Animal Spirits and Informational Externalities in an Endogenous-Timing Investment Game: an Experimental Study*

Asen Ivanov
Virginia Commonwealth University

Dan Levin
The Ohio State University

James Peck
The Ohio State University

October 5, 2010

Abstract

We experimentally test an endogenous-timing herding model in which subjects observe their cost of investing and a signal correlated with the common investment return. Investment exceeds the Nash benchmark in all of our treatments, reminiscent of Keynes' notion of "animal spirits." The effect of overinvestment on best-response profits due to the informational externality can be positive or negative, depending on the environment. Subjects are highly heterogeneous in their propensities to invest. Initial overinvestment is followed by underinvestment with the tendency towards overinvestment dominating.

*This work is supported by the NSF under Grant No. SES-0417352. Any opinions, findings and conclusions or recommendations expressed are those of the authors and do not necessarily reflect the views of the NSF. We thank the John Glenn Institute at Ohio State for their support.
aivanov@vcu.edu, levin.36@osu.edu, peck.33@osu.edu.

"Most, probably, of our decisions to do something positive, the full consequences of which will be drawn out over many days to come, can only be taken as a result of animal spirits–of a spontaneous urge to action rather than inaction, and not as the outcome of a weighted average of quantitative benefits multiplied by quantitative probabilities." J. M. Keynes (The General Theory, p. 161)

1 Introduction

The quotation displayed above offers the viewpoint that, when investors face a decision for which they cannot accurately formulate outcomes and probabilities, then investors may have an urge or bias in favor of action. Moreover, Keynes thought that these "animal spirits" can benefit the economy. For example, if a coordination failure puts the economy in a recession due to low aggregate demand, animal spirits can increase investment, thus increasing aggregate demand and pulling the economy out of the recession. In this paper, we consider a different channel through which animal spirits could affect the economy.

We experimentally implement a series of endogenous-timing herding games, along the lines of Chamley and Gale (1994) and Levin and Peck (2008). In this environment, there are no payoff externalities, so investment by others does not affect one's payoff directly. However, there can be informational externalities–players who wait may infer something about investment returns by observing investment intensity in previous rounds. Thus, an urge to invest due to animal spirits will alter this informational externality, in either a positive or negative direction, depending on which types of subjects are induced to invest (i.e., types with favorable or unfavorable information about investment returns).[1] In dynamic environments, it is also possible that an initial investment surge due to animal spirits could be stifled by an insufficient response due to "passive spirits" in later rounds.

In the current paper, we address experimentally the following questions:

[1] See also Peck and Yang (forthcoming), who briefly discuss the implications of animal spirits for creating beneficial informational externalities.

1. Is there overinvestment relative to the Nash equilibrium (NE)? If overinvestment is pervasive, we will label it as animal-spirits-driven.

2. How do the informational externalities generated in the lab compare to those predicted by the NE?

3. What behavioral biases, and what features of the environment, affect this comparison?

4. Is there subject heterogeneity in terms of the generation of informational externalities?

We also compare the average profits in the lab to the NE benchmark, and address the question of how features of the environment affect this comparison.

In our experiment, each subject receives either a favorable (1) or unfavorable (0) private signal about the common investment return (we call this the common-value signal), and also observes a private cost of investing. Our games differ according to the market size (two vs. ten players per market) and according to the cost-structure (one-cost vs. two-cost games). In the one-cost games, the investment cost is the same for all subjects–it is high (H) in the high-cost game and low (L) in the low-cost game. In the two-cost games, the cost is either H or L and can differ across subjects. A subject's type consists of her common-value signal and her investment cost. Thus, possible types are $(0, H)$, $(0, L)$, $(1, H)$, and $(1, L)$.

We find round-1 overinvestment in all games. In this sense, we confirm the presence of animal spirits. However, we also find that subjects exhibit passive spirits in round 2, in the sense that they respond more conservatively to observed round-1 investment than a player in NE or a player best responding to the empirical frequencies. This is in line with the notion that animal spirits are not universally present and, depending on the situation, may even give way to passive spirits. For our games, the interplay between animal spirits and passive spirits leads to higher overall investment than the NE benchmark–thus, animal spirits dominate.

The comparison between informational externalities in the lab and in the NE benchmark differs across environments.[2] Subjects respond to their type in a sensible

[2]To make this comparison, we will develop a measure of the aggregate size of the informational externalities.

way, but exhibit behavioral biases that affect informational externalities (as compared to NE) in systematic ways. One notable bias is that subjects fail to take into account the possibility to defer profitable investment in round 1 in order to free-ride on the information provided by round-1 investment by stronger types. Another notable bias is that, when deciding in round 1 whether to invest or wait, subjects fail to realize that the market size affects the informativeness of market activity that can be observed if one waits. As a result, lab behavior in round 1 is insensitive, in contrast to NE predictions, to two subtle features of the game that affect the option value of waiting– the cost structure of the game (which determines whether some types can free-ride on other, stronger, types) and the market size. It is in large markets that animal spirits and subjects' failure to appreciate the effects of market size on round-1 investment incentives can lead to the greatest informational externalities relative to NE.

Animal spirits sometimes provides stronger informational externalities than the outcome of Nash play, so one can imagine overall economic efficiency rising as a result. However, in all of our treatments, average profits are lower than NE profits. The most costly mistake subjects are making is to invest in round 1 despite a negative expected value of doing so. Another systematic mistake is that subjects who wait until round 2 tend to be overly conservative at that point.

Turning to individual-level behavior, we find that there is significant heterogeneity across subjects in terms of the generation of informational externalities. Certain individuals seem to have an urge to invest in round 1–this is in line with the notion that animal spirits are individual-specific, rather than a universal property of human nature. Moreover, the proportion of individuals with animal spirits who generate a large informational externality in round 1 is large relative to the amount of investment we would see in round 1 in the NE of large markets.

The layout of the paper is as follows. Section 2 reviews the relevant literature. 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 some concluding remarks.

2 Literature Review

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 that lead to inefficiencies. Anderson and Holt (1997) provide the first experimental tests of herding models with exogenous timing and demonstrate the frequent occurrence of information cascades.

The theoretical literature on herding with endogenous timing is pioneered by Chamley and Gale (1994), who assume that all players have the same investment cost. Each of our one-cost games is a special case of their model. Levin and Peck (2008) consider the more general case that allows for cost heterogeneity.

The first experimental work on endogenous timing herding models is done by Sgroi (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. The NE prediction is that all subjects eventually guess, so it is difficult to address questions of overinvestment and animal spirits.

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. While this strategy is not consistent with NE, it may well be a best response to the empirical distribution of strategies.[3] The choice of a subject with signal 3 or -3 to guess in round 1 can also be viewed as the result of an urge to action over inaction, in line with our theme of animal spirits.

The experimental work closest to the current paper is Ivanov, Levin, and Peck

[3] Ziegelmeyer et al (2005) do not report the empirical distribution in a precise enough way for us to perform the computation, and the decision is very close.

(2009). That paper studies behavior in the one-cost and two-cost games described above with two players per market. Unlike the current paper, Ivanov, Levin, and Peck (2009) does not focus on the issue of market-level informational externalities. This issue is more interesting in large markets. Instead the authors focus on individual-level behavioral issues regarding how people gather information from, or anticipate gathering information from, others' behavior over time. The study considers prominent behavioral theories–the level-k model (see Nagel (1995), Stahl and Wilson (1994, 1995), Crawford and Iriberri (2005)), cursed equilibrium (see Eyster and Rabin (2005)), and Quantal Response Equilibrium (see McKelvey and Palfrey (1995, 1998))–as possible explanations of behavior. The authors conclude that a new rule-of-thumb-based explanation fits the data better: rather than best-responding (possibly with noise) to (possibly incorrect) beliefs as in the aforementioned behavioral theories, subjects seem to be following rules of thumb based on different insights about the game.

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.[4]

There is also a large literature on sunspots and animal spirits, emanating from Keynes' General Theory. The theoretical sunspots literature includes Shell (1977), Cass and Shell (1983), Azariadis (1981), Peck (1988), and many others. In a sunspot equilibrium, events unconnected to the economic fundamentals affect economic outcomes like prices and allocations, by affecting the equilibrium beliefs of agents. In many cases, there is a continuum of Pareto ranked equilibria, indexed by the confidence agents have about (say) the return from investing rather than consuming; hence, the connection to Keynes' phrase about animal spirits. Howitt and McAfee (1992) draw a distinction between sunspots, which can affect beliefs about inflation, and animal spirits, which can affect labor markets and aggregate demand, due to a coordination externality. Shleifer (1986) explores the role of animal spirits in coordinating expectations and causing cyclical fluctuations. See also the references in Benhabib and Farmer (1999) for papers that perform simulations using sunspots or animal spirits models that perform at least as well as more traditional real business cycle models.

[4]See Cipriani and Guarino (2005) and Drehmann et al (2005) for experimental studies.

In the sunspots and previous animal spirits literatures, agents are fully rational, in effect computing "a weighted average of quantitative benefits multiplied by quantitative probabilities." When we use the phrase animal spirits to describe some of our subjects, we refer to subjects who cannot compute the expected profit from investing and waiting, but who have a gut feeling that investing is the right thing to do.[5] We capture an aspect of what Keynes was talking about that is missing from this literature. Our subjects face investment decisions that are complex and probably outside their experience, and it is nearly certain that they cannot perform a mathematical computation to determine their actions. This is a good feature of our experiment, because many of the most important real world investment choices are also outside the decision makers' experiences. "If human nature felt no temptation to take a chance, no satisfaction (profit apart) in constructing a factory, a railway, a mine or a farm, there might not be much investment merely as a result of cold calculation," Keynes wrote (The General Theory, p. 150). For a view along similar lines, see also Akerlof and Shiller (2009, Introduction and Ch. 1).

3 Theoretical Framework

Our theoretical framework is based on the general model in Levin and Peck (2008). There are n risk-neutral[6] players or potential investors. Let $Z \in \{0, 10\}$ denote the true investment return, common to all investors, with $Pr(Z = 0) = Pr(Z = 10) = \frac{1}{2}$. Each player $i = 1, ..., n$ 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. The accuracy of the signal is given by the parameter

[5]Two interesting experimental studies try to find sunspot fluctuations. Marimon, Spear, and Sunder (1993) induce cyclical movements with a cyclical pattern of generation size in an overlapping generations model, then remove the generation-size fluctuations. In some sessions, cyclical movements in economic activity continues. Duffy and Fisher (2005) provide evidence that sunspots can matter, but only when there is a common language for interpreting the sunspot signal. Thus, subjects can coordinate on an announcement that prices will go up, but not on a flashing red or yellow box on the computer screen. Our experiment has little connection to these experiments. We are not interested in generating fluctuations or inducing subjects to respond to particular extrinsic random variables. Any animal spirits are brought into the lab by the subjects themselves.

[6]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 5.7.

$\alpha \in [\frac{1}{2}, 1]$:

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

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, X_i.

Each player i also privately observes a second signal, representing the idiosyncratic cost of undertaking the investment, c_i. 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.

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 an investment probability to each history and type.

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 the NE behavior for rounds 1 and 2.[7] 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, in NE the option value of waiting equals the value of waiting and investing iff the expected payoff of doing so is positive in the next round (and otherwise never investing). For the case of $n = 2$, the explicit derivation of the NE can be found in Ivanov, Levin, and Peck (2009). 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 Figures 2 (for $n = 2$) and 4 (for $n = 10$).[8] Note that type $(1, H)$ chooses to wait in round 1 despite a positive expected profit of investing.

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 panels of Figures 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 Figures 1 (for $n = 2$) and 3 (for $n = 10$) (only the columns corresponding to types $(0, H)$ and $(1, H)$ apply). Note that type $(1, H)$ chooses to invest with probability less than one in round 1 despite a positive expected profit of investing. Moreover, because of the large number of players and the corresponding potential for large informational externalities when $n = 10$, the required probability of round-1 investment that induces indifference between investing and waiting for the $(1, H)$ type

[7] 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 Ivanov et al. (2009) for details).

[8] 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.

in the NE is very low–only 0.07.

Notice that the round-1 probability of investment for the $(1, H)$ (and to a lesser extent for the $(1, L)$) type differs a lot depending on the number of players and cost structure (one-cost vs. two-cost) of the game. This probability is at or near zero in the two-cost games and the ten-player high-cost game where $(1, H)$ players can free-ride on the information provided by round-1 investment activity by $(1, L)$ players and/or the market is large so that the investment activity one can observe if one waits in round 1 is very informative. On the other hand, this probability is 0.49 in the two-player high-cost game where neither of these conditions holds.

Also, note that, in the NE of all games, 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. 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.[9] [10]

4 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). Studying overinvestment and informational externalities is most interesting in large markets, but we include the two-player data to demonstrate that behavior does not depend strongly on the market size.[11]

[9] $(0, H)$ 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.

[10] 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 iff it is profitable and otherwise never invest.

[11] The data from the two-player treatments is analyzed in Ivanov, Levin, and Peck (2009). They also had another treatment, Treatment 3, that was used to distinguish between different explanations

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.[12] 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.[13] [14]

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.[15]

Average earnings for the R2/A2/R10/A10 were $26.04/$26.49/$25.71/$26.32. Including the reading of instructions, 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

of behavior. We do not consider Treatment 3 here because, from the point of view of informational externalities, it adds little to the analysis.

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

[13]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.

[14]For 2 sessions of the R2 and of the A2, the investment game was followed by some lottery problems and a questionnaire. Ivanov, Levin, and Peck (2009) used these to test their interpretation of behavior. We do not use the data from the lottery problems and the questionnaire in our current study.

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

cash began, no more questions were allowed. See the Appendix for our Instructions from the R10[16] 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.

5 Results

We will focus on actual behavior only in rounds 1 and 2. The right panel in Figure 1 summarizes actual behavior in rounds 1 and 2 in the $A2$. The top row shows, for each 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; again, the subscripts in parentheses show 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 Figures 2, 3, and 4 analogously summarize actual behavior in rounds 1 and 2 in the $R2$, $A10$, and $R10$, respectively.

5.1 Basic Rationality

Let us start by checking if subjects' aggregate behavior satisfies some basic requirements on rationality. First note that, in the aggregate, subjects in all treatments respond to their investment cost and common-value signal. This can be seen when we consider investment in round 1, and when we consider investment in any round.

The higher the expected profit from investment given a subject's type, the more likely she is to invest in round 1. In particular, in all treatments the frequencies with which subjects of type $(0, H)$, $(0, L)$, $(1, H)$, and $(1, L)$ invest in round 1 are strongly increasing in the given order. In addition, the higher the expected payoff

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

from investment given a subject's type, the more likely she is to invest eventually (i.e. during some round). In particular, the frequencies with which subjects of type $(0, H)/(0, L)/(1, H)/(1, L)$ invest are 17/49/65/95 percent in the $A2$, 20/39/68/95 percent in the $R2$, 21/63/66/98 percent in the $A10$, and 28/48/67/91 in the $R10$.

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:

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

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 1 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 are more likely to invest eventually, 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.*

5.2 Actual Behavior vs. NE Behavior For Each Type

Let us compare the actual and NE behavior of each type in rounds 1 and 2. The differences in behavior we observe will allow us to explain any differences between actual and NE aggregate investment levels, informational externalities, and profits.

First, let us consider round-1 behavior. The NE investment probability for types $(0, H)$ and $(0, L)$ is 0. In the actual data, there is almost no investment by $(0, H)$ players (5-7 percent) and small but nonnegligible investment by $(0, L)$ players (11-17 percent). The NE investment probability for type $(1, H)$ changes dramatically across treatments–it is either zero or near zero in the $R2$, $A10$, and $R10$, or 0.49 in the $A2$. However, actual behavior across treatments is remarkably stable–the probability of round-1 investment for this type is 0.35-0.39. The NE probability of investment for the $(1, L)$ type is 0.75 in the $A10$ and 1 in the remaining treatments. The actual probability of investment for this type is stable across treatments–it is in the range 0.68-0.79.

Let us now compare NE and actual round-2 behavior for each type in each treatment. First, consider the $A2$ and $R2$. Observe that, whenever the NE prescribes for a given type to invest with probability 0 after a given value of k_1, the frequency of actual investment tends to be relatively small–it does not exceed 0.256. Notice also that, $(0, L)$ and $(1, H)$ types do not respond to round-1 investment by the other player as much as in the NE–the frequency of investment for these types after history (1) does not exceed 0.68. Thus, for types $(0, L)$ and $(1, H)$, round-2 behavior in the lab is more conservative than round-2 behavior in NE, in the sense that overinvestment after history (0) is less than underinvestment after history (1). In addition, because in round 2 best-response behavior given the empirical frequencies of play and NE behavior coincide, $(0, L)$ and $(1, H)$ types are also being more conservative than someone who best responds to actual play. Type $(1, L)$ subjects are also being overly conservative in round 2 given that, both at history (0) and history (1), they invest with frequency considerably less than 1 even though investing dominates waiting.

In the $A10$ and $R10$, we observe for all types a similar tendency for being overly conservative in round 2. Casual inspection of Figures 3 and 4 suggests that, in order to invest in round 2, subjects tend to require a higher value of k_1 than in NE. To formalize the latter observation, 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.[17] Figure 5 reports, for each type in the $A10$ and $R10$, $\overline{k}_1$ along with the estimated standard error. We also extrapolate the corresponding NE cutoffs, $\overline{k}_1^{NE}$, and best-response cutoffs, $\overline{k}_1^{BR}$.[18] For types $(0, L)$ and $(1, H)$ in the $A10$ and for all types in the $R10$, the actual cutoff is significantly higher (at the 5-percent level) than

[17] $\overline{k}_1$ may not be an integer. This does not seem to be an important issue and we ignore it.

[18] To compute $\overline{k}_1^{NE}$, denote by $\widetilde{k}_1$ the highest level of round-1 investment for which the NE probability of investment is less than 0.5. Let p_1 and p_2 be the NE probabilities of investment after histories $(\widetilde{k}_1)$ and $(\widetilde{k}_1 + 1)$, respectively. Then $\overline{k}_1^{NE} = \widetilde{k}_1 + \frac{(0.5 - p_1)}{p_2 - p_1}$. For $\overline{k}_1^{BR}$, it is the average of the highest k_1 for which not investing is a best response and the lowest k_1 for which investing is a best response given the empirical frequencies of play.

both the NE cutoff and best-response cutoff.

Thus, in all treatments there is a tendency for overly conservative behavior in round 2 (more on this in section 5.6). Interestingly, we find that the overall level of actual investment is higher than the NE prediction (see section 5.3), in spite of the conservative behavior in round 2.

Note also that, for each type, round-2 behavior seems similar between the $A2$ and $R2$ (see Figures 1 and 2) as well as between the $A10$ and $R10$ (see Figures 3 and 4; also compare $\overline{k}_1$ between both panels in Figure 5). More formally, we test, for each type, the hypothesis that the estimated coefficients in equation (1) are equal between the $A2$ and $R2$ as well as between the $A10$ and $R10$. We cannot reject this hypothesis (even at the 10-percent level) except for type $(1, H)$ in the $A10$ and $R10$ (p-value is 0.019). This is in spite of the fact that, for each type, there are differences in the NE round-2 behavior between the $A10$ and $R10$–for example, there are differences in $\overline{k}_1^{NE}$ for each type.[19]

Let us summarize our main findings regarding how behavior in rounds 1 and 2 compares to NE behavior.

Result 2

(i) The NE round-1 investment probability for types $(0, H)$ and $(0, L)$ is 0. In the actual data, 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 NE round-1 investment probability for type $(1, H)$ is zero or near zero in the $R2$, $A10$, and $R10$ and 0.49 in the $A2$. In the actual data, the frequency of round-1 investment for this type is 0.35-0.39.

(iii) The NE probability of investment for the $(1, L)$ type is 0.75 in the $A10$ and 1 in the remaining treatments. In the actual data, the frequency of round-1 investment for this type is 0.68-0.79.

(iv) In all treatments, there is a tendency for overly conservative behavior in round

[19]As mentioned, for type $(1, H)$ we find significant differences in behavior in the lab between the $A10$ and $R10$. However, note that, for this type, $\overline{k}_1$ is greater in $A10$ than in $R10$ (5.54 vs. 4.58, respectively) while $\overline{k}_1^{NE}$ is smaller in $A10$ than in $R10$ (1.26 vs. 2.76, respectively).

2 (relative to NE behavior and to best-response behavior given the empirical frequencies of play).

We can also state:

Result 3 *For each type:*

(i) Round-1 behavior is similar across treatments, even when cost structure and market size differences cause the NE investment probabilities to be very different (most notably in the case of type $(1, H)$*);*

(ii) Round-2 behavior is similar between the $A2$ *and* $R2$ *as well as between the* $A10$ *and* $R10$.

Thus, by result 1, subjects respond to factors that directly affect the expected value of investing, such as their private information or observed investment activity from round 1. However, by result 3, behavior in the lab is insensitive to two subtle considerations that affect the option value of waiting, namely the cost structure (which determines whether $(1, H)$ players can free-ride on the information provided by $(1, L)$ players) and the market size (which affects the strength of any statistical inference based on observed market activity).

A useful way of summarizing round-1 behavior in the lab, which highlights the insights from results 1, 2, and 3, is given in the following:

Result 4 *Regardless of cost structure and market size, subjects invest in round 1 with a small (5-7 percent)/nonnegligible (11-17 percent)/substantial (35-39 percent)/large-but-considerably-less-than-one (68-79 percent) frequency when type* $(0, H)/(0, L)/(1, H)/(1, L)$, *i.e. when the expected profit of investing based on one's type is strongly negative/slighly negative/slightly positive/strongly positive.*

5.3 Aggregate Investment

We have established that there is overinvestment in round 1, which we identify with animal spirits. However, it is possible that animal spirits induce overinvestment (relative to NE) in round 1, but passive spirits in subsequent periods stifle initial overinvestment. Let us now evaluate the extent to which the overall effect is animal-spirits-driven overinvestment over the course of the games, relative to 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 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. In the ten-player games, TNE investment provides an upper bound on 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.

We compute TNE investment *ex post* for each market trial, i.e., given the realization of the gross return, the common-value signals, and the costs. 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.

Table 2 shows the actual investment frequencies for all rounds in the two-player games and for rounds 1-3 in the ten-player games along with TNE investment frequencies. 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 $(1, L)$ should invest in round 1 and 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. 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 5 *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-investment is especially pronounced (and the difference is statistically significant) in the* $R2$ *and in the high-cost game in the* $A10$*.*

Thus, we observe animal spirits in the sense that there is systematic overinvestment relative to NE.

5.4 Informational Externalities

Having established the presence of overinvestment in all treatments, are the markets more or less informative than the NE benchmarks? We define a measure of (the aggregate size of) the informational externalities associated with a distribution of strategies, and we compute this measure given the empirical distribution as well as given the degenerate distribution in which all subjects play the NE. We will see what aspects of behavior in the lab and what features of the environment affect the comparison.

We define our measure of the informational externalities as $IE(\mu) = \Pi^{BR}(\mu) - \underline{\Pi}$, where μ is a distribution of strategies in the population[20], $\Pi^{BR}(\mu)$ is the ex ante expected profit of a player who best-responds to μ, and $\underline{\Pi}$ is the ex ante expected profit of someone who behaves optimally given her private information but ignores others' behavior.[21] Thus, $IE(\mu)$ shows the optimal profits of an investor who knows μ 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.[22]

[20] We will abuse notation by letting a strategy s also denote a distribution over strategies which is degenerate on s.

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

[22] When we compare actual investment or profits to the NE benchmarks, we consider ex post NE in order to eliminate possible noise due to the realizations of investment returns, signals, and costs. However, in computing information externalities, the appropriate calculation of best responses is ex ante.

Let μ^a be the actual distribution of strategies in the population and let s^{NE} denote the NE strategy. Ideally, we would like to compare $IE(\mu^a)$ and $IE(s^{NE})$. However, this is impractical for two reasons. First, we do not observe individuals' strategies so that it is difficult to estimate μ^a. Therefore, we replace μ^a with the average strategy in the population $\overline{s}^a$ which, for each history and type, gives the average probability of investment that results from μ^a.[23] Our estimate of $\overline{s}^a(h_r; X_i, c_i)$ will simply equal the observed frequency of investment at history h_r by type (X_i, c_i).

Second, at histories beyond round 2 in the ten-player games, we typically have too few observations to be able to reliably estimate $\overline{s}^a(h_t; X_i, c_i)$ so that we cannot compute best-response behavior beyond round 2. Thus, in the ten-player games, instead of $\Pi^{BR}(\overline{s}^a)$, we will use truncated best-response (TBR) profits, $\Pi^{TBR}(\overline{s}^a)$. These are defined as the ex ante expected profits that arise if a player best-responds to $\overline{s}^a$ in rounds 1 and 2, and, in round 3, invests iff it is profitable to invest (given $\overline{s}^a$ in rounds 1 and 2) and otherwise never invests.[24] Based on this, we will use $IE^T(\overline{s}^a) = \Pi^{TBR}(\overline{s}^a) - \underline{\Pi}$ as a proxy for $IE(\mu^a)$. In the two-player games, best-responding subjects never invest after round 3, so $IE^T(\overline{s}^a) = IE(\overline{s}^a)$ holds.

Similarly, we use $IE^T(s^{NE}) = \Pi^{TBR}(s^{NE}) - \underline{\Pi} = \Pi^{TNE} - \underline{\Pi}$ as a proxy for $IE(s^{NE})$ (where $\Pi^{TBR}(s^{NE})$ is defined analogously to $\Pi^{TBR}(\overline{s}^a)$). This definition offers a fair comparison to our approximation of the actual informational externalities, $IE^T(\overline{s}^a)$, and also avoids a nearly intractable computation involving hundreds of histories.

Table 3 shows $\underline{\Pi}$, as well as TBR and TNE profits for each game in each treatment. From the information in the table, $IE^T(\overline{s}^a)$ and $IE^T(s^{NE})$ can readily be computed for each game in each treatment. Based on this, we can state:

Result 6

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

(ii) $IE^T(\overline{s}^a)$ *is smaller than* $IE^T(s^{NE})$ *in the ten-player two-cost game (0.31 vs. 0.497).*

[23] Formally, $\overline{s}^a$ is defined by $\overline{s}^a(\emptyset; X_i, c_i) = \int s(\emptyset; X_i, c_i) d\mu^a$ and $\overline{s}^a(k_1, \ldots, k_r; X_i, c_i) = \int s(k_1, \ldots, k_r; X_i, c_i) \prod_{l=1}^{r-1} [1 - s(k_1, \ldots, k_l; X_i, c_i)] d\mu^a$. Here, the integration is done over strategies s. The terms under the integrals give the probability that a player using s ends up investing after the given history.

[24] We compute TBR profits using the one-step property.

(iii) In the remaining games, the magnitudes of the informational externalities in the actual markets and in NE are similar.

Thus, informational flows in the actual experimental markets can be larger, smaller, or similar in size to those in NE. Before turning to the aspects of behavior and features of the environment that affect this comparison, first 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 (i.e. by type $(1, H)$ or $(1, L)$) 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 (i.e. by type $(0, H)$ or $(0, L)$) tends to decrease informational flows because it weakens this inference. These effects tend to be largest in round 1.

The aspects of behavior in the lab described in result 1(i) (that says that round-1 behavior depends on one's types), result 3(i) (that says that round-1 behavior is independent of cost structure and market size), and result 4 (that says that type $(0, H)/(0, L)/(1, H)/(1, L)$ invests in round 1 with a small /nonnegligible/substantial/large-but-considerably-less-than-one frequency) play a key role in determining how $IE^T(\overline{s}^a)$ and $IE^T(s^{NE})$ compare. Based on these aspects of behavior, the following three features of the environment systematically affect the impact of animal spirits on informational externalities relative to NE.

1. 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 while it never invests in round 1 in NE.[25]

2. Presence of strong type with favorable common-value signal and low cost (type $(1, L)$ in our settings):

[25] Although there is no such type in our experiment, conceivably, the presence of a type with a favorable common-value signal but with a high enough investment cost so that the expected profit of investing in round 1 is only slightly negative, would 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.

The presence of such a type often has a direct negative effect on informational externalities relative to NE because, in NE, such a type often (at least for markets that are not very large) invests with probability 1 while, in the actual data, it invests with a probability that falls considerably short of 1.

The presence of such a type also has an indirect positive effect on informational externalities relative to NE in games where there are players with favorable common-value signal but with a higher cost (i.e., in games with 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 result 3(i), such players in the lab do not do so and, by result 4, 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.[26]

3. Market size:

   In a larger market, a given investment probability by players with a favorable common-value signal leads to a higher option value of waiting because players who wait observe a larger, and hence more informative, sample. Thus, the NE round-1 investment probabilities for types with a favorable common-value signal are lower while, by result 3(i), behavior in the lab is unaffected by market size. Thus, a larger market size has a positive effect on the informational externalities relative to NE.

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

[26]The indirect effect could also be negative if there is a type 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.

5.5 Average Profits

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 decompose the expected average profit on the market given a distribution of strategies in the population μ, $\pi(\mu)$, into two components (plus a constant) both of which are related to informational externalities. In particular, we can write:

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

The first component is $IE(\mu)$ 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) by either investing too early or not investing (early enough). Because the cost of mistakes is strictly positive outside of NE, the only way for actual markets to exceed NE efficiency is to attain an increase in the size of informational externalities relative to NE that is larger than the cost of mistakes. Given that the informational externalities in the lab are larger than those in NE only in the ten-player high-cost game, this is the only game where we have any hope of observing higher market efficiency than in NE.

In order to compare market efficiency in the lab with the NE benchmark, we consider actual average profits and compare them to TNE profits. Note that, TNE profits provide a lower bound on NE profits from an ex ante perspective, and they coincide with NE profits for the two-player games (both ex ante and ex post). However, in order to eliminate any differences between actual and TNE profits which are due to noise in the realization of the gross return, the common-value signals, and the costs, we compute TNE profits from an ex post perspective for each market trial.

Table 4 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 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, the increase in informational externalities due to animal spirits is nearly equal to the cost of mistakes, so that the difference does not attain statistical significance.

Let us summarize:

Result 7 *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.*

What kind of mistakes are players making in our experiment? In particular, what are the most important, i.e. costly, ways in which subjects are failing to best respond to the empirical frequencies?

Table 5 shows the difference between the expected profits of investing in round 1 and the expected profit (given $\overline{s}^a$) of waiting in round 1 and following TBR play afterwards. Combining the information from Figures 1-4 and Table 5, we see that decisions to invest in round 1 by $(0, H)$ players are very costly but, fortunately, players commit such mistakes infrequently. Decisions to invest in round 1 by $(0, L)$ players are also quite costly and occur with nonnegligible frequency. In the ten-player games, decisions to invest in round 1 by $(1, H)$ players are quite costly and occur quite frequently. Players also commit mistakes in round 2, predominantly by being too conservative.

5.6 Subject Heterogeneity

We now turn to whether subjects are heterogeneous, in terms of the generation of informational externalities. The issue is important for our understanding of animal spirits. First, if our overinvestment results are due to all subjects investing with a higher probability in round 1 than the NE benchmarks, then animal spirits (at least in our context) would be a universal property of human nature. On the other hand, if there is substantial heterogeneity in round-1 overinvestment across subjects, then animal spirits would be a property of certain individuals who have a greater urge to action than others.

Second, animal spirits need not manifest themselves in the same way in all individuals. For example, some individuals may overinvest with a high probability in round 1 only when they have a favorable common-value signal–these individuals generate large informational externalities; at the same time, other individuals may overinvest regardless of their common-value signal–these individuals add a lot of noise which dilutes informational flows on the market.

Here we classify subjects according to their probabilities of investing in round 1 for each type. In particular, we assume each player belongs to group $g \in \{1, \dots, G\}$. Each group of players, g, is characterized by a four-tuple $(q^g_{0H}, q^g_{0L}, q^g_{1H}, q^g_{1L}) \in [0,1]^4$ where the first/second/third/fourth component represents the probability of round-1 investment for type $(0,H)/(0,L)/(1,H)/(1,L)$. For each possible number of groups, G, we will estimate via maximum likelihood $(q^g_{0H}, q^g_{0L}, q^g_{1H}, q^g_{1L})$ for each $g \in \{1, \dots, G\}$. Then, based on the Bayesian Information Criterion (BIC) (see Schwarz (1978)), we will choose G.

In order to write down the likelihood function, let I_i be a 24-tuple, such that its t^{th} component records what subject i did in round 1 in period t. Also, let T_i be a 24-tuple, such that its t^{th} component records subject i's type in period t. Denote by $I_{i0H}/I_{i0L}/I_{i1H}/I_{i1L}$ the number of times subject i invested when type $(0,H)/(0,L)/(1,H)/(1,L)$ and by $X_{i0H}/X_{i0L}/X_{i1H}/X_{i1L}$ be the number of times subject i had type $(0,H)/(0,L)/(1,H)/(1,L)$. Then, the probability of observing I_i, given T_i and the subject's group g equals:

$$\begin{aligned} Pr(I_i|T_i, g; q^g_{0H}, q^g_{0L}, q^g_{1H}, q^g_{1L}) = \; & (q^g_{0H})^{I_{i0H}}(1-q^g_{0H})^{X_{i0H}-I_{i0H}}(q^g_{0L})^{I_{i0L}}(1-q^g_{0L})^{X_{i0L}-I_{i0L}} \\ & (q^g_{1H})^{I_{i1H}}(1-q^g_{1H})^{X_{i1H}-I_{i1H}}(q^g_{1L})^{I_{i1L}}(1-q^g_{1L})^{X_{i1L}-I_{i1L}} \end{aligned}$$

Letting p_g be the ex ante probability of subject i belonging to group g, we can obtain the probability of i's round-1 investment choices before her group is known:

$$Pr(I_i|T_i;\{p_g\}_{g=1}^G,\{q_{0H}^g,q_{0L}^g,q_{1H}^g,q_{1L}^g\}_{g=1}^G)=\sum_{g=1}^G p_g Pr(I_i|T_i,g;q_{0H}^g,q_{0L}^g,q_{1H}^g,q_{1L}^g)$$

Finally, the probability of all subjects' round-1 behavior is given by:

$$Pr(\{I_i\}_{i=1}^n|\{T_i\}_{i=1}^n;\{p_g\}_{g=1}^G,\{q_{0H}^g,q_{0L}^g,q_{1H}^g,q_{1L}^g\}_{g=1}^G)=\prod_{i=1}^n Pr(I_i|T_i;\{p_g\}_{g=1}^G,\{q_{0H}^g,q_{0L}^g,q_{1H}^g,q_{1L}^g\}_{g=1}^G) \quad (3)$$

Treating G as fixed, we estimate $\{p_g\}_{g=1}^G$ and $\{q_{0H}^g,q_{0L}^g,q_{1H}^g,q_{1L}^g\}_{g=1}^G$ by maximizing (3) for each treatment. We repeat this estimation for $G=1,2,\dots$ and we compute the BIC for each G. We select the optimal G as the one that minimizes the BIC. The optimal G equals 4/4/4/5 in the $A2/R2/A10/R10$. Table 6 reports the estimates of $\{p_g\}_{g=1}^G$ and $\{q_{0H}^g,q_{0L}^g,q_{1H}^g,q_{1L}^g\}_{g=1}^G$ for the optimal G in each treatment. Columns 8 and 9 report $q_0^g=\frac{1}{2}(q_{0H}^g+q_{0L}^g)$ and $q_1^g=\frac{1}{2}(q_{1H}^g+q_{1L}^g)$, which are relevant from the point of view of how much a group of subjects contributes to informational externalities in a given treatment.

Based on Table 6, we see a similar picture emerging in all treatments. Around 61-86 percent of subjects fall in a group with a relatively low q_0^g and high q_1^g (groups 1 and 2 in the $A2$, groups 1, 2, and 3 in the $R2$, groups 1 and 2 in the $A10$, and groups 1,2, and 3 in $R10$) and thus generate a relatively strong informational externality. The remaining subjects fall in groups that generate a weak-to-modest informational externality. We can state:

Result 8 *There is considerable subject heterogeneity in the generation of informational externalities. Based on their round-1 behavior, 61-86 percent of subjects are in groups generating a relatively strong informational externality while the remaining subjects are in groups generating a weak-to-modest informational externality.*

Our finding of strong subject heterogeneity in round 1 behavior indicates that animal spirits are individual-specific and not universal to human nature. An interesting question is whether our finding of passive spirits in round 2 reflects selection bias, where subjects with strong animal spirits are unlikely to wait until round 2. An alternative and equally interesting explanation is that passivity in round 2 holds across the board, so that animal spirits is both heterogeneous (in round 1) and situation specific (round 1 vs. round 2). By comparing the round-2 behavior of subjects with above average investment tendencies in round 1 and subjects with below average investment tendencies in round 1, we observe no tendency for the former group to be less passive in round 2 (details available from authors). Thus, passivity in round 2 seems to hold across the board.

In the appendix, we take a different angle on subject heterogeneity. In particular, we consider whether subjects who differ in terms of their cognitive abilities, as measured by their SAT scores, also differ in terms of their round-1 behavior and in their propensity to exploit informational externalities when in round 2.[27]

5.7 Departures From Risk Neutrality

So far, we have been conducting our analysis under the assumption of risk neutrality. One could argue, however, that overinvestment in round 1 may be driven by risk loving[28] and overly conservative behavior in round 2 may be driven by risk aversion. Thus, behavior that we attribute to animal spirits or passive spirits could be due to nothing more than risk loving or risk aversion, respectively. Although departures from risk neutrality could play a role in the lab, we don't believe they play an important role for explaining behavior.

First, even under alternative attitudes to risk, NE behavior depends on the cost structure and market size. Yet, by result 3, actual behavior does not. This is more consistent with behavior being driven by gut feelings that are not "the outcome of a weighted average of quantitative benefits multiplied by quantitative probabilities."

Second, if we assume that attitudes to risk are driving behavior, our finding of pervasive round-1 overinvestment would suggest that risk loving dominates risk aversion. However, this stands in stark contrast with individual choice studies, according

[27]In our analysis so far, we have for simplicity been ignoring the possibility of learning as sessions progress. In the appendix, we explicitly consider learning.

[28]Assuming CARA utility, risk loving does indeed (weakly) increase round-1 investment probabilities for each type in NE.

to which risk aversion is widespread and risk loving is rare.[29]

Third, in order to explain why subjects who overinvest in round-1 are overly conservative in round 2, one needs to assume that they are risk-loving in round 1 and risk-averse in round 2–clearly a very awkward explanation.

6 Concluding Remarks

We observe overinvestment relative to the NE benchmarks in all of our games, which we identify with animal spirits. The effect of animal spirits on the strength of the informational externalities depends on the environment. In the ten-player high-cost game, all of the features of the environment suggest that animal spirits generate larger informational externalities than the NE benchmark: (i) there are no type $(0, L)$ subjects present whose investment would weaken the information content of market activity, and (ii) since the market is large, the NE investment probability for type $(1, H)$ is only 0.071, so the benchmark informational externality is low. We conjecture that large experimental markets (with, say, 50+ subjects) will generate larger informational externalities than the NE benchmark. The reason is that, in such markets, the NE requires that only the most favorably inclined subjects invest in round 1 and, moreover, that they do so with a low probability. If animal spirits are present, so that a significant minority of subjects with favorable common-value signal invest in round 1, then we will see massive overinvestment by such subjects relative to NE and an associated dramatic increase of information flows relative to NE.

References

Akerlof, George A and Robert J. Shiller (2009). "Animal Spirits: How Human Psychology Drives the Economy, and Why It Matters for Global Capitalism," Princeton and Oxford: Princeton University Press.

Anderson, Lisa R and Charles A.Holt (1997). "Information Cascades in the Laboratory," *American Economic Review*, vol. 87, no. 5, pp. 847-62.

[29]See Holt and Laury (2002).

Avery, Christopher and Peter Zemsky (1998). "Multidimensional Uncertainty and Herd Behavior in Financial Markets," *American Economic Review*, vol. 88, no. 4, pp. 724-48.

Azariadis, Costas (1981). "Self-Fulfilling Prophecies," *Journal of Economic Theory*, vol. 25, no. 3, pp. 380-96.

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

Benhabib, Jess, Roger E. A. Farmer (1999). "Indeterminacy and Sunspots in Macroeconomics," *Handbook of Macroeconomics*, vol. 1A., pp. 387-448.

Bikhchandani, Sushil, David Hirshleifer, and Ivo Welch, (1992). "A Theory of Fads, Fashion, Custom, and Cultural Change in Informational Cascades," *Journal of Political Economy*, vol. 100, no. 5, pp. 992-1026.

Cass, David and Shell, Karl (1983). "Do Sunspots Matter?," *Journal of Political Economy*, vol. 91, no. 2, pp. 193-227.

Chamley, Christophe and Douglas Gale (1994). "Information Revelation and Strategic Delay in a Model of Investment," *Econometrica*, vol. 62, no. 5, pp. 1065-85.

Cipriani, Marco and Antonio Guarino (2005). "Herd Behavior in a Laboratory Financial Market," *American Economic Review*, 95(5): 1427-1443.

Crawford, Vincent P. and Nagore Iriberri (2007). "Level-k Auctions: Can a Nonequilibrium Model of Strategic Thinking Explain the Winners Curse and Overbidding in Private-Value Auctions?" *Econometrica*, 75(6): 172170.

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

Duffy, John, Eric O. Fisher (2005). "Sunspots in the Laboratory," *American Economic Review*, vol. 95, no. 3, pp. 510-29.

Eyster, Erik and Matthew Rabin (2005). "Cursed Equilibrium," *Econometrica*, 73(5): 1623-1672.

Fischbacher, Urs (2007). "Z-Tree: Zurich Toolbox for Ready-Made Economic Experiments." *Experimental Economics*, 10(2): 17178.

Holt, Charles A. and Susan K. Laury (2002). "Risk Aversion and Incentive Effects," *American Economic Review*, vol. 92, no. 5, pp. 1644-55.

Howitt, Peter and R. Preston McAfee (1992). "Animal Spirits," *American Economic Review*, vol. 82, no. 3, , pp. 493-507.

Ivanov, Asen, Dan Levin, and James Peck (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.

Keynes, J. M. (1964). "The General Theory of Employment, Interest, and Money," Harcourt Brace Jovanovich, New York/London, First Harbinger Edition.

Levin, Dan and James Peck (2008). "Investment Dynamics with Common and Private Values," *Journal of Economic Theory*, vol. 143, no. 1, pp. 114-39.

Marimon, Ramon, Stephen E. Spear, and Shyam Sunder (1993). "Expectationally Driven Market Volatility: An Experimental Study," *Journal of Economic*

Theory, vol. 61, no. 1, October, pp. 74-103.

McKelvey, Richard D. and Thomas R. Palfrey (1995). “Quantal Response Equilibria for Normal Form Games,” *Games and Economic Behavior*, 10: 6-38.

McKelvey, Richard D. and Thomas R. Palfrey (1998). “Quantal Response Equilibria in Extensive Form Games,” *Experimental Economics*, 1: 9-41.

Nagel, Rosemarie (1995). “Unraveling in Guessing Games: An Experimental Study,” *American Economic Review*, 85: 1313-1326.

Peck, James (1988). “On the Existence of Sunspot Equilibria in an Overlapping Generations Model,” *Journal of Economic Theory*, vol. 44, no. 1, pp. 19-42.

Schwarz, Gideon (1978). “Estimating the dimension of a model,” *Annals of Statistics*, 6(2): 461-464.

Sgroi, Daniel (2003). “The Right Choice at the Right Time: A Herding Experiment in Endogenous Time,” *Experimental Economics*, 6: 159-180.

Shell, Karl (1977). “Monnaie et allocation intertemporelle” mimeo, Sminaire Roy-Malinvaud, Centre National de la Recherche Scientifique, Paris.

Shleifer, Andrei (1986). “Implementation Cycles,” *Journal of Political Economy*, vol. 94, no. 6, pp. 1163-90.

Stahl, Dale O. and Paul W. Wilson (1994). “Experimental Evidence on Players’ Models of Other Players,” *Journal of Economic Behavior and Organization*, vol. 25, no. 3, pp. 309-27.

Stahl, Dale O. and Paul W. Wilson. (1995). “On Players’ Models of Other Players:

Theory and Experimental Evidence," *Games and Economic Behavior*, 10: 218-254.

Ziegelmeyer, Anthony, Kene Boun My, Jean-Christophe Vergnaud, and Marc Willinger (2005). "Strategic Delay and Rational Imitation in the Laboratory," Max Planck Institute of Economics Discussion Paper on Strategic Interaction 200535.

7 Appendix: SAT Scores and Behavior in Rounds 1 and 2

In this section, we consider whether subjects who differ in terms of their cognitive abilities, as measured by their SAT scores, also differ in terms of their round-1 behavior and in their propensity to exploit informational externalities when in round 2.

We start with round-1 behavior. We focus on the $(1, H)$ type, because this type faces the most interesting tradeoff between investing in round 1 and waiting to observe others' round-1 behavior.[30] We estimate the following random effects probit model for $(1, H)$ subjects in each treatment:[31]

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

The estimated marginal effect of the SAT score is -0.00075/-0.00130/-0.00046/-0.00022 in the $A2/R2/A10/R10$ and is significant (at the 5-percent level) in the $A2$ and $R2$.[32] Thus, subjects with higher SAT scores are less likely to invest in round 1 (significant in the $A2$ and $R2$). Waiting in round 1 is a best response (given $\overline{s}^a$) in the $R2$, $A10$, and $R10$ but not in the $A2$. Ivanov, Levin, and Peck (2009) find evidence that the decision to wait by a $(1, H)$ subject has more to do with whether the subject has the foresight to realize that there is an option value to waiting than with a calculation of the exact option value. This would explain why subjects with higher SAT scores are more likely to wait whether or not this is optimal.

Turning to round-2 behavior, we ask whether subjects with higher SAT scores are more responsive to observed round-1 behavior in round 2. We estimate the following

[30] Types $(0, H)$ and $(0, L)$ have a negative expected profit based on their private information and should thus always wait in round 1. For type $(1, L)$ in the two-player treatments, there is no information to be gained by waiting because, even if the other subject has an unfavorable common-value signal, it is still worth it to invest. In the ten-player treatments, there is no informational advantage in waiting for this type in the NE and given the empirical frequencies of play.

[31] We have the SAT scores for around two-thirds of all subjects. The estimations in this section are performed only for these subjects.

[32] If we use GPA instead of SAT, we obtain significance in the $A10$ as well. Note also that the estimated effects are rather large in size: to obtain the difference in predicted probability between two subjects who are two standard deviations apart in their SAT scores, we need to multiply the estimated coefficients by 319.

random effects regression for each type in each treatment:

$$y = \beta_0 + \beta_1 k_1 + \beta_2 SAT + \beta_3 k_1 SAT + v + u \tag{4}$$

where y is a dummy that equals 1 if subject i invested in round 2, v is an individual-specific error term and u is an *iid* idiosyncratic shock.[33]

Note that almost all estimates of β_2 are negative and almost all estimates of β_3 are positive.[34] Based on this, we can state:

Result 9 *We find evidence for subject heterogeneity in terms of how they exploit informational externalities. In particular:*

(i) Subjects with higher SAT scores are more likely to wait in round 1 when type $(1, H)$ *(whether or not this is a best-response).*

(ii) Subjects with higher SAT scores tend to be less likely to invest in round 2 after low values of k_1 *(*β_2 *tends to be negative) and tend to be more responsive to* k_1 *(*β_3 *tends to be positive).*

8 Appendix: Learning

In our 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 right-hand side variable:

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

[33]The estimation uses only observations where the subject did not invest in round 1. Also, we use a regression rather than a probit because the size of the marginal effect of the interaction term $k_1 SAT$ has no natural interpretation within the latter. Probit estimation leads to similar results in terms of the signs of all effects and significance levels.

[34]Although many of the estimates in the ten-player games are not significantly different from 0, the fact that all 8 estimates of β_2 have the same sign and 7 out of 8 estimates of β_3 have the same sign strongly suggests that there is a systematic effect of SATs on round-2 behavior. Note also that to obtain a sense of the differences in behavior between two subjects who are two standard deviations apart in their SAT scores, we need to multiply the estimated coefficients by 319.

Figure 6 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 figure, 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 $A10$.

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

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

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.

Let us summarize our main findings.

Result 10 *There is a reduction in the probability of round-1 investment for type $(1, H)$ over the course of the session in all treatments. This reduction is substantial in size and statistically significant in the $A2$, $R2$ and $R10$. In round 2, we find little evidence for changes in behavior over the course of the session.*

9 Appendix: Figures and Tables

A2 Treatment

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

Figure 1: Round 1 and Round 2 Frequencies of Investment in the $A2$. The subscripts on the actual frequencies show the number of decisions that were made at each history.

R2 Treatment

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

Figure 2: Round 1 and Round 2 Frequencies of Investment in the $R2$. The subscripts on the actual frequencies show the number of decisions that were made at each history.

A10 Treatment

	Nash				Actual			
	(0,H)	**(0,L)**	**(1,H)**	**(1,L)**	**(0,H)**	**(0,L)**	**(1,H)**	**(1,L)**
after ∅	0	0	0.071	0.746	0.046(239)	0.152(217)	0.394(241)	0.787(263)
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)	-

Figure 3: Round 1 and Round 2 Frequencies of Investment in the $A10$. The subscripts on the actual frequencies show the number of decisions that were made at each history; "-" indicates that a given history never occurred.

R10 Treatment

	Nash				Actual			
	(0,H)	**(0,L)**	**(1,H)**	**(1,L)**	**(0,H)**	**(0,L)**	**(1,H)**	**(1,L)**
after ∅	0	0	0	1	0.058(241)	0.169(231)	0.348(250)	0.681(238)
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)

Figure 4: Round 1 and Round 2 Frequencies of Investment in the $R10$. The subscripts on the actual frequencies show the number of decisions that were made at each history; "-" indicates that a given history never occurred.

	$(\mathbf{0},\mathbf{H})$	$(\mathbf{0},\mathbf{L})$	$(\mathbf{1},\mathbf{H})$	$(\mathbf{1},\mathbf{L})$
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 1: Marginal Effect of k_1 on the probability of investment in round 2. (*/**/*** indicates significance at the 10/5/1 percent level.)

	A10			
	(0,H)	**(0,L)**	**(1,H)**	**(1,L)**
$\overline{k}_1$	8.71	5.33	5.54	2.93
	(2.94)	*(0.37)*	*(1.29)*	*(0.84)*
$\overline{k}_1^{NE}$	3.5	4.5	1.26	2.66
$\overline{k}_1^{BR}$	4.5	4.5	1.5	3.5

	R10			
	(0,H)	**(0,L)**	**(1,H)**	**(1,L)**
$\overline{k}_1$	6.85	4.86	4.58	3.26
	(0.82)	*(0.39)*	*(0.37)*	*(0.76)*
$\overline{k}_1^{NE}$	4.5	3.5	2.76	0.5
$\overline{k}_1^{BR}$	4.5	3.5	2.5	0.5

Figure 5: Round 2 cutoffs.

	(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 2: Investment Frequencies

	High-Cost Game	Low-Cost Game	Two-Cost Game
$\underline{\Pi}$	0.250	1.750	1
A2 NE ex ante	0.250	2.034	-
A2 BR ex ante	0.250	1.947	-
R2 NE ex ante	-	-	1.073
R2 BR ex ante	-	-	1.084
A10 TNE ex ante	0.295	2.380	-
A10 TBR ex ante	0.590	2.214	-
R10 TNE ex ante	-	-	1.497
R10 TBR ex ante	-	-	1.310

Table 3: TBR vs. TNE profits (in ECU)

	(0,H)	(0,L)	(1,H)	(1,L)	Overall High-Cost Game	Overall Low-Cost Game	Overall Two-Cost Game
A2 NE 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 NE 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 4: Average actual profits vs. TNE profits (in ECU)

	(0,H)	(0,L)	(1,H)	(1,L)
A2	-3.500	-0.894	0.050	0.350
R2	-3.500	-0.727	-0.110	0.350
A10	-3.593	-1.423	-0.525	0.123
R10	-3.588	-1.081	-0.548	0.309

Table 5: Expected profit of investing in round 1 minus expected profit from waiting in round 1 (in ECU) given empirical frequencies of play

Treat-ment	Group	p_g	q^g_{0H}	q^g_{0L}	q^g_{1H}	q^g_{1L}	q^g_0	q^g_1	Log-likelihood
A2									-857.3348
	1	0.16	0.00	0.03	0.92	0.98	0.01	0.95	
	2	0.45	0.01	0.13	0.32	0.90	0.07	0.61	
	3	0.25	0.03	0.01	0.14	0.41	0.02	0.27	
	4	0.14	0.28	0.52	0.44	0.73	0.40	0.59	
R2									-538.3518
	1	0.39	0.00	0.04	0.18	0.93	0.02	0.55	
	2	0.24	0.08	0.07	0.56	0.61	0.08	0.59	
	3	0.23	0.22	0.37	0.76	1.00	0.29	0.88	
	4	0.14	0.00	0.00	0.03	0.35	0.00	0.19	
A10									-348.8657
	1	0.38	0.03	0.03	0.71	0.97	0.03	0.84	
	2	0.28	0.01	0.05	0.10	0.94	0.03	0.52	
	3	0.19	0.00	0.00	0.20	0.28	0.00	0.24	
	4	0.15	0.20	0.71	0.28	0.69	0.45	0.48	
R10									-359.0541
	1	0.32	0.04	0.24	0.66	1.00	0.14	0.83	
	2	0.13	0.03	0.00	0.44	0.60	0.02	0.52	
	3	0.20	0.00	0.00	0.00	0.92	0.00	0.46	
	4	0.21	0.00	0.04	0.03	0.19	0.02	0.11	
	5	0.14	0.30	0.63	0.44	0.35	0.46	0.39	

Table 6: Groups of players based on round-1 behavior

	$(\mathbf{0}, \mathbf{H})$	$(\mathbf{0}, \mathbf{L})$	$(\mathbf{1}, \mathbf{H})$	$(\mathbf{1}, \mathbf{L})$
A2 $\boldsymbol{\beta_2}$	-0.00048***	0.00007	-0.00043	-0.00093*
A2 $\boldsymbol{\beta_3}$	0.00022	0.00074***	0.00099**	0.00197***
R2 $\boldsymbol{\beta_2}$	-0.00046***	-0.00054**	-0.00079***	0.00103**
R2 $\boldsymbol{\beta_3}$	0.00033	0.00117***	0.00115***	-0.0001
A10 $\boldsymbol{\beta_2}$	-0.00036	-0.00065	-0.00085	-0.00204
A10 $\boldsymbol{\beta_3}$	0.00000	0.00004	0.00027	0.00055*
R10 $\boldsymbol{\beta_2}$	-0.00024	-0.0001	-0.0011*	-0.00141
R10 $\boldsymbol{\beta_3}$	0.00005	-0.00002	0.00037**	0.0003

Table 7: Effects of SAT scores on round-2 investment probabilities. */**/*** indicates significance at the 1/5/10 percent level.

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

Figure 6: Round 1 Investment and Learning.

10 Appendix: Instructions for R2

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:

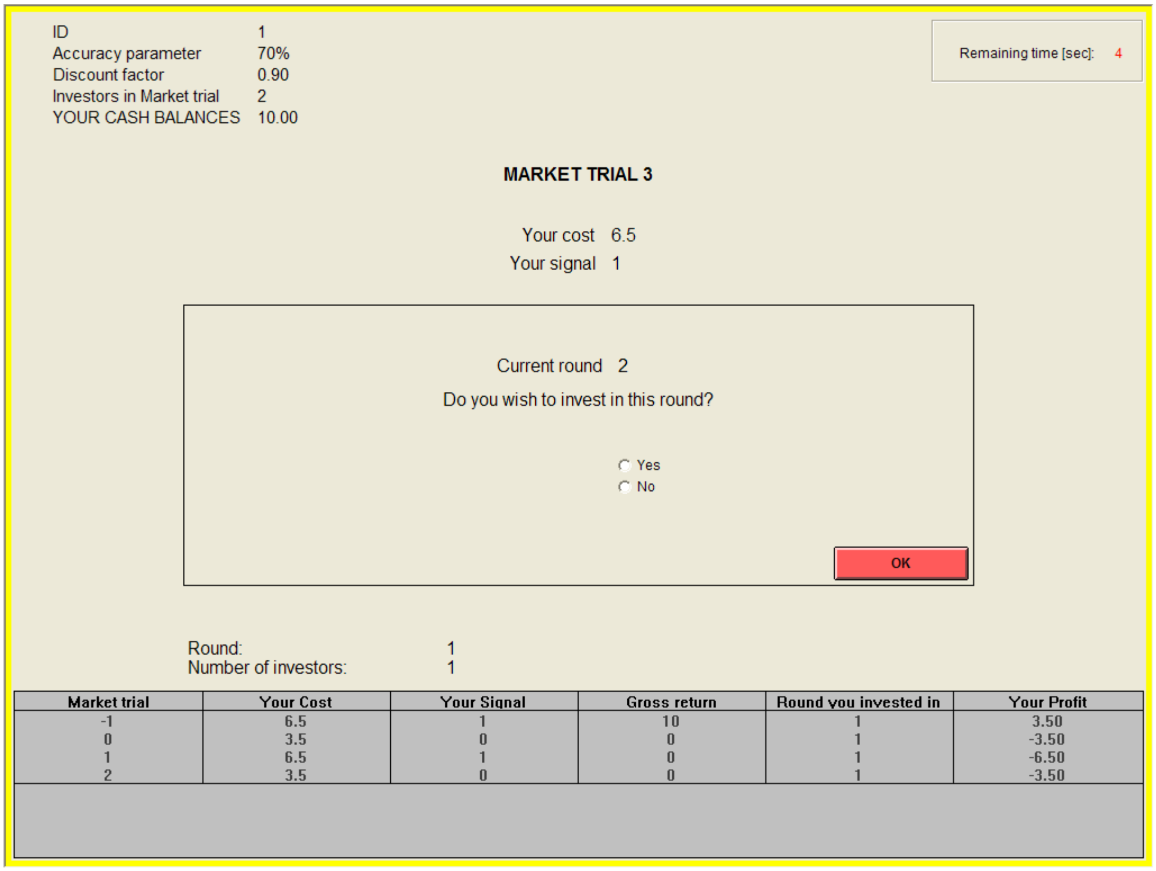

The screen in the R2, A10, and R10 looks the same (except that in the ten-player games we would have "Investors in Market trial 10").