Behavioral Biases in Endogenous-Timing Herding Games: an Experimental Study*

Asen Ivanov[†,‡] Dan Levin[§] James Peck[¶]

November 30, 2012

Abstract

We experimentally study behavior in an endogenous-timing herding game. We find that subjects respond to their type and to observed investment activity in a sensible way, but there are also substantial departures from Nash Equilibrium. Some departures can be viewed as mere noise in decision making while other departures represent systematic biases reflecting subjects' failure to appreciate subtle aspects of the game.

Keywords: endogenous-timing herding game, biases, experiment

*We thank the editor, associate editor, and two referees, whose suggestions improved the paper. We are grateful to seminar participants at the University of the Basque Country and Maastricht University. This work is supported by the NSF under grants SES-0417352 and SES-1031101 (for Dan Levin and James Peck) and under grant SES-1030467 (for Asen Ivanov). An earlier version of the paper was circulated under the title, "Animal Spirits and Informational Externalities in an Endogenous-Timing Investment Game: An Experimental Study." Any opinions, findings and conclusions or recommendations expressed are those of the authors and do not necessarily reflect the views of the NSF.

†Department of Economics, Virginia Commonwealth University, 301 W. Main Street, Richmond, VA 23220, USA.

‡Corresponding author. Present address: School of Economics and Finance at Queen Mary, University of London; 327 Mile End Road, London E1 4NS, United Kingdom; tel. +44 (0)20 7882 5886; a.ivanov@qmul.ac.uk

§Department of Economics, The Ohio State University, 1945 N. High Street, Columbus, OH 43210, USA; levin.36@osu.edu

¶Department of Economics, The Ohio State University, 1945 N. High Street, Columbus, OH 43210, USA; peck.33@osu.edu

1 Introduction

We consider an endogenous-timing herding model with the following features: (i) There is an uncertain, exogenously determined, common gross return for investment. (ii) Each player's type consists of a pair of private signals, one about the investment return and the other about investment cost. (iii) There are multiple rounds; in round 1, players simultaneously decide whether to invest, in round 2 those who did not invest in round 1 observe how many other players invested in round 1, and decide whether to invest in round 2, and so on. Investment is irreversible, and can be done at most once. (iv) There is a cost to waiting (due to discounting).

In this environment, investment by others does not affect one's payoff directly so that there are no payoff externalities. However, there can be informational externalities, as the history of investment by other players may reveal something about their private information, and hence about the investment return, to players who wait. In such settings, important issues arise regarding how well the market aggregates information and yields efficient outcomes.

Chamley and Gale (1994) address these issues theoretically for the case where all players have the same investment cost so that types are one-dimensional. Levin and Peck (2008) extend the analysis to two-dimensional types with heterogenous investment costs. Both studies find that the Nash equilibrium (NE) outcome is often inefficient.

To understand the relevant forces, consider the optimal behavior of some player i given her type and given others' strategies. In each round prior to investing, i needs to compare the expected payoff from investing with the option value of waiting. To compute the expected payoff from investing, i needs to take into account her type and, after round 1, to draw correct inferences about the investment return based on the observed investment history. To compute the option value of waiting, i needs to look ahead to the possible realizations of others' investment activity and the inferences that can be drawn from each realization. Clearly, if i's expected payoff of investing is negative she waits. However, even if it is positive but less than the option value of waiting, she still waits in order to free-ride on the information provided by others' investment activity. Such free-riding leads to inefficiency in NE.

The forces shaping i's optimal behavior become especially apparent when one considers how i's computations depend on the structure of the game. For example, i's computations are sensitive in a subtle but crucial way to two features, the set of possible types and the size of the market. The set of types affects the informational content of any given observed investment history. It also determines whether i might be able to free-ride on other subjects, e.g. on subjects with stronger types in the sense of higher expected payoff from investing. Increasing the market size enhances the option value of waiting in round 1 by virtue of sample-size effects–if i waits, she gets to observe a larger sample coming from the same distribution conditional on the investment return.[1] As a result, NE behavior is critically sensitive to such features of the game.

The main goal of our paper is to experimentally study systematic departures from NE.[2] To this end, we implement a series of endogenous-timing investment games in the lab. The games differ by the number of players in the market (two or ten), and by the cost-structure (one-cost or two-cost). In the one-cost games all players have identical costs–high in the high-cost games and low in the low-cost games. In the two-cost games, the cost of investment is either high or low and can differ across

[1]After round 1, the sets of possible histories are different on markets of different sizes so that comparisons are more difficult to make.

[2]In the appendix, we also explore the implications of these departures for informational externalities and market efficiency.

subjects.

We find that subjects respond to their type and to observed investment activity in a sensible way. Thus, behavior satisfies some basic requirements of rationality. Nevertheless, subjects' behavior departs from NE in important ways. One departure from NE occurs because subjects sometimes invest despite a negative expected payoff from doing so. Another departure occurs because subjects do not always invest when investing dominates waiting.

Although these departures have an important impact on informational externalities and efficiency, they might be viewed as nothing more than noise that inevitably moves the empirical frequencies of play away from NE boundary values of 0 or 1. However, subjects also exhibit four systematic behavioral biases. First, they exhibit insensitivity to market size by failing to appreciate the relevant sample-size effects. Second, subjects are insensitive to opportunities to free-ride on other subjects. Third, subjects fail to appreciate differences in the informational content of observed market activity between the one-cost and the two-cost games. Fourth, subjects are excessively conservative, in the sense that they are more reluctant to invest in response to market activity, than a player in a NE or a player best responding to the empirical frequencies.

The first generation of herding models assumes exogenous timing, i.e., that agents are exogenously placed in a queue and must sequentially decide whether or not to invest (see Banerjee (1992) or Bikhchandani, Hirshleifer, and Welch (1992)). The seminal result in these models is the possibility of information cascades, which lead to inefficiencies as a result of a failure to aggregate private information. Anderson and Holt (1997) provide the first experimental tests of herding models with exogenous timing and demonstrate the frequent occurrence of information cascades.[3]

The first experimental work on endogenous timing herding models is done by Sgroi

[3] See Weizsäcker (2010) for a meta-analysis of exogenous-timing herding experiments.

(2003) and Ziegelmeyer et al. (2005). In Sgroi (2003), subjects receive two draws from one of two urns: a "red" urn which contains two red (R) balls and one white (W) ball, and a "white" urn which contains two W balls and one R ball. The possible signals are either strongly red (RR), strongly white (WW), or neutral (RW or WR). In each round, subjects either guess which urn the draws came from, or wait (at a cost) to observe others' guesses. In Ziegelmeyer et al (2005), each of two subjects receives a signal that is randomly drawn from the set of integers between -4 and 4, and the subjects have to guess whether the sum of the two signals is positive or negative. Although the NE predicts perfect identification (where signals 4 and -4 guess in round 1, signals 3 and -3 guess in round 2, etc.), a more common strategy is for signals of absolute value 3 or 4 to guess in round 1. Our study goes beyond Sgroi (2003) and Ziegelmeyer et al (2005), by identifying behavioral biases that explain departures from NE. The previous papers did not vary the market size or the set of possible types. We show that subjects do not respond to these features of the environment, in contrast to the theoretical predictions.

The experimental work closest to the current paper is Ivanov, Levin, and Peck (2009, henceforth ILP), which studies behavior in the two-player one-cost and two-cost games mentioned above. ILP focuses on individual-level behavioral issues. The authors consider prominent behavioral theories–the level-k model (see Nagel (1995), Stahl and Wilson (1994, 1995), and Crawford and Iriberri (2007)), cursed equilibrium (Eyster and Rabin (2005)), and Quantal Response Equilibrium (McKelvey and Palfrey (1995, 1998))–as possible explanations of behavior. They conclude that rather than best-responding (possibly with noise) to (possibly incorrect) beliefs as in the aforementioned behavioral theories, subjects seem to be following simple, computation-free rules of thumb based on different insights about the game.

The rules of thumb identified in ILP are interesting from the point of view of individual-level cognition in dynamic settings where people gather information from, or anticipate gathering information from, others' behavior over time. However, they do not map readily into behavior for markets with more than two players. The current paper takes a different angle on behavior by identifying aggregate-level behavioral patterns that apply more generally to endogenous-timing herding games. The two approaches are complementary. Moreover, the finding in the current paper that aggregate-level behavior does not respond to more subtle features of the game is consistent with the finding in ILP that, at the individual level, subjects use simple, computation-free rules of thumb.[4]

The layout of the paper is as follows. Section 3 defines the games and presents the Nash equilibria. The experimental design is explained in Section 4. Our experimental results are presented in Section 5. Section 6 offers concluding remarks. The appendix considers some additional important issues.

2 Theoretical Framework

Our theoretical framework is based on the model in Levin and Peck (2008). There are n risk-neutral players.[5] Let $Z \in \{0, 10\}$ denote the true gross investment return, common to all players, with $Pr(Z = 0) = Pr(Z = 10) = \frac{1}{2}$. Each player i observes a signal correlated with the investment return, $X_i \in \{0, 1\}$, which we call the common-value signal of player i. We assume that signals are independent, conditional on Z.

[4]Another strand of the literature, pioneered by Avery and Zemsky (1998), retains the exogenous-timing assumption but allows the investment cost (or asset price) to change as new information is revealed. See Cipriani, Guarino (2005) and Drehmann et al. (2005), and Park and Sgroi (2012) for experimental studies.

[5]Our derivations of the Nash equilibria below, as well as our analysis of behavior in the lab, will be conducted under the assumption of risk-neutrality. We discuss violations of this assumption in section 4.3.

The accuracy of the signal is given by the parameter $\alpha \in [\frac{1}{2}, 1]$:

$$Pr(Z = 0 \mid X_i = 0) = Pr(Z = 10 \mid X_i = 1) = \alpha.$$

When $\alpha = \frac{1}{2}$, common-value signals have no informational content at all, and when $\alpha = 1$, the common-value signal fully reveals Z. Thus, the parameter α effectively captures the informativeness of the common-value signal. In addition, each player i privately observes a second signal representing the idiosyncratic cost, c_i, of undertaking the investment. The cost c_i is independent of all other variables, and distributed according to a distribution function defined over the support, $[\underline{c}, \overline{c}]$. Impatience is measured by the discount factor, $0 < \delta < 1$. If player i has cost c_i and the state is Z, her profits are zero if she does not invest, and $\delta^{r-1}(Z - c_i)$ if she invests in round $r = 1, 2, ...$

Here is how the game proceeds. First, each player observes her private information, or type, (X_i, c_i). Let k_r be the number of players who invest in round r. For $r = 1, 2, ...$, each player observes the history of investment prior to round r, h_r, where $h_1 = \emptyset$ and, for $r \geq 2$, $h_r = (k_1, \ldots, k_{r-1})$. Players not yet invested simultaneously decide whether to invest in round r. We require that a player can invest at most once. In these settings, a strategy s is a function which assigns, for each type, an investment probability to each history of observed investment.

Although Levin and Peck (2008) consider continuous cost distributions, our experimental design considers a discrete distribution containing either one point ($c_i = L$ or $c_i = H$ but is the same for all subjects in a given game) or two points ($c_i = L$ or $c_i = H$ and can differ across subjects in a given game). This simplifies the decision making required of subjects and simplifies the data analysis. At the same time, it maintains the essential tradeoff between the incentive to delay and gain information by observing investment activity, versus the associated shrinkage of the (expected) pie due to discounting. For the remainder of the paper, we restrict attention to the parameter values, $n = 2$ or $n = 10$, $\delta = 0.9$, $\alpha = 0.7$, $L = 3.5$, and $H = 6.5$. Note that given these parameters, the expected profit of investing in round 1 is negative for types with $X_i = 0$ (-3.5 for type $(0, H)$ and -0.5 for type $(0, L)$) and positive for types with $X_i = 1$ (0.5 for type $(1, H)$ and 3.5 for type $(1, L)$).

We now define the games relevant to our experiment. We also compute NE behavior for rounds 1 and 2.[6] This computation relies on the "one-step property," which is proved by Levin and Peck (2008) for the case of a continuum of possible cost realizations. According to this property, for a type that invests with positive probability at a given history in NE, the option value of waiting at that history equals the expected payoff from waiting and investing in the next round if and only if the expected payoff of doing so is positive (and otherwise never investing). For the case of $n = 2$, the explicit derivation of the NE can be found in ILP. For the case of $n = 10$, NE behavior is computed numerically.

Two-Cost Games:

There are two equally likely cost realizations, $L = 3.5$ and $H = 6.5$. Thus, we have four possible types of players based on the common-value signal and the cost: $(0, H)$, $(0, L)$, $(1, H)$, and $(1, L)$. NE behavior in rounds 1 and 2 is shown in the left panels of Tables 2 (for $n = 2$) and 4 (for $n = 10$).[7]

Low-Cost Games:

There is only one possible cost realization, 3.5. Thus, we have two possible types of players: $(0, L)$ and $(1, L)$. NE behavior in rounds 1 and 2 is shown in the left

[6]For $n = 10$, the number of histories quickly becomes intractable beyond round 2. For $n = 2$, there are no decisions to invest beyond round 2 in NE and relatively few such decisions in the actual data (see ILP for details).

[7]In the two-cost games and in the two-player low-cost game (see below), type $(1, L)$ never reaches round 2 in NE. In these cases, we show the sequentially rational round-2 investment probabilities for type $(1, L)$ assuming others behave according to NE.

Two-Player One-Cost Games (A2 Treatment)

	Nash				Actual			
	(0,H)	**(0,L)**	**(1,H)**	**(1,L)**	**(0,H)**	**(0,L)**	**(1,H)**	**(1,L)**
Round 1	0	0	0.492	1	0.048 (589)	0.128 (564)	0.380 (545)	0.779 (574)
Round 2								
- after (1)	0	1	1	1	0.155 (103)	0.510 (206)	0.663 (80)	0.831 (71)
- after (0)	0	0	0	1	0.055 (458)	0.112 (286)	0.256 (258)	0.518 (56)

Table 1: Round-1 and round-2 frequencies of investment in the two-player one-cost games ($A2$ treatment). The subscripts on the actual frequencies show the number of decisions that were made at each history.

Two-Player Two-Cost Game (R2 Treatment)

	Nash				Actual			
	(0,H)	**(0,L)**	**(1,H)**	**(1,L)**	**(0,H)**	**(0,L)**	**(1,H)**	**(1,L)**
Round 1	0	0	0	1	0.072 (459)	0.113 (415)	0.353 (465)	0.771 (433)
Round 2								
- after (1)	0	1	1	1	0.216 (116)	0.452 (115)	0.679 (112)	0.892 (37)
- after (0)	0	0	0	1	0.058 (310)	0.111 (253)	0.228 (189)	0.597 (62)

Table 2: Round-1 and round-2 frequencies of investment in the two-player two-cost game ($R2$ treatment). The subscripts on the actual frequencies show the number of decisions that were made at each history.

panels of Tables 1 (for $n = 2$) and 3 (for $n = 10$) (only the columns corresponding to types $(0, L)$ and $(1, L)$ apply).

High-Cost Games:

There is only one possible cost realization, 6.5. Thus, we have two possible types of players: $(0, H)$ and $(1, H)$. NE behavior in rounds 1 and 2 is shown in the left panels of Tables 1 (for $n = 2$) and 3 (for $n = 10$) (only the columns corresponding to types $(0, H)$ and $(1, H)$ apply).

Ten-Player One-Cost Games (A10 Treatment)

	Nash				Actual			
	(0,H)	(0,L)	(1,H)	(1,L)	(0,H)	(0,L)	(1,H)	(1,L)
Round 1	0	0	0.071	0.746	$0.046_{(239)}$	$0.152_{(217)}$	$0.394_{(241)}$	$0.787_{(263)}$
Round 2								
- after (9)	1	1	1	1	-	$0.000_{(1)}$	-	-
- after (8)	1	1	1	1	-	$1.000_{(5)}$	-	$1.000_{(5)}$
- after (7)	1	1	1	1	-	$0.714_{(7)}$	-	$0.875_{(8)}$
- after (6)	1	1	1	1	$0.000_{(2)}$	$0.655_{(29)}$	$0.500_{(2)}$	$0.909_{(11)}$
- after (5)	1	1	1	1	$1.000_{(1)}$	$0.419_{(31)}$	$0.444_{(9)}$	$0.737_{(19)}$
- after (4)	1	0	1	1	$0.174_{(23)}$	$0.259_{(27)}$	$0.368_{(19)}$	$0.333_{(3)}$
- after (3)	0	0	1	0.763	$0.071_{(42)}$	$0.125_{(32)}$	$0.357_{(42)}$	$0.667_{(3)}$
- after (2)	0	0	1	0	$0.086_{(35)}$	$0.036_{(28)}$	$0.143_{(21)}$	$0.250_{(4)}$
- after (1)	0	0	0.322	0	$0.037_{(81)}$	$0.083_{(24)}$	$0.185_{(27)}$	$0.333_{(3)}$
- after (0)	0	0	0	0	$0.068_{(44)}$	-	$0.154_{(26)}$	-

Table 3: Round-1 and round-2 frequencies of investment in the ten-player one-cost games ($A10$ treatment). The subscripts on the actual frequencies show the number of decisions that were made at each history; "-" indicates that a given history never occurred.

Table 5 shows the cutoff values for k_1 above which each type invests in round 2. We are interested in the actual cutoffs, the NE cutoffs, and the cutoffs based on best-responding to the actual behavior of the other subjects. The NE cutoffs are complicated by the fact that the NE sometimes involves mixing. As a useful summary statistic of the NE round-2 behavior of each type in the ten-player games, we interpolate the cutoff value of the number of observed round-1 investments that induces investment with probability 0.5 in round 2. In particular, denoting this cutoff value by $\overline{k}_1^{NE}$, we set $\overline{k}_1^{NE} = \frac{(p_2-0.5)}{p_2-p_1}\widetilde{k}_1 + \frac{(0.5-p_1)}{p_2-p_1}(\widetilde{k}_1+1)$, where $\widetilde{k}_1$ is the highest level of round-1 investment for which the round-2 NE probability of investment is less than 0.5, p_1 is the NE probability of investment after history $(\widetilde{k}_1)$, and p_2 is the NE probability of investment after history $(\widetilde{k}_1+1)$. That is, $\overline{k}_1^{NE}$ is a weighed average of $\widetilde{k}_1$ and $\widetilde{k}_1+1$, where the weights depend on how close p_1 and p_2 are to 0.5. The third

Ten-Player Two-Cost Game (R10 Treatment)

	Nash				Actual			
	(0,H)	**(0,L)**	**(1,H)**	**(1,L)**	**(0,H)**	**(0,L)**	**(1,H)**	**(1,L)**
Round 1	0	0	0	1	0.058(241)	0.169(231)	0.348(250)	0.681(238)
Round 2								
- after (9)	1	1	1	1	-	-	-	-
- after (8)	1	1	1	1	-	-	-	-
- after (7)	1	1	1	1	0.500 (2)	1.000 (1)	-	-
- after (6)	1	1	1	1	0.429 (7)	1.000(3)	0.583(12)	0.667 (6)
- after (5)	1	1	1	1	0.300(20)	0.385(13)	0.571 (7)	1.000 (5)
- after (4)	0	1	1	1	0.167(42)	0.379(29)	0.467(45)	0.438 (16)
- after (3)	0	0	0.655	1	0.000(50)	0.186(43)	0.220(41)	0.550 (20)
- after (2)	0	0	0.020	1	0.034(59)	0.051(59)	0.087(46)	0.400 (20)
- after (1)	0	0	0	1	0.000(34)	0.000(26)	0.000(6)	0.167 (6)
- after (0)	0	0	0	0	0.077(13)	0.056 (18)	0.000(6)	0.333 (3)

Table 4: Round-1 and round-2 frequencies of investment in the ten-player two-cost game ($R10$ treatment). The subscripts on the actual frequencies show the number of decisions that were made at each history; "-" indicates that a given history never occurred.

	A10					R10			
	(0,H)	**(0,L)**	**(1,H)**	**(1,L)**		**(0,H)**	**(0,L)**	**(1,H)**	**(1,L)**
$\overline{k}_1$	8.71	5.33	5.54	2.93	$\overline{k}_1$	6.85	4.86	4.58	3.26
	(2.94)	*(0.37)*	*(1.29)*	*(0.84)*		*(0.82)*	*(0.39)*	*(0.37)*	*(0.76)*
$\overline{k}_1^{NE}$	3.5	4.5	1.26	2.66	$\overline{k}_1^{NE}$	4.5	3.5	2.76	0.5
$\overline{k}_1^{TBR}$	4.5	4.5	1.5	3.5	$\overline{k}_1^{TBR}$	5.5	3.5	3.5	1.5

Table 5: Round-2 actual, Nash, and truncated best-response cutoffs.

row in the left and right panels of Table 5 shows $\overline{k}_1^{NE}$ for each type in the one-cost ten-player games and the two-cost ten-player game, respectively. (The actual cutoffs, $\overline{k}_1$, and the best-response cutoffs, $\overline{k}_1^{TBR}$, in the remaining rows are explained further below.)

The NE of the games described above exhibit the following theoretical predictions:

(1) In each game, the higher the expected profit from investment given a subject's type, the more likely she is to invest in round 1–the frequencies with which subjects of type $(0, H)$, $(0, L)$, $(1, H)$, and $(1, L)$ invest in round 1 are (weakly) increasing in the given order.

(2) For each type in each game, the round-2 probability of investment is (weakly) increasing in the number of subjects who are seen to invest in round 1.

(3) In all games, $(0, H)$ and $(0, L)$ players never invest in round 1 because that entails a negative expected profit.

(4) In the two-player games, $(1, L)$ players invest in round 1 with probability 1–in fact, investing in round 1 dominates waiting.[8] In the two-cost ten-player game, $(1, L)$ players also invest in round 1 with probability 1.

(5) Sensitivity to market size:

For a given cost structure (one-cost or two-cost game), the round-1 probability of investment for each type is (weakly) smaller in the larger market because the investment activity that can be observed by waiting is more informative by virtue of sample-size effects. Notably, type $(1, H)$ invests in round 1 with probability 0.49 in the high-cost two-player game and with probability 0.07 in the high-cost ten-player game; type $(1, L)$ invests in round 1 with probability 1 in the low-cost two-player game and with probability 0.75 in the low-cost ten-player game.

(6) Sensitivity to opportunities to free-ride:

For a given market size ($n = 2$ or $n = 10$), the round-1 probability of investment

[8] Investing in round 1 dominates waiting because, even if a $(1, L)$ player were to find out for sure in round 2 that the other player has an unfavorable common-value signal, it would still be optimal to invest.

for type $(1, H)$ is smaller in the two-cost game than in the high-cost game. The reason is that, in the two-cost game, type $(1, H)$ can free-ride on the stronger $(1, L)$ type.

When $n = 10$, the round-1 probability of investment for type $(1, L)$ is smaller in the low-cost game than in the two-cost game. The reason is that, in the two-cost game, only half of the players with $X_i = 1$ are type $(1, L)$, so that the opportunities for free-riding are limited; on the other hand, in the low-cost game, all players with $X_i = 1$ are type $(1, L)$, so that $(1, L)$ players end up free-riding on each other and the NE involves mixing.

(7) Sensitivity to the informational content of market activity:

When $n = 10$, $\overline{k}_1^{NE}$ for types $(0, H)$ and $(1, H)$ is higher in the two-cost game than in the high-cost game. The reason is that, in the two-cost game, players with a low investment cost invest with a high probability in round 1 if $X_i = 1$. Thus, relatively low round-1 investment makes the presence of type $(0, L)$ more likely, which is bad news. In the high-cost game, no type invests with a high probability in round 1, so that the same level of round-1 investment can still be viewed as good news.

When $n = 10$, $\overline{k}_1^{NE}$ for types $(0, L)$ and $(1, L)$ is lower in the two-cost game than in the low-cost game. The reason is that, in the two-cost game, half of the players with $X_i = 1$, namely the $(1, H)$ players, do not invest in round 1 despite their favorable common-value signal. Thus, relatively low round-1 investment can still be viewed as good news. In the low-cost game, only one-quarter of the players with $X_i = 1$, namely the $(1, L)$ players who wait as part of the NE mixing, do not invest in round 1. Thus, the same level of round-1 investment may no longer be viewed as good news.

The main goal of our paper is to check to what extent actual behavior in the lab conforms to the NE features of behavior described in equilibrium predictions (1)-(7). Note that equilibrium predictions (1)-(2) capture aspects of behavior that are about basic rationality. On the other hand, the remaining equilibrium predictions, especially equilibrium predictions (5)-(7), are based on more subtle considerations about the logic of the game.

Note that some departures from NE may still be optimal given the empirical frequencies of play. Therefore, we will also wish to consider the best-response behavior of a player given that all other players use the average empirical strategy in the population, s^a, which, for each history and type, gives the empirical frequency of investment at that history by that type.[9] [10] However, in the ten-player games, we will not compute exact best-responses.[11] Instead we compute truncated best-response (TBR) as best-response behavior under the restriction that it does not involve any investment after round 3. TBR behavior and best-response behavior coincide in the two-player games. In the ten-player games, TBR profits provide a lower bound for best-response profits, though we believe the difference is very small.

As a useful summary statistic of the TBR round-2 behavior of each type in the ten-player games, we interpolate the cutoff value of the number of observed round-1 investments that induces investment with probability 0.5 in round 2. In particular, denoting this cutoff value by $\overline{k}_1^{TBR}$, we set it equal to the average of the highest k_1

[9] If there are no observations at some history h_r in our experiment, so that the relevant frequency is not defined, we set $s^a(h_r; X_i, c_i) = s^{NE}(h_r; X_i, c_i)$, where s^{NE} is the NE strategy.

[10] More generally, one could look at the best-response behavior and profits given the distribution of strategies in the population, rather than assuming that everyone uses the average empirical strategy. However, for our purposes, this would be impractical because we do not observe subjects' strategies. In addition, requiring that a best-responder knows not just the empirical frequencies of play, but the whole distribution of strategies, may be too demanding.

[11] For example, we cannot compute exact best response behavior given s^a because, at each history beyond round 2, we have few observations so that the empirical frequencies are either not defined or are not reliable. Thus, it is not clear how a player, say, in round 4 should respond to observed investment in round 3.

for which TBR prescribes not investing and the lowest k_1 for which TBR prescribes investing. Thus, $\overline{k}_1^{TBR}$ is the TBR analogue of $\overline{k}_1^{NE}$. The fourth row in the left and right panels of Table 5 shows $\overline{k}_1^{TBR}$ for each type in the one-cost ten-player games and the two-cost ten-player game, respectively.

3 Experimental Design

The experiment consisted of the two-player random two-cost treatment (R2), the two-player alternating one-cost treatment (A2), the ten-player random two-cost treatment (R10), and the ten-player alternating one-cost treatment (A10). The R2 consisted of four sessions (78 participants in total). The A2 also consisted of four sessions (96 participants in total). We conducted two sessions of the R10 (42 participants in total) and two sessions of the A10 (54 participants in total).[12]

In the R10, each session consisted of 2 practice periods and 24 periods in which subjects played for real money. At the start of each period, subjects were randomly and anonymously matched in two groups of ten to form separate ten-player markets that bore no relation to each other.[13] In any given market, subjects played the two-cost game.

The A10 was identical to the R10, except that subjects played the low-cost game in odd numbered periods and the high-cost game in even numbered periods. The R2/A2 was identical to the R10/A10 except that there were two players per market.[14] [15]

[12]The data from the two-player treatments is analyzed in ILP. The analysis in this paper addresses new questions not considered in ILP, including the effect of market size and the other behavioral biases we identify.

[13]Thus, in any given period, only twenty subjects played. The rest, who were randomly chosen each period, sat out.

[14]To guarantee that the trials ended, without changing the equilibria, subjects in all treatments were told that the game ended after either all subjects had invested or there were two consecutive rounds with no investment.

[15]For 2 sessions of the R2 and of the A2, the investment game was followed by some lottery problems and a questionnaire. ILP use these to test their interpretation of behavior. We do not use

In the two-player/ten-player treatments, subjects were given an initial cash balance of 20/30 experimental currency units (ECU). In addition, they could gain or lose ECU in each trial, which were added to or subtracted from their cash balances. At the end of the session in the two-player treatments, ECU were converted into dollars at a rate of \$0.6 per ECU. In the ten-player games, the exchange rate was \$0.5 per ECU. Subjects were paid the resulting dollar amount or \$5, whichever was greater. If a subject's cash balances fell below 0 at any point during the session, that subject was paid \$5 and was asked to leave.[16]

Average earnings for the R2/A2/R10/A10 were \$26.04/\$26.49/\$25.71/\$26.32. Sessions lasted between 1 hour 45 minutes and 2 hours.

Subjects in all treatments were undergraduate students at The Ohio State University (OSU). The sessions were held at the Experimental Economics Lab at OSU. At the beginning of each session, the experimenter read the instructions aloud as subjects read along, seated at their computer terminals. Subjects were invited to ask questions during the instructions and after the practice periods. Once play for real cash began, no more questions were allowed. See the Appendix for our Instructions from the $R10$[17] and a printout of the screen seen by a subject in the A2 with cost 6.5 and signal 1, who is deciding whether to invest in round 2 after the other subject has invested in round 1.

4 Results

We will focus on actual behavior mostly in rounds 1 and 2. The right panel in Table 1 summarizes actual behavior in rounds 1 and 2 in the $A2$. The top row shows, for each

the data from the lottery problems and the questionnaire in our current study.

[16]This occurred for three/three/one/zero subjects in the R2/A2/R10/A10.

[17]The instructions for the other treatments are similar and are available upon request.

type, the actual frequencies with which subjects invested in round 1; the subscripts in parentheses show the number of decisions that these frequencies are based on. For example, in round 1 subjects made 564 decisions as type $(0, L)$ and 12.8 percent of these decisions were decisions to invest in round 1. The remaining rows show, for each type and each history of round-1 investment, the actual frequencies with which subjects invested in round 2 and the number of decisions that these frequencies are based on. For example, in round 2 subjects made 258 decisions as type $(1, H)$ after observing 0 investments in round 1; 25.6 percent of these 258 decisions were decisions to invest. The right panels in Tables 2, 3, and 4 analogously summarize actual behavior in rounds 1 and 2 in the $R2$, $A10$, and $R10$, respectively.

Before we proceed, let as make a brief note on statistical methodology. In the statistical analysis below, we treat the unit of observation as the subject-trial pair, taking into account individual-specific unobserved effects but not session-level effects.[18] To the extent that within-session correlation exists, our econometric tests will presume "too many" observations and will tend to reject the null hypothesis too often. We are on safe ground, since our main results are about an *absence* of differences in behavior across treatments. Thus, we fail to reject the null of no difference across treatments despite using a possibly overly sensitive test.[19]

4.1 Basic Rationality

Let us start by checking if subjects respond to obvious incentives. ILP show that, for the two-player markets, in the aggregate subjects respond to their investment cost and common-value signal in a sensible way, corresponding to equilibrium prediction (1). The same is true for the ten-player markets (see the first row in Tables 3 and 4).

[18] In the statistical analysis in the Appendix, there are some exceptions to this rule. These are explicitly mentioned.

[19] Further details are provided in the appendix.

	$(\mathbf{0},\mathbf{H})$	$(\mathbf{0},\mathbf{L})$	$(\mathbf{1},\mathbf{H})$	$(\mathbf{1},\mathbf{L})$
A2	0.085**	0.424***	0.439***	0.430***
R2	0.143***	0.380***	0.544***	0.388***
A10	0.021**	0.150***	0.064***	0.108**
R10	0.040***	0.099***	0.165***	0.129**

Table 6: Marginal Effect of k_1 on the probability of investment in round 2. (*/**/*** indicates significance at the 10/5/1 percent level.)

We now move on to the question of whether subjects respond to the behavior of the other subjects in their trial. We estimate the following random effects probit model for each type in each treatment:

$$Pr(\text{subject invests in round } 2|k_1, v, \text{subject has not invested in round } 1) = \Phi(\beta_0+\beta_1 k_1+v) \tag{1}$$

where k_1 is the number of players who invested in round 1, v is an individual-specific random effect, and $\Phi(\cdot)$ is the standard normal cdf. Table 6 shows the estimated marginal effects of a one-unit increase in k_1 on the probability of investment in round 2. All marginal effects are significant.

Let us summarize our results so far:

Result 1 *In the aggregate, for all treatments, (i) types with higher expected profits are more likely to invest in round 1, and (ii) for each type, there is a positive and statistically significant marginal effect of k_1 (the number of subjects who invest in round 1) on the probability of investment in round 2.*

4.2 Departures from NE

We now consider systematic departures from NE in rounds 1 and 2. Because some departures from NE may still be optimal given the empirical frequencies of play, we

	(0,H)	(0,L)	(1,H)	(1,L)
A2	-3.500	-0.812	0.022	0.350
R2	-3.500	-0.727	-0.110	0.350
A10	-3.624	-1.428	-0.556	0.063
R10	-3.608	-1.170	-0.598	0.301

Table 7: Expected profit of investing in round 1 minus expected profit from waiting in round 1 and behaving according to TBR thereafter (in ECU).

will also consider TBR behavior given the empirical frequencies.

4.2.1 Round-1 Behavior

At various points below, we will consider, for each type, incentives for investing in round 1 versus waiting given the empirical frequencies of play. Table 7 shows, for each type in each treatment, the expected payoff of investing in round 1 minus the expected payoff of waiting in round 1 and behaving according to TBR thereafter. We call this difference the TBR net payoff from investing (in round 1).[20]

The next result summarizes our findings regarding round-1 behavior in the lab based on Tables 1, 2, 3, and 4.

Result 2 *In round 1 in all treatments of our experiment:*

(i) there is almost no investment by $(0, H)$ players (5-7 percent) and small but nonnegligible investment by $(0, L)$ players (11-17 percent);

(ii) the frequency of round-1 investment for type $(1, H)$ is in the range 0.35-0.39;

(iii) the frequency of round-1 investment for type $(1, L)$ is in the range 0.68-0.79.

(iv) for each type, the difference in round-1 investment between any two treatments is not statistically significant.[21]

[20]Note that, on average, a player is each type six times per session. Thus, each number in the table needs to be multiplied by six in order to obtain an idea of incentives for the whole session.

[21]Let $d_{A2}, d_{R2}, d_{A10}, d_{R10}$ be treatment dummies and v be an individual-specific random effect.

Thus, by Result 2(i) subjects, in contrast with equilibrium prediction (3), make mistakes in round 1 by sometimes investing despite a negative expected payoff from doing so. The frequency of mistakes depends on the expected payoff from investing and is nonnegligible when this expected payoff is only slightly negative.

By Result 2(iii), $(1, L)$ subjects, in contrast with equilibrium prediction (4), also make mistakes by failing to always invest in round 1 in the two-player games and the two-cost ten-player game. In the two-player games, investing strictly dominates waiting. In the two-player games and the two-cost ten-player game, the TBR net payoffs from investing are strictly positive (see Table 7).

These departures from NE could be viewed as noise that inevitably moves the empirical frequencies of play away from NE boundary values of 0 or 1. However, Result 2 also provides evidence that subjects exhibit the following behavioral biases that are at odds with equilibrium predictions (5) and (6).

Insensitivity to Market Size:

In contrast with equilibrium prediction (5), Result 2 indicates that, for a given cost-structure of the game, round-1 behavior in the lab is insensitive to the size of the market. Notably, the round-1 probability of investment of type $(1, H)$ is slightly higher in the ten-player high-cost game than in the two-player high-cost game (0.39 vs. 0.38, respectively) even though it is much lower in NE (0.07 vs. 0.49, respectively); the round-1 probability of investment of type $(1, L)$ is slightly higher in the ten-player low-cost game than in the two-player low-cost game (0.79 vs. 0.78, respectively) even though it is quite a bit lower in NE (0.75 vs. 1, respectively).

Further, for a given cost structure, the behavior of the various types is insensitive to the size of the market despite differences in the TBR incentives to invest (see

To check for statistical significance, for each type we run the following random effects probit $Pr(\text{subject invests in round } 1 | d_{A2}, d_{R2}, d_{A10}, d_{R10}, v) = \Phi(\beta_1 d_{A2} + \beta_2 d_{R2} + \beta_3 d_{A10} + \beta_4 d_{R10} + v)$ and, for any $i, j \in \{1, 2, 3, 4\}, i \neq j$, we test the hypothesis $\beta_i = \beta_j$.

Table 7). First, consider $(1, H)$ subjects. In the high-cost games, the difference between the TBR net payoffs from investing in the two-player and in the ten-player game is $|-0.56-0.02| = 0.58$ ECU. In the two-cost games, the difference between the TBR net payoffs from investing in the two-player and in the ten-player game is $|-0.60-(-0.11)| = 0.49$ ECU.

In the case of $(1, L)$ subjects, in the low-cost games the TBR incentives differ by a modest but nonnegligible amount ($|0.35-0.06| = 0.29$ ECU). In the case of $(0, L)$ subjects, for a given cost structure, the investment frequency in the lab is actually a bit larger in the ten-player market even though the TBR net payoffs from investing are actually quite a bit smaller. All this evidence suggests that subjects fail to appreciate the relevant sample-size effects.

Insensitivity to Opportunities to Free-Ride:

In contrast with equilibrium prediction (6), we observe that, for a given market size, round-1 behavior in the lab is insensitive to the cost structure of the game. Notably, for a given market size, the round-1 probability of investment of type $(1, H)$ is very close between the high-cost and two-cost games (and any difference is not statistically significant by 2(iv)) even though the NE probability of investment differs, especially when $n = 2$; the round-1 probability of investment of type $(1, L)$ is higher in the ten-player low-cost game than in the ten-player two-cost game (0.79 vs. 0.68, respectively) even though the reverse is true in NE (0.75 vs. 1, respectively).

Further, for a given market size, the behavior of the various types is insensitive to the cost structure of the game despite differences in the TBR incentives to invest (see Table 7). In the case of $(1, H)$ subjects, the difference between the TBR net payoffs from investing in the one-cost two-player and in the two-cost two-player game is modest but nonnegligible ($|-0.11-0.02| = 0.13$ ECU). In the case of $(1, L)$

subjects, the TBR incentives differ by $|0.30 - 0.06| = 0.24$ ECU between the low-cost and the two-cost ten-player game. The corresponding number for $(0, L)$ subjects is $|-1.43 - (-1.17)| = 0.26$ ECU. All this evidence suggests that subjects are insensitive to opportunities to free-ride on the information provided by other subjects' behavior.

4.2.2 Round-2 Behavior

Our first goal in this section is to check to what extent behavior in the ten-player games in the lab conforms to equilibrium prediction (7). To do that, we consider again the random effects probit model given in equation (1). Based on this, we compute $\overline{k}_1$ as the value of k_1, such that the predicted probability of investment for the average subject (i.e. for one with $v = 0$) equals 0.5.[22] Thus, $\overline{k}_1$ is the actual-behavior analogue of $\overline{k}_1^{NE}$ and $\overline{k}_1^{TBR}$, which were introduced in section 2.

Table 5 reports, for each type in the ten-player games, $\overline{k}_1$ along with the estimated standard error. We can state the following result.

Result 3

(i) For each type, $\overline{k}_1$ is similar in magnitude between the $A10$ and $R10$ and any difference is not statistically significant.[23]

(ii) For each type in the $A10$ and $R10$, $\overline{k}_1$ is greater than $\overline{k}_1^{NE}$; the difference is statistically significant (at the 5-percent level) for types $(0, L)$ and $(1, H)$ in the

[22] That is, $\overline{k}_1 = -\frac{\widehat{\beta}_0}{\widehat{\beta}_1}$, where $\widehat{\beta}_0$ and $\widehat{\beta}_1$ are the estimates of β_0 and β_1 from equation (1).

[23] To test for statistical significance between $\overline{k}_1$ from the $A10$ and $\overline{k}_1$ from the $R10$, we run the following version of the probit model given in equation (1) jointly on the data from the $A10$ and the $R10$: $Pr(\text{subject invests in round } 2|k_1, d_{A10}, d_{R10}, v, \text{subject has not invested in round } 1) = \Phi(\beta_0 d_{A10} + \beta_1 d_{A10} k_1 + \gamma_0 d_{R10} + \gamma_1 d_{R10} k_1 + v)$, where d_{A10}/d_{R10} is a dummy that equals 1 in the $A10/R10$. (Thus, we're allowing for treatment-specific constant and slope coefficient, which is effectively the same thing as running the model in equation (1) separately on the data from the $A10$ and the $R10$.) We test the hypothesis that $\overline{k}_1$ from the $A10$ equals $\overline{k}_1$ from the $R10$ by testing $\frac{\beta_0}{\beta_1} = \frac{\gamma_0}{\gamma_1}$. Later, we will test more general hypotheses of the form $\overline{k}_1$ from the $A10$ minus $\overline{k}_1$ from the $R10$ equals some constant c. This will be done by testing $\frac{\beta_0}{\beta_1} - \frac{\gamma_0}{\gamma_1} = c$.

$A10$ and for all types in the $R10$. For each type in the $A10$ and $R10$, except for type $(1, L)$ in the $A10$, $\overline{k}_1$ is greater than $\overline{k}_1^{TBR}$;[24] *the difference is statistically significant (at the 5-percent level) for types $(0, L)$ and $(1, H)$ in the $A10$ and for types $(0, L)$, $(1, H)$, and $(1, L)$ in the $R10$.*

Result 3 provides evidence that subjects exhibit two additional behavioral biases.

Insensitivity to the Informational Content of Market Activity:

By Result 3(i), round-2 behavior in the ten-player games in the lab is insensitive to the cost structure of the game. This is true even when NE and TBR behavior vary with cost-structure. Most notably, $\overline{k}_1$ for type $(1, L)$ is slightly higher in the $R10$ than in the $A10$, even though both $\overline{k}_1^{NE}$ and $\overline{k}_1^{TBR}$ are lower by at least 2 units in the $R10$. We can reject the hypothesis that $\overline{k}_1$ in the $A10$ minus $\overline{k}_1$ in the $R10$ is at least 2 (p-value is 0.019; one-sided test). For type $(1, H)$, $\overline{k}_1$ is higher in the $A10$ than in the $R10$, even though both $\overline{k}_1^{NE}$ and $\overline{k}_1^{TBR}$ are lower in the $A10$ by at least 1.5 units. We can reject the hypothesis that $\overline{k}_1$ in the $R10$ minus $\overline{k}_1$ in the $A10$ is at least 1.5 (p-value is 0.039; one-sided test). Thus players fail to appreciate the different informational content of observed investment activity in the ten-player one-cost and two-cost games.

Excessive Conservatism in Updating from Market Activity:

This bias follows directly from Result 3(ii). Note that it is consistent with a general finding from the exogenous-timing herding literature that subjects underappreciate the informational content of observed investment activity (see Weizsäcker (2010) for a meta-analysis).

[24]For type $(1, L)$ in the $A10$, $\overline{k}_1$ is lower than $\overline{k}_1^{TBR}$, but the difference is not statistically significant.

4.3 Departures from Risk Neutrality

So far, we have been conducting our analysis under the assumption of risk neutrality. To what extent are our conclusions robust to alternative attitudes to risk?

First, round-1 investment by $(0, H)$ and $(0, L)$ players could be due to risk-loving. A failure by $(1, L)$ players to always invest in round 1 even when investing dominates waiting (under risk neutrality) could be due to risk aversion.[25]

Second, three of our biases–insensitivity to market size, to opportunities to free-ride, and to the informational content of observed market activity–are about how behavior fails to respond to changes in the structure of the game. The logic for why behavior should respond to changes in the set of possible types and changes in market size does not hinge on a particular attitude to risk. Thus, a failure to respond cannot be explained by departures from risk neutrality.

Third, subjects' excessive conservatism in round 2 could be explained by risk aversion. One might even conjecture that more risk averse subjects wait in round 1, so that self-selection would foster more risk averse behavior in round 2. To test this conjecture, we split subjects into two equal-sized groups depending on their frequency of investment in round 1 (averaged over types). We then check whether subjects in the low-round-1-investment group are more conservative in round 2 than subjects in the high-round-1-investment group. We find no evidence at all in favor of this hypothesis. This finding casts a doubt on the view that attitudes to risk are a driving force behind behavior in our experiment.[26]

[25]Nevertheless, our conclusions in the Appendix regarding how these types' behavior affects informational externalities and market efficiency in the lab relative to the risk-neutral NE remain valid.

[26]In the appendix, we address two additional issues. First, we compare aggregate investment levels in the lab vs. in NE. Second, we analyze learning over the course of our sessions, an issue that, for simplicity, we have avoided in the main analysis.

5 Concluding Remarks

We study experimentally behavior in a series of endogenous-timing investment games. We find that subjects respond to their type and to observed investment activity in a sensible way, but also make mistakes. Some mistakes can be viewed as noise while other mistakes represent systematic biases reflecting subjects' failure to appreciate subtle aspects of the environment. From a more general perspective, our results can be viewed as lending support to a particular notion of bounded rationality. According to this notion, people respond to incentives when these incentives, as well as the appropriate response, are fairly obvious. However, people may miss more subtle considerations that play a crucial role in the theory.

Herding models have been used to study adoptions of new technologies or behavior in financial markets. Entry models typically assume fully rational agents, and finance models assume some people are fully rational and others are noise traders. Our study suggests that people exhibit systematic biases such as the four we identify here. Ideally, real-world applications should take into account the way that people actually behave, rather than making convenient noise-trader assumptions.

References

Anderson, L.R., Holt C.A., 1997. Information Cascades in the Laboratory. American Economic Review, vol. 87, no. 5, pp. 847-62.

Avery, C., Zemsky, P., 1998. Multidimensional Uncertainty and Herd Behavior in Financial Markets. American Economic Review, vol. 88, no. 4, pp. 724-48.

Banerjee, A.V., 1992. A Simple Model of Herd Behavior. Quarterly Journal of Economics, vol. 107, no. 3, pp. 797-817.

Bikhchandani, S., Hirshleifer, D., Welch, I., 1992. A Theory of Fads, Fashion, Custom, and Cultural Change in Informational Cascades. Journal of Political Economy, vol. 100, no. 5, pp. 992-1026.

Chamley, C., Gale, D., 1994. Information Revelation and Strategic Delay in a Model of Investment. Econometrica, vol. 62, no. 5, pp. 1065-85.

Cipriani, M., Guarino, A., 2005. Herd Behavior in a Laboratory Financial Market. American Economic Review, 95(5): 1427-1443.

Crawford, V.P., Iriberri, N., 2007. Level-k Auctions: Can a Nonequilibrium Model of Strategic Thinking Explain the Winner's Curse and Overbidding in Private-Value Auctions? Econometrica, 75(6): 1721–70.

Drehmann, M., Oechssler, J., Roider, A., 2005. Herding and Contrarian Behavior in Financial Markets: An Internet Experiment. American Economic Review, 95(5), 1403-1426.

Eyster, E., Rabin, M., 2005. Cursed Equilibrium. Econometrica, 73(5): 1623-1672.

Fischbacher, U., 2007. Z-Tree: Zurich Toolbox for Ready-Made Economic Experiments. Experimental Economics, 10(2): 171-78.

Gigerenzer, G., Selten, R. (Eds.), 1999. Bounded Rationality: The Adaptive Toolbox. The MIT Press, Cambridge, MA and London, England.

Gigerenzer, G., Todd, P.M., ABC Research Group, 1999. Simple Heuristics That Make Us Smart. Oxford University Press.

Ivanov, A., Levin, D., Peck, J., 2009. Hindsight, Foresight, and Insight: An Experimental Study of a Small-Market Investment Game with Common and Private Values. American Economic Review, vol. 99, no. 4, pp. 1484-1507.

Levin, D., Peck, J., 2008. Investment Dynamics with Common and Private Values. Journal of Economic Theory, vol. 143, no. 1, pp. 114-39.

McKelvey, R.D., Palfrey, T.R., 1995. Quantal Response Equilibria for Normal Form Games. Games and Economic Behavior, 10: 6-38.

McKelvey, R.D., Palfrey, T.R., 1998. Quantal Response Equilibria in Extensive Form Games. Experimental Economics, 1: 9-41.

Nagel, R., 1995. Unraveling in Guessing Games: An Experimental Study. American Economic Review, 85: 1313-1326.

Park, A., Sgroi, D., 2012. Herding, Contrarianism and Delay in Financial Market Trading. European Economic Review 56: 1020-1037

Sgroi, D., 2003. The Right Choice at the Right Time: A Herding Experiment in Endogenous Time, Experimental Economics, 6: 159-180.

Simon, H., 1956. Rational Choice and the Structure of the Environment. Psychological Review, 63, 129-138.

Stahl, D.O., Wilson, P.W., 1994. Experimental Evidence on Players' Models of Other Players. Journal of Economic Behavior and Organization, vol. 25, no. 3, pp. 309-27.

Stahl, D.O., Wilson, P.W., 1995. On Players' Models of Other Players: Theory and Experimental Evidence. Games and Economic Behavior, 10: 218-254.

Weizsäcker, G., 2010. Do we follow others when we should? A simple test of rational expectations. American Economic Review, 100, pp. 2340-2360

Ziegelmeyer, A., My, K.B., Vergnaud, J.-C., Willinger, M., 2005. Strategic Delay and Rational Imitation in the Laboratory. Max Planck Institute of Economics Discussion Paper on Strategic Interaction 2005–35.

6 Appendix: A Note on Statistical Methodology

In the statistical analysis in the main text, we treat the unit of observation as the subject-trial pair, taking into account individual-specific unobserved effects but not session-level effects. To the extent that session-level effects are present, the effective number of observations will be somewhere in between the number of subject-trial pairs and the number of sessions. However, our main results center on the conclusion that subjects do not respond to incentives reflecting subtle aspects of the game. In other words, we show that the null hypothesis of no differences in behavior across treatments cannot be rejected. Thus, by potentially overstating the number of effective observations, we are using an overly sensitive test that makes rejecting the null hypothesis even more likely.

To further address this issue, for each type in each treatment, we tested the hypothesis that there are no session-level effects on round-1 and round-2 behavior. Regarding round-1 behavior, we ran the following random effects probit for each type in each treatment:

$$Pr(\text{subject invests in round 1}|\{d_s\}_{s=1}^{S}, v) = \Phi(\sum_{s=1}^{S} \beta_s d_s + v),$$

where v is an individual-specific random effect, S is the number of sessions in the given treatment, and d_s is a session dummy that equals 1 in session s. Then we test the hypothesis $\beta_1 = \ldots = \beta_S$, which gives us 16 hypothesis tests (4 treatments $\times$ 4 types). We can reject the null at the 5-percent level in only one instance, which is in line with what one would expect in the absence of session-level effects.

Regarding round-2 behavior in the $A10$ and $R10$, we ran the following random effects probit for each type in each of these two treatments:

$$Pr(\text{subject invests in round } 2|d_1, d_2, k_1, v, \text{subject has not invested in round 1}) =$$
$$\Phi(\beta_0 d_1 + \beta_1 d_1 k_1 + \gamma_0 d_2 + \gamma_1 d_2 k_1 + v).$$

This model allows for session-specific constants and slope coefficients. We test the hypothesis $\frac{\beta_0}{\beta_1} = \frac{\gamma_0}{\gamma_1}$, which tests whether the cutoff $\overline{k}_1$ is the same in both sessions. This gives us 8 tests (2 treatments $\times$ 4 types). We can reject the null at the 5-percent level in only one instance, which is in line with what one would expect in the absence of session-level effects.

We also stepped back and considered whether any plausible sources of within-session correlation might be present in our experiment. There were no differences in how sessions of the same treatment were conducted, because we conducted all sessions by adhering strictly to the same protocol with the same personnel. The sessions were conducted using the same subject pool. We also feel that it is very unlikely that a small group of outlier subjects might contaminate the behavior of the other subjects in their session. The reason is that, in our experiment, a subject merely decides whether to invest or not; given that her type (and subject number) is never observed by others, neither decision appears particularly odd to the other players. We feel that it is extremely unlikely that correlation is introduced through supergame effects, because in our games there are no payoff externalities. It would take a fairly sophisticated understanding of the game to even figure out how to reward or punish other players by affecting information flows, and a sophisticated player would surely realize that such rewards or punishments would not be recognized as such by the other players.

7 Appendix: Informational Externalities and Market Efficiency

In endogenous-timing herding games, a crucial question is how well the market aggregates information and yields efficient outcomes. In the NE of the games we consider, information aggregation is suboptimal because of free-riding. In particular, in the NE of each game, investment by the $(1, H)$ type, as well as by the $(1, L)$ type in the ten-player low-cost game, is too low from an efficiency point of view.[27] The joint-profit-maximizing probability of investment is 1 for the $(1, L)$ type in all games and 0.746/0.77/0.783/0.568 in the two-player high-cost/two-player two-cost/ten-player high-cost/ten-player two-cost game for the $(1, H)$ type.[28] [29]

An important question is how the departures from NE that we identified earlier affect information aggregation and market efficiency relative to the NE benchmark. Note that this question may not have an unequivocal answer as the answer may depend on the structure of the game.[30] In this appendix, we explore how the interaction between departures from NE and the structure of the game affects informational externalities and market efficiency.

To study the effect of departures from NE on informational externalities, we develop a measure of the latter. Our measure equals the expected profits of a player who optimally exploits both her private information and the information available on the market minus the expected profits of a player who optimally exploits her private

[27] We measure efficiency by average profits.

[28] $(0, H)$ and $(0, L)$ players should never invest in round 1 because (i) they have a negative expected profit of doing so and (ii) any investment by these types only dilutes the informational externalities.

[29] In computing the joint-profit-maximizing behavior, we are assuming symmetric strategies. For the ten-player games, we restricted attention to strategies where players invest in round 2 if and only if it is profitable and otherwise never invest.

[30] The idea that different environments may cater differently to bounded rationality goes back at least to Herbert Simon (e.g see Simon (1956); see also Gigerenzer and Selten (1999) as well as Gigerenzer et al. (1999)).

information but does not observe market activity. This measure captures the informational flows on the market that can potentially be exploited by a shrewd investor who understands the market well.

In our experiment, informational externalities are either smaller or larger than in NE, depending on how the specific structure of the game caters to different departures from NE. For example, in larger markets subjects' insensitivity to market size tends to enhance informational externalities relative to NE. This occurs because, while investment in the lab is insensitive to market size, in the NE players with favorable private information respond to increases in market size by increasing their propensity to delay investment, thereby reducing informational externalities.

We also compare market efficiency in the lab with the NE benchmark. Efficiency is related to informational externalities, but the two notions are distinct because players could be making mistakes in exploiting informational externalities or their private information. Given that in NE players make no mistakes, a necessary (but not sufficient) condition for increased efficiency in the lab is that the informational externalities exceed those in NE. However, even in the one game where informational externalities in the lab exceed the NE benchmark, average profits fail to exceed NE profits.

Before we proceed, note that, because we have computed NE behavior only for rounds 1 and 2 in the ten-player games, we cannot compute NE profits. Instead, we define truncated Nash equilibrium (TNE) in which players behave according to the NE in rounds 1 and 2, and, in round 3, invest if it is profitable to invest (given others' NE behavior in rounds 1 and 2) and otherwise never invest. In the two-player games, NE and TNE coincide. For types $(1, H)$ and $(1, L)$ in the ten-player games, TNE profits equal NE profits.[31] For types $(0, H)$ and $(0, L)$ in the ten player games, and

[31]It is straightforward to show that this follows from the one-step property.

hence for each of the ten-player games as a whole, TNE profits provide a lower bound for NE profits, though we believe the difference is very small.

7.1 Informational Externalities

Are the markets in the lab more or less informative than the NE benchmarks? To answer this question, we start by defining a measure of the informational externalities on a market, and we compute this measure given the empirical frequencies of play as well as given NE play. Later, we will explore how the structure of the game and subjects' departures from NE affect the comparison.

Given a strategy s, we define our measure of the informational externalities as $IE(s) = \Pi^{BR}(s) - \underline{\Pi}$, where $\Pi^{BR}(s)$ is the expected profit of a player who best-responds to a population of players each of whom uses s, and $\underline{\Pi}$ is the expected profit of someone who behaves optimally given her private information but does not observe others' behavior.[32] Thus, $IE(s)$ shows the expected profits of a player who knows s and best responds to it, over and above what she can earn based solely on her private information. Thus, it is a measure of the informational flows on the market that can potentially be exploited by a shrewd investor.

Ideally, we would like to compute and compare the informational externalities in NE, $IE(s^{NE})$, and given the empirical frequencies of play, $IE(s^a)$. However, as explained above, we cannot compute exact NE or best-response profits. Therefore, as a proxy for $IE(s^{NE})$, we use $IE^T(s^{NE}) = \Pi^{TBR}(s^{NE}) - \underline{\Pi}$ and, as a proxy for $IE(s^a)$, we use $IE^T(s^a) = \Pi^{TBR}(s^a) - \underline{\Pi}$. ($\Pi^{TBR}(s)$ is the expected profit of someone who plays according to TBR given that all other players use strategy s.)

Table 8 shows $\underline{\Pi}$, as well as $\Pi^{TBR}(s^a)$ and $\Pi^{TBR}(s^{NE})$ for each game in each treatment. From the information in the table, $IE^T(s^a)$ and $IE^T(s^{NE})$ can readily be

[32] In our games, the latter kind of player invests in round 1 for types $(1, H)$ and $(1, L)$ and otherwise never invests.

	High-Cost Game	**Low-Cost Game**	**Two-Cost Game**
$\underline{\Pi}$	0.250	1.750	1
A2 $\mathbf{\Pi^{TBR}}(\mathbf{s^{NE}})$	0.250	2.034	-
A2 $\mathbf{\Pi^{TBR}}(\mathbf{s^{a}})$	0.250	1.906	-
R2 $\mathbf{\Pi^{TBR}}(\mathbf{s^{NE}})$	-	-	1.073
R2 $\mathbf{\Pi^{TBR}}(\mathbf{s^{a}})$	-	-	1.084
A10 $\mathbf{\Pi^{TBR}}(\mathbf{s^{NE}})$	0.295	2.380	-
A10 $\mathbf{\Pi^{TBR}}(\mathbf{s^{a}})$	0.590	2.214	-
R10 $\mathbf{\Pi^{TBR}}(\mathbf{s^{NE}})$	-	-	1.497
R10 $\mathbf{\Pi^{TBR}}(\mathbf{s^{a}})$	-	-	1.344

Table 8: $\underline{\Pi}$, $\Pi^{TBR}(s^a)$, and $\Pi^{TBR}(s^{NE})$ (in ECU)

computed for each game in each treatment. Based on this, we can state:

Result 4

(i) $IE^T(s^a)$ *is much larger than* $IE^T(s^{NE})$ *in the ten-player high-cost game (0.34 vs. 0.045).*

(ii) $IE^T(s^a)$ *is smaller than* $IE^T(s^{NE})$ *in the two-player low-cost game (0.156 vs. 0.284), the ten-player low-cost game (0.464 vs. 0.63), and the ten-player two-cost game (0.344 vs. 0.497).*

(iii) In the remaining games, $IE^T(s^a)$ *and* $IE^T(s^{NE})$ *are very similar.*

Thus, informational flows in the actual experimental markets can be larger, smaller, or similar in size to those in NE. Note that, at any point in the game, a *ceteris paribus* increase in the probability of investment by some type with favorable common-value signal tends to increase informational flows because it strengthens the inference that anyone who invested has $X_i = 1$ and anyone who didn't invest has $X_i = 0$. Similarly, a *ceteris paribus* increase in the probability of investment for some type with unfavorable common-value signal tends to decrease informational flows because it weakens this inference. These effects tend to be largest in round 1.

Based on our findings about how behavior in the lab departs from NE, the following features of the game can be expected to systematically affect the comparison between informational externalities in the lab and in NE.

Possible presence of type with unfavorable common-value signal but with low enough investment cost so that expected profit of investing in round 1 is only slightly negative (type $(0, L)$ in our settings):

The presence of such a type has a negative effect on informational externalities relative to NE–such a type invests, unsurprisingly, with nonnegligible frequency in round 1 in the lab (see Result 2(i)) while it never invests in round 1 in NE.

Possible presence of type with favorable common-value signal but with high enough investment cost so that expected profit of investing in round 1 is only slightly negative:

Although there is no such type in our experiment, the presence of such a type can be expected to have a positive effect on informational flows relative to NE–such a type would invest, presumably, with nonnegligible frequency in round 1 in the lab while it would never invest in round 1 in NE.

Possible presence of "strong" type with favorable common-value signal and low cost in small markets (type $(1, L)$ in our two-player games):

When the market is small, the information that can be acquired by waiting is limited so that, in NE, a strong type would invest in round 1 with probability 1. On the other hand, in the lab, such a type would invest with a probability that falls considerably short of 1 (see Result 2(iii)). Thus, the presence of a strong type on a small market has a direct negative effect on informational externalities in the lab relative to NE.

The presence of a strong type also has an indirect positive effect on informational externalities relative to NE when there are types with favorable common-value signal but with a higher cost (type $(1, H)$ in our settings). This indirect effect occurs because, in NE, players with a favorable common-value signal but with a higher cost can free-ride on the information provided by round-1 investment by the strong type while, by subjects' insensitivity to opportunities to free-ride, such players in the lab do not do so and, by Result 2(ii), go ahead and invest with a substantial frequency. Thus, in games where there are players with favorable common-value signal but with a higher cost, the direct and indirect effects work in opposite directions so that the overall effect is ambiguous.

The indirect effect could also be negative if there are types with an unfavorable common-value signal but with a low enough investment cost so that the expected profit of investing in round 1 is positive (there is no such type in our experiment)–the latter type could free-ride on strong types in NE (thus improving informational flows) while it is unlikely to do so in actual play.

Possible but relatively unlikely presence of a strong type with favorable common-value signal and low cost (type $(1, L)$ in our two-cost games):

When the ex ante probability of the strong type is relatively small, players with this type cannot free-ride on each other in round 1 in NE.[33] As a result, this type invests with probability 1 in NE while, in the actual data, it invests with a probability that falls considerably short of 1 (see Result 2(iii)). This leads to a negative effect

[33]What is meant by a "relatively small" ex ante probability of the strong type depends largely on the market size. For example, in our two-cost ten-player game, an ex ante probability of 0.25 for type $(1, L)$ is sufficiently small so that $(1, L)$ players cannot free-ride on each other in round 1.

on informational externalities in the lab relative to NE. The possible indirect effects mentioned above also remain valid.

Market Size:

Fixing all types' round-1 behavior, the option value of waiting is higher on a larger market because players who wait observe a larger, and hence more informative, sample coming from the same distribution conditional on the investment return. In response to the changing incentives to wait in round 1, NE behavior adjusts accordingly. For example, in our one-cost games, the NE round-1 investment probabilities by the type with a favorable common-value signal decrease as the market size increases, partially offsetting the informational advantages of a larger sample. In the two-cost environment, the round-1 investment probability by type $(1, L)$ would also decrease in sufficiently large markets. In general, as Chamley and Gale (1994) and Levin and Peck (2008) show, the NE adjustments may be such that even as the market size goes to infinity, round-1 investment may not reveal very much about the investment return. In contrast, given subjects' insensitivity to market size, round-1 behavior in the lab does not adjust in a way that offsets the informational advantages of the larger sample. As a result, a larger market size has a positive effect on the informational externalities in the lab relative to NE.

In our ten-player high-cost game, there are no $(0, L)$ subjects and the market size is fairly large. As a result, markets in the lab are more informative than the Nash benchmark.

7.2 Market Efficiency

Informational externalities have to do with the profits of a shrewd investor. But this is not the whole story–unless play is in NE, some players are not best responding to the empirical frequencies of play and average actual profits will be below best-response profits. Thus, larger informational externalities do not necessarily imply increased market efficiency.

To shed light on the connection between informational externalities and market efficiency, we assume for simplicity that all players on a given market are using the same strategy s and we decompose the expected profit on this market, $\pi(s)$, into two components (plus a constant) both of which are related to informational externalities. In particular, we can write:

$$\pi(s) \equiv \underbrace{\Pi^{BR}(s) - \underline{\Pi}}_{IE(s)} - \underbrace{(\Pi^{BR}(s) - \pi(s))}_{\text{cost of mistakes}} + \underline{\Pi} \tag{2}$$

The first component is $IE(s)$ and, thus, captures the size of the informational externalities. The second component, the cost of mistakes, shows to what extent players are failing to exploit the informational externalities (as well as their private information).

Outside of NE, it is instructive to distinguish between two kinds mistakes. The first kind occurs either when players with favorable common-value signal invest despite a higher option value of waiting or when players with unfavorable common-value signal wait despite a higher expected payoff from investing. Although such mistakes increase the cost of mistakes, they also increase informational externalities. The efficiency-maximizing investment probabilities balance the two effects optimally. For example, in our ten-player high-cost game, the round-1 efficiency-maximizing investment probabilities for types $(0, H)$ and $(1, H)$ are 0 and 0.78, respectively. Any additional investment by type $(1, H)$ increases informational externalities but does not contribute to efficiency because too many players are not around after round 1 and hence cannot take advantage of the increased informational externalities.

All other mistakes unequivocally hurt efficiency both by diluting informational externalities and increasing the cost of mistakes. In our experiment, important instances of such mistakes occur when: (a) subjects with unfavorable common-value signal invest in round 1 despite a negative expected value of doing so; (b) subjects with favorable common-value signal fail to invest in round 1 even though the expected payoff from doing so exceeds the option value of waiting; and (c) in round 2, subjects display excessive conservatism in responding to observed round-1 investment.

To compare market efficiency in the lab with the NE benchmark, we compute TNE average profits *ex post* for each market trial, i.e., given the realization of the investment return, the common-value signals, and the costs. The ex post calculations allow us to eliminate any differences between actual and TNE profits that are due to noise in the realization of the investment return, the common-value signals, and the costs.

Note that, because the cost of mistakes is strictly positive outside of NE, the only way for actual markets to exceed NE efficiency is to generate an increase in informational externalities relative to NE that is larger than the cost of mistakes. Thus, our only hope of observing increased efficiency in the lab relative to NE is in the ten-player high-cost game. In this game, not only are informational externalities larger in the lab, but the probability of round-1 investment by type $(1, H)$ is closer in the lab (0.39) than in NE (0.07) to the efficiency-maximizing probability (0.78).

Table 9 shows, for each treatment, the average actual profits per subject as well as ex post TNE profits. All numbers are broken down by type and game. The table shows that, for each type in each game, average actual profits are lower than ex post

	(0,H)	(0,L)	(1,H)	(1,L)	Overall High-Cost Game	Overall Low-Cost Game	Overall Two-Cost Game
A2 TNE ex post	0.000	0.397	0.442	3.364	0.212	1.894	-
A2 actual, all rounds	-0.335	-0.124	0.392	3.156	0.014	1.530	-
R2 TNE ex post	0.000	0.324	0.489	3.936	-	-	1.166
R2 actual, all rounds	-0.612	0.003	0.178	3.733	-	-	0.801
A10 TNE ex post	0.109	1.557	0.555	4.056	0.333	2.926	-
A10 actual, all rounds	-0.482	0.730	0.467	3.973	-0.006	2.507	-
R10 TNE ex post	0.427	0.926	1.07	3.769	-	-	1.543
R10 actual, all rounds	-0.283	0.06	0.386	3.478	-	-	0.906

Table 9: Average actual profits vs. ex post TNE profits (in ECU)

TNE profits. We test the hypothesis that expected actual profits equal expected TNE profits. We do this via a two-tailed paired t-test in which each pair of observations consists of the actual observed aggregate level of profits in one market trial and the aggregate level of profits in the corresponding ex post TNE. The difference is significant in all two-player games (p-value is 0.008/0.000/0.000 in the high-cost/low-cost/two-cost game) as well as in the low-cost and two-cost ten-player games (p-value is 0.005/0.000 in the low-cost/two-cost game). However, it is not significant in the high-cost ten-player game (p-value is 0.1378). Thus, in this game, the increase in informational externalities relative to NE is nearly equal to the cost of mistakes, so that the difference does not attain statistical significance.

Let us summarize:

Result 5 *Average actual profits are lower than ex post TNE profits for all types and for all games. The difference is statistically significant in all games, except in the high-cost ten-player game.*

An important question is whether efficiency in actual markets may actually rise above efficiency in NE on very large markets. This is possible, given that a larger

	(0,H)	(0,L)	(1,H)	(1,L)	Overall High-Cost Game	Overall Low-Cost Game	Overall Two-Cost Game
A2 actual, all rounds	0.168	0.486	0.648	0.953	0.399	0.723	-
A2 NE ex post	0.000	0.401	0.624	1.000	0.300	0.703	-
R2 actual, all rounds	0.198	0.393	0.675	0.952	-	-	0.553
R2 NE ex post	0.000	0.200	0.290	1.000	-	-	0.367
A10 actual, rounds 1-3	0.155	0.521	0.618	0.966	0.387	0.765	-
A10 TNE ex post	0.051	0.384	0.324	0.941	0.188	0.689	-
R10 actual, rounds 1-3	0.203	0.394	0.616	0.882	-	-	0.525
R10 TNE ex post	0.194	0.273	0.443	1	-	-	0.478

Table 10: Investment Frequencies

market size enhances informational externalities on actual markets relative to the NE benchmark. However, if players are insensitive to the informational content of market activity, they may not be able to fully exploit this.

8 Appendix: Aggregate Investment

In this appendix we compare aggregate investment levels between markets in the lab and markets in NE. Note that since we have computed the NE actions only for rounds 1 and 2 in the ten-player games, we cannot compute exact NE investment. Instead we consider ex post TNE investment in each market.[34] In the two-player games, NE and TNE coincide. In the ten-player games, ex post TNE investment provides an upper bound on ex post NE investment in the first three rounds, since in the NE, players may prefer to wait beyond round 3 even if it is profitable to invest in round 3.

Table 10 shows the actual investment frequencies for all rounds in the two-player

[34]The ex post calculations allow us to eliminate any differences between actual and TNE investment which are due to noise in the realization of the gross return, the common-value signals, and the costs.

games and for rounds 1-3 in the ten-player games along with TNE investment frequencies.[35] All numbers are broken down by type and game. Actual investment exceeds ex post TNE investment for each game as well as for each type in each treatment, except for $(1, L)$ in the $A2$, $R2$, and $R10$ (where $(1, L)$ always invests in round 1 in NE so that overinvestment is impossible). Overinvestment relative to TNE is mainly driven by overinvestment in round 1 by $(0, L)$ players and (except in the two-player high-cost game) by $(1, H)$ players. In the two-player games, investment by $(0, H)$ players also contributes significantly towards overinvestment. Players' excessive conservatism in round 2 provides a countervailing force which, however, is insufficient to eliminate or reverse overinvestment. Overinvestment relative to TNE is especially pronounced in the $R2$ and in the high-cost game in the $A10$, largely due to overinvestment in round 1 by $(1, H)$ players.

For each game, we test the hypothesis that expected actual investment (in all rounds for the two-player games and in rounds 1-3 for the ten-player games) is different from expected TNE investment. We do this via a two-tailed paired t-test in which each pair of observations consists of the actual level of investment in one market trial and the level of investment in the corresponding ex post TNE for the same market trial. The differences are significant (at the 5-percent level) in all games, except in the low-cost game in the $A2$ and in the $R10$.

Let us summarize our findings regarding investment:

Result 6 *Actual investment (in all rounds for the two-player games and in rounds 1-3 for the ten-player games) exceeds ex post TNE investment for each game as well as for each type in each treatment (except for* $(1, L)$ *in the* $A2$*,* $R2$*, and* $R10$*). Over-*

[35]In the ten-player low-cost game, the actual investment frequency for all rounds equals 0.82; 61/85/93 percent of all investment occurs by round 1/2/3. In the ten-player high-cost game, the actual investment frequency for all rounds equals 0.44; 51/78/89 percent of all investment occurs by round 1/2/3. In the ten-player two-cost game, the actual investment frequency for all rounds equals 0.58; 54/78/90 percent of all investment occurs by round 1/2/3.

investment is especially pronounced (and the difference is statistically significant) in the $R2$ and in the high-cost game in the $A10$.

9 Appendix: Learning

In our main analysis, we have for simplicity been ignoring the possibility of learning as sessions progress. In this section, we explicitly consider learning. First, we focus on changes in round-1 behavior across time. In order to do that, we run, for each type in each treatment, the following random effects probit with period, t, as a right-hand side variable:

$$Pr(\text{subject invests in round } 1|t, v) = \Phi(\beta_0 + \beta_1 t + v)$$

Table 11 shows, for each type in each treatment, the p-value for the hypothesis that $\beta_1 = 0$ (i.e. that there is no effect of the period number on round-1 behavior), the predicted probability of investment in rounds 1 and 24, as well as the NE probability of investment. From the table, we see that there is a reduction in the probability of round-1 investment for type $(1, H)$ in all treatments. This reduction is substantial in size and statistically significant in the $A2$, $R2$ and $R10$. Note that the observed shift in the behavior of $(1, H)$ players over time is towards the NE in the $R2$, $A10$, and $R10$, but away from the NE in the $A2$. We also observe a substantial and statistically significant reduction in the round-1 probability of investment for $(0, H)$ and $(0, L)$ players in the $A10$.

Importantly, note that in period 24 we do not observe any dilution of the behavioral biases that were identified earlier based on aggregate round-1 behavior from periods 1-24. Regarding insensitivity to market size: for a given cost-structure, the predicted probability of round-1 investment in period 24 does not differ between the two-player game and the ten-player game in a statistically significant way (at the 5-percent level) for any type. Regarding insensitivity to opportunities to free-ride: for a given market size, the predicted probability of round-1 investment in period 24 does not differ between the one-cost game and the two-cost game in a statistically significant way (at the 5-percent level) for any type; the only exception is for type $(0, H)$ in the ten-player high-cost and ten-player two-cost games.[36]

We also investigate whether there is any learning in round 2. In particular, we run the probit model:

$$\begin{aligned} &Pr(\text{subject invests in round } 2|k_1, t, v, \text{subject has not invested in round } 1) = \\ &\qquad = \Phi(\beta_0 + \beta_1 k_1 + \beta_2 t + \beta_3 k_1 t + v) \end{aligned}$$

Our estimate of β_2 is significant at the 5-percent level only for the $(0, H)$ type in the $A2$ and for the $(1, L)$ type in the $R10$. Our estimate of β_3 is significant at the 5-percent level only for the $(1, L)$ type in the $R10$. Given that we have 16 estimates of β_2 and 16 estimates of β_3 (one for each type in each treatment), it is not unexpected for a few of the estimates to be significant. Thus, we find little evidence that learning eliminates the behavioral biases that were identified earlier based on aggregate round-2 behavior from periods 1-24.

[36] Here is how we test for statistically significant differences in the predicted probability of round-1 investment in period 24 for a given type between two games. First, we estimate the random effects probit model $Pr(\text{subject invests in round } 1|t, d, v) = \Phi(\beta_0 + \beta_1 t + \beta_3 d + \beta_4 dt + v)$, where d is a dummy that equals 1 in one of the games and 0 in the other game. This model allows for a separate constant and slope coefficient in each game. Then, we test the hypothesis $\beta_3 + \beta_4 24 = 0$.

A2

	(0,H)	(0,L)	(1,H)	(1,L)
p-value for H0: $\beta_1 = 0$	0.253	0.915	0.003	0.760
Predicted prob. of investment in round 1 for $t = 1$	0.022	0.058	0.473	0.860
Predicted prob. of investment in round 1 for $t = 24$	0.003	0.055	0.234	0.843
NE prob. of investment in round 1	0	0	0.492	1

R2

	(0,H)	(0,L)	(1,H)	(1,L)
p-value for H0: $\beta_1 = 0$	0.147	0.440	0.000	0.229
Predicted prob. of investment in round 1 for $t = 1$	0.049	0.023	0.475	0.942
Predicted prob. of investment in round 1 for $t = 24$	0.009	0.009	0.145	0.882
NE prob. of investment in round 1	0	0	0	1

A10

	(0,H)	(0,L)	(1,H)	(1,L)
p-value for H0: $\beta_1 = 0$	0.037	0.017	0.27	0.167
Predicted prob. of investment in round 1 for $t = 1$	0.153	0.213	0.442	0.825
Predicted prob. of investment in round 1 for $t = 24$	0.005	0.072	0.338	0.738
NE prob. of investment in round 1	0	0	0.071	0.746

R10

	(0,H)	(0,L)	(1,H)	(1,L)
p-value for H0: $\beta_1 = 0$	0.653	0.949	0.006	0.562
Predicted prob. of investment in round 1 for $t = 1$	0.054	0.154	0.441	0.671
Predicted prob. of investment in round 1 for $t = 24$	0.074	0.15	0.228	0.714
NE prob. of investment in round 1	0	0	0	1

Table 11: Round-1 Investment and Learning.

10 Appendix: Instructions for R10

This is an experiment on decision-making in investment markets. The National Science Foundation has provided funds for conducting this research. The instructions are simple, and if you follow them carefully and make good decisions, you may earn a CONSIDERABLE AMOUNT OF MONEY, which will be PAID TO YOU IN CASH at the end of the experiment.

Every participant in the experiment is guaranteed a payment of at least $5, independent of their performance in the experiment. All monetary values in the experiment, such as investment costs, investment returns, and account balances, are written in experimental currency units (EC). Your balance of ECs at the end of the experiment will be converted to US dollars at the exchange rate of $0.50 for each EC. Because your decisions may involve losses, we will endow you with a starting cash balance of 30 ECs. Your gains (losses) during the experiment will be added to (subtracted from) your cash balance. However, if your cash balance falls below zero, you will no longer be allowed to continue. At the end of the experiment you will receive in cash your end of experiment balance of ECs converted to US dollars, or $5, whichever is greater.

1. In this experiment we will create a sequence of *market trials*. In each given *market trial*, the participants will act as potential *investors*. Each potential investor will have to decide *whether*, and *when*, s/he wishes to invest, based on the information s/he is provided (and which we will explain later).

2. In the experimental *session* today we will have between 20-25 market trials. Each market trial has several *rounds*. The initial round is round 1, the next is round 2, and so on. In each round you and the other potential investors in your market trial will have to decide (simultaneously) whether to invest in that round or not. The decision to invest is irreversible. Any potential investor who has not yet invested will be told how many of the other potential investors have invested during each previous round of that trial.

3. In each trial, the market in which you are a potential investor has several more potential investors besides yourself. In a typical session we will recruit (about) 24 students. The computer will randomly select 2 groups of 10 students, with the remaining students selected to sit out that trial. Each such group of 10 constitutes a separate market trial that has no relation to the other market. A given market trial keeps the same matched students over the several rounds of that

market trial. However, after the market trial is over, the computer randomly rematches students to form two new market trials. This matching procedure makes it very unlikely that you will be matched with the same group of students from one trial to the next.

4. The structure of information.

Information about investment cost: Each potential investor will know, before each market trial starts, her/his investment cost for that trial. There are two possible levels of investment cost: low cost, C_L=3.5 and high cost, C_H =6.5. Each potential investor will be assigned one of the two cost levels with equal probability (1/2). In other words, in your market trial, you will know your investment cost, and that the investment cost of each other potential investor is equally likely to be either 3.5 or 6.5.

Information about investment *gross* returns: The computer assigns a *gross* return to every market trial. The *gross* return remains the same for all rounds of the same market trial, and is completely uncorrelated with your investment cost. The computer randomly picks the *gross* return to be either 10 or 0, with equal probabilities. Once the gross return is picked, high or low, it is the same for all potential investors in your trial, and it remains the same for all rounds of that trial. You will NOT observe whether the gross return for your trial is high or low. Instead, each potential investor will be given her/his own signal, which takes the value of either 0 or 1. Signals are 70% accurate, in the following sense:

If the gross return is 10, you have a 70% chance of observing signal 1 and a 30% chance of observing signal 0. If the gross return is 0, you have a 70% chance of observing signal 0 and a 30% chance of observing signal 1.

Each potential investor’s signal is related to the gross return, but the computer randomizes separately for each potential investor, so different potential investors can receive different signals. For example, suppose there are 10 potential investors in a trial. If the gross return is 10, then on average there are 7 people who receive signal 1 and 3 people who receive signal 0. If the gross return is 0, then on average there are 7 people who receive signal 0 and 3 people who receive signal 1. However, the actual numbers can vary.

The signal for each potential investor is chosen at the beginning of the trial and remains the same for all rounds of that trial. Each potential investor observes her/his own signal, but not the signal of the other potential investors in that trial.

Observing your signal may help you better predict the likelihood that the gross return in your market trial is high or low.

Information about other investors in your market: You will NOT be told the signals of the other potential investors in your market trial. However, you will be informed about how many other potential investors have already invested, and during which rounds. If this information reveals something about others' signals, it could improve your decision about if and when to invest.

You are not allowed to reveal or discuss your information with other students or look at another student's screen (this will be strictly monitored and violators will be removed from the experiment).

5. **The structure of the game**.

Once you are randomly assigned to a market trial, you privately observe your cost and your signal, which remain constant for that market trial. The other potential investors observe their cost and signal. In round 1, you are asked to decide if you wish to invest. If you do not invest in round 1, you are informed about how many other potential investors invested in round 1, and you are asked if you wish to invest in round 2. If you have not invested by round 2, we move to round 3, and so on. Once you have decided to invest, there are no more decisions to make in that market trial. That is, an investment decision in a given trial is irreversible. You cannot disinvest or invest a second time. After two consecutive rounds in which no one in your trial invests, that trial is over.

In order to make good decisions, you must understand how your gains and losses are determined. This will be carefully explained below.

Once a market trial is over, the whole process starts again. The computer randomly selects two groups of 10 to form a new set of market trials, those selected will be assigned an investment cost and a signal, etc.

If we start with more than 20 students, then some of you will be sitting out from time to time. For example, with 24 students, 4 students are randomly selected to sit out each trial, but since the computer performs a new randomization each time, it is very unlikely that anyone will sit out very often.

If a student's cash balance falls below zero, then that student must stop playing, so one fewer student will sit out each trial. If the number of remaining students falls below 20, then everyone will be assigned to a market trial and you will be told how many potential investors are in your trial.

Your screen will inform you of the trial number, and the round number within the trial.

How your gains (discounted net returns) or losses are determined.

If you invest, your gains from that trial are the discounted difference between the gross return and your investment cost. Let us illustrate what this means by using a simple example. Suppose that in the current market trial your investment cost is 3.5. If you decide to invest in round 1, then your gains are: 6.5 if the gross return is 10 (10 – 3.5 = 6.5) or – 3.5, a loss of 3.5, if gross return is 0 (0 – 3.5 = -3.5). Note that gains or losses in round 1 are not discounted; they are just the difference between the market gross return and your investment cost. For each round that you wait, your gains or losses are discounted by a factor of 0.9, as shown in the following table.

Discounted Net Returns when Cost is 3.5

Round that you invest	If return is 10 (high)	If return is 0 (low)
1	6.5	- 3.5
2	5.85	- 3.15
3	5.26	- 2.84

There are several important things to note here:

(i) If for whatever reason you have decided not to invest at all in a particular market trial, you will earn zero for that market trial.

(ii) You will not be told the actual gross return during a market trial. After each trial is over, the gross return is revealed and you will learn your discounted net gains or losses, which will be added to, or subtracted from, your cash balances.

(iii) It is up to you to decide if and when to invest. Clearly, your investment cost and your signal can affect your decision. Observing the activity of the other potential investors in your trial might indirectly yield useful information about the gross return, by telling you something about the other potential investors' signals.

6. **Information on the computer screen.** Throughout the experimental session, the computer screen will show your ID number and current cash balances, in the upper left corner. The upper left corner of the screen will also remind you of the number of potential investors in each trial (usually 10), the discount factor (0.9), and the "accuracy parameter" of your signal (70%).

At the beginning of each round of each market trial, you will see the number of the market trial, your cost of investment (either 3.5 or 6.5), and your signal (0 or 1). This information stays the same during the trial. In the middle of the screen, you will see the current round number. At the bottom of the screen, you will see a "history" of investment in previous rounds of that trial. For example, if the history lists 2 investors in round 1 and 3 investors in round 2, then the total number of investors during the first two rounds is 5.

In the shaded area at the very bottom, you will also see your personal statistics from your previous trials. (If you are listed as investing in round -1, this means that you never invested during that trial. A row of all -1s means that you sat out during that trial.)

You will have 10 seconds to think about whether to invest in that round. At that time, boxes marked "YES" and "NO" will appear, and you should mark a box to indicate whether you want to invest or not. Please make your choice within 5 seconds.

At the end of the market trial, you will see a screen that tells you the market trial number, your investment cost, your signal, the actual gross return, and your net discounted gains or losses from that trial. You will also see your personal statistics from your previous trials.

7. We will start the session with two practice "dry runs" that do not count towards your earnings, at which point we will stop and answer additional questions. At the end of the experiment, while we are calculating your earnings, we ask that you answer the short questionnaire on your computer.

8. Are there any questions?

11 Appendix: Screen Printout

Printout of the screen seen by a subject in the A2 with cost 6.5 and signal 1, who is deciding whether to invest in round 2 after the other subject has invested in round 1:

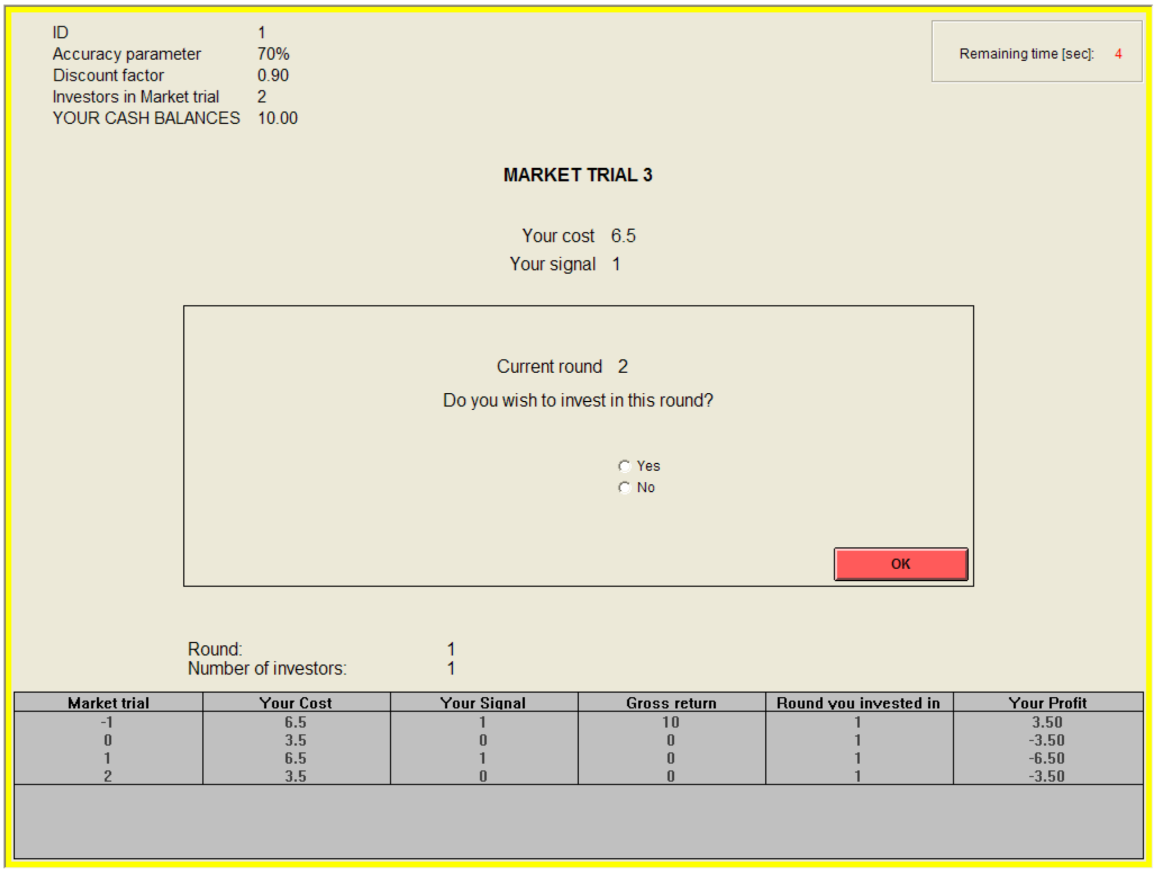

The screen in the R2, A10, and R10 looks the same (except that in the ten-player games we would have "Investors in Market trial 10").