Limited depth of reasoning and failure of cascade formation in the laboratory

Dorothea Kübler
Humboldt University Berlin[1]

Georg Weizsäcker
Harvard University[2]

September 12, 2000

[1] Institut f. Wirtschaftstheorie I, Spandauer Str.1, 10178 Berlin, Germany. Email address: kuebler@wiwi.hu-berlin.de

[2] Harvard Business School, Sherman Hall, Boston, MA 02163, USA. Email address: weizsack@fas.harvard.edu

Abstract

We examine the robustness of information cascades in laboratory experiments. Apart from the situation in which each player can obtain a signal for free (as in the experiment by Anderson and Holt, 1997, American Economic Review), the case of costly signals is studied where players decide whether to obtain private information or not, at a small but positive cost. In the equilibrium of this game, only the ...rst player buys a signal and chooses an urn based on this information whereas all following players do not buy a signal and herd behind the ...rst player. The experimental results show that too many signals are bought and the equilibrium prediction performs poorly. To explain these observations, the depth of the subjects' reasoning process is estimated, using a statistical error-rate model. Allowing for di¤erent error rates on di¤erent levels of reasoning, we ...nd that the subjects' inferences become signi...cantly more noisy on higher levels of the thought process, and that only very short chains of reasoning are applied by the subjects.

1 Introduction

In simple cascade games, the players sequentially choose one out of two alternatives, after receiving private signals about the pro…tability of the two options, and after observing the choices of all preceding players. While the signals are not revealed to subsequent players, the latter may be able to infer the information observed by their predecessors from the decisions that were made. As a consequence, Bayesian Nash Equilibrium implies the possibility (depending on the sequence of signals) that rational herding occurs, i.e. that players disregard their own private information and follow the decisions of previous players. In this case, no further information is revealed, and an "information cascade" develops, with all players choosing the same option. For example, if the decision problem is whether or not to invest in a new technology, such herding behavior may create fads, where many potential investors decide to invest in the technology without much further pondering over their private pieces of information.[1]

From a behavioral perspective, one can ask whether such a reasoning process would be applied by actual decision makers. This appears particularly doubtful in situations where relatively deep levels of reasoning are needed, by which we mean that decisions are determined after several steps of using the knowledge about the knowledge ... about the others' rationality. However, most of the existing experimental tests of cascade games seem to support the theoretical predictions. Anderson and Holt (1997) report that in cases where a player should, in equilibrium, disregard her own signal, most subjects do so and indeed follow the others' decisions; a result which has since been replicated by Hung and Plott (1999) and other researchers.[2]

We modify the experimental design by Anderson and Holt (1997) by introducing a separate stage for each player, at which she is asked whether or not she wants to receive a signal, at a small but positive cost. This modi…ed game can be viewed as a "hard" test for Bayesian rationality, in the sense that the equilibrium prediction is much more extreme: In Bayesian Nash equilibrium, the …rst player buys a signal and all other players blindly follow the …rst player's decision, independently of the chance moves which determine the signals: After the …rst player's choice, no further signals are bought, and cascades occur with certainty. (In the example of technology

[1] Following Banerjee (1992) and Bikhchandani, Hirshleifer, and Welch (1992), the economic literature on rational herding has grown considerably over the last decade. For a survey see Bikhchandani, Hirshleifer, and Welch (1998).

[2] For extensions and discussion see also the experimental papers by Nöth and Weber (1999) and Huck and Oechssler (2000) .

adoption given above, all except one agent will avoid incurring costs of information acquisition, even if they are arbitrarily small, and the technology will be uniformly accepted or rejected.)

The experimental results are not in line with these predictions. While not all of the subjects acting as …rst players buy a signal, the signal acquisition in later stages is excessive, and in sum far too many signals are bought. Cascades often do not form at all, or are fragile in that after a cascade has started subjects buy signals and choose in opposition to their predecessors' decisions. The predictive value of Bayesian Nash Equilibrium is much lower in the modi…ed game than it is in a control treatment without the cost, which is comparable to Anderson and Holt's design.

A natural candidate to explain this signal acquisition behavior are errors. In this more complex game, subjects may simply err when making their decisions, given their updated beliefs. A complementary and perhaps more convincing explanation goes one step further in the reasoning process: Subjects may not trust their predecessors to reveal their information as prescribed in equilibrium (e.g. because of errors), and hence prefer to buy signals themselves. According to this hypothesis, it would help the subjects to know whether or not their predecessors bought signals. We tested this possibility by including another treatment, the "high information treatment", in which subjects were given the information who of the previous subjects had obtained a signal. It turns out, however, that even more signals are bought under this treatment, and the prediction of Bayesian Nash Equilibrium – which is identical under both treatments – performs even worse.[3]

To explain these observations, we conduct a depth-of-reasoning analysis. I.e., we employ a statistical model which takes all levels of thinking about thinking ... about others' behavior into account, and allows us to make inferences about the subjects' updated beliefs after observing a given choice history. Estimating parameters which capture the error rates on all levels of the reasoning process, we are able to disentangle the di¤erent "anomalies" that can arise in long chains of reasoning, and obtain an estimate of the actual depth of reasoning in the subject pool. The model, which is based on the Agent Quantal Response Equilibrium by McKelvey and Palfrey (1998), also allows us to apply straightforward statistical tests to answer a number of questions concerning the reasoning process.

Depth-of-reasoning analyses have been conducted by several experimentalists (see the work by Nagel, 1995, Sefton and Yavaş, 1996, and Ho,

[3] A similar treatment has been run independently by Kraemer, Nöth, and Weber, 2000, with similar results.

Camerer, and Weigelt, 1998), but they all investigate normal-form game play.[4] We argue that cascade games are especially well suited for an analysis of depth of reasoning, and particularly so because they are extensive-form games: First, subjects do not face problems of calculating a fixed point or limit point in the strategy space, which typically arises in (behavioral) models of normal-form game play. Second, the extensive structure clearly defines the chains of reasoning that a player has to go through. Third, the cascade games under investigation are relatively long (six players), implying that with enough data we are able to obtain a complete picture over the full length of the reasoning process (under the assumptions of the statistical model). Fourth and finally, cascade games have the property that although they are games in extensive form, no backward-induction reasoning is involved when thinking about other players' decisions. Thus, the results do not depend on the subjects' ability to solve a game backwards, which is often doubted.

The estimation results suggest that the subjects' depth of reasoning is very limited, and that the reasoning gets more and more imprecise on higher levels: Subjects attribute a significantly higher error rate to their opponents as compared to their own, and this imbalance gets more extreme when considering the responses on the next level, i.e. when they think about the error rate that others, in turn, attribute to their opponents. More strikingly, the reasoning process ends after these two steps, although several more steps would be possible and profit-increasing in the games.

The subjects' signal acquisition behavior can be explained along the lines of these estimation results. In the treatment with cost, subjects do not trust their opponents' decisions and excessively buy signals if there are only few preceding players. With more predecessors, they tend to follow the others more, as they expect that several of these predecessors may have made an informed decision. However, they do not reason far enough to realize that other subjects also sometimes rely on third players' decisions. Therefore, in later stages of the games, they behave as if many of the preceding players made an informed decision, regardless of the history. In the high information treatment, where they learn about the signal acquisition of their predecessors, they are often surprised how little informed the others were, and hence tend to buy even more signals.

The next section contains the experimental design and procedures. Sec-

[4] Relatedly, Stahl and Wilson (1994, 1995), Costa-Gomes, Crawford, and Broseta (2000), Goeree and Holt (2000), and Weizsäcker (2000) all estimate models of normal-form game play behavior which allow for a limited depth of reasoning.

tion 3 presents the results of the different treatments in summary statistics, and Section 4 the statistical depth-of-reasoning analysis. Section 5 concludes.

2 Experimental design and procedure

2.1 Experimental design

This section contains a basic description of the four experimental treatments. We start by presenting the main treatments, Games HC and LC ("high cost" and "low cost", respectively), which involve a cost of obtaining a signal, but are otherwise almost identical to the baseline experiment conducted by Anderson and Holt (1997).

Game HC/LC:

- Nature draws one of two possible states of nature, $\omega \in \{A, B\}$, with commonly known probability $\frac{1}{2}$. Nature's draw is not disclosed to the players. Each state of nature represents an urn, where urn A contains two balls labelled a and one ball labelled b, and urn B contains two balls labelled b and one ball labelled a.
- 6 players play in an exogenously given order, as follows: In stage $n; n = 1; ...; 6$; the nth player

  1. observes the $(n - 1)$ urn choices made by the previous players,
  2. decides whether or not to obtain a private draw from the urn ω (a signal, with possible realizations $s_n \in \{a; b\}$), at a cost K, where K equals \$1.50 in Game HC and \$0.50 in Game LC, and
  3. chooses one of two possible urns, A or B.

  If the player's urn choice coincides with the true urn ω, she gets a fixed prize of U = \$12, and nothing otherwise.
- After all decisions are made, ω is announced and payoffs are realized.

We restrict attention to the case of signals being not too expensive relative to the possible prize U that subjects receive if they correctly predict the urn. More specifically, both \$1.50 and \$0.50 are below $\frac{1}{6}$ of \$12. Under this condition, the prediction of any Perfect Bayesian Nash equilibrium of the game is for the first player to obtain a signal, and for all subsequent players not to buy a signal and simply to follow the first player's choice. To see this, notice that the second player, knowing that the first player obtained a signal, cannot do better than following the first player's action, even if she obtains the opposite signal herself. Therefore, it is optimal for her not to buy a signal and to follow the first player blindly. The same logic applies to all subsequent players.

As the equilibrium prediction critically hinges on the players relying on the first player to optimally have obtained a signal, one can ask whether the specific uncertainty about previous signal acquisitions, which is not present in the baseline game by Anderson and Holt (1997), causes deviations from equilibrium play in the experiment. In order to examine this hypothesis, we also conducted a high information treatment, Game HCHI.

Game HCHI:

All stages are as in Game HC, except that the nth player, before making her own decisions, also observes whether or not each of the previous $(n - 1)$ players obtained a signal.

With the additional information given in Game HCHI, the equilibrium prediction remains unchanged, as compared to Games HC and LC: In equilibrium the players know each other's actions in Games HC and LC, so no new information is revealed. But the subjects' possible uncertainty about whether or not previous subjects made an informed decision is removed. Hence, if this uncertainty alone drives non-equilibrium behavior in Games HC and LC, deviations should be reduced in Game HCHI.

Finally, a control treatment was conducted without the cost, as in Anderson and Holt's (1997) experiment:

Game NC:

All stages are as in Game HC/LC, except that players can obtain signals for free, i.e. $K = 0$.

In contrast to Anderson and Holt's design, where players receive their signal automatically, Game NC includes a stage for each player where she is explicitly asked whether she wants to obtain a signal. This modification was introduced in order to make Game NC comparable to the other treatments: The structure of the games is the same, and the instructions could be held essentially identical.[5]

In any Perfect Bayesian Nash equilibrium of Game NC (of which there are a multitude, depending on how subjects break ties if indifferent between

[5] The instructions are given in the Appendix.

their possible decisions), cascades occur with positive probability: If, for example, the third player receives a private signal a, but the two preceding players both chose B, almost all equilibria would prescribe for her to disregard her own signal and also choose B.[6] Assuming a speci…c tie-rule, one can then observe how many of the subjects' choices are consistent with the equilibrium path prescribed by the corresponding equilibrium.[7] Importantly, the equilibrium prediction here is di¤erent from games HC, LC, and HCHI, as more signals are obtained.

2.2 Experimental procedure

The experiment was run in the Computer Lab for Experimental Research at Harvard Business School in four sessions between March and June 2000, using the software z-Tree. At the beginning of each session, two draws from physical urns were made as a demonstration. After that, all obtained signals were displayed on the subjects' computer screens. The subjects in each session were anonymously divided into groups of six players who stayed together over the entire session and played the games with player roles randomly changing after each round. In sessions 1 and 2, the subjects played Games HC, NC, and HCHI, and in sessions 3 and 4, subjects only played Game LC. To ensure at least partially that di¤erences in behavior between games are not due to learning or other e¤ects that arise because of the order in which the games are played, we switched the order of Games HC and NC. Half of the subjects in sessions 1 and 2 faced the order [NC, HC, HCHI], the other half [HC, NC, HCHI]. Within each session, half of the subjects played according to each order to control for session e¤ects. Game HCHI was always played at the end of the session because if a subject plays HCHI …rst,

[6] This is not true if the equilibrium prescribes for the second player to always follow the …rst player, regardless of his (the second player's) signal. If, however, the equilibrium tie-rule involves any positive probability for the second player to follow his own signal if it contradicts the …rst player's decision, then two preceding B's are su¢ cient for the third player to disregard her own a signal. (Therefore, the tie-rule matters for the equilibrium prediction of this game.)

[7] Anderson and Holt (1997) consider the tie-rule "Follow your own signal if indi¤erent". To simplify the analysis, we will restrict attention to a corresponding tie-rule for Game NC: "If indi¤erent concerning the urn choice, follow your own signal if you observed one, and randomize otherwise. Concerning the signal acquisition decision, always obtain a signal unless it is strictly optimal to follow the previous players' choices regardless of the signal, in which case you randomize between obtaining a signal and not." Consideration of other tie-rules would not change the equilibrium predictions in most cases, although in some cases it would (cf. Footnote 6). Notice that in Games HC, LC, and HCHI, the equilibrium-path prediction does not rely on a speci…c tie-rule.

she may transfer the information on how many subjects bought signals to the other games, which could distort the results. Each game was played for 15 rounds in a row, with one unpaid practice round immediately preceding each game.[8]

Overall, 66 subjects (mostly undergraduate students from universities in the Boston area) participated in the experiment: 24 in session 1, 12 in session 2, 18 in session 3, and 12 in session 4. Given the number of rounds chosen, this implies that games HC, LC, and HCHI were played 90 times each and game LC was played 75 times, yielding a total of 4140 decisions (12 per round). At the end of the experiment, three payo¤-relevant rounds in sessions 1 and 2 (one per treatment) and one payo¤-relevant round in sessions 3 and 4 were randomly determined by drawing from a stack of 15 numbered cards. The earnings from these rounds were added to a show-up fee of $16.[9] The subjects were identi…ed by code numbers only and received their total earnings in cash directly after the experiment.

3 Results: Summary statistics

Figures 1 through 4 summarize how well Bayesian Nash Equilibrium predicts the behavior in the four experimental treatments. In each of the …gures, the …rst column reports the relative frequency of subjects making the signal acquisition decision prescribed on the equilibrium path, at each of the six stages. Likewise, the second column shows the frequency of subjects following the equilibrium-path urn decision, at each stage. It is important to notice that these frequencies are not conditioned on the histories of the games at the respective stages – histories which in many cases are inconsistent with equilibrium play of the previous subjects.

The third column, in contrast, summarizes in how many of the rounds all decisions up to the respective stage are on the equilibrium path, including both signal acquisitions and urn choices. In the course of the six stages, this proportion decreases quite dramatically, particularly in the treatments with

[8] In session 3, the subjects played Game LC for another 15 rounds, which had not been announced to them before. To increase comparability between the di¤erent treatments, we decided not to include these data in the analysis and only used those from the …rst 15 rounds. Subjects in sessions 1 and 2 were, likewise, not told what would happen after the …rst (and second) set of 15 rounds.

[9] The relatively high show-up fee was chosen in order to increase average earnings, which became necessary because sessions 1 and 2 lasted for about two and a half hours. At the beginning of the experiment, only $8 was announced as a show-up fee. The additional payment of $8 was announced in sessions 1 and 2 right before game HCHI started. In sessions 3 and 4 it was announced at the end of the experiment.

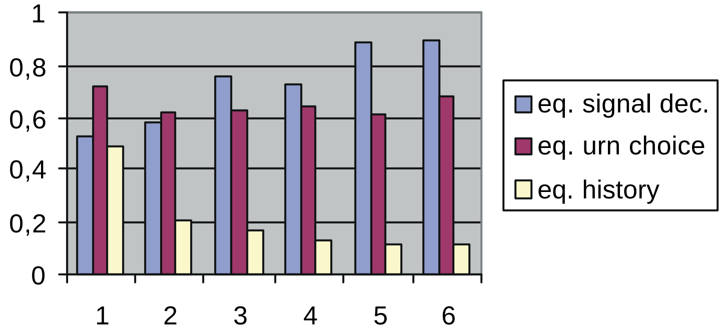

Figure 1: Decisions consistent with Perfect Bayesian Nash Equilibrium in Game HC.

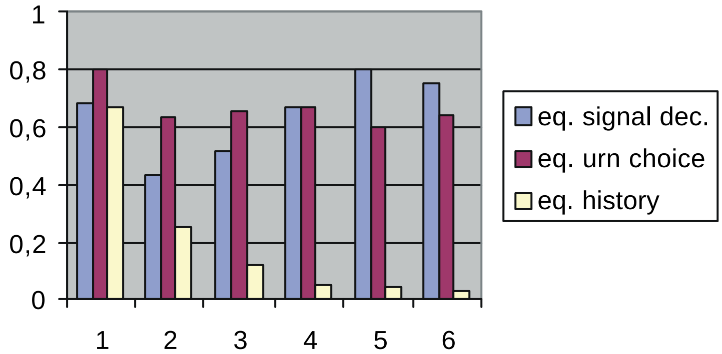

Figure 2: Decisions consistent with Perfect Bayesian Nash Equilibrium in Game LC.

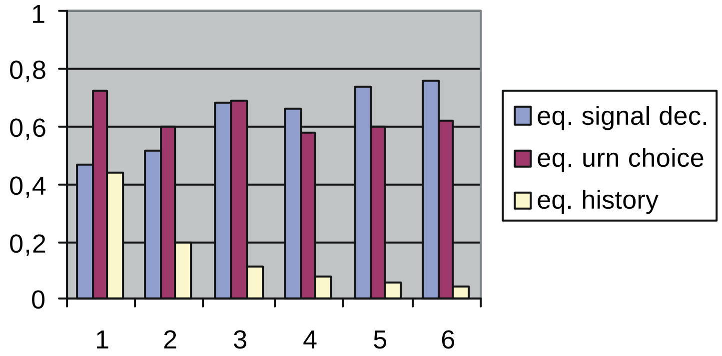

Figure 3: Decisions consistent with Perfect Bayesian Nash Equilibrium in Game HCHI.

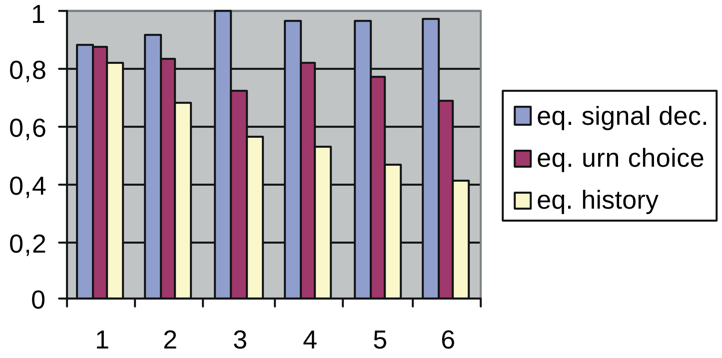

Figure 4: Decisions consistent with Perfect Bayesian Nash Equilibrium in Game NC.

a signal cost. In treatments HC, LC, and HCHI respectively, only 12%, 3%, and 4% of all games display equilibrium behavior of all six participants. In the NC-treatment, all six players make equilibrium decisions in 41% of the 90 rounds. Hence, with a positive cost the equilibrium prediction performs much worse, according to these aggregate numbers.

Next we ask whether the failure of the equilibrium prediction in the cost treatments must be attributed to subjects taking the wrong signal acquisition decision or to out-of-equilibrium urn choices. Table 1 summarizes the number of signal acquisition decisions consistent with the equilibrium prediction (…rst row) and the number of cases in which both decisions follow the equilibrium path (second row), each divided by the total number of decisions after an equilibrium history or after what could be an equilibrium history.[10] Notice that in treatment HCHI the proportions of observed equilibrium signal decisions and of observed equilibrium signal and urn decisions are the lowest. Thus, providing the subjects with information about who of the preceding players saw a signal does not lead to more decisions consistent with equilibrium behavior. In the high-cost treatment HC, signal decisions and urn choices follow the equilibrium path slightly more often than in the low-cost treatment LC. By far the highest number of equilibrium decisions occurs in the no-cost treatment NC where only 5% of all signal decisions and 13% of combined signal and urn decisions are not in line with equilibrium play.

Yet, the equilibrium prediction for Game NC di¤ers from the other three games in that herding should occur less often. So it may still be the case that rational herding occurs as rarely in Game NC as it does in the treatments with cost, relative to the number of equilibrium herding decisions. However, this is not true: The third row reports in how many situations after an equilibrium-path history subjects either decided not to see their signal at all or to disregard their own signal "correctly", divided by the number of situations in which such herding is prescribed in equilibrium. Again, the NC treatment is closest to equilibrium play. In 84% of all potential herding decisions (on an equilibrium path) subjects actually disregard their signal or do not wish to see it and follow their predecessors.[11] The cost treatments are

[10] Since it is not clear what can be assumed about o¤-equilibrium beliefs, for the numbers reported in Table 1 we only consider decisions following a history that could be part of an equilibrium. Histories can still be included after certain out-of-equilibrium decisions, as long as the latter do not lead to an observable history that cannot be part of an equilibrium. E.g., in any of the games, if the …rst player observes a signal a but chooses urn B, the second player's decision would still be included in the analysis.

[11] In the corresponding treatment by Anderson and Holt (1997), 70% of the subjects

ordered such that high costs and no information lead to more equilibrium herding decisions than either low costs or high information.[12]

Data:	HC	LC	HCHI	NC
eq. signal dec.	272/375=0.73	215/316=0.68	104/166=0.63	475/500=0.95
eq. signal & urn dec.	240/375=0.64	196/316=0.62	87/166=0.52	433/500=0.87
eq. herding dec.	196/285=0.69	146/241=0.61	46/76=0.61	80/93=0.84

Table 1: Frequencies of decisions as predicted in equilibrium, given that the observed history could be part of an equilibrium.

While the equilibrium is often socially ine¢cient (in cases of false cascades where all players choose the wrong urn), in none of the four treatments did the observed deviations from equilibrium play increase the overall e¢ciency. Expressed in percentages of the total payments that would have been received in equilbrium, total earnings in the four treatments were: 81.3% in HC, 96.8% in LC, 78.6% in HCHI, and 91.4% in NC.[13]

Finally, consider Figures 1 to 4 again, and in particular the columns for the …rst stages of the games. As these columns never reach 1, the …rst players deviate from equilibrium play, either when deciding whether to buy a signal or when choosing the urn or both. In particular, only 53% of all players in the …rst round in treatment HC decide to see their signal, 68% in treatment LC, 47% in treatment HCHI, and 89% in treatment NC. Since there is no uncertainty about others' behavior involved, these decisions may

followed the equilibrium prescription to herd in such herding situations. However, the relative number of rounds in which complete equilibrium play occurred was higher in their experiment, at 60%, as compared to 41% in Game NC. Also, the number of rounds in which herding decisions should occur in equilibrium and do occur in the laboratory, diveded by all rounds in which herding should occur, are 0.65 in Game NC and 0.73 in the games conducted by Anderson and Holt (1997). These numbers are calculated on the basis of the tie rules given in Footnote 7, and not counting decisions as herding decisions when players are indi¤erent in equilibrium and disregard their own signal.

[12] Only in Game NC Bayesian Nash Equilibrium allows for obtaining a signal in situations where herding is prescribed. Hence, the numbers reported in the third row of Table 1 for the treatments with cost do not include situations where a subject buys a signal after an equilibrium history and then, contradicting her own signal, follows her predecessors. However, the numbers of these cases are relatively small for the treatments with cost: In Game HC, a subject bought a signal in a herding situation which contradicted the equilibrium urn decision in 10 cases, and disregarded it in 6 cases. In LC, 13 out of 17 such signals were disregarded, in HCHI 0 out of 5.

[13] These numbers are calculated as expected payo¤s, not just considering the rounds that were selected to be payo¤ relevant for the subjects.

be viewed as mistakes, at least if subjects are risk-neutral money maximizers. An obvious question is whether anticipating this relatively high number of apparent mistakes rationalizes some of the behavior of players at later stages. Below, a model is estimated to determine – among other things – whether players expect other players to deviate from money-maximizing decisions and whether they expect them to do so even more often than is actually observed. For this analysis, all data can be used, not just the decisions following a history which looks like equilibrium play.

4 A statistical depth-of-reasoning analysis

In this section, we present and estimate an error-rate model which allows us to make inferences about the subjects' reasoning processes. The model uses logistic response functions to determine choice probabilities, but specifies separate parameters for the response rationality on each level of reasoning, i.e. 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.

We will first present the behavioral assumptions describing the single-person decision process of a subject who decides at stage n. Let β_n be the probability of the event that the true urn is A, given the nth subject's information before she has the opportunity to see a signal. Also, let $\hat{\beta}_n(s_n;\beta_n)$ be the subject's updated probability of A, after observing a private signal $s_n \in \{a;b\}$, or after deciding not to buy a signal, which will be denoted by $s_n = 0$.[14] The expected payoff from choosing A, after buying a signal with realization s_n, is then given by $\pi(A;s_n;\beta_n) = \hat{\beta}_n(s_n;\beta_n)U-K$, the payoff from choosing B is $\pi(B;s_n;\beta_n) = (1-\hat{\beta}_n(s_n;\beta_n))U-K$: If the subject has not bought a signal, K is not subtracted.

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

[14] Using Bayes' rule, it holds that $\hat{\beta}_n(a;\beta_n) = \frac{\frac{2}{3}\beta_n}{\frac{2}{3}\beta_n+\frac{1}{3}(1-\beta_n)}$ and $\hat{\beta}_n(b;\beta_n) = \frac{\frac{1}{3}\beta_n}{\frac{1}{3}\beta_n+\frac{2}{3}(1-\beta_n)}$. If no signal is bought, no updating can occur, so $\hat{\beta}_n(0;\beta_n) = \beta_n$.

probability

$$\Pr(A; s_n; \beta_n; \lambda_1) = \frac{\exp(\lambda_1 \pi(A; s_n; \beta_n))}{\sum_{j=A,B} \exp(\lambda_1 \pi(j; s_n; \beta_n))}$$

and to choose B with the remaining probability mass.

When deciding whether to buy a signal or not, subjects are assumed to anticipate their own decision probabilities when choosing an urn, to calculate the expected payoffs from their two options accordingly, and to decide logistically: Letting $\bar{\pi}(k; \beta_n)$), $k = y; z$, be the subject's expected payoffs from buying (denoted by y) and not buying (denoted by z), respectively,[15] the probability of buying a signal is given by

$$\Pr(y; \beta_n; \lambda_1) = \frac{\exp(\lambda_1 \bar{\pi}(y; \beta_n))}{\sum_{k=y,z} \exp(\lambda_1 \bar{\pi}(k; \beta_n))}.$$

This two-step decision process is an immediate application of the logit Agent Quantal Response Equilibrium defined by McKelvey and Palfrey (1998) to the present single-person decision problem. As usual in such logistic-choice models, the parameter λ_1 captures the response precision of the decision maker: The higher λ_1, the more "rational" are the decisions. As λ_1 approaches infinity, decision probabilities become arbitrarily close to an optimal pair of responses, given the prior β_n; if $\lambda_1 = 0$, behavior is completely random. At the same time, for any $\lambda_1 > 0$, the probability of making a non-optimal decision decreases with the expected relative loss from this decision.[16]

Now consider the question how subjects make use of their predecessors' decisions when forming their prior beliefs β_n. 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.[17] Analogously, when a subject considers the reasoning that others apply when

[15] These expected payoffs are given by $\bar{\pi}(y; \beta_n) = (\beta_n(\frac{2}{3}\Pr(A; a; \beta_n; \lambda_1) + \frac{1}{3}\Pr(A; b; \beta_n; \lambda_1)) + (1-\beta_n)(\frac{2}{3}\Pr(B; a; \beta_n; \lambda_1) + \frac{1}{3}\Pr(B; b; \beta_n; \lambda_1)))U - K$ and $\bar{\pi}(z; \beta_n) = \beta_n \Pr(A; 0; \beta_n; \lambda_1) + (1-\beta_n)\Pr(B; 0; \beta_n; \lambda_1)$. For all estimates, expectations over the payoff-relevant rounds were used, i.e. all dollar amounts were divided by 15.

[16] A common interpretation is that λ_1 captures the impact of computational errors made by the subjects. For a random-utility justification of Quantal Response Equilibrium models and further discussion see McKelvey and Palfrey (1995, 1998) as well as Anderson, Goeree, and Holt (1999).

[17] Using this specification, it is possible to test the "rational expectations" assumption $\lambda_1 = \lambda_2$; i.e. that on average subjects have a correct perception of the randomization processes of others.

thinking about third subjects, we allow for a third parameter λ_3, which she supposes each of her predecessors attributes to each of their predecessors. For even longer chains of reasoning, additional higher-level parameters are used. With these attributed parameters, the subject determines her belief β_n via Bayes' rule and the formulae for the decision probabilities given in the previous paragraphs.

Since the longest chains of reasoning in the games involve five steps of thinking about other subjects, the resulting model includes six parameters altogether: λ_1 through λ_6. It is essential, however, that higher-level parameters are only applied when a player goes through chains of reasoning of the according length, and not when she directly considers the decision of others who decided several steps before herself. For example, player 3 attributes the precision parameter λ_2 to the decisions of both previous decision makers, because she uses both players' urn choices directly when forming her updated belief. She also attributes the parameter λ_3 to player 1, but only when she considers how player 2 thinks about player 1's decision. Subsequent players go through increasingly complex updating procedures, with different precision parameters for higher levels of reasoning.[18] Using this set of parameters and starting with $\beta_1 = 0.5$, one can recursively construct the players' updated probabilities that A is the true urn, for any history of observed choices.[19]

The model contains a number of special cases that can be tested using the experimental data. If all parameters are equal, we have the logit Agent Quantal Response Equilibrium applied to the entire game, which prescribes for the subjects to know the error rate of other subjects, on all levels of reasoning.[20] If all parameters are infinite, Perfect Bayesian Nash Equilibrium is predicted. Of particular interest are those cases in which one of the parameters is equal to zero, because this reflects the limit in the depth of reasoning. E.g., if $\lambda_2 = 0$ holds, then players behave as if responding to random behavior by all other players, since no information is inferred

[18] Analogously, later players also have to consider that several of their predecessors' chains of reasoning skip player positions when forming their priors. E.g., player 4 has to consider how player 3 considers player 1's behavior directly. Such a chain is treated in the analysis as a chain of length 2 (not of length 3), so λ_3 is applied there.

[19] The updating procedure, which only relies on Bayes' rule, is not presented in more detail here for the sake of brevity. While the procedure is completely analogous for treatments HC, LC, and NC, it differs for the HCHI treatment in that the additional information about signal acquisitions is also taken into account.

[20] In the context of normal-form games, this assumption has been tested using related behavioral models by both Goeree and Holt (2000) and Weizsäcker (2000), and has uniformly been rejected for a large number of games.

from previous decisions. If the first two parameters are strictly positive but $\lambda_3 = 0$ holds, then players only make direct inferences from their predecessors' choices, and do not take into account that these predecessors also think about third players when making their decisions. Similar statements apply to cases in which higher-level parameters vanish. Hence, the length of the reasoning process in the subject pool is reflected by the first parameter that is indistinguishable from zero in the estimation results.

Data:	pooled	HC	LC	HCHI	NC
λ_1	10.45	11.36	8.19	12.97	10.84
	(0.000, 0.000)	(0.000, 0.139)	(0.000, 0.941)	(0.000, 0.000)	(0.000, 0.000)
λ_2	5.94	8.12	8.31	4.71	3.77
	(0.000, 0.000)	(0.000, 0.000)	(0.000, 0.000)	(0.000, 0.000)	(0.000, 0.000)
λ_3	1.65	1.29	2.44	0.00	0.00
	(0.795, 0.194)	(0.994, 0.641)	(0.905, 0.575)	(1.000, 0.970)	(1.000, 0.996)
λ_4	0.00	0.00	0.61	-	-
	(1.000, 0.968)	(1.000, 0.975)	(1.000, 0.908)		
λ_5	-	-	373.32	-	-
			(1.000, 0.981)		
λ_6	-	-	0.00	-	-
			(0.996, -)		
ll^*	-2045.9389	-518.1821	-470.2698	-523.7681	-464.9299

Table 2 : Response precisions estimated from the experimental data. Numbers in parentheses are (i) the marginal level of significance for the parameter to be different from zero, and (ii) the marginal level of significance for the parameter to be different from the parameter on the next-higher level.

Table 2 reports the results of the maximum-likelihood estimation of the model, for the four separate data sets and the pooled data. The table also contains the levels of significance for each parameter to be distinguishable (i) from zero and (ii) from the parameter on the next-higher level of reasoning, which are obtained using appropriate likelihood-ratio tests. An empty cell in the table ("-") indicates that the parameter is not identified. This happens if at the maximum value of the likelihood function a lower-level parameter is estimated to be zero, so beyond this level of reasoning no information is used when decisions are made.

The estimates show a clear distortion in the subjects' perception of their opponents: With only two insignificant exceptions, the response parameters decrease from one level of reasoning to the next, in all four data sets.[21] In particular, a comparison of the estimates for λ_1 and λ_2 shows that subjects on average attribute a lower response precision to their opponents than they have themselves. More strikingly, there is a large gap under all four treatments between the estimated response precisions on the next level, as λ_2 differs from λ_3 at a high significance level. The parameter λ_3, in turn, cannot be distinguished from zero in any of the data sets.[22]

Taken together, the results suggest that the subjects apply only short chains of reasoning, and that the perceived response precisions get lower and lower on higher levels of reasoning. As an interpretation, apart from the possibility of an underestimation of the opponents' response rationality, one may think of these biases as evidence that the subjects' reasoning gets more and more "fuzzy" on higher levels. Within the assumptions of the model, this evidence can serve as an explanation of the observed deviations from equilibrium play in the games, and in particular of the observed signal acquisition behavior. While the subjects tend to "distrust" the response rationality of their opponents (as λ_1 exceeds λ_2), and hence often prefer to buy signals themselves, they also behave as if disregarding the fact that their opponents at least sometimes use the information that is conveyed by third subjects' decisions. Subjects fail to realize that other subjects may have had good reasons not to obtain a signal in later stages of the games. In the high information treatment, they learn how little information is accumulated in the course of the game, which induces them to buy signals with an even higher probability.

5 Conclusions

The paper investigates cascade formation with costly signals. The experimental data exhibit substantial divergence from equilibrium play. In particular, players who have to decide early (but not first) buy too many signals, whereas players who decide toward the end of the games seem confident that previous decisions were based on private signals, hence buy less signals themselves, and herd. When players are informed about who saw a signal and who did not, they tend to buy even more signals than when they have

[21] The hypothesis that all six parameters are equal is rejected on high levels of significance, for each of the data sets.

[22] The hypothesis that the parameter values decrease with a constant ratio between one parameter and the next, as suggested by the model of Goeree and Holt (2000), can only be rejected for the data of the LC treatment, at a 5% level of significance. For the pooled data, the hypothesis is accepted ($p = 0.266$).

to form beliefs about who of their predecessors saw a signal. We explain this finding by limited depth of reasoning, using an error-rate model that allows for false beliefs about the opponents' behavior. The estimation results suggest that players do not consider what their predecessors thought about their respective predecessors, so they do not understand that some of the decisions they observe have been herding decisions, not based on any private information.

These results can perhaps help to assess the value of Bayesian Nash predictions in situations where social learning is possible. Fads may well occur – not because decision makers follow the equilibrium reasoning, but rather because they tend to believe that previous decision makers were informed, and hence follow the majority.

On a more general level, it may be interesting to compare the estimated length of the subjects' reasoning process with the results of previous studies, cited in the Introduction, that investigate lengths of reasoning in different experimental games. In contrast to these studies, we employ a random-utility (or quantal-response) model of behavior, with incomplete information about the others' randomization processes, and draw all our conclusions from the estimations of unobservable parameters. Despite these differences in the estimation approaches, there is a congruence in results with most of the earlier work: The average subject does not make more than two steps of reasoning.

References

[1] Anderson, S.P., Goeree, J.K., and Holt, C.A. (1999): The logit equilibrium: A perspective on intuitive behavioral anomalies. Mimeo, University of Virginia.

[2] Anderson, L.R., and Holt, C.A. (1997): Information cascades in the laboratory. American Economic Review 87(5): 847-862.

[3] Banerjee, A.V. (1992): A simple model of herd behavior. Quarterly Journal of Economics 107: 797-817.

[4] Bikhchandani, S., Hirshleifer, D., and Welch, I. (1998): Learning from the behavior of others: Conformity, fads, and informational cascades. Journal of Economic Perspectives 12: 151-170.

[5] Bikhchandani, S., Hirshleifer, D., and Welch, I. (1992): A theory of fads, fashion, custom, and cultural change as informational cascades. Journal of Political Economy 100: 992-1026.

[6] Costa-Gomes, M., Crawford, V., and Broseta, B. (2000): Cognition and behavior in normal-form games: An experimental study. Mimeo, University of California, San Diego.

[7] Goeree, J.K., and Holt, C.A. (2000): A model of noisy introspection. Mimeo, University of Virginia.

[8] Ho, T., Camerer, C., and Weigelt, K. (1998): Iterated dominance and iterated best response in experimental 'p-beauty contests'. American Economic Review 88(4): 947-969.

[9] Huck, S. and Oechssler, J. (2000): Informational cascades in the laboratory: Do they occur for the right reasons? Journal of Economic Psychology (forthcoming).

[10] Hung, A.A. and Plott, C.R. (1999): Information cascades: replication and extension to majority rule and conformity rewarding institutions. American Economic Review (forthcoming).

[11] Kraemer, C., Nöth, M., and Weber, M. (2000): Information aggregation with costly information and random ordering: Experimental evidence. Mimeo, University of Mannheim.

[12] McKelvey, R.D., Palfrey, T.R. (1995): Quantal response equilibria for normal-form games. Games and Economic Behavior 10: 6-38.

[13] McKelvey, R.D. and Palfrey, T.R. (1998): Quantal response equilibria for extensive form games. Experimental Economics 1: 9-41.

[14] Nagel, R. (1995): Unravelling in guessing games: An experimental study. American Economic Review 85: 1313-1326.

[15] Nöth, M. and Weber, M. (1999): Information Aggregation with Random Ordering: Cascades and Overcon…dence. Mimeo, University of Mannheim.

[16] Sefton, M., and Yavaş, A. (1996): Abreu-Matsushima mechanisms: Experimental evidence. Games and Economic Behavior 16: 280-302.

[17] Stahl, D.O. and P.W. Wilson (1994): Experimental Evidence on players' models of other players. Journal of Economic Behavior and Organization 25: 309-327.

[18] Stahl, D.O. and P.W. Wilson (1995): On players' models of other players: Theory and experimental evidence. Games and Economic Behavior 10: 218-254.

[19] Weizsäcker, G. (2000): Ignoring the rationality of others: Evidence from experimental normal-form games. Mimeo, Harvard University.

Appendix

A Instructions

[Instructions for groups playing the sequence [HC, NC, HCHI].]

This is an experiment in decision making. The Harvard Business School has provided funds for the experiment. Your earnings will be paid to you privately, in cash, at the end of the experiment. Only by coming here, you have already earned a show-up fee of $8. It is very important to us that you do not communicate with the other participants in the room (except via the computer terminal in front of you). Therefore, please remain silent at all times during the experiment. If you have a question, please raise your hand, and we will come to your desk.

This experiment consists of several separate rounds. You will be randomly matched with 5 other participants in the room. This group of 6 persons will be the same over all rounds of the experiment.

In each round you will be asked to predict a randomly chosen urn. A new urn is chosen before each round. It is equally likely that urn A or urn B will be chosen. Urn A contains 2 chips marked "a" and 1 chip marked "b". Urn B contains 1 chip marked "a" and 2 chips marked "b". If you correctly predict the urn, you win a fixed prize.

To help you determine which urn has been selected, each person has the option to see one chip, drawn at random, from the urn. The result of this draw will be your private information and should not be shared with the other participants. After each draw, we will return the chip to the urn before making the next private draw. Each person can have at most one private draw. This is done on the computer. When it is your turn to obtain your draw, you will be asked on the computer screen whether you want to obtain the draw. To obtain the draw you have to pay a fee. If you decide to see the draw, click on the button marked "Yes" that will appear on your computer screen. If you don't want to see your draw, click on the button marked "No". If you clicked on the "Yes"-button, your window will read "The draw is: b" if the chip that the computer has randomly drawn for you is marked "b". And it will read "The draw is: a" if the chip that the computer has randomly drawn for you is marked "a".

After you have or have not seen you draw, you will be asked to input the letter of the urn (A or B) that you think is more likely to have been used. But you will also see the choice of all persons in your group who made a decision before you. For all persons in your group, the same urn will be used. (But remember that a new urn will be chosen before each round.) The order in which you get to decide (among the participants in your group) is randomly determined every round. The …rst person who has to make a decision (in your group) sees no other decisions. The second person sees which urn the …rst person chose. The third person sees the choice of the …rst and the second person, and so forth. This process will be repeated until all 6 people have made decisions. Finally, we will inform everyone of the urn that was actually used.

Your earnings are determined as follows: If your decision matches the urn that was actually used, you earn $12. Otherwise you get nothing. If you decided to see your draw, you have to pay $1.50, which will be deducted from your total earnings. The …rst part of the experiment consists of 15 rounds. After the experiment, we will randomly determine one of these 15 rounds that will actually count for your earnings. All other rounds are not relevant for your earnings. If you decided to obtain a draw in any of these other periods, you will also not have to pay for it. Your earnings in the selected round will be added to your show-up fee of $8.

At the end of each round, your choice will be recorded on your screen along with the actual urn used, and your payo¤ for the round. (This round payo¤ will be shown in all 15 rounds, not just in the one that determines your earnings.)

(The following paragraphs in italics will also be read aloud to you.)

Before we begin the actual experiment, we will go through a demonstration. We will show how the actual urn is chosen, and the process by which the draws are made. Note that, in the actual experiment, the computer will choose the actual urn, as well as make the draws for each person.

We have two bags (the urns), labeled "A" and "B". Urn A contains two chips marked "a" and one chip marked "b". Urn B contains two chips marked "b" and one chip marked "a". Now we will ‡ip a coin, to determine which urn is chosen. If the result of the coin ‡ip is heads, then urn A will be used, and if the result of the coin ‡ip is tails, then urn B will be used. We will now draw a chip for the …rst person who wants to see his or her draw. If this were not just a demonstration, then this person would see the letter marked on the chip on the screen, and make a decision by checking the button for urn A or urn B on the screen.

Then, we will draw a chip for the next person who wants to see his or her draw. If this were not a demonstration, this person would see the letter marked on the chip on the computer screen, also see the decision of the …rst person on the screen, and make a decision (A or B) him- or herself.

Are there any questions before we begin?

Please do not open other windows on the computer while the experiment is running.

Before we begin the actual rounds, we will go through one practice round, for which you will not be paid. After this practice round, the 15 rounds of the …rst part of the experiment start.

[The following part of the instructions was distributed after the …rst 15 rounds were over.]

That concludes the …rst part of the experiment. For the second part, the procedure is the same, but you (and all other participants in your group) will now be able to see the draw for free. So, nothing will be deducted from your earnings when you see your draw. Remember that only one randomly chosen round (out of 15) is relevant for your earnings.

Before the second part begins, we will again have one round of practicing the game, which does not count for your earnings. After that, the new 15 rounds will start.

[The following part of the instructions was distributed after the second 15 rounds were over.]

That concludes the second part of the experiment. For the third part, it again costs $1.50 to see your draw. The procedure is the same, but you will now see whether the persons deciding before you have seen their draws or not. On your screen you will now read for example "The 1st player's decision was: B. The 1st player did obtain a draw." if the …rst player chose urn B and decided to see his draw. You will read "The 1st player's decision was: B. The 1st player did not obtain a draw." if the …rst player chose urn B and decided not to see the draw. Similarly, this information will appear on your screen for all persons who made decisions prior to you. When it is your turn, you will be asked whether you want to see your draw. After seeing or not seeing your draw, you will have to make a decision (urn A or urn B).

As before, we will have one practice round before the 15 rounds of the third part begin. Again, one randomly determined round out of all 15 rounds will count for your earnings.