# Information Aggregation with Random Ordering: Cascades and Overconfidence*

Markus Nöth[a,b] and Martin Weber[b]

December 17th, 1999

**Abstract**

In economic models, it is usually assumed that agents aggregate their private and all available public information correctly and completely. In this experiment, we identify subjects' updating procedures and analyze the consequences for the aggregation process. Decisions can be based on private information with known quality and observed decisions of other participants. In this setting with random ordering, information cascades are observable and agents' overconfidence has a positive effect on avoiding a non-revealing aggregation process but it reduces welfare in general.

JEL: C92, D8

keywords: aggregation, Bayes' rule, cascades, experiment, overconfidence

---

*The authors gratefully acknowledge the financial support for this research which was provided by the Deutsche Forschungsgemeinschaft (grants No381/1 and We993/7-2). Carlo Kraemer and Tobias Kremer programmed the software for this project. Helpful comments were received from Rachel Croson, Wolfgang Gerke, Charles Holt, Susanne Prantl and participants at the Economic Science Association 1998 meeting, the European Finance Association 1999 meeting and at the Wharton Finance Micro Lunch seminar.

[a]Wharton School; University of Pennsylvania; 2300 Steinberg Hall-Dietrich Hall; Philadelphia, PA 19104-6367; USA

[b]Lehrstuhl für Allgemeine Betriebswirtschaftslehre, Finanzwirtschaft, insbes. Bankbetriebslehre; Universität Mannheim; 68131 Mannheim; Germany; noeth@bank.BWL.uni-mannheim.de; weber@bank.BWL.uni-mannheim.de

## Information Aggregation with Random Ordering: Cascades and Overconfidence

In most economic models, it is assumed that agents apply rules of conditional probability (Bayes' rule) to decide based on private and public information. With sequential decision making without a pricing mechanism, Bayesian updating leads to the development of information cascades. In an information cascade, an agent takes an identical action for all private signals because they cannot overrule the available public information. Information cascades as modeled by Banerjee (1992), Bikhchandani, Hirshleifer, and Welch (1992) and (1998), or Welch (1992) are stable as soon as two consecutive decisions are identical because it is assumed that the information quality for all agents is identical and that all agents act rationally. Even if there exists a distribution of information qualities the results would not change in the limit as long as the information is positively correlated with the true value (Lee 1993). Clustering of decisions or herding can also occur because of endogenous timing decisions and waiting costs as in Gul and Lundholm (1995) and in Zhang (1997), or because of exogenous incentives (Scharfstein and Stein 1990).

Informational cascades and herding models are typically used to explain clustering of decisions. But in financial markets with endogenously determined prices, cascades, i.e. bubbles and crashes, can arise only under specific circumstances as Avery and Zemsky (1998) showed. One possible restriction are transaction costs as in Lee (1998). Due to the transaction costs, small (rational) biases can accumulate. Triggered by a "rare event", all previously unrevealed information will be aggregated leading to a crash or jump of asset prices. Alternatively, arbitrage restrictions such as short trading horizons or limited funds can prevent arbitrageurs from correcting mispricing as Dow and Gorton (1994) and Shleifer and Vishny (1997) demonstrate. In both models, rational agents cannot trade to profit from obvious mispricings because they may have to liquidate their positions before the true value is revealed. Allen, Morris, and Postlewaite (1993) showed the same effect by imposing short-selling restrictions. Even in efficient markets, noise traders can create risk that results in price deviations from rational expectations (De Long, Shleifer, Summers, and Waldmann 1990a). If systematic non-rational behavior of some traders is known, rational traders can profit by anticipating the future order flow (De Long, Shleifer, Summers, and Waldmann 1990b). In all these models, agents possessing information are assumed to be rational. Thus, the question how systematic biases and random irrational behavior influences the aggregation process is not addressed.

With our experiment we want to evaluate the structure of agents' updating behavior. The experimental method is chosen mainly because it is possible to control all major parameters, to vary the available information, and to repeat identical situations to account for potential learning effects. Furthermore, no restrictions are imposed on how participants use private and public information. As a result, theoretical predictions and actual behavior can be compared to evaluate and to explain observed differences. More specifically, it is possible to distinguish between rational herd behavior and non-Bayesian behavior. Our experimental setting will also demonstrate why huge swings in opinions or asset prices might be observed although no new information seems to be available. Individual overconfidence within cascades is identified to be the most likely reason since we can eliminate many other explanations that are not consistent with the observed decisions.

To keep our experimental design as simple as possible we focus on sequential decisions with random ordering of the agents who decide once in every round. At the end of each round, uncertainty about the true value is resolved to allow for controlled learning. Our design is an extension of Anderson and Holt (1997).[1] We introduce two instead of one signal qualities because a simple counting heuristic leads in a design with a uniform signal quality to the same observed behavior as using Bayes' rule.[2] The signal quality is part of the private information and known with certainty. This modification increases the complexity of the decision problem sufficiently to eliminate the success of simple heuristics. In addition, it reflects economic situations more appropriately because agents usually do not receive identical signal qualities. Different information qualities increase on the one hand the information content of observed decisions but introduce on the other hand uncertainty about others' information.[3] Thus, there is enough room for identifiable non-Bayesian updating behavior. Finally, two signal qualities reduce the likelihood that agents have to randomize their decision because of inconclusive private and public information.[4]

Potential cascades can collapse in our design if an agent receives high quality information or if somebody believes more in her private information than justified by Bayes' rule. Note that putting more weight on the own private information might be a signal for overconfidence but it can be a (rational) reaction to others' behavior, too. Whereas we can distinguish between superior information and overconfidence since we know the signal distribution, it is rather difficult observing only others' predictions. As a result, the aggregation process can switch from a cascade to a reverse cascade and vice versa either because of superior information or because of undetected overconfidence.

In markets, overconfidence can cause speculation because traders are "certain" that they have superior skills or information. As a result, information mirages can develop in which the price process looks as if new information exists (Camerer and Weigelt 1991). Smith, Suchanek, and Williams (1988) and Porter and Smith (1994) investigated experiments in which huge bubbles occurred that are based mostly on overconfident speculators.[5] Overconfidence and other results from individual experiments in psychology and economics have been incorporated in market models recently.[6] Daniel, Hirshleifer, and Subrahmanyam (1998) and Gervais and Odean (1998) use individuals' overconfidence to explain overreaction and volatility changes. Although it is legitimate to use results from individual decision making to build market models, two potential problems have to be addressed. First, individual behavior varies largely and may not be as stable as assumed in models. Second, agents might identify or anticipate others' behavior and try to act accordingly. Whether this attempt offsets or increases the effect of non-Bayesian behavior on information aggregation is an open question. The answer depends on whether others' errors

[1] Huang and Plott (1999) replicated and extended Anderson and Holt (1997) to investigate the effect of different reward mechanisms on the evolution of cascades.

[2] About one third of the participants used the counting heuristic in a modified asymmetric design when Bayesian updating would lead to the alternative prediction (Anderson and Holt 1997).

[3] This uncertainty together with the above mentioned uncertainty about others' behavior creates composition uncertainty as in Avery and Zemsky (1998).

[4] Anderson and Holt (1997) assume that agents follow their own signal in this situation. This assumption is justified since a small probability for incorrect updating by other participants would lead to this prediction instead of randomizing.

[5] See Camerer (1989) for an overview about earlier research to explain bubbles and fads.

[6] Camerer (1995) provides an overview over individual decision making.

are correctly identified or anticipated, or not. For these two questions, we want to find some answers within our experimental setting.

The most important result is that participants do not make their predictions using Bayes' rule but they employ identifiable heuristics, which put too much weight on private information. The heuristics are based on overconfidence. As a consequence a relatively large number of potential cascades collapse or do not develop at all. However, participants are able to increase the number of correct predictions significantly above their private information level despite their own and others' updating mistakes using specific heuristics which improve predictions especially if private information is not very reliable. In relative terms, more ex post incorrect (reverse) cascades collapse. But in absolute numbers, ex post correct cascades are destroyed more often than reverse cascades due to systematic mistakes. As a result, groups' welfare decreases compared to the situation in which all participants use Bayes' rule.

We proceed with the experimental design and procedures. In section 2, we will present the information aggregation theory for this experiment. Section 3 contains the main results and an analysis of observed cascades and their survival. The final section 4 contains a summary and some ideas about design extensions.

# 1 Design and Procedures

As mentioned, we extended the experimental design of Anderson and Holt (1997) in two respects - two different information qualities and using computers to increase the number of repetitions per experimental session. With more observations within one session it is now possible to analyze whether individuals or the whole groups learn during a session. Based on private and public information each of the six participants in a session has to predict in every round whether state $A$ or state $B$ occurs. Public information consists of all predictions that are already made within a round. Thus, each of the six subjects in a session faces the following situation:

- She receives an independently drawn private signal about the true state. In addition, the signal's strength indicates the probability that this signal is correct.
- She observes other participants' predictions about the true state, i.e. state $A$ or $B$, made by those acting before her. However, she cannot identify neither other participants' private signals and strengths with certainty[7] nor the identity of these other participants since predictions were submitted anonymously and the participants' ordering was determined randomly for each round. Obviously, the first participant in each round has no public information.
- As a result, each agent must decide how to aggregate private and public information in order to predict the true state and maximize her own profits.

[7]In some situations it is possible to infer the signals' strength of the immediate predecessor assuming Bayesian updating.

At the beginning of each round the state is determined. Both states $(A, B)$ occur with the same probability $\left(p_A = p_B = \frac{1}{2}\right)$. Then, the ordering of all six subjects is fixed randomly for this round. Finally, private signals $(i_X)$ with $i \in \{a, b\}$ and $X \in \{W, S\}$ are generated independently for each agent in a two step procedure depending on the realized state:

1. The signals' strength is drawn first. It is either *weak* or *strong* with probability $p_W = p_S = \frac{1}{2}$.
2. If the signal is *strong* $(X = S)$, the private information $i \in \{a, b\}$ is correct with probability $p(A \mid a_S) = p(B \mid b_S) = \frac{4}{5}$.

   If the signal is *weak* $(X = W)$, the private information $i \in \{a, b\}$ is correct with probability $p(A \mid a_W) = p(B \mid b_W) = \frac{3}{5}$.

Thus, even the *weak* signal contains some information about the realized state.[8] Each participant receives in addition to the public information one of four possible private signals: $i_X \in \{a_S, a_W, b_S, b_W\}$. Based on this information, the participant has to decide about her own prediction for which she will get 300 cu if the prediction is correct and 100 cu otherwise. This information structure is common knowledge because it is explained as part of the instructions (see appendix).

The experiment was performed using software that was developed specifically for this experiment. The screen of a subject at position IV before she made her decision is shown in Figure 1.

Figure 1: Screenshot at position IV

This screenshot shows all available information for the participant who has to submit a prediction at position IV. The private information for this participant is $a_S^4$, i.e. she receives a strong signal that is correct with probability $p(A \mid a_s) = \frac{4}{5}$. The public information consists of

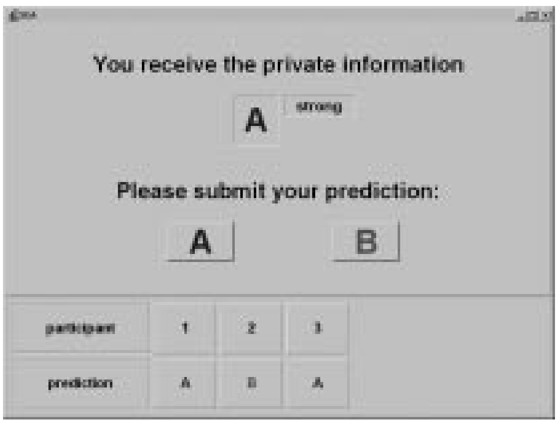


[8]The probabilities associated to strong and weak signals are selected to satisfy the following restrictions. First, the difference of information quality between *strong* and *weak* signals should be as large as possible to increase the value of public information. Second, the weak signal should be considerably more informative than having no information at all. Third, the strong signal should not contain too much information since otherwise no information aggregation task would remain. Fourth, we wanted to have on average the same information content as in Anderson and Holt (1997) who chose $p(\text{state} \mid \text{signal}) = \frac{2}{3}$. Finally, the probabilities should be some "prominent" number such that subjects understand the design easily within 20 minutes.

The observable predictions $(A_1, B_2, A_3)$ can be used to update the own private information $a_S$ (*A strong*). The rational updating procedure assuming rationality for the first three participants is analyzed in section 2. Predicting the state might be easier in some situations (e.g. first three predictions: $B_1, B_2, B_3$; own signal: $b_S^4$) than in others (e.g. first three predictions: $A_1, A_2, B_3$; own signal: $a_W^4$). As a result, the time between getting the signal and predicting the state might depend on the complexity of the individual problem. Because other agents might try to learn by evaluating the length of this time interval, the new signal is delayed by a minimum of three seconds and by a maximum of seven seconds to generate a noisy "time" signal. This procedure is public information. Figure 2 illustrates the procedure within one round. Note that subjects face no time restrictions if they have to submit their prediction.

Figure 2: Procedure within one round

The decision procedure within one round is illustrated in this figure. After the state is determined, the signals' strengths and information for each subject are drawn. In addition, the ordering is randomly fixed for the round. Then, each subject receives the private information as soon she has to submit a prediction. There is no time limit for submitting a prediction, which will become public knowledge. The next subject receives her private information after a random delay of three to seven seconds.

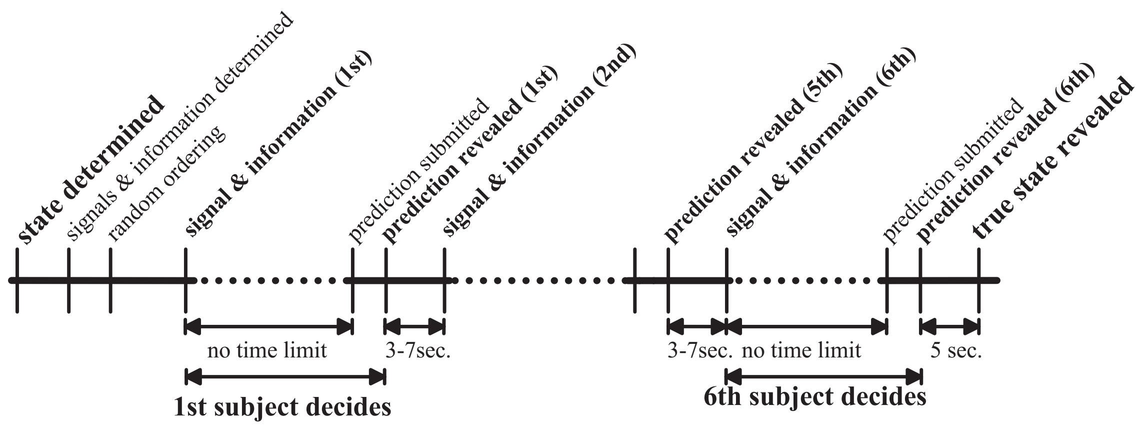


Each session lasted about 110 minutes and consisted of at least 74 rounds (maximum: 86 rounds) of which the first three periods were part of the instructions and thus not paid. The relatively large number of rounds per session enables us to evaluate the data with respect to learning. Moreover, the questionnaire, which subjects filled out at the end of the experiment, will help to distinguish between systematic non-Bayesian behavior and random errors since participants were asked to describe their decision heuristics at the end of the experiment. 126 subjects participated in this experiment (=21 sessions). They were recruited from undergraduate and graduate business administration courses at the University of Mannheim, Germany, and had no previous experience with this experiment. Each session lasted about two hours. All earned currency units were converted to Deutsche Mark (DM) and rounded up to the next DM at the end of each session. Participants earned on average 31.79 DM with a minimum of 27.00 DM and a maximum of 36.00 DM.

## 2 Rational Bayesian Strategy

The obvious benchmark to analyze the experimental data is based on Bayes' rule (BR) assuming that every participant acts accordingly in every situation. By contrast, public information contains no information under the alternative private information (PI) assumption. Agents believing that PI is optimal are very overconfident because they consider others' decisions as being completely useless under all circumstances.

Under both benchmarks, a subject at position I should predict according to her signal since this is always better than random guessing. Thus, the first participant should predict state $A$ if she has received an $a_x^1$-signal and state $B$ otherwise.[9]

The second participant who observes the first prediction (e.g. $A_1$) can infer using BR that the predicted state will occur with probability $p\left(A \mid A_1\right) = \frac{7}{10}$. If she receives a strong signal she should predict according to her private information. Thus, in this situation BR and PI lead to the same prediction. However, a weak signal is dominated by the first participant's prediction. Suppose that the first prediction is $A_1$. It is obvious that signals $a_S^1$ and $a_W^1$ imply a prediction of state $A_1$. Observing a first prediction $A_1$ the private information $b_W^2$ cannot lead to a prediction $B_2$ in this situation because the first decision is based with the same probability either on signal $a_W^1$ or $a_S^1$, which are at least as informative as $b_W^2$. Figure 3 shows the possible prediction paths up to position III with the respective probabilities that state $A$ will occur.

The prediction history $h_{id}^2 = A_1A_2$ with $p\left(A \mid A_1A_2\right) = \frac{\frac{7}{10}\cdot\frac{3}{5}\cdot\frac{1}{2}}{\frac{7}{10}\cdot\frac{3}{5}\cdot\frac{1}{2}+\frac{3}{10}\cdot\frac{2}{5}\cdot\frac{1}{2}} = \frac{7}{9} \cong 0.778$ leads to the same prediction pattern at position III, i.e. only a $b_S$ signal can prevent the development of a cascade at this stage. The posterior probabilities are calculated as follows: $p\left(A \mid A_1A_2i_x^3\right) = \frac{p\left(A_1A_2i_x^3|A\right)*p(A)}{p(A_1A_2i_x^3|A)*p(A)+p(A_1A_2i_x^3|B)*p(B)}$. Because signal $i_x^3$ is independently drawn from signals $i_x^1$ and $i_x^2$, it follows immediately: $p\left(A \mid A_1A_2i_x^3\right) = \frac{p(A_1A_2|A)*p\left(i_x^3|A\right)*p(A)}{p(A_1A_2|A)*p(i_x^3|A)*p(A)+p(A_1A_2|B)*p(i_x^3|B)*p(B)}$. Summing up the calculations, the posterior probabilities at position III with history $h_{id}^2 = A_1A_2$ (see figure 3) are: $p\left(A \mid h_{id}^2a_S^3\right) = \frac{14}{15} \cong 0.933$, $p\left(A \mid h_{id}^2a_W^3\right) = \frac{21}{25} = 0.840$, $p\left(A \mid h_{id}^2b_W^3\right) = \frac{7}{10} = 0.700$, and $p\left(A \mid h_{id}^2b_S^3\right) = \frac{7}{15} \cong 0.467$.

Observing the other possible history $h^2 = A_1B_2$, all remaining participants know that the second decision is based on $b_S^2$. As a result, only $a_S^3$ can lead to prediction $A_3$. Thus, the prediction history $h^3 = (A_1, B_2, A_3)$ implies two contradicting strong signals at positions II and III which neutralize each other. The probability for state $A$ is the same as after position I observing a prediction $A_1$ $\left(p\left(A \mid A_1B_2A_3\right) = \frac{7}{10} = p\left(A \mid A_1\right)\right)$. The prediction paths displayed in figure 4 are based on histories $h_{id}^3 = A_1A_2A_3$ (*), on $h^3 = A_1A_2B_3$ (**) or on $h^3 = A_1B_2B_3$ (***).

After three identical predictions (*) one should always predict the same state regardless of the own private signal, i.e. an information cascade arises rationally. State $B$ is predicted at position III after two A-predictions (**) only if a $b_S^3$ signal has been drawn. As a result, it is rational to predict the state

[9] Some notation: probabilities are always implicitly expressed in relation to state $A$. The general position index is denoted by $y \in \{1, 2, ...6\}$. $i_x^y$ is the private information with $i \in \{a, b\}$ and $x \in \{S, W\}$. Decisions are denoted as $D_y$ with $D \in \{A, B\}$. $h^y$ are histories of predictions that can be observed at position $y + 1$ before making a prediction. $h_{id}^y$ refer to identical predictions, i.e. by convention $h_{id}^y = A_1...A_y$.

Figure 3: Some prediction paths at positions I, II and III

The decision situations at positions I, II and III are displayed depending on the private signal $i_X \in \{a_S, a_W, b_W, b_S\}$ and the observable decision history. Based on probabilities for state $A$ the rational decision is shown. Moreover, the posterior probabilities for an observed decision are provided. History $h^3 = A_1B_2A_3$ leads to the same posterior probability (0.700) as history $h^1 = A_1$. The star symbol indicates this circle.

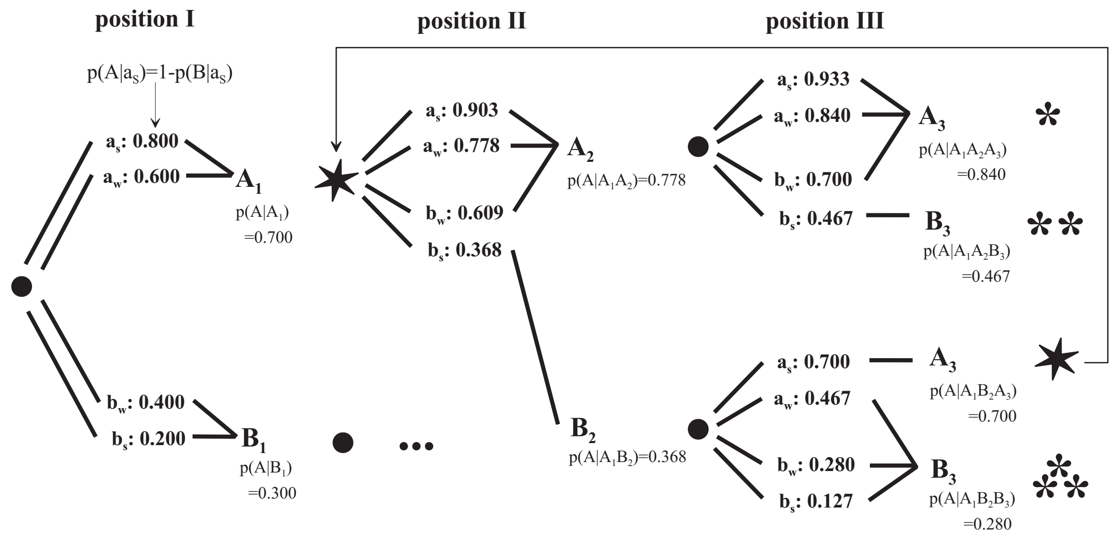


according to the own private signal at position IV. At positions V and VI no obvious heuristic can be provided. Note that an information cascade starts always after three consecutive identical predictions. In addition, public and private information lead with two exceptions at position VI to unambiguous predictions in contrast to Anderson and Holt (1997).

A potentially interesting situation arises if the a posteriori probabilities for both states are close to 50%. If participants are just slightly uncertain whether observable predictions are reliable or whether the probability of individual mistakes is greater than zero, they will put more weight on their own private information that might lead to a collapse of an information cascade. The data will show however that it is possible to distinguish between "rational" adjustments in the updating procedure and "irrational" overconfidence.

Figure 4: Some prediction paths at positions IV, V and VI

The decision situations at positions IV, V and VI are displayed in this figure depending on the private signal $i_X \in \{a_S, a_W, b_W, b_S\}$ and the decision histories based on histories $h^3_{id} = A_1A_2A_3$ (*), on $h^3 = A_1A_2B_3$ (**) or on $h^3 = A_1B_2B_3$ (***). Based on probabilities for state $A$ the decision is shown. Moreover, the posterior probabilities for an observed decision are provided. Situations, which occur identically at different positions, are marked with a special symbol (e.g. a star).

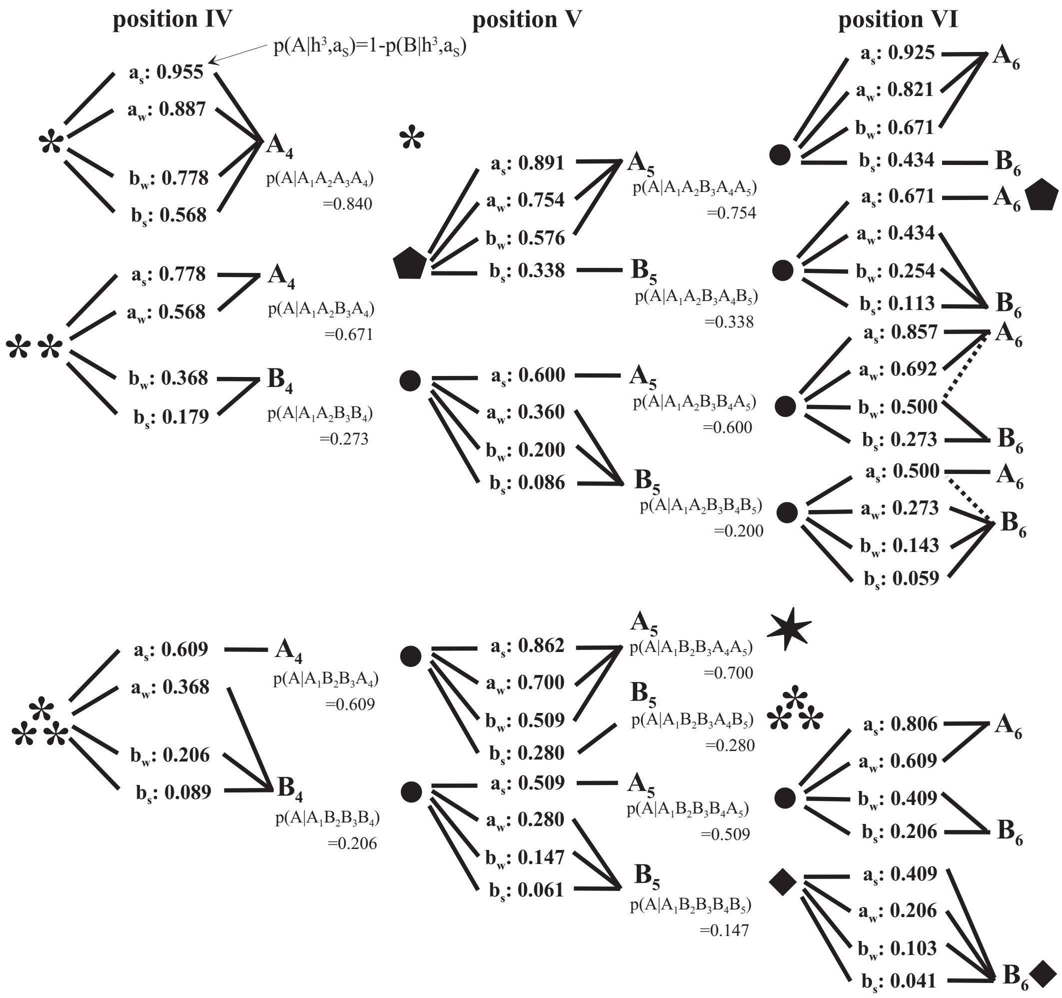

## 3 Results

Our analysis is based on the experimental data from 21 sessions with a total of 126 subjects. In 1639 rounds, they submitted 9834 individual predictions. 107 of the 126 participants made more correct predictions than they would have made based only on their private information. On average, subjects were able to predict correctly in 3.78 rounds ($\bar{\sigma} = 0.47$) in which their own signal was wrong.[10] Using the available public information, subjects predicted in 3.78 rounds against their own ex-post wrong signal. This significant improvement (t-statistic=8.1, $\alpha < 0.001$) was achieved during the whole session. Learning within the whole group is not observable since comparing the results of the first fifteen rounds with those of the last fifteen rounds does not reveal a significant difference. In addition, individuals' behavior is stable, i.e. systematic deviations from rationality do not disappear or worsen.

To understand the development of cascades it is necessary to analyze the first three predictions within each round since these decisions have a crucial influence on the results of the round. Moreover, it is easier to identify plausible reasons for deviations from rational behavior. Then, we proceed with the analysis of cascades and reverse cascades. This includes the extraction of behavioral regularities and the identification of their effect on welfare.

At position I within a round, a participant can base her decision only on her private information and on her knowledge about the information structure. It is obvious that she should predict the state indicated by her private information since even a weak signal has a higher probability than random guessing. Note that the risk attitude or beliefs about others' behavior do not influence the prediction at position I because only two states exist and the prediction is irreversible. Table 1 shows the aggregated predictions classified as "Bayesian" or "non-Bayesian" depending on the signal's strength.

Table 1: First Prediction

The first predictions of each round are shown for all 1639 rounds depending on the signal's strength (strong/weak). In addition, each decision is classified as "Bayesian" or "non-Bayesian". Since only the own private information is available at position I, the predicted state should be the one indicated by the private information. In this case, the prediction is classified as "Bayesian".

| position I | strong | | weak | | $\sum$ | |
|---|---|---|---|---|---|---|
| | obs. | in % | obs. | in % | obs. | in % |
| Bayesian | 747 | 97.0 | 746 | 85.8 | 1493 | 91.1 |
| non-Bayesian | 23 | 3.0 | 123 | 14.2 | 146 | 8.9 |
| all | 770 | | 869 | | 1639 | |

As table 1 shows, about 91% of all first predictions are made according to the first participant's signal. There are only 23 (3.0%) predictions against a strong signal, but 123 (14.2%) predictions against a weak

[10] Subjects participated on average in 78 rounds. Based on the probabilities for strong and weak signals, they received 23.4 (=30%) wrong signals.

signal. A plausible (but not rational) explanation can be found for 14 of these 23 decisions based on a strong signal: in the previous round they have predicted the ex post wrong state although this might have been rational for them. 47 of the 123 predictions against a weak signal occurred after an ex post wrong prediction in the previous round.[11] Of the remaining 74 non-Bayesian predictions, twelve (six) occurred within the first (last) ten rounds of an experimental session. Thus, there does not exist any indication that these predictions should be attributed to inexperience or to boredom confirming our result of no learning.[12]

However, gambler's fallacy can explain about half of the predictions against the own private information if the prediction in the previous round has been correct. These subjects believe that the probability for both states is changing based on the observed history of state realizations in previous rounds: subjects predict against their own signal more often if the private information indicates the state which occurred in the previous round(s) even though they have submitted a correct prediction. Suppose a subject receives the private information $a_W^1$. In addition, she has observed and has correctly predicted in the previous $n$ rounds ($n \geq 1$) state $A$. In this situation 37 of the 74 predictions against the own weak private information occurred. The same happened in five of the nine similar cases with a strong signal. In addition, nobody predicted against the own private information if this person's prediction in the previous round has been correct and the private information indicates the other state for this round. The remaining 37 predictions against the own private information at position I cannot be explained since they exhibit no regularity.

The predictions at position II are based on the observed prediction at position I and on the own private information. Moreover, the information structure is public knowledge and can be used for updating probabilities. In table 2 the predictions are classified as "Bayesian" or "non-Bayesian" depending on both, the signal's strength and on the first prediction $D_1$.

More than 97% of all predictions at position II are made based on a strong private information regardless of the first prediction. The predictions against the own strong information are the result of random errors. If the private weak information confirms the first prediction, about 91% of the participants decide to follow the own information and thus the first prediction. The remaining 9% of the predictions are submitted against the own weak private information *and* the first prediction $(D_1 = i_W^2)$. As at position I, gambler's fallacy, random errors and a reaction to the own ex post wrong prediction in the previous round explain some of these predictions. It is notable however that the number of deviations is almost three times as high as at position I with a strong signal although the probabilities are almost the same $\left(p\left(A \mid A_1 a_W^2\right) = \frac{7}{9} \text{ vs. } p\left(A \mid a_S^1\right) = \frac{4}{5}\right)$.

Although one should predict against the own weak signal that is contradicting the first prediction $(D_1 \neq i_W^2)$ based on Bayes' rule, 49.3% of all decisions follow the own signal. It is obvious that such a

[11] Two predictions against the own weak signal occurred in round 1. The median of predictions against the own weak signal at position I is five. This number varies between zero and ten except for one session, in which 18 out of 46 predictions were made against the own weak signal.

[12] The result that subjects did not learn is not too surprising given the limited information they received at the end of each round. They could only compare the prediction sequence and their own signal with the outcome but the underlying signal sequence was not revealed.

Table 2: Second prediction

In this table the second predictions of each round are displayed for all 1639 rounds depending on the second information ($i^2$), its strength (strong/weak) and on the round's first prediction ($D_1$). In addition, each decision is classified as "Bayesian" or "non-Bayesian" assuming rationality of the first decider. The decision should be based only on the second signal if this signal is strong. A weak signal implies the same decision as the first one regardless of the signal. Thus, if $D_1 \neq i_W^2$ it is rational to follow the first prediction. Results are given in percentage of column total. Rational herding, i.e. following the previous decision against the own private information, can occur only with a weak signal (*italics*). Predictions that can be caused by overconfidence are denoted in **bold**.

| position II | $D_1 = i^2$ | | $D_1 \neq i^2$ | | |
|---|---|---|---|---|---|
| | strong | weak | strong | weak | $\sum$ |
| Bayesian | 97.1 | 90.8 | 97.7 | *50.7* | 78.3 |
| non-Bayesian | 2.9 | 9.2 | 2.3 | **49.3** | 21.7 |

deviation cannot be explained using the above mentioned reasons especially since the probability for the correct state is about the same as having a weak signal at position I (60.9% vs. 60.0%). The only difference is that it requires a prediction against the own private information at position II. Obviously, participants put too much weight on their private information compared to the public information, which clearly indicates the existence of overconfidence.[13] Note that gambler's fallacy would increase the proportion of "Bayesian" predictions because agents would then predict against their own private signal.

One might argue at this point that the observed deviations are the result of a more sophisticated updating procedure, i.e. taking a certain amount of mistakes at position I into account. Both random ordering and anonymous predictions prevent conditioning the decision at position II on the identity of the person predicting at position I. As a consequence error rates include beliefs about individual error rates as well as their distribution within the group. Prediction errors with a strong signal ($\epsilon$) and errors with a weak signal ($\theta$) must be high enough such that $p\left(A \mid \tilde{A}_1\right) \leq \frac{3}{5}$ to justify a prediction of state $B$ based on information $\tilde{A}_1 b_W^2$. Thus, $p\left(A \mid \tilde{A}_1\right) = \frac{\frac{1}{2}\left[\frac{1}{2}\left[(1-\epsilon)\frac{4}{5}+(1-\theta)\frac{3}{5}\right]\right]}{\frac{1}{2}\left[\frac{1}{2}\left[(1-\epsilon)\frac{4}{5}+(1-\theta)\frac{3}{5}\right]+\frac{1}{2}\left[(1+\epsilon)\frac{1}{5}+(1+\theta)\frac{2}{5}\right]\right]} \leq \frac{3}{5}$. Using some algebra leads to the conclusion: $\epsilon \geq \frac{5}{11} - \frac{12}{11}\theta$. Figure 5 shows the area of error rate combinations in which the same prediction as at position I is rational even if the own weak signal indicates the other state. Note that an error rate of 0.5 is equivalent to assuming that *all* predictions at position I are random.

The actual average error rates (see table 1) are $\epsilon = 0.03$ and $\theta = 0.142$ for strong and weak signals at position I, respectively. Error rates vary between 0 and 0.114 for $\epsilon$ and between 0 and 0.391 for $\theta$ between sessions. error rates are only in one session almost sufficiently high enough to justify a prediction according to the own weak signal at position II. On average, error rates could have been 2.5 times higher than the actually observed ones, before it would have been rational to deviate from predicting the state suggested by Bayesian updating.

[13] Anderson and Holt (1997) report about 15% deviations in which subjects predict according to their own signal but should follow the crowd without evaluating possible reasons.

Figure 5: Error rates at position II

An agent at position II who receives a weak signal which indicates the other, not yet predicted state $\left(i_W^2 \neq D_1\right)$, should predict against her own signal as long as the anticipated error rates at position I are less than $\epsilon$ and $\theta$ after receiving a strong or weak signal, respectively. An error rate of 0.5 is equivalent to the assumption that *all* predictions with the associated strength are randomly made at position I.

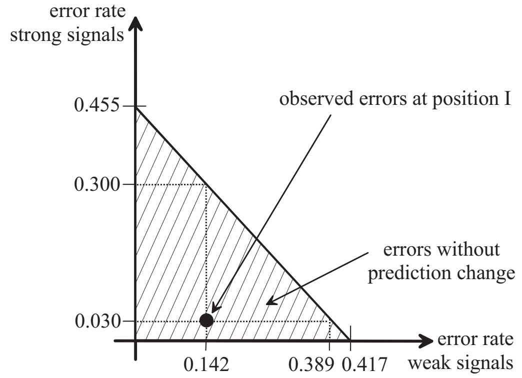


These results are a clear indication of overconfidence because agents' believe that the others made more mistakes than they actually did. An alternative explanation would be regret aversion. Regret averse people suffer an additional utility loss if they predict against their own signal and this turns out to be ex post wrong. To avoid this, agents put a higher weight on their own information than is rationally appropriate. Regret aversion and overconfidence are closely related in our experimental setting since both biases lead to overweighing of the own private information. However, only overconfidence is consistent with gambler's fallacy because the decision maker believes in her superior prediction ability even if this implies to predict against the own information. Overconfidence is also consistent with subjects' answers in the final questionnaire. Only very few subjects mentioned that they adjusted their predictions to account for others' potential errors. They simply believed in their (wrong) decision heuristics and the observed behavior is not the result of random mistakes. Other alternative explanations such as conformity and representativeness would enforce predictions according to Bayesian updating.

Summing up, potential cascades collapse relatively often at position II because subjects assign more weight to their own weak private information. Although we have collected only subjects' predictions and not their probability judgments that lead to these predictions it is possible to identify some more precise decision heuristics. In addition to the observed behavior, subjects' answers in the final questionnaire reveal that a lot of subjects followed their own signal if they had to decide at positions II without considering the first prediction. As a consequence, two contradicting predictions at the beginning of a round would contain no additional information compared to the situation without public information since the distinction between strong and weak signals at position II is lost. Moreover, this heuristic demonstrates that agents use simple heuristics which may often lead to Bayesian-like predictions but not always.

This fact has an important impact on some decisions at position III which are shown in table 3 depending on the observed history of predictions and on the private information. The classification of predictions assumes that all agents use Bayes' rule to aggregate information.

Table 3: Third prediction

In this table the third predictions of each round are displayed for all 1639 rounds depending on the signal ($i^3$), its strength (strong/weak) and on the first two decisions ($h^2$) within this round, which are either identical $\left(h_i^2 d \in \{A_1A_2, B_1B_2\}\right)$ or not $\left(h^2 \in \{A_1B_2, B_1A_2\}\right)$. In addition, each decision is classified as "Bayesian" or "non-Bayesian" assuming rationality of the first two deciders. Rational herding, i.e. following the previous decision against the own private information, can occur only with a weak signal (*italics*). Predictions, which can be caused by overconfidence, are denoted in **bold**. "Irrational" herding is marked with a star (*).

| | $h_{id}^2 \in \{A_1A_2, B_1B_2\}$ | | | | n=958 |
|---|---|---|---|---|---|
| position III | $D_2 = i^3$ | | $D_2 \neq i^3$ | | |
| | strong | weak | strong | weak | $\sum$ |
| Bayesian | 98.8 | 95.3 | 85.3 | *76.4* | 88.5 |
| non-Bayesian | 1.2 | 4.7 | 14.7* | **23.6** | 11.5 |
| | $h^2 \in \{A_1B_2, B_1A_2\}$ | | | | n=681 |
| | $D_2 = i^3$ | | $D_2 \neq i^3$ | | |
| | strong | weak | strong | weak | $\sum$ |
| Bayesian | 95.9 | 84.2 | 98.5 | *26.3* | 78.1 |
| non-Bayesian | 4.1 | 15.8 | 1.5 | **73.7** | 21.9 |

At position III it is rational to predict always according to the own strong private information. With a weak signal one should follow the immediate predecessor. Decisions, which are based on a strong signal, are almost always in line with Bayes' rule. The only notable exception can be observed if $(h_{id}^2)$ and $D_2 \neq i_S^3$ when 14.7% follow the crowd by predicting against their own private information. This "irrational" herding is consistent with conformity and assumed errors in observed predictions. In addition, it is in line with the stated heuristics of the questionnaire because the prediction $A_2$ is based in this case only on signals $a_W^2$ and $a_S^2$ which overrule the information $b_S^3$ using Bayes' rule. However, the same reasoning as well as anticipated error rates cannot explain the substantial deviation from predictions based on a weak signal in the same situation. The only explanation for this behavior is overconfidence, i.e. assigning almost no weight to the first two predictions.

If the first two predictions disagree and the own weak information contradicts the last observed prediction $\left(D_2 \neq i_W^3\right)$ about three quarters of the predictions (73.7%) follow the own signal. This evidence can be explained with the anticipation of error rates and with the stated heuristic in the questionnaire, which is based on overconfidence. A noteworthy 26.3% of Bayesian predictions indicate as in the previous situation with $h_{id}^2$ that public information is not completely ignored as suggested by the Private Information hypothesis.

In summary, predictions at position III are mostly consistent with overconfident agents who put too much weight on their own information. The attempt to "correct" for errors contained in the public information is an indication for overconfidence at position III because it assumes a degree of sophistication at position II that contradicts the assumption of errors at both positions I and II. Moreover, predictions at positions IV, V and VI confirm also that overconfidence is the reason for deviations from rational Bayesian updating.

Starting at position IV and assuming that everybody uses Bayes' rule, the own private information does not change the prediction, i.e. to follow the crowd. Only the assumption of position independent error rates without attempts to correct for these errors at earlier positions can change the information content of the publicly observable predictions enough to justify a prediction against the crowd based on a strong signal. If the agent at position IV believed that every predecessor decided only based on her own private information, there would be even more reason to follow the crowd $\left(p = \frac{343}{451}\right)$ than under the assumption of Bayes rule $\left(p = \frac{21}{37}\right)$. As a result, 503 complete cascades[14] and 139 complete reverse cascades should be observed based on three identical predictions at positions I through III. It is obvious that the overconfident behavior at positions II and III destroyed potentially complete (reverse) cascades. If the first three predictions were made using Bayes law, 723 complete cascades and 220 reverse cascades would have occurred. Table 4 shows how many (reverse) cascades survive until the end of the round. In addition, the private information, which is responsible for the collapsing cascade, is provided.

Table 4: Survival of cascades and reverse cascades

In this table the number of (reverse) cascades are shown that survived until this position. In addition, the collapse initiator's private information $i_x^y$ with $x \in \{S, W\}$ in relation to the most recent prediction $D_{y-1}$ is provided.

| | cascades | | | reverse cascades | | |
|---|---|---|---|---|---|---|
| position $y$ | IV | V | VI | IV | V | VI |
| start | 503 | 422 | 372 | 139 | 91 | 77 |
| $i_S^y = D_{y-1}$ | 2 | 0 | 2 | 0 | 0 | 0 |
| $i_W^y = D_{y-1}$ | 4 | 0 | 8 | 5 | 0 | 0 |
| $i_S^y \neq D_{y-1}$ | 65 | 40 | 35 | 38 | 13 | 16 |
| $i_W^y \neq D_{y-1}$ | 10 | 10 | 9 | 5 | 1 | 1 |
| end | 422 | 372 | 318 | 91 | 77 | 59 |

Of the 503 complete cascades, which should occur after position III with there identical predictions, only 318 (-36.8%) are actually completed at the end of a round. About 75% of the cascades collapse due to a strong private signal that indicates the opposite state. These collapses result in a welfare loss especially if the remaining participants follow this prediction. Reverse cascades collapse relatively more often (-57.6%). Their number is reduced from a potential of 139 after position III to 59 completed reverse cascades at the end of a round. The higher collapse rate for reverse cascades is not surprising due to the information structure. Since all private information depends on the realized state it is more likely that strong and weak signals indicate the correct state. As a direct consequence, the likelihood of a strong

[14] A (reverse) cascade is complete if all six subjects predict the same state and this prediction is ex post correct (wrong).

signal that contradicts the developing reverse cascade is higher than the likelihood for a contradicting strong signal within a cascade. Together with overconfidence the result is explained.

Collapsing reverse cascades increase welfare, i.e. overconfidence can be beneficial. But since the absolute number of collapsed cascades (185) is larger than that of the collapsed reverse cascades (80), the overall effect of overconfidence on information aggregation is negative. Participants were obviously scared by the prospect of encountering a reverse cascade and therefore tried to avoid it although this was costly. After they had received their payment, some participants were asked to guess how often complete cascades occurred in relation to reverse cascades. The most common answer was "close to 1:1" although the relation was more than 5:1 as the results in table 4 show.

The analysis of the prediction behavior in potential (reverse) cascades is the next step. Looking only at those cases in which an unanimous prediction history exists has one advantage: The stable BR benchmark allows to identify hints about the updating procedure. Table 5 provides the percentage of conforming predictions within potentially complete (reverse) cascades.

Table 5: Confirming predictions within complete (reverse) cascades

The percentage of conforming predictions after observing an unanimous history ($h^y_{id}$) until position $y$ is displayed dependent on the signal's strength ($W, S$), on the private information $i$ and on whether a cascade (C) or a reverse cascade (RC) is developing. Following previous predictions ($D_{y-1}$) is rational except in those cases marked with a star (*).

| | predictions (in %) after observing $h^y_{id}$ | | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|---|
| position | II | | III | | IV | | V | | VI | |
| | C | RC | C | RC | C | RC | C | RC | C | RC |
| $i^y_S = D_{y-1}$ | 98 | 94 | 99 | 100 | 99 | 100 | 100 | 100 | 98 | 100 |
| $i^y_W = D_{y-1}$ | 90 | 92 | 96 | 94 | 96 | 89 | 100 | 100 | 87 | 95 |
| $i^y_W \neq D_{y-1}$ | 51 | 50 | 80 | 72 | 91 | 79 | 90 | 96 | 93 | 95 |
| $i^y_S \neq D_{y-1}$ | 3* | 2* | 18* | 11* | 35 | 39 | 50 | 46 | 61 | 47 |

As expected, differences in prediction behavior in cascades and reverse cascades do not exist since the agents do not know in which cascade situation they are. Deviation from the rational prediction is almost non-existent starting at position III if the own private signal confirms the observed previous predictions.[15] Agents with a weak contradicting signal deviate in more than 10% of all cases until position V. The increasing percentages demonstrate that agents do not ignore public information completely but they put more weight on their own information than it is rational. The same pattern can be observed looking at predictions based on strong contradicting private information. From positions IV to VI the prediction percentage increases by about ten percentage points with each confirming prediction.

As mentioned before, learning did not occur in this experiment since the prediction pattern did not change within a session. The decision time (excluding the random delays) decreases significantly if the

[15]The drop at position VI cannot be explained. 52 out of 60 possible cascades (=87%) occurred. The remaining eight cascades collapsed in six different sessions.

first ten rounds are compared with the last ten rounds. This decrease is due to subjects' experience because there is no evidence that previous outcomes have an influence on the decision time even after experiencing a reverse cascade. Thus, we can conclude that subjects used their decision heuristics and they did not modify them systematically during the session. In some sense this confirms the notion of overconfidence because subjects were (over-)confident predicting optimally.

We have established that agents are on average overconfident. Now, two questions remain. First, the consequences for the information aggregation process must be quantified. Second, the performance of agents' stated heuristics will be evaluated. To answer the first questions, we compare the observed data with our two benchmarks, Bayes' Rule (BR) and Private Information (PI). Within both scenarios it is assumed that all agents use the same decision rules, i.e. under PI everybody uses only her private information and disregards publicly observable predictions in all situations completely. For the second question, we generate a third benchmark, Heuristic Updating (HU). This updating procedure is derived from a combination of heuristics which subjects provided after the experiment:

- Predict according to your own **strong** signal if you are at position I, II, III or IV. In addition, use it at positions V and VI if more than one deviation is observable. Otherwise, follow the majority.
- Predict according to your own **weak** signal if your are at position I or II. In addition, use the own signal only if no majority exists and the last two subjects have not predicted the same state. Otherwise, follow the majority.

The heuristic is consistent with the notion of overconfidence since it puts more weight on the private information than on the public information which is the major difference to BR. Therefore, the decisions reveal the basic information better or more obviously than under BR. But, the information quality decreases because the distinction between strong and weak signals is no longer possible in some situations. We calculate $\frac{X-PI}{BR-PI}$ with $X \in \{BR, PI, HU\}$ as a measure for efficiency. Observed predictions lead to an efficiency of 62.8% whereas using exclusively the heuristic increases efficiency significantly to 88.9%. In other words, agents would have earned more if they had used their own heuristics. This heuristic is obviously a reasonable response to others' behavior as long as everybody is not completely discarding public information. It is more robust than Bayes' rule because it is easy and incorporates the possibility of others' errors. Moreover, it avoids most of the "painful" reverse cascades at an efficiency loss of about 10% and it explains why learning does not occur. In four (of 21) sessions subjects were not able to predict better that PI, i.e. the earning would have been higher even with only predicting according to the own information.

## 4 Conclusion

The purpose of this experiment was to study information aggregation with two different qualities of information and to identify how the individual updating process influences the aggregation process. Aggregation was observed since almost all participants predicted better than based only on their own private information. Agents' overconfidence provides the only consistent explanation for the observed deviations from Bayes' rule. Other explanations, such as advanced error correction, regret aversion and gambler's fallacy are inconsistent with the data. Overconfident prediction behavior led to fewer than expected cascades and reverse cascades. Although individual behavior reduced relatively more reverse cascades than correct cascades the (absolute) effect on welfare was significantly negative. The collapse of information cascades initiated sometimes new cascades.

Based on this experiment several extensions will provide further insights about how information is aggregated in groups. Eventually, these will then lead to market situations in which prices might provide additional information about the precision of private information. The next step to evaluate the updating procedure is to extract probability judgments right before participants submit their predictions. Another modification of this baseline experiment is the choice whether participants want to buy private information for a fixed cost. This will answer the question whether participants can distinguish between informative and uninformative decisions in a rather simple environment. A crucial feature of markets is the possibility to decide at which time one would like to take an action. An endogenous timing decision can have two effects on the aggregation process. On the one hand it can improve aggregation especially if participants with higher quality of information have an incentive (e.g. to avoid a waiting cost) to move earlier than those with weak signals who gain more by observing public information. But on the other hand overconfidence can lead to situations in which agents move too fast based on their private information and thus create misleading public signals for the others. Finally, our simple setting can be extended by a pricing mechanism and by allowing simultaneous or repeated decisions.

# Appendix

## Sequential Information Processing Experiment
## Instructions

Thank you for your participation in this experiment of economic decision making. The money for your payment has been provided by the Deutsche Forschungsgemeinschaft. This session will probably last about two hours. Please follow these instructions very carefully, in order to earn as much money as possible. You can always ask questions until the end of the test rounds.

### Information structure and course of a round

In this experiment you shall predict the occurring state in each round based on your private signal and the existing public information. The ordering of the six participants is determined randomly in each round.

Two states, marked “A” (white ball) and “B” (black ball), can occur. The state is being determined by random draw from an urn, which contains ten “A”-balls and ten “B”-balls, i.e. both states occur with the same probability $\left(p = \frac{1}{2}\right)$.

If state $A$ occurred, the private signal will be determined for each participant as follows:

First the strength of the signal has to be determined by draw from an urn, which contains ten “strong” and ten “weak” signals, i.e. the possibility of the signal being strong (S) or weak (W) is equal $\left(p = \frac{1}{2}\right)$ (see left big urn).

The signal is now being determined, dependant on its strength, by a draw from another urn:

- The ”strong” urn contains four “A”-signals and only one “B”-signal (small urn, top-left).
  ⇒ The ratio of “A”- and “B”-signals is 4:1.
- The weak urn contains three “A”-Signals and two “B”-signals (small urn, lower-left).
  ⇒ The ratio of “A”- and “B”-signals is 3:2.

The following figure illustrates the procedure:

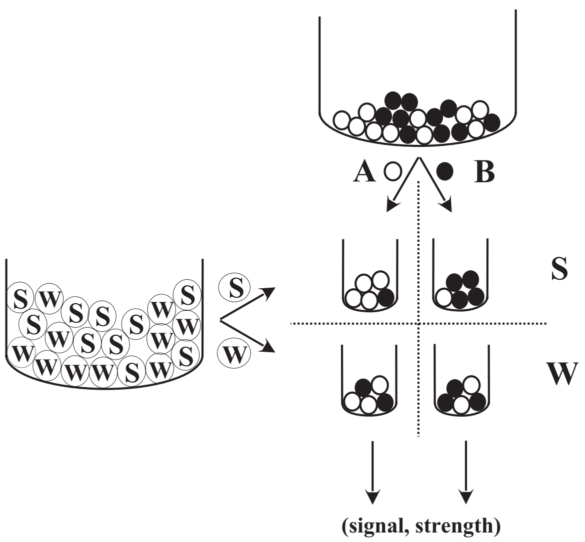

First of all, the computer determines the order, in which the predictions have to be submitted. When it is your turn, you first see your private signal, as well as the accompanying strength of signal. Then you are asked to submit your prediction. Submitted predictions are public information, i.e. the following participants can observe the predictions of all predecessors in addition to their own signal (down on the monitor). However, they cannot infer neither the underlying signal nor the accompanying signal's strength. The identification of the participants is not possible either. Your position within a round is display as a red number.

**Attention:** An additional information cannot be inferred from the reaction time of the acting participant since the computer enforces a random delay of at least three and not more than seven seconds before passing on the private signal.

As soon as all six participants have made their decision, the occurred state will be announced and a further round (with new information) begins.

### Test Rounds

Before you will earn money with your predictions, you will become better acquainted with the procedures in three unpaid test rounds. During these test rounds you can always ask questions about the information structure and the course of the experiment.

### Payment

You will participate in at least 25 and at most 100 rounds, in which you will be paid according to the correctness of your predictions. For each correct prediction you will receive 300 currency units (cu) , for each wrong prediction only 100 cu. At the end of the experiment the total payoff for all six participants will be converted in Deutsche Mark (DM) according to the expected hourly earnings of 16 DM. With the resulting exchange rate for this session your earnings will be converted in DM (and rounded up to the next DM).

**Example:**

- You have submitted 27 correct and 8 wrong predictions in 35 rounds: 8900 cu.
- All six participants have earned with their predictions: 43200 cu.
- The experiment (instructions and test rounds included) has lasted 2 hours.

Consequently, the exchange rate is computed as $\frac{43200cu}{16\frac{DM}{h}*6*2h} = 225\frac{cu}{DM}$.

As a result you earned 39.56 DM and you will receive 40.00 DM.

If you have any questions, now or during the test rounds, you can ask them in the next three minutes as well as during the three test rounds.

## Final questionnaire

This questionnaire can help us to understand your decisions better and to generate new experiment ideas. The more precisely you formulate your statements, the better we can use them.

1. Which decision rule (or heuristic) have you used to make your predictions?
2. Has your behavior changed during the experiment? If so, why?

3. How strong, depending on the decision time (first position, second position, etc.), have you weighted your signal compared to the decisions that were already public/known?
4. Would you like to decide again at the end of a period? If applicable, how often and why would you predict against your own information?
5. What would you do, if you could decide when to submit your prediction, instead of doing this in a predetermined order?
6. How would you change your behavior, if you lose money by waiting for a longer time?

## References


Allen, F. H., S. Morris, and A. Postlewaite, 1993, "Finite Bubbles with Short Sale Constraints and Asymmetric Information," *Journal of Economic Theory*, 61, 206–229.

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

Avery, C., and P. Zemsky, 1998, "Multidimensional Uncertainty and Herd Behavior in Financial Markets," *American Economic Review*, 88, 724–748.

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.

Bikhchandani, S., D. Hirshleifer, and I. Welch, 1998, "Learning from the Behavior of Others: Conformity, Fads, and Informational Cascades," *Journal of Economic Perspectives*, 12, 151–170.

Camerer, C. F., 1989, "Bubbles and Fads in Asset Prices," *Journal of Economic Surveys*, 3, 3–41.

Camerer, C. F., 1995, "Individual Decision Making," in John H. Kagel, and Alvin E. Roth (ed.), *The Handbook of Experimental Economics* . pp. 587–703, Princeton University Press, Princeton, NJ, 1 edn.

Camerer, C. F., and K. Weigelt, 1991, "Information Mirages in Experimental Asset Markets," *Journal of Business*, 64, 463–493.

Daniel, K., D. Hirshleifer, and A. Subrahmanyam, 1998, "A Theory of Overconfidence, Self-Attribution, and Security Market Under- and Overreaction," *Journal of Finance*, 53, 1839–1885.

De Long, J. B., A. Shleifer, L. H. Summers, and R. J. Waldmann, 1990a, "Noise Trader Risk in Financial Markets," *Journal of Political Economy*, 98, 703–738.

De Long, J. B., A. Shleifer, L. H. Summers, and R. J. Waldmann, 1990b, "Positive Feedback Investment Strategies and Destabilizing Rational Speculation," *Journal of Finance*, 45, 375–395.

Dow, J., and G. Gorton, 1994, "Arbitrage Chains," *Journal of Finance*, 49, 819–849.

Gervais, S., and T. Odean, 1998, "Learning to be Overconfident," mimeo.

Gul, F., and R. Lundholm, 1995, "Endogenous Timing and the Clustering of Agents' Decisions," *Journal of Political Economy*, 103, 1039–1066.

Huang, A. A., and C. R. Plott, 1999, "Information Cascades: Replication and an Extension to Majority Rule and Conformity Rewarding Institutions," forthcoming: *American Economic Review*.

Lee, I. H., 1993, "On the Convergence of Informational Cascades," *Journal of Economic Theory*, 61, 395–411.

Lee, I. H., 1998, "Market Crashes and Informational Avalanches," *Review of Economic Studies*, 65, 741–759.

Porter, D. P., and V. L. Smith, 1994, "Stock Market Bubbles in the Laboratory," *Applied Mathematical Finance*, 1, 111–127.

Scharfstein, D. S., and J. C. Stein, 1990, "Herd Behavior and Investment," *American Economic Review*, 80, 465–479.

Shleifer, A., and R. W. Vishny, 1997, "The Limits of Arbitrage," *Journal of Finance*, 52, 35–55.

Smith, V. L., G. L. Suchanek, and A. W. Williams, 1988, "Bubbles, Crashes, and Endogenous Expectations in Experimental Spot Asset Markets," *Econometrica*, 56, 1119–1151.

Welch, I., 1992, "Sequential Sales, Learning and Cascades," *Journal of Finance*, 47, 695–732.

Zhang, J., 1997, "Strategic Delay and the Onset of Investment Cascades," *RAND Journal of Economics*, 28, 188–205.