

Drehmann, Mathias; Oechssler, Jörg; Roeder, Andreas

Working Paper

Herding and Contrarian Behavior in Financial Markets: An Internet Experiment

Bonn Econ Discussion Papers, No. 25/2002

Provided in Cooperation with:

Bonn Graduate School of Economics (BGSE), University of Bonn

Suggested Citation: Drehmann, Mathias; Oechssler, Jörg; Roeder, Andreas (2002) : Herding and Contrarian Behavior in Financial Markets: An Internet Experiment, Bonn Econ Discussion Papers, No. 25/2002, University of Bonn, Bonn Graduate School of Economics (BGSE), Bonn

This Version is available at:

<https://hdl.handle.net/10419/22846>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

BONN ECON DISCUSSION PAPERS

Discussion Paper 25/2002

Herding and Contrarian Behavior in Financial Markets - An Internet Experiment

by

Mathias Drehmann, Jörg Oechssler,
Andreas Roeder

April 2003



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

The Bonn Graduate School of Economics is
sponsored by the

Deutsche Post  World Net

MAIL EXPRESS LOGISTICS FINANCE

Herding and Contrarian Behavior in Financial Markets - An Internet Experiment*

Mathias Drehmann Jörg Oechssler[†]
Andreas Roeder

*Department of Economics
University of Bonn*

This version: April 2003
Previous version: November 2002

Abstract

We report results of an internet experiment designed to test the theory of informational cascades in financial markets (Avery and Zemsky, *AER*, 1998). More than 6000 subjects, including a subsample of 267 consultants from an international consulting firm, participated in the experiment. As predicted by theory, we find that the presence of a flexible market price prevents herding. However, the presence of contrarian behavior, which can (partly) be rationalized via error models, distorts prices, and even after 20 decisions convergence to the fundamental value is rare. We also study the effects of transaction costs and the expectations of subjects with respect to future prices. Finally, we report some interesting differences with respect to subjects' fields of study.

JEL-classification: C92, D8, G1

Key words: informational cascades, herding, contrarians, experiment, internet.

*We are grateful to Jörg Breitung, Georg Weizsäcker, and Anthony Ziegelmeyer for detailed comments. We also thank seminar participants at Amsterdam, Bonn, Frankfurt, Harvard, Mannheim, NYU, Zurich, the Royal Economic Society meetings 2003 in Warwick, the ESA meetings 2002 in Strasbourg, and the ENDEAR meetings 2002 in Jena for lively discussions and suggestions. We thank Frauke Lammers, Susanne Niethen, and Nina Wessels for very helpful discussions throughout the project, and McKinsey & Company for financial support. We also thank Deutsche Forschungsgemeinschaft, grant OE-198/1/1. Finally, we are very grateful to Ben Greiner who did an excellent job programming the experiment.

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

1 Introduction

While the popular press often blames “lemming-like” behavior of investors for the observed exuberance in financial markets, the question whether herding actually occurs in real financial markets, remains for various reasons very difficult to resolve with field data. When investors are observed taking the same action, this may be due to herding. But it may also be caused by investors following the same information. Or, the clustering of actions may simply be incidental. In empirical work, those explanations are difficult to disentangle since the private information of investors is, in general, unobservable. In order to contribute to this question, we conducted a large-scale internet experiment based on a sequential asset market as described by Avery and Zemsky (1998). An experiment offers an opportunity to directly test herding theories since all fundamentals and the private information of agents are under control of the experimenter.

Several sources of rational herding are known to the theoretical literature. For example, when market participants’ payoffs depend directly on the behavior of others, herd behavior is natural. Such payoff externalities cause 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).¹ Even if such payoff externalities are absent, however, herd behavior may be observed in markets through a process of information transmission. Models based purely on informational externalities were pioneered by Bikhchandani, Hirshleifer and Welch (1992) (henceforth BHW), Welch (1992), and Banerjee (1992). An *informational cascade* is said to occur when it becomes rational to ignore one’s own private information and instead follow the predecessors’ decisions. Since no further information is revealed once an informational cascade has started, inefficiencies occur even though each individual is behaving rationally.

Theories of rational herding or informational cascades are, however, not directly applicable to financial markets. Market prices are a powerful mechanism which, in theory, efficiently aggregate private information of traders. In particular, Avery and Zemsky (1998) (henceforth AZ) have shown that informational cascades cannot occur in a simple sequential asset market because a flexible market price incorporates all publicly available information. Hence, rational traders should always follow their private signal and thereby reveal their information.² Note

¹Within our experiment we also implemented treatments, which look at the effects of reputation or payoff externalities in herding decisions. At the moment, these treatments are evaluated and will be reported in a companion paper.

²AZ show that herding may occur in the presence of multi-dimensional uncertainty, even though informational

that in this class of sequential trade models, traders are only allowed to buy or sell once, and hence classical price bubbles driven by traders, who think they can resell the asset before the bubble bursts, are not possible.³

In reality, herding may nevertheless occur due to the likely existence of boundedly rational traders who may be plagued by a variety of biases and follow more or less plausible rules of thumb. Imitation, trend chasing, momentum trading strategies, and the like are all alternative possible sources for herd behavior in financial markets. Finally, there are strategies advocated by popular guide books and analysts that should counteract herd behavior.⁴ In particular, “contrarian” or “value strategies” call for buying assets with low prices relative to some fundamental value like earnings, dividends, historical prices etc. (for empirical evidence on the profitability of such strategies, see e.g., Lakonishok et al., 1994, or La Porta et al., 1998).

The internet experiment we report on in this paper was designed to address the following objectives. First, we want to test the theory of informational cascades in financial markets as introduced by AZ by taking the theory at face value and implementing a design that exactly matches the theoretical set-up. The main test for this theory is whether herding is actually prevented by market prices. Second, we are interested in filtering out empirical regularities that may explain possible deviations from the theory. In particular, we want to find out whether traders follow their own signal (which is rational if all others are rational too), whether they engage in herd behavior, or whether they follow contrarian strategies by trading against their signal and the market. Third, we study the effect of transactions costs which, from a theoretical point of view, can prevent convergence of the price to its true value. Finally, an important issue is also what expectations subjects hold with respect to future prices, in particular, whether traders understand the (theoretical) martingale property of prices. For this purpose, we asked some subjects to bet on the final price of the asset.

More than 6000 subjects participated in our experiment, in which a substantial amount of prize money was at stake. The subject pool was exceptionally educated with more than 13% holding a Ph.D. and another 31% being Ph.D. students. Almost half the subjects were educated

cascades remain impossible. See also Cipriani and Guarino (2001) for other modifications of the AZ model that make herds possible.

³For a survey of theories of rational bubbles and herding, see e.g., Brunnermeier (2001). For a survey of experimental research on bubbles in asset markets, see e.g., Sunder (1995). For recent work in this area, see e.g., Hommes et al. (2002), or Hey and Morone (2002).

⁴See e.g. the investment classic *Contrarian Investment Strategies* by Dreman (1979, 1998) or his column “The Contrarian” in Forbes Magazine.

in natural sciences, mathematics, or engineering. We also conducted a control experiment with 267 consultants of an international consulting firm. The main treatments in this experiment were variants of the basic model by AZ in which a market price aggregates all publicly available information. Traders received a private signal and could observe the past history of prices and in most treatments additionally the decisions of their predecessors. The large number of participants in the experiment allowed us to introduce a variety of modifications of the basic model. For example, we explored many different combinations of a priori probabilities and signal precisions to check for robustness. We also looked at the effects of two different levels of transaction costs. For comparison, we also conducted treatments without market prices corresponding to the basic model of BHW. Finally, an important benchmark is a treatment in which subjects could not only observe the decisions of their predecessors but also their private signals. In this treatment, doubts about the rationality of others cannot be an issue.

Our experiment complements a large empirical literature with field data. Beginning with Lakonishok et al. (1992) researchers were analyzing the tendency of fund managers, security analysts (Welch, 2000), or investment newsletters (Graham, 1999) to herd (for surveys, see e.g., Bikhchandani and Sharma, 2000, Hirshleifer and Teoh, 2003, or Daniel et al., 2002). However, as Hirshleifer and Teoh (2003) note, it will always be difficult to empirically disentangle the mixture of reputational effects, informational effects, direct payoff externalities, and imperfect rationality.

Following Anderson and Holt (1997) there is by now a well established experimental literature on cascade and herding models.⁵ However, to our knowledge there is only one other experiment on cascades in financial markets with flexible prices which has been conducted by Cipriani and Guarino (2002).⁶ Their design follows closely that by Anderson and Holt (1997) in that subjects could observe other subjects taking decisions, and private information was framed as a draw from an urn. Some of our treatments are very close to their experiment, and we will comment at several places in this paper on their results.

Our paper is structured as follows. In the next section we discuss the theoretical predictions

⁵See e.g. Allsopp and Hey (1999), Ashiya and Doi (2001), Bounmy et al. (1998), Goidel and Shields (1994), Huck and Oechssler (2000), Hung and Plott (2001), Krämer et al. (2000), Kübler and Weizsäcker (2001), Nöth and Weber (2002), Nöth et al. (2001), Oberhammer and Stiehler (2001), Plott and Yang (1997), Sadiraj et al. (1999), Willinger and Ziegelmeyer (1998).

⁶When we conducted our experiment, we were not aware of the paper by Cipriani and Guarino (2002). However, it is clear that their experiment has precedence.

for the basic BHW model without prices and for the AZ model with market prices. In Section 3, we describe the experimental design, in particular the different treatments, the recruitment, the characteristics of the subject pool, and the implementation on the internet.⁷

The results of the experiment are presented in Section 4. Maybe the most important result is that we find no evidence of herding or imitative behavior in the presence of a flexible market price. While this aspect is consistent with the AZ model, other theoretical predictions of the AZ model find no support in the data. Recall that the AZ model predicts that all subjects follow their private information. In the experiment this happens only in between 50 and 70 % of cases. Clearly, such behavior yields substantial deviations of actual prices from *theoretical* prices which would obtain if everyone behaved rationally. We find that, on average, actual prices are less extreme than theoretical prices which implies that volatility in the actual market is lower than it should be theoretically. In light of the popular view that irrational traders are partly responsible for the large swings observed in financial markets, this is an interesting finding. While we do not observe herding, we find considerable support for the existence of “contrarian” behavior. When the price of asset A is high, subjects often buy asset B even if their own private information and the decisions of their predecessors favor asset A , and vice versa. Since we find that contrarian behavior can be profitable at very low or very high prices, we explore the possibility that subjects have doubts about the rationality of others and consequently mistrust their decisions. We find that error models (as in the quantal response models of McKelvey and Palfrey, 1995 and 1998), which explicitly take into account the possibility of mistakes, are partly able to rationalize contrarian behavior.

The large number of participants allows us further to conduct a number of interesting comparisons of behavior with respect to demographics, fields of studies etc. There seems to be no significant difference between male and female subjects, or between subjects with and without college education. Ph.D.’s and Ph.D. students, however, are slightly more in line with theoretical predictions. Maybe it does not come as a surprise that when we look at selected fields of studies, physicists and mathematicians perform best in terms of “rationality” (i.e. performance according to theory) and psychologists worst. However, since “rational” behavior

⁷Conducting experiments on the internet is still novel. For first experiences, see e.g., Forsythe et al. (1992, 1999), Lucking-Reiley (1999), Anderhub et al. (2001), Charness et al. (2001), Shavit et al. (2001), Bosch-Domènech et al. (2002), and Güth et al. (2002). For technical issues, see e.g., Greiner et al. (2002). The internet is also used to provide a platform to run economic experiments for interactive learning (Holt, 2002).

is only profitable when other subjects also behave rationally, good performance in terms of “rationality” does not imply good performance in terms of profits. And indeed, the ranking in terms of profits is just the opposite: psychologists are best and physicists are worst.⁸ Finally, it is reassuring that the consultants in our control experiment did not behave significantly different from the subjects in the main experiment, which is important for the outside validity of our experiment. Section 5 contains a conclusion.

2 Theoretical predictions

Consider a number of investors who have to decide sequentially whether to invest in one of two assets, A or B . For simplicity each investor can only buy either one unit of asset A or one unit of asset B (sometimes we also allow for the possibility that no trade occurs). Investors are risk neutral and have the same a priori beliefs regarding the probabilities of success of the two investments. Specifically, only one asset is successful and worth 10 units at the end of the period while the other is worth 0. Each investor believes a priori that asset A 's probability of success is $P(A)$.

The timing is as follows. Investors move sequentially in some exogenous order with each investor moving only once. Before deciding what to buy, each investor receives a private, informative signal a or b regarding the success of the assets. The signal's precision is $P(a|A) = P(b|B) > 0.5$, which is the conditional probability that signal $s \in \{a, b\}$ is given when the true state is $S \in \{A, B\}$. For all investors, the signal is identically and independently distributed conditional on the true state. This is commonly known. Each investor can observe the decisions of all his predecessors.

We consider two principal versions of this model: one in which the prices for the two assets are fixed (and normalized to zero) and one in which the prices are market prices that reflect publicly available information. The version with zero prices is equivalent to the basic model studied by BHW. The model with market prices has been studied by AZ.

2.1 The BHW model

All investors can invest either in asset A or B – but not both – at zero cost. Clearly, an investor with t predecessors will choose A if and only if the conditional probability that A is successful

⁸Luckily, economists' performance turns out to be closer to psychologists than to physicists.

given all private and public information $P(A|H_t, s)$ is greater than $1/2$,⁹ where H_t denotes the observable history of the decisions of all predecessors up to round t , and $s = a, b$ the private signal.

The difficulty lies in the interpretation of decisions of predecessors. Assuming that all predecessors are perfectly rational Bayesians, an investor who is a Bayesian himself follows his private signal and thereby reveals it, unless an informational cascade has started. If a signal can be deduced from the chosen action, it is called an imputed signal.

A cascade on asset S , a S -cascade, starts when an investor should buy asset S regardless of his own signal, i.e. when $P(S|H_t, s) > 1/2$, for $s = a, b$. Depending on the a priori probabilities and the signal precisions, this requires a different number of (imputed) a or b signals. In all cases, however, the onset of a cascade depends only on the net number of signals $\Delta = \#a - \#b$ which can be imputed from the history of decisions (i.e., Δ is defined net of the signal of the current investor).

We demonstrate the calculations for our main probability combination 55-60, that is, $P(A) = 0.55$ and $P(a|A) = P(b|B) = 0.6$. The first investor should always follow his own signal since even if he receives a b -signal $P(B|b) = 1 - P(A|b) > 1/2$ holds. Hence, the signal of the first player can be imputed from his action. If the first investor chooses A , the second should already disregard his own signal: even with a b signal, the second investor should choose A since

$$P(A|ab) = \frac{P(ab|A)P(A)}{P(ab|A)P(A) + P(ab|B)P(B)} = 0.55, \quad (1)$$

which is the a priori probability for A . The two signals a and b cancel out and the decision should follow the a priori probability. In this case the third player cannot impute the signal of the second player and thus faces a similar decision problem as the second player. Hence after one A , an A -cascade starts, i.e. when $\Delta \geq 1$. Likewise it can be shown that a B -cascade must start for $\Delta \leq -2$. If all agents are rational Bayesians, a cascade is never broken once it started, and information accumulation stops (i.e. $-2 \leq \Delta \leq 1$ at all times unless errors occur).

In an experiment, one can hardly assume that all subjects are rational Bayesians, let alone that all subjects believe that all *other* subjects are rational. In particular, one has to make provisions for the fact that irrational behavior may be unambiguously observed (as when the second subject chooses B following an A by the first subject in the example above).

⁹In most of our treatments, ties in expected profit cannot occur. When a tie-breaking rule is required, it is explicitly mentioned below.

To account for possible non-Bayesian behavior, we assume that subjects impute signals in the following way. If a decision is not in obvious contradiction to Bayes' rule, the imputed signal equals the decision unless a cascade has started, in which case no signal can be extracted from the respective decision. For a decision which, given the history of imputed signals, obviously violates Bayes' rule, we considered two variants: (a) successors ignore the decision of the deviator and (b) subjects assume that the deviator followed his private signal. As it turns out, the empirical truth lies somewhere in the middle but both assumptions yield qualitatively the same results. In the following, we only reports results based on rule (b). We say that an agent is "rational under common knowledge of rationality" (henceforth: *ruck*) if he follows Bayesian updating with respect to the imputed signal history and his own signal. In the following, it is important to remember that *ruck* is based on the assumption that all predecessors are rational. It may be perfectly rational to deviate from *ruck* if one believes that one's predecessors are not rational.

Of course, no ambiguity with regard to rationality arises when not only decisions but also signals of others are observable. In this case, an optimal decision is purely a matter of calculating conditional probabilities. As mentioned above, we also consider this setting in one of our treatment.

2.2 The Avery/Zemsky model

To keep the experiment as simple as possible, we consider the simplest version of the AZ model (cf. Avery and Zemsky, 1998, Section I), which is the BHW model enriched by a flexible price. In this simple model, the price is set by a market maker who efficiently incorporates all publicly available information.¹⁰ The crucial question is how the existence of a market price changes the possibilities for herding.

Let p_t denote the market price of asset A in round t and assume that a successful asset pays out 10 units in the end. Hence,

$$p_t = 10P(A|H_t). \tag{2}$$

The price of B is always equal to $10 - p_t$ since $P(A|H_t) = 1 - P(B|H_t)$.

¹⁰In contrast to AZ's general model which is in the spirit of Glosten and Milgrom (1985), the simple model does not incorporate uninformed traders and therefore has no bid-ask spread. With only informed traders, setting a bid-ask spread to ensure a zero profit condition for the market maker would lead to a market breakdown. However, the results of the simple model carry over to a more complex world with informed and uninformed traders and a market maker setting bid-ask spreads (see AZ, Proposition 3). As in the experiment the market maker was played by the computer, the possibility of losses was not an issue.

The decision of an investor is straightforward. An investment in A is profitable in expectation if and only if

$$10P(A|H_t, s) - p_t > 0 \tag{3}$$

that is, if and only if $s = a$. Likewise, an investment in B is profitable if and only if $s = b$. In other words, each investor follows his private signal. All information is revealed, and therefore it is incorporated into the price immediately after each decision. This implies that the price is semi-strong efficient, i.e., at any point in time the price incorporates all publicly available information. The price is a martingale with respect to public information, i.e. $E(p_{t+1}|H_t) = p_t$ for all t , and one cannot take advantage of the knowledge of historical price movements to earn superior returns. As everyone follows his signal, rational herding cannot occur.

Note that not trading is never optimal unless one introduces transaction costs because subjects always have an informational advantage over the market maker. At some prices, transaction cost induce rational agents not to trade irrespective of their signal. Once such a price is reached, the market breaks down since no rational agent will buy anymore and market prices remain constant. Thus, in the presence of transaction costs the market price need not converge to the true value of the asset and even in the long run substantial mispricing is possible. In the presence of a flexible market price, we say that an agent is *ruck* if he follows his signal, or does not trade if transaction costs are too high, respectively.¹¹

Again, the problem becomes more complicated when investors cannot be fully confident that their predecessors behaved rationally. Suppose instead the investor believes that (some) prior decisions were taken randomly. A *regression to the mean* argument implies that high prices are likely to be overvalued and low prices undervalued. Thus, it may pay for an investor to trade against the market and against one's own signal. Such investors are called *contrarians* even though rational contrarians would never occur in our setting if all investors were known to be rational.

¹¹With respect to the price setting rule, we assume that the market maker holds the belief that all decisions are formed according to *ruck* unless this can be unambiguously rejected. Given an irrational no trade decision, no signal can be imputed, and the price remains constant. If the market maker unambiguously observes an irrational A or B decision (this can only happen in the presence of transaction cost when no trade is rational) the market maker holds the belief that the decision reflects the agent's signal.

3 Experimental design

More than 6000 subjects participated in our online experiment which was available for a period of about six weeks in the spring of 2002 on our web site <http://www.a-oder-b.de> which is German for *a-or-b*. Subjects decided in sequence and were able to observe the actual decisions of prior participants in their respective groups. In general, the group size was 20. 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.

Common to all treatments are the following features. Subjects had to choose between investment opportunities A and B (in some price treatments, there was also the option of choosing neither which we label N). Only one of the two could be successful and, if so, would pay 10 “Lotto-Euros”. The unsuccessful investment paid nothing. Subjects were told the a priori probabilities $P(S)$, which varied among our treatments. Furthermore, they were told that they would receive a tip by an investment banker which was reliable with a specified probability $P(s|S)$, which also varied among our treatments. Subjects were informed that all prior investors 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).

In the next subsection we introduce the details of the different treatments.

3.1 Treatments

Given the large number of participants we were able to explore a variety of different questions, information conditions, probability combinations, etc.. In this paper, we focus on treatments that are relevant to financial markets, that is, treatments which follow the basic set up of AZ where a market price exists which reflects all public available information. For comparisons, we also include two treatments without prices which follow the basic model of BHW.

Table 1 lists the main features of all treatments. Treatment names with market prices start with P followed by $+D$ when, additional to the price history, the decisions of all prior investors are observable, or $-D$ when only the price history is observable. Hence, in treatments $+D$ subjects could observe which effects decisions had on prices. A $-N$ denotes treatments in which the “no trade” option was absent, i.e. subjects were forced to buy either A or B . If all agents act in line with theory the different price treatments without transaction costs have the same theoretical prediction: everyone should follow his own signal. Whether past decisions are

observable or not is irrelevant since the price history fully reveals past decisions such that the decision history yields no additional information. Furthermore, the no-trade option N should not alter the result because in the absence of transaction costs, not trading is never optimal. In treatments where transaction costs are present $+T_i$ is added to the treatment name. There were two levels of transaction costs: $i \in \{0.1, 0.5\}$

Table 1: Treatments

treatment		
group	treatment	description
P	$P+D$	Price, all decisions observable*
	$P-D$	Price, no decisions observable
	$Pe+D$	$P+D$ with explicit formulation*
$P-N$	$P+D-N$	Price, all decisions observable, N not possible*
	$P-D-N$	Price, no decisions observable, N not possible
$P+T_i$	$P+D+T_i$	$P+D$ + transactions costs of size i
	$Pe+D+T_i$	$Pe+D$ + transactions costs of size i
	BHW	Bikhchandani/Hirshleifer/Welch
	$BHW+AS$	BHW + all signals observable

Note: N denotes the option of not trading. * these treatments also offered the option of placing a bet on the final price in the respective group.

To explain the notion of an efficient market price to subjects is not a simple task. To check whether a more explicit description of the price process makes a difference we included treatments which additionally mentioned that prices are conditional expected values given the history of decisions. Such treatments are denoted by an additional e .¹² Finally, treatments without prices are denoted by BHW . The no-price treatment in which also *all signals* of predecessors were observable is denoted by $BHW+AS$.

While in the price treatments cascades should never happen regardless of the probability parameters of the model, in the no-price treatments the likelihood of cascades crucially depends on the a priori probability of the true state and the precision of the private signals. We have therefore looked at a number of different probability combinations shown in Table 2. For each probability combination, the last two columns in Table 2 give the minimum (maximum) net number of imputed signals necessary for the start of an A cascade (B cascade), respectively.

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.

¹²We have also used two versions of the instructions, an old and a new version, the latter being more precise about the independence of the investment bankers’ tips. See Section 4.10 for more on this.

Table 2: Probabilities

a priori probability $P(A)$	signal precision $P(a A)$	ΔA -cascade	ΔB -cascade
55	60	1	-2
51	55	1	-2
55	80	1	-2
50*	66	2	-2
60*	60	1	-3
60	51	-9	-12
60	55	-1	-4

Note: * tie-breaking rule: follow own signal when indifferent.

In the price treatments subjects received additionally an endowment of 11 for each task to avoid losses because they had to pay the market price for their investment (which could vary between 0 and 10) and in some treatments the transaction costs of 0.1 or 0.5. Thus, the payoff from each task was $11 - \text{market price} + 10$ (if successful) $-$ transaction costs (if applicable).

From the perspective of efficient market theory it is not only interesting whether the price incorporates all information efficiently but also whether subjects perceive the price to be a martingale. Therefore, treatments $P+D$, $Pe+D$ and $P+D-N$ offered an additional chance to make money by placing a bet \hat{p}_{T+1} on the market price p_{T+1} for asset A in this group, i.e. the market price after the last player in the group had made his decision.¹³ This “price bet” was rewarded as follows:

$$\max [5 - |\hat{p}_{T+1} - p_{T+1}|, 0]. \quad (4)$$

3.2 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

¹³The bet was offered to subjects after they took their decision but without revealing the new price.

in a lottery for 11 prizes of 1000 Euro each. While the odds in those lotteries were fixed in advance and known to subjects, they were diminishing over time. In phase I of the experiment, 1409 subjects played with high powered incentives where each of 40000 lottery tickets had an equal chance of winning one of 5 prizes of 1000 Euros. Since subjects played on average for about 15 minutes, they were making an expected hourly “wage” of 14.19 Euros, which is comparable to a very good student job and to pay in laboratory experiments. In phase II, each of 90000 lottery tickets had an equal chance of winning one of another 5 prizes of 1000 Euros. Finally, in phase III, all remaining 1162 subjects competed for the last 1000 Euros which amounted to almost no monetary incentive. This payment scheme was due to the fact that an unexpected large number of subjects participated in our experiment. But it also gives us the chance to test the role of incentives in such a setting. The 11 winners were notified by mail 2 weeks after the experiment ended and their prize money was paid through bank transfers.

Additionally, there was a control group of 267 consultants of 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 consultants of this firm. 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.

3.3 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).

- Insert Figure 1 here -

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

¹⁴788 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).

¹⁵Given that each time when 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

Considering the number of Ph.D. students and Ph.D.’s we believe we succeeded in recruiting a fairly bright subject pool.

3.4 Implementation

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

¹⁶It will never be possible to completely prevent clever people from playing more than once. However, we are confident that not many such attempts were successful, and given the size of the subject pool, those few probably do not matter much. It turned out that less than 5% of the participants entered an invalid email address (including unintentional typos), and these subjects did not behave significantly different (see Section 4.9).

¹⁷A 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.

4 Results

For the evaluation of the results we shall consider the following 3 measures. (1) Average *ruck* is defined as the fraction of subjects who behaved rationally under the assumption of common knowledge of rationality of predecessors. (2) The fraction of cases in which subjects followed their own signal is denoted by *own*. (3) Finally, the actual market price p_t is compared to the *theoretical price* p_t^* , which would have resulted had all subjects decided rationally, and to the full-information price p_t^F , which would result if the market maker could directly observe the signals. Note that in the absence of transaction cost, $p_t^* = p_t^F \forall t$.

Before we present the results it might be useful to collect the theoretical hypotheses for our various treatments.

1. In all P and $P-N$ treatments, subjects should follow their own signal $\rightarrow ruck = own = 1$.
2. There should not be any difference between any of the treatments in P and $P-N$, regardless of whether prior decisions are observable or not ($+D$ or $-D$) and whether the option N was available or not.
3. Actual prices p_t should match theoretical prices p_t^* .
4. The no trade option (N) should only be observed in the transaction cost treatments ($P+T$) and even there only in some well defined circumstances.
5. In treatment BHW , subjects should follow the cascade behavior described in the last two columns of Table 2 *if* they believe that their predecessors are rational $\rightarrow ruck = 1$.
6. In treatment $BHW+AS$, subjects should follow the cascade behavior of Table 2 *regardless* of what they believe about others $\rightarrow ruck = 1$.
7. Different prior probabilities and signal precisions should not alter average *ruck*.

4.1 Preliminary data analysis

As stated above the number of subjects that decided in sequence was usually 20 (it was 10 for treatment $BHW+AS$ and the control experiment with consultants). Table 4 gives the number of groups that participated in our experiment, separately for each combination of treatments and probabilities. Since reporting results for each variation would be tedious, we checked first

Table 4: Number of groups for treatments

treatment	probability combination							total
	50-66	51-55	55-60	55-80	60-51	60-55	60-60	
<i>P+D</i>	12	6	26+15*	8	-	6	4	77
<i>P-D</i>	11	6	20+9*	1	-	-	3	50
<i>Pe+D</i>	12	6	12	6	-	6	-	42
<i>P+D-N</i>	16	6	18	-	-	-	-	40
<i>P-D-N</i>	17	6	18	-	-	-	-	41
<i>P+D+T_{0.1}</i>	12	6	20	8	-	6	4	56
<i>P+D+T_{0.5}</i>	12	6	26+24*	2	-	-	4	74
<i>Pe+D+T_{0.1}</i>	-	-	-	6	-	6	-	12
<i>Pe+D+T_{0.5}</i>	12	6	12	-	-	-	-	30
<i>BHW</i>	12	17	65+15*	8	12	6	15	150
<i>BHW+AS</i>	12	18	70+9*	2	12	-	12	135
total	128	83	287+72*	41	24	30	42	701

Note: As a convention the first number in x-y is the prior and the second the signal precision used in the experiment. *denotes groups in the control experiment with consultants.

which variants of our treatments can be grouped together and pooled. We did this by comparing treatments with respect to the variables *ruck* and *own* both by non-parametric MWU tests and logit regressions taking each group as one observation. The following summarizes the results of those tests:

- The phase of the experiment (recall that incentives were different in phases I, II, and III) had no significant effect. This implies that – at least in this experiment – incentives seem to matter little as compared to the intrinsic motivation of subjects to perform well.¹⁸
- The stage of the task (whether a task was the first, second, or third a subject performed in) did not matter except in transaction cost treatments (see Section 4.7). This shows that learning effects do not play a significant role which is in line with our expectations given that subjects did not receive any feedback until the end of the game.
- The two versions of the instructions (see Footnote 12) made no significant difference. Also, a more explicit formulation for the price formation process (*e*) was irrelevant which, given that a large fraction of the subject pool has a mathematical background (see Figure 1), is reassuring as to that subjects understood how prices were formed.

¹⁸Camerer and Hogarth (1999) provide a survey of studies which look at the effects of monetary incentives.

- Treatments that belong to the same treatment group (see Table 1) were not significantly different. In particular, the observability of decisions of predecessors did not matter which implies that subjects understood how their predecessors' decisions are reflected in the price history.

Henceforth, unless otherwise stated, we will pool data from treatment variants that did not significantly differ. In the following, we will present the results from the main experiment. Results from the control experiment with consultants are reported in Section 4.10.

4.2 Summary statistics

In this subsection, we present summary statistics on the main variables of interest, *ruck* and *own*. Recall that theory predicts $ruck = 1$ for all treatments, and $own = 1$ in price treatments. Table 5 shows average *ruck* for 5 treatment groups and the probability combinations 50-66 and 55-60.¹⁹ Average *ruck* in *BHW* reaches 78% in 50-66 but only 66% in 55-60. Surprisingly, in *BHW+AS*, where there is no need to worry about the rationality of others, *ruck* reaches only 76% and 72%, respectively.

Table 5: Average *ruck* and *own*

	prob. comb.	<i>BHW</i>	<i>BHW+AS</i>	<i>P</i>	<i>P-N</i>	<i>P+T</i>
<i>ruck</i>	50-66	.78	.76	.59	.68	.53
		(.09)	(.12)	(.11)	(.09)	(.13)
	55-60	.66	.72	.54	.64	.50
		(.11)	(.14)	(.11)	(.09)	(.10)
<i>own</i>	50-66	.75	.69	.59	.68	.54
		(.12)	(.12)	(.11)	(.09)	(.13)
	55-60	.75	.74	.54	.64	.50
		(.13)	(.14)	(.11)	(.09)	(.10)

Note: Standard deviations in parentheses.

In the *P* and *P+T* treatments, however, *ruck* is even lower, partly because an additional mistake can be made, namely not to trade. In fact, *N* was chosen on average 20% of the time in the *P* treatments, which was never *ruck*.²⁰ With transaction cost, *N* was chosen 26% of the

¹⁹Since in the *BHW* treatments the observed *ruck* varies with different probability combinations, we compare treatments only for our main probability combinations 50-66 and 55-60.

²⁰These findings are in line with Cipriani and Guarino (2002). In their flexible price treatment (given their probability combination 50-70) they report that 62.7% (23.8%) of subjects acted rationally (did not trade) which is comparable to our treatments *P* given the probability combination 50-66 where the respective numbers are 59.4% and 24.4%.

time but only 12% of the no trade decisions were *ruck*.²¹ In the *P-N* treatments, in which no trade was not an option, average *ruck* reaches 68% and 64%, respectively.

For both probability combinations, non-parametric MWU-tests reveal that *ruck* is significantly different between all treatments except in the following cases.²² For 50-66, the differences between *BHW* and *BHW+AS*, and *BHW+AS* and *P-N* are not significant. For 55-60, the difference between *BHW* and *P-N* is not significant. These findings show that it is not so much the different treatments but rather the additional *N* option which lowers rationality. This is supported by the finding that, given probability combination 50-66 (55-60), average *ruck* among subjects who bought either *A* or *B* is 71% (67%) in treatments *P* as well as in *P+T*.

Table 5 also reveals a curious pattern with respect to the percentage of subjects who followed their own signals (*own*). In the *BHW* treatments, 75 % of subjects follow their own signal (often in contrast to what they should do). On the other hand, in the price treatments only between 50% and 68% of subjects follow their own signal, when following the own signal is always rational. According to MWU tests, the difference between treatment *BHW* and each price treatment is significant at the 1% level except between *BHW* and *P-N* for 50-66 where the p-value is 0.07. The differences between *BHW* and *BHW+AS* are not significant.

Finally, we can take a more detailed look how *ruck* (which, in the absence of transaction costs, coincides with *own*) varies for different probability combinations. Table 6 shows *ruck* levels for treatment *P*. Taking our main treatment 55-60 as base, only 50-66 shows a significant difference at the 5% level according to MWU tests.²³ Hence, results from the price treatments seem to be fairly robust across probability combinations.

Table 6: *ruck* in treatment *P* across probability combinations

Prob. combinations	50-66	51-55	55-60	55-80	60-55	60-60
<i>ruck</i>	.59	.56	.54	.53	.51	.60
	(.11)	(.13)	(.11)	(.11)	(.13)	(.10)

Note: Standard deviations in parentheses.

²¹Risk aversion cannot really account for this behavior since we employed a binary lottery procedure which – at least theoretically – should induce risk neutrality (see Roth and Malouf, 1979, or Berg et al., 1986).

²²All at a 1% significance level except between *P* and *P+T* where it was 5%.

²³Differences across probability combinations look similar in treatment *P-N* but not all combinations were played for this treatment group.

4.3 Actual versus theoretical prices

The key question with respect to the efficient market hypothesis is whether prices accurately reflect the information in the market. Clearly, informational efficiency presupposes that individual traders act rationally on their information and, as we have seen in the last subsections, this is not always the case in our experiment. Consequently, in this subsection we want to look at how strongly prices are distorted relative to two theoretical benchmarks. We define the *theoretical price* p_t^* as the price that would have resulted if all traders had behaved according to theory (i.e. *ruck*). For P and $P-N$, the theoretical price is equivalent to the *full-information price* p_t^F , i.e. the price the market maker would set if he could directly observe the private information of traders. In treatment $P+T$, the theoretical and full-information prices may differ as in this treatment it can be rational to choose N .

While in the long-run the full-information price will converge to the true value of the asset, this is not necessarily the case in the short or medium-run. Also note that in treatment $P+T$ even in the long-run the theoretical price need not converge to the true value because once it is rational to choose N , it is also rational for all successors to choose N , and hence the theoretical price remains stuck at the respective level which may not even be in the direction of the true value.

- Insert Figure 2 here -

To judge how well actual prices incorporate available information, we first look at convergence of final prices to full information final prices. Figure 2 shows the fractions of actual final prices p_{T+1} , i.e. prices after the decision of the last player, which were within an interval of ± 1 (± 0.5 , respectively) of the final full-information prices, p_{T+1}^F . As can be seen, final prices are rarely close to the final full information price. In $P-N$ the final price is within ± 1 of the full information price in only 21% of groups. The option of not trading makes things worse since now only 14% of end prices are close to full information prices. As can be expected this percentage, is even lower with transactions costs.²⁴ To have a closer look at convergence properties, we focus on treatment P for the remainder of this subsection. Results are qualitatively the same for $P-N$. The effects of transaction costs on convergence are discussed in more detail

²⁴However, only the difference between $P-N$ and $P+T$ (P) for ± 0.5 is significant at the 1% (10%) level of a MWU-test.

in Section 4.7.

Figure 3 depicts the distribution of the deviation ($p_{T+1}^* - p_{T+1}$) of the final price of asset A from the final theoretical price of asset A separately for states where asset A is successful (left panel) and where asset B is successful (right panel). Recall that the price of B is 10 minus the price of A . An important observation is that theoretical prices tend to be more extreme than actual prices. In other words, actual prices undershoot: when state A (B) is true, the actual final price of A is often too low (high).

- Insert Figure 3 here -

In a next step, in Figure 4 we consider the empirical distribution of actual prices (upper panel) and theoretical prices (lower panel) over time.²⁵ Unless the signal precision is very high (as with probability combination 55-80), theoretical prices are concentrated around the a priori price. The distribution of actual prices matches this feature well but the distributions are even more concentrated. Extreme prices are rarely observed. It is interesting to see that even though a high fraction of subjects behave seemingly irrationally, the volatility in the actual markets is less than the volatility in the theoretical markets. This holds for all probability combinations and is significant at a 1% level in all cases except 55-80 and 60-60 where it is not significant.

- Insert Figure 4 here -

That deviations between the theoretical and actual price can be severe is illustrated in the left panel of Figure 5 which shows the average deviation ($p_t^* - p_t$) across theoretical prices p_t^* (not only final ones). The graph clearly indicates that actual prices undershoot and are less extreme than the theoretical prediction. Since the same signal imbalance has very different effects on the price level in the different probability combinations, we also plot the average deviation for a given theoretical signal imbalances Δ in the right panel of Figure 5, which shows this relationship even more clearly.

- Insert Figure 5 here -

Finally, over time deviations become more severe. The left panel of Figure 6 shows that the average absolute price deviation increases. Note that by construction of the price mechanism,

²⁵Actual (theoretical) prices include the decision (signal) of the current subject.

deviations in the early rounds must be small. More interestingly, there is no inverted U-shape form which would indicate convergence over time. Rather the opposite is the case (see right panel of Figure 6).²⁶ In the next subsection we shall consider possible explanations for those deviations.

- Insert Figure 6 here -

4.4 Possible explanations: imitation, contrarians, and regression to the mean

An obvious candidate explanation for observed deviations from the theoretical prices is herd behavior. As explained in Section 2.2, rational herding is impossible in our framework. However, if subjects indeed behave like “lemmings”, they would imitate prior decisions and thus produce herd-like behavior. We find, however, that imitation plays no significant role.

The first evidence against imitation stems from inspecting Figures 3 through 6, which show that on average actual prices are less extreme than theoretical prices. Yet, imitation would predict the opposite because if, for example, an early investor in asset A induced later investors to buy A even though they got a b signal, this would quickly drive the actual price above the theoretical one. Likewise, an early investor in B should drive down the price of A more quickly than the theoretical price since imitation would yield more buyers of B than justified by private information.

To check more rigorously for imitation on the individual level, we analyze whether subjects’ choices are influenced by the decisions of their direct predecessors. If a subject is more likely to choose A after (say) three subjects immediately before him also chose A , we would take that as an indication for imitative behavior. Let $pred$ measure the number of direct predecessors who chose an identical action. The variable is positive if this action was A and negative if the action was B .²⁷ In Table 7, we report results from logit regressions for treatment $P-N$ in order to test whether the 0-1 variable “choice of A ” is influenced by $pred$.²⁸ Besides $pred$, the following explanatory variables were used: (i) $\tilde{\Delta}$ represents the net number of A decisions

²⁶The observed divergence over time may explain the higher levels of convergence reported by Cipriani and Guarino (2002) who consider groups of 12 subjects.

²⁷Alternatively, we included dummies if the number of A or B predecessors is greater than a certain threshold. This change was inconsequential.

²⁸In treatments P we ran ordered logits with the endogenous variable taking the values 1 (choice of A), 0 (no trade), and -1 (choice of B), and obtained qualitatively the same results as in $P-N$. Since, from a theoretical point of view, N should never be chosen, it is not perfectly clear whether A , B and N can indeed be ordered. Hence, we also ran multinomial logits which also qualitatively yielded the same results. The independence of irrelevant alternatives underlying the multinomial model is fulfilled.

Table 7: Logit analysis: choice of A in treatment P - N

obs. used	$R1$ all	$R2$ only obs. with $\tilde{\Delta} = 0$	$R3$ only obs. with $ruck = 0$
$\tilde{\Delta}$	-0.354** (0.042)	—	—
$dhint$	1.549** (0.118)	2.058** (0.226)	—
$pred$	-0.017 (0.036)	-0.022 (0.081)	-0.240** (0.048)
Observations	1479	408	495
-2 Log-Likelihood	1775.20	469.99	659.11
R^2 (Nagelkerkes)	0.226	0.278	0.071

Note: ** significant at 1%-level; * significant at 5%-level; standard deviations in parentheses; all probability combinations used; all regressions include dummies for probability combinations and a constant.

up to round t , i.e. $\tilde{\Delta} = \#A$ decisions $- \#B$ decisions. The variable $\tilde{\Delta}$ is a one-to-one mapping of the actual price but is easier to aggregate across probability combinations. (ii) $dhint$ is a dummy variable that is 1 if the private information of the subjects in question was a and 0 otherwise. Finally, all regressions included dummies for the various probability combinations and a constant. Regression $R1$ includes all observations. Regression $R2$ includes only observations where $\tilde{\Delta} = 0$; in these cases, imitation cannot be discouraged by extreme prices.²⁹ Finally, regression $R3$ includes only observations in which subjects deviated from their private signal. Note that this implies that $dhint$ cannot be included in $R3$ as an explanatory variable.

None of the regressions produce a positive coefficient for the imitation variable $pred$. In regressions $R1$ and $R2$, $pred$ is negative but not significantly different from zero. In $R3$, the coefficient is even significantly negative. The variable $\tilde{\Delta}$ and the dummy for private information $dhint$ have the expected signs and are significant at the 1% level. It seems that in the presence of a flexible market price, imitation cannot explain our data.³⁰

If imitation is not the right story to explain the deviations between the theoretical and actual price, what is? We suggest a story based on contrarian behavior which can be justified by a

²⁹Regressions for $\tilde{\Delta} = -1$ and $\tilde{\Delta} = 1$ qualitatively yielded the same result as $R2$.

³⁰A regression analysis confirms that there is also no imitation with respect to the no-trade decision N .

regression to the mean argument. As explained above, we say that a subject is a *contrarian* if he trades against his signal and against the market, or equivalently if he receives an a (b) signal at a price for A which is strictly above (below) the a priori price, i.e., 10 times the a priori probability for A , and deviates from his signal to buy B (A) instead. Such contrarian behavior can only be (ex post) rational if the trader is convinced that prior traders irrationally drove the price to an extreme.

Figure 7 gives evidence in favor of contrarian behavior. The four panels show average frequencies of *ruck* and for trading against one's signal for treatment $P+D$ and probabilities 55-60.³¹ The higher the price, the more likely subjects are to trade against their signals and the lower *ruck*. Average *ruck* drops from around 80% at low prices to around 30% at high prices. While trading against one's signal shows exactly the opposite trend, the propensity not to trade remains roughly constant at around 20% across prices. Contrarian behavior is also strongly suggested by Figures 3 through 5.

- Insert Figure 7 here -

These findings are confirmed in our regression analysis (see again Table 7). High prices for A , or equivalently, high $\tilde{\Delta}$, lower the probability of choosing asset A at a 1% significance level in regression $R1$.³² Moreover, in regression $R3$, $pred$ has a significant negative effect which indicates that subject have a tendency to trade against the decisions of their predecessors.

The question arises, however, whether it is a good idea to be a contrarian? There is a paradox here since contrarian behavior is profitable if and only if prices are overvalued, but if there are many contrarians, prices are not overvalued. We will present two possible explanations for this apparent puzzle. The first is purely based on contrarian behavior, the second on random behavior by some traders and regression to the mean.

The first explanation is best illustrated by an actual price path in our experiment (see Figure 8). Suppose the price of A falls below the a priori price and there is one contrarian. This price gap is preserved in all future rounds unless someone else deviates from his signal. Suppose next that due to a series of a signals the price of A rises above the a priori price just as in the

³¹Other treatments and probability combinations look similar.

³²Cipriani and Guarino (2002) run similar regressions. In the absence of transaction costs, they do not find a significant effect of $\tilde{\Delta}$. In their transaction costs treatment, $\tilde{\Delta}$ has the expected sign and is significant at the 5%-level.

example. Now a trader would actually be better off being a contrarian in this later rounds. Thus, contrarian behavior is compatible with overshooting, which is sometimes observed in the data (see Figure 3).

- Insert Figure 8 here -

The second explanation assumes that some traders behave like noise traders, in particular, they choose A and B with equal probabilities. This implies that whenever the actual price is very high or very low, it is likely that this was driven by noise traders. In other words, whenever the actual price is extreme, the theoretical price is likely to be less extreme. Vice versa, whenever the theoretical price is extreme, the actual price is likely to be less extreme, which is simply a regression to the mean argument. Given that traders anticipate the random behavior of noise traders, they should be contrarians because a low price for A yields a buying opportunity even if the own private information is favoring B . Given that there are some contrarians among the traders, actual prices will on average be less extreme than theoretical ones, just as observed in our data (see Figure 4).

4.5 Does it pay to be a contrarian?

The previous subsection shows that contrarian behavior has at least the potential to be profitable. But is it actually profitable? A first aggregate look at the data suggests that this is not the case. Table 8 lists average profits (excluding the fixed payment of 11 in the price treatments and the profit from the price bet) of subjects depending on whether subjects were *ruck* or not. In all cases, *ruck* yields higher average profits. The last column of Table 8 contains Pearson correlation coefficients between *ruck* and profits. All correlation coefficients are positive and significantly different from zero at the 1% level.³³

We say a subject is a *potential contrarian* if he faces a price for asset S strictly higher than the a priori price and receives a signal in favor of S , where $S \in \{A, B\}$. In Table 9, we compare the average profits of contrarians with that of potential contrarians who did not become contrarians either because they followed their signal or because they did not trade at all. In treatment groups P and $P+T$, contrarians did worse than both other groups. However, in $P-N$ contrarians fared better than *ruck* subjects.

³³This also holds in BHW and $BHW+AS$. On average, *ruck* players earned 6.1 in BHW and $BHW+AS$ whereas others earned 4.4 in BHW and 3.4 in $BHW+AS$.

Table 8: Profits and *ruck*

treatment	av. profits of subjects classified as		corr. coeff.
	<i>ruck</i> = 1	<i>ruck</i> = 0	
<i>P</i>	1.3	-0.5	0.20**
<i>P-N</i>	1.0	-0.4	0.13**
<i>P+T</i> _{0.1}	1.1	-0.4	0.18**
<i>P+T</i> _{0.5}	0.6	-0.5	0.13**

Note: pooled over all probability combinations, profits exclude fixed payments and the profit from the price bet, ** indicates significance at the 1% level.

Table 9: Profits of (potential) contrarians

treatment	profits of potential contrarians who were		
	contrarians	rational	no-traders
<i>P</i>	-0.6	0.9	0
<i>P-N</i>	0.1	-0.5	-
<i>P+T</i> _{0.1}	-0.7	0.2	0
<i>P+T</i> _{0.5}	-0.9	0.5	0

Note: pooled over all probability combinations, profits exclude fixed payments and the profit from the price bet.

It seems plausible that contrarian behavior is sensible if actual prices are extreme. To check this, we compare the profit of a contrarian with the *counterfactual* profit he would have received had he played according to theory. Since $p_t = (10 - \text{price of } B)$, it is straightforward to show that, in the absence of transaction costs, the counterfactual payoff is always (-1) times the actual profit.

Figure 9 shows the average profits of contrarians across actual prices for our main probability combination 55-60 in treatments *P-N* (left panel) and *P* (right panel). For extreme prices a contrarian is actually more successful than a counterfactual, *ruck* subject would be.

- Insert Figure 9 here -

For other probability combinations the evidence is more mixed. If we pool over all probabilities and consider prices that deviate by more than two from the a priori price p_0 , we get the results presented in Table 10. In treatment *P*, contrarians always do worse than their counterfactuals whereas in treatment *P-N* contrarians do better for prices that are more than two away from p_0 .

Table 10: Contrarians' average profits in different price regions

	$p_t - p_0 \geq 2$	$2 \geq p_t - p_0 \geq -2$	$p_t - p_0 \leq -2$
<i>P</i>	-0.3	-0.7	-0.6
<i>P-N</i>	0.4	-0.2	0.9

Note: pooled over all probability combinations, profits exclude fixed payments and profits from the price bet. p_0 is the a priori price.

4.6 Two error models to explain contrarian behavior

As discussed in Section 4.4 contrarian behavior can be optimal if traders believe that others make mistakes. In this subsection, we explicitly incorporate the possibility that subjects are aware that others make mistakes. We shall do so under two different assumptions.

In the simple error model (*s*) it is assumed that each subjects believes that each of his predecessors followed his signal with a constant probability x and erroneously traded against his signal with probability $1 - x$. The average *ruck* in error model (*s*) is denoted by *ruck-s*, where we call an agent *ruck-s* if he is a Bayesian with respect to the decision history, a given x and his signal. We assume $x = 0.70$ which resembles the fraction of subjects in treatments *P* and *P-N* who followed their own signal given they chose *A* or *B*.³⁴ However the model is very robust to the exact specification of x . Error model (*s*) is rather simplistic in that it is assumed that (i) subjects neglect the fact that predecessors might have taken into account the errors of their predecessors, and (ii) the error probability is constant across price levels whereas Figure 7 clearly shows that average *ruck* is decreasing in the price level.

Error model (β) relaxes these two assumptions: based on the concept of quantal response equilibrium of McKelvey and Palfrey (1995, 1998) and analogous to Anderson and Holt (1997), error parameters β_t are estimated recursively for each round t taking into account that predecessors have reacted to possible earlier errors. Hence, this model assumes a full level of reasoning.³⁵

The estimation is done in the following way. It is assumed that subjects decide according to a logistic model with independent shocks to the expected payoff difference between assets *A*

³⁴It is also close to the level of *ruck* in *BHW+AS* in which there is no ambiguity about the definition of rationality.

³⁵Kübler and Weizsäcker (2003) show in an experiment that is similar to treatment *BHW* that full reasoning of this type will typically not be observed. However, since our two error models are extreme cases and yield (as we will see) very similar results, we suppose that a “level of reasoning” analysis would not yield very different results either.

and B . For reasons of tractability, we only estimate this error model for treatment group $P-N$. Formally, the probability that subject i decides to buy asset A (which we denote by “ $D_i = A$ ”) is given by

$$P(D_i = A|H_t, s_i) = \frac{1}{1 + e^{-\beta_t(\pi_i^A - \pi_i^B)}}, \quad (5)$$

where π_i^S is the expected profit of buying asset $S \in \{A, B\}$. The parameter β_t characterizes the sensitivity to payoff differences. Subjects buy randomly if $\beta_t \rightarrow 0$ but play rational best replies if $\beta_t \rightarrow \infty$. Since expected profits in round 1 only depend on the realization of the signal, the estimation of β_1 is straightforward with a logit regression. The estimation of subsequent error parameters is more involved in that expected profits in a certain round depend on the error parameters of all previous rounds which implies path-dependency.

To estimate β_2 , we first calculate the probability $P(D_1 = S|s)$ that in round 1 the subject chose asset S in case he received signal $s \in \{a, b\}$ taking the error parameter β_1 into account. In a second step, this information can be used to calculate

$$P(D_1 = A|A) = P(D_1|a) \cdot P(a|A) + P(D_1|b) \cdot P(b|A). \quad (6)$$

Hence, if $D_1 = A$, the posterior that asset A is successful is given by

$$P(A|D_1 = A) = \frac{P(D_1 = A|A) \cdot P(A)}{P(D_1 = A|A) \cdot P(A) + P(D_1 = A|B) \cdot P(B)}. \quad (7)$$

Combining this with private signals, one can calculate expected profits for second round players. With those, β_2 can be estimated yielding a new prior for player 3, and so on for all subsequent rounds.³⁶ Average *ruck* in error model (β) is denoted *ruck- β* .

Even though the requirements on rationality and depths of reasoning are quite different in the two models, Table 11 shows that the average *ruck* levels are remarkably similar.³⁷ We believe that this is due to the fact that in our case errors do not change the optimal actions for most prices even though expected profits do change.

This can be seen in Figure 10 which shows average expected profits of buying asset A given an a or b signal depending on the assumptions with respect to errors of predecessors (treatment

³⁶In an iterative process, we also estimated a constant parameter β across all rounds: in each iteration, expected profits were calculated by using the error parameter obtained in the previous estimation, and we observed convergence of β after approximately 20 iterations. The results from this (simplified) logistic error model are qualitatively the same as those reported.

³⁷Note that while the standard model without errors is nested both in error models (s) and (β), error model (s) is not nested in error model (β), and hence *ruck-s* > *ruck- β* , as in $P-N$, is possible.

Table 11: Average *ruck* under different error models

	<i>ruck</i>	<i>ruck-s</i>	<i>ruck-β</i>
<i>P</i>	0.55	0.59	-
<i>P-N</i>	0.66	0.69	0.67

Note: pooled over all probability combinations.

P-N, 55-60).³⁸ In the standard model, where agents hold the belief that all agents follow their signal, it is always optimal to follow one’s signal as well. Not so with error model (*s*) or (*β*): Figure 10 shows that at low or moderate prices for asset *A* (*B*), agents should optimally follow an *a* (*b*) signal but at high prices contrarian behavior is optimal.

- Insert Figure 10 here -

Given Figure 10, it is not surprising that average *ruck* at moderate prices is similar under all three definitions of rationality. Only at very high and very low prices *ruck* increases if one allows for the possibility of errors (see Figure 11).³⁹

- Insert Figure 11 here -

These results have implications for the discussion about “noise traders” in market microstructure models. As discussed earlier, in a model where there are only informed traders and where a market maker sets a bid-ask spread, the no-trade theorem would apply. In order to ensure that even in the presence of a zero-profit condition for the market maker market failure is avoided, most market microstructure model introduce noise traders ad hoc.⁴⁰

First, our results suggest that “noise” seems to emerge automatically due to the irrationality of some of the decisions of the traders. This becomes especially obvious in our full information treatment *BHW+AS* where, pooled over all probability combinations, average rationality is still only 73%. This observation together with the fact that even under the error models *ruck* does not exceed this value, suggests that the problem is not so much extracting the relevant

³⁸Note that expected profits in error model (*β*) are path-dependent, and hence they are not necessarily equal for players which encounter the same realized price.

³⁹Logit regressions of *ruck-s* and *ruck-β* on a constant and the price deviation from the a priori expected value of *A* reveal that the price deviation is not significant for *ruck-s* and has a small negative impact on *ruck-β*, which is just significant at the 5% level.

⁴⁰To justify the presence of noise traders it has been argued that (rational) noise traders act due to liquidity or hedging needs (see e.g. Ausubel, 1990) or due to incentives arising from optimal delegation contracts for portfolio managers (see e.g., Dow and Gorton, 1997). Following Black (1986) there is a strand of the literature in which noise traders are seen as traders which trade on noise as if it was information or as agents who just act randomly.

information from the decision of predecessors but rather processing it correctly. Second, our data provide support for the hypothesis that *each* agent decides rationally with a certain probability as opposed to the hypothesis that some of the subjects always decide rationally while others always decide irrationally. To check for this, 240 of our subjects played three times *BHW+AS*, which does not require assumptions with respect to the rationality of predecessors. Table 12 depicts how many of these subjects made zero, one, two, or three rational decisions. A two-sided Kolmogorov–Smirnov test reveals that this distribution is not significantly different (at any reasonable level) from a distribution which would result if each of these subjects had always decided rationally with probability 0.7 (as assumed in error model (*s*)).

Table 12: Frequency of *ruck* when playing *BHW+AS* three times

	0	1	2	3
frequency in actual data	.067	.129	.371	.429
expected if rational with prob. 0.7	.027	.189	.441	.343

Note: This includes only the 240 subjects who played *BHW+AS* three times; pooled over all probability combinations.

4.7 Transaction costs

An important question for financial markets is how transaction costs might influence the efficiency of markets.⁴¹ Treatments with transaction costs were identical to *P+D* except that traders had to pay a cost of 0.1 (in the *P+D+T_{0.1}* and *Pe+D+T_{0.1}* treatments) or a cost of 0.5 (in the *P+D+T_{0.5}* and *Pe+D+T_{0.5}* treatments). If all traders are rational, this can imply that not trading is rational independent of the signal.

Table 13 shows the relationship between transaction costs and *ruck* and as expected, transaction costs make the choice of *N* more frequent, 23% and 28%, respectively, rather than 19% without transaction costs. However, only 3% of the no trade decisions are justified if subjects have to pay a small cost of 0.1. This percentage increases to 18% for large transaction cost (0.5).⁴²

However, this does not explain the difference between average no trading in *P* and *P+T*.

⁴¹In a related model, Lee (1998) shows that herding and informational avalanches can occur on financial markets with transaction costs.

⁴²Cipriani and Guarino (2002) consider a setting where the successful asset pays out 100 and transaction costs equal 9 (comparable to transaction cost of 0.9 in our framework). They report that in this case more than 80% of no trade-decision are consistent with rationality.

Table 13: Transaction costs and *ruck*

	<i>ruck</i>	against own signal	no trade	<i>ruck</i> if N
P	0.55	0.25	0.19	0
$P+T_{0.1}$	0.52	0.25	0.23	0.03
$P+T_{0.5}$	0.50	0.23	0.28	0.18

Note: pooled over all probability combinations.

One may conclude from this that transaction costs have an emotional effect. Indeed, in the transaction cost treatments we find the only significant effect of the *stage* of the task (whether a task was the first, second, or third a subject performed in), which allows for an interesting interpretation. $P+T$ treatments were only played on stages 2 and 3. We find a strong effect that N was chosen more frequently when a subject encountered a $P+T$ treatment on stage 3 for the first time and played P without transaction cost on stage 2.⁴³ The obvious interpretation is that subjects felt compelled to react to the change in the rules by overreacting to transaction costs. Whereas when $P+T$ treatments appeared already on stage 2, no such overreaction could be observed and there is no significant difference between average N for P and $P+T$ on this stage.

4.8 Is price perceived as a martingale?

The theoretical price in the model without transaction cost is a martingale with respect to public information. From our previous discussion we already know that realized prices in our experiment strongly diverged from the theoretical prediction. However, for an efficient market it is not only important to see whether prices efficiently incorporate all information but also whether subjects perceive the price to be a martingale.

If all players are rational, they know that the price tomorrow reflects their decision D_t in this period, i.e. they know that $E(A|H_t, D_t) = p_{t+1}$. Given that a rational trader only buys in the direction of his signal, the price tomorrow is $E(A|H_t, s)$. The law of iterated expectations implies that the best guess of a trader for the price p_{T+1} after the last trader has decided must also be $E(A|H_t, s) = 10P(A|H_t, s)$.⁴⁴ As laid out in Section 3.1, in three of our treatments

⁴³The frequency of N is 0.11 in $P+T_{0.1}$ on stage 2 but 0.26 on stage 3. Likewise, in $P+T_{0.5}$ it is 0.21 on stage 2 and 0.30 on stage 3. Both differences are significant at the 1% level of a MWU test.

⁴⁴Note that if, in line with error model (s) of Section 4.6, subjects assume that all of their successors follow their respective signal with a constant probability, the price remains a martingale as the expected price process in such an error model is equivalent to the expected price process in a Glosten and Milgrom (1985) world with noise traders.

we asked subjects to enter a bet \hat{p}_{T+1} on the endprice which was remunerated according to equation (4). If the price is perceived as a martingale, we should find that the price bet just equals the lead price, i.e., that $\hat{p}_{T+1} = p_{t+1}$.

Figure 12 shows the distribution of the difference between lead price and price bet. About 14% of bets equal exactly the lead price p_{t+1} ; 31% of all price bets are within a 0.3 interval of the lead price (about 10% just named the current price p_t). This gives at best moderate support for the hypothesis that subjects perceive the price to be a martingale. And indeed, it would be strange for those subjects who are contrarians to believe in the martingale property.

- Insert Figure 12 here -

To check whether actual prices behaved as a martingale, we performed a runs test (see e.g. Campbell et al., 1997). Under the null hypothesis (martingale property) both, signals and decisions of subjects should be history independent. For the test we pooled observations over groups from a given probability combination but separately for whether the true state was A or B . The results in Table 14 show that for signals we can never reject the null hypothesis of random choice (as it should be since signals were generated by i.i.d. processes).

Table 14: Runs test: number of runs

true value		probability combination					
		50-66	51-55	55-60	55-80	60-55	60-60
A	signal	286 (364/328)	189 (227/154)	357 (418/299)	39 (86/26)	26 (30/27)	7 (10/3)
	decision	370** (309/283)	210* (200/181)	399** (368/349)	73** (59/53)	36* (30/27)	7 (8/5)
B	signal	225 (173/342)	82 (64/91)	404 (336/508)	22 (16/60)	68 (62/66)	41 (32/47)
	decision	305** (255/260)	88 (82/73)	468** (413/431)	51** (36/40)	80** (67/61)	46 (40/39)

Note: treatments P and $P-N$; in treatment P no-trade decisions are omitted; ($\#A/\#B$) denotes the number of A/B signals or decisions, respectively; ** (*) randomness rejected at 1% (10%) significance level.

However, decisions (and therefore also prices) violate the martingale property for most of our probability combinations. In all cases except in 60-60, where the number of observations is insufficient, and in 51-55 when the true value is B , the runs test rejects the hypothesis of history independent choices, mostly at the 1% level. Furthermore, randomness is always violated in a

particular direction, namely there are more runs than predicted by theory. This again is evidence against herd behavior as herds would yield fewer runs than predicted by theory. Contrarian behavior, on the other hand, is compatible with too many runs.

4.9 The influence of personal characteristics

With the large number of participants in our experiment, it is possible to investigate the behavior of a variety of subgroups. For the price treatments, Table 15 compares average *ruck* of male subjects, female subjects, subjects who are current or former students (“college”), subjects who have never attended college (“no college”), subjects holding a Ph.D. and current Ph.D. students. To test for differences, we ran logit regressions at the individual level to explain the rationality variable *ruck* by dummies for the above subgroups. Controlling for the duration of play, the age of the subject and its position in the group, we find that neither the sex of the subject nor the college dummy are significant at reasonable levels. However, there is evidence that the 45% of our subjects studying for or holding a Ph.D. have a significantly higher *ruck* (with *p*-values of 0.01 and 0.066, respectively).⁴⁵

Table 15: *ruck* of subgroups in price treatments

male	female	college	no college	Ph.D.	Ph.D. stud.	all
0.56	0.56	0.56	0.53	0.58	0.58	0.56

Note: Data from treatment groups *P*, *P-N*, and *P+T_i*; pooled over all probability combinations.

As a large part of our subject pool had some university education, we asked subjects for their major field of study as it should be interesting whether there are differences in behavior. And indeed, there are some differences. First, we look at the full-information treatment *BHW+AS* where one should ignore the decisions of others. In this treatment, the skill of judging the irrationality of others is irrelevant while the ability to deal with probabilities is decisive. In Figure 13 we see that the group of mathematicians perform best followed by engineers, computer scientists and economists. Worst are linguists, biologists, lawyers and most of all pharmacists.⁴⁶

- Insert Figure 13 here -

⁴⁵Controlling for the field of studies reveals that these effects are mainly driven by Ph.D.’s and Ph.D. students from the sciences, medicine and engineering. Surprisingly, there is no significant Ph.D. effect in economics. We also included a dummy to reflect whether a subject had provided an invalid email address (see Footnote 17). This dummy is not significant.

⁴⁶Note that in this treatment expected profit are strictly increasing in *ruck*.

However, in the price treatments *ruck* is not necessarily a virtue if others are irrational. What matters are the profits subjects made through their choices. Figure 14 depicts average *ruck* and profits (excluding fixed payments and payoffs from the price bet) for selected fields of study. As expected from Figure 13, with respect to *ruck* physicists and mathematicians perform above average. Business and (surprisingly) economics students are slightly below average and psychologists performed worst. However, if one looks at profits, the ranking is almost exactly reversed. Physicists and mathematicians do much worse than psychologists and economists. Figure 14 suggests that sometimes an intuition for the partly irrational behavior of others seems to be more important than being able to apply Bayes' rule.⁴⁷

- Insert Figure 14 here -

4.10 Are consultants different?

A common (and justified) critique against experiments in economics is that, with few exceptions, they rely on a subject pool consisting only of economics students. While this can be changed with relatively little effort, the reliance on students in general is often dictated by financial and practical constraints. This makes it all the more important that outside validity is checked when one has the rare opportunity to conduct experiments with professionals in business. We were able to conduct the experiment on the same platform with 267 consultants of an international consultancy.⁴⁸ Do their results differ?

Table 16: *ruck* of general subject pool and consultants

	general subject pool	consultants
old text	.53 (.09)	.54 (.17)
new text	.56 (.14)	.58 (.12)

Note: Treatment *P+D*, probability combination 55-60. Standard deviations in parentheses. Number of groups: general/old text: 8, consultants/old text: 9, general/new text: 18, consultants/new text: 6.

⁴⁷Regressions (logit for *ruck*, OLS for *profits*) on dummies for the above fields of study, Ph.D. status, duration of play, age and sex of the subject, and its position in the group reveal that the field dummies have the right sign but only the profits of economists and psychologists are significantly higher than those of physicists (with *p*-values of 0.08 and 0.05, respectively). Note that profits may vary from -10 to +10 and standard deviations are rather large.

⁴⁸The groups size in this control experiment was always limited to 10 and sometimes 5. However, testing revealed that this has no significant effect.

Originally we used a different text version in the control experiment.⁴⁹ To verify whether consultants behaved differently all four combinations of general subject pool / consultants and old text / new text were played in treatment $P+D$ and probability combination 55-60 (which was the only probability combination played in the control treatment with consultants). Table 16 shows average *ruck* for these combinations. For both text versions the consultants have slightly higher values of *ruck* but none of the differences is significant at the 5% level of a MWU-test (two-sided). Thus, the control group of consultants does not behave significantly different from our general subject pool.

5 Conclusion

In this paper, we present results of a large scale internet experiment based on a sequential asset market with privately informed traders. Avery and Zemsky (1998) predict that in such markets herd behavior should not be observed because all trade decisions are immediately incorporated into the market price which, consequently, reflects all public information. And indeed, as predicted, we do not find evidence for herding or imitative behavior in our experiment. However, in contrast to theory, subjects do not always follow their private information but frequently act as contrarians, i.e., they trade against the market and their own signal. To explain this behavior, we study two error models which allow for the possibility that subjects have doubts about the rationality of others and consequently mistrust their decisions. These error models are able to rationalize contrarian behavior at relatively low or high prices. In fact, a very successful strategy in the price treatments would have been to follow one's signal for moderate prices and to be a contrarian for extreme prices. Both error models produce results which are not significantly different from each other. This is quite surprising given their very different requirements with respect to the depth of reasoning of subjects. Furthermore, in the standard model as well as in the error models, rationality seems to be "bounded" by the level subjects achieved in the full-information treatment $BHW+AS$.

Our experiment complements a large empirical literature on herding. Our results confirm the finding of the bulk of this literature that herd behavior driven by informational externalities does not seem to be an important force in financial markets. To the contrary, one could even

⁴⁹The new text was more precise about the independence of the investment bankers' tips. In the general phase, the two different text versions did not yield significantly different results, and hence the groups with old text were always included in the previous results.

argue that the observed contrarian behavior, which we find sometimes to be profitable, has a stabilizing effect as it implies that agents tend to differentiate their investments from those of their predecessors. Of course, this does not rule out herding in financial markets based on explanations other than purely information based ones, as for example reputation concerns, payoff externalities, etc.. To disentangle these factors is an important task for future empirical and experimental research.

References

- Allsopp, L. and J. D. Hey**, “Two Experiments to Test a Model of Herd Behavior,” *Experimental Economics*, 2000, *3*(2), 121–136.
- Anderhub, V., R. Müller, and C. Schmidt**, “Design and Evaluation of an Economic Experiment Via the Internet,” *Journal of Economic Behaviour and Organization*, 2001, *46*(2), 227–47.
- Anderson, L. R. and C. A. Holt**, “Information Cascades in the Laboratory,” *American Economic Review*, 1997, *87*(5), 847–862.
- Ashiya, Masahiro and Takero Doi**, “Herd Behavior of Japanese Economists,” *Journal of Economic Behavior and Organization*, 2001, *46*, 343–346.
- Ausubel, L.M.**, “Partially-Revealing Rational Expectations Equilibria in a Competitive Economy,” *Journal of Economic Theory*, 1990, *50*, 93–126.
- Avery, Christopher and Peter Zemsky**, “Multidimensional Uncertainty and Herd Behavior in Financial Markets,” *American Economic Review*, 1998, *88*(4), 724–48.
- Banerjee, A. V.**, “A Simple Model of Herd Behavior,” *Quarterly Journal of Economics*, 1992, *107*(3), 797–817.
- Berg, J. E., A. D. Lane, J. W. Dickhaut, and J. R. O’Brien**, “Controlling Preferences for Lotteries on Units of Experimental Exchange,” *Quarterly Journal of Economics*, 1986, *101*, 281–306.
- Bikhchandani, S. and S. Sharma**, “Herd Behavior in Financial Markets: A Review,” *IMF Policy Working Papers WP/00/48*, 2000.
- , **D. Hirshleifer, and I. Welch**, “A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades,” *Journal of Political Economy*, 1992, *100*(5), 992–1026.
- Black, Fischer**, “Noise,” *Journal of Finance*, 1986, *41*(3), 529–543.
- Bosch-Domenech, A., J. Montalvo, R. Nagel, and A. Satorra**, “One, Two, (Three), Infinity, ...: Newspaper and Lab Beauty-Contest Experiments,” *American Economic Review*, 2002, *92*(5), 1687–1701.
- Bounmy, K., J. Vergnaud, M. Willinger, and A. Ziegelmeyer**, “Information Externalities and Learning with Sequential Interactions,” in Cohendet, Llerena, Stahn, and Umbhauer, eds., *The Economics of Networks: Interaction and Behaviours*, Springer-Verlag, 1998, pp. 291–305.

- Brunnermeier, Markus**, *Asset Pricing under Asymmetric Information: Bubbles, Crashes, Technical Analysis, and Herding*, Oxford: Oxford University Press, 2001.
- Camerer, C. and R. Hogarth**, “The Effects of Financial Incentives in Experiments: A Review and Capital-Labor-Production Framework,” *Journal of Risk and Uncertainty*, 1999, 19(1-3), 7–42.
- Campbell, J., A. Lo, and A. MacKinley**, *The Econometrics of Financial Markets*, Princeton: Princeton University Press, 1997.
- Charness, G., E. Haruvy, and D. Sonsino**, “Social Distance and Reciprocity: The Internet Vs. The Laboratory,” *mimeo*, 2001.
- Cipriani, Marco and Antonio Guarino**, “Herd Behavior and Contagion in Financial Markets,” *mimeo*, *New York University*, 2001.
- and —, “Herding and Price Convergence in a Laboratory Financial Market,” *mimeo*, *New York University*, 2002.
- Daniel, K., D. Hirshleifer, and S. Teoh**, “Investor Psychology in Capital Markets: Evidence and Policy Implications,” *Journal of Monetary Economics*, 2002, 49(1).
- Diamond, D. and P. Dybvig**, “Bank Runs, Deposit Insurance and Liquidity,” *Journal of Political Economy*, 1983, 91(3), 401–19.
- Dow, J. and G. Gorton**, “Noise Trading, Delegated Portfolio Management, and Economic Welfare,” *Journal of Political Economy*, 1997, 105(5), 1024–1050.
- Dreman, D.**, *Contrarian Investment Strategy*, New York: Random House, 1979.
- , *Contrarian Investment Strategies: The Next Generation*, New York: Simon and Schuster, 1998.
- Forsythe, R., F. Nelson, G. Neumann, and J. Wright**, “Anatomy of an Experimental Political Stock Market,” *American Economic Review*, 1992, 82, 1142–1161.
- , **T. Rietz, and T. Ross**, “Wishes, Expectations and Actions: A Survey on Price Formation in Election Stock Markets,” *Journal of Economic Behavior and Organization*, 1999, 39, 83–110.
- Glosten, Lawrence and Paul Milgrom**, “Bid, Ask and Transaction Prices in a Specialist Market with Heterogeneously Informed Traders,” *Journal of Financial Economics*, 1985, 14(1), 71–100.
- Goidel, R. K. and T. G. Shields**, “The Vanishing Marginals, the Bandwagon, and the Mass Media,” *The Journal of Politics*, 1994, 56(3), 802–810.
- Graham, J.R.**, “Herding Among Investment Newsletters: Theory and Evidence,” *Journal of Finance*, 1999, 54, 237–268.
- Greiner, B., H. Jacobsen, and C. Schmidt**, “The Virtual Laboratory Infrastructure for Online Economic Experiments,” *mimeo*, 2002.
- Güth, W., C. Schmidt, and M. Sutter**, “Fairness in the Mail and Opportunism on the Internet - a Newspaper Experiment on Ultimatum Bargaining,” *German Economic Review*, *forthcoming*, 2003.

- Hey, J. and A. Morone**, “Do Markets Drive Out Lemmings - or Vice Versa?,” *mimeo*, 2002.
- Hirshleifer, D. and S.H. Teoh**, “Herd Behavior and Cascading in Capital Markets: A Review and Synthesis,” *European Financial Management*, 2003, *9(1)*, 25–66.
- Holt, Charles**, “Webgames and Strategy: Recipes for Interactive Learning,” *mimeo*, 2002.
- Hommes, C., J. Sonnemans, J. Tuinstra, and H. Van de Velden**, “Expectations and Bubbles in Asset Pricing Experiments,” *CeNDEF Working Paper Series WP-02-05*, 2002.
- Huck, S. and J. Oechssler**, “Informational Cascades in the Laboratory: Do They Occur for the Right Reasons?,” *Journal of Economic Psychology*, 2000, *21*, 661–671.
- Hung, A. A. and C. R. Plott**, “Information Cascades: Replication and an Extension to Majority Rule and Conformity Rewarding Institutions,” *American Economic Review*, 2001, *91(5)*, 1508–20.
- Krämer, C., M. Nöth, and M. Weber**, “Information Aggregation with Costly Information and Random Ordering: Experimental Evidence,” *mimeo*, 2000.
- Kübler, D. and G. Weizäcker**, “Limited Depth of Reasoning and Failure of Cascade Formation in the Laboratory,” *Review of Economic Studies*, *forthcoming*, 2003.
- Lakonishok, J., A. Shleifer, and R. Vishny**, “The Impact of Institutional Trading on Stock Prices,” *Journal of Financial Economics*, 1992, *32*, 23–43.
- , — , and — , “Contrarian Investment, Extrapolations, and Risk,” *Journal of Finance*, 1994, *49*, 1541–1578.
- Lee, In Ho**, “Market Crashes and Informational Avalanches,” *Review of Economic Studies*, 1998, *65*, 741–59.
- Lucking-Reiley, D.**, “Using Field Experiments to Test Equivalence Between Auction Formats: Magic on the Internet,” *American Economic Review*, 1999, *89(5)*, 1063–1080.
- McKelvey, R.D. and T.R. Palfrey**, “Quantal Response Equilibria for Normal-Form Games,” *Games and Economic Behavior*, 1995, *10*, 6–38.
- and — , “Quantal Response Equilibria for Extensive Form Games,” *Experimental Economics*, 1998, *1*, 9–41.
- Nöth, M. and M. Weber**, “Information Aggregation with Random Ordering: Cascades and Overconfidence,” *Economic Journal*, *forthcoming*, 2002.
- , **C. F. Camerer, C. R. Plott, and M. Weber**, “Information Aggregation in Experimental Asset Markets: Traps and Misaligned Beliefs,” *mimeo*, 2001.
- Oberhammer, C. and A. Stiehler**, “Does Cascade Behavior in Information Cascades Reflect Bayesian Updating? An Experimental Study,” *mimeo*, 2001.
- Plott, C., J. Wit, and W. Yang**, “Parimutuel Betting Markets as Information Aggregation Devices: Experimental Results,” *Economic Theory*, 2003, *22*, 311–351.
- Porta, R. La, J. Lakonishok, A. Shleifer, and R. Vishny**, “Good News for Value Stocks: Further Evidence on Market Efficiency,” *Journal of Finance*, 1997, *52(2)*, 859–874.
- Roth, A. and M. Malouf**, “Game-Theoretic Models and the Role of Information in Bargaining,” *Psychological Review*, 1979, *86*, 574–594.

- Sadiraj, K., A. Schram, and S. Van Wijnbergen**, “Rational Behavior Vs. Herd Behavior: Theoretical and Experimental Analysis,” *mimeo*, 1999.
- Scharfstein, David and Jeremy Stein**, “Herd Behavior and Investment,” *American Economic Review*, 1990, *80(3)*, 465–79.
- Shavit, T., D. Sonsino, and U. Benzion**, “A Comparative Study of Lotteries Evaluation in Class and on the Web,” *Journal of Economic Psychology*, 2001, *22*, 483–491.
- Sunder, Shyam**, “Experimental Asset Markets: A Survey,” in John Kagel and Alvin Roth, eds., *Handbook of Experimental Economics*, Princeton: Princeton University Press, 1995, pp. 445–500.
- Welch, I.**, “Sequential Sales, Learning, and Cascades,” *Journal of Finance*, 1992, *47(2)*, 695–732.
- , “Herding Among Security Analysts,” *Journal of Financial Economics*, 2000, *58*, 369–396.
- Willinger, M. and A. Ziegelmeyer**, “Are More Informed Agents Able to Shatter Information Cascades in the Lab?,” in Cohendet, Llerena, Stahn, and Umbhauer, eds., *The Economics of Networks: Interaction and Behaviours*, Springer, 1998, pp. 291–305.

FIGURES




Figure 1: Composition of the subject pool (“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)




Figure 2: Fraction of final prices within ± 1 and ± 0.5 of the final full information price (pooled over all probability combinations)

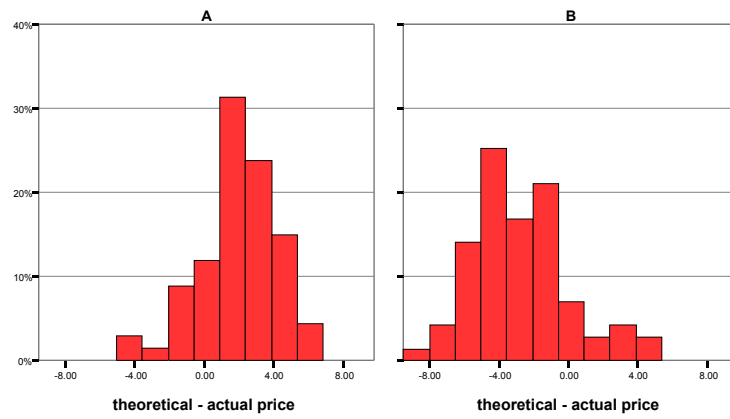


Figure 3: Distribution of the deviation of the final price of asset A from the final theoretical price of asset A for states where asset A is successful (left panel) and where B is successful (right panel) (treatment P , pooled over all probability combinations)

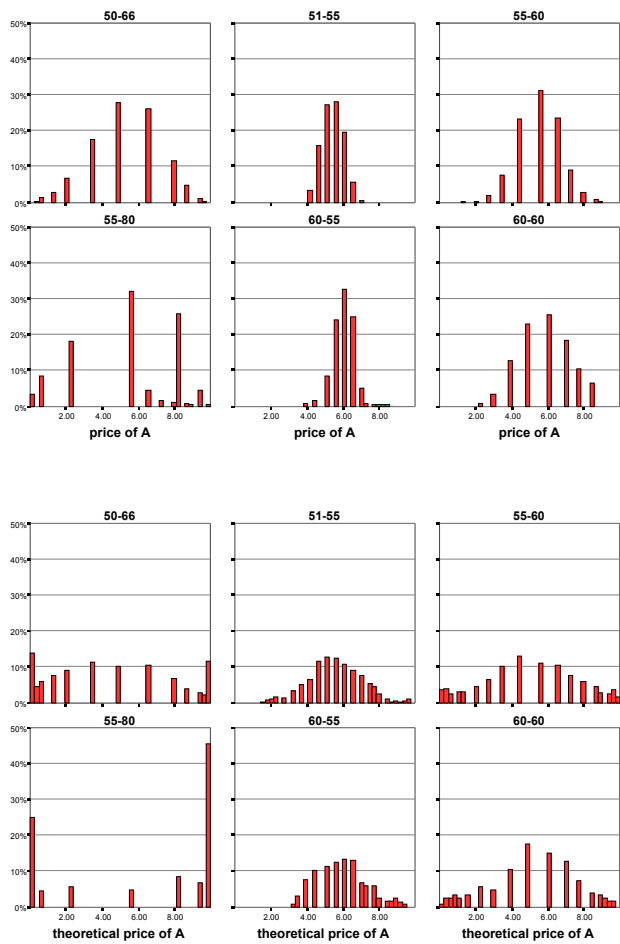


Figure 4: Distribution of actual prices (upper panel) and theoretical prices (lower panel) over all periods. (treatment P)




Figure 5: Average theoretical minus actual price given a theoretical price (left panel) and average theoretical minus actual signal imbalance Δ given a theoretical signal imbalance (right panel) (treatment P , pooled over all probability combinations)




Figure 6: Average absolute price deviations (left panel) and the fraction of actual prices within ± 1 of the theoretical price (right panel) (treatment P , pooled over all probabilities)




Figure 7: Average *ruck* (upper panels) and average buying against the own signal (right panels) given an a signal (left panels) or a b signal (right panels) across actual price of A (treatment P , probability combination 55-60, more extreme prices than those shown are omitted due to a low number of observations.)




Figure 8: Example of an actual price path (treatment $P-N$, probability combination 51-55)




Figure 9: Profits of contrarians across prices in treatments $P-N$ (left panel) and P (right panel) (probability combination 55-60, counterfactual profits are (-1) times the displayed profits).




Figure 10: Average expected net profits of buying asset A given an a signal (left panel) or a b signal (right panel) for the standard model, error model (s) and error model (β) (treatment $P-N$, probability combination 55-60)




Figure 11: Average *ruck* in the standard model (*ruck*), error model (*s*) (*ruck-s*) and error model (β) (*ruck- β*) across prices in treatments *P-N* (left panel) and *P* (right panel) (probability combination 55-60)




Figure 12: Distribution of the difference between lead price and price bet (treatments *P+D*, *Pe+D* and *P+D-N*, pooled over all probability combinations)




Figure 13: *ruck* in treatment *BHW+AS* (pooled over all probability combinations, only fields with more than 10 participants are included, overall average of *ruck* is 0.72)




Figure 14: Ruck (left panel) and profits (right panel) for selected fields of study (treatments P , $P-N$ and $P+T$, all probability combinations, profits exclude fixed payments and payoffs from the price bet)

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 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 (screen 3 below). Subsequently, the actual experiment began and Screen 4 below provides an example of the first of three rounds (treatment *BHW*). Screen 5 below provides an example of a price treatment played in the second round (treatment *P+D*), and screen 6 displays the text for the price bet. Round 3 had the same basic structure, and therefore we omit an example of this round.

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 11,000 Euros which are distributed among 11 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: Round 1

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

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

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

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

Was the decision difficult? Independent of your decision, what do you think is the probability of A being the successful asset? [] %.

After the third round you'll find out whether your decision was correct. Let's move on the next round.

Screen 5: Round 2

In this round you receive an endowment of 11 Lotto-Euro. The basic structure remains the same as in round 1. (In case you want to review the central features of round 1 please click [\[here\]](#).) This time you have to decide in which of the two risky shares (A or B) you want to invest. Only one share will yield a profit of 10 *Lotto-Euro*, the other one will be worthless. *Share A* is successful in 55% of all cases, *share B* in 45% of all cases. As in round 1 the successful share was determined by chance before the first player of this group played.

In contrast to round 1, you - as every player in the group - have to pay the current share price if you decide to invest in a share. Share prices are determined by supply and demand such that outside investors, who can observe the history of trades but not the tips given by the investment bankers, have no incentive to trade, i.e. an outside investor could not expect to profit from buying or selling one of the shares [only in treatments *e*: ... because the price of share A (B) is equal to the conditional expected value of A (B) given the decisions of all your predecessors]. The role of outside investors is played by the computer.

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

The current price of share A is 6.47 LE. The current price of share B is 3.53 LE. The profit in this round is given by:

your endowment (11 LE)
 - price of the respective share
 + stock profit (10 or 0 LE)

You can also decide not to invest. In this case, you just keep your endowment of 11 LE.

Like your predecessors, you can *observe* the *price history* and the *history of decisions* in your group. You are the [2nd] investor in this group. Your predecessors in this group were facing the same problems as you, and one after another they have purchased the shares shown below at the price valid at that point in time:

Investor no.	1	
Decision	A	
Price of A	5.50	
Price of B	4.50	What do you choose? [A], [B] or [No trade].

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.

Screen 6: Round 2 (price bet)

Now that you have made your decision, you have the possibility to earn an extra 5 LE by correctly predicting the price of A which is in place after the last player in this group has made his decision (but before the true value of the shares is revealed). If you not perfectly predict this price the extra 5 LE will be reduced by 1 LE for each LE your prediction deviates from the final price.

In the worst case you will not receive additional LE. There are still [18] players deciding after you in this group.

As explained above, the price of share A is determined by supply and demand such that outside investors, who can observe the history of decisions but not the tips, have no incentive to trade. To refresh your memory, you find the history of prices and decisions below:

Investor no.	1	
Decision	A	
Price of A	5.50	
Price of B	4.50	Final price of share A: []