Appendix to

Herding and Contrarian Behavior in Financial Markets - An Internet Experiment

By Mathias Drehmann, Jörg Oechssler, and Andreas Roider

(available at: http://www.aeaweb.org/aer/contents/)

This Appendix contains experimental instructions (Appendix A), descriptions of the various pricing rules used in the experiment (Appendix B), average values of *ruck*, *own*, and *own** for all employed probability combinations (Appendix C), and details of the estimation of the error model (Appendix D).

Appendix A: Instructions

Once connected to our website *www.a-oder-b.de*, there was first a general overview on the experiment (Screen 1 below). Then, subjects where asked to provide some personal information (Screen 2 below). Only if all information was provided, subjects were allowed to continue and learn their player number as well as the monetary incentives in the current phase of the experiment (Screen 3 below). Note that the number of lottery tickets and the prizes mentioned below relate to Phase I of the experiment. Subsequently, the actual experiment began and Screen 4 below provides an example of the first of three stages (treatment *BHW*). Screen 5 below provides an example of a price treatment played in the second stage (treatment *P+D*). Stage 3 had the same basic structure, and therefore we omit an example of this stage.

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

At the end of this Appendix A, we additionally provide translations of the descriptions of the pricing rules in treatments *P-N+AS*, *Pt*, *Pf+D-N*, and $P\beta$*+D-N*.

Screen 1: Introduction

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

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

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

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

Screen 2: Request of personal information

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

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

Screen 3: Player number and incentives

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

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

Screen 4: Stage 1

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

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

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

Investor no.	1	2	3
Decision	B	A	B

What do you choose? [A] or [B].

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

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

Screen 5: Stage 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.

Description of the Respective Pricing Rules in Other Treatments

Recall that in the treatments to be described below (i) subjects did not have the option not to trade (and hence the fifth paragraph of Screen 5 above was omitted in these treatments), and (ii) these treatments were played in all three rounds.

Treatments *P-N+AS*

In the instructions for these treatments, subjects were notified in the sixth paragraph of Screen 5 above that also the history of tips was observable. In addition, the following replaced the second paragraph of Screen 5 above:

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 all the tips that your predecessors have received from their respective investment bankers (but cannot observe the tip you have received), have no incentive to trade, i.e., an outside investor could not expect to profit from buying or selling one of the shares. The role of outside investors is played by the computer.

Treatments *Pt*

In the *Pt* treatments the following replaced the second paragraph of screen 5 above:

You - as every player in the group - have to pay the current share price if you decide to invest in a share. To be precise, the share price depends on the difference (#A-#B) between the number of predecessors who bought A and the number of predecessors who bought B. The following table lists the share price for all possible differences (#A-#B):

$\#A-\#B$	0	1	2	3	4	5	6	7	8	9
share price A	5.50	6.47	7.33	8.05	8.61	9.03	9.33	9.54	9.69	9.79
share price B	4.50	3.53	2.67	1.95	1.39	0.97	0.67	0.46	0.31	0.21
$\#A-\#B$	0	-1	-2	-3	-4	-5	-6	-7	-8	-9
share price A	5.50	4.49	3.52	2.66	1.95	1.39	0.97	0.67	0.45	0.31
share price B	4.50	5.51	6.48	7.34	8.05	8.61	9.03	9.33	9.55	9.69

$\#A-\#B$	10	11	12	13	14	15	16	17	18	19
share price A	9.86	9.90	9.94	9.96	9.97	9.98	9.99	9.992	9.994	9.996
share price B	0.14	0.10	0.06	0.04	0.03	0.02	0.01	0.008	0.006	0.004
$\#A-\#B$	-10	-11	-12	-13	-14	-15	-16	-17	-18	-19
share price A	0.21	0.14	0.09	0.06	0.04	0.03	0.02	0.012	0.008	0.006
share price B	9.79	9.86	9.91	9.94	9.96	9.97	9.98	9.988	9.992	9.994

Three examples shall illustrate the above table. Suppose you were the fifth participant in a group.

If three of your predecessors chose share A and one share B, $\#A-\#B$ would be 2, in which case the price of share A would be 7.33, and the price of share B would be 2.67.

However, had three of your predecessors chosen share B and only had chosen share A, $\#A-\#B$ would be -2. The share price of A would be 3.52 and the price of share B would be 6.48.

In case two of your predecessors had chosen A and two had chosen B, $\#A-\#B$ would be 0. The share price of A would be 5.50 and that of share B would be 4.50. These are also the prices faced by the first participant of each group.

Treatment *Pf+D-N*

The pricing rule in this treatment was described in exactly the same way as in treatments *Pt* above with the only difference that the share prices in the table and in the examples were adapted to take account of the assumption that subjects followed their signal only with a certain probability (see Appendix B below).

Treatment *Pβ+D-N*

In treatment *Pβ+D-N* the following replaced the second paragraph of screen 5 above:

You - as every player in the group - have to pay the current share price if you decide to invest in a share. The share price is set by the computer. The computer sets the price in a way of a seller who maximizes his profit. It is assumed that the following holds:

1. The seller knows the above described basic structure of the situation.
2. When the seller sets the prices at which he is willing to trade, he can observe only the decisions of the predecessors in the respective group. In particular, the seller does not know any of the tips of the investment bankers.
3. However, the seller has ample experience. Based on the behavior of participants in similar earlier experiments he is trying to estimate which influence the history of prices, the earlier decisions of participants and the tips of investment bankers have on the decisions of participants.
4. The seller is facing competition by many other sellers who possess the same properties as himself, and who also want to trade in shares A and B.

Appendix B: Pricing Rules

In this appendix we explain in more detail how prices were set in each of our price treatments. The computer served as market maker. Prices were set based on the observable history of decisions – not based on the history of signals (except in treatment group *P-N+AS* where prices were based on the history of signals since they *were* observable in this treatment group). At the end of this appendix we provide examples to illustrate the various pricing rules. In all treatments the price of A encountered by the first player was $10 \cdot P(A)$.

1. *P* and *P-N*: In all treatments in these treatment groups it was assumed that all observed decisions were formed according to *ruck*, i.e., that all subjects followed their own signal. The only possible observable deviation from this is a no–trade decision in treatment *P* (in violation of *ruck)*. In this case the out-of-equilibrium belief was that a deviation could have been committed by subjects with either signal. Hence, no signal could be imputed, and the price remained unchanged. Given this assumption, it is straightforward to show (by a simple application of Bayes' rule) that the price p_t only depends on the net number n_t^D of A decisions up to date t, i.e., the number of decisions to buy A minus the number of decisions to buy B observed by the market maker:

$$p_t \equiv 10P(A|H_t) = 10P(A|n_t^D) = \frac{P(A)}{P(A) + (1 - P(A)) \cdot P(a|A)^{-n_t^D} \cdot (1 - P(a|A))^{n_t^D}}. \quad \text{(B1)}$$

2. *P-N+AS* and full information price p_t^*: In both treatments in treatment group *P-N+AS* it was assumed that the market maker can observe the history of signals. In analogy to the reasoning above, the price in this case only depends on the net number n_t^{AS} of a signals up to date t, i.e., the number of a signals minus the number of b signals observed by the market maker:

$$p_t \equiv 10P(A|n_t^{AS}) = \frac{P(A)}{P(A) + (1 - P(A)) \cdot P(a|A)^{-n_t^{AS}} \cdot (1 - P(a|A))^{n_t^{AS}}}. \quad \text{(B2)}$$

   Our theoretical benchmark (the full information price p_t^*) is calculated in exactly the same way, and hence does not depend on the specific treatment under consideration.

3. *Pf+D-N*: In contrast to treatment *P-N* in both error treatments the market maker and subjects take into account that subjects may make mistakes. In *Pf+D-N* the price was set as

in treatment *P-N* with the only difference that it was assumed that subjects followed their signal only with probability 0.6542 (which is the empirical frequency of *own* in treatment group *P-N+AS*). Given this assumption it is straightforward to show that:

$$p_t \equiv 10P(A|H_t) = 10P(A|n_t^D) = \frac{P(A)}{P(A) + (1 - P(A)) \cdot \widehat{P}(a|A)^{-n_t^D} \cdot (1 - \widehat{P}(a|A))^{n_t^D}}, \quad \text{(B3)}$$

where $\widehat{P}(a|A) \equiv 0.6542 \cdot P(a|A) + (1 - 0.6542) \cdot (1 - P(a|A))$.

4. *Pβ+D-N*: Again, in this treatment the market maker and subjects take into account that subjects do not always follow their signals. However, in contrast to *Pf+D-N*, in the present treatment the market maker does not assume that the probability of following one's signal is independent of the history of decisions. Rather, based on the concept of quantal response equilibrium of McKelvey and Palfrey (1995, 1998), the market maker assumes that, for a given signal, the probability with which a subject chooses to buy asset A is given by a logit function that depends on the difference in payoffs between the two alternatives (A or B):

$$P(D_i = A|H_t, s_i) = \frac{1}{1 + e^{-\beta(\pi_i^A - \pi_i^B)}}, \quad \text{(B4)}$$

where it is assumed that, due to his experience in market making, the market maker knows the error rate β (for as to how β was derived see below). Hence, if the payoff difference is large, the correct action is chosen with high probability. If it is small, mistakes are more likely.

Pricing rule: Given (B4), the market maker sets prices in the following way. For the first player the market maker states a price equal to $10 \cdot P(A)$. Since for the first player the payoff difference from choosing A rather than B only depends on $P(A)$ and the realization of the signal, the calculation of the price that the second player faces is straightforward: the market maker uses (B4) to calculate $P(s_1 = a|H_1)$, then calculates $P(A|H_1)$ (see (D2) below), and sets the price equal to $p_2 = 10P(A|H_1)$.

$P(A|H_1)$ is then taken as the prior for the second player. Using this prior and the second player's expected payoff difference $\pi^A - \pi^B$ as a function of the second player's signal, the market maker calculates a new $P(A|H_2)$ and a new price $p_3 = 10P(A|H_2)$. $P(A|H_2)$ is then the prior for the third player, and so on.

The error rate β was estimated from the data of treatments *P-N+AS* where subjects did not need to worry about the behavior of their predecessors. As in these treatments the expected payoff differences $\pi_i^A - \pi_i^B$ only depended on the (observable) signal history, for all players, we calculated the expected payoff differences from the signal histories and estimated (B4) with a logit regression. We obtained $\beta = 0.5552$ (standard deviation = 0.09), and based the price calculations in treatment *Pβ+D-N* on this value. An alternative to estimating a single β would have been to estimate different β's for each position in the group. However, we feel that the assumption that the market maker knows the (single) error rate β is already rather demanding.

Examples

To illustrate the pricing rules Tables 1 and 2 contain examples how prices were calculated for a given history of up to 5 decisions and signals in our main probability combination 55-60.

Table 1: Price Formation: An Example Without the No-Trade Option

Investor #	1	2	3	4	5	6	
Signal	*a*	*a*	*b*	*a*	*b*	*a*	
Decision	*A*	*B*	*B*	*A*	*A*	*B*	
P-N	5.50	6.47	5.50	4.49	5.50	6.47	5.50
P-N+AS	5.50	6.47	7.33	6.47	7.33	6.47	7.33
Pf+D-N	5.50	5.80	5.50	5.19	5.50	5.80	5.50
Pβ+D-N	5.50	5.99	5.52	5.03	5.53	6.03	5.56
Full information price p_t^*	5.50	6.47	7.33	6.47	7.33	6.47	7.33

Note: Probability combination 55-60.

Table 2: Price Formation: An Example With the No-Trade Option

Investor #	1	2	3	4	5	6	
Signal	*a*	*a*	*b*	*a*	*b*	*a*	
Decision	*A*	*N*	*B*	*A*	*A*	*N*	
P	5.50	6.47	6.47	5.50	6.47	7.33	7.33
Full information price p_t^*	5.50	6.47	7.33	6.47	7.33	6.47	7.33

Note: Probability combination 55-60.

Appendix C: Average *ruck*, *own,* and *own**

The following table displays average values of *ruck* and *own* for all employed probability combinations. For treatments *BHW* and *BHW+AS*, as a benchmark, the table additionally displays the average values of *own**, which is the fraction of subjects that would have followed their own signal if all subjects had behaved rationally.

Table 3: Average *ruck*, *own*, and *own**

prob.	P	P-N	P-N+AS	Pf+D-N	Pβ+D-N	BHW			BHW+AS		
comb.	ruck	ruck	ruck	ruck	ruck	ruck	own	own*	ruck	own	own*
55-60	.54	.65	.65	.68	.70	.66	.75	.59	.72	.74	.68
	(.12)	(.11)	(.15)	(.10)	(.07)	(.11)	(.13)	(.12)	(.14)	(.14)	(.13)
50-66	.59	.68	_	_	_	.78	.75	.62	.76	.69	.69
	(.11)	(.09)	_	_	_	(.09)	(.12)	(.16)	(.12)	(.12)	(.18)
51-55	.56	.71	_	_	_	.69	.73	.56	.72	.77	.64
	(.13)	(.11)	_	_	_	(.08)	(.12)	(.13)	(.13)	(.15)	(.14)
55-80	.53	_	_	_	_	.78	.81	.70	.83	.73	.79
	(.11)	_	_	_	_	(.09)	(.07)	(.14)	(.24)	(.09)	(.02)
60-51	_	_	_	_	_	.84	.57	.50	.77	.62	.49
	_	_	_	_	_	(.10)	(.13)	(.12)	(.14)	(.16)	(.15)
60-55	.51	_	_	_	_	.78	.62	.53	_	_	_
	(.13)	_	_	_	_	(.09)	(.13)	(.10)	_	_	_
60-60	.60	_	_	_	_	.69	.69	.59	.73	.63	.68
	(.10)	_	_	_	_	(.11)	(.11)	(.14)	(.16)	(.17)	(.23)

Notes: Standard deviations are given in parentheses. The variable *own** denotes the fractions of subjects that would have followed their own signal if all subjects had behaved rationally. Note that in all treatments except *BHW* and *BHW+AS* we have *ruck=own* and *own**=1 by definition.

Appendix D: Details of the Estimation of the Error Model

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 and B. For reasons of tractability, we only estimate the 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

$$P(D_i = A|H_t, s_i) = \frac{1}{1 + e^{-\beta_t(\pi_i^A - \pi_i^B)}}, \quad \text{(D1)}$$

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 $\beta_t \to 0$ but play rational best replies if $\beta_t \to \infty$. Since the expected profits of the first player 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

$$P(D_1 = A|A) = P(D_1|a) \cdot P(a|A) + P(D_1|b) \cdot P(b|A). \quad \text{(D2)}$$

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

$$P(A|D_1 = A) = \frac{P(D_1 = A|A) \cdot P(A)}{P(D_1 = A|A) \cdot P(A) + P(D_1 = A|B) \cdot P(B)}. \quad \text{(D3)}$$

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.

Below we provide our estimates of the β_t's. As it is interesting to compare our results to those of Anderson and Holt (1997), note that the level of the estimated coefficients depends on the payoff resulting from a correct decision. To illustrate this, assume that p denotes the probability of asset A being successful. In Anderson and Holt (1997) a correct decision yielded a payoff of 2, and hence the expected payoff difference was given by $\pi^A - \pi^B = 2(p - (1 - p))$. In our experiment, besides the payoff resulting from a correct decision (which was equal to 10), the expected payoff difference depended also on the current market price of A. Hence, the expected payoff difference was given by $\pi^A - \pi^B = 2[p(10 - p_A) - (1 - p)p_A]$. Given that the actual price in treatment P-N was on average equal to 5.44, our estimates have to be multiplied by (approximately) 5 to make them comparable.

Table 4: Coefficients of the Error Model

	Coefficient	Standard deviation
β_1	.58**	.12
β_2	.08	.09
β_3	.52**	.13
β_4	.31**	.10
β_5	.30**	.09
β_6	.27**	.09
β_7	.55**	.13
β_8	.41**	.11
β_9	.23*	.09
β_{10}	.37**	.11
β_{11}	.32**	.10
β_{12}	.24*	.10
β_{13}	.35**	.10
β_{14}	.31**	.11
β_{15}	.43**	.10
β_{16}	.19*	.09
β_{17}	.27**	.09
β_{18}	.39**	.10
β_{19}	.37**	.10
β_{20}	.35**	.10

** indicates significance at the 1-percent level.
* indicates significance at the 5-percent level.