Herding with and without Payoff Externalities - An Internet Experiment*

Mathias Drehmann†
Bank of England

Jörg Oechssler
University of Heidelberg

Andreas Roider‡
University of Bonn

February 2005

Abstract

Most real world situations that are susceptible to herding are also characterized by direct payoff externalities. Yet, the bulk of the theoretical and experimental literature on herding has focused on pure informational externalities. In this paper we experimentally investigate the effects of several different forms of payoff externalities (e.g., network effects, first-mover advantage, etc.) in a standard information-based herding model. Our results are based on an internet experiment with more than 6000 subjects, including a subsample of 267 consultants from an international consulting firm. We also replicate and review earlier cascade experiments. Finally, we study reputation effects (i.e., the influence of success models) in the context of herding.

JEL–classification: C92, D8.

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

*We are grateful to Jeremy Bulow, Muriel Niederle, Garth Saloner, Katja Seim, Catherine Tucker, and Jeffrey Zwiebel for helpful conversations, and seminar participants at Stanford for comments. We thank Frauke Lammers, Susanne Niethen, and Nina Wessels for very useful discussions throughout the project, and McKinsey & Company for financial support. Financial support was also provided by SFB/TR 15, GESY. The third author gratefully acknowledges the hospitality of Stanford Graduate School of Business and a DFG research fellowship. Finally, we are very grateful to Ben Greiner who did an excellent job programming the experiment.

†The views expressed in this paper are those of the authors and do not necessarily reflect those of the Bank of England.

‡Corresponding author: Andreas Roider, Department of Economics, Wirtschaftspolitische Abteilung, University of Bonn, Adenauerallee 24-42, 53113 Bonn, Germany, roider@uni-bonn.de.

1 Introduction

Motivation Whether one considers investment decisions, bank runs, fashion markets, or even the choice of restaurants, herding behavior seems to be ubiquitous in human decision processes. Despite the bad reputation that herding behavior sometimes enjoys, it can 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 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.[1] In these models externalities are created only through the information that can be deduced from observed actions.[2] Other explanations of herd behavior are based on payoff externalities, which seem to be 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.

However, in many real world situations both informational externalities as well as payoff externalities seem to be present at the same time. For example, the choice of software or hardware is often described as a situation with network externalities.[3] While one's choice of such a product may very well convey information about its quality to later potential users, it is also the case that the more users adopt the same system, the easier becomes interaction with them. For such positive payoff 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 payoff 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

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

[2] For surveys of this literature see e.g., Bikhchandani, Hirshleifer, and Welch (1998) and Gale (1996), for a generalization see Smith and Sorensen (2000), and finally for a recent textbook treatment see Chamley (2004).

[3] See e.g., Katz and Shapiro (1985), or Church and Gandal (1992). Based on internal documents made public in the antitrust trial, Bresnahan (2004) provides an overview of network effects in the software industry as perceived by Microsoft.

same restaurant, cashier, or parking lot but where one is not bothered by people arriving later). Finally, there are situations where one is punished for taking the same action as a predecessor but is rewarded for successors. Avant-gardist and fashion leaders fall into this category as well as the unlucky participants in snowball systems or chain letters.

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 additionally payoff externalities are present. And third, the paper investigates the importance of reputation for cascade behavior, or in other words, the influence of role models.

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 that 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 has a background in the natural sciences. (4) Finally, we deviate from the usual laboratory setting by utilizing the internet for our experiment.[4]

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

Finally, our paper aims to study the effect of subjects' reputation on the (herd) behavior of

[4] 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, Müller, and Schmidt (2001), Charness, Haruvy, and Sonsino (2001), Shavit, Sonsino, and Benzion (2001), Bosch-Domenech, Montalvo, Nagel, and Satorra (2002), and Güth, Schmidt, and Sutter (2003). For technical issues, see e.g., Greiner, Jacobsen, and Schmidt (2002). The internet has also been used to provide a platform to run economic experiments for interactive learning (see e.g., Holt, 2002).

later decision-makers. We do this by informing subjects of the cumulative payoffs their predecessors have achieved in earlier unrelated rounds. We hypothesize that subjects are primarily influenced by the subject with the highest earlier payoff (i.e., the highest reputation), despite the fact that this information is irrelevant in a rational Bayesian model.

Related literature There is by now a large (mostly theoretical) literature on network effects.[5] Surprisingly, there is, however, only a rather small theoretical literature on the interplay between information cascades and network effects, and the few papers that do exist differ in some important aspects from Bikhchandani, Hirshleifer, and Welch (1992) (see e.g., Choi, 1997; Vergari, 2004; Frisell, 2003; Jeitschko and Taylor, 2001; Corsetti, Dasgupta, Morris, and Shin, 2004; Dasgupta, 2000).[6] A combination of Bikhchandani, Hirshleifer, and Welch's (1992) seminal model with payoff externalities does not seem to have been treated theoretically. Given the relative scarcity of theoretical work on the interplay between information cascades and network effects it is not surprising that there is also almost no experimental work in this area.[7] The only experimental paper introducing payoff externalities in the Bikhchandani, Hirshleifer, and Welch (1992) framework we are aware of is Hung and Plott (2001).[8] They study treatments in which subjects are rewarded if a majority of decisions was correct or if the subject's action agreed with the majority, respectively. The externalities in our experiment are, however, of a different form.

The remainder of the paper is structured as follows. In Section 2 we describe the basic experimental settings. In Section 3 we derive the theoretical predictions for the various treatments. While we find a multitude of equilibria in a treatment with network effects, there do not seem to exist

[5] For a recent survey of this literature, see e.g., Farrell and Klemperer (2004). For a textbook treatment, see e.g., Shy (2001).

[6] In Choi's (1997) model herding is not driven by private information but by the interplay of risk aversion and network effects (see also Vergari, 2004). Frisell (2003) considers a waiting game, where two firms with private information regarding the most profitable niche have to decide about entry into a market with horizontally differentiated products. While Jeitschko and Taylor (2001) study an investment game where randomly matched agents play pairwise coordination games, Corsetti, Dasgupta, Morris, and Shin (2004) explore the influence of a large trader in a model of speculative currency attacks with private information. Finally, Dasgupta (2000) studies a model relatively close in structure to Bikhchandani, Hirshleifer, and Welch (1992). However, he considers continuous signals and a network externality of a relatively extreme form: agents may realize a positive profit only if all agents coordinate on the same action.

[7] Beginning with Anderson and Holt (1997), there is by now a large experimental literature on information cascades in the absence of payoff externalities. This literature will be (partly) reviewed in Section 5.1 below.

[8] Guarino, Huck, and Jeitschko (2003) provide an experimental test of the above mentioned paper by Jeitschko and Taylor (2001). See also Schotter and Yorulmazer's (2004) experimental study of bank runs, where both asymmetric information and (negative) payoff externalities are present.

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 novel, we explain how we resolved the issues of recruitment, payment of subjects, and the implementation on the internet.

Section 5 contains our results. Section 5.1 deals with replication of earlier cascade experiments. Besides presenting our own results, we review and compare results from 12 earlier experiments that are scattered in the literature. Compared to these studies we find that subjects behave less frequently in line with theory. We attribute this observation to the fact that we consider longer decision sequences as well as asymmetric priors (i.e., one of the alternatives is more likely to be successful from an ex-ante perspective). Separately considering various subgroups of our diverse subject pool reveals that there are no significant differences across educational background or sex of the subjects, but that consultants have a somewhat larger tendency to follow their own signal. Section 5.2 presents the results from the treatments with payoff externalities. We study uniformity, volatility, and predictability of behavior in these treatments. We find that subjects seem to behave myopically (i.e., they tend to take only the decisions of their respective predecessors, but not the behavior of their successors, into account).[9] Finally, Section 5.3 deals with reputation effects in the basic Bikhchandani et al. model. While from a theoretical point of view reputation should not matter, we find that subjects' behavior is significantly influenced by the behavior of the predecessor with the highest reputation, and these "success models" were on average indeed more likely to pick the successful alternative. Section 6 concludes. Instructions of the experiment are contained in an Appendix.

2 The Experiment

In 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 that was reliable with probability $P(a|A) =$

[9] For related empirical evidence on the apparent lack of forward-looking behavior in the presence of payoff externalities see Tucker (2004).

$P(b|B) = 0.6$. Sessions with these probabilities are denoted by 55-60. In some treatments we conducted additional sessions with the probability combination 50-66, which is the one most often used in the literature (see e.g., Anderson and Holt, 1997).[10]

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.[11] In a first version a subject's payoff depends exclusively on his own decision. This version is equivalent to the basic model studied by Bikhchandani, Hirshleifer, and Welch (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 will be discussed in more detail in Section 5.3.

In a 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.

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 capture network externalities. Examples are the choice of software or mobile phone operators. There, it is not the quality of the product alone on which a choice should be based. As the utility from such products is increasing in the number of adopters, it is also important which product is selected by the majority of other consumers.

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

[10] The probability combination 50-66 has, however, the disadvantage of requiring a tie–breaking assumption in many cases.

[11] In a companion paper we focus on treatments with market prices for the investment opportunities A and B (as in Avery and Zemsky, 1998). Those treatments were conducted in the same experiment (see Drehmann, Oechssler, and Roider, 2005).

Table 1: Treatments

treatment	description		# of groups
BHW	Bikhchandani/Hirshleifer/Welch		63/12/15*
BHW+AS	*BHW* +	signals of all predecessors 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; ** denotes that there were 12 groups with $x = 0.4$ and 6 with $x = 1$; in treatments *Follower* and *Early bird* x is always 0.4.

In treatment *Early bird* subjects have to pay x for each predecessor who chose the same action as they. This kind of payoff externality is typical for situations where overcrowding is an issue as in restaurants, movie theaters, beaches, etc.

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, avant-gardist, and the participants in snowball systems or chain letters.

3 Theoretical Predictions

As mentioned above, there does not seem to exist in the literature a theoretical treatment of a model combining the Bikhchandani, Hirshleifer, and Welch (1992) model with payoff externalities. When subjects' payoffs are either independent of others or depend only on the behavior of predecessors, (unique) equilibrium predictions are straightforward to obtain by backward induction. This is the case in treatments *BHW* and *Early bird*. At first sight it may be surprising that for the other treatments either there exist multiple pure–strategy equilibria or none seem to exist at all. However, it is well known from the literature on network effects that this may happen as the strategic situation may well resemble those of coordination or mis–coordination games. This feature lights up in treatments *Network*, *Follower*, and *Hipster*.

Table 2 presents a non–exhaustive list of candidates for pure strategy (perfect Bayesian) equilibria given probability combination 55-60.

Table 2: Candidate equilibria for probability combination 55-60

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, 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 (for details see Bikhchandani, Hirshleifer, and Welch, 1992, or Drehmann, Oechssler, and Roider, 2004). 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 again be solved by backward induction. It turns out that cascades happen but they are endogenously broken once sufficiently many predecessors have chosen 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 perfect Bayesian equilibria with suitably chosen off–equilibrium beliefs.[12] 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 Nash equilibria, and we conjecture that no pure strategy equilibria exist. As an example consider the *uniform* equilibrium candidate in treatment *Follower*. If the first players receives and follows a b signal,

[12] This holds for $x = 0.4$. For $x = 1$, other equilibria exist.

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.

Finally, given the potential complexity of the equilibrium strategies in some of the payoff-externality treatments, there might be a tendency for players to behave myopically (i.e., to ignore the behavior of their respective successors altogether). If subjects indeed behave myopically, it turns out that treatment *Follower* yields the same prediction as *BHW*. Likewise, *Early bird* and *Hipster* become indistinguishable from each other. Below we shall also test this behavioral assumption, which is not uncommon in the literature on network effects.

4 Experimental Design

More than 6000 subjects participated in our internet 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.[13] 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. The last column of Table 1 lists the number of groups that participated in our experiment, separately for each combination of treatments, probabilities, and whether subjects came from the general subject pool or the control experiment with the consultants.

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 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 because negative payments are obviously very difficult to enforce in any experiment, let alone an internet experiment.

[13] 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).

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. Importantly, 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 a 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 a Phase II, each of 90000 lottery tickets had an equal chance of winning one of another 5 prizes of 1000 Euros. Finally, in a 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 effects of monetary incentives on behavior 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.[14] Table 3 lists some of the main characteristics of the combined subject pool (including the control experiment with consultants).

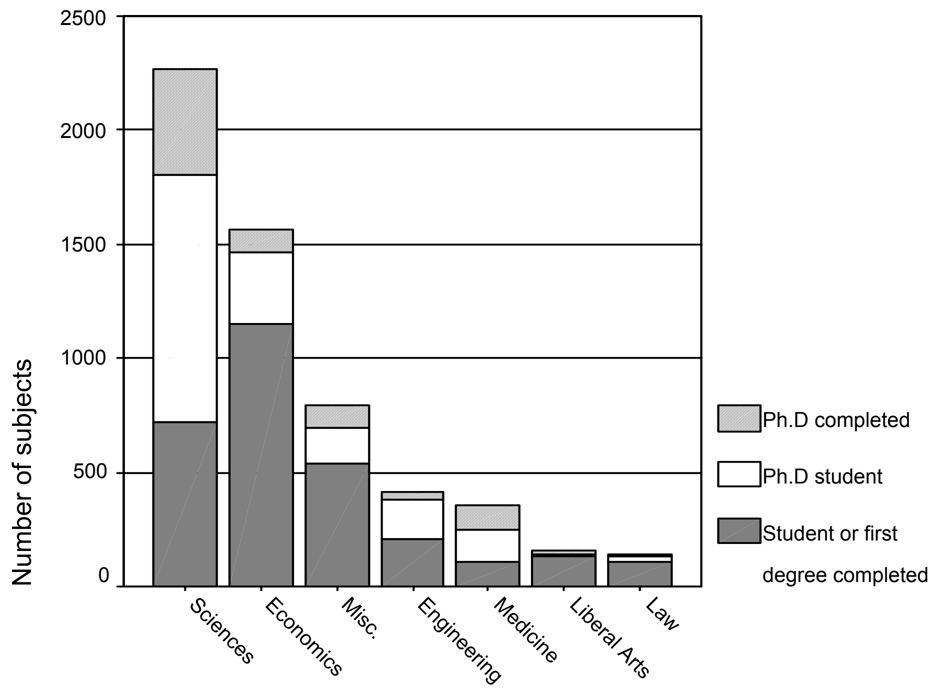

Figure 1: Composition of the subject pool. (Note: "Sciences" includes physics, chemistry, mathematics, and computer science; "Economics" includes economics, business administration, and related subjects; "Medicine" includes medicine, psychology, and dentistry; "Liberal Arts" includes all languages, history, 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 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.[15] Considering the number

[14] 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).

[15] 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.

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

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 (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.[16]

After entering the personal information, subjects were randomly placed in a currently active group,[17] 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.

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

[17]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.

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.[18] (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. comb.	all subjects *ruck*	all subjects *casc*	all subjects *own*	all subjects *own**	subjects on pot. eq. path *ruck*	subjects on pot. eq. path *casc*	subjects on pot. eq. path *own*	subjects on pot. eq. path *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 of decisions.

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*). To construct Table 4 we have pooled data from all phases (recall that different incentives were provided in different phases) and all stages (whether a task was first, second or third) as it turned out that neither the phase of the experiment nor the stage of the task had a significant influence on those results according to

[18] For a decision that, given the history of imputed signals, obviously violates Perfect Bayesian equilibrium, we let players suppose that the deviator followed his private signal. Dominitz and Hung (2004), who privately elicit the beliefs of all subjects after all decisions, report that subjects indeed seem to act on this assumption.

MWU-tests. Table 4 lists the above defined measures for all subjects and for 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.[19] 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.[20] Those numbers are lower than those found previously in most of the literature. In the following, we provide a brief overview of earlier experiments on information cascades and offer some explanations for the interesting observed behavioral differences.

Table 5 lists the results of all cascade experiments implementing the basic setup of BHW that we were able find in the literature.[21] 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 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.[22] To test for this we 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, Palfrey, Rogers, and McKelvey (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 (which again coincides with the length of sequences considered in many of the earlier studies). On those paths, values of *ruck* and *casc*

[19] Note that after a subject had made his decisions, we also asked for his prediction regarding the probability of A being successful (this prediction was, however, not renumerated). Interestingly, despite the relatively low value of *ruck*, like Dominitz and Hung (2004) we find that in over 90% of cases subjects chose actions that were consistent with their probability judgement.

[20] Recently, Kuebler and Weizsaecker (2004) have provided some evidence that longer cascades tend to be more stable indicating that, counter to theory, subjects perceive decisions in a cascade to be informative. They consider subjects already in a cascade and sort them according to the number of their cascade-predecessors. In five out of the six studies they review, average *casc* is lower in the first half of observation than in the second half. This is also the case in our experiment, albeit to a smaller extent.

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

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

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	
Dominitz/Hung (2004)	"replication" treatment	50-66	10		.88
Goeree et al. (2004)		50-66	20		.64
Hung/Plott (2001)	individualistic	50-66	10		.77
Kübler/Weizsäcker (2004)	NC	50-66	6		.78
Oberhammer/Stiehler (2001)		50-60	6	.86	.73
Stiehler (2003)	only equilibrium histories	50-60	6	.97	.93
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 $y\%$.

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 more and consequently rely more readily 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 higher (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,[23] and the number for *casc* is substantially lower compared to the general subject pool. It seems that consultants are more reluctant to rely

[23] 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.

on the decisions of others. Interestingly, Alevy, Haigh, and List (2003) 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.[24]

Finally, the level of payoffs may 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 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.[25] While we do not observe a significant difference in *ruck* between Phases I and II, 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 do not find significant differences between subjects holding a Ph.D., Ph.D. students, or others, and between male and female subjects. 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.

We call the difference between the number of A and B decisions the *decision imbalance*, and Figure 2 depicts the distribution of decision imbalances after the last player. This variable captures how many subjects in a given group make the same choice, and hence it measures uniformity.[26] Treatment *Network* (where a subject's payoff is the larger, the more players make the same choice) clearly stands out as the only treatment in which extreme imbalances occur. That is, in this treatment subjects often coordinate on the same choice, and, as we will see below, observational

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

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

[26] Note that one could also interpret this variable as the "market shares" of A and B.

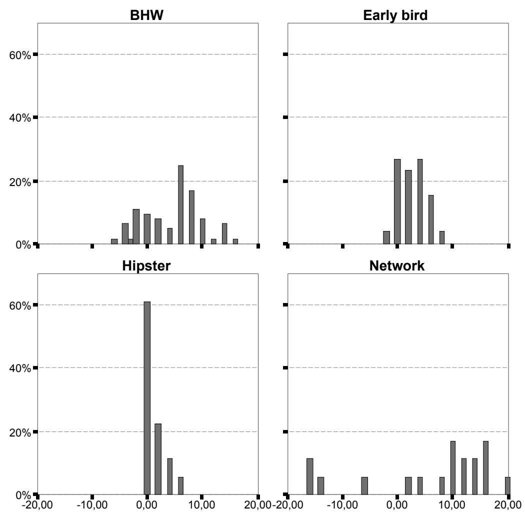

Figure 2: Distribution of decision imbalances after the last period (pooled over x).

learning early in the process is frequently pivotal for the outcome. On the other hand, treatments *Early bird* and *Hipster* (where one's payoff is the lower, the more players have made the same choice in the past) produce very balanced distributions centered around 0: 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 in *BHW,* in 7.7% in *Follower*, and in 0% in *Network*. 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 of possible behavioral differences in the payoff-externality treatments is the number and length of *runs* in the data. A run is a sequence of consecutive subjects who made the same decision. Hence, the number of runs captures whether a group exhibits relative stability or whether it is characterized by a rapid succession of short-lived fads. If there are positive payoff externalities (as in *Network*) we would expect longer (and therefore fewer) runs, i.e., runs should not be as fragile. When it is harmful to have many predecessors who made the same choice, runs

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.

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).[27] 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.

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? Above it has already become clear that in treatment *Hipster* an equal split is very likely: there, in 61.1% of cases the decision imbalance after the last player is exactly zero.[28] Indeed, in this treatment decision imbalances do in general not move too far away from zero: looking at decision imbalances pooled over all players (not just the last one) reveals that *Hipster* produces both the lowest mean (.84) and the lowest standard deviation (1.78) across all treatments.

In order to uncover potential predictability in the remaining treatments, 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. 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. Given that subjects in *Hipster* frequently split

[27] Treatment *BHW+AS* is excluded since all groups in this treatment consisted only of 10 subjects.
[28] This holds true in 50% (83%) of cases when $x = 0.4$ ($x = 1$).

50:50, it is not surprising that in this treatment it is hard to predict which alternative is chosen (slightly) more often.

Table 7: Predictability of majority choice

treatment	x	correlation between the sign of the decision imbalance after player 20 and after player... 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* for $x = 0.4$ 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 4 subjects deviate from the equilibrium path. In this sense, of the 12 groups with $x = 0.4$, 6 groups can be classified as one of the equilibria listed in Table 2, namely, 2 as *stubborn*, and 4 as *uniform* or *stubborn*.[29]

5.3 Reputation Effects

Treatment *BHW+R* is identical to treatment *BHW* except that subjects were able to observe not only the actions of their predecessors but also the cumulative payoffs (denoted in "Lotto-Euros") from the two decisions those subjects made in the first two (unrelated) stages of the experiment.[30] In a rational Bayesian model, this extra information is irrelevant. However, we suspected that subjects would rely more on the decisions of the predecessors with the highest payoffs (the "success models"). That is, subjects with higher payoffs have a better reputation and are imitated more often.

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

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

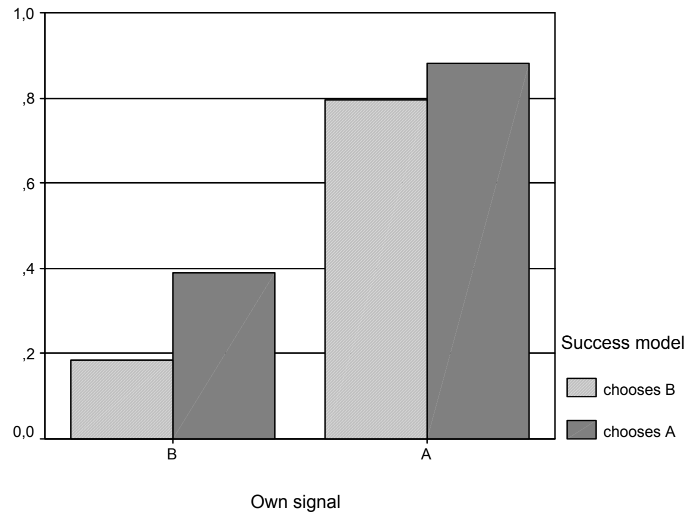

Figure 3: Fraction of subjects choosing A depending on the choice of the predecessor with the highest reputation and on the own signal. (Note: Including only subjects that had at least one predecessor and including only cases where the predecessor with the highest reputation was unique (547 out of 551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

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 the choice of A according to MWU–tests (at the 5% level).[31]

From an ex-post perspective, did it make sense for subjects to follow the respective success model? While on average subjects chose the successful alternative in 56% of cases, success models did so in 62% of cases. Hence, these subjects were indeed (somewhat) more successful in picking the right alternative, and imitating their behavior made sense.[32]

6 Conclusion

In a large–scale internet experiment we investigated information cascade models with and without payoff externalities. Reassuringly, while our subject pool is quite diverse (with large fractions of subjects having a background in the natural sciences and holding or studying for a Ph.D.), various

[31] This result is supported by a logit regression. Even when variables like a subject's signal and the signal imbalance are included, the decision of the predecessor with the highest payoff has a significant influence on the decision to choose a certain action.

[32] In a similar spirit, based on theoretical work by Celen and Kariv (2004), Celen, Kariv, and Schotter (2003) experimentally study the role of advice in an information cascade model. While in theory advice and actions are equally informative, Celen, Kariv, and Schotter (2003) find that (i) subjects appear to be more willing to follow their predecessors' advice than their predecessors' actions, and that (ii) the presence of advice is welfare-enhancing.

subgroups of subjects do not seem to behave significantly different with respect to the main research questions.

For the base treatment without payoff externalities, compared to earlier results in the literature we find a substantially lower percentage of subjects who behaved according to theory. We explain this deviation through a combination of 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 or experimental (see 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 shorter and more 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 also found to have strong effects on the predictability of the majority decision. With strong network effects, already after the second player (of 20) the majority decision can be predicted with great reliability.

Finally, one treatment in 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, J., M. HAIGH, AND A. LIST (2003): "Information Cascades: Evidence from a Field Experiment with Financial Market Professionals," *mimeo, University of Maryland.*

ANDERHUB, V., R. MÜLLER, AND C. SCHMIDT (2001): "Design and Evaluation of an Economic Experiment Via the Internet," *Journal of Economic Behaviour and Organization*, 46(2), 227–47.

ANDERSON, L., AND C. HOLT (1997): "Information Cascades in the Laboratory," *American Economic Review*, 87(5), 847–862.

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

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

BIKHCHANDANI, S., D. HIRSHLEIFER, AND I. WELCH (1992): "A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades," *Journal of Political Economy*, 100(5), 992–1026.

BIKHCHANDANI, S., D. HIRSHLEIFER, AND P. WELCH (1998): "Learning from the Behavior of Others: Conformity, Fads, and Informational Cascades," *Journal of Economic Perspectives*, 12, 151–170.

BOSCH-DOMENECH, A., J. MONTALVO, R. NAGEL, AND A. SATORRA (2002): "One, Two, (Three), Infinity, ...: Newspaper and Lab Beauty-Contest Experiments," *American Economic Review*, 92(5), 1687–1701.

BRESNAHAN, T. (2004): "Network Effects and Microsoft," *mimeo, Stanford University.*

CELEN, B., AND S. KARIV (2004): "Observational Learning Under Imperfect Information," *Games and Economic Behavior*, 47(1), 72–86.

CELEN, B., S. KARIV, AND A. SCHOTTER (2003): "The Advice Puzzle: An Experimental Study of Social Learning Where Words Speak Louder Than Actions," *mimeo, New York University.*

CHAMLEY, C. (2004): *Rational Herds.* Cambridge University Press, Cambridge.

CHARNESS, G., E. HARUVY, AND D. SONSINO (2001): "Social Distance and Reciprocity: The Internet Vs. The Laboratory," *mimeo.*

CHOI, J. (1997): "Herd Behavior, the 'Penguin Effect,' and the Suppression of Informational Diffusion: An Analysis of Informational Externalities and Payoff Interdependency," *Rand Journal of Economics*, 28 (3), 407–425.

CHURCH, J., AND N. GANDAL (1992): "Network Effects, Software Provision, and Standardization," *Journal of Industrial Economics*, 40, 85–104.

CIPRIANI, M., AND A. GUARINO (2004): "Herd Behavior in a Laboratory Financial Market," Working Paper, University College London.

CORSETTI, G., A. DASGUPTA, S. MORRIS, AND H. SHIN (2004): "Does One Soros Make a Difference? A Theory of Currency Crises with Large and Small Traders," *Review of Economic Studies*, 71(1), 87–114.

DASGUPTA, A. (2000): "Social Learning with Payoff Complementarities," *mimeo, Yale University.*

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

DOMINITZ, J., AND A. HUNG (2004): "Homogeneous Actions and Heterogeneous Beliefs: Experimental Evidence on the Formation of Information Cascades," *mimeo, Carnegie Mellon University.*

DREHMANN, M., J. OECHSSLER, AND A. ROIDER (2005): "Herding and Contrarian Behavior in Financial Markets - An Internet Experiment," *American Economic Review, forthcoming.*

FARRELL, J., AND P. KLEMPERER (2004): "Coordination and Lock-In: Competition with Switching Costs and Network Effects," *mimeo, Nuffield College.*

FORSYTHE, R., F. NELSON, G. NEUMANN, AND J. WRIGHT (1992): "Anatomy of an Experimental Political Stock Market," *American Economic Review*, 82, 1142–1161.

FORSYTHE, R., T. RIETZ, AND T. ROSS (1999): "Wishes, Expectations and Actions: A Survey on Price Formation in Election Stock Markets," *Journal of Economic Behavior and Organization*, 39, 83–110.

FRISELL, L. (2003): "On the Interplay of Informational Spillovers and Payoff Externalities," *Rand Journal of Economics*, 34 (3), 582–592.

GALE, D. (1996): "What Have We Learned from Social Learning?," *European Economic Review*, 40(3-5), 617–628.

GOEREE, J., T. PALFREY, B. ROGERS, AND R. MCKELVEY (2004): "Self-Correcting Informational Cascades," *California Institute of Technology, Social Science Working Paper 1197.*

GREINER, B., H. JACOBSEN, AND C. SCHMIDT (2002): "The Virtual Laboratory Infrastructure for Online Economic Experiments," *mimeo, Max Planck Institute for Research into Economic Systems Jena.*

GUARINO, A., S. HUCK, AND T. JEITSCHKO (2003): "Can Fear Cause Economic Collapse? Insights from an Experimental Study," *mimeo, University College London.*

GÜTH, W., C. SCHMIDT, AND M. SUTTER (2003): "Fairness in the Mail and Opportunism on the Internet - a Newspaper Experiment on Ultimatum Bargaining," *German Economic Review*, 4(2), 243–265.

HOLT, C. (2002): "Webgames and Strategy: Recipes for Interactive Learning," *mimeo, University of Virginia.*

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

JEITSCHKO, T., AND C. TAYLOR (2001): "Local Discouragement and Global Collapse: A Theory of Coordination Avalanches," *American Economic Review*, 91(1), 208–224.

KATZ, M., AND C. SHAPIRO (1985): "Network Externalities, Competition and Compatibility," *American Economic Review*, 75(3), 424–440.

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

Kuebler, D., and G. Weizsaecker (2004): "Are Longer Cascades More Stable?," *mimeo, London School of Economics and Political Science.*

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

Oberhammer, C., and A. Stiehler (2001): "Does Cascade Behavior in Information Cascades Reflect Bayesian Updating? An Experimental Study," Working Paper, Humboldt University Berlin.

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

Schotter, A., and T. Yorulmazer (2004): "On the Severity of Bank Runs: An Experimental Study," *mimeo, New York University.*

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

Shy, O. (2001): *The Economics of Network Industries.* Cambridge University Press, Cambridge.

Smith, L., and P. Sorensen (2000): "Pathological Outcomes of Observational Learning," *Econometrica*, 68(2), 371–398.

Stiehler, A. (2003): "Do Individuals Recognize Cascade Behavior of Others? An Experimental Study," *mimeo, Max Planck Institute for Research into Economic Systems Jena.*

Tucker, C. (2004): "Network Effects and the Role of Influence in Technology Adoption," *mimeo, Stanford University.*

Vergari, C. (2004): "Herd Behavior, Strategic Complementarities, and Technology Adoption," *mimeo, CORE.*

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

Willinger, M., and A. Ziegelmeyer (1998): "Are More Informed Agents Able to Shatter Information Cascades in the Lab?," in *The Economics of Networks: Interaction and Behaviours*, ed. by P. Cohendet, P. Llerena, H. Stahn, and G. Umbhauer, pp. 291–305. Springer-Verlag, New York.

Ziegelmeyer, A., J. Bracht, F. Koessler, and E. Winter (2002): "Fragility of Information Cascades: An Experimental Study Using Elicited Beliefs," *mimeo, University of Strasbourg.*

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 to 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 the previous round that is successful.

As in round 1, every participant receives a tip from his investment banker that 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 5 Lotto-Euros. 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 [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 to the next round.