



Herding with and without payoff externalities — an internet experiment

Mathias Drehmann ^{a,1}, Jörg Oechssler ^{b,*}, Andreas Roider ^c

^a Bank of England

^b University of Heidelberg, Germany

^c University of Bonn, Germany

Received 2 December 2004; received in revised form 28 December 2005; accepted 6 April 2006
Available online 10 July 2006

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, of which more than 2400 participated in the treatments reported here, 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.

© 2006 Elsevier B.V. All rights reserved.

JEL classification: C92; D8

Keywords: Information cascades; Herding; Network effects; Reputation; Experiment; Internet

1. Introduction

1.1. 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

* Corresponding author. Department of Economics, University of Heidelberg, Grabengasse 14, 69117 Heidelberg, Germany.
E-mail address: oechssler@uni-hd.de (J. Oechssler).

¹ The views expressed in this paper are those of the authors and do not necessarily reflect those of the Bank of England.

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 et al. (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.² In these models, externalities are created only through the information that can be deduced from observed actions.³ 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.⁴ 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 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 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 Ponzi (or pyramid) schemes or chain letters.

The present paper is part of a project that in total involved more than 6000 subjects participating in an internet experiment. In a companion paper (Drehmann et al., 2005), we consider treatments in which the investment opportunities have a variable market price (as in Avery and Zemsky, 1998). The current paper deals with settings in which, as in most of the literature on herding, prices for investment opportunities are constant. 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 many more subjects 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,

² 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.

³ For surveys of this literature, see, e.g., Bikhchandani et al. (1998) and Gale (1996); for a generalization, see Smith and Sorensen (2000); and for a recent textbook treatment, see Chamley (2004).

⁴ 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.

including 267 consultants from an international consulting firm. More than 40% of our subjects hold a PhD or are currently enrolled in a PhD 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.⁵

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 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).

1.2. Related literature

There is by now a large (mostly theoretical) literature on network effects.⁶ 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 et al. (1992) (see, e.g., Choi, 1997; Vergari, 2004; Frisell, 2003; Jeitschko and Taylor, 2001; Corsetti et al., 2004; Dasgupta, 2000).⁷ A combination of Bikhchandani et al.'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.⁸ To our knowledge, the only experimental paper introducing payoff externalities in the Bikhchandani et al. (1992) framework is Hung and Plott (2001).⁹ 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 treat-

⁵ Methodological and organizational issues arising in an internet setting are discussed in Section 4.

⁶ For a recent survey of this literature, see, e.g., Farrell and Klempere (2004). For a textbook treatment, see, e.g., Shy (2001).

⁷ 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 pair-wise coordination games, Corsetti et al. (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 et al. (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.

⁸ 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.

⁹ Guarino et al. (in press) 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.

ments. In Section 4, we discuss the experimental procedures and methodological issues connected to the implementation on the internet.

Sections 5–7 contain 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 provide evidence that this difference is driven by 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. To further investigate the influence of asymmetric priors, in Section 5.2, we report results on the subjects' beliefs, which we elicited as part of the experiment.

Section 6 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).¹⁰ Finally, Section 7 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 8 concludes.

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 $P(a|A)=P(b|B)=0.6$. 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).¹¹

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 Appendix A 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. 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 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 will be discussed in more detail in Section 7. In a second version of the above model, we introduce four different forms of payoff

¹⁰ For related empirical evidence on the apparent lack of forward-looking behavior in the presence of payoff externalities, see Tucker (2004).

¹¹ A slight difference to the literature is that we used a signal reliability of 0.66 (to be consistent with our usual two-digit representation of probabilities), whereas most of the literature used exactly 2/3. All probability combinations with a prior of 0.5, however, have the disadvantage of requiring a tie-breaking assumption in many cases.

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. In those cases, it is not the quality of the product alone on which a choice should be based. Since the utility derived 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.

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 because 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-gardists, and the participants in Ponzi (or pyramid) schemes or chain letters.

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.

3. Theoretical predictions

As mentioned above, there does not seem to exist in the literature a theoretical treatment of a model combining the Bikhchandani et al. (1992) model with payoff externalities. When subjects’ payoffs are either independent of others or depend only on the behavior of predecessors, (unique)

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 the same alternative	12/6**
Follower	BHW+receive x for each follower who chooses the same alternative	26
Early bird	BHW+pay x for each predecessor who chose the same alternative	26
Hipster	BHW+pay x for each predecessor and receive x for each follower who chooses the same alternative	12/6**

Notes: $*x/y/z$ 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, Network, Follower, Early bird, and Hipster, 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.

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

Notes: 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.

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 there may exist multiple pure-strategy equilibria and equilibrium strategies may take a rather complicated form. However, it is well known from the literature on network effects that this may happen as the strategic situation resembles 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.

In treatments BHW and BHW+AS, there is a unique perfect Bayesian equilibrium outcome, which depends in a simple way on the net number of a signals Δ that can be imputed from the actions of predecessors and the own signal (for details, see Bikhchandani et al., 1992). We call this the bhw equilibrium. It can easily be verified 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 (see Appendix B). 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.

In treatments Follower and Hipster, none of the candidates listed in Table 2 are Perfect Bayesian equilibria for 55–60. As an example consider the uniform equilibrium candidate in treatment Follower. If the first player 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 .¹³ However, more complicated equilibria exist. For example, in Appendix C, we prove that the following forms a perfect Bayesian equilibrium for Follower, 55–60 ($x=0.4$). Players 1, ..., 15 choose A regardless of their signal. If any of these players chose B , successors hold the belief that such a deviator got an a signal. Players 16, ..., 20 behave and form beliefs as in equilibrium bhw, i.e., for players 16, ..., 20 an A cascade starts if $\Delta \geq 1$, a B cascade starts if $\Delta \leq -2$, and otherwise these players follow their signal.

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

¹³ For probability combination 50–66, uniform would be an equilibrium because the last player would be indifferent between following his signal and following the action of the first player.

Given the complexity of equilibrium strategies in some of the payoff-externality treatments, we believe a more reasonable theoretical benchmark is to suppose that players behave myopically (i.e., that they ignore the behavior of their respective successors altogether). Indeed, this assumption is not uncommon in the literature on network effects. If subjects 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 the validity of this behavioral assumption.

4. Experimental design and methodological issues

The experiment was conducted on the internet and was available for a period of about six weeks in the spring of 2002 on our website <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 20,¹⁴ and subjects were asked to make decisions in three independent groups.¹⁵

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, utilizing the internet to conduct economic experiment is still novel. Therefore, in the following, we discuss some of the advantages, challenges, and organizational issues that arise in an internet setting.¹⁶

Clearly, the most important advantage of an internet experiment is the possibility to attract a large and diverse subject pool, which allows to study a greater number of treatment variations, to obtain more robust test statistics, and to investigate the outside validity of previous results. In total, more than 6000 subjects participated in our internet experiment and of those subjects more than 2400 participated in the treatments reported in the present paper.¹⁷ Notably, this includes a control group of 267 consultants from McKinsey and Company, an international consulting firm, who participated in the experiment on the same website a couple of weeks before the start of the actual experiment. These subjects knew that all other subjects were also consultants.¹⁸

Table 3 lists some of the main characteristics of the combined subject pool (including the control experiment with consultants). In contrast to most experiments in economics, our subjects come from a broad range of fields. Fig. 1 shows the frequencies of the main subject groups. Each bar in Fig. 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 PhD. Considering the number of PhD students and PhDs, we believe we succeeded in recruiting a fairly bright subject pool.¹⁹

¹⁴ 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).

¹⁵ For a more detailed account of the implementation, the recruiting and payment of subjects, and other practical issues of our experiment, we refer the reader to our companion paper (Drehmann et al., 2005).

¹⁶ For experiments that have been conducted over the internet, see, e.g., Forsythe et al. (1992), Forsythe et al. (1999), Lucking-Reiley (1999), Shavit et al. (2001), Bosch-Domenech et al. (2002), Güth et al. (2003), and Bossaerts and Plott (2004). For technical issues, see, e.g., Greiner et al. (2002). Anderhub et al. (2001), and Charness et al. (in press) contain detailed discussions of methodological issues.

¹⁷ The remaining subjects participated in other treatments reported in our companion paper (Drehmann et al., 2005).

¹⁸ 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.

¹⁹ Given that each time that we sent out emails to PhD students and post-docs to advertise the experiment, there was immediately a peak in access to our web page, 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 PhD	13.7
% current PhD students	31.3

The fact that the internet allows for larger subject pools is driven by a number of factors. For example, subjects' transaction costs are substantially lower. Rather than having to be physically present in a laboratory at a predetermined time, subjects can participate from their home or office at a time of their choice; the only prerequisite being access to the web. This factor also seems to have been pivotal in attracting the relatively large number of consultants. These subjects were recruited by an internal email to all German McKinsey consultants, and about a third of those addressed participated. This more generally suggests that internet experiments might be a particularly promising way to conduct experiments with professionals.

Also, the internet's potential for widespread visibility of the experiment facilitates finding an external sponsor who, in our case, agreed to advertise the experiment (through a number of newspaper ads and posters)²⁰ and to fund prizes.²¹ In the experiment, each Lotto-Euro was a ticket in a lottery to win one of our main prizes. In total there were 11 prizes of 1000 Euros 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. Consultants had the chance to win eight vouchers for a nice dinner for two in a restaurant each worth 150 Euros.²²

Finally, apart from a larger and more diverse subject pool, the internet also increases the flexibility of the experimenter. For example, through an online administration of the experiment, new variations of treatments can quickly and easily be created and run in response to behavior observed so far.

While the internet offers a number of advantages, it also poses some challenges. First, it implies a certain loss of control for the experimenter. For example, it is hard to rule out that some subjects were in fact groups of people deciding together.²³ Also, some subjects might be less attentive than they would be in a laboratory, which could lead to higher variances in behavior on the internet as reported, for example, by [Anderhub et al. \(2001\)](#) and [Bosch-Domenech et al.](#)

²⁰ An additional (very effective) way of attracting subjects was sending out emails to addresses we collected from websites of science and economics departments at German universities (see also [Lucking-Reiley, 1999](#)), although some people complained about spamming issues.

²¹ Importantly, McKinsey neither influenced the design of the experiment nor did they have access to the data. Only the instructions were edited with their help to make them more accessible to a general audience, which seems to be common in newspaper or internet experiments (see, e.g., [Bosch-Domenech et al., 2002](#); [Güth et al., 2003](#)). For example, subjects were told that they would receive a tip by some "investment banker".

²² We chose this lottery procedure rather than to pay all subjects in order to reduce transaction costs on our part.

²³ However, note that, by allocating subjects who logged on at about the same time to different treatments, we aimed to rule out "observational learning" in case two subjects sat next to each other in a computer pool.

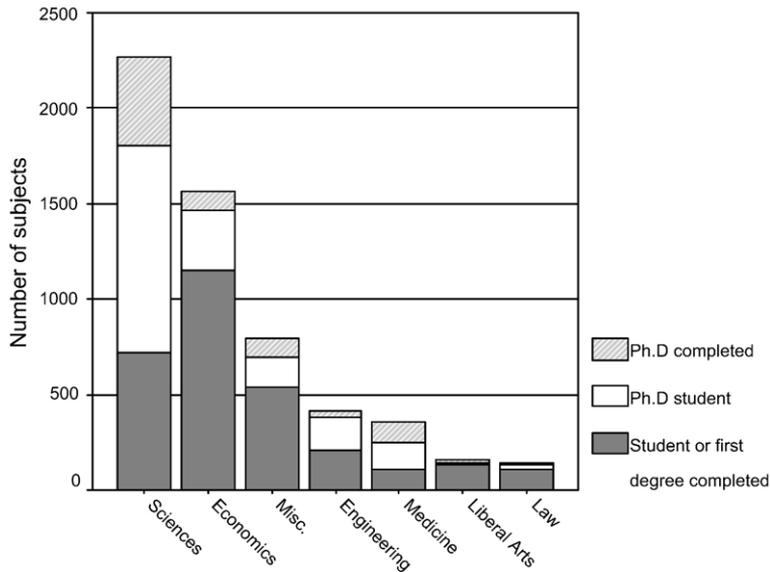


Fig. 1. Composition of the subject pool. (Note: “sciences” includes physics, chemistry, mathematics, and computer science; “economics” includes economics, business administration, and related subjects; “medicine” includes medicine, psychology, and dentistry; “liberal arts” includes all languages, history, and pedagogy. “Misc.” stands for miscellaneous fields).

(2002).²⁴ Finally, subjects might attempt to play more than once. To counter such behavior, we ensured that subjects with the same name–postal code combination or the same email address could only play once (providing a correct postal address was necessary for being able to win one of the prizes). The program also set cookies to prevent using the same computer twice.

A second concern is to make the rules of the game and the promised financial rewards credible. To this end, our web page was linked to the Laboratory for Experimental Research in Economics at the University of Bonn and to the sponsor, McKinsey and Company. Also, we ensured subjects of the confidentiality of their data by posting a legally binding “protection of privacy” statement that was approved by McKinsey lawyers.

Finally, from a practical perspective, the large (and in general ex-ante unknown) number of participants gives rise to various organizational questions. For example, due to the discrete nature of the advertising (e.g., ads appeared once a week), access to our web page was characterized through longer periods of low traffic and short peaks in access (with, for example, over 500 subjects participating in one day). While up to 40 subjects could play simultaneously, the program had a priority for filling up still uncompleted groups first, so as to maximize the number of completed groups, i.e., groups with 20 players.

To summarize, while conducting experiments on the internet poses challenges (which to a certain degree can be counteracted by the experimenter through careful design), it opens up some interesting new possibilities, like access to larger and more diverse subject pools, and greater flexibility for the experimenter. Moreover, our experiment was particularly suitable for the internet. First, as decision-making was sequential there was no need for subjects to be logged on at the same time. Second, the decision problem was rather easy to understand, and the duration of

²⁴ Interestingly, in settings where fairness considerations might come into play, Güth et al. (2003) and Charness et al. (in press) find significant differences between behavior in the lab and on the internet.

play was relatively short (15 min on average), such that a possible lack of attention of subjects was probably less of a problem. This makes us confident that our results are comparable to earlier experiments on information cascades conducted in the laboratory, and indeed, if one looks at similar experimental design parameters as in earlier laboratory studies, our findings are roughly the same (see Section 5.1 below).

5. Information-based herding

5.1. Replication and robustness

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.²⁵ (2) The fraction of cases in which subjects rationally decided against their own signal if they are in a cascade is denoted by *cas*. Arguably, *cas* 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 reports those measures 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 stages (whether a task was first, second or third) because it turned out that the stage of the task did not have a significant influence on those results according to MWU tests. Table 4 lists the above defined measures for all subjects and for the subset of players who are on a *self-confirming equilibrium path*. We say that a subject is on a self-confirming equilibrium path as long as this subject cannot tell whether a deviation from a perfect Bayesian equilibrium has occurred.²⁶ That is, the observed history of decisions must be such that there exists some history of signals under which the observed decisions are in accordance with *ruck*.

First, a look at all subjects shows that from a theoretical perspective subjects rely too heavily on their own private signals as *own* is weakly above *own** in all cases. As a result, in our main treatment, BHW 55–60, subjects act in accordance with theory in only 66% of cases. Even more dramatic is the picture with respect to *cas*. Only in 34% of cases did subjects decide against their signal but in accordance with Bayesian updating.²⁷ 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 differences in behavior.

Table 5 lists the results of all cascade experiments implementing the basic setup of Bikhchandani et al. (1992) that we were able find in the literature.²⁸ While the experiments differ with respect to a number of design issues, most notable the number of players in a sequence and the probability

²⁵ For a decision that cannot be reconciled with a perfect Bayesian equilibrium for any history of signals, we let players suppose that the deviator followed his private signal. Dominitz and Hung (2004) provide evidence in support of this assumption.

²⁶ This is related to the concept of a self-confirming equilibrium (see e.g., Fudenberg and Levine, 1998).

²⁷ Recently, Kübler and Weizsäcker (2005) 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 *cas* 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.

²⁸ We thank Lisa Anderson and Charlie Holt for kindly providing their data.

Table 4
BHW treatments

Subject pool	Treatment	Prob.	All subjects				Subjects on self-confirming equilibrium path			
			Comb.	Ruck	Casc	Own	Own*	Ruck	Casc	Own
General	BHW	55–60	0.66	0.34	0.75	0.59	0.86	0.74	0.74	0.90
	BHW+AS	55–60	0.72	0.41	0.74	0.68	–	–	–	–
Consultants	BHW	55–60	0.68	0.16	0.85	0.66	0.90	1.00	0.83	0.93
	BHW+AS	55–60	0.78	0.52	0.69	0.60	–	–	–	–
General	BHW	50–66	0.78	0.45	0.75	0.62	0.95	0.78	0.84	0.77
	BHW+AS	50–66	0.76	0.59	0.69	0.69	–	–	–	–

Note: The average length of self-confirming 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 self-confirming equilibrium history of decisions.

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 6.²⁹ To test for this, we look at the decisions of our first six subjects in each group. And indeed, for the first six 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 self-confirming equilibrium path (right panel of Table 4). On average, self-confirming equilibrium paths have length 6 (which, incidentally, coincides with the length of sequences considered in many of the earlier studies). 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 self-confirming equilibrium path. Interestingly, on self-confirming 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.³⁰ 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,³¹ 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. (in press) 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.

²⁹ Huck and Oechssler (2000) show that ruck values are substantially lower when decisions are more complex.

³⁰ It might also be the case that the medium per se has a certain effect (see, e.g., Shavit et al., 2001).

³¹ 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.

Table 5
Previous cascade experiments

Study	Treatment	Prob.	Group size	Ruck	Casc
Alevy et al. (2003)	Symmetric, students	50–66	5	0.95	0.89
Anderson and Holt (1997)	Symm., no public sig.	50–66	6	0.92	0.73
Anderson (2001)	2\$	50–66	6		0.70
Cipriani and Guarino (2005)	Fixed price	50–70	12	0.83	
Dominitz and Hung (2004)	“Replication” treatment	50–66	10		0.88
Goeree et al. (2004)		50–66	20		0.64
Hung and Plott (2001)	Individualistic	50–66	10		0.77
Kübler and Weizsäcker (2004)	NC	50–66	6		0.78
Oberhammer and Stiehler (2001)		50–60	6	0.86	0.73
Stiehler (2003)	Only equilibrium histories	50–60	6	0.97	0.93
Willinger and Ziegelmeyer (1998)	Treatment 1	50–60	6		0.64
Ziegelmeyer et al. (2002)	Blue line, experiments 1 and 2	55–66	9		0.69

Notes: Only studies that implement the Bikhchandani et al. (1992) 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 self-confirming 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 $y\%$.

This hypothesis is supported by the significantly higher numbers for ruck (78%) and casc (45%) in 50–66.³² For example, a subject 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. And indeed, if the first player chose A and the second player got a b signal, ruck values are very low, namely 18% (compared with ruck values of 66% for 55–60 overall). Even first players make substantially more mistakes in 55–60. When we consider first players who received a b signal in 50–66, 88% follow their signal, as they should. In 55–60, on the other hand, only 54% of those receiving a b signal chose B , which is indicative of the difficulty many subjects had in making the correct Bayesian updating from the prior of $P(A)=0.55$.³³ All this indicates that probability combination 55–60 is indeed harder than 50–66, which in combination with the first explanation above might be responsible for the values observed in Table 4.

Finally, we shall investigate whether our non-standard subject pool can account for the observed differences. As discussed above, relative to laboratory experiments, where subjects are usually (economics) students, our subject pool is out of the ordinary in that the subjects are unusually educated, come from diverse fields of studies, and include professional. Do those differences contribute to different values of ruck and casc?

To study this question, we first isolate the subgroup of subjects who are undergraduate economics or business students.³⁴ When we compare values for ruck and casc, no clear differences emerge relative to the general subject pool. In 55–60 (50–66), ruck for this subgroup of economics and business students is 69% (82%), which is slightly higher than in the general subject pool in our experiment. On the other hand, casc is lower in 50–66 (20%) but higher in 55–60 (39%). Of course, we cannot exclude selection effects through the internet but this applies to the laboratory setting too.

We also do not find significant differences between subjects holding a PhD, PhD 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). Also, the difference

³² With respect to ruck (casc) the difference is significant at the 1% (10%) level.

³³ This also shows up when we asked subjects about their beliefs regarding the true state (see Section 5.2).

³⁴ The number of decisions in BHW 55–60 and 50–66 are 170 and 28, respectively.

Table 6
Beliefs and rationality

		Probability combination	
		50–66	55–60
% of decisions in line with beliefs		0.93	0.84
Ruck depending on whether	Yes	0.81	0.69
Decision in line with beliefs	No	0.41	0.50

Note: Treatment BHW.

between BHW and BHW + AS is more pronounced for consultants than for the general subject pool. This suggests that consultants may have less trust in the rationality of their predecessors and, therefore, rely less on predecessors' decisions.

5.2. Beliefs and decisions

After subjects had made their decisions, we asked them for their prediction regarding the probability of *A* being successful (this prediction was, however, not remunerated). Table 6 depicts the percentage of decisions that were in line with stated beliefs (i.e., predictions),³⁵ and, like Dominitz and Hung (2004) in a laboratory setting, we find that overall this held true in more than 90% of cases. Additionally, Table 6 reports average ruck values separately for decisions that were, respectively were not, in line with stated beliefs. Indeed, if subjects' actions and their stated beliefs did not agree, ruck values were substantially lower (41% vs. 81% and 50% vs. 69%, respectively). Whether subjects who did this were simply more confused or inattentive, one cannot know, but it certainly had an effect.

Given that beliefs and decisions agree in more than 90% of cases, but given ruck values as low as they are, it follows that beliefs were often mistaken. Thus, it is interesting to investigate what kind of mistakes subjects made in forming their posteriors. In order to avoid the difficulty of determining subjects' beliefs about their predecessors' rationality, we concentrate on the first player in each group, where the task seems to be simple: take the given prior and update it given the private signal. Table 7 shows the average ruck values of first players in each group depending on their private signal (where we focus on those players whose decisions were in line with their stated beliefs).

In all of the cases considered in Table 7, the optimal action is to simply follow one's signal. However, the task seems to be much more difficult in the case of receiving a *b* signal in 55–60 than in all other cases. In fact, more than 50% of subjects chose the wrong action. Comparing stated beliefs to the true posterior probability, we see a clear pattern: subjects generally make too small an adjustment from the prior in the direction of their signal.³⁶ In the case of 55–60 and a *b* signal, this leads to 52% of subjects choosing *A* instead of *B*. The phenomenon may be explained as a consequence of the well-known anchoring effect.

6. Payoff externalities

Fig. 2 presents a first view on how the various payoff externalities influence behavior. For each group, we call the difference between the number of *A* decisions and the number of *B* decisions up to

³⁵ An *A* (*B*) decision is said to be in line with the stated belief if this belief is weakly above (below) 0.5.

³⁶ See also Dominitz and Hung (2004).

Table 7
Beliefs, true posteriors, and behavior of first players

		Signal					
		<i>a</i>			<i>b</i>		
		Belief	True prob.	Ruck	Belief	True prob.	Ruck
Probability	50–66	0.64	0.66	1.00	0.41	0.34	0.88
Combination	55–60	0.60	0.65	0.97	0.49	0.45	0.48

Note: Only players in position 1 whose decisions were in line with their stated beliefs; treatment BHW.

and including player t the *decision imbalance after player t* . Fig. 2 depicts the distribution of decision imbalances after the last player, where each group represents one observation. The decision imbalance after the last player simply measures the net number of A decisions in a group, i.e., $\#A - \#B$. Extreme values (i.e., -20 or $+20$) would indicate that all players in a given group took the same decision, a value of zero indicates the same number of A and B decisions in the respective group.³⁷

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 learning early in the process is frequently pivotal for the outcome.³⁸ 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 distributions concentrated near 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. All other pair-wise comparisons show differences that are significant at least at the 5% level.

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 should be shorter and more frequent. Table 8 lists the average number and length of runs per group for our treatments (separate for A and B runs).³⁹ 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 .

³⁷ Note that one could also interpret this variable as the “market shares” of A and B .

³⁸ Given the multiplicity of equilibria for treatment Network, it is interesting which, if any, of the “simple” equilibria discussed in Table 2 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 four subjects deviate from the equilibrium path. In this sense, of the 18 groups, 11 groups can be classified as one of the equilibria listed in Table 2, namely, three as uniform, seven as uniform or stubborn, and one as reverse or stubborn.

³⁹ Treatment BHW+AS is excluded since all groups in this treatment consisted only of 10 subjects.

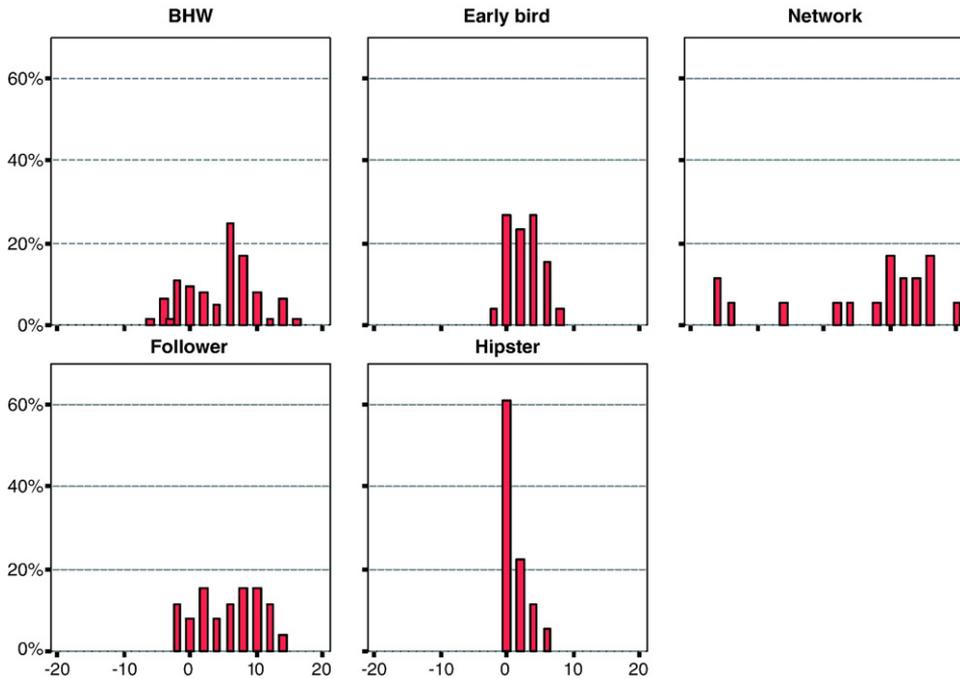


Fig. 2. Distribution of decision imbalances after the last period (pooled over x). (Note: The decision imbalance after player t is defined as the difference between the number of A decisions and the number of B decisions in a certain group up to and including player t).

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.⁴⁰ 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 (0.84) 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 9 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 50:50, it is not surprising that in this treatment it is hard to predict which alternative is chosen (slightly) more often.

7. Reputation effects

Treatment BHW+ R is identical to treatment BHW except that subjects were able to observe (additionally to the actions of their predecessors) also their predecessors' cumulative payoffs

⁴⁰ This holds true in 50% (83%) of cases when $x=0.4$ ($x=1$).

Table 8
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.

(denoted in “Lotto-Euros”) from the two decisions those subjects made in the first two (unrelated) stages of the experiment.⁴¹ 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.

Fig. 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).⁴² Note that the increase is less drastic for a signals than for b signals. However, the frequency of A given an a signal is already at 80%, which leaves little room for reputation to increase it even further. On the other hand, the frequency of A given a b signal starts very low (below 20%). Here, the impact of the decision of a player with high reputation can be and is substantial.

From an ex-post perspective, did it make sense for subjects to follow the respective success model? A first look at overall values of ruck and casc in treatment BHW + R suggests that there is no significant difference to treatment BHW. On average, ruck is equal to 0.66 (which is exactly the same as in BHW), while casc is equal to 0.31. However, success models were indeed somewhat more successful in picking the right alternative. While on average subjects chose the successful alternative in 56% of cases, success models did so in 62% of cases. Thus, on average, imitating their behavior made sense. However, imitation of success models sometimes also discouraged ruck behavior. Fig. 4 shows ruck values depending on whether a subject was already in a cascade (and therefore should ignore his signal) or not. The two bars on the right show that in cascade-situations ruck values decrease when the private signal of the current subject was in line with the action of its success model. On the other hand, when not yet being in a cascade, ruck increases when the private signal is confirmed by the action of the success model.

Finally, the susceptibility to success models does not seem to be gender- or age-dependent. Controlling for the subjects’ sex and age in the regression discussed in footnote 41 reveals no significant effect.

8. Conclusion

In a large-scale internet experiment, we have investigated information cascade models with and without payoff externalities. Reassuringly, while our subject pool is quite diverse (with large

⁴¹ Recall that each subject had to make three decisions (stages 1 through 3). BHW + R was always played on stage 3.

⁴² 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.

Table 9
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.

fractions of subjects having a background in the natural sciences and holding or studying for a PhD), various 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) and the number of subjects deciding in sequence (20 vs. 6). 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 have 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

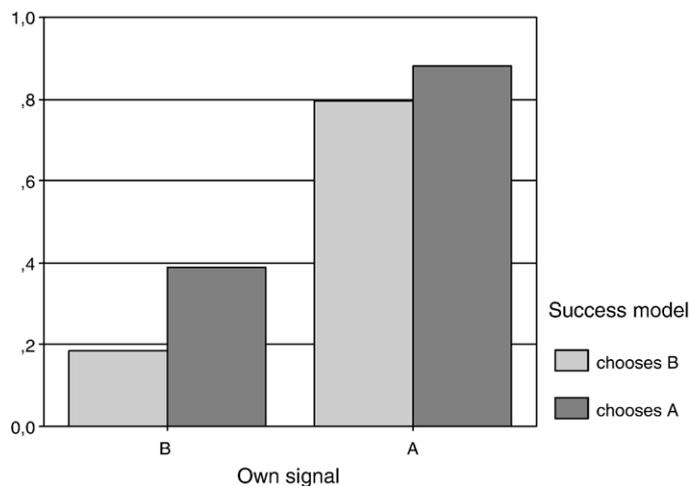


Fig. 3. Fraction of subjects choosing *A* depending on the choice of the predecessor with the highest reputation and on the own signal. (Note: Including only subjects that had at least one predecessor and including only cases where the predecessor with the highest reputation was unique (547 out of 551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

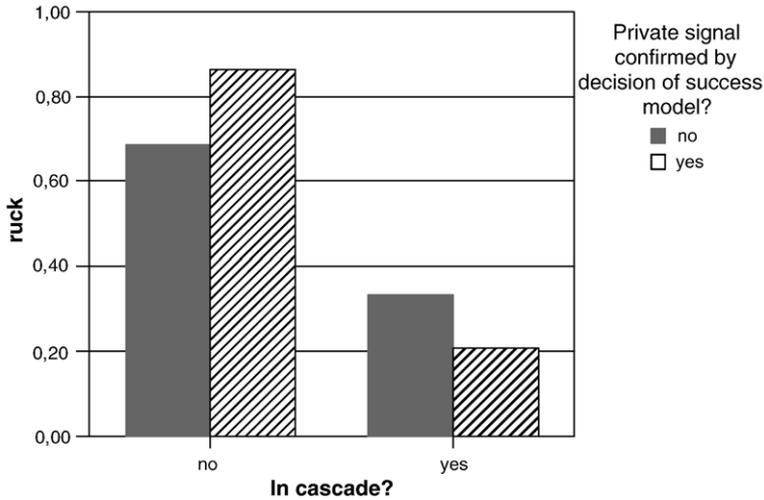


Fig. 4. Average ruck depending on being in a cascade and on the choice of the predecessor with the highest reputation. (Note: We say that a private signal is confirmed if it coincides with the decision of the success model. Including only subjects that had at least one predecessor and including only cases where the predecessor with the highest reputation was unique (547 out of 551 cases)).

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 our 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 had earned in previous 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 influenced the choice behavior of later subjects.

In a nutshell, our paper shows that information cascades happen, but probably less frequently than predicted by theory and earlier experiments. In particular, this seems to hold true when subjects face longer decision sequences or asymmetric priors. When payoff externalities are added to informational ones, herding is influenced in the predicted way: with positive externalities there is more herding, with negative externalities there is less herding. Moreover, in the presence of payoff externalities, we find evidence for myopic behavior; an assumption that is frequently made in the literature on network effects. Finally, we find that reputation indeed matters. Subjects tend to follow successful role models in their decisions.

Acknowledgement

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, Wissenschaftszentrum Berlin (WZB), the EARIE meetings 2005 in Porto, the Verein für Socialpolitik meetings

2005 in Bonn, and the Workshop on Informational Herding Behavior 2005 in Copenhagen for suggestions. Two referees provided very constructive comments. We thank Frauke Lammers, Susanne Niethen, and Nina Wessels for very useful discussions throughout the project, and McKinsey and Company for financial support. Financial support was also provided by SFB/TR 15, GESY. The third author gratefully acknowledges the hospitality of the Stanford Graduate School of Business and financial support by the German Research Foundation. Finally, we are very grateful to Ben Greiner who did an excellent job programming the experiment.

Appendix A. 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 were 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.

A.1. 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 and 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 5000 Euros which are distributed among five 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 min.

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!

A.2. 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, PhD status, age, and sex]

A.3. 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 sent 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.

A.4. 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”].

A.5. 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 3×0.4 LE because three of your predecessors have chosen B. If you would choose A, you would have to pay 1×0.4 LE.

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.

Appendix B. Treatment Early bird

As in treatment BHW, rational behavior in treatment Early bird (under the assumption that all of one’s predecessors behaved rationally as well) depends on the net number of a signals Δ_t that

Table 10
Threshold values for treatment Early bird

Δ_t	$P(A \Delta_t)$	D_{t-1}	$\hat{D}_t(D_{t-1})$
5	0.90	5	0.60
4	0.86	4	0.58
3	0.80	3	0.56
2	0.73	2	0.54
1	0.65	1	0.52
0	0.55	0	0.50
-1	0.45	-1	0.48
-2	0.35	-2	0.46
-3	0.27	-3	0.44
-4	0.19	-4	0.42
-5	0.14	-5	0.40

player t can impute from the decisions of his predecessors and his own signal (see Section 3). However, additionally, rational behavior by player t also depends on the decision imbalance D_{t-1} after player $t-1$ (which is defined as the difference between the number of A decisions and the number of B decisions up to and including player $t-1$, i.e., $D_{t-1} := \#A_{t-1} - \#B_{t-1}$). In particular, player t will choose A if and only if

$$10P(A|\Delta_t) - 0.4\#A_{t-1} > 10(1 - P(A|\Delta_t)) - 0.4\#B_{t-1} \Leftrightarrow P(A|\Delta_t) > \hat{D}_t(D_{t-1}) := 0.5 + (0.4D_{t-1})/20. \tag{1}$$

Note that in the present treatment ties in expected profits do not occur.

Table 10 presents the relevant values for $P(A|\Delta_t)$ and $\hat{D}_t(D_{t-1})$.

Suppose that player 1 receives an a signal (which in treatment BHW 55–60 would trigger an A cascade up to player 20). As player 1 has no predecessors he chooses A because $P(A|1) = 0.65 > 0.50 = \hat{D}_1(0)$. Independent of their own signals, players 2 and 3 choose A as well because $P(A|0) = 0.55 > 0.54 = \hat{D}_2(2)$ holds. However, player 4’s decision is again informative because he will only choose A if he receives an a signal (because $0.73 > 0.56$). If one, however, supposes that player 4 receives a b signal, the herd is endogenously broken because choosing B is optimal (because $0.55 < 0.56$). In this case, by a similar logic, player 5’s decision will be informative as well. However, a herd may start again at a later stage.

Appendix C. Treatment Follower

Proposition. The following is a pure strategy perfect Bayesian equilibrium for treatment Follower (55–60, $x=0.4$): players 1, ..., 15 choose A regardless of their signal. If any of these players chose B , successors hold the belief that such a deviator got an a signal. Players 16, ..., 20 behave and form beliefs as in equilibrium BHW, i.e., for players 16, ..., 20 an A cascade starts if $\Delta \geq 1$, a B cascade starts if $\Delta \leq -2$, and otherwise these players follow their signals.

Proof. (i) Preliminaries: The probability that B is correct given a b signal is given by

$$P(B|b) = \frac{P(b|B)P(B)}{P(b|B)P(B) + P(b|A)P(A)} = \frac{0.6*0.45}{0.6*0.45 + 0.4*0.55} = 0.551 \tag{2}$$

Furthermore, let $P(b|b)$ be the probability that some other player gets a b signal given that the first player received a b signal:

$$P(b|b) = P(b|B)P(B|b) + P(b|A)P(A|b) = 0.6*0.551 + 0.4*(1-0.551) = 0.5102 \quad (3)$$

And hence, $P(a|b) = 1 - 0.5102 = 0.4898$ and $P(a|b)^2 = 0.4898^2 = 0.2399$. Likewise, let $P(a|ab)$ be the probability that some other player gets a a signal given there have been an a and a b signals:

$$P(a|ab) = P(a|B)P(B|ab) + P(a|A)P(A|ab) = 0.4*0.45 + 0.6*0.55 = 0.51 \quad (4)$$

(ii) Player 15 (and a fortiori each of his predecessors) does not have an incentive to deviate if he receives a b signal: in equilibrium his payoff is at least

$$\begin{aligned} \pi(A|b) &\geq 10P(A|b) + 5xp(a|b) + 4xp(b|b)p(a|b)^2 \\ &= 10(1-0.551) + 5*0.4*0.4898 + 4*0.4*0.5102*0.4898^2 = 5.6654 \end{aligned} \quad (5)$$

where the latter two terms stem from the probability that player 16 gets an a signal (and hence starts an A cascade, resulting in a payment of $5x$) respectively the probability that players 16–18 get signals baa , which also starts an A cascade (resulting in a payment of $4x$). A deviation to B would only yield $10p(B|b) = 5.51$ to player 15 because all followers will choose A .

(iii) Player 16 (and a fortiori each of his followers) does not have an incentive to deviate from the BHW equilibrium path. Intuitively, a deviation motivated by the desire to cause a herd by one's followers is not profitable because the number of followers does not suffice to compensate for a lower chance of having picked the correct alternative. First, if following the equilibrium strategy would cause a herd by one's followers, there is no incentive to deviate. In that case, both the informational and the payoff externality point in the same direction, which makes a deviation unprofitable. Formally, this holds true (a) if $\Delta \leq -2$, or (b) if $\Delta \geq 1$, or (c) if $\Delta = -1$, and player 16 has received a b signal, or (d) if $\Delta = 0$ and player 16 has received an a signal, where Δ is defined excluding player 16's signal. Second, in the following, we show that in the remaining two cases (e) $\Delta = 0$ and player 16 has received a b signal, and (f) $\Delta = -1$ and player 16 has received an a signal, this player has indeed an incentive to follow his own information. First, consider case (e): in equilibrium player 16's payoff is at least

$$10p(B|b) + p(b|b)4x = 5.551 + 0.5102*4*0.4 = 6.3673 \quad (6)$$

while a deviation to A would yield at most

$$10p(A|b) + 4x = 10(1-0.551) + 4*0.4 = 6.09 \quad (7)$$

and, hence, there is no incentive to deviate. Finally, consider case (f): in equilibrium player 16's payoff is at least

$$\pi(A|ab) \geq 10P(A|ab) + 4xp(a|ab) = 10*0.55 + 4*0.4*0.51 = 6.316 \quad (8)$$

while by deviating he receives at most

$$\pi(B|ab) \leq 10P(B|ab) + 4x = 10*0.45 + 4*0.4 = 6.1 \quad (9)$$

which concludes the proof. \square

References

- Alevy, J., Haigh, M., List, J., in press. Information cascades: evidence from a field experiment with financial market professionals. *Journal of Finance*.
- Anderhub, V., Müller, R., Schmidt, C., 2001. Design and evaluation of an economic experiment via the internet. *Journal of Economic Behavior and Organization* 46 (2), 227–247.
- Anderson, L.R., 2001. Payoff effects in information cascade experiments. *Economic Inquiry* 39 (4), 609–615.
- Anderson, L., Holt, C., 1997. Information cascades in the laboratory. *American Economic Review* 87 (5), 847–862.
- Avery, C., Zemsky, P., 1998. Multidimensional uncertainty and herd behavior in financial markets. *American Economic Review* 88 (4), 724–748.
- Banerjee, A., 1992. A simple model of herd behavior. *Quarterly Journal of Economics* 107 (3), 797–817.
- Bikhchandani, S., Hirshleifer, D., Welch, I., 1992. A theory of fads, fashion, custom, and cultural change as informational cascades. *Journal of Political Economy* 100 (5), 992–1026.
- Bikhchandani, S., Hirshleifer, D., Welch, I., 1998. Learning from the behavior of others: conformity, fads, and informational cascades. *Journal of Economic Perspectives* 12, 151–170.
- Bosch-Domenech, A., Montalvo, J., Nagel, R., Satorra, A., 2002. One, two, (three), infinity,...: newspaper and lab beauty-contest experiments. *American Economic Review* 92 (5), 1687–1701.
- Bossaerts, P., Plott, C., 2004. Basic principles of asset pricing theory: evidence from large-scale experimental financial markets. *Review of Finance* 8, 135–169.
- Bresnahan, T., 2004. Network effects and Microsoft. mimeo, Stanford University.
- Chamley, C., 2004. *Rational Herds*. Cambridge University Press, Cambridge.
- Charness, G., Haruvy, E., Sonsino, D., in press. Social distance and reciprocity: the internet vs. the laboratory. *Journal of Economic Behavior and Organization*, forthcoming.
- 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., Gandal, N., 1992. Network effects, software provision, and standardization. *Journal of Industrial Economics* 40, 85–104.
- Cipriani, M., Guarino, A., 2005. Herd behavior in a laboratory financial market. *American Economic Review* 95 (5), 1427–1443.
- Corsetti, G., Dasgupta, A., Morris, S., Shin, H., 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., Dybvig, P., 1983. Bank runs, deposit insurance and liquidity. *Journal of Political Economy* 91 (3), 401–419.
- Dominitz, J., Hung, A., 2004. Homogeneous actions and heterogeneous beliefs: experimental evidence on the formation of information cascades. mimeo, Carnegie Mellon University.
- Drehmann, M., Oechssler, J., Roider, A., 2005. Herding and contrarian behavior in financial markets: an internet experiment. *American Economic Review* 95 (5), 1403–1426.
- Farrell, J., Klemperer, P., 2004. Coordination and lock-in: competition with switching costs and network effects. mimeo, Nuffield College.
- Forsythe, R., Nelson, F., Neumann, G.R., Wright, J., 1992. Anatomy of an experimental political stock market. *American Economic Review* 82, 1142–1161.
- Forsythe, R., Rietz, T.A., Ross, T.W., 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.
- Fudenberg, D., Levine, D., 1998. *The Theory of Learning in Games*. MIT Press, Cambridge.
- Gale, D., 1996. What have we learned from social learning? *European Economic Review* 40 (3–5), 617–628.
- Goeree, J., Palfrey, T., Rogers, B., McKelvey, R., 2004. Self-correcting informational cascades. California Institute of Technology, Social Science Working Paper, vol. 1197.
- Greiner, B., Jacobsen, H.-A., Schmidt, C., 2002. The virtual laboratory infrastructure for online economic experiments. In: Kremer, K., Macho, V. (Eds.), *Forschung und Wissenschaftliches Rechnen 2002: Beiträge Zum Heinz-Billing-Preis*. Gesellschaft für Wissenschaftliche Datenverarbeitung, Göttingen, pp. 59–73.
- Guarino, A., Huck, S., Jeitschko, T., in press. Averting economic collapse and the solipsism bias. *Games and Economic Behavior*.
- Güth, W., Schmidt, C., Sutter, M., 2003. Fairness in the mail and opportunism on the internet—a newspaper experiment on ultimatum bargaining. *German Economic Review* 4 (2), 243–265.

- Huck, S., Oechssler, J., 2000. Informational cascades in the laboratory: do they occur for the right reasons? *Journal of Economic Psychology* 21, 661–671.
- Hung, A., Plott, C., 2001. Information cascades: replication and an extension to majority rule and conformity rewarding institutions. *American Economic Review* 91 (5), 1508–1520.
- Jeitschko, T., Taylor, C., 2001. Local discouragement and global collapse: a theory of coordination avalanches. *American Economic Review* 91 (1), 208–224.
- Katz, M., Shapiro, C., 1985. Network externalities, competition and compatibility. *American Economic Review* 75 (3), 424–440.
- Kübler, D., Weizsäcker, G., 2004. Limited depth of reasoning and failure of cascade formation in the laboratory. *Review of Economic Studies* 71 (2), 425–441.
- Kübler, D., Weizsäcker, G., 2005. Are longer cascades more stable? *Journal of the European Economic Association* 3, 330–339.
- 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., Stiehler, A., 2001. Does cascade behavior in information cascades reflect Bayesian updating? An experimental study. mimeo, Humboldt University Berlin.
- Scharfstein, D., Stein, J., 1990. Herd behavior and investment. *American Economic Review* 80 (3), 465–479.
- Schotter, A., Yorulmazer, T., 2004. On the severity of bank runs: an experimental study. mimeo, New York University.
- Shavit, T., Sonsino, D., Benzion, U., 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., Sorensen, P., 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., Ziegelmeyer, A., 1998. Are more informed agents able to shatter information cascades in the lab? In: Cohendet, P., Llerena, P., Stahn, H., Umbhauer, G. (Eds.), *The Economics of Networks: Interaction and Behaviours*. Springer-Verlag, New York, pp. 291–305.
- Ziegelmeyer, A., Bracht, J., Koessler, F., Winter, E., 2002. Fragility of information cascades: an experimental study using elicited beliefs. mimeo, University of Strasbourg.