

Averting Economic Collapse and the Solipsism Bias*

Antonio Guarino, Steffen Huck and Thomas D. Jeitschko[†]

September 1, 2005

Abstract

We study the behavior of experimental subjects who have to make a sequence of risky investment decisions in the presence of network externalities. Subjects follow a simple heuristic—investing after positive experiences and reducing their propensity to invest after a failure. This result contrasts with the theoretical findings of Jeitschko and Taylor (2001) in which even agents who have only good experiences eventually stop investing because they account for the fact that others with worse experiences will quit. This can trigger sudden economic collapse—a coordination avalanche—even in the most efficient Bayesian equilibrium. In the experiment, subjects follow their own experiences and disregard the possible bad experiences of others—thus exhibiting behavior that we term “solipsism bias.” Solipsism results in sustained investment activity and thus averts complete collapse.

KEYWORDS: Coordination, Coordination Avalanche, Economic Collapse, Experimental Economics, Network Externalities, Solipsism Bias

JEL CODES: C7, C9, D8, E0, G0

*We wish to thank Dirk Engelmann, Wieland Müller, Hans Normann, Pedro Rey Biel, and Georg Weizsäcker for many helpful comments and suggestions. Financial support from the Leverhulme Trust and the ESRC via ELSE is gratefully acknowledged.

[†]Guarino and Huck are at the Department of Economics and ELSE, University College London. Jeitschko is at the Department of Economics, Michigan State University. Correspondence to Steffen Huck, Department of Economics & ELSE, University College London, Gower Street, London WC1E 6BT, United Kingdom, email: *s.huck@ucl.ac.uk*.

1 Introduction

Consider a situation in which agents can invest repeatedly in a risky project. The project can be lucrative or not, depending on the state of the world. The outcome of the project also depends on the decisions of other agents, as there are network externalities or complementarities in actions. Agents only observe the outcomes of their own investments and then decide whether to invest again. As Jeitschko and Taylor (2001), who first analyzed this class of situations, have shown, a global economic collapse—a “coordination avalanche”—can occur: In such an event, agents become discouraged and stop investing even after having experienced *only positive outcomes*.

At a first glance this result appears counter-intuitive. But, in fact, it follows from a rather simple logic: Agents have to consider the possibility that there are other agents who, as time passes, have bad experiences. After some bad experiences these agents will decide not to invest. Because these agents drop out, the positive network externality is not big enough for agents who had just slightly better experiences than those with the worst experiences. These agents, therefore, will drop out as well, and so on. Jeitschko and Taylor show that this chain can go on in such a way that *all* agents *simultaneously* abandon the investment. The local discouragement triggers and propagates a coordination avalanche. Even if the state of the world is good (and all investment opportunities are profitable), the sheer possibility that other agents may abandon the project induces everyone to do so. Notice that this is not the result of coordination failure. Indeed, the collapse of investment occurs in the Pareto-efficient equilibrium of the game.

In addition to its obvious macro-economic implications, the coordination avalanche theory has many other interesting applications. Jeitschko and Taylor, for instance, show that the avalanche can explain organizational meltdowns: Agents may simultaneously abandon an efficient organization only because each one fears that his colleagues may do so. Similarly, the fear that other agents may upgrade their computers may induce a sudden

and inefficient shift to a new program, in order to keep compatibility with friends and coworkers. In recent years it has also been argued that it is difficult to explain some financial crises by looking at the fundamentals of the economy and that, instead, they seem triggered just by panic. Given that network externalities are important in these markets, the avalanche theory can also be thought of as a contributing factor in such crises.

In this paper we discuss some implications of the Jeitschko and Taylor model through an experimental study. In particular, we consider an environment in which features of a coordination avalanche are present in order to ascertain whether this phenomenon is observed in the laboratory. One of the central question of this undertaking is whether subjects indeed abandon a risky investment after observing only good outcomes.

The answer of our experimental study to this questions is negative. Subjects in the laboratory seem to follow a simple behavioral rule that is in stark contrast with the perfect Bayesian equilibrium solution of the game: They invest after they receive good outcomes. Once they have a bad experience, their propensity to invest drops, on average, from almost 100 percent to less than 60 percent. While this behavior is not in line with individual payoff maximization (after failing, subjects should quit investing in order to maximize their payoffs and even after successes they should stop investing at a certain point), continued investment provides a public good because of the network externality. Positive network externalities last longer in the laboratory than they do in theory and, consequently, sudden economic collapse is avoided. Because of this, the heuristic allows subjects to avoid a collapse of investment activity and, when there is a high level of coordination, they earn payoffs similar to those in the Pareto-efficient equilibrium.

While our experimental results run counter to the equilibrium of the theoretical model, they highlight an interesting anomaly in subjects' decision making. Subjects receiving good economic outcomes fail to realize that others might have experienced different results—a phenomenon that we term “solipsism bias.”

In Section 2, we discuss the model and its Pareto-efficient equilibrium. In Section 3, we illustrate the experiment and the experimental design. In Section 4, we present our results. Finally, in Section 5 we discuss our results, and relate them to other observations in laboratory experiments.

2 Theory

2.1 The model

We present a simplified version of the model by Jeitschko and Taylor. Consider an economy with a continuum of agents with mass normalized to 1. In each period, $t = 1, 2, \dots, T$, agents are randomly paired and called upon to play a simultaneous-move, two-by-two investment game, akin to a stochastic version of the classic stag-hunt game.

At the outset of each period each agent receives an endowment of ω (> 0). This can either be consumed completely (strategy choice “NO”), in which case the agent’s payoff is simply ω and the period effectively ends for the agent; or a portion c ($0 < c < \omega$) of the endowment can be invested (strategy choice “INVEST”).

In the case that the agent chooses “INVEST,” his payoff depends on two things: the choice of the agent he is matched with; and a chance move by nature. If his partner chooses “NO,” the investment of c is forfeited and the agent obtains a payoff of $\omega - c$. If his partner chooses “INVEST,” with probability p the investment is a success, yielding a (net) payoff of $(\omega - c) + S$. Otherwise (with probability $1 - p$), the investment fails, yielding a payoff of $\omega - c$.

Denote the measure of the population that invests by m . An agent’s per-period expected payoff for action $x \in \{\text{INVEST}, \text{NO}\}$, $u(x)$, is given by

$$\begin{aligned} u(\text{NO}) &= \omega, & \text{and} \\ u(\text{INVEST}) &= (\omega - c) + mpS. \end{aligned} \tag{1}$$

Agents observe their own payoffs, but not other agents’ payoffs or strat-

egy choices (i.e., m), nor nature’s chance moves.

Provided that the success rate of investments, p , and the payout of successful investment, S , are sufficiently high compared to the cost of investing, c , (i.e., $pS > c$), the one-shot game has two symmetric pure strategy equilibria: with all playing “INVEST” ($m = 1$) Pareto-dominating the all play “NO” ($m = 0$) equilibrium. The one-shot game is then a game of pure coordination.

Now suppose that while ω, c and S are known values, there is uncertainty about the magnitude of the intrinsic success rate of investments, p . In particular, suppose that p is fixed over time, but takes one of two possible values, \underline{p} or \bar{p} , reflecting an underlying good or bad state of the economy. Suppose further that $0 \leq \underline{p} < c/S < \bar{p} \leq 1$ so that simultaneous investments by all agents are expected to be profitable in the good state, but not in the bad state. Lastly, let the expectation of the state of the economy be such that $\mathbf{E}[p] > c/S$ so that initially investment is collectively rational.

Assume that all agents have the correct common prior, so that the game is one of pure coordination—at least initially, at time $t = 1$. Assume further that it is common knowledge that all agents tacitly coordinate on the good action in the first period of play. Then, nevertheless, in the first period some agents may experience failed investment outcomes while others do not. Those with successful investment outcomes will attach more weight to the likelihood that the value of the intrinsic success rate of investments is \bar{p} ; those with failures find it more likely that it is \underline{p} .

Jeitschko and Taylor refer to learning about p due to investment success as an *informational effect*. However, they point out that there is also a *marginal network externality* that agents must account for. To appreciate this, note that since $\underline{p} < c/S$, if any positive mass of agents puts sufficiently strong weight on the state being bad, they will choose to play “NO,” and complete coordination, in which $m = 1$, can no longer be achieved. All agents can determine after how many failed investment attempts a Bayesian agent, who assumes successful coordination amongst the population, will no

longer invest. At this point, in addition to updating beliefs about p , all agents must form beliefs about m . Thus, agents update their beliefs on the success rate of investments (the informational effect), and also account for a (negative) marginal network externality, namely, how attrition affects the size of m .

As agents assess the decrease in m , they expect lower returns from investing for any given beliefs about the state of the world. Consequently, it may be the case that agents with fairly optimistic beliefs regarding p begin playing “NO,” due to the imputed negative network externality caused by those who have become pessimistic about p . A point may even come where agents with uniformly good experiences, i.e., those who have never had an investment fail, realize that there likely is sufficient attrition in the population to no longer warrant an investment. In other words, while their beliefs about p are strictly increasing (due to their good experiences), their estimate of m (which is weakly decreasing) may dominate so that the composite expected success-rate of investments, mp , falls below the critical threshold. At this point, the only action supportable in equilibrium is the one in which all agents play “NO” and $m = 0$.

It is important to note that whether such a point is in fact reached is a function of prior beliefs and need not depend on the true state of the world. In fact, even if $p = \bar{p} \equiv 1$, this type of complete attrition may occur. This is demonstrated in the parameterization chosen below.

2.2 Parameterization and the Pareto-efficient perfect Bayesian equilibrium

We turn now to illustrate the features of the model for the particular parameters that we chose for the experiment. Agents are repeatedly randomly matched for a total of 5 periods. At the beginning of each period they are newly endowed with $\omega = \mathcal{L}1.70$. The cost of undertaking an investment is $c = \mathcal{L}1.64$, and successful investments pay out $\mathcal{L}2.00$ (per agent). There is no discounting. In the bad state of the world, the success rate is $\underline{p} \equiv 0.8$,

whereas in the good state is $\bar{p} \equiv 1$. Both states are *ex ante* equally likely.

We are interested in determining the longest string of playing “INVEST,” under complete coordination. Complete coordination means that an agent plays “INVEST” whenever it is individually rational to do so provided that all others with equivalent or more optimistic beliefs also play “INVEST.” That is, we are looking for a Pareto-efficient perfect Bayesian equilibrium. Since $\mathbf{E}[p] = 0.9$, if all agents invest in the first period, they receive an expected payoff of

$$u(\text{INVEST}) = (\pounds 1.70 - \pounds 1.64) + (.9)(1)(\pounds 2.00) = \pounds 1.86,$$

which is greater than the $\pounds 1.70$ obtained from playing “NO.” So all agents invest in the first period.

Under this assumption, if an agent experiences a failure in the first period, the intrinsic success rate of investments is revealed to this agent to be $p = \underline{p}$ (because there are no failures with $\bar{p} \equiv 1$). Thus, agents with a failure in Period 1 will have full information and, as is easily verified, will cease to invest—regardless of the number of people in the population who do invest.

Agents who experience a success will not know the true state of the world, but will increase the weight that they attach to the state being good. In particular, the Bayesian estimate of the state of the world being good after a success in the first period (when $m = 1$) is $5/9$. As they cannot be sure that no-one has experienced a failure, however, they will anticipate that there may be some agents who quit investing in the second period. Hence, assuming that all agents with a success in the first period continue investing, an agent with a first period success expects with probability $5/9$ to have another success *for sure* in the second period when investing. However, with the complementary probability of $4/9$, he expects an effective success rate of only 64% (the probability of being matched with someone who invests *and* experiencing a success with that person, when the true p is $\underline{p} \equiv 0.8$). This yields a composite expected success rate $\mathbf{E}[pm|\text{success}]$ of 0.84 and hence $u(\text{INVEST}|\text{success}) = \pounds 1.74$. Since $\pounds 1.74 > \pounds 1.70 = u(\text{NO})$, all

agents that experience a success in the first period also invest in the second period.

Agents who in the second period experience a failure also know that the true state of the world is the bad state, for either they were matched with someone who invested, yet their investment failed; or they were matched with someone who had previously observed that the state was bad, and hence chose to play “NO.”

On the other hand, those with a success in the second period will be even more optimistic that the state is the good state. While the Bayesian update puts a probability of $25/41$ on the state that all agents have been successful with their investments throughout the game, agents must also account for the fact that with probability $16/41$ the state is actually bad and attrition has spread through two periods of play. In the latter instance, the probability of being matched with someone who has had consecutive good experiences is only 64% and the expected success if both were to invest is, of course, only $\underline{p} \equiv 0.8$. Hence, the expected payoff when investing if all with a history of two good experiences invest (while others properly refrain from investing) is only about $\mathcal{L}0.88$, yielding $u(\text{INVEST}|2 \text{ successes}) \approx \mathcal{L}1.60$. So, despite being rather optimistic about the state of the world after two consecutive successes, it is no longer worth investing, because in the event that the state is bad, attrition will have spread so far as to make investments very unlikely to pay off.

Needless to say, once all agents stop playing “INVEST” in the third period, attrition is complete, no more learning takes place and there is no reason to play “INVEST” in any of the subsequent periods either.

This yields the following Proposition:

Proposition (Pareto Efficient Perfect Bayesian Equilibrium) *The strategy profile that specifies that all agents play “INVEST” until they have either experienced a failure in the previous period or have reached the third period, and thereafter play “NO,” constitutes a Pareto efficient*

perfect Bayesian equilibrium (PEPBE).

An implication of the PEPBE of the Proposition is that all agents play “NO” in Periods 3 through 5, regardless of their experiences in the first two periods, and regardless of the true value of p .¹

3 Experimental design and procedures

3.1 The experiment

We conducted the experiment in the ELSE laboratory at University College London. We recruited subjects from the College’s undergraduate population across all disciplines. They had no previous experience with this or similar experiments. For each session of the experiment we recruited 8 students. In total, we recruited 120 subjects to run 15 sessions (five sessions for each of three treatments).

At the beginning of the sessions, we handed out written instructions to all subjects (see Appendix *B*). Subjects were made aware that they all had the same instructions. Afterwards, subjects saw a simulated run of the experiment on their computer screens in order for them to understand all the steps that they had to go through during the experiment, the timing of the decisions that they had to take, and the information that they would have received after each decision. Finally, we asked if there were any clarifying questions.

Each session consisted of a series of 15 rounds. Each round was a sequential 5-period game as described in the previous section. Let us illustrate the procedures for each round. At the beginning of the round the software chose, with equal probability of 0.5, whether the investment that subjects

¹The PEPBE of the Proposition is not unique. Any prescription that involves all playing “NO” in any three of the five periods, and prescribes maximal cooperation (in the history-dependent sense above) in the other two, is also a PEPBE. We do not consider multiple PEPBE a concern for the questions of interest as the Proposition gives a very clear focal point. Nevertheless, we do address this issue in our treatments (see Section 3.2).

were going to make would be successful with probability 1 or with probability 0.8 only.² We explained this to students by writing in the instructions that the computer would randomly choose, by “flipping a coin,” between two urns, Urn 1 or Urn 2. Urn 1 contained only green balls and Urn 2 contained 80 green and 20 red balls. The chosen urn, i.e., the chosen probability distribution, was used for all participants and for all periods in this round. Subjects were not informed of which distribution had been selected.

After the probability distribution (urn) was chosen, subjects had to decide whether or not to invest in a project for five times in sequence. For each decision in the sequence, subjects had to act simultaneously, i.e., not knowing what other subjects decided to do. After all subjects made their first decision, for each subject who had invested, the computer drew a ball from the chosen urn to decide whether the investment was successful or not. A green ball meant that the investment was successful, and a red ball that it was a failure. The draws were independent, i.e., keeping the urn metaphor, balls were drawn with replacement.

After all participants had made their first decision, they were randomly matched into pairs by the software. There was an equal probability for each subject of being matched with anyone of the other seven participants. The matching co-determined the payoff of subjects who had decided to invest. A subject would earn money by investing under two conditions: that the computer had drawn a green ball; and—at the same time—that the subject she was matched with had decided to invest as well.

After all 8 subjects had made their decision, they were informed of their payoff. We gave subjects £1.70 for each of the five decisions. They could use that money to pay for the cost of the investment or just keep it, if they decided not to invest. If they decided to invest, they had to pay a cost of £1.64. If the investment was successful and the match invested too, a subject

²Our choice of $\bar{p} \equiv 1$ for the good state of the world has two advantages. First, it makes the updating process easy. Second, given that we are primarily interested in the behavior of agents who had good experiences, it maximizes the number of interesting observations in our experiment.

earned £2; and nothing otherwise. Therefore, in the case of an investment, the subject had two possible payoffs: $£1.70 - £1.64 + £2 = £2.06$ in the first case; and $£1.70 - £1.64 + £0 = £0.06$ in the other.

Subjects were informed only of their payoff, and nothing else. Therefore, a subject could not distinguish whether the low payoff was due to bad luck, i.e., the red color of the ball, or due to the decision of the other participant. After learning about their payoffs, participants had to make their second investment decision. They knew that the same distribution (urn) previously chosen would be used again, but that new balls would be drawn for everyone investing.

Subjects had to repeat this decision for 5 periods in a round. After the fifth period of the round, each participant could see on the screen what urn the computer had chosen in this round, the color of the balls that he was assigned for each investment decision, and the decisions of the other participants he was matched with in each of the five periods. The computer also showed the total payoff for that round, i.e., the sum of the payoffs obtained in each period.

After the first round was over, we repeated the same procedure for the second round: At the beginning the computer chose again one of the two distributions and students made their five decisions. The same procedure followed for a total of 15 rounds. Subjects were made aware that they would be paid (in private, immediately after the experiment) on the basis of a show up fee and of their payoff in three rounds randomly chosen, one from each third of the experiment.

3.2 Treatments

We ran three treatments for this experiment. The procedures for the first treatment (“Treatment A,” from now on) were as described in the previous section. In addition to that, we asked subjects to write down their beliefs, i.e., what their expected value of the investment was. In particular, in each period, we asked students how much, at most, they would have paid for

the investment. They could choose a number between $\mathcal{L}0$ and $\mathcal{L}2$ with ticks of 1 penny. We did so in order to understand how beliefs evolved over time. Notice, in fact, that the decision of a subject to invest (not to invest) only tells us that he thought the investment was worth more (less) than its cost ($\mathcal{L}1.64$) but it does not reveal the precise value he attaches to the investment. We did not provide monetary incentives for this task as the true value of the investment depends on both a subject’s posterior belief about the chosen urn and on his expectation about others’ investment behavior. As we cannot know the latter (without eliciting these expectations through another mechanism), it is impossible to rely on a simple incentive mechanisms for this task.³

In another treatment (“Treatment *B*”) we included a feature that would help subjects to coordinate on the Pareto optimal equilibrium. Remember that our primary interest is in observing the behavior of subjects after some successful experiences. After all, this is a key implication of the theoretical model: Agents may quit investing because of fear that others do so (despite having made only positive experiences on their own). If no one invests all the time, there is no possibility of such observations. Therefore, to help subjects to coordinate on the PEPBE described in the Proposition, we forced subjects in Treatment *B* to invest in Period 1. We disabled the “don’t invest” button for that period. Therefore, in this treatment, subjects actually had to make only four decisions, from the 2nd to the 5th period. In this treatment we also elicited beliefs, although in a different way than in Treatment *A*. Instead of asking subjects in each period how much, at most, they would have paid for the investment, we asked them to write down the probability with which they expected the subject they would be matched with to invest. For this

³Interestingly and somewhat surprisingly, Sonnemans and Offerman (2001) show for a task where correct beliefs can be specified that subjects’ actual beliefs can be as accurately elicited without incentives as with a quadratic scoring rule. It is also worth mentioning that, despite the absence of an incentive mechanism, in our experiment, subjects’ answers were fairly consistent: Indeed, in 85% of the cases in which they invested, they reported a belief higher than 1.64; in 95% of the cases in which they did not invest, they reported a belief lower than 1.64.

task we did provide monetary incentives: indeed, we paid them on the basis of a standard quadratic scoring function.⁴

In a third treatment (“Treatment *C*”) we tried to create coordination on the Pareto optimal equilibrium by adding an extra device. Not only did we force subjects to invest in Period 1, as in Treatment *B*, we also told them that whenever they decided not to invest, they would not have the opportunity to invest again in later periods of that round. We made the decision of not investing being definitive, so that subjects would think seriously whether it was the right decision or not. This eliminates all PEPBE outcomes, except for the one described in the Proposition, so the PEPBE of the game is unique.⁵

In what follows we shall mainly focus on treatments *A* and *B*. For brevity’s sake, a detailed analysis of treatment *C* is relegated to the appendix.

4 Results

We start by describing the aggregate investment decisions over time in Treatment *A*. Table 1 shows the proportion of investment decisions for all five periods averaged over all rounds. These proportions are compared with those of the Pareto optimal equilibrium, given the realization of the chance moves, i.e., the frequency of the two urns.

Table 1 reveals two important features of the data—departing from the Pareto efficient equilibrium investment rates. First, there is a large amount of coordination failure. Fewer than half of the decisions in the first period are

⁴Subjects chose on the screen a value between 0% and 100%, indicating the probability of an investment. The computer would automatically show the complementary probability of not investing, so that subjects could always see the entire vector of probabilities. We paid them using the following formula: $[10 - (\text{truth} - \text{prediction})^2 / 1000]$ pence. We chose to pay them a small amount for this task to be sure that the belief elicitation did not turn the game into one of pure coordination, in which subjects would have an incentive to coordinate on a specific strategy to gain more money instead of normally playing the game (see on this, e.g., Nyarko and Schotter, 2002).

⁵Note that in this treatment we did not elicit beliefs.

Aggregate Data A	Period 1	Period 2	Period 3	Period 4	Period 5
PEPBE					
Investment Rates	1	0.91	0	0	0
Observed					
Investment Rates	0.44	0.38	0.33	0.29	0.28
Observations	600	600	600	600	600

Table 1: PEPBE and observed investment rates in Treatment A.

Only Successes A	Period 2	Period 3	Period 4	Period 5
PEPBE				
Investment Rates	1	0	N/A	N/A
Observed				
Investment Rates	0.94	0.93	0.94	1.00
Observations	134	74	49	33

Table 2: PEPBE and observed investment rates for subjects who have only experienced successes in Treatment A.

investments. Second, there is no sharp drop in investment rates in Period 3, in which all subjects should stop investing. Instead the decline in investment rates is slow and steady over all five periods.

To explain this pattern we looked at subjects' decisions after histories of successful or unsuccessful investments. Table 2 shows the investment rates of subjects who had only experienced successes in the previous periods. Table 3 shows the same for subjects who had experienced a failure in the previous period.⁶ From these two tables, a very simple pattern emerges: Subjects keep investing as long as they are successful. Once they experience a failure, their propensity to invest drops, on average, to less than 60 percent.

This behavioral pattern is also reflected in the belief data that we elicited.

⁶Note that in this table we are pooling different histories, as we are not distinguishing whether the failure at time $t - 1$ was the first one, or followed previous failures. We prefer to pool such histories in order to base our results on a larger number of observations. Furthermore, our choice does not affect our results in a significant way. Subjects reacted to a failure in the previous period in a similar way to that shown in the table, independently of what happened in the history until $t - 2$. For instance, if we consider histories in which subjects had a failure only at time $t - 1$, the investment rates are 0.58, 0.52, 0.8 and 0.69.

Previous Failure A	Period 2	Period 3	Period 4	Period 5
PEPBE	0	0	N/A	N/A
Investment Rates				
Observed	0.58	0.57	0.58	0.60
Investment Rates				
Observations	132	116	92	96

Table 3: PEPBE and observed investment rates for subjects who have experienced a failure in the previous period in Treatment A.

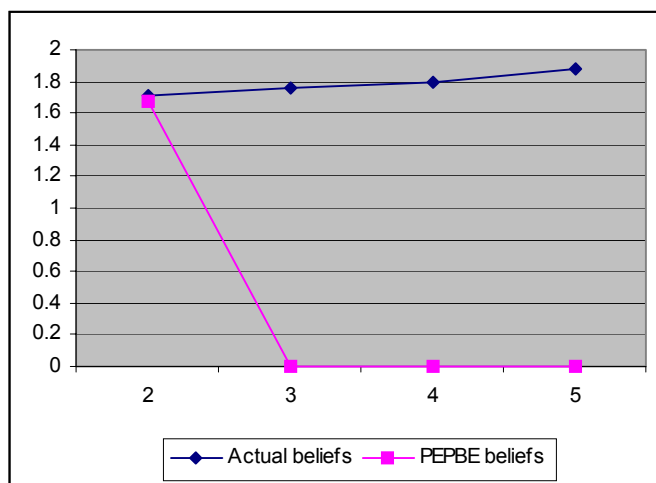


Figure 1: PEPBE and observed average beliefs about the expected value of an investment, given a history that consists of successes only (Treatment A).

In Figure 1, we plot the PEPBE beliefs (downward sloping curve) as well as the reported beliefs for subjects who were always successful when investing in the previous periods (upward sloping curve). In the PEPBE the expected value of the investment decreases over time. In Period 3 it becomes lower than the cost and, therefore, from Period 3 on, in equilibrium, agents attach a value of 0 to the investment, as all agents decide not to invest. The reported beliefs are in stark contrast to this. In Period 2 on average subjects valued the investment at £1.7, almost identical to the full coordination value of £1.68. After experiencing successes, however, subjects are more and more optimistic, becoming increasingly confident that the true state is the good one and that other subjects will keep investing, too.

A pattern of increasing confidence might be expected if there is initial uncertainty about the level of coordination in the population, yielding lower expected returns at first, and then subsequent increasing confidence as fears of lacking coordination are dispelled. Such a sequence of beliefs, however, cannot exceed the most optimistic beliefs in the PEPBE. As reported beliefs rise above those of the PEPBE, it is clear that subjects are updating in the *wrong direction*.

While these data draw a rather coherent picture of subjects' behavior in Treatment *A*, they are not ideally suited for comparison with the PEPBE given in the Proposition. There is simply too much coordination failure early on. In Treatment *B*, to help subjects to coordinate, we forced subjects to invest in the first period.

Table 4 shows aggregate investment rates over time for this treatment. Tables 5 and 6 show the investment rates for subjects who had only successes and for subjects who experienced a failure in the preceding period.⁷

Several observations are in order. First, the coordination in the first period has a strong effect for the later periods. Once coordinated, subjects

⁷As with Treatment *A*, also in this treatment subjects reacted to a failure in the previous period in a similar way to that shown in the table, independently of what happened in the history until $t - 2$. For instance, if we consider histories in which subjects had a failure at time $t - 1$ only, the investment rates are 0.44, 0.49, 0.48, 0.65.

Aggregate Data B	Period 1	Period 2	Period 3	Period 4	Period 5
PEPBE					
Investment Rates	1	0.91	0	0	0
Observed					
Investment Rates	1	0.66	0.60	0.52	0.49
N	600	600	600	600	600

Table 4: PEPBE and observed investment rates in Treatment B.

Only Successes B	Period 2	Period 3	Period 4	Period 5
PEPBE				
Investment Rates	1	0	N/A	N/A
Observed				
Investment Rates	.68	.96	.97	.99
Observations	538	278	208	172

Table 5: PEPBE and observed investment rates for subjects who have only experienced successes in Treatment B.

Previous Failure B	Period 2	Period 3	Period 4	Period 5
PEPBE				
Investment Rates	0	0	N/A	N/A
Observed				
Investment Rates	0.44	0.50	0.41	0.48
Observations	62	97	101	92

Table 6: PEPBE and observed investment rates for subjects who have experienced a failure in the previous period in Treatment B.

keep coordinating and the average investment rate in Period 2 of Treatment *B* is almost twice as high as the average investment rate in Period 2 of Treatment *A*. This result is similar to earlier findings on the importance of the first few periods for the entire trajectory of play (see, for example, van Huyck, Battalio, and Beil, 1991, or van Huyck, Cook, and Battalio, 1997).

Second, the decline in investment rates mirrors the decline observed in Treatment *A* in that it is steady and slow. Moreover, there is no sudden drop-off in the investment rate in Period 3. Third, subjects seem to apply the same decision rules as in Treatment *A*. They keep investing as long as investments are successful.⁸ Once discouraged by a failure, the investment rate drops to slightly smaller numbers than in Treatment *A*.

This behavioral pattern, and, in particular, the investment rates after successful experiences, is easily explained if, again, we look at elicited beliefs. Recall that in this treatment we elicited subjects' beliefs on the probability of being matched with another participant investing in the project. In Figure 2, we plot the PEPBE beliefs (downward sloping curve) as well as the reported beliefs for subjects who were always successful when investing in the previous periods (upward sloping curve).

In the PEPBE the expected value of the investment is higher than the cost in Periods 1 and 2, and lower in the successive periods. In the first period each agent attaches probability 1 to his random match deciding to invest. In the second period this probability goes down to 0.91 since one has to take into account that (if the state of the world is bad) some agents could experience a failure at time 1. From Period 3 onwards, the expected value of the investment becomes lower than the cost and, therefore, in equilibrium, agents attach 0 probability to the other agents deciding to invest. The reported beliefs are in stark contrast to this. In Period 2, after experiencing

⁸It is worth mentioning that in Table 5 the figure for Period 2 is not strictly comparable to those for the other periods. Indeed, it refers to the investment rate after being successful in a previous *enforced* investment, while the others refer to the investment rates after *voluntary* (successful) investment. In other words, the small figure for that period reflects the unwillingness of some subjects to coordinate on investment during the experiment.

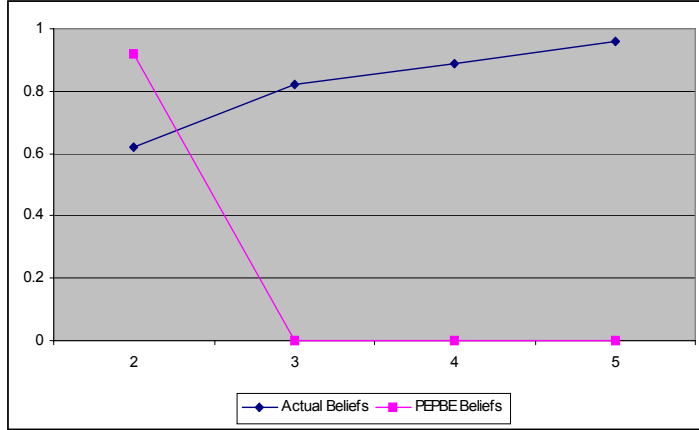


Figure 2: PEPBE and observed average beliefs about m , given a history that consists of successes only (Treatment B).

the first success, subjects attach on average probability 0.62 to their match deciding to invest. After experiencing more successes, they become more and more optimistic, becoming increasingly confident that other subjects would keep investing too.⁹¹⁰ These beliefs are higher than those of the PEPBE, making it clear that subjects are unable to account for other participants' possible bad experiences.

The data that we have described so far refer to the entire experiment and one may wonder whether there is any time trend in the data, perhaps toward the PEPBE. Analyzing the last few rounds separately, however, we see that behavior is remarkably stable over time. Subjects' actions are consistent with the simple behavioral rule throughout the experiment. In Tables 7 and 8 we report the investment decisions in the last five rounds of the fifteen-round sessions. Even at the end of the experiment, subjects in both

⁹Note that the graph starts from Period 2. Indeed, given that in Period 1 we obliged subjects to invest, and this was common knowledge, the belief of others investing was trivially 1. Therefore, we elicited beliefs starting from Period 2.

¹⁰Note that, while the belief data from Treatment B have the advantage of being elicited with incentives, we cannot check whether beliefs are consistent with actions since we do not observe subjects' beliefs about the underlying state of the economy.

Treatment	Only Successes (last 5 rounds)	Period 2	Period 3	Period 4	Period 5
A	Investment Rate	0.89	0.95	1	1
	Observations	35	21	14	9
B	Investment Rate	0.66	0.94	1	1
	Observations	172	85	61	56

Table 7: Observed investment rates for subjects who have only experienced successes in the last five rounds of the two treatments.

Treatment	Previous Failure (last 5 rounds)	Period 2	Period 3	Period 4	Period 5
A	Investment Rate	0.32	0.61	0.80	0.67
	Observations	41	37	36	35
B	Investment Rate	0.39	0.38	0.39	0.47
	Observations	28	32	28	19

Table 8: Observed investment rates for subjects who have experienced a failure in the previous period.

treatments invested almost always after having experienced successes;¹¹ and significantly reduced their propensity to invest after a failure.

All the previous analysis refers to data aggregated over all participants. To shed more light on subjects' behavior, we also looked at individual decisions. Table 7 illustrates how subjects reacted to a successful investment. In particular, it divides subjects according to their different propensities to invest after a success. Behavior is remarkably homogeneous in both treatments: Almost 70% of subjects in Treatment *A* and almost 60% in Treatment *B* re-invest in more than 80% of all cases after a success.¹² Table 8 shows how subjects reacted to a failure. Behavior here is clearly more heterogeneous. The distribution of investment rates is almost uniform in both

¹¹The same consideration of footnote 8 applies to the figures for Treatment *B* in Table 7.

¹²After histories of successes only, there is even more homogeneity in individual behavior. Indeed, 86% of subjects in Treatment *A* and 60% in Treatment *B* re-invest in more than 80% of the time after such histories.

Percentage of Investments after a Success	Treatment A	Treatment B
[0, 20)	0	0.15
[20,40)	0.07	0.10
[40,60)	0.16	0.10
[60,80)	0.10	0.08
[80,100]	0.68	0.58

Table 9: Individual decisions: subjects are divided according to the percentage of the time in which they invested after a success in the previous period.

Percentage of Investments after a Failure	Treatment A	Treatment B
[0, 20)	0.09	0.20
[20,40)	0.33	0.28
[40,60)	0.18	0.13
[60,80)	0.21	0.23
[80,100]	0.18	0.15

Table 10: Individual decisions: subjects are divided according to the percentage of the time in which they invested after a failure in the previous period.

treatments.¹³

As we said, we also ran a third treatment. To obtain coordination on the PEPBE, not only did we oblige subjects to invest in Period 1, we also imposed that, after deciding not to invest, subjects were not able to invest in subsequent periods. While our devices were successful in obtaining a very high level of coordination, subjects' behavior conditional on histories of successes or failures is essentially identical to that explained so far. Therefore, we do not report the results here, and refer the reader to Appendix A for a detailed illustration.

¹³At the very extremes of the distribution we find that in Treatment A 17.5% of subjects never invest, while there is no one who always invested. In Treatment B 7.5% of subjects never invest, while 10% always invested.

5 An interpretation of the findings

5.1 The solipsism bias

The data from the two treatments draw a clear picture of subjects' behavior. After a success, they keep investing—which is quite at odds with the Pareto efficient equilibrium. After a failure, they reduce their propensity to invest. Some subjects do this very drastically, while others reduce their investment rates in a more moderate manner—again, at odds with the PEPBE. Together, by relying on such simple backward-looking heuristics, subjects avoid economic collapse. It is worthwhile to note that by investing more than they should, subjects provide a substantial public good for all the others because of the rather big externalities of investments. As a consequence of this, the subjects in Treatment *B* earn almost as much as they would in the PEPBE (£8.24 versus £8.75).¹⁴

Basically, it is subjects' inability to apply equilibrium reasoning that “protects” them from experiencing a collapse. Subjects seem unable to put themselves in others' shoes, to imagine that others may have had a different past than they did. In particular, subjects who were successful appear to have difficulties in realizing that others might have experienced failures. This anomaly — that we introduce here and hope to add to the canon of anomalies in experimental and behavioral economics — is related to what has been observed in other experiments, the closest being, perhaps, those on “herding” and “informational cascades” (as introduced in the seminal contributions by Banerjee, 1992, and Bikhchandani, Hirshleifer, and Welch, 1992).

Experimental studies on cascades have sometimes found support for the Bayesian equilibrium prediction and sometimes not. Recent evidence in a

¹⁴In Treatment *C* (where there is even more coordination on the investment action, see the appendix) the difference between actual and theoretical payoffs is even smaller: £8.57 versus £8.67 of the PEPBE. In Treatment *A*, the large level of coordination failure lowered the average payoff to £7.30, versus a theoretical level of £8.76. Note that the PEPBE payoffs are slightly different across treatments, due to the different empirical frequencies of the urns.

paper by Kübler and Weizsäcker (2004) suggests that the same behavioral rule can explain both the theory’s successes and its failures. The key feature of the behavioral rule Kübler and Weizsäcker identify is that subjects basically fail to understand that others whom they observe and from whom they learn have also learned from others whom they, in turn, had observed. Put differently, subjects do extract information from their predecessors’ actions, but they do not take into account that their predecessors did the same.

A similar failure of understanding the experience and reasoning of others also causes the discrepancy between equilibrium and observed play in our experiment. We call it the “solipsism bias.” Solipsism refers to a philosophical school of thought that argues that the self can only identify itself as real and that the self can only be aware of its own experiences. The Oxford English Dictionary defines solipsism as “the view or theory that the self is the only object of real knowledge or the only thing really existent” which is amazingly paraphractical of what we observe. Successful investors do not envision a different set of experiences, a state after experiencing a failure. Such a solipsism bias is distinctly different from limited depth of reasoning (see, for example, Nagel 1995, Camerer, Ho, and Chong 2004, in addition to Kübler and Weizsäcker, 2004). Models of limited depth of reasoning entail as a crucial feature that agents have beliefs about others against to which they best respond, i.e., they are assumed to have a mental model of others. In our framework this would likely affect the timing of collapse, but would not lead to a persistent continuation of investment activity.

The solipsism bias—as we define it—may also account for other experimental observations. For example, in Carroll, Bazerman, and Maury’s (1988) study of the so-called acquiring-a-company problem (which, in essence, is a no-trade problem). Carroll *et al.* suggest that the observed acquisitions that should not take place in theory are due to the tendency to “ignore the cognitions of others.” A similar tendency is apparent in Sovik’s (2004) data on the no-betting conjecture.¹⁵

¹⁵Cooper and Kagel (2004) argue that fictitious play is also very much ignoring that

In order to substantiate the applicability of solipsism bias to our findings, we briefly consider whether our results can be explained by simply arguing that the theoretical model of Jeitschko-Taylor is not robust to minor changes in behavioral assumptions, so that our data may be a reflection of other common and well-known behavioral deviations from equilibrium behavior.

5.2 Robustness of the Jeitschko-Taylor model

Besides the specific features of the model that we have replicated in the experiment, the Jeitschko-Taylor model assumes that agents do not make mistakes, that they update in a Bayesian fashion, that they are able to coordinate, and that they are risk neutral. One may wonder whether the model is robust to small deviations from these assumptions. If not, our results could simply be due to the fact that subjects in an experiment do make mistakes, are not perfect Bayesian updaters, are unable to coordinate perfectly and are not necessarily risk neutral. In this section we show that, indeed, small deviations from the Jeitschko-Taylor assumptions do not explain our results.

Trembles and belief distortions. In an experiment one has to expect subjects to make mistakes. If mistakes happen and this is common knowledge, does the PEPBE still hold? For vanishing mistakes we know, of course, that the PEPBE is robust, since it is trembling-hand perfect in the sense of Selten (1975). But if trembles and errors in updating are bounded away from zero, what can we conclude?

As explained in the description of the model, in order to come to an estimate of the expected success rate of investments, pm , subjects must weigh both an informational/learning effect (concerning p) and a marginal network externality (concerning m). What is important for the analysis is whether or not subjects recognize the latter, even as they become more optimistic

others have a state of mind, too: “A fictitious play learner treats his opponents like a fixed statistical distribution rather than forming a model of how his opponents make decisions.” However, they also identify some subjects whom they call “sophisticated learners” and who appear to have models of how others form decisions.

about the underlying state of the world due to good experiences. Whether the exact evaluation and weighing of the two opposing effects is done according to Bayesian updating is secondary, in fact, not even particularly material for the question of abandoning the project after good experiences. Thus, note that Bayes' rule is used only to update one's beliefs about p , whereas one's assessment of the size of m is based on equilibrium reasoning on the basis of one's updated beliefs about p . That is, updating one's beliefs about m requires one to imagine others' experiences and that they will react to these possibly different experiences.




Figure 3: Belief evolution under different rules

In order to assess how deviation from Bayesian learning about p affects optimal behavior, Figure 3 depicts alternative beliefs that a subject attaches to the probability that the good urn was drawn after having consecutive good experiences. Specifically, the lines starting at the true prior that $p = \bar{p} \equiv 1$ (i.e., 0.5) trace four different paths for alternative updating rules. The constant line refers to subjects who do not update—despite successes

in Periods 1 through 5, they continue to attach the same probability to the state being the good state. The next higher curve shows the belief evolution of a Bayesian updater. The third highest of these lines depicts the beliefs of someone who updates by assuming that the good state is t -times as likely as the bad state after observing $t - 1$ good outcomes. The remaining updating rule depicted is that of a subject who, after a success, always reduces the (residual) likelihood of the bad state by half.

In contrast to the alternative beliefs about p , the line starting in the origin gives the critical threshold level that one's beliefs about the likelihood of the good state must exceed in order to make it worthwhile to continue to invest—provided that one's equilibrium reasoning about m is correct, given one's beliefs about p . That is, regardless of how one updates beliefs about p , so long as one properly assesses the evolution of m (given one's beliefs about p), a subject should continue to invest so long as his beliefs about p are greater than those given by the line starting in the origin.

Notice that the first three updating rules discussed result in no-one investing in the third and subsequent periods, i.e., all of these updating rules are below the critical threshold for continued play in the third period. The only updating rule depicted that does not yield collapse in the third period is the highest, which represents a rather dramatic built-up of confidence after good outcomes.

Of course, one might argue that rather than equilibrium reasoning about m , the relevant benchmark is reasoning that properly anticipates population responses to experiences. Indeed, we observe that even after a failure, a substantial portion of agents continue to invest — just under 60% in Treatment *A* and about 45% in Treatment *B* (see Tables 3 and 6). Accounting for this leads to a decrease in the critical threshold level (given by the line starting in the origin, in Figure 3). However, even under the supposition of full coordination (which is not obtained in either treatment), when assuming that subjects show a propensity to invest of 95% following a success (which is somewhat greater than observed, see Tables 2 and 5), and allowing for

the greater investment rate of 60% following a failure in Treatment B , the lower bound obtained on the critical threshold in the third period remains substantially above the Bayesian update (i.e., .64 vs. .61).

On balance, we find that only the most drastic deviations from Bayesian updating could explain the deviations from equilibrium play that we observe. Many plausible non-Bayesian updating rules would also predict sudden economic collapse after period 2, even if we consider the underlying responses in m in line with the observed data.

Mis-coordination. Perhaps the most plausible source of departure from the PEPBE is coordination failure. In fact, we observed such a failure in one of our treatments. In Treatments B and C , however, we used some devices to reduce the risk of coordination failure. While these devices worked well in avoiding coordination failures, subjects' reactions to experiences were still remarkably similar to those observed in the other treatment and remarkably different from those of the PEPBE. Moreover, it should be noted that, on the whole, the departure from the PEPBE that we observe in the experiment is that subjects exhibit a tendency to over-invest after good experiences, something that cannot be accounted for by coordination failure.

Risk aversion. Another possibility is that the theoretical prediction simply does not apply as subjects in the lab may not be risk neutral. From Rabin's (2000) calibration theorem, however, we can safely deduce that in our experiment all expected utility maximizers would behave as if risk neutral. This follows from the observation that, since utility depends on final wealth, utility functions must be either locally flat in the range of experimental payoffs (which only affect final wealth to a negligible degree) or globally bizarre. In other words, as long as one wants to maintain the assumptions of expected utility theory to find a theoretical benchmark, the equilibrium computed assuming risk neutrality tells us how rational agents would behave, regardless of their global risk attitudes. Any other benchmark would have to rely on some variant of non-expected utility theory.

Loss aversion and uncertainty aversion. Finally, it is worth considering

whether loss or uncertainty aversion can explain some of our results. Should loss or uncertainty aversion matter in some way, however, it would lead towards the risk-dominant strategy and hence towards under-investment—a point that may also be made if attitudes towards risk should matter. In the experiment, on the contrary, however, the main result is over-investment by agents with good experiences.

Thus, we are left to conclude that none of the alternative explanations considered here suffice to give as stark a behavioral pattern as does the supposition that subjects display a solipsism bias. That is, subjects regularly fail to appreciate how the experiences of others that differ from their own experiences will affect others' actions, and, thus, affect the expected payoffs in the game. In other words, subjects with good experiences make their investment decisions as though their own experiences are the only ones that matter.

6 Conclusion

In Jeitschko and Taylor's model economic collapse can be inevitable—even under the best of circumstances, in the best Pareto efficient equilibrium, even in a booming economy where all agents experienced good returns on all investments all the time, a collapse occurs. Agents quit investing not just because of their own bad experiences but also because they are afraid that others might have had bad experiences or even because they are afraid that others are afraid that others may have had bad experiences, and so on—as this happens because agents account for others' (unobserved yet possible) experiences, a collapse can occur even when nobody actually has a bad experience.

In this paper we investigate whether features of sudden meltdown are observed in the laboratory. The answer is that they are not. Subjects are persistently optimistic as long as their own experiences are good. Our analysis suggests that this behavior is best explained by noting subjects'

inability to acknowledge alternative—contradictory—experiences from their own positive experiences. We term this ‘solipsism bias.’ Solipsism bias explains the observed behavior by having the effect of raising confidence and preventing economic breakdown. In contrast, alternative explanations, such as trembles, failures in Bayesian updating, or risk, loss, or uncertainty aversion on part of the subjects, are shown not to be able to account for subjects’ behaviors.

References

- [1] Banerjee, Abhijit. “A Simple Model of Herd Behavior,” *Quarterly Journal of Economics*, 1992, 107 (3), 787-818.
- [2] Bikhchandani, Sushil, David Hirshleifer, and Ivo Welch. “A Theory of Fads, Fashion, Custom and Cultural Change As Informational Cascades,” *Journal of Political Economy*, 1992, 100 (5), 992-1027.
- [3] Camerer, Colin F., Teck-Hua Ho, and Juin-Kuan Chong. “A Cognitive Hierarchy Model of Games,” *Quarterly Journal of Economics*, 2004, 119 (3), 861-898.
- [4] Carroll, John S., Max H. Bazerman, and R. Maury. “Negotiator Cognitions: A Descriptive Approach to Negotiators’ Understanding of Their Opponents,” *Organizational Behavior and Human Decision Processes*, 1988, 41, 352-370.
- [5] Cooper, David J., and John Kagel. “Learning and Transfer in Signaling Games,” *mimeo*, 2004, Ohio State University.
- [6] Jeitschko, Thomas D., and Curtis Taylor. “Local Discouragement and Global Collapse: A Theory of Coordination Avalanches,” *American Economic Review*, 2001, 91 (1), 208-224.

- [7] Kübler, Dorothea, and Georg Weizsäcker. “Limited Depth of Reasoning and Failure of Cascade Formation in the Laboratory,” *Review of Economic Studies*, 2004, 71(2), 425-44.
- [8] Nagel, Rosemarie. “Unraveling in Guessing Games: An Experimental Study,” *American Economic Review*, 1995, 85 (5), 1313-1326.
- [9] Nyarko, Yao, and Andrew Schotter. “An Experimental Study of Belief Learning Using Elicited Beliefs,” *Econometrica*, 2002, 70, pp. 971-1005.
- [10] Rabin, Matthew. “Risk Aversion and Expected-Utility Theory: A Calibration Theorem,” *Econometrica*, 2000, 68(5), 1281-1292.
- [11] Selten, Reinhard. “Reexamination of the Perfectness Concept for Equilibrium Points in Extensive Games,” *International Journal of Game Theory*, 1975, 4 (1), 25-55.
- [12] Sonnemans, Joep, and Theo Offerman. “Is the Quadratic Scoring Rule Really Incentive Compatible?,” *mimeo*, 2001, University of Amsterdam.
- [13] Sovik, Ylva. “Strength of Dominance and Depths of Reasoning,” *mimeo*, 2004, Ministry of Trade and Industry, Norway.
- [14] van Huyck, John B., Raymond C. Battalio, and Richard O. Beil. “Strategic Uncertainty, Equilibrium Selection, and Coordination Failure in Average Opinion Games,” *Quarterly Journal of Economics*, 1991, 106 (3), 885-910.
- [15] van Huyck, John B., Joseph P. Cook, and Raymond C. Battalio. “Adaptive Behavior and Coordination Failure,” *Journal of Economic Behavior and Organization*, 1997, 32 (4), 483-503.

Aggregate Data C	Period 1	Period 2	Period 3	Period 4	Period 5
PEPBE					
Investment Rates	1	.90	0	0	0
Observed					
Investment Rates	1	0.94	0.87	0.77	0.66
N	600	600	600	600	600

Table 11: PEPBE and observed investment rates in Treatment B.

Only Successes C	Period 2	Period 3	Period 4	Period 5
PEPBE				
Investment Rates	1	0	N/A	N/A
Observed				
Investment Rates	.94	.97	.98	.94
Observations	540	440	361	290

Table 12: PEPBE and observed investment rates for subjects who have only experienced successes in Treatment B.

APPENDIX A: RESULTS FOR TREATMENT C

In Treatment C , we utilized the devices discussed above to help subjects to coordinate: in particular, we forced subjects to invest in the first period and did not allow them to invest again in the same round after deciding not to invest.

Table 11 shows aggregate investment rates over time for this treatment. Tables 12 and 13 show the investment rates for subjects who had only successes and for subjects who experienced a failure in the preceding period.¹⁶

Our devices were very successful in achieving coordination. The coordination in the first period has a strong effect for the later periods. As in Treatment B , once coordinated, subjects keep coordinating. The decline in investment rates mirrors the decline observed in the other two treatments, in that it is steady and slow. Moreover, there is no sudden drop-off in the

¹⁶As with the other treatments, also in this treatment subjects reacted to a failure in the previous period in a similar way to that shown in the table, independently of what happened in the history until $t-2$. For instance, if we consider histories in which subjects had a failure only at time $t-1$, the investment rates are 0.87, 0.69, 0.57 and 0.63.

Previous Failure C	Period 2	Period 3	Period 4	Period 5
PEPBE				
Investment Rates	0	0	N/A	N/A
Observed				
Investment Rates	0.87	0.64	0.53	0.60
Observations	60	84	101	98

Table 13: PEPBE and observed investment rates for subjects who have experienced a failure in the previous period in Treatment B.

Only Successes C (last 5 rounds)	Period 2	Period 3	Period 4	Period 5
Investment Rate	0.97	0.94	0.99	1
Observations	179	141	112	88

Table 14: Observed investment rates for subjects who have only experienced successes in the last five rounds of the two treatments.

investment rate in Period 3. Subjects seem to apply the same decision rules as in the other treatments. They keep investing as long as investments are successful. Once discouraged by a failure, the investment rate drops to numbers slightly higher than in treatments *A* and *B*.

Analyzing the last few rounds separately, we see that behavior is, again, remarkably stable over time. Subjects' actions are consistent with the simple behavioral rule throughout the experiment. Even at the end of the experiment, subjects in both treatments invested almost always after having experienced successes; and significantly reduced their propensity to invest after a failure.¹⁷

As with the other treatments, we also looked at individual decisions. Behavior is very homogeneous after successful experiences: More than 80% of subjects re-invest in more than 80% of all cases after a success. After

¹⁷Notice that it is slightly odd that 87% of subjects keep investing after experiencing a failure in the first period, where the correct inference should now be that the state is bad and investment does not pay regardless of what the others do. However, we are not too concerned with this, since play conditioned on one's previous experience is remarkably stable across treatments, especially after experiencing success.

Previous Failure C (last 5 rounds)	Period 2	Period 3	Period 4	Period 5
Investment Rate	0.86	0.60	0.44	0.58
Observations	39	56	46	49

Table 15: Observed investment rates for subjects who have experienced a failure in the previous period.

Percentage of Investments after a Success	Treatment C
[0, 20)	0.025
[20,40)	0.025
[40,60)	0.025
[60,80)	0.075
[80,100]	0.85

Table 16: Individual decisions: subjects are divided according to the percentage of the time in which they invested after a success in the previous period.

a failure, behavior is more heterogeneous. While there is hardly anybody with an investment rate of less than 20%, the distribution of investment rates above 20% is almost uniform—with a slight bias towards optimistic behavior.

Percentage of Investments after a Failure	Treatment C
[0, 20)	0.075
[20,40)	0.15
[40,60)	0.175
[60,80)	0.2
[80,100]	0.4

Table 17: Individual decisions: subjects are divided according to the percentage of the time in which they invested after a failure in the previous period.

APPENDIX B: INSTRUCTIONS FOR TREATMENT A

Welcome to our experiment!

Please be quiet during the entire experiment. Do not talk to your neighbours and do not try to look at their screens. Simply concentrate on what you have to do.

If you have a question, please raise your hand. We will come to you and answer it privately.

You are participating in an economics experiment in which you interact with seven other participants for 15 rounds. Depending on your choices, the other participants' choices and some luck you can earn a considerable amount of money. You will receive the money immediately after the experiment. Notice that all participants have the same instructions.

The experiment

What you have to do

The experiment consists of a series of 15 rounds. In each round you have to make five decisions. More specifically, you will be asked 5 times in sequence if you want to make an investment or not. If you make an investment, you will have to pay a cost of £1.64. The investment can be

either successful or not. If it is successful, you can earn £2. After each investment decision, you will be informed about your payoff.

What determines whether the investment is successful or not

Imagine two urns. Urn 1 contains only green balls. Urn 2 contains 80 green and 20 red balls. At the beginning of each round the computer will randomly choose one of the two urns by tossing a coin. That is, both urns are equally likely to be chosen. However, we will not tell you which urn the computer has chosen!

Once an urn has been chosen for a round this urn will be used to determine the success of all investment decisions within this round. (Remember there are 5 such decisions in each round.) After each investment decision the computer will randomly draw one ball of the chosen urn for each participant. If you are assigned a green ball, your investment is successful, if it is red it is unsuccessful.

Notice that if the computer initially chooses urn 1, the investment will always be successful. On the other hand, if it chooses urn 2, the investment will, on average, be only successful 80 times out of 100. Notice that there will be separate draws for each of your investment decisions and for each participant. That is, it is possible that an investment is successful for you and unsuccessful for another participant or vice versa.

How do you pay for an investment?

We give you £1.70 each time you have an investment opportunity. You may use this money to pay for the cost of the investment or you can just keep it if you decide not to invest.

What do you earn if you decide to invest?

If the investment is unsuccessful, its value is zero — regardless of what other subjects decide to do. Therefore, in this case, if you decided to invest, you lose £1.64 from the £1.70 you received from us. And you keep 6 pence.

If, on the contrary, the investment is successful, its value depends on the decisions of other subjects. After you have made your decision (invest or not) the computer will randomly match you with one of the other 7 subjects in the experiment. If this subject has decided to invest too, then you will receive £2. But if the other participant has decided not to invest, you will receive nothing. Therefore, in the former case your total payoff from this investment decision is what you have kept from the money we gave you $(1.70 - 1.64) = 6$ pence plus the pound you get from the successful investment. So, your total would be £2.06. In the latter case (where you have been matched with somebody who didn't invest) your payoff is $(1.70 - 1.64) + 0 = 6$ pence. In other words, you will earn money by investing if the investment is successful for you and if the other participant matched with you has also decided to invest. Notice that although you can lose £1.64 by making an investment, your payoff for each investment will never be negative, as we give you £1.70 for each investment opportunity.

What do you earn if you decide not to invest?

If you decide not to invest, you will neither earn nor lose anything. You will just keep your £1.70.

Procedures for each round

Remember that the experiment is organized into different rounds and that within each round you will have to make five investment decisions. So, let us summarize what happens within each round.

1) At the beginning of each round the computer randomly chooses one of the two urns, urn 1 or urn 2. The chosen urn will be used for all participants and all investment decisions in this particular round. But you will not be told which urn has been chosen.

2) Now you make your first decision: either invest or not.

3) If you have invested in 2), the computer draws a ball from the urn that was chosen in stage 1). The ball is drawn and then replaced, so that the total number of balls in the urn is always the same. The colour of the ball that was chosen for you determines whether your investment is successful or not. If it is green: success; if red: failure.

4) You will be randomly matched with another participant. If you haven't invested this is irrelevant for you. If you have invested and if your investment was successful in 3), your payoff depends on whether the other participant you have been matched with has also invested. If he has, you earn £2, if he hasn't you get nothing.

5) You will be informed about what has happened to your investment, i.e., about your payoff.

6) Now you make your second investment decision. Notice that the same urn that has determined your success previously will be used again. But a new ball will be drawn from that urn.

In other words, from 6) onwards everything is the same as from 2) on. This will be repeated 5 times. Your total payoff from the 5 decisions is just the sum of all payoffs you earned for each decision.

Once the first round is over, you will be informed of what urn the computer chose, the colour of the balls that it assigned to you, the decisions of the other participants that you were matched with and your payoff. Then,

we will repeat the same procedure for the second round at the beginning of which the computer will choose again one of the two urns. And we will repeat the same procedures for the whole sequence of 5 investment decisions.

One more question

Whenever you have to decide whether to invest or not, we will also ask you another question: How much, at most, would you pay for the investment?

You know that the true cost of the investment is £1.64. Therefore, whenever you decide to invest, it's clear that you think the investment is worth an amount bigger than £1.64. What is this amount? That is, up to which cost would you be willing to pay for the investment? Similarly, whenever you decide not to invest, clearly you believe that the investment is worth an amount lower than £1.64. Again we ask you what this amount is. What is the highest cost for which you would be willing to invest?

Note that if you are successful the investment is worth £2, as this is the amount that we pay you. So if you were sure of obtaining success, clearly the investment would always be worth £2. On the other hand, if you were sure of being unsuccessful, then clearly the investment would be worth 0, as in case of failure we pay you nothing. Therefore, the true value of the investment cannot be more than £2 and cannot be less than £0. Given that you do not know with certainty what the true situation is, the value of the investment has to be between 0 and 2. For this reason, on your screen you can select any number between £0.00 and £2.00.

Examples of per round payoff

Example 1

Suppose that for this round the computer has chosen urn 1. Therefore all balls are green. Suppose you make the following decisions:

INVEST, INVEST, INVEST, NO, NO

Suppose, finally, that in the first two times you are matched with someone who invested, but in the third you are matched with someone who did not invest. Your payoffs would be:

$$206+206+6+170+170=758.$$

In this round, you earn money with the first 2 investments but you lose money with the third, despite all balls were green, because your match decided not to invest. The fourth and fifth times you just kept your 1.70 pence.

Example 2

Suppose for this round the computer has chosen urn 2. Therefore 80 out of 100 balls are green. Suppose you make the following decisions:

INVEST, INVEST, INVEST, INVEST, NO

Suppose, finally, that you received a green ball in the first 3 cases but not in the fourth. Suppose, finally, that you were always matched with someone who invested. Your payoffs would be:

$$206+206+206+6+170=794.$$

In this round, you earn money with the first 3 investments but you lose money with the fourth, because you got the red ball, i.e., your investment was not successful. The fifth time you just kept your 1.70 pence.

Final payment

For the simple fact that you showed up in time for the experiment you earn £4. The rest of the payment depends on how you perform. The computer will randomly choose one round out of the first 5 rounds, one among the 6th through the 10th and one among the 11th through the 15th. Your payment will depend on how you performed in the selected rounds. We will sum up your payoffs in these three rounds and divide the sum by 3. Your final payment will be equal to this amount plus the £4 for showing up.