Bonn Econ Discussion Papers

Discussion Paper 15/2004

Herding with and without Payoff Externalities - An Internet Experiment

by

Mathias Drehmann, Jörg Oechssler, Andreas Roider

August 2004

Bonn Graduate School of Economics
Department of Economics
University of Bonn
Adenauerallee 24 - 42
D-53113 Bonn

The Bonn Graduate School of Economics is
sponsored by the

Deutsche Post World Net

MAIL EXPRESS LOGISTICS FINANCE

Herding with and without Payoff Externalities - An Internet Experiment[*]

Mathias Drehmann[†]
Bank of England

Jörg Oechssler[‡]
Department of Economics
University of Bonn

Andreas Roider[§]
Department of Economics
Stanford University and University of Bonn

August 23, 2004

Abstract

Most real world situations which are susceptible to herding are also characterized by direct payoff externalities. Yet, the bulk of the theoretical and experimental literature focuses on pure informational externalities. In this paper we study several different forms of payoff externalities that interact with a standard herding model. More than 6000 subjects, including a subsample of 267 consultants from an international consulting firm, participated in an internet experiment. We also replicate and review earlier cascade experiments. Finally, we study reputation effects in the context of herding.

JEL–classification: C92, D8.

Key words: information cascades, herding, network effects, experiment, internet.

[*] We thank Frauke Lammers, Susanne Niethen, and Nina Wessels for very helpful discussions throughout the project, and McKinsey & Company for financial support. Financial support was also provided by SFB/TR 15, GESY. Finally, we are very grateful to Ben Greiner who did an excellent job programming the experiment.

[†] Mathias.Drehmann@bankofengland.co.uk

[‡] Corresponding author: Department of Economics, University of Bonn, Adenauerallee 24-42, 53113 Bonn, Germany, oechssler@uni-bonn.de

[§] roider@stanford.edu

1 Introduction

Herding behavior seems to be ubiquitous in human decision processes and it can actually be justified as a rational response to uncertainty and informational asymmetries in the environment. Several sources of rational herding have been described in the theoretical literature. Information cascade[1] models pioneered by Bikhchandani, Hirshleifer and Welch (1992), Welch (1992), and Banerjee (1992) show that herding may occur even when individuals' payoffs do not in any way depend on the behavior of others. Externalities are created only through the information that can be deduced from the observed actions. Other models are additionally based on payoff externalities that are widespread in practice. For example, herding of analysts or fund managers in models of reputational herding (e.g., Scharfstein and Stein, 1990), or herd behavior of depositors in bank runs (e.g., Diamond and Dybvig, 1983) may be explained by such models.

The purpose of the current paper is threefold. First, the paper presents a broad–scaled replication of existing laboratory studies on information cascade models. Second, the paper aims to extend this literature by experimentally studying various settings in which payoff externalities are present. And third, the paper investigates the importance of reputation for cascade behavior. In other words, this is a study of the influence of role models on herd behavior.

With respect to replication, our experiment differs from earlier cascade experiments in that (1) we replicate those experiments with more than 6000 subjects, many more than usual. (2) The large number of subjects allows us to test a number of variations, which may potentially be important (e.g., longer sequences of decisions). (3) Instead of the usual undergraduate student population, we use a diverse subject pool, including 267 consultants from an international consulting firm. More than 40% of our subjects hold a Ph.D. or are currently enrolled in a Ph.D. program. A majority of subjects have a background in the natural sciences. (4) Finally, we deviate from the usual laboratory setting by utilizing the internet for our experiment.[2]

[1] An *information cascade* is said to occur when it becomes rational to ignore one's own private information and instead follow the predecessors' decisions. Since no further information is revealed once an information cascade has started, inefficiencies occur even though each individual is behaving rationally.

[2] Arguably, for many people who buy and sell goods on the internet, use internet banking and brokerage services etc., the internet is probably by now a very natural setting for decision making. Nevertheless, conducting experiments on the internet is still novel. For experiments that have been conducted over the internet, see e.g., Forsythe et al. (1992, 1999) , Lucking-Reiley (1999), Anderhub et al. (2001), Charness et al. (2001), Shavit et al. (2001), Bosch–Domènech et al. (2002), and Güth et al. (2002). For technical issues, see e.g., Greiner et al. (2002). The internet is also used to provide a platform to run economic experiments for interactive learning (Holt, 2002).

The second purpose of our paper is to study payoff externalities. The bulk of previous experimental studies has focused on settings without payoff externalities. However, many, if not most, real world examples of herding have a payoff externality component. We will consider positive payoff externalities, which should reinforce herding, as well as negative externalities, which should slow down herding. Since we assume sequential decision making, one can further differentiate between externalities that apply only to predecessors, only to followers, or to both, but possibly in different ways.

There are many real world examples for these kinds of payoff externalities. For example, the choice of software or computers is often described as a situation with positive externalities. The more users adopt the same system as oneself, the easier becomes the interaction with them. For the network externalities to materialize it does not matter (much) whether other users have already adopted the system or whether they will do so soon. Timing may be important, however, in other situations. Consider the choice of a research area. If one aims at maximizing the number of citations, one should enter a new research area as citations can obviously only be directed at older papers. A new research area may turn out to be a dead end, though.

Negative externalities may be caused by overcrowding (e.g., in restaurants, supermarket check–out counters, parking lots etc. where one's utility decreases with the number of predecessors who chose the same restaurant, cashier, or parking lot but where one is not bothered by people who arrive later). Finally, there are situations where one is punished for taking the same action as a predecessor but is rewarded for successors. Avantgardists and fashion leaders fall into this category as well as the unlucky participants in snowball systems or chain letters.

Finally, the third purpose of our paper is to study the role of reputation for herding or cascade behavior. We do this by telling subjects the cumulative payoffs their predecessors have achieved in earlier unrelated rounds. The hypothesis is that subjects are influenced primarily by the subject with the highest earlier payoff despite the fact that this information is irrelevant in a rational Bayesian model.

In Section 2 we describe the basic experimental settings. Despite the large literature on network economics (see e.g. Shy, 2001, for a textbook treatment), there is a very small literature on the interplay between information cascades and network effects. Two exceptions are Choi (1997) and Frisell (2003) who, however, both deal with versions of herding models that differ in some important aspects from the Bikhchandani et al. (1992) model. A combination of the Bikhchandani et al. model with payoff externalities does not seem to be have been treated theoretically.[3] Therefore, in Section 3 we derive the theoretical predictions for the various treatments. In this context it is interesting that we find a multitude of equilibria in the treatment with network effects whereas there do not seem to exist pure strategy equilibria in treatments where only the followers cause positive payoff externalities.

In Section 4 we describe the experimental procedures in detail. Since internet experiments are still relatively new, we explain how we resolved the issues of recruitment, payment of subjects, and the implementation on the internet. Section 5 contains the results. The first part in this section deals with replication of earlier cascade experiments. Besides presenting our own results, we review and compare results from 10 experiments that are scattered in the literature. The second part presents the results from the treatments with payoff externalities. And finally, the third part deals with reputation effects in the basic BHW model. Section 6 concludes. Instructions of the experiment are contained in an Appendix.

2 The experiment

Subjects had to choose sequentially between two "investment opportunities" A and B. Only one of the two could be successful and, if so, would pay 10 "Lotto–Euros". The unsuccessful alternative paid nothing. Subjects were told the a priori probability that investment A was successful, $P(A) = 0.55$ (and consequently, $P(B) = 0.45$). Furthermore, they were told that they would receive a tip by an investment banker which was reliable with probability $P(a|A) = P(b|B) = 0.6$. Sessions with these probabilities are denoted by 55-60. In some treatments we conducted additional sessions with the probability combination 50-66.[4]

Subjects were informed that all prior subjects in their group had received a tip by other investment bankers and that these tips were independent of theirs (see the Appendix for a translation of the instructions). Subjects were able to observe the decisions of their predecessors but, in general, not their signals.

We consider two principal versions of this model:[5] one in which the payoff depends exclusively on

[3] There is one experimental paper with payoff externalities in a Bikhchandani et al. (1992) framework we are aware of. Hung and Plott (2001) study treatments in which subjects are rewarded if a majority of decisions was correct or if their respective action agreed with the majority, respectively. The externalities in our experiment are, however, of a different form.

[4] The probability combination 50-66 was most often used in the literature (see e.g., Anderson and Holt, 1997). However, it has the disadvantage of requiring a tie–braking assumption in many cases.

[5] In a companion paper we focus on treatments with market prices for the investment opportunities A and B (as in

one's own decision. This version is equivalent to the basic model studied by Bikhchandani et al. (1992) and is denoted by *BHW*. For comparison, we also include a treatment *BHW+AS* in which additional to predecessors' actions also all their signals were observable. There is also a "reputation" treatment *BHW+R,* which we will be described in more detail in Section 5.3.

In the second version of the above model we introduce four different forms of payoff externalities, i.e., we consider treatments in which payoffs also depend directly (positively or negatively) on the decisions of others. Table 1 lists the main features of all treatments.

Table 1: Treatments

treatment	description	# of groups
BHW	Bikhchandani/Hishleifer/Welch	63/12/15*
BHW+AS	*BHW* + all signals observable	70/12/9*
BHW+R	*BHW* + cumulative payoffs of predecessors observable (reputation)	29
Network	*BHW* + receive x for each group member who chooses same alternative	12/6**
Follower	*BHW* + receive x for each follower who chooses same alternative	26
Early bird	*BHW* + pay x for each predecessor who chose same alternative	26
Hipster	*BHW* + pay (receive) x for each predecessor (follower) who chose same alternative	12/6**

Note: * x/y/z denotes x groups with probability combination 55-60, y groups with 50-66, and z groups with consultants (also 55-60). In treatments *BHW+R*, *Follower* and *Early bird* the probability combination is 55-60; in treatments *Network* and *Hipster* x is either 0.4 or 1; in treatments *Follower* and *Early bird* x is always 0.4. ** denotes that there were 12 groups with $x = 0.4$ and 6 with $x = 1$.

In treatment *Network* subjects receive an amount x for each other subject in their group that chooses the same action. This payoff structure is supposed to model network externalities. Examples are the choice of software or mobile phone operators where it is not only important to choose the best product but also to choose the product that is chosen by the majority of other consumers because the utility from such a product is increasing in the number of adopters.

In treatment *Followers* subjects receive an amount x only for those subjects that decide later in their group and choose the same action. Examples for such one–sided network externalities are choices on software that is only upwards compatible or the choice of a research topic by a scientist who is

Avery and Zemsky, 1998). Those treatments were conducted in the same experiment (see Drehmann, Oechssler, and Roider, 2004).

concerned about the number of citations to his work. Clearly there can be no citations from papers that have already been published.

In treatment *Early birds* subjects have to pay x for each predecessors who chose the same action as they. This kind of payoff externality is typical for situations where overcrowding is an issue as with restaurants, movie theaters, beaches etc. But also bank runs fall into this category.

Finally, treatment *Hipster* is a combination of *Follower* and *Early bird* as subjects receive x for each follower who chooses the same action but have to pay x for each predecessor with the same action. Examples include fashion leaders, avantgardists, and the participants in snowball systems or chain letters.

3 Theoretical predictions

Despite the simplicity of our experimental setting, theoretical predictions are surprisingly difficult to make in some of our treatments. Table 2 presents a non–exhaustive list of candidates for pure strategy (perfect Bayesian) equilibria given probability combination 55-60.

Table 2: Candidate equilibria

candidate	first player's strategy	strategies of players 2 through 20
BHW	follow own signal	A if $\Delta \geq 1$; B if $\Delta \leq -2$; otherwise follow own signal
uniform	follow own signal	follow action of player 1
reverse	follow own signal	choose opposite of player 1
stubborn	choose A	choose A

Note: Δ denotes the net number of a signals ($\#a$ signals $-$ $\#b$ signals) that can be imputed from the actions of predecessors; in treatment *BHW+AS,* Δ denotes the net number of directly observed a signals.

In treatments *BHW* and *BHW+AS* there is a unique perfect Bayesian equilibrium (see Bikhchandani et al. 1992, or Drehmann, Oechssler, and Roider, 2004), which depends in a simple way on the net number of signals Δ that can be imputed from the actions of predecessors and the own signal. We call this the *BHW* equilibrium. It can easily be checked that this equilibrium does not exist for any of the treatments with payoff externalities.

Since in treatment *Early bird* payoffs depend only on the actions of predecessors and the own action, the game can be solved by backward induction. It turns out that cascades happen but they are endogenously broken once sufficiently many predecessors chose the same action. From this point on, actions may reveal signals again, which may, in turn lead to a new cascade. In comparison to the *BHW* treatment, where cascades once started last until the end of the group, we should see shorter cascades in treatment *Early bird*.

Treatment *Network* allows for a multiplicity of equilibria. All of the candidates *uniform, reverse,* and *stubborn* can be supported as a perfect Bayesian equilibrium with suitably chosen off–equilibrium beliefs. More complex equilibria, in which players 2 through 20 act differently, also exist. Finally, in treatments *Follower* and *Hipster* none of the candidates listed in Table 2 are a Nash equilibrium and we conjecture that no pure strategy equilibrium exists. As an example consider the *uniform* equilibrium candidate in treatment *Follower*. If the first players receives and follows a b signal, all subsequent players are supposed to play B. However, the last player, who does not have any followers, wants to deviate if he receives an a signal because his a signal and the first player's b signal cancel and we are back to the a priori probability, which with 0.55 is in favor of A.

However, note that if one assumes that players are myopic, i.e., that they ignore all future actions, treatment *Follower* yields the same prediction as *BHW*. Likewise, *Early bird* and *Hipster* become indistinguishable from each other. Below we shall also test this prediction.

4 Experimental design

More than 6000 subjects participated in our online experiment which was available for a period of about six weeks in the spring of 2002 on our web site *http://www.a-oder-b.de* which is German for *a-or-b*. Subjects decided in sequence and were able to observe the actual decisions of prior participants in their respective groups. In general, the group size was 20.[6] Subjects were asked to make decisions in three independent groups, thus in total there were more than 18000 decisions. We call the first decision stage 1, the second stage 2, etc.

Payoffs in "Lotto–Euros" were calculated as follows. If a subject chose the correct investment, he received 10 Lotto–Euros. This was the final payoff for this task in the *BHW* treatments. In the treatments with payoff externalities, once all subjects in the respective group had decided, the payoffs of the subjects were raised or lowered by the amount of the respective payoff externalities. In treatments

[6] Except in two cases: in the general subject pool, the group length in treatment *BHW+AS* was 10; with the consultants, the average group length in treatment *BHW* (*BHW+AS*) was 7 (8).

with negative externalities (*Early bird* and *Hipster*) subjects additionally received an endowment of 5 (if $x = 0.4$) or 10 (if $x = 1$) for each task to avoid losses since negative payments are obviously very difficult to enforce in any experiment, let alone an internet experiment.

4.1 Recruiting and payment

The experiment was announced in several ads in the science section of the largest German weakly newspaper *Die Zeit,* two popular science magazines, and two national student magazines. Posters were distributed at most sciences faculties at German universities. Finally, emails were sent to Ph.D. students and postdocs in science and economics departments at 35 universities in Germany. The web site *www.a-oder-b.de* was linked to the Laboratory for Experimental Research in Economics at the University of Bonn and to the sponsor McKinsey & Company to demonstrate that the experiment had a proper scientific background and that the promised financial rewards were credible.

All payoffs in the experiment were denoted in "Lotto–Euro". Each Lotto–Euro was a ticket in a lottery to win one of our main prizes. In total there were 11 prizes of 1000 Euro each. The odds in those lotteries were fixed in advance and known to subjects. Each subject, when logging in on our website was told explicitly the odds per lottery ticket for winning one of our main prizes. Thus, maximizing the probability of winning one of the prizes was equivalent to maximizing the number of lottery tickets. All winners were notified by mail, and their prize money was paid through bank transfers.

In Phase I of the experiment, 1409 subjects played with high powered incentives where each of 40000 lottery tickets had an equal chance of winning one of 5 prizes of 1000 Euros. Since subjects played on average for about 15 minutes, they were making an expected hourly "wage" of 14.19 Euros, which is comparable to a very good student job and to pay in laboratory experiments. In Phase II, each of 90000 lottery tickets had an equal chance of winning one of another 5 prizes of 1000 Euros. Finally, in Phase III, 1162 subjects competed for the remaining 1000 Euros. Only in this Phase III of the experiment, where almost no monetary incentives were provided, subjects did not know how many lottery tickets were issued in the respective phase. This payment scheme was due to the fact that an unexpected large number of subjects participated in our experiment. But it also gives us the chance to test the role of incentives in such a setting.

Additionally, there was a control group of 267 consultants from McKinsey & Company, an international consulting firm, who participated in the experiment on the same web site a couple of weeks before the start of the actual experiment. The subjects of the control experiment were recruited by an internal email to all German McKinsey consultants. Subjects knew that all other subjects were also consultants. About a third of those addressed participated. These subjects had the chance to win 8 vouchers for a nice dinner for two in a restaurant each worth 150 Euros.

4.2 Subject pool

In total, 6099 subjects finished our experiment of which 5832 subjects participated in the main experiment and 267 in the control experiment with consultants.[7] Table 3 lists some of the main characteristics of the combined subject pool (including the control experiment with consultants).

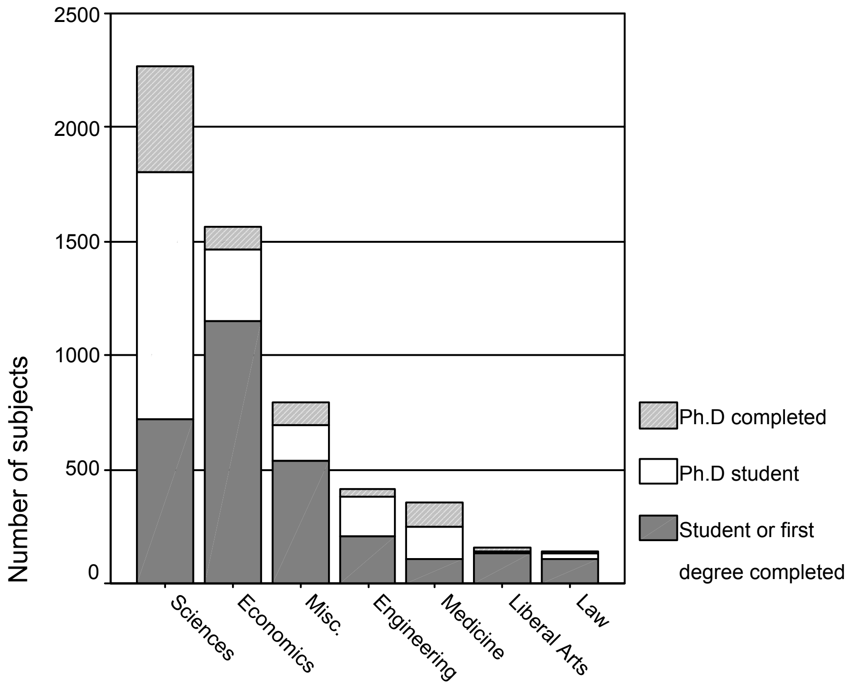

Figure 1: Composition of the subject pool. (Note: "Sciences" includes physics, chemistry, mathematics, and computer science; "Economics" includes economics, business administration and related subjects; "Medicine" includes medicine, psychology, and dentistry; "Liberal Arts" includes all languages, history, and pedagogy. "Misc." stands for miscellaneous fields.)

In contrast to most experiments in economics, our subjects come from a broad range of fields. Figure 1 shows the frequencies of the main subject groups. Each bar in Figure 1 shows the number of subjects

[7] 788 individuals logged on but did not finish the experiment. Their decisions were not included in the history H_t since they did not face monetary incentives (payment was conditional on finishing all three stages of the experiment).

Table 3: Properties of the subject pool

Average age	28.3
% of female subjects	27.8
% completed (at least) first university degree	56.9
% current students	36.4
% non–students	6.7
% completed Ph.D.	13.7
% current Ph.D. students	31.3

who study for or have finished a first degree, the number of subjects who currently are Ph.D. students, and the number of subjects who have finished a Ph.D.[8] Considering the number of Ph.D. students and Ph.D.'s we believe we succeeded in recruiting a fairly bright subject pool.

4.3 Implementation

When arriving on our web site, subjects read a screen that introduced the general problem and the rules of the game. Subsequently, subjects were asked for some personal information, like name, mailing address, email, field of study, age etc., and subjects were only allowed to play if all information requested was actually provided. This was also a measure to prevent subjects from playing twice: in order to win in the lottery, one had to give a correct mailing address, and the program ensured that the same name-postal code combination as well as the same email address could only play once. We also used cookies to prevent using the same computer twice.[9]

After entering the personal information, subjects were randomly placed in a currently active group,[10] and had to make their first decision. Afterwards they were randomly placed in another active group for the second task and then in a third group for the final task. No feedback about results was given until the subject had completed all three tasks, and even then they were only told how many "Lotto–Euro" they had won. Usually the tasks for each subject came from different treatments. Finally, we asked subjects for voluntary feedback as to how they formed their decision, and 687 subjects sent response emails. The last column of Table 1 lists the number of groups that participated in our experiment, separately

[8] Given that each time that we sent out emails to Ph.D. students and post-docs to advertise the experiment, there was immediately a peak in access to our webpage, one can be confident in these numbers.

[9] It will never be possible to completely prevent clever people from playing more than once. However, we are confident that not many such attempts were successful, and given the size of the subject pool, those few probably do not matter much.

[10] A group was *active* when it was neither full nor closed (i.e., when another subject was active in this group). We also ensured that subjects who logged on at about the same time were allocated to different treatments to prevent "observational learning" in case two subjects sat next to each other in a computer pool.

for each combination of treatments, probabilities, and whether subjects came from the general subject pool or the control experiment with the consultants.

5 Results

5.1 Replication

To make our results comparable to earlier experimental studies we shall concentrate on the following three measures. (1) Average rationality under common knowledge of rationality (*ruck*), which is defined as the fraction of subjects who behaved according to a Perfect Bayesian equilibrium under the assumption that all predecessors are commonly known to be Bayesians.[11] (2) The fraction of cases in which subjects rationally decided against their own signal if they are in a cascade is denoted by *casc*. Arguably, *casc* is a harder test for cascade theories since *ruck* includes all the cases in which subjects (rationally) follow their own signal. (3) The fraction of cases in which subjects followed their own signal is denoted by *own*. For comparison, we also report the equilibrium value of *own*, denoted by *own**, that would have obtained had all subject behaved according to *ruck* (as defined above).

Table 4: **BHW treatments**

subject pool	treatment	prob.	all subjects				subjects on pot. eq. path			
		comb.	*ruck*	*casc*	*own*	*own**	*ruck*	*casc*	*own*	*own**
general	*BHW*	55-60	.66	.34	.75	.59	.86	.74	.74	.90
	BHW+AS	55-60	.72	.41	.74	.68	-	-	-	-
consultants	*BHW*	55-60	.68	.16	.85	.66	.90	1.00	.83	.93
	BHW+AS	55-60	.78	.52	.69	.60	-	-	-	-
general	*BHW*	50-66	.78	.45	.75	.62	.95	.78	.84	.77
	BHW+AS	50-66	.76	.59	.69	.69	-	-	-	-

Note: The average length of potential equilibrium paths is 6 in treatment *BHW* 50-66, and 3 in the two remaining cases; as in treatment *BHW+AS* signals of predecessors were public information, we do not differentiate in this case whether or not a subject observed a potential equilibrium history.

Table 4 lists those measure for our *BHW* and *BHW+AS* treatments, and for our subsample with consultants (who also played treatments *BHW* and *BHW+AS*). First, note that the phase of the experiment (recall that different incentives were provided in different phases) does not have a significant influence on those results according to MWU-tests. Neither does the stage of the task (whether a task was first, second or third) matter. Table 4 lists the above defined measures for all subjects and for

[11] For a decision which, given the history of imputed signals, obviously violates Perfect Bayesian equilibrium, we let players assume that the deviator followed his private signal.

Table 5: Previous BHW cascade experiments

study	treatment	prob.	group size	*ruck*	*casc*
Alevy et al. (2003)	symmetric, students	50-66	5	.95	.89
Anderson/Holt (1997)	symm., no public sig.	50-66	6	.92	.73
Anderson (2001)	2$	50-66	6		.70
Cipriani/Guarino (2004)	fixed price	50-70	12	.83	
Goeree et al. (2004)	$T = 20$	50-66	20		.64
Hung/Plott (2001)	individualistic	50-66	10		.77
Kübler/Weizsäcker (2004)	NC	50-66	6		.78
Oberhammer/Stiehler (2001)		50-66	6	.86	.73
Willinger/Ziegelmeyer (1998)	treatment 1	50-60	6		.64
Ziegelmeyer et al. (2002)	blue line, exp. 1&2	55-66	9		.69

Note: Only studies that implement the BHW model are included. In some cases, values for ruck or casc could not be determined from the information given in the respective papers. Also, in some cases it was unclear whether all observations were counted or only those on a potential equilibrium path. The probability combination (prob.) is given as x-y, which denotes an a priori probability for A of $x\%$ and a signal precision of $y\%$.

those that are on a *potential equilibrium path.* We say that subjects are on a potential equilibrium path as long as there is no prior decision that obviously violates behavior in a Perfect Bayesian equilibrium from the viewpoint of a player who cannot observe the private signals of predecessors.

First, a look at all subjects shows that from a theoretical perspective subjects rely too heavily on their own private signal as *own* is weakly above *own** in all cases. As a result, in our main treatment, *BHW*, 55-66, subjects act in accordance with theory in only 66% of cases. Even more dramatic is the picture with respect to *casc.* Only in 34% of cases did subjects decide against their signal but in accordance with Bayesian updating. Those numbers are lower than those found previously in most of the literature. In the following, we provide a brief overview over the earlier experiments on information cascades and offer some explanations for the observed behavioral differences.

Table 5 lists the results of all cascade experiments implementing the basic setup of BHW that we could find in the literature.[12] While the experiments differ with respect to a number of design issues, most notable the number of players in a sequence and the probability combinations, most values of *ruck* and *casc* are roughly comparable and are substantially higher than those in our experiment.

What could account for those differences? One possible explanation may be that decisions are more difficult on average when 20 subjects decide in sequence rather than the usual 6.[13] To test for this we

[12] We thank Lisa Anderson and Charlie Holt for kindly providing their data.

[13] Huck and Oechssler (2001) show that *ruck* values are substantially lower when decisions are more complex.

look at the decisions of our first 6 subjects in each group. And indeed, for the first 6 subjects *ruck* is 82%, which is closer to the numbers found in the literature. Additional support for this hypothesis is provided by the results of Goeree et al. (2004). They also consider sequences of 20 subjects and report one of the lowest values for *casc* (see Table 5). Another aspect emerges when we consider subjects on a potential equilibrium path (right panel of Table 4). On average, potential equilibrium paths have length 6. On those paths, values of *ruck* and *casc* are very high (and comparable to those reported in Table 5), which indicates that subjects become confused as soon as they observe deviations from a potential equilibrium path. Interestingly, on potential equilibrium paths, subjects do rely less often on their private information than predicted by theory under probability combination 55-60.

A second possible explanation is that subjects simply mistrust the behavior of their predecessors on the internet so much that they rather rely on their own signals. However, this consideration should not matter in treatment *BHW+AS* where all signals of predecessors were observable and the payoff maximizing decision is simply a matter of forming conditional expectations. Yet, the measures *ruck*, *own,* and *casc* are not substantially better (even though for probability combination 55-60 both for the general subject pool and the consultants *ruck* and *casc* are significantly higher in *BHW+AS* at the 1% level according to MWU–tests). Also, in the control experiment with consultants, where all participants had a relatively good idea about the types of their predecessors, the reliance on the own signal is even more pronounced,[14] and the number for *casc* is substantially lower compared to the general subject pool. It seems that consultants are more reluctant to rely on the decisions of others. Interestingly, Alevy et al. (2004) also find in their experiment that professional traders have lower *ruck* and *casc* values than college students.

A third possible explanation is that the probability combination 55-60 (with asymmetric prior) is more difficult than the (symmetric) combination 50-66, which was often used in the literature. For example, a subjects with a b signal on the second position should already ignore his signal if the first subject chose A for 55-60 but not for 50-66. This hypothesis is supported by the significantly higher numbers for *ruck* (78%) and *casc* (45%) in 50-66.[15]

Finally, the level of payoffs can play a decisive role in complex decision problems. For example, Anderson (2001) shows that errors decrease substantially when the payoff for a correct decision is

[14] This cannot be explained by a higher *own* in equilibrium. For consultants, given the random draw of signals, equilibrium *own* would have been 0.66 whereas for *BHW* 55-66 it would have been 0.59.

[15] With respect to *ruck* (*casc*) the difference is significant at the 1% (10%) level.

increased from 0 to 2$. In our experiment subjects earned about 1.25 Euros for a correct decision in phase I and 0.55 Euros in phase II.[16] While we do not observe a significant difference in *ruck* between phase 1 and 2, it is possible that the lower payoffs in combination with the first and the third explanation above are responsible for the values observed in Table 4.

It is also interesting to test whether different subject characteristics influence *ruck* and *casc*. However, we find no significant difference between male and female subjects or between subjects holding a Ph.D., Ph.D. students, or others. In treatment *BHW* the McKinsey consultants differ from the general subject pool by showing significantly higher values for *own* and lower values for *casc* (at the 1% respectively 5% level according to MWU-tests).

5.2 Payoff externalities

Figure 2 presents a first view on how the various payoff externalities influence behavior.

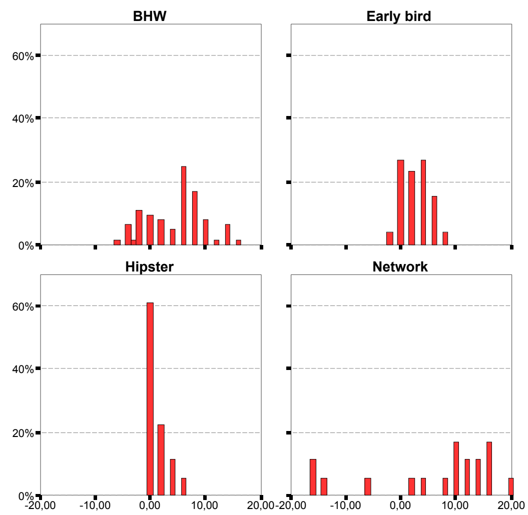

Figure 2: Distribution of decision imbalances after the last period (pooled over x)

We call the difference between the number of A and B decisions the "decision imbalance". Figure 2 shows the distribution of decision imbalances after the last player. Treatment *Network* clearly stands out as the only treatment in which extreme imbalances occur. On the other hand, treatments *Early bird*

[16] Half the groups in *BHW* 50-66 were played in phase I and half in phase II.

and *Hipster* produce very balanced distributions centered around 0.[17] Kolmogorov–Smirnov tests reveal that the distribution for *Network* is significantly different from all other treatments except *Follower* (at the 5% level or better). Recall that if subjects are myopic, there should not be any difference between *Early bird* and *Hipster,* and between *BHW* and *Follower.* And indeed, there are no significant differences in the distributions of decision imbalances for those two pairs. In all other pair-wise comparison one cannot reject that the distributions are significantly different (at the 5% level or better).

A second indicator is the number and length of runs in the data. A run is a sequence of consecutive subjects who made the same decision. If there are positive payoff externalities (as in *Network*) we would expect longer (and therefore fewer) runs. With negative externalities, runs should be shorter and more frequent. Table 6 lists the average number and length of runs per group for our treatments (separate for A and B runs).[18] As expected, *Network* has the lowest number and highest average length of runs. Again in accordance with myopia, *BHW* and *Follower* seem to show runs of similar (medium) length and frequency. The shortest and most frequent runs are found for *Early bird* and *Hipster.* The fact that B runs are shorter on average in all but one case might be explained by the higher a priori probability for A.

Table 6: Number and length of runs

treatment	x	number of runs	average length A runs	average length B runs
Network	0.4	8.55	3.51	1.39
Network	1	4.50	4.38	4.50
Follower	0.4	9.58	2.62	1.53
Early bird	0.4	10.38	2.14	1.70
Hipster	0.4	11.92	1.79	1.56
Hipster	1	12.33	1.68	1.57
BHW	-	9.84	2.45	1.59

Note: Probability combination 55-60; general subject pool.

A third interesting aspect of the data is predictability. For example, is it possible to predict early on which product will capture a larger slice of a market? Hence, we ask whether one can forecast the majority decision in a group after observing the first n players. Table 7 shows correlations between the sign of the decision imbalance after player $n = 2, 5, 10, 15$ and the sign of the decision imbalance after player 20. Note that a decision imbalance is positive if a majority of subjects chose A, and vice versa.

[17] The decision imbalance is exactly 0 in 26.9% of cases in *Early bird* and in 61.1% of cases in Hipster, whereas the same holds only in 9.2% of cases for *BHW,* in 7.7% for *Follower*, and in 0% for *Network*.

[18] Treatment *BHW+AS* is excluded since all groups in this treatment consisted only of 10 subjects.

Treatment *Network* with $x = 1$ shows the highest predictability. Already after the second player the correlation is 0.86 and significant at the 5% level. *Follower* and *Network* with $x = 0.4$ also show high correlations. In contrast, for *Hipster* correlations of early rounds with the final imbalance is sometimes negative. Even very late in the game, the final outcome cannot be reliably predicted.

Table 7: Predictability of majority choice

		correlation between the sign of the decision imbalance after player 20 and after player...			
treatment	x	2	5	10	15
Network	0.4	0.28	0.52*	0.67**	1.0***
Network	1	0.86**	0.93***	0.93***	0.93***
Follower	0.4	0.21	0.71***	0.91***	0.80***
Early bird	0.4	0.20	0.32	0.71***	0.65***
Hipster	0.4	−0.30	−0.71**	0.13	0.71**
Hipster	1	−1.0*	0.32	−0.32	0.45
BHW	-	0.08	0.23*	0.61***	0.75***

Note: Probability combination 55-60; general subject pool; *** significant at 1%-level; ** significant at 5%-level; * significant at 10%-level.

Given the multiplicity of equilibria for treatment *Network* it is interesting which, if any, of those equilibria can be observed in the data. We classify the decisions of a group of 20 subjects as in accordance with an equilibrium if at most 2 (for $x = 1$) or at most 4 (for $x = 0.4$) subjects deviate from the equilibrium path. In this sense, 5 of the 6 groups for $x = 1$, can be classified as one of the equilibria listed in Table 2. In particular, there are 3 groups that can be classified as playing a *uniform* equilibrium, 1 as *uniform* or *stubborn*,[19] and 1 group even seems to be playing the slightly strange *reverse* equilibrium. Of the 12 groups with $x = 0.4$, 6 groups can be classified as one of the listed equilibria, namely, 2 as *stubborn*, and 4 as *uniform* or *stubborn*.

5.3 Reputation effects

Treatment *BHW+R* was the same as *BHW* except that subjects were able to observe not only the actions of their predecessors but also the cumulative payoffs from the two decisions those subjects made in the first two (unrelated) stages of the experiment.[20] In a rational Bayesian model, this extra information is irrelevant. However, we suspected that subjects would rely more on the decision of

[19] If the first subjects receives an a signal, the two equilibria are indistinguishable.

[20] Recall that each subjects had to make three decisions (stage 1 through 3). *BHW+R* was always played on stage 3.

predecessors with the highest payoffs (the "success models"). That is, subjects with higher payoffs have a better reputation and are imitated more often.

Figure 3 shows that subjects are indeed influenced by the decision of the predecessor with the highest reputation (i.e., the highest cumulative payoff in stages 1 and 2). Regardless of the own signal, an A decision by this predecessor significantly increases the frequency of A according to MWU–tests, at the 5% level.

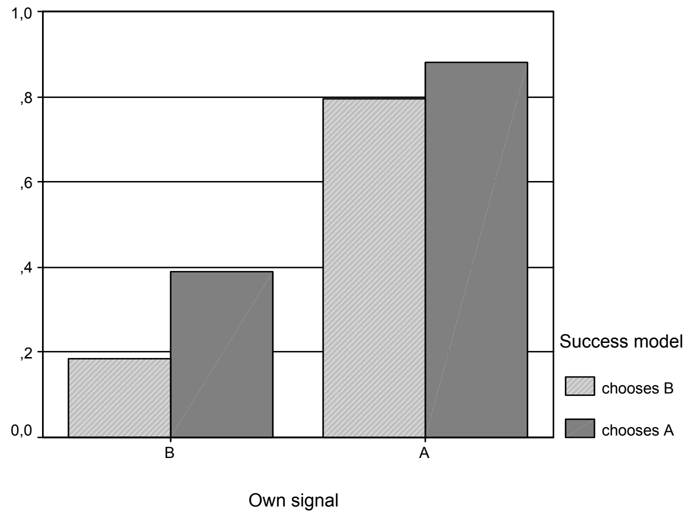

Figure 3: Fraction of subjects choosing A depending on the choice of the most successful predecessor (and own signal). (Note: Including only subjects that had at least one predecessor and including only cases when the predecessor with the highest reputation was unique (547 out of 551 cases).)

This result is supported by a logit regression where the probability of subject t choosing A is explained by (1) subject t's signal, (2) a dummy variable da_{t-1} that is equal to one if the history Δ_{t-1} (expressed by the net number of imputed a signals) is such that the subject should choose A irrespective of his signal and that is equal to zero otherwise, (3) a dummy variable db_{t-1} that is equal to one if, given Δ_{t-1}, the subject should choose B irrespective of his signal and that is equal to zero otherwise, (4) the decision of the predecessor with the highest payoff, Rep_{t-1}, and (5) a constant. Table 8 shows that, even though the private signal and the history should summarize all relevant information, the coefficient of Rep_{t-1} is positive and significant at the 1% level, which demonstrates that subjects seem to follow success models.

From an ex-post perspective, did it make sense to follow the respective success model? On average

Table 8: Logit analysis: choice of A in treatment *BHW+AS*

	signal	da_{t-1}	db_{t-1}	Rep_{t-1}
coefficient	2.571**	−.055	.069	.457**
(standard deviation)	(.221)	(.233)	(.344)	(.125)

Note: 547 observations (only subjects with at least one predecessor and a unique highest-reputation predecessor are included); -2 Log-Likelihood = 555.95; R^2 (Nagelkerkes) = 0.383; ** significant at 1%-level; regression includes a constant.

subjects chose the successful alternative in 56% of cases. However, success models did so in 62% of cases, so that these subjects were indeed (somewhat) more successful in picking the right alternative.

6 Conclusion

In a large–scale internet experiment we investigated information cascade models with and without payoff externalities. For the base treatment without payoff externalities we found a substantially lower percentage of subjects who behaved according to theory as compared to earlier results in the literature. We explained this deviation through a combination of the experimental setting (internet vs. lab), the probability combination (asymmetric prior vs. symmetric prior), the number of subjects deciding in sequence (20 vs. 6), and the level of payoffs. While our results do not question the fact that information cascades do happen in experiments, they certainly show that cascades – depending on the setting – may be rarer and shorter than predicted by theory and suggested by earlier experiments.

Surprisingly, there is only a very small literature on the interplay between information cascades and payoff externalities, either theoretical (see e.g. Choi, 1997; Frisell, 2003) or experimental (e.g. Hung and Plott, 2001). We studied several different forms of payoff externalities, positive and negative ones and those that apply to all group members or only to predecessors or followers. The experimental results are by and large compatible with the theoretical predictions. With positive externalities (network effects) cascades become longer and more robust, whereas with negative externalities they become short and fragile. In most cases we could not reject the hypothesis that subjects behaved myopically as treatments that have the same theoretical solution under myopia yield very similar results. The form of payoff externality was found to have strong effects also on the predictability of the majority decision. With strong network effects, already after the second player (of 20) the majority decision can reliably be predicted.

Finally, this experiment was designed to test reputation effects in the framework of a cascade model. Reputation of a player was presented as the cumulative payoff the player earned in previous and unrelated rounds. Subjects could observe these payoffs, and we found strong support for the hypothesis that the decision of the player with the highest reputation significantly influences the choice behavior of later subjects.

References

Alevy, E. Jonathan, Michael S. Haigh, and A. John List, "Information Cascades: Evidence from a Field Experiment with Market Professionals," *mimeo: University of Maryland*, 2003.

Anderhub, V., R. Müller, and C. Schmidt, "Design and Evaluation of an Economic Experiment Via the Internet," *Journal of Economic Behaviour and Organization*, 2001, *46(2)*, 227–47.

Anderson, L. R., "Payoff Effects in Information Cascade Experiments," *Economic Inquiry*, 2001, *39(4)*, 609–15.

Banerjee, A. V., "A Simple Model of Herd Behavior," *Quarterly Journal of Economics*, 1992, *107(3)*, 797–817.

Bikhchandani, S., D. Hirshleifer, and I. Welch, "A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades," *Journal of Political Economy*, 1992, *100(5)*, 992–1026.

Bosch-Domenech, A., J. Montalvo, R. Nagel, and A. Satorra, "One, Two, (Three), Infinity, ...: Newspaper and Lab Beauty-Contest Experiments," *American Economic Review*, 2002, *92(5)*, 1687–1701.

Charness, G., E. Haruvy, and D. Sonsino, "Social Distance and Reciprocity: The Internet Vs. The Laboratory," *mimeo*, 2001.

Choi, Jay Pil, "Herd Behavior, the "Pinguin Effect", and the Suppression of Informational Diffusion: An Analysis of Informational Externalities and Payoff Interdependency," *RAND Journal of Economics*, 1997, *28(3)*, 407–425.

Cipriani, Marco and Antonio Guarino, "Herd Behavior in a Laboratory Financial Market," *mimeo, University College London*, 2004.

Diamond, D. and P. Dybvig, "Bank Runs, Deposit Insurance and Liquidity," *Journal of Political Economy*, 1983, *91(3)*, 401–19.

Drehmann, Mathias, Jörg Oechssler, and Andreas Roider, "Herding and Contrarian Behavior in Financial Markets - An Internet Experiment," *mimeo, University of Bonn*, 2004.

Forsythe, R., F. Nelson, G. Neumann, and J. Wright, "Anatomy of an Experimental Political Stock Market," *American Economic Review*, 1992, *82*, 1142–1161.

___, T. Rietz, and T. Ross, "Wishes, Expectations and Actions: A Survey on Price Formation in Election Stock Markets," *Journal of Economic Behavior and Organization*, 1999, *39*, 83–110.

Frisell, Lars, "On the Interplay of Informational Spillovers and Payoff Externalities," *RAND Journal of Economics*, 2003, *34(3)*, 582–592.

Goeree, Jacob K., Thomas R. Palfrey, Brian W. Rogers, and McKelvey Richard D., "Self-Correcting Information Cascades," *mimeo, California Institute of Technology*, 2004.

Greiner, B., H. Jacobsen, and C. Schmidt, "The Virtual Laboratory Infrastructure for Online Economic Experiments," *mimeo*, 2002.

Güth, W., C. Schmidt, and M. Sutter, "Fairness in the Mail and Opportunism on the Internet - a Newspaper Experiment on Ultimatum Bargaining," *German Economic Review, forthcoming*, 2003.

Holt, Charles, "Webgames and Strategy: Recipes for Interactive Learning," *mimeo*, 2002.

Hung, A. A. and C. R. Plott, "Information Cascades: Replication and an Extension to Majority Rule and Conformity Rewarding Institutions," *American Economic Review*, 2001, *91(5)*, 1508–20.

Kübler, D. and G. Weizsäcker, "Limited Depth of Reasoning and Failure of Cascade Formation in the Laboratory," *Review of Economic Studies*, 2004, *71(2)*, 425–441.

Lucking-Reiley, D., "Using Field Experiments to Test Equivalence Between Auction Formats: Magic on the Internet," *American Economic Review*, 1999, *89(5)*, 1063–1080.

Oberhammer, C. and A. Stiehler, "Does Cascade Behavior in Information Cascades Reflect Bayesian Updating? An Experimental Study," *mimeo*, 2001.

Scharfstein, David and Jeremy Stein, "Herd Behavior and Investment," *American Economic Review*, 1990, *80(3)*, 465–79.

Shavit, T., D. Sonsino, and U. Benzion, "A Comparative Study of Lotteries Evaluation in Class and on the Web," *Journal of Economic Psychology*, 2001, *22*, 483–491.

Shy, Oz, *The Economics of Network Industries*, Cambridge: Cambridge University Press, 2001.

Welch, I., "Sequential Sales, Learning, and Cascades," *Journal of Finance*, 1992, *47(2)*, 695–732.

Willinger, M. and A. Ziegelmeyer, "Are More Informed Agents Able to Shatter Information Cascades in the Lab?," in Cohendet, Llerena, Stahn, and Umbhauer, eds., *The Economics of Networks: Interaction and Behaviours*, Springer, 1998, pp. 291–305.

Ziegelmeyer, Anthony, Juergen Bracht, Frederic Koessler, and Eyal Winter, "Fragility of Information Cascades: An Experimental Study Using Elicited Beliefs," *mimeo: University of Strasbourg*, 2002.

Appendix: Instructions

Once connected to our website *www.a-oder-b.de*, there was first a general overview on the experiment (screen 1 below). Then, subjects where asked to provide some personal information (screen 2 below). Only if all information was provided, subjects were allowed to continue and learn their player number as well as the monetary incentives in the current phase of the experiment (screen 3 below). Note that the number of lottery tickets and the prizes mentioned below relate to Phase I of the experiment. Subsequently, the actual experiment began. Screen 4 below provides an example of the first of three stages (treatment *BHW*), and we point out how these instructions were altered in case of treatments *BHW+AS* and *BHW+R*. Screen 5 below provides an example of a treatment with payoff externalities played in the second stage (treatment *Early bird*). The other treatments with payoff externalities were explained in a similar fashion. As each stage had the same basic structure, we do not provide an example of the third stage.

Subjects also had at all times the option of opening a pop–up window that contained a summary of the main features of the set–up. All phrases emphasized in this translation were also emphasized in the original web page.

Screen 1: Introduction

A game-theoretic experiment Are you a good decision-maker? We challenge you! Professor J. Oechssler together with the "Laboratorium for Experimental Research in Economics" at the University of Bonn aims to test various scientific theories through the online–experiment "A-or-B". Financial support is provided by the consultancy McKinsey & Company.

Attractive prizes By participating in the experiment you support the scientific work of the University of Bonn. At the same time you participate in a lottery for a total of 5,000 Euros which are distributed among 5 of the participants. The more thorough your decisions are, the greater your chances of winning. Of course you will also need some luck. The game takes approximately 15 minutes.

The experiment The experiment consists of three rounds. In every round you'll be assigned to a group and you - as well as every other member of your group - will have to take an investment decision. Without background knowledge the decision would be pure speculation. However, all players in a group will receive tips by investment bankers. Each group member gets a tip from a different investment banker. The investment bankers are experienced but can't make perfect predictions. The reliability of the tip is the same for every investment banker. As additional information, each player can observe the decisions of his predecessors in his group.

For each correct decision you will earn a predetermined amount of Lotto–Euros. After the third round, the Lotto–Euros you earned will be converted into lottery tickets on a one-to-one basis. Hence, the better your investment decisions, the higher your chances of winning. The experiment ends on June 7, 2002. The winners of the lottery will be notified after June 16, 2002 via ordinary mail. Now, let's begin the experiment!

Screen 2: Request of personal information

Welcome to the online-experiment "A-or-B". Please note that you can only play once. Before the game starts, we would like to ask you for some personal information. Of course, the results of the game will be kept separately from your personal information and will be analyzed anonymously. The mail address is only needed to notify the winners. Information on your field of studies, age, sex, etc. are only used for scientific purposes. Detailed information regarding data protection may be found here [Link].

[Data entry fields for last name, first name, address, email, student status, field of studies, year of studies, Ph.D. status, age and sex]

Screen 3: Player number and incentives

Thank you for providing the requested information. Your player number is: [player number]. Your player number, the number of lottery tickets you won, and additional information regarding the experiment will be automatically send to your email address after you have completed the experiment.

In this phase of the experiment, a total of 40,000 lottery tickets will be distributed, and 5 participants can win 1000 Euros each. Every lottery ticket has the same chance of winning.

Screen 4: Stage 1

You have to make an important investment decision: there are two risky assets (A and B). Only *one* asset will be successful and pay out *10 Lotto-Euros* (LE). The other asset will yield no profit at all. The successful asset was determined randomly before the first player of this group played. Hence, *the same asset* is successful for all players in your group. *Without additional information* you can rely on the fact that in *55%* of cases *asset A* is successful while in *45%* of cases *asset B* is successful.

Each participant in your group faces the same problem as you do: he has to choose between the assets and receives a tip from his respective investment banker. The reliability of the tips is the same for all investment bankers, and the tips of the investment bankers are independent of each other. The tip of each investment banker is correct in *60% of the cases*, i.e. in 100 cases where asset A (respectively B) is successful, in 60 cases the investment banker gives the correct tip A (respectively B) while in 40 cases the tip is not correct. The tip of your investment banker is: [B]

While each participant only knows the tip of his own investment banker, you - as every player in your group - can observe the decisions of the respective predecessors. Which players are assigned to which group is random and will differ from round to round. You are the [4th] investor in this group. One after another, your predecessors have made the following decisions:

Investor no.	1	2	3	
Decision	B	A	B	What do you choose? [A] or [B].

Was the decision difficult? Independent of your decision, what do you think is the probability of A being the successful asset? [ ] %.
After the third round you'll find out whether your decision was correct. Let's move on the next round.

[In case of treatment *BHW+AS*, in addition to the decisions also the tips of the predecessors were displayed, and the third paragraph of Screen 4 was replaced by: "You - as every player in your group - can observe the decisions of the respective predecessors and the tips that they have received from their respective investment bankers. Which players are assigned to which group is random and will differ from round to round. You are the [4th] investor in this group. One after another, your predecessors have made the following decisions and have received the following tips:"].

[In case of treatment *BHW+R*, in addition to the decisions also the cumulative payoffs of the predecessors earned in the respective other two stages were displayed, and the third paragraph of Screen 4 was replaced by: "While each participant only knows the tip of his own investment banker, you - as every player in your group - can observe the decisions of the respective predecessors. In addition, each participant can observe how many Lotto–Euros their respective predecessors have earned in their respective other two stages. Which players are assigned to which group is random and will differ from round to round. You are the [4th] investor in this group. One after another, your predecessors have made the following decisions and have earned the following amount of Lotto–Euros on their respective two other stages:"].

Screen 5: Round 2

Another investment decision has to be made. The basic structure remains the same as in round 1. (In case you want to review the central features of round 1 please click [here].) Again, there are two risky assets (A and B). Only *one* asset will be successful and pay out *10 Lotto-Euros* (LE). In *55%* of cases it is *asset A* that is successful. As in the first round the successful asset was determined randomly before the first player of this group played. Hence, it is not necessarily the same asset as in round that is successful.

As in round 1, every participant receives a tip from his investment banker which is correct in 60% of all cases. This time, your investment banker recommends: [A]

In contrast to round 1, each participant has to *pay 0.4 LE for each of his predecessors* in his group who has selected *the same asset as himself* - independent of whether his decision to choose A respectively B turns out to be successful, or not.

Consider the following example: suppose you were the fifth participant in a group and your predecessors had made the choices BABB. If you also would choose B, you would have to make a payment of $3 \times 0.4LE$ because three of your predecessors have chosen B. If you would choose A, you would have to pay $1 \times 0.4LE$.

In order to be able to make these payments you receive an endowment of 11 Lotto-Euro. Once the above payments have been deducted, you can keep the remainder.

While each participant only knows the tip of his own investment banker, you - as every player in your group - can observe the decisions of the respective predecessors. You are the [5th] investor in this group. One after another, your predecessors have made the following decisions:

Investor no.	1	2	3	
Decision	B	A	B	What do you choose? [A] or [B].

Was the decision difficult? Independent of your decision, what do you think is the probability of A being the successful asset? [ ] %.
After the third round you'll find out whether your decision was correct. Let's move on the next round.