \left(\mu_{in}^{\beta}\right)^2\right). \tag{1}$$

The payoff to subject i when state β was randomly chosen is, analogously,

$$\Pi_{\beta} = 0.25 - \frac{1}{8}\left(\left(1-\mu_{in}^{\beta}\right)^2 + (\mu_{in}^{\alpha})^2\right) = 0.25\left(1 - (\mu_{in}^{\alpha})^2\right). \tag{2}$$

It can easily be demonstrated that this reward function provides an incentive for risk-neutral subjects to reveal their true beliefs about the randomly chosen state (see Murphy and Winkler, 1970 for more details).[13]

The payoffs from the assessment task were all received at the end of the experiment.[14] We made sure that the amount of money that could potentially be earned in the assessment part of the experiment was not large in comparison to the game being played. In this respect, the maximum amount that could be earned in the assessment task of Experiment 2 was only 33.75 French francs as compared to the theoretical expected payoff of the decision task: 90 French francs for a low informed subject and 100 French francs for a high informed subject. Table 3 on the following page summarizes the features of both experiments.

[12] In the experiment μ_{in}^{α} and μ_{in}^{β} were keyed in as numbers in [0,100], so are divided by 100 to get probabilities.

[13] While payoffs are maximized by a truthful revelation of beliefs, reporting equal probabilities for each state of Nature would guarantee the largest minimal payment ("secure" stated beliefs). Risk aversion could induce subjects to behave in such a way. We did neither observe a bias toward flat beliefs' vectors (risk averse subjects) nor toward extreme beliefs' vectors (risk loving subjects) in our data.

[14] For the sake of understanding, instead of presenting Equations (1) and (2) to the subjects, we included in the instructions a table summarizing the respective payoff depending on the couple of beliefs reported and the chosen state.

	Elicitation of beliefs	Line	Total number of subjects	Number of rounds per session	Total number of decisions	Total number of reported beliefs
Experiment 1 (3 sessions)	No	Low	27	15	405	0
		High	21	15	315	0
Experiment 2 (3 sessions)	Yes	Low	27	15	405	3645
		High	21	15	315	1890

Table 3: Experimental design.

4 Results

We analyze our results by first checking whether information cascades develop consistently in the laboratory and whether the conformity that cascades cause is brittle. Then, we study the dynamics of elicited beliefs and their consistency with the actions. Finally, we estimate a statistical error-rate model.

In the model of Section 2, theoretical predictions are clear: the equilibrium outcome is unique. Nevertheless, whereas equilibrium decisions for high informed players are unique whatever the history of observed decisions, low informed agents' equilibrium strategies rely on interpretations of observable off the equilibrium path decisions. As there is no unique prediction off the equilibrium path, we only consider decisions following a history that could be part of an equilibrium outcome.[15]

A *cascade situation* is a situation where an action (A or B) constitutes an established pattern, and a subject's signal does not coincide with the established pattern. Let n_{i-1} be the number of a signals less the number of b signals that can be inferred from an equilibrium history h^{i-1}. Formally, player i in the low line (respectively player j in the high line) is in a cascade situation if either $n_{i-1} = 1$ and $t_i = b$ (respectively $t_j = b_S$), or $n_{i-1} = -2$ and $t_i = a$ (respectively $t_j = a_S$). Given a cascade situation in period i, a *cascade behavior* is observed in the low line if player i chooses action A when $n_{i-1} = 1$ and action B when $n_{i-1} = -2$. Similarly, given a cascade situation in period j, a *cascade break* is observed in the high line if player j chooses action B when $n_{i-1} = 1$ and action A when $n_{i-1} = -2$.

4.1 Emergence of Information Cascades: Decision Data in the Low Line

We denote by n_{CS} the total number of cascade situations. The relative frequency of cascade behavior is the ratio $\frac{n_{CB}}{n_{CS}}$, where n_{CB} is the total number of cases in which a low informed subject,

[15] Histories are still included after non-equilibrium decisions as long as those decisions do not lead to an history that cannot be part of an equilibrium outcome. We obtained the same results by including off the equilibrium histories and by assuming that agents associate each action contradicting an equilibrium decision with the corresponding signal.

placed in a cascade situation, has chosen in contradiction with his signal. The relative frequencies of cascade behavior for each session in both experiments are given in Table 4. Though the cascade phenomenon is replicated in our experiments, it is not very strong since cascade behavior is observed 69% of the time overall.[16] It is useful to point out that random decisions yield a relative frequency of cascade behavior equal to 50%. According to χ^2 tests, the observed behavior in cascade situations is significantly different from both the theoretical predicted behavior and a random behavior based on flip coins at the 5 percent level.

Experiment	Session	Cascade behavior
	1	85%
1	3	83%
	5	70%
	Average	79%
	2	61%
2	4	44%
	6	73%
	Average	59%

Table 4: Relative frequencies of cascade behavior in equilibrium histories.

Since averaging over all cascade situations does not give a clear picture of subjects' behavior in the low line, we examine how subjects' mimetism depends on the strength of the majority of observed choices. Figure 1 on the following page represents the relative frequencies of cascade behavior as a function of $|n^A - n^B|$, where n^d is the number of actions $d \in \{A, B\}$ taken up to the current period. After an A action, only 24% of the subjects followed the trend when they received a contradictory signal. After two similar actions not canceled out by previous choices the relative frequency of cascade behavior increases markedly and reaches 64%. When the absolute difference between the number of A and B decisions attains 7, the proportion of cascade behavior is identical to the theoretical one: 100% of the subjects followed the established pattern.[17]

4.2 Fragility of Information Cascades: Decision Data in the High Line

We define the relative frequency of cascade break as the ratio $\frac{n_{SC}}{n_{CS}}$, where n_{SC} is the total number of cases where a high informed subject, placed in a cascade situation, has chosen in accordance with his signal and in contradiction with the majority of previous actions. The relative frequencies of cascade break for each session in both experiments are given in Table 5 on the next page. Overall,

[16]These results are comparable to Anderson and Holt's (1997) findings. Apparently, there is a difference between the relative frequency of cascade behavior in Experiment 1 and the relative frequency of cascade behavior in Experiment 2. However, applying a robust rank-order test (for a description of this test see Siegel and Castellan, 1988, p. 137), we cannot reject the null hypothesis of no difference between the two relative frequencies at any conventional significance levels. This result is confirmed by the estimations reported in Subsection 4.4.

[17]In the last period however, the relative frequency of cascade behavior decreases. Whereas we have no convincing explanation for such an anomaly, subjects' elicited beliefs also reflect this "end-game" behavior (see the next subsection).

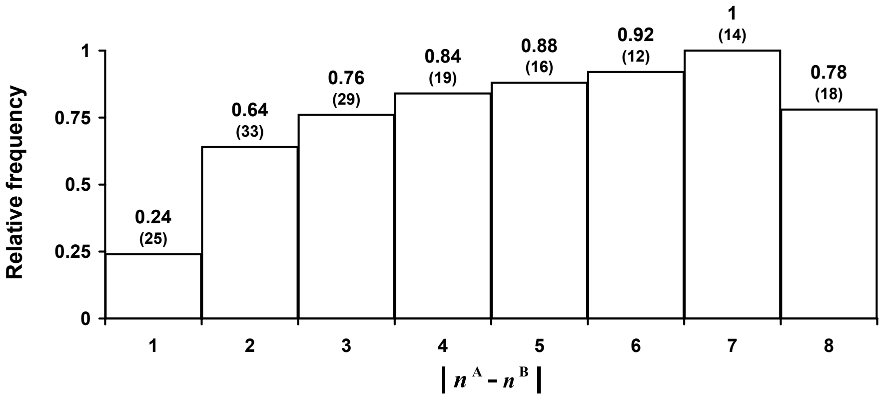

Figure 1: Relative frequencies of cascade behavior as a function of $|n^A - n^B|$ in equilibrium histories (numbers of cascade situations are indicated in parentheses).

in only one third of the situations, a subject who had a strong signal contradicting the established pattern broke the cascade.[18]

Experiment	Session	Cascade break
1	1	27%
	3	26%
	5	53%
	Average	35%
2	2	20%
	4	42%
	6	32%
	Average	31%

Table 5: Relative frequencies of cascade break in equilibrium histories.

Figure 2 on the following page shows the relative frequencies of cascade breaks depending on the strength of the majority of past decisions. At least half of the information cascades are broken in case an imbalance of at most 3 actions in one direction is observed whereas, on average, less than one six of the information cascades are broken as soon as an imbalance of at least 4 actions in one direction is observed.

4.3 Dynamics of Beliefs

In this subsection we look at the dynamics of elicited beliefs. We show that, in contradiction with the theory but in accordance with actual behavior, subjects do not consider choices belonging to

[18]Applying a robust rank-order test, we cannot reject the null hypothesis of no difference between the relative frequency of cascade break in Experiment 1 and the relative frequency of cascade break in Experiment 2 at any conventional significance levels. Again, this result is confirmed by the estimations reported in Subsection 4.4.

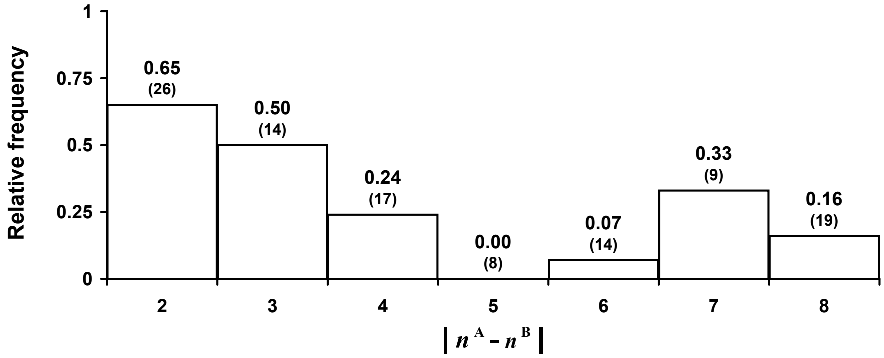

Figure 2: Relative frequencies of cascade break as a function of $|n^A - n^B|$ in equilibrium histories (numbers of cascade situations are indicated in parentheses).

an information cascade as uninformative.

Elicited beliefs in the first period

Table 6 summarizes subjects' average prior beliefs as well as their beliefs after having received a weak signal in the first period, i.e., before having observed any history of actions (standard deviations are given in brackets).[19] As actual beliefs are very close to theoretical ones, a conclusion might be that subjects took in the information structure and were, on average, able to apply Bayes' rule when thinking about others' behaviors is not necessary.

	No signal 675 obs. (both lines)	Signal a 26 obs. (low line)	Signal b 19 obs. (low line)
Session 2	54% (0.12)	78% (0.17)	33% (0.22)
Session 4	51% (0.12)	62% (0.18)	36% (0.10)
Session 6	55% (0.14)	61% (0.15)	37% (0.22)
Average elicited beliefs	53% (0.13)	66% (0.18)	35% (0.18)
Theoretical beliefs	55%	71%	38%

Table 6: Average prior and updated elicited beliefs without history.

History dependent elicited beliefs

Figure 3 on the following page represents the evolution of subjects' beliefs, before being endowed with a private signal, in a cascade with an established pattern of A actions (respectively B actions) when the depth $n^A - n^B$ (respectively $n^B - n^A$) increases. From a theoretical point of view, as public information stops accumulating once a cascade has started, an agent's beliefs stay constant whatever the number of similar actions observed. Clearly, the dynamics of stated beliefs for subjects

[19] Recall that subjects in the high line received their signal only after period 2.

without private information do not reflect this theoretical feature. On the contrary, subjects' beliefs increase when the depth increases and the dynamics of actual beliefs are very close to the dynamics of "Private Information" (PI) beliefs, i.e., to the dynamics of beliefs that agents would have if it was mutually known that all agents follow their own signal.

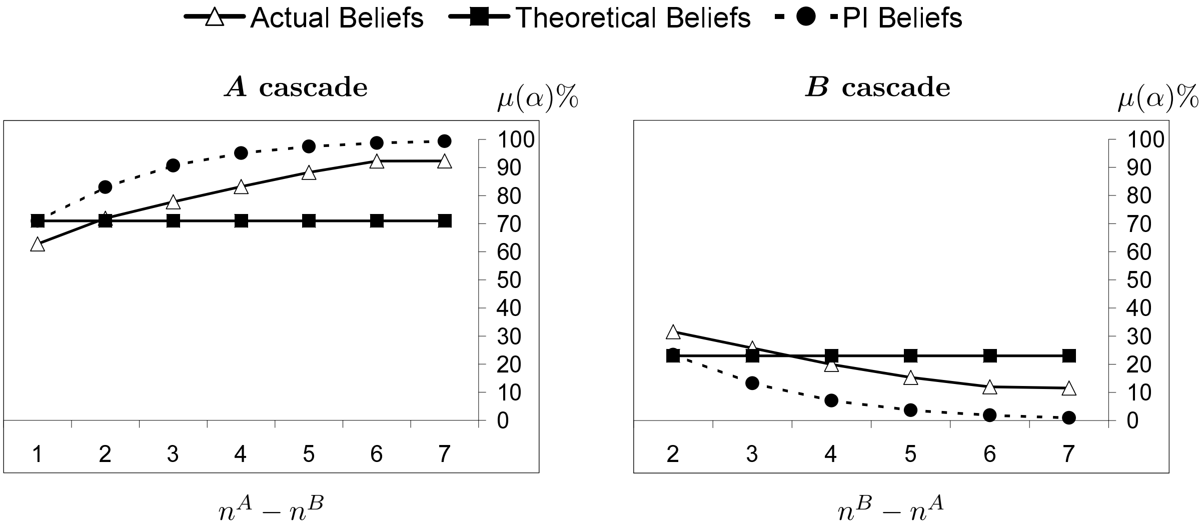

Figure 3: Dynamics of beliefs in A and B cascades without private signal (low and high line).

Besides, because we are able to measure beliefs directly we can directly check whether subjects' decisions are consistent with their beliefs, i.e., whether subjects maximize their expected payoff given their beliefs. Table 7 gives the relative frequency of actions in accordance with subjects' beliefs for each session, respectively in the low and high line. On average, beliefs and actions are consistent about 97% of the time. Thus, there is a clear indication that subjects have rationally linked their stated beliefs to their actions. Given this observation we come back to the cascade phenomenon by linking together subjects' beliefs and behaviors in cascade situations.

	Low line	High line	Total
Session 2	99%	100%	99%
Session 4	95%	95%	95%
Session 6	96%	100%	98%
Average	96%	98%	97%

Table 7: Percentages of subjects' actions consistent with their beliefs.

The observation that the relative frequency of cascade behavior increases with the depth of a cascade seems to be explained by the fact that agents' beliefs do not stay constant in a cascade. To see the link between the relative frequency of cascade behavior and the way agents update their beliefs we represent in Figure 4 on the following page the dynamics of the low line subjects' beliefs when the depth of similar actions increases and their (low quality) signal contradicts the majority of previous actions.[20] Figure 4 shows that when the majority of previous decisions are A actions and a subject received a weak b signal, then his belief becomes greater than 50% only after a depth of 2, and his beliefs is largely greater than 50% with a depth larger than 3. The same phenomenon appears with a majority of previous B decisions, albeit beliefs cross 50% only after a depth of 3 (instead of 2 in a sequence with a majority of A). This explains what we observed in Figure 1: cascade behaviors are very low with a depth of 1, and very high with a depth greater than 3.

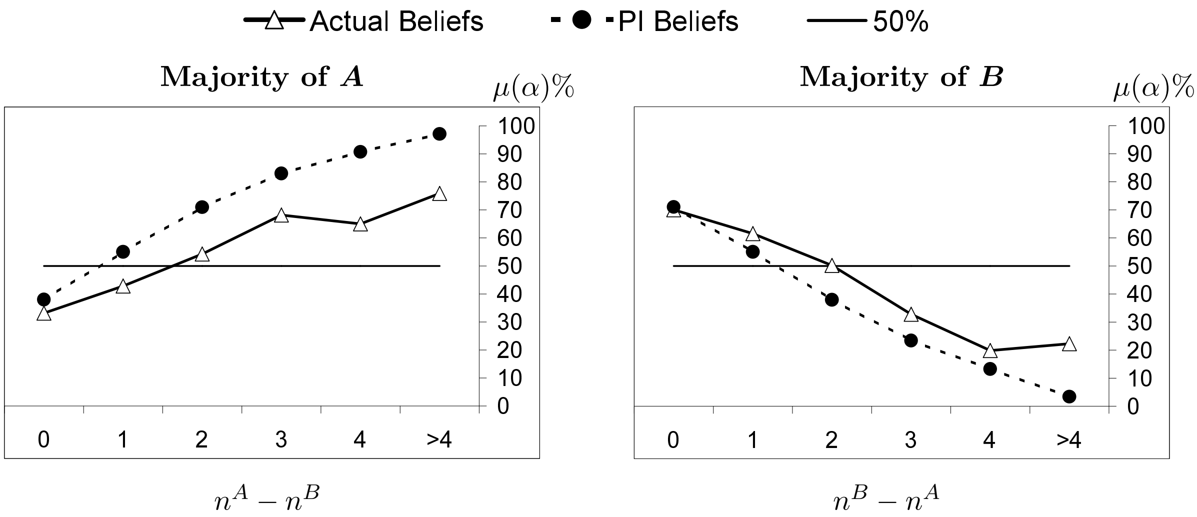

Figure 4: Dynamics of beliefs with a weak contradictory signal (low line).

Similarly, the fact that the relative frequency of cascade break decreases with the depth of a cascade can also be explained by the way agents update their beliefs. Figure 5 on the following page shows the dynamics of high subjects' beliefs when the depth of similar actions increase but when their (high quality) signal contradicts the majority of previous actions. In both cases (majority of previous A actions and majority of previous B actions) we see that agents believe more in the state corresponding to the majority of previous actions than in their own contradictory signals after a depth of 3. This explains what we observed in Figure 2: when the depth of the cascade in larger or equal to 3, agents in a cascade often do not break the cascade even with a strong contradictory signal.

[20]For depths strictly greater than 4 we only report the average of subjects' beliefs since the number of observations falls from 117 with depth 0 to 15 with depth 5. In this way, we collected 54 observations for depths strictly greater than 4.

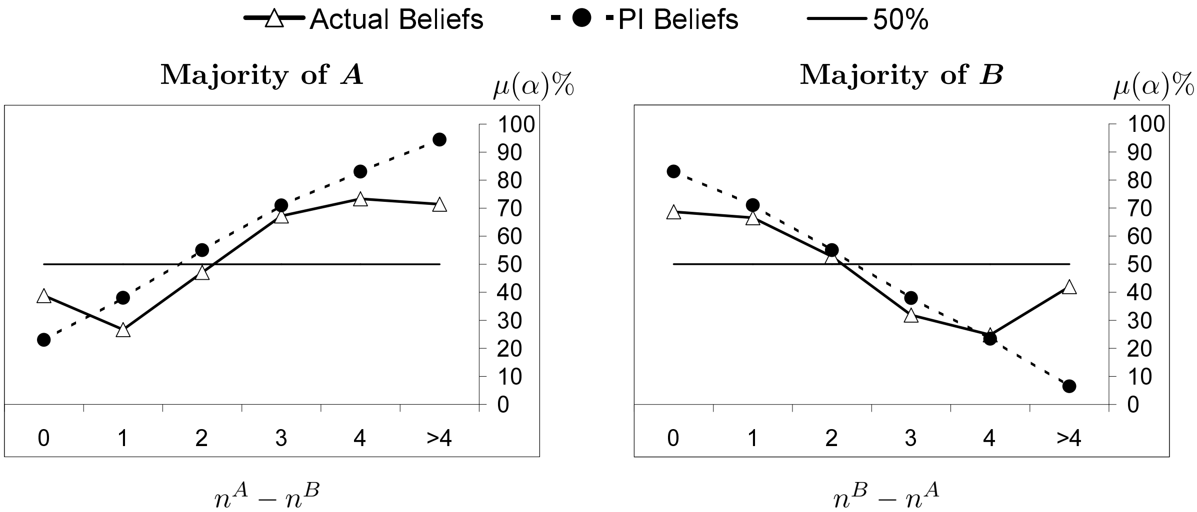

Figure 5: Dynamics of beliefs with a strong contradictory signal (high line).

4.4 A Statistical Error-Rate Analysis

In this subsection, following Kübler and Weizsäcker (2004), we present and estimate an error-rate model that uses logistic response functions to determine choice probabilities and specifies separate parameters for the response rationality on each level of reasoning. That is, it allows for different error rates at each step of thinking about thinking ... about others' behavior. In particular, the model does not impose the assumption that subjects have a correct perception of other subjects' error rates, or that they have a correct perception of other subjects' perceptions of third subjects, and so on.

Subjects are assumed to employ a logistic choice function with precision parameter $\lambda_1 \geq 0$ when making their choices, i.e., to choose A with probability

$$\Pr(A \mid t_i, h^{i-1}, \lambda_1) = \frac{\exp\left(\lambda_1\left(15\mu_i(t_i, h^{i-1}) - 5\right)\right)}{\exp\left(\lambda_1\left(15\mu_i(t_i, h^{i-1}) - 5\right)\right) + \exp\left(\lambda_1\left(15(1 - \mu_i(t_i, h^{i-1})) - 5\right)\right)},$$

and to choose B with the complementary probability, where $h^0 = \emptyset$ is the empty history. As usual in such logistic-choice models (McKelvey and Palfrey, 1995, 1998), the parameter λ_1 captures the response precision of the decision maker; the higher λ_1, the more "rational" are the decisions.

It is assumed that subjects are aware that all other subjects follow the logistic decision process described above, with the exception that they attribute a possibly different precision parameter to the decisions of their opponents: λ_2 instead of λ_1. Analogously, when a subject considers the reasoning that others apply when thinking about third subjects, we allow for a third parameter λ_3. For even longer chains of reasoning, additional higher-level parameters are used (see Kübler and Weizsäcker (2004) for more details). The length of the reasoning process is reflected by the first parameter that is indistinguishable from zero. For instance, if λ_1 is indistinguishable from zero, then players tend to behave randomly, and if λ_2 is indistinguishable from zero players are not able to make inference from predecessors' decisions. Note that under the restriction that the data generating process is correctly specified by Bayesian Nash Equilibrium, we should obtain estimates of $\lambda_1, \ldots, \lambda_9$ that are all ∞.

Table 8 (Table 9, respectively) contains the estimation results on the data from subjects in the low line (high line, respectively) for the two experiments and the pooled data, where "LL" is the value of the log likelihood. An empty cell in the table ("—") indicates that the parameter cannot be estimated well as the lower parameter is indistinguishable from 0. Table 10 also contains the estimation results of the restricted econometric model corresponding to the logit Quantal response equilibrium model of McKelvey and Palfrey (1998) in which players know the common rate of errors everyone makes. That is $\lambda_i = \lambda^C$ for each element $i = 1, 2, \ldots, 9$.

To construct hypotheses tests, we conduct a simulation study. We generate a set of simulated choices, based on the maximum likelihood estimates and the sample size used in our experiment. We re-estimate the model and obtain a vector of maximum likelihood estimates of the simulated data. We repeat the process a large number of times. We obtain a distribution that resembles

	Pooled Data	Experiment 1	Experiment 2
λ_1	**3.49**	**3.55**	**3.51**
	[3.07 4.18]	[2.95 4.66]	[2.95 4.40]
	[3.11 4.06]	[3.02 4.47]	[2.97 4.26]
λ_2	**2.52**	**3.62**	**1.81**
	[2.02 3.51]	[2.49 17.30]	[1.40 2.72]
	[2.03 3.45]	[2.56 8.16]	[1.46 2.48]
λ_3	**1.81**	**2.01**	**0.97**
	[0.87 4.12]	[1.04 5.93]	[0.52 2.47]
	[1.00 2.86]	[1.16 4.71]	[0.59 2.06]
λ_4	2.30	1.65	-2.61
	[-1.75 30.97]	[-2.15 20.00]	[-20.00 1.87]
	[-1.30 24.67]	[-1.95 19.60]	[-19.50 1.48]
$\lambda_{5,6,7,8,9}$	—	—	—
$-$LL	207.43	92.91	111.20

Table 8: Response precisions estimations and the 95% and 90% confidence intervals from the experimental data for the low line. Boldface parameters are significantly different from 0.

the sampling distribution of the maximum likelihood estimator. The standard errors are obtained by the sample standard deviation. Confidence intervals are obtained by deleting values of the appropriate number of estimates from the sorted array of maximum likelihood estimates. Using the data of experiment 1, experiment 2 or the pooled data from the low line or the high line we conclude that the parameters λ_1, λ_2 and λ_3 are different from zero, and λ_4 is not different from zero, except in experiment 2 in the high line where only λ_1 and λ_2 are different from zero but λ_3 is not different from zero. Using any data from the low or high line we conclude that the parameter λ_1 is not different from the parameter λ^C estimated with a common error rate. However, the parameter λ_2 is different from λ^C at the 10% level of significance in any data from the low line, and the data from experiment 2 in the high line. The parameter λ_3 is also different from λ^C with the pooled data of both lines, in experiment 2 of the low line and in experiment 1 of the high line. Finally, except in experiment 1 in the low line we reject the hypothesis that the parameter λ_1 is equal to the parameter λ_2 based on the 95% confidence interval around the estimate of λ_1.

	Pooled Data	Experiment 1	Experiment 2
λ_1	**1.97**	**2.42**	**1.63**
	[1.59 2.32]	[1.99 3.96]	[1.22 2.06]
	[1.66 2.29]	[2.08 7.38]	[1.34 2.00]
λ_2	**5.34**	**3.87**	**6.43**
	[3.35 17.32]	[2.50 8.54]	[2.86 89.30]
	[3.58 9.99]	[2.68 7.38]	[3.12 35.90]
λ_3	**0.86**	**0.92**	-0.32
	[0.47 1.55]	[0.38 2.21]	[-21.42 1.25]
	[0.53 1.43]	[0.50 1.85]	[-13.77 0.88]
λ_4	0.60	0.84	—
	[-20.02 3.17]	[-8.38 20.00]	
	[-20.00 2.96]	[-5.68 19.01]	
$\lambda_{5,6,7,8,9}$	—	—	—
$-$LL	161.80	66.78	90.60

Table 9: Response precisions estimations and the 95% and 90% confidence intervals from the experimental data for the high line. Boldface parameters are significantly different from 0.

	Low Line			High Line		
	Pooled Data	Experiment 1	Experiment 2	Pooled Data	Experiment 1	Experiment 2
λ^C	3.44	3.85	3.11	1.98	2.27	1.77
$-$LL	212.40	95.17	115.88	204.16	87.39	114.95

Table 10: Response precisions estimations of the restricted econometric model (logit Quantal response equilibrium) from the experimental data for the low and high lines.

5 Conclusion

In this paper we investigated some aspects of subjects' behaviors and beliefs in an innovative information cascade experiment. Two lines of subjects were considered. In the low line, a sequence of nine subjects, each of them endowed with a 2/3 quality signal, publicly guessed a randomly chosen state of Nature. In the high line, a sequence of seven subjects, each of them endowed with a 4/5 quality signal, was given the same task by observing the choices made in the low line. Whereas rational herd behavior predicts a high occurrence of information cascades in the low line, optimal behavior always consists in following private information in the high line. Besides, in one of our two experiments, we directly elicited subjects' beliefs. Since the conformity of followers in a cascade has no informational value, beliefs should stay constant and should be upset by more informative signals.

Although herd behavior emerged in most rounds, we do not conclude that actual choices are consistent with all characteristics of the information cascade theory. In particular, our data analyses show that the uniformity stemming from previous decisions becomes more robust as the number of similar actions increases. Most importantly, the main contribution of our experimental analysis is to show that conformity is not brittle since cascades were rarely shattered by more informed subjects. Consequently, the benefit of diverse information sources was lost due to the subjects' behavior in the high line. The rationale of these claims are summarized below.

First, a close look at the dynamics of cascade behavior reveals that subjects follow an established pattern only when a sufficiently pronounced imbalance of decisions in one direction is observed. Second, once an information cascade starts and develops for several periods, more informed subjects rarely go against the established pattern of decisions. Such an imitative behavior is strengthened by the depth of the information cascade, i.e., by the weight of the majority emerging from previous actions. Third, the dynamics of subjects' elicited beliefs, which was shown to be extremely consistent with subjects' behavior, in not constant with the establishment of a majority of similar previous actions. This feature explains both the dynamics of cascade behavior and cascade break.

We scrutinized the observed discrepancy between actual choices and rational herd behavior by considering two versions of a statistical error-rate model, one corresponding to the logit agent Quantal response equilibrium (McKelvey and Palfrey, 1995, 1998) and another allowing for different error rates on different levels of reasoning (Kübler and Weizsäcker, 2004). Errors in decision making that are negatively correlated with the cost of choosing the incorrect decision have a theoretical role in the source of information. Indeed, when agents believe that others make some errors, a decision off the equilibrium path increases sharply public information since it tends to be interpreted as a revelation of the deviant's signal. Then, this error counters the herd externality and it curbs the emergence of an informational cascade. As a consequence, when a pattern of decisions is established, individuals should realize that public information is less and less rich but, because they think that others tend to go against the trend when their signal contradicts the established pattern, they never consider those decisions as uninformative.

References

Anderson, L. and C. Holt (1997): "Information Cascades in the Laboratory," *American Economic Review*, 87, 847–862.

Banerjee, A. (1992): "A Simple Model of Herd Behavior," *Quarterly Journal of Economics*, 107, 797–818.

Bikhchandani, S., D. Hirshleifer, and I. Welch (1992): "A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades," *Journal of Political Economy*, 100, 992–1026.

——— (1998): "Learning from the Behavior of Others: Conformity, Fads, and Informational Cascades," *Journal of Economic Perspectives*, 12, 151–170.

Boun My, K. (2000): "A VB Software for Experiments on Information Cascades and Beliefs," BETA, CNRS.

Celen, B. and S. Kariv (2004): "Distinguishing Informational Cascades from Herd Behavior in the Laboratory," *American Economic Review*, 94, 484–498.

Chamley, C. (2003): *Rational Herds: Economic Models of Social Learning*, Cambridge University Press.

Huck, S. and J. Oechssler (2000): "Informational Cascades in the Laboratory: Do they Occur for the Right Reasons?" *Journal of Economic Psychology*, 21, 661–671.

Huck, S. and G. Weizsäcker (2002): "Do Players Correctly Estimate What Others Do? Evidence of Conservatism in Beliefs," *Journal of Economic Behavior and Organization*, 47, 71–85.

Hung, A. and C. Plott (2001): "Information Cascades: Replication and an Extension to Majority Rule and Conformity Rewarding Institutions," *American Economic Review*, 91, 1508–1520.

Kübler, D. and G. Weizsäcker (2004): "Limited Depth of Reasoning and Failure of Cascade Formation in the Laboratory," *Review of Economic Studies*, 71, 425–441.

——— (2005): "Are Longer Cascades More Stable?" *Journal of the European Economic Association*, 3, 330–339.

McKelvey, R. and T. Palfrey (1995): "Quantal Response Equilibria for Normal Form Games," *Games and Economic Behavior*, 10, 6–38.

——— (1998): "Quantal Response Equilibria for Extensive Form Games," *Experimental Economics*, 1, 9–41.

McKelvey, R. D. and T. Page (1990): "Public and Private Information: An Experimental Study of Information Pooling," *Econometrica*, 58, 1321–1339.

Murphy, A. H. and R. L. Winkler (1970): "Scoring Rules in Probability Assessment and Evaluation," *Acta Psychologica*, 34, 273–286.

Nöth, M. and M. Weber (2003): "Information Aggregation with Random Ordering: Cascades and Overconfidence," *The Economic Journal*, 113, 166–189.

Nyarko, Y. and A. Schotter (2002): "An Experimental Study of Belief Learning Using Elicited Beliefs," *Econometrica*, 70, 971–1005.

Offerman, T., J. Sonnemans, and A. Schram (1996): "Value Orientations, Expectations and Voluntary Contributions in Public Goods," *Economic Journal*, 106, 817–845.

Siegel, S. and N. J. Castellan (1988): *Nonparametric Statistics for the Behavioral Sciences*, McGraw-Hill, New York.

Sonnemans, J. and T. Offerman (2001): "Is the Quadratic Scoring Rule really incentive compatible?" CREED working paper, University of Amsterdam.

Ziegelmeyer, A. (2001): "Informational Externalities and Sequential Interactions: Experimental Evidence on Social Learning," Unpublished Ph.D. thesis. Louis Pasteur University, Strasbourg.