LSE Research Online

Article (refereed)

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

Dorothea Kübler and Georg Weizsäcker

Limited depth of reasoning

and failure of cascade formation in the laboratory[*]

by

Dorothea Kübler[†]   Georg Weizsäcker[‡]

[†]Institute for Economic Theory I, Department of Economics, Humboldt University Berlin, Spandauer Str.1, 10178 Berlin, Germany. Email address: kuebler@wiwi.hu-berlin.de

[‡]Department of Economics, Harvard University, and 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 or not to obtain private information, at a small but positive cost. In the equilibrium of this game, only the first player buys a signal and makes a decision based on this information whereas all following players do not buy a signal and herd behind the first 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 different error rates on different levels of reasoning, we find that the subjects' inferences become significantly more noisy on higher levels of the thought process, and that only short chains of reasoning are applied by the subjects.

3

1 Introduction

In simple cascade games, the players sequentially choose one out of two alternatives, after receiving private signals about the profitability 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.

Following the papers by Banerjee (1992) and Bikhchandani, Hirshleifer, and Welch (1992), models of information cascades have been used to explain a great number of economic phenomena. They include consumers herding behind other consumers' purchasing decisions, herding among security analysts and mutual fund managers, herding among bank customers resulting in bank runs, waves of mergers and waves of takeovers, herding among economic forecasters, adoption of certain medical procedures by doctors imitating other doctors, potential employers not hiring a candidate with a history of joblessness, etc.[1]

From a behavioral perspective, one can ask whether the reasoning process underlying Bayesian Nash Equilibrium in cascade games is applied by actual decision

---

[1] For surveys see Bikhchandani, Hirshleifer, and Welch (1998) as well as Ivo Welch's homepage with an annotated bibliography (http://welch.som.yale.edu/cascades/).

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, initial 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 (2001). Other researchers find only limited support for Bayesian Nash Equilibrium in these games. In particular, Nöth and Weber (1999) identify a tendency of subjects to follow their own signals in situations where equilibrium prescribes to follow one's predecessors. Huck and Oechssler (2000) confront their subjects with related single-person decision tasks and find that Bayes' rule is systematically violated. Bracht, Koessler, Winter, and Ziegelmeyer (2000) study a number of different counting heuristics to learn whether and how subjects base their decisions on counts of their predecessors' decisions.

We modify the experimental design of 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 modified 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 first player buys a signal, chooses the according urn, and *all* subsequent players blindly follow the first player's decision. Thus, after the first player's choice, no further signals are bought, and cascades occur with certainty.

The experimental results are not in line with these predictions. While not all of the subjects acting as first players buy a signal, signal acquisitions in later stages are excessive, such that overall far too many signals are bought. Subjects tend to follow the majority of preceding urn choices, but only if this majority is strong enough. Cascades rarely start after the first player's urn choice, as subjects at the second and third stages often buy signals and, if appropriate given their signal, choose in opposition to their predecessor's decision. The predictive value of Bayesian Nash Equilibrium is much lower in the game with costly signals than in a control treatment where signals are costless, which is comparable to Anderson and Holt's design.

A natural candidate to explain the excessive signal acquisitions are errors. Subjects may simply err or tremble 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 in this treatment, and the prediction of Bayesian Nash Equilibrium – which is identical in both treatments – performs worse.[2]

To explain these observations, we conduct a depth-of-reasoning analysis. I.e., we

---

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

employ a statistical model (based on work by McKelvey and Palfrey, 1998) 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 various "anomalies" that can arise in long chains of reasoning, and to obtain an estimate of the actual depth of reasoning in the subject pool.

Depth-of-reasoning analyses have been conducted by several experimentalists (see e.g., the papers by Nagel, 1995, Sefton and Yavaş, 1996, and Ho, Camerer, and Weigelt, 1998), but they all investigate normal-form game play.[3] We argue that cascade games are especially well suited for an analysis of depth of reasoning, largely because they are extensive-form games, and because a player's payoff is independent of what other players do: 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, because cascade games are extensive-form games in which a player's payoff is not affected by later players' actions, the results do not depend on the subjects' ability to solve a game backwards,

---

[3]Relatedly, Stahl and Wilson (1995), Goeree and Holt (2000), Costa-Gomes, Crawford, and Broseta (2001), and Weizsäcker (2002) estimate models of normal-form game play behavior which allow for a limited depth of reasoning.

which is often doubted.

The model-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 payoff-increasing in the games. In other words, the subjects learn from observing their predecessors' decisions, but they fail to realize that other subjects also learn from observing their respective predecessors.

The subjects' signal acquisition behavior can be explained along the lines of these estimation results. In the treatment with cost, many 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 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 acquisitions of their predecessors, they are often surprised about how little information previous urn decisions were based on, and hence tend to buy even more signals.

The next section describes the experimental design and procedures. Section 3

8

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 $t, t = 1, ..., 6$, the tth player

  1. observes the $(t-1)$ 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 $s_t \in \{a, b\}$), 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.

In both cost treatments, the ratio of the signal cost to the possible prize, K/U, is below one sixth. 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. The same logic applies to all subsequent players. Thus, cascades always occur, independent of the signal realizations, and all games can be used for an analysis of herding behavior.

As the equilibrium prediction critically hinges on the players relying on the first player to 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 conducted a high information treatment, game HCHI.

Game HCHI:

All stages are as in game HC, except that the tth player, before making her own decisions, also observes whether or not each of the previous $(t - 1)$ 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 strategies 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 with costless signals, 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 at which 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 (see the supplementary appendix)[4].

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

---

[4]The supplementary appendix can be found at http://www.restud/org.uk/supplements/htm.

equilibria would prescribe for her to disregard her own signal and also choose B.[5] Assuming a specific tie-rule, one can then observe how many of the subjects' choices are consistent with the equilibrium path prescribed by the corresponding equilibrium.[6] Importantly, the equilibrium prediction here is different from games HC, LC, and HCHI, as more signals are obtained. Figure 5 in Appendix A illustrates the possible equilibrium paths for game NC.

2.2   Experimental procedure

The experiment was run in four sessions at the Computer Lab for Experimental Research at Harvard Business School, using the software Z-Tree (see Fischbacher, 1999). At the beginning of each session, two draws from physical urns were made as a demonstration. Afterwards, 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 during the entire session and played the games with player roles randomly changing after each round. Table 1 shows that in sessions 1 and

---

[5]This is not true if the equilibrium prescribes for the second player to always follow the first 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 first player's decision, then two preceding B's are sufficient for the third player to disregard her own a signal.

[6]Anderson and Holt (1997) consider the tie-rule "Follow your own signal if indifferent". To simplify our analysis, we restrict attention to an analogous tie-rule for game NC: "If indifferent 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 player's choice 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 prediction in most cases. Notice that in HC, LC, and HCHI, the equilibrium prediction does not rely on a specific tie-rule.

2, the subjects played games HC, NC, and HCHI, and in sessions 3 and 4, subjects only played game LC.

Table 1 : Experimental sessions.

session	# of subjects	# of groups	# of rounds per treatment	treatment order
1	24	4	15	HC, NC, HCHI (2 groups); NC, HC, HCHI (2 groups)
2	12	2	15	HC, NC, HCHI (1 group); NC, HC, HCHI (1 group)
3	18	3	15 (+15)[7]	LC
4	12	2	15	LC

Each game was played for 15 rounds, preceded by an unpaid practice round. Subjects in sessions 1 and 2 were not told what would happen after the first and second set of 15 rounds. To ensure at least partially that differences in behavior between games are not due to learning, we switched the order of games HC and NC within the first two sessions (see last column of Table 1). Game HCHI was always

---

[7] 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 different treatments, we decided not to include these data in the analysis and only used the first 15 rounds. Behavior in the second part of the session was very similar to the first part.

13

played at the end of the session, to prevent subjects from transferring information about how many subjects bought signals to the other games, which could distort the results.

Overall, 66 subjects, mostly undergraduate students from universities in the Boston area, participated in the experiment. Given the number of rounds chosen, this implies that games HC, NC, and HCHI were played 90 times each and game LC was played 75 times, yielding a total of 4140 decisions. At the end of each session one payoff-relevant round per treatment was randomly determined by a draw from a stack of 15 numbered cards. The earnings from these rounds were added to a show-up fee of $16. Average earnings were $36.47 in sessions 1 and 2, $29.20 in session 3, and $22.80 in session 4. The subjects were identified by code numbers only and received their total earnings in cash directly after the experiment.

3   Results: Descriptive statistics

Figures 1 through 4 summarize to what extent Bayesian Nash Equilibrium predicts the behavior in the four experimental treatments. For treatments HC, LC, and HCHI in particular, they display how often first players bought a signal and chose the indicated urn, and how often second to sixth players did not buy a signal and followed their predecessor's urn choice. For treatment NC, the prescribed equilibrium-path decisions can be taken from Figure 5. For all four treatments, only decisions are considered that follow an equilibrium-path history of previous play. The first (white) column reports the relative frequency of subjects making the signal acquisition decision prescribed on the equilibrium path at each stage, contingent on the observation

of equilibrium behavior by the previous subjects. That is, for HC, LC, and HCHI it shows the proportions of subjects buying a signal at stage 1 and the proportion of subjects not buying a signal at stages 2 to 6. Likewise, the second (shaded) column shows the relative frequency of subjects following the equilibrium path in both the signal acquisition and the urn decision. For games HC, LC, and HCHI, the column thus displays the proportion of subjects following their signal at the first stage, and the proportion of subjects following their predecessor's choice without buying a signal at all later stages. The third (black) column represents the cumulated equilibrium decisions. It shows in how many rounds *all decisions up to* (and including) the respective stage follow the equilibrium path, in both signal acquisitions and urn choices.

Note that restricting the decisions to those arising along an equilibrium-path history has different implications in the four games. For games HC and LC, all decisions following histories that contain only urn A choices or only urn B choices are considered in the construction of the corresponding figures. For game HCHI, an additional requirement is that the first player bought a signal and the others did not. The equilibrium paths of game NC can be different from those in the other games, as summarized in Figure 5.[8]

---

[8]Histories are still 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. This allows for unobservable deviations in signal acquisitions (in games HC, LC, and NC). Also, if for example the first player observes a signal a but chooses urn B, the second player's decision would still be included in the graphs. However, for the third column, representing the cumulation of equilibrium behavior, *all* observable and unobservable decisions must be in equilibrium.

HC

Figure 1: Frequencies of decisions that are consistent with Perfect Bayesian Nash Equilibrium in Game HC, conditional on equilibrium-path play up to the respective stage.

LC

Figure 2: Frequencies of decisions that are consistent with Perfect Bayesian Nash Equilibrium in Game LC, conditional on equilibrium-path play up to the respective stage.

HCHI

Figure 3: Frequencies of decisions that are consistent with Perfect Bayesian Nash Equilibrium in Game HCHI, conditional on equilibrium-path play up to the respective stage.

NC

Figure 4: Frequencies of decisions that are consistent with Perfect Bayesian Nash Equilibrium in Game NC, conditional on equilibrium-path play up to the respective stage.

Consider the results of treatment HC, in Figure 1. The white column indicates for the first stage that 52% of all signal acquisition decisions are in line with the equilibrium prediction (i.e., the subjects decided to see a signal). Of the observed urn choices at stage 1, almost all were as predicted when the subject had obtained a signal. This can be taken from the second column of the figure, which is almost of the same size as the first. It is also evident from the figure that the number of equilibrium signal acquisition decisions increases at later stages. As the equilibrium predicts not to see a signal at stages 2 to 6, the white columns show that subjects acting as player 2 buy too many signals whereas later players rarely buy signals if previous play is consistent with Bayesian Nash Equilibrium. This suggests that most subjects follow the majority of urn choices once enough people have chosen the same urn. Apparently, they do not consider the first player's choice as strong enough evidence, but herd after two or more identical urn choices.

Similar observations hold for treatment LC where, however, more signals are bought at early stages. Concerning treatment HCHI, notice that the proportions of observed equilibrium signal decisions and of observed equilibrium signal and urn decisions at stages 2 to 6 are smaller than in treatment HC (with the single exception of signal acquisitions at stage 4). 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. The results of the no-cost treatment NC are closest to the equilibrium prediction: The proportion of signal and urn decisions consistent with equilibrium is quite stable, at a level of at least 80% in all stages.[9]

---

[9] Notice, however, that taking an equilibrium signal decision in treatment NC is less difficult than in the three treatments with costly signals, because seeing a signal is always an optimal decision.

These observations are corroborated by the third (black) column, which shows the cumulative proportions of equilibrium play. In the course of the six stages, this proportion decreases quite dramatically in the treatments with 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 treatment NC, all six participants play according to the equilibrium in 41% of the 90 rounds. Hence, with a positive signal cost the equilibrium prediction performs worse, according to these aggregate numbers.

For games HC, LC, and HCHI the figures reveal how often subjects actually herd, i.e., how often they "correctly" follow their predecessors. In the equilibria of these games, herding should occur in all stages after stage 1, as the prediction for players 2 through 6 is to blindly follow suit. In roughly two thirds of the cases where subjects observed an equilibrium-path history they played in accordance with this prediction (69% in HC, 61% in LC, and 61% in HCHI.) In game NC, however, obtaining a signal is always optimal, so a different measure of the propensity to herd is needed. Consider the cases in treatment NC where a player saw a signal contradicting the equilibrium prediction for her urn choice, after an equilibrium-path history. Subjects herded (i.e. disregarded their signal) in 78% of these cases (63/81).[10] Additionally, in

(Cf. Footnote 6 for the tie rule applied.) Note also that players *must* see a signal in equilibrium if no information cascade has started yet. This explains why some of the white columns in Figure 4 are below 100%.

[10] In the corresponding treatment by Anderson and Holt (1997), 70% of the subjects followed the equilibrium prescription to herd in such herding situations. The percentage of rounds in which complete equilibrium play occurred was higher in their experiment, at 60%, as compared to 41% in Game NC. This difference can be explained by the fact that in our design, subjects also had to

14% (32/224) of the cases in which the signal realization is irrelevant for the optimal urn choice (again, after an equilibrium history) did a subject not obtain a signal and followed the predecessors' urn choices.

While the equilibrium is often socially inefficient (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 efficiency, computed as the sum of all players' payoffs, relative to the equilibrium. Expressed in percentages of the total payments that would have been received in equilibrium (given nature's initial draw and the signal realizations), total earnings in the four treatments were: 81.3% in HC, 96.8% in LC, 78.6% in HCHI, and 91.4% in NC. Earnings are closer to the equilibrium in game LC than in NC although there are much less equilibrium decisions in LC than in NC. This is explained by the fact that too many players buy a signal in LC, which increases total earnings as signals are not very costly and additional information is revealed.[11]

We now include behavior off the equilibrium path in the descriptive data analysis. Generally, analyzing off-equilibrium behavior is difficult in cascade experiments due to the sheer number of different possible histories of signal and urn decisions.

decide whether to see a signal or not.

[11] Criticizing the above calculations, it can be argued that our notion of efficiency is only appropriate in the laboratory, where there are no effects on third parties who are not part of the game. When such externalities are present, stopping a wrong cascade can potentially generate large welfare gains. Bikhchandani, Hirshleifer, and Welch (1992) provide the historical example of doctors performing tonsillectomies on a routine basis, merely because other doctors have done the same before. In this case, the wrong cascade had very serious negative externalities (children were injured during the procedure, tonsils have been found to be a defense against infections, etc.).

However, transition matrices can be used to organize the data by pooling histories with identical numbers of urn A and B choices (Tables 3 to 6 in Appendix A). These matrices indicate *(i)* the proportion of signals bought, and *(ii)* the proportion of subjects disregarding their own signal in favor of the urn most frequently chosen by the predecessors, after a history with a given number of urn "A" and "B" choices. (Urns are interchangeable in our symmetric setup, and the urn chosen more often than the other is called urn "A". The urn less frequently chosen is called urn "B". Quotation marks are used to indicate this change in notation.)

The tables show that the greater the difference between the number of previous urn "A" and urn "B" choices, the less signals are bought. This holds both on and off the equilibrium path. For example, the second row of Table 3 (for game HC) shows that the relative frequency of obtained signals decreases from 47% after one urn "A" and one urn "B" choice, to 3% after four urn "A" choices and one urn "B" choice. In addition, in game NC (the only treatment where the number of herding decisions after seeing a signal is large enough to draw any conclusions) the greater the difference between the number of urn "A" and urn "B" choices, the more subjects disregard their own signal.

To check whether these behavioral patterns changed over the course of the experimental sessions, we also computed transition matrices for earlier and later rounds separately by splitting the data between rounds 1-8 and 9-15. No significant changes in the transition probabilites are discernible in any of the four treatments. (The tables are not included in the paper.)

Note that the transition matrices in Tables 3 to 6 control neither for the order

of urn "A" and urn "B" decisions in the histories considered, nor for the knowledge about previous signal acquisitions in treatment HCHI. We therefore conducted probit regressions, which are presented in the supplementary appendix. To sum up the findings, the order of previous "A" and "B" choices has little or no impact on later decisions, whereas the difference between the number of "A" and "B" choices significantly affects behavior. This is consistent with the observed tendency to follow the majority of urn choices once it is strong enough. Furthermore, in treatment HCHI subjects clearly take into account whether the urn choices of their predecessors are based on a signal or not. Urn choices that do not follow signal draws are essentially disregarded by later players.

Finally, consider Figures 1 to 4 again, and in particular the columns for the first stage of each game. As these columns never reach 1, some first players deviate from equilibrium play, either when deciding whether to buy a signal or when choosing the urn. In particular, 48% of all first players in treatment HC decide not to see a signal, 31% in treatment LC, 53% in treatment HCHI, and 11% in treatment NC.[12] Since there is no uncertainty about others' behavior involved, these decisions may be viewed as mistakes, at least if subjects are considered to be risk-neutral money maximizers. An obvious question is whether anticipating these apparent mistakes rationalizes some of the behavior at later stages. Below, a model is estimated to determine – among other things – whether players expect other players to deviate

---

[12]The number of equilibrium deviations of first players in NC does not differ much from Anderson and Holt's results. In their experiment, 10% of the subjects in the first stage did not follow their private signal. This happened in about 7% of all cases where first players saw a signal in our experiment.

from money-maximizing decisions and whether they expect them to do so as often as is actually observed. For this analysis, the decision data need not be separated according to the histories of previous play, but all data can be used simultaneously.

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 t. Let α_t be the probability of the event that the true urn is A, given the tth subject's information before she has the opportunity to see a signal. Let $\widetilde{\alpha}_t(s_t, \alpha_t)$ be the subject's updated probability of A, after observing a private signal $s_t \in \{a, b\}$, or after deciding not to buy a signal, which will be denoted by $s_t = 0$.[13] Also, denote by $c_t \in \{A, B\}$ the subject's urn choice. Her expected payoff from choosing A, after buying a signal with realization s_t, is given by $\widetilde{u}(A, s_t, \alpha_t) = \widetilde{\alpha}_n(s_t, \alpha_t)U - K$, and the payoff from choosing B is

---

[13] Using Bayes' rule, it holds that $\widetilde{\alpha}_t(a, \alpha_t) = \frac{\frac{2}{3}\alpha_t}{\frac{2}{3}\alpha_t + \frac{1}{3}(1\text{-}\alpha_t)}$ and $\widetilde{\alpha}_t(b, \alpha_t) = \frac{\frac{1}{3}\alpha_t}{\frac{1}{3}\alpha_t + \frac{2}{3}(1\text{-}\alpha_t)}$. If no signal is bought, no updating can occur, so $\widetilde{\alpha}_t(0, \alpha_t) = \alpha_t$.

23

$\widetilde{u}(B, s_t, \alpha_t) = (1 - \widetilde{\alpha}_t(s_t, \alpha_t))U - K$. 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 probability

$$\Pr(A; s_t, \alpha_t, \lambda_1) = \frac{\exp(\lambda_1 \widetilde{u}(A, s_t, \alpha_t))}{\sum_{c_t = A, B} \exp(\lambda_1 \widetilde{u}(c_t, s_t, \alpha_t))}$$

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: Let $\overline{u}(b_t, \alpha_t))$, $b_t \in \{"Buy", "Don't\ Buy"\}$, be the subject's expected payoffs from buying and not buying respectively.[14] The probability of buying a signal is then given by

$$\Pr("Buy"; \alpha_t, \lambda_1) = \frac{\exp(\lambda_1 \overline{u}("Buy", \alpha_t))}{\sum_{b_t} \exp(\lambda_1 \overline{u}(b_t, \alpha_t))}.$$

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 α_t; if $\lambda_1 = 0$, behavior

---

[14]These expected payoffs are given by $\overline{u}("Buy", \alpha_t)) = (\alpha_t(\frac{2}{3}\Pr(A; a, \alpha_t, \lambda_1) + \frac{1}{3}\Pr(A; b, \alpha_t, \lambda_1)) + (1 - \alpha_t)(\frac{2}{3}\Pr(B; a, \alpha_t, \lambda_1) + \frac{1}{3}\Pr(B; b, \alpha_t, \lambda_1)))U - K$ and $\overline{u}("Don't\ Buy", \alpha_t) = \alpha_t \Pr(A; 0, \alpha_t, \lambda_1) + (1 - \alpha_t)\Pr(B; 0, \alpha_t, \lambda_1)$. For all estimates, expectations over the payoff-relevant rounds were used, i.e., all dollar amounts were divided by 15.

is completely random. Also, for any $\lambda_1 > 0$, the probability of making a non-optimal decision decreases with the expected loss from this decision.[15]

Now consider the question how a subject makes use of her predecessors' decisions when forming her prior belief α_t. It is assumed the subject is aware that all other subjects follow the logistic decision process described above, with the exception that she attributes a possibly different precision parameter to the decisions of her opponents: λ_2 instead of λ_1. (This is similar to the model estimated in Weizsäcker, 2002.) Thereby, the "rational expectations" assumption of the Quantal Response Equilibrium, that subjects are informed about the error rate of their opponents, is avoided and can be tested.

Analogously, when a subject considers the reasoning that others apply when thinking about third subjects, we allow for a third parameter λ_3, which she supposes each of her predecessors attributes to each of his or her predecessors. For even longer chains of reasoning, additional higher-level parameters are used. 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. Using this set of parameters, and starting with $\alpha_1 = 0.5$, one can recursively construct the players' updated probabilities that A is the true urn, for any history of observed choices (see the supplementary appendix).

Note that the subscript of λ indicates the number of iterations made when thinking about how others think about how others..., not the stage at which the player has to

---

[15] 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 e.g. McKelvey and Palfrey (1995, 1998).

make a decision.[16] Also, it is important to notice 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. As another example, player 6 attributes λ_2 to all five previous decision makers. He also considers how each of them considers his or her respective predecessors, and attributes the corresponding higher-order parameters to these steps of reasoning. E.g., when player 6 considers how player 3 considers player 1's urn decision, he attributes λ_3 to player 1 and λ_2 to player 3. When player 6 considers how player 3 considers how player 2 considers player 1's decision, then λ_4 is attributed to player 1, λ_3 to player 2, and λ_2 to player 3.

The model contains a number of special cases that can be tested using the experimental data. If all six parameters are equal, we have the logit Agent Quantal Response Equilibrium applied to the entire game. It prescribes that the subjects know the error rate of the other subjects, on all levels of reasoning.[17] If all parameters are infinite, Perfect Bayesian Nash Equilibrium is predicted. Of particular interest are

---

[16]Anderson and Holt (1997) also conducted an analysis of their data based on the logit Agent Quantal Response Equilibrium, but assumed rational expectations of players and a fixed set of different λ-parameters at different stages.

[17]In the context of normal-form games, this assumption has been tested using related behavioral models by both Goeree and Holt (2000) and Weizsäcker (2002), and has uniformly been rejected for a large number of games.

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 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 their 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.

Some special cases of the model can be interpreted as behavioral heuristics that players might apply in the games. In particular, when $\lambda_1 \rightarrow \infty$ and $\lambda_2 = 0$, players use the rule "follow your own signal" as they perceive other players to be randomizing.[18] Note also that when $\lambda_1 = \lambda_2 \rightarrow \infty$ and $\lambda_3 = 0$, players apply a counting heuristic in the cascade games we consider. I.e., they follow the majority of urn decisions, and if there is no majority, they buy a signal and follow it. The counting heuristic is a best response if a player believes that her predecessors do not learn anything from the actions of their predecessors and if she therefore supposes that her predecessors buy a signal and follow it.

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

---

[18]See the papers cited in Footnote 3 for related evidence in normal-form games, as well as Beard and Beil (1994) and Huck and Weizsäcker (2002) for games of extensive form.

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.

Table 2 : Response precisions estimated from the experimental data.

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

Note: 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.

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. The hypothesis that all six parameters are equal is rejected on high levels of significance, for each of the data sets. 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.[19] More strikingly, in all four treatments there is a large gap between the estimated response precisions of the next levels, as λ_2 significantly differs from λ_3. The parameter λ_3, in turn, cannot be distinguished from zero in any of the data sets.[20]

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. This points at a consistent underestimation of the opponents' response rationality. As an alternative interpretation, one may think of these biases as evidence that the subjects' reasoning gets more and more fuzzy on higher levels.

Along these lines, one can explain the observed deviations from equilibrium play

[19]Note that the high λ_1-parameter is compatible with the relatively small number of subjects buying a signal in the first stage. The reason is that the expected payoff from buying a signal is not much higher than from not buying a signal, especially in the high-cost treatments HC and HCHI. However, with a high value of λ_1, a person is likely to follow her signal if she has bought one, because payoff differences with respect to urn choice are larger.

[20]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 ($p = 0.266$).

in the games, and in particular of the observed signal acquisition behavior. First, the subjects distrust the response rationality of previous players (as λ_1 exceeds λ_2), and hence tend to buy signals themselves. Second, they behave as if disregarding the fact that their predecessors often 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, and therefore do not recognize herding behavior. Thus, they consider every urn choice of their predecessors as (about) equally informative and follow the majority. However, in the high information treatment, they learn how little information is accumulated in the course of the game, which induces them to buy signals at later stages with an even higher probability.

To test the robustness of the statistical results, we also considered three variations of the above model estimations, one allowing for more general risk attitudes, another for learning effects, and the third for subject heterogeneity. The results of the estimations are shortly summarized in the following.

Concerning the question of risk considerations, we followed the analyses by Goeree and Holt (2000) and Goeree, Holt, and Palfrey (2002) who incorporate constant-relative-risk-aversion utilities into related models of probabilistic choice, instead of assuming risk neutrality. The according generalization of our model estimations leaves the main results untouched, as the λ-parameters decrease from one level to the next, and λ_3 is still indistinguishable from zero in all data sets.[21]

---

[21] To keep the paper short, we do not specify the details of the estimations and results here. In two treatments (HC and HCHI) we find significant evidence of risk aversion, in one treatment (LC) risk-loving behavior. However, a caveat is that when estimating risk attitudes from experimental data, the results crucially rely on the mental frame or status-quo point that subjects are assumed to

Now consider the question whether the observed belief distortions are stable over the 15 rounds of the games. Expressed in terms of the statistical model, there are two possible ways in which subjects may learn: The λ-parameters could increase, and they could lie closer together in later rounds. In order to investigate these issues, we again partitioned the data sets into two subsets each, one containing only decisions made earlier in the games (rounds 1-8), the other only later decisions (rounds 9-15), and reestimated the error-rate model using these subsets of data separately. As in the descriptive analysis of Section 3, no significant evidence of learning can be discerned (with significance levels above $p = 0.2$ in all four treatments, using likelihood ratio tests). Tables 8 and 9 in the supplementary appendix show that both in earlier and in later rounds of the games the estimated λ-parameters are smaller on higher levels of reasoning, and that, again, λ_3 cannot be distinguished from zero in any of the data subsets.

Estimations of the λ-parameters for each subject separately, summarized in Table 10 in the supplementary appendix, confirm the previous results. Of the 36 subjects who played three games (HC, HCHI, and NC) 30 subjects have a λ_1-parameter which exceeds λ_2, and for 31 out of 36 subjects λ_2 exceeds λ_3. These relations are similar when considering single games. Also, the table shows that more than half of the subjects have an estimated $\lambda_3 = 0$, which again supports our finding from the aggregated dataset.

have. The reason is that neither all contingencies within the experiment, nor the subjects' "outside" wealth can be considered. See Rabin (2000) for a more rigorous discussion. In our risk-attitude analysis, subjects were assumed to view each round of the experiment separately.

31

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. We explain these findings 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 systematically misperceive other players in two ways. First, they attribute an error rate to their opponents that is higher than their own. This bias leads them to rely too little on their predecessors, and hence to acquire too many signals themselves. Second, players do not consider what their predecessors thought about their respective predecessors. Thus, they do not understand that some of the decisions they observe have been herding decisions, not based on any private information. Many players therefore follow the majority of urn choices, once this majority is sufficiently strong.

The results of the model estimation provide a unified explanation for both, the herding behavior observed in some earlier cascade experiments, as well as the deviations that we and other researchers discuss. Along these lines, the results can perhaps help to assess the value of Bayesian Nash predictions in other 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. It may be worth noting that if such behavior is prevalent, the order of the previous players' decisions is generally

32

irrelevant for the outcome of a cascade game, because each predecessor is viewed as if following only his or her own signal. (This conclusion is also supported by the regression analysis presented in the supplementary appendix.) The independence of the order of choices would hold even if some players had better private information than others.

On a more general level, it seems worthwhile to compare the estimated length of the subjects' reasoning process with the results of previous studies investigating lengths of reasoning in experimental games (see the citations in the introduction). 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 our conclusions from the estimations of unobservable parameters. Despite these differences in the estimation approaches, our results are consistent with most of the earlier work: The average subject does not make more than two steps of reasoning.

References

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

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

Beard, T.R. and Beil, R.O. (1994): Do people rely on the self-interested maximization of others? An experimental test. *Management Science* 40: 252-262.

Bikhchandani, S., Hirshleifer, D., and Welch, I. (1992): A theory of fads, fash-

ion, custom, and cultural change as informational cascades. *Journal of Political Economy* 100: 992-1026.

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.

Bracht, J., Koessler, F., Winter, E., and Ziegelmeyer, A. (2000): Behaviors and beliefs in information cascades. Mimeo, Hebrew University of Jerusalem.

Costa-Gomes, M., Crawford, V., and Broseta, B. (2001): Cognition and behavior in normal-form games: An experimental study. *Econometrica* 69: 1193-1235.

Fischbacher, U. (1999): Z-Tree, Zurich Toolbox for Readymade Economic Experiments. Working Paper 21, Institute for Empirical Research in Economics, University of Zurich.

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

Goeree, J.K., Holt, C.A., and Palfrey, T.R. (2002): Quantal response equilibrium and overbidding in private value auctions. *Journal of Economic Theory* 104: 247-272.

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.

Huck, S., and Oechssler, J. (2000): Informational cascades in the laboratory:

Do they occur for the right reasons? *Journal of Economic Psychology* 21: 661-671.

Huck, S., and Weizsäcker, G. (2002): Do players correctly estimate what others do? Evidence of conservatism in beliefs. *Journal of Economic Behavior and Organization* 47: 71-85.

Hung, A.A., and Plott, C.R. (2001): Information cascades: replication and extension to majority rule and conformity rewarding institutions. *American Economic Review* 91: 1508-1520.

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

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

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

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

Nöth, M., and Weber, M. (1999): Information Aggregation with Random Ordering: Cascades and Overconfidence. Mimeo, University of Mannheim.

Rabin, M. (2000). Risk aversion and expected-utility theory: A calibration theorem. *Econometrica* 86: 1281-1292.

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

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

Weizsäcker, G. (2002): Ignoring the rationality of others: Evidence from experimental normal-form games. *Games and Economic Behavior*, forthcoming.

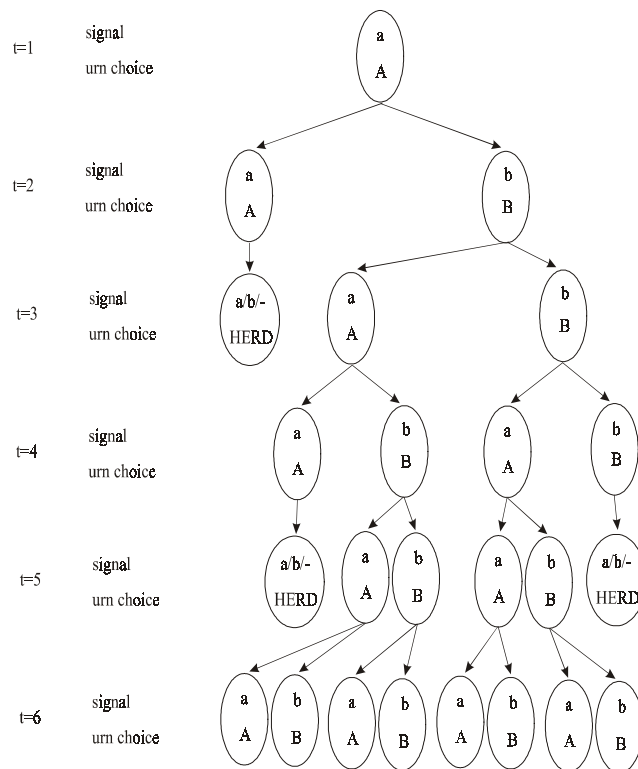## A   Figure 5 and Tables 3-6

Figure 5: Equilibrium paths in game NC, starting with signal "a". "HERD" indicates that current and all subsequent players follow the previous urn choice.

Table 3: Transition matrix for HC

"A" "B"	0	1	2	3	4	5
0	47/90 (.52) —	40/90 (.44) 2/17 (.12)	7/59 (.12) 2/3 (.67)	7/49 (.14) 3/5 (.60)	5/45 (.11) 1/3 (.33)	1/41 (.0 (no case
1		15/32 (.47) —	17/41 (.41) 4/10 (.40)	2/33 (.06) (no cases)	1/33 (.03) (no cases)	
2			4/12 (.33) —	8/16 (.50) 0/5 (.00)		

Note : *Number in first row of each cell shows proportion of players who bought a signal after corresponding choices of "A" and "B" by predecessors. Number in second row shows proportion of players choosing "A" after seeing signal "b". "A" is the urn more frequently chosen (which can be A or B in the experiment).*

Table 4: Transition matrix for LC

"A" "B"	0	1	2	3	4	5
0	52/75 (.69) —	43/75 (.57) 3/20 (.15)	17/49 (.35) 3/5 (.60)	9/45 (.20) 4/5 (.80)	5/39 (.13) 3/3 (1.00)	4/36 (. 4/4 (1.
1		19/26 (.73) —	16/30 (.53) 2/7 (.29)	7/26 (.27) 1/4 (.25)	6/21 (.29) 2/2 (1.00)	
2			3/10 (.30) —	9/18 (.50) 4/6 (.67)		

Note : *See Table 3.*

Table 5: Transition matrix for HCHI

”A” ”B”	0	1	2	3	4	5
0	42/90 (.47) –	43/90 (.48) 1/18 (.06)	17/56 (.30) 1/6 (.17)	13/45 (.29) 1/6 (.17)	6/34 (.18) 0/2 (.00)	3/30 0/1
1		12/34 (.35) –	18/45 (.40) 1/7 (.14)	14/42 (.33) 1/7 (.14)	10/37 (.27) 0/3 (.00)	
2			4/14 (.29) –	10/23 (.43) 0/6 (.00)		

Note : See Table 3.

Table 6: Transition matrix for NC

”A” ”B”	0	1	2	3	4	
0	80/90 (.89) –	83/90 (.92) 4/23 (.17)	54/60 (.90) 13/20 (.65)	43/49 (.88) 19/21 (.90)	37/46 (.80) 14/20 (.70)	33/4 12/1
1		30/30 (1.00) –	39/41 (.95) 3/19 (.16)	23/27 (.85) 4/9 (.44)	24/28 (.86) 6/8 (.75)	
2			15/17 (.88) –	21/22 (.95) 2/11 (.18)		

Note : See Table 3.