

THE RIGHT CHOICE AT THE RIGHT TIME: A HERDING EXPERIMENT IN ENDOGENOUS TIME

DANIEL SGROI

Abstract. This paper examines experimental evidence relating to herd behaviour in situations when subjects can learn from each other, and can delay their decision. Subjects acted rationally, gaining from observational learning, despite penalties for delay. Cascades were ubiquitous and reverse-cascades occurred in which incorrect decisions made by early decision-makers produced herds on the incorrect choice. The major departure from rationality came when subjects realized they had chosen incorrectly despite following the majority view. This led many to add extra delay to future decision-making. It is argued that this may be due to certain cognitive biases, and is likely to make matters worse, making it all the more important that policy-makers attempt to minimize the chance of reverse-cascades.

1. Introduction

Observing the decisions of others can often be helpful when making a choice. In fact, optimal behaviour often requires a careful analysis of others' actions. There will be times when others' behaviour seems to contradict an individual's own private information and it is here that an information cascade may result. This can occur when early information is such that some individuals opt for a particular action which is enough to convince later movers to disregard their own information. A great deal therefore depends upon the information possessed by early movers in such situations. This process of social learning in sequence was the basis of much of the early herding literature initiated by Banerjee (1992) and Bikhchandani et al (1992) for discrete action spaces and Lee (1993) for more general action spaces. This paper instead focuses on the alternative case when individuals can try to rectify the difference between their own information and the observed actions of others by waiting. The actions are now no longer exogenous and sequential instead they become endogenous to the model itself. This paper will present new experimental work relating to such endogenous-time herding models pioneered by Chamley and Gale (1994) for discrete time and Gul and Lundholm (1995) for continuous time.

The primary goal of the experiment is to test for the practical application of existing endogenous-time models, which allows players to receive their private signals and wait for as

Date: 24th March, 2000.

Key words and phrases. learning, herding, delay, experiments. J.E.L. Classification Codes: C91, D82, D83.

The author is grateful to the E.S.R.C. and Nuffield College for financial support. The author would like to thank Chris Avery, Michael Bacharach, Alan Beggs, Paul Klemperer, Sophocles Mavroudis, John Morgan and Daniel Zizzo for useful comments and Marc Chapman, David Gill and Rosemary Sgroi for their role as laboratory assistants during the experiments.

long as they wish. In this way players have to counterbalance their desire to observe the actions of others, infer further useful information, and then make a more informed choice against the discounting of payoffs which occurs over time. At a more basic level, as with most experiments, the central notion is (Bayesian) rationality. Here this comes through examining whether people respond appropriately to private signals and publicly observed actions.

This paper continues with a series of brief literature reviews in Section 2, listing some major theoretical and experimental herding papers. Section 3 details the experimental design, which following Holt and Anderson (1997) is as simple as possible. Section 4 examines optimal behaviour within this design and gives the optimal strategies as a function of the initial signal and observed actions. Section 5 presents some analysis of the results which are in general very supportive of rationality and herding. In all cases decisions were made quickly with delay very close to suggested optimal levels and despite rationality, the occasional incorrect herd (or reverse cascade) could not be avoided. Section 6 introduces an alternative condition: the subjects are now informed what the correct action was after each game is played. This should make no difference if all subjects behave rationally. In a reverse cascade many subjects disregard their own good information to follow those with bad information, and it is just after such an event that revealing the correct choice is the most devastating. Having discovered that an incorrect choice was made despite following the prevalent view many subjects respond by delaying decisions in future games and moving away from rationality. It seems here that suboptimal actions derived from rational play, can seed irrationality in the minds of players, possibly in response to the failure of rational play to produce success. In another sense certain well researched cognitive biases might be seen to be at work as subjects, having recently witnessed an unlikely event, exaggerate its importance for future decision-making. Section 7 concludes by stressing that the paper does support the general rationality of subjects while still producing herding and in some cases even herding on the wrong action, and also by focusing on the ending in Section 6.

2. The Herding Literature

Let us begin with a simple example of a herd. Consider two restaurants, A and B. Now consider a stream of agents arriving sequentially at the doors of the two restaurants. They all have a private signal (perhaps a newspaper review or information about the restaurants passed on from a friend). They consider all private information to be of the same quality - their own signal is not a priori any better than anyone else's. They can add to the information contained within their signal by observing the action (not the signal) of their predecessors. The first agent has only his (informative) signal to guide him, so will go where his signal indicates, say into restaurant A. The next agent has both his private signal and also the public information relating to the action of the first agent, which in this case perfectly reveals his signal - it clearly indicated that restaurant A was superior to restaurant B. Let us assume that agent 2 also has a signal favouring restaurant A. He will therefore also enter restaurant A. The third agent arrives, observes the entrance of the first two agents. Let us consider what happens if his signal suggests that restaurant B is superior. He has observed the actions of his two predecessors, and can infer that the first mover had a signal favouring restaurant A. The second mover is difficult: he may

also have had a signal favouring restaurant A, which would account for his actions. However, he may have had a signal favouring restaurant B, but this would render his net information neutral - so he might still go for restaurant A. To deal with this we assume that a player will go for his own information if his net information is neutral. This was the assumption in Banerjee, and is also justified by the experimental evidence in Anderson and Holt and works against any herding by making players favour their own information. The third player can now infer two signals suggesting restaurant A is superior. This swamps his own signal and he will (rationally) go for restaurant A. He is the first to be trapped in what Bikhchandani et al call an informational cascade. With no new information, agent 4 will of course now also enter restaurant A even if his signal suggests that he should not, as will agent 5, and so on, all agents herding into restaurant A.

It is of course perfectly possible that restaurant B is in fact superior. Consider, for example, a criterion for superiority based on the number of positive reviews. It might be that restaurant A was considered superior by 40% of reviewers and restaurant B by 60%, but a predominance of those who had read reviews by those favouring restaurant A early on might well produce a cascade in its favour. Note that there is a huge loss of information which would eventually reveal the superiority of restaurant B if signals and not actions were observable. Banerjee named this the herd externality.

This example involves a discrete (binary) action space, sequential ordering and no opportunity to delay decision-making. With no opportunity to delay, discount factors play no role. The next subsection looks at the literature which grew out of the simple ideas evident within this example.

2.1. Theoretical Literature. The herding literature in its present form dates back to Banerjee and Bikhchandani et al who first developed simple models of rational herding (or informational cascades) demonstrating that it would often be in an individual's best interest to disregard his own information and join a clustering of previous actions. In this sequential and de facto discrete framework they both showed that social suboptima could result with agents herding into an action and therefore never revealing their informative signals for later agents to use. Banerjee did use a continuous probability distribution, but its degenerate nature effectively made the problem a discrete one. The restaurant example is a particularly simple example of this form of discrete action-space sequential herding.

Lee considered more general modelling specifications looking at different action spaces and concluding that much would depend upon whether it was possible to map signals into the action space one-to-one so that signals reveal actions. This would clearly produce something closer to a normal learning model and we would expect later agents to get the decision correct in the sense of opting for the superior choice.

This type of decision-making in which agents are rational and can learn from the decisions of others was then taken into the endogenous-time world by Chamley and Gale who considered investment decisions in which agents could delay to observe the actions of others, so the ordering became endogenous to the model. They used a discrete-time model and found effects similar

to those in Banerjee in that there was clustering on a particular action and also at a particular time, and furthermore the amount of delay together with the potential for errors to be made even by fully rational Bayesian agents produces social suboptima. Gul and Lundholm developed a similar model, but instead opted for continuous time and found that the timing of decision making in a richer action space would perfectly reveal signals and therefore ensure that the optimal choice was made with minimal delay. Just as with the sequential models the richness of the action space seems to be all-important in determining whether optimal results ensue.

2.2. Experimental Literature. Very little experimental work has been published on the herding phenomenon. Anderson and Holt produce results showing that sequential herding does seem to occur even with few subjects in the laboratory. They suggest a very simple experiment design which is worth detailing here as it will form the basis of the more complex experimental design in the next section.

Consider two urns, A and B, each containing three balls. Urn A contains two red and one white ball, and urn B contains two white and one red ball. The contents of one urn (randomized with odds 50:50) are emptied into a container. This process is not seen by test subjects though they are fully aware of the original contents of the urns and that with 50:50 odds one urn has been emptied into the container. Next we allow a sequence of test subjects to arrive at the container and select one ball from the container, note its colour and replace it. After selecting a ball they must predict which urn was used, winning a prize for a correct answer. A red ball is suggestive of urn A and a white ball of urn B, but neither is completely revealing. All signals are therefore of the same quality. Once subjects have made their decision it is noted on a board which is clearly visible by all the subjects in the room. Anderson and Holt found that herding occurred consistently in the laboratory where other sociological incentives to go along with the crowd can be controlled. Some decision sequences resulted in reverse cascades where initially misrepresentative signals started chains of incorrect decisions not broken by more representative signals gained later. Cascades were roughly split between reverse and normal cascades. In 12 sessions cascades formed in 87 periods of 122 in which they were possible. Individuals generally used information efficiently and followed the decisions of others when it was rational. They did find that there were errors which tended to make subjects rely more on their own private signals. Anderson and Holt felt they could explain this by factoring in the positive probability of an error in decision-making. Subjects would then slightly favour their own signals to the possibly erroneous decisions of others. They also found that the most prevalent systematic bias was the tendency for some of their subjects to rely on the simple counting of signals rather than the use of Bayes' rule for updating where these implied different decisions. The main conclusion to be taken from Anderson and Holt is that the most basic herding literature does seem to have some predictive power in the laboratory, and that this justifies further experimental work looking at the later theoretical work in herding.

The Anderson and Holt experiment was based entirely on sequential discrete models in the style of Banerjee and Bikhchandani et al. As a result they did not consider endogenous-time or discounting, which are examined in this paper.

3. Experiment Design

The experiment was designed to be as simple as possible and yet still capture the main themes of endogenous-timing in herding models. It was essential to avoid any form of collusion, and to this end throughout the experiment, no form of communication between subjects was allowed. Appendix A gives the full experimental text. 54 subjects took part, each playing four games, for a grand total of 216 different games. The data is fully described in Section 5.

3.1. First Step: Signals. There are two urns, one red and one white. The red urn contains two red and one white ball and the white urn contains two white and one red ball and this is known to the test subjects. The contents of one urn is emptied into a container, the probability of this being the red or the white urn standing at 50%. Again, this is known to the test subjects. Subjects win a prize for correctly guessing which urn was emptied. The subjects arrive in sequence at the container and draw two balls, one at a time, replacing the ball after each draw. This gives them a signal as to which urn was emptied, a red ball suggesting the red urn, a white ball suggesting the white urn. Since they have two draws they have signals of different quality. There are three sorts of signal: a strong signal in favour of the red urn (two red); a strong signal in favour of the white urn (two white); and an indifferent signal conveying no information (one red, one white). Subjects then return to their seats, and without communication wait until all other subjects have drawn their signals.

3.2. Second Step: The Prize. Each subject is told that he or she will be awarded a prize of \$Z for guessing correctly which urn was emptied. They are also told that they may wait as long as they wish before coming to a conclusion, but that for every minute waited their potential prize will fall steadily over time. To make all this clear the subjects are given a table which allows them to check easily what their payoff will be at every time interval. Ten seconds before each minute elapses the time is announced. A large timer is clearly visible throughout the experiment. When everyone is ready, and the rules are thoroughly understood, the timer is started.

3.3. Third Step: Decision Making. The experiment then enters the decision-making step. To capture the notion of discrete time the subjects each receive private forms upon which to make their choices and at the end of each period (a minute in length) the game pauses and the form is viewed by the experimenter or an assistant. The choices are: red, white or wait. After noting any positive choices (red or white) on a board clearly visible to all, the next period is initiated and play continues, pausing again at the end of the second period, and so on. After making a choice other than wait subjects leave the main experiment room. When all have made a decision other than wait all subjects return to the main experiment room and the correct urn is revealed. The game definitely ends after 15 periods and this is made clear to the subjects, so if waiting continues to period 14 all will know that they must decide red or white in period 15 or receive no prize. At the end of this stage the debriefing will start and payments will be made to them based upon their performance following the scheme. Each game gives a certain \$1.50 and a prize of between \$0 and \$3.50. In practice 4 games were played in a period of

between 90 minutes and 2 hours. This gives an average expected payment of about \$8 (about \$12) per hour.

4. Optimal Behaviour

This section looks at what the optimal behaviour resulting from the experimental design might be expected to look like.

The notation XY is used to denote the two draws by a subject where X and Y can be R (red) or W (white). The italicized red and white refer to the red urn and white urn respectively. Consider any given signal actually representing the probability that the correct urn is red. In this case we have three signals: RR is an 80% signal; RW is a 50% signal; and WW is a 20% signal. All signals are now on the interval $(0; 1)$. Define a signal strength s^i_R such that if $s^i \geq s^i_R$ a subject i will definitely go for red in the first period. Similarly define a signal strength s^i_W such that a subject i will choose white. The area in the middle is a zone of uncertainty such that our subject will wait. Our task is to find out if the available signal strengths produce clear cut optimal actions in these regions.

Consider choosing red in the first period on the basis of a WR signal. This signal is strong, providing an expected payoff of \$4. Opting for white would provide an expected payoff of \$1. Waiting is a more complex case. If you wait you are likely to obtain more information. Even if the other subjects do nothing, they still reveal that they do not have strong enough signals to choose red or white in the first period. For example, waiting until the second time period and observing all 8 other subjects going white. This would clearly imply that the 8 other subjects did not have RR signals, and our subject might regret selecting red.

Start with the assumption that movement on the first period requires a strong signal, so we have the following equilibrium mapping:

Signal RR \rightarrow Action red; Signal WW \rightarrow Action white; Signal WR or RW \rightarrow Action wait

We need to check that a subject would be correct to follow this course of action given that the other subjects also do so (a Nash check). We remove the certain payoff of \$1.50 per game, since this is guaranteed and consider the bonus for a correct guess. The cost of delaying an action given the observed signal RR is \$0.20, that is the difference between choosing red now and red next period, multiplied by the probability that red is actually the correct choice. This probability is 0.8 given an observed signal of RR . Now we must look at the benefit of delay to determine if the cost exceeds the benefit of delay when the initial signal is RR . We will go through one calculation, before looking at the general calculation. Having elected to choose red, the subject now gets to observe the first period decisions of the other 8 players, and there are 45 different possibilities. Of these 45 certain combinations will induce regret about having chosen red. For example, if he should observe all 8 other subjects selecting white. From the hypothesized equilibrium this implies a total set of inferable draws of 16 white balls and 2 red balls, that is a net 14 white balls. This provides a probability in favour of the urn being white of $\frac{2^{14}}{2^{14}+1}$, with only a probability of $\frac{1}{2^{14}+1}$ in favour of red. Therefore subject i will now see his

expected payoff bonus shrink from $0.8 \times \$3.50 = \2.80 , to $\$0.000214$. He should certainly feel considerable regret. The benefit of delay in this case is just under $\$3.25$. However, he will not expect to see all 8 other players go for white! In fact, if he receives an initial RR signal he will expect to see very few people acting as if they had WW signals. In terms of real options theory, this notion of regret is none other than the real option value of the decision to wait, destroyed through choosing positive action. This is similar to the real option in an optimal stopping problem: for more on this, see Dixit (1993).

Appendix B calculates the general benefit of delay incorporating this option value and shows that for a subject with a RR or WW signal the cost of delay is $20p$, whereas the benefit of delay is only $0.5p$, therefore the subject should choose red or white respectively in the first period. By contrast, Appendix B also shows that the cost of delay for a subject with an initially neutral signal (RW or WR) is $12.5p$ whereas the benefit of delay is $38.5p$. Since this exceeds the cost of delay, the subject with a neutral signal should wait. Therefore, we see that the candidate equilibrium is actually reasonable, since the full mapping from signal to strategy is consistent with optimal behaviour. With this candidate equilibrium shown to form an optimal set of (Nash) actions, we can say immediately that all subjects with strong signals should move immediately and all subjects with neutral signals should wait one period. Having observed all the strong signals after one period all the subjects with weak signals should then make their decision in the second period to avoid further discounting, since there is no further benefit to waiting.

To summarize, optimal expected actions are: all with RR choose red in the first period; all with WW choose white in the first period; and all with neutral signals of R and W should wait one period, then select based on the majority of first-period choices. In a tie they should simply randomize in the second period. There is no benefit from waiting any further since all the useful information has been revealed.

5. Results

This section evaluates the raw data to be found in tabular form in Appendix C. Glancing at the data a number of points stand out:

1. The games ended quickly, in two or at most three periods.
2. Almost all players waited at least one period when their signals were not very conclusive.
3. Initial movement was almost always based on strong signals.
4. In almost all cases what occurred looked like herding.
5. Most "herds" were on the right action, though one was on the wrong action.

Based on the calculated optimal strategies in Section 4 we can assign optimal actions to the various observed signals in the experiment.

To address one immediate concern, it is fairly clear from the optimal behaviour calculations in Appendix B that the problem is not a trivial one. Drawing RR from the urn and choosing sensibly involves finding the cost of delay and then comparing this with the likelihood of seeing any useful information if you wait. The size of the calculation indicates that it might be unreasonable to expect a subject to work this out without a calculator in a single minute. However,

the scale of the cost-benefit differential means that precise measures are not necessary. Since 20p is 40 times larger than 0:5p it should be fairly obvious to anyone with a basic mathematical ability that moving immediately is the right thing to do if you have a strong signal. Since 38p is three times larger than 12:5p it should also be reasonable to expect subjects to wait if they receive a mixed signal. In fact, even using a simple rule of thumb, it is fairly clear that a completely uninformative signal should not produce immediate action with such a mild discount rate. Since the optimal actions are so clear cut we could reasonably hope for optimal behaviour despite the complexity of the calculations required.

This section splits the optimal actions into two benchmarks. Initially we consider a thought experiment in which we simply forced all subjects to play according to the optimal strategy set given their signals. This generates a set of ex ante optimal actions and a corresponding set of payoffs. Secondly, we consider each player's actions and compare these with the expected optimal actions given the observed actions of others. In this way if one player has deviated a second player might choose a rational option but this might diverge from the ex ante set of optimal actions when we assumed no such deviation. In this sense we have a set of ex post optimal actions to compare with observed actions. We can also calculate the payoffs associated with the ex ante and ex post optimal actions and the observed payoffs.

5.1. Ex Ante Predicted Choices. Here we concentrate on the actions that would be taken were all subjects to act according to the optimal strategies in Section 4. It is interesting to note that in only one case did a subject exceed the expected ex ante optimal payoff, but here one of the optimal rules involves randomization and so there is a 50% chance that a subject following the optimal strategy would have done better. Table 1 summarizes the actual payoffs compared with the payoffs that would have ensued if all subjects behaved optimally.

Table 1: Average Subject Payments, \$

	Day 1	Day 2	Day 3	Day 4	Day 5	Day 6	Overall
Observed Actual ¼	14.23	18.69	16.06	17.56	16.00	16.11	16.44
Ex Ante Optimal ¼	15.44	18.86	18.00	17.64	16.89	16.35	17.20
Absolute % difference	7.84	0.90	10.78	0.45	5.27	1.47	4.42

The average ex ante optimal payoff for day 1 was \$15:44, which is a little higher than the actually observed average of \$14:13. On day 2 the optimal average would have been \$18:86, whereas the observed actual average was \$18:69. For day 3 the observed figure was \$16:06 and the ex ante optimal figure was \$18, and similarly for days 4, 5 and 6 observed averages were slightly below the optimal figure. Overall, the average ex ante optimal payoff would have been \$17:20, whereas the observed average was a little lower at \$16:44. The percentage difference is only 4:4%.

Table 2: Percentage of Predicted Choices

	Day 1	Day 2	Day 3	Day 4	Day 5	Day 6	Overall
Colour	81	100	86	100	94	97	93
Time	89	86	89	92	89	94	90
Time (+1 period) ¹	94	100	97	100	100	100	99
Time (+2 periods) ¹	100	100	100	100	100	100	100

Note: ¹ +X period(s) refers to a choice of time within X period(s) of the predicted time.

In order to examine how close the subjects got to the predicted ex ante behaviour we split the data into two parts, the correct colour choice and the correct time of action. Within the optimal action sets there exist randomizations, at points of indifference. It is clearly difficult to determine whether subjects used randomizations, so with this in mind, any colour choice at a point of indifference is taken to be optimal (since either colour would be optimal). Details are given in Table 2. What we find is that overall 93% colour choices were ex ante optimal reactions to signals. In terms of time choices, 90% of decisions were made when predicted, 9% of decisions were delayed by one extra period and only 1% of decisions were delayed by an extra two periods, with no delays beyond two extra periods.

What we have is a picture of day 1 as departing mildly from ex ante optimality in terms of colour choice, and having a close to the ex ante optimal pattern of timing decisions. Day 2 looks much like optimality in terms of colour choice and in terms of timing. This seems to provide some evidence for the ability of subjects to choose correctly but a little more slowly than if fully optimal. More than 90% of decisions made by subjects were exactly as predicted by time and colour which seems very high given the complexity of the decision. This rises to 93% when we add a lag of a single period.

5.2. Rational Responses. We no longer impose rationality on all subjects to find our benchmark. Instead we require that the subjects respond rationally given the assumption that all others are rational. This is perhaps the more natural way to examine the data. For example, given the early decision by some with good signals to wait in game 2 of day 1 and the resulting decision to go for the wrong choice by those with neutral signals, the behaviour of later movers within a reverse-cascade is not an irrational phenomenon. We would still expect a strong signal (RR or WW) to result in optimal first period movement in the relevant colour. However, a neutral signal (WR or RW) would result in waiting until the second period and then basing a decision on the majority choice in the first period. This would factor in observed behaviour, rather than ex ante optimal behaviour, and help solve the problem of one-off irrational choices rendering the decisions of all later movers sub-optimal, when they are actually responding rationally to this deviation. In this sense we are looking at ex post optimality. Appendix C also provides the actions which rational agents should have played in response to the observed actions of the other subjects, while this subsection examines some aggregate findings based a comparison of the raw data with the ex post optimal actions.

Table 3: Average Subject Payments, \$

	Day 1	Day 2	Day 3	Day 4	Day 5	Day 6	Overall
Observed Actual $\frac{1}{4}$	14.23	18.69	16.06	17.56	16.00	16.11	16.44
Ex Post Optimal $\frac{1}{4}$	14.51	18.86	16.92	16.91	16.17	15.99	16.56
Absolute % di®erence	1.96	0.90	5.06	(-)3.84	1.03	(-)0.77	2.26

As before we can compare the optimal payo[®], with our alternative measure of optimality, to the observed payo[®], and this is done in Table 3. Since we are considering a weaker form of optimality as expected the percentage di®erence is on average lower standing at only 2:26%. this ¯gure is an average of the absolute percentage di®erence, the simple average is lower at 0:72%. Note that where the optimal strategy calls for a coin ¯ip, and where that ¯ip results in the correct actions being taken, it is easily possible to exceed the expected optimal payo[®].

Table 4: Percentage of Predicted Choices

	Day 1	Day 2	Day 3	Day 4	Day 5	Day 6	Overall
Colour	92	100	97	100	100	100	98
Time	89	86	89	92	89	94	90
Time (+1 period) ¹	94	100	97	100	100	100	99
Time (+2 periods) ¹	100	100	100	100	100	100	100

Note: ¹ +X period(s) refers to a choice of time within X period(s) of the predicted time.

Table 4 suggests that the choice of colour and time of action was very close to our candidate set of rational actions. As before no one was more than 2 periods slower than they should have been and almost all were at most 1 period slower. Despite the complexity of the decision 90% of subjects behave exactly as predicted by time and colour and 98% behave as predicted with only one extra period of delay.

6. An Alternative Condition

This section details the results of making a slight change in the experimental design. For two days of the experiment the subjects were told what the correct choice was after each game was played. This alternative should make no di®erence to optimal play.

6.1. The Availability Hypothesis. Despite the assertion that the revelation of the correct option after each game should make no di®erence to optimal play, and the fact that subjects seemed to act close to optimally the minor change in experiment design did in fact have a significant change. In four days of the experiment reverse cascades were experienced, and in all cases games were played after the reverse cascades. In the games with the new condition in the incorrect choices of subjects was made clear there seems to have been a dramatic result - directly after a reverse cascade, subjects slow their decision-making right down and even seem to trust their own signals less. While subjects were not told that a reverse cascade occurred, for many this will have been obvious on discovering that despite following the majority view of which choice was correct, they chose incorrectly. In effect the subjects seem to behave more

optimally as long as unlikely events do not take place. A reverse cascade is sufficiently unlikely to shake their belief in their own optimality. This might well be related to certain commonly observed cognitive biases, such as the availability heuristic which seems to be a common feature in experimental psychology, for example see Akerlof and Yellen (1987). In general, this naturally occurring heuristic seems to result in subjects relying too heavily on salient information which is easily retrievable from memory. This might result in subjects over-exaggerating the likelihood of unlikely events having seen them occur recently. Many who are not confident in their own ability to correctly determine the relevant probabilities might set aside their beliefs and rely on observed outcomes, putting aside probabilities in favour of frequencies. According to the representativeness heuristic, subjects act as if stereotypes are more common than they should. This might also apply in combination with the availability heuristic to exaggerate the dangers of reverse cascades.

6.2. Examining the Data. The data given in Appendix C, in Table C7 and Table C8, details the results of the experiment for the alternative condition. Table 5 below summarizes the data for days on which reverse cascades were obtained.

Table 5: Observed Behaviour Before and After Reverse Cascades

Day	Correct Choice Revealed?	Average Deviation from Optimum	
		Before reverse cascade game	After reverse cascade game
1	No	+0.11 periods	+0.06 periods
3	No	+0.11 periods	+0.11 periods
7	Yes	+0.06 periods	+0.44 periods
8	Yes	+0.06 periods	+0.39 periods

There is clearly a remarkable difference between the results for the alternative condition. There is a deviation from optimum of an extra 6-11% delay in most cases, but this rises dramatically to about 40% after subjects have been informed of their payoff just after a reverse cascade. There is no appreciable change if they are not informed, as in the standard condition. This is an interesting result and not easily explained. A reverse cascade is a low probability event with unfortunate consequences for payoffs and should subjects incorrectly exaggerate the likelihood of this event there are clearly many possible irrational responses. Should subjects react by ignoring the signals of others they would be expected to reduce delay and simply act immediately based on their private signal, this was clearly not the case. Alternatively they might respond by doubting the significance of even their own private signal which might lead to a more random pattern of actions. Glancing at the raw data, this was clearly also not the case. The actual response seems to have been simply extra delay. This might be based on a reduction in confidence in signals resulting in a desire to wait longer and attempt to gain more information from other subjects. While this is clearly not a rational action if other subjects are sticking with the optimal set of strategies and indeed is also not rational if others also wait, it is a feasible irrational response. It might then be possible to conclude that the revelation to agents that they chose incorrectly, based on entering a low probability reverse cascade might,

when combined with the commonly observed availability heuristic, result in a deviation from rational play.

It is interesting to note that when a reverse cascade occurred in game 2 on day 8, by game 4 the average deviation from optimum fell back to +0.22 periods. This is the only example of a day in which a reverse cascade occurred early enough in the sequence of four games to see any change in behaviour more than one game later and where subjects were told the correct choices between games. This suggests, though is by no means conclusive, that the deviation from rationality which occurs after realizing a reverse cascade occurred might well be a short-run phenomenon.

7. Conclusion

Some general findings should be stressed. The experimental games should have ended in two periods if the players were fully rational. They did in fact end in three periods in a few cases, but often did end in two periods. Strictly this implies less than rational behaviour, however given the complexity of the required mental calculations this is close enough to suggest something like rational behaviour. In fact 85%-90% of the time subjects did just as expected, and this figure rises to about 100% with the addition of up to 1 period of lag above the optimal time of action. It should be stressed that it was made clear to the subjects that they had up to 15 periods in which to make a decision, so there was clearly a good deal of thought involved in acting so rapidly. Virtually all subjects decided to wait a period if they received an uninformative signal. This suggests that despite the discount factor virtually all subjects appreciated the value of waiting to observe others behaviour. The results certainly seem to clearly support the theoretical literature with the single proviso that some slight extra delay might be seen in practice, perhaps when a subject fears that others may not be rational. In this sense the results here do not conflict with the results in Anderson and Holt for the simpler sequential case.

Herds are a reality. This is clearly confirmed by the data, which even produced a reverse cascade. The cascades occurred at times slightly more slowly than they should with fully optimal behaviour, but more often did occur by the second period as predicted. The reverse cascade was a particularly interesting case since it is difficult to explain this kind of phenomenon without reference to social learning. This could only reasonably happen if subjects were playing close attention to the behaviour of their peers and making a serious effort to update probabilities. This is also a clear warning that sensible observation can still result in poor payoffs.

It seems reasonable to suggest that if laboratory subjects with small prizes behave much as they should in theory then large firms or governments with so much more at stake are even more likely to learn from their peers and herd as a result. This is of course a subjective point.

An interesting additional finding comes from a close examination of the alternative condition given in Section 6. It seems to be the case that when subjects realize they have chosen incorrectly despite having followed the majority view, they respond through extra delay in the following game. This delay seems to fall back a little in any later game. This would suggest that when faced with the realization of a low probability reverse cascade subjects may respond with a higher degree of irrationality. This is clearly an important result since it suggests that not

only may very suboptimal choices be made despite rationality in agents in a herding model, but in practice, this might lead to a burst of irrationality which clearly increases the chance of further suboptimal choices being made in future games. This would seem to make it even more important that policy makers keep a close watch for the dangers of reverse cascades in situations where herds are possible.

Appendix A: Subject Instructions

\Thank you for attending today's experiment. I shall start by laying down a simple rule about silence, then detail of the experiment and then explain the prizes you can win. Feel free to ask questions. We shall then run a practice and you will have another opportunity to ask questions. When I am satisfied that everyone is ready we will begin the main experiment which will consist of what I shall call a game, repeated four times.

I would ask that you do not talk during the experiment except when I invite questions. You can raise your hand at any time to attract the attention of myself or one of my assistants who can deal with minor problems such as mislaying your pen. Otherwise please remain silent. Any attempt to do anything other than following the experimental instructions will be punished through a reduction in your prize.

I shall ask you to perform a number of tasks. I shall detail some of my aims at the end of the experiment, but leave you to consider the significance as you see fit. Much of the action will revolve around the simple observation of coloured balls and decisions you will have to make. I shall not suggest any methodology, how you make your decisions is very much up to you.

These are the devices I shall use, a green urn and six balls, three red, three white. I am now placing two red balls and one white ball into the red bag. This bag is coloured red to help you remember that it contains a majority of red balls. Similarly I am now placing the remaining balls, two white and one red, into the white bag. When the first game begins I shall randomly choose one of these bags and empty it into the urn outside this room. You will not be able to see which bag was emptied. I shall then call you in sequence and you will leave this room and enter the adjoining corridor where you will stand with your back to the urn. The urn will be covered but I shall raise the cover when I am satisfied that you cannot see the contents of the urn and your hand will be directed into the urn. You will then take a single ball from the urn and you may remember the colour of that ball. You will then drop the ball back into the urn. I shall shake the urn and you will, as before, draw once again from the urn. You will then be asked to leave the corridor, and be taken into the larger room. In the larger room you will be taken to a seat by an assistant and before you will be three items: a pen, a form and a table. The form must be completed as follows. You will see a space for your name, and your signal. By signal I mean the colour of the balls which you saw. The form also includes a list of time periods with a space next to each period in which you will be asked to write your choice. At this point you will simply wait until all the other subjects have drawn from the urn and entered the larger room. When all of you have entered the larger room I shall enter and remind you to write your name and signal on the form. I shall also remind you to keep silent and not attempt to move out of your seat until you are asked to leave. You will note that you are not able to

see any of the other subjects' forms and you should not attempt to do so. The table lists the prize you could win based upon your actions. I shall detail this later.

After I have re-stated these rules I shall initiate the first time period. You will have one full minute in which to write one of three alternatives on your form alongside the first time period listed there. You will be told 10 seconds before the minute ends and a clock is clearly visible. The timing on this clock will be final. You may write: "red", "white", or "wait". A colour indicates you believe the urn contains the contents of the respective bag (red or white). If you write wait you are indicating that you do not yet wish to make a decision and wish to try again in the next period. An assistant will come to each of you and observe your form. If a colour is written there you will be asked to leave, taking with you your form. You will hand that form to me as you leave. I shall then note the colour choices that were made on the clearly visible board. For example, if two of you choose to write a colour, one red and one white, you will be asked to leave. You will give me your forms as you leave and I shall write on the board alongside the first period, "W, R", to indicate the two colour choices. Those who write "wait" will simply stay in their seats, but may observe the board, consult their tables, re-read their forms and do whatever else they wish, in silence, until the next period is announced. When the room has been emptied of those who made their decision I shall start the timer again and as before you will have a minute to make your decision. This will continue until either all have decided or 15 minutes have elapsed. Then the game will end and I shall go with the last remaining subjects back to the small room. Those who left earlier will have been directed to the small room by an assistant and asked to wait there quietly. There will be another assistant there to ensure silence is observed.

When we all find ourselves back in the little room I shall rerun the game again. I shall once again go into the corridor and randomize which bag is emptied into the urn and we will go on as before. We will do this four times. After this I shall announce what the correct choices were, calculate the prizes you are due and for those who are interested, explain what I expected to observe.

Now I shall explain the structure of the prizes. After each game has ended I shall have a note of your four choices and when you made them. I shall also have a note of the correct choice, i.e. the bag which was emptied into the urn. This will totally determine your prize. You will get a basic prize of \$1.50 for participation in each game, and a further prize of between \$0 and \$3.50 based on your performance. If you guessed correctly in the first period you will get \$3.50. If you wait until the second period and then guess correctly you will get \$3.25. The prize will continue to fall by 25p every period. Therefore if you guess correctly in the final period, period 20, you will get only the basic \$1.50 for taking part. If you guess incorrectly at any time you will also get just the basic \$1.50 for taking part. This is all detailed on the tables in the larger room and you can consult this at any time. It is also listed on the board in this room and in the larger room. There are also copies of the tables in this room alongside your seats.

As the game is repeated four times you will receive a definite \$6 and potentially up to \$20. We have up to 2 hours to complete the experiment including the time I have used in explaining these rules. Before we start the experiment we will run through one practice game. This will

be run just as I have described the main games except for two changes. You will not receive a prize for your actions, and we will have another opportunity for questions at the end of the practice. It will also give me the opportunity to double-check that you understand what is expected of you. You now have some time to examine the tables and ask any questions."

Appendix B: Costs and Benefits of Delay

This appendix calculates the costs and benefits of delay for a subject in the experiment with identical discount factors. Start by assuming the subject has already observed a private RR signal. This then leads him to expect others to be more likely to have also observed signals from the red urn, since his own signal biases him in that direction. He should consider all 45 alternative sets of signal observations, calculate the likelihood of each one taking place (which will be biased by his initial beliefs based on his own RR signal) and then evaluate the expected value given the possible signal. This will in some cases produce regret, where the subject feels he should really have gone for white. The total regret is the total of all the expected values of opting for red, given that white seems more likely after waiting one period, weighted by the likelihood of observing these pro-white sequences of signals. The expression begins with $\frac{1}{3}$ to the power equal to the number of possible R signals and $\frac{2}{3}$ to the power equal to the number of W signals. This is then multiplied by the number of combinations which yield this set of signals and by the probability of white being correct, remembering we have already observed a RR. This set of calculations is then summed and the total is multiplied by the second period payoff[®].

The actual sets of possible actions which would lead to regret are:

f8WWg, f7WW; 1RRg, f6WW; 2RRg, f5WW; 3RRg, f7WW; 1WRg, f6WW; 1RR; 1WRg, f5WW; 2RR; 1WRg, f6WW; 2WRg, f5WW; 1RR; 2WRg, f4WW; 2RR; 2WRg, f5WW; 3WRg, f4WW; 1RR; 3WRg, f4WW; 4WRg, f3WW; 1RR; 4WRg, f3WW; 5WRg, f2WW; 6WRg.

Summing each possibility respectively produces:

$$\begin{aligned} & \frac{1}{3} \frac{16}{2^{14+1}} + \frac{1}{3} \frac{14}{3} \frac{2^2}{3} 8 \frac{2^{10}}{2^{10+1}} + \frac{1}{3} \frac{12}{3} \frac{2^4}{3} 28 \frac{2^6}{2^{6+1}} + \frac{1}{3} \frac{10}{3} \frac{2^6}{3} 56 \frac{2^2}{2^{2+1}} + \frac{1}{3} \frac{15}{3} \frac{2^1}{3} 8 \frac{2^{12}}{2^{12+1}} + \frac{1}{3} \frac{13}{3} \frac{2^3}{3} 56 \frac{2^8}{2^{8+1}} \\ & + \frac{1}{3} \frac{11}{3} \frac{2^5}{3} 168 \frac{2^4}{2^{4+1}} + \frac{1}{3} \frac{14}{3} \frac{2^2}{3} 28 \frac{2^{10}}{2^{10+1}} + \frac{1}{3} \frac{12}{3} \frac{2^4}{3} 168 \frac{2^6}{2^{6+1}} + \frac{1}{3} \frac{10}{3} \frac{2^6}{3} 420 \frac{2^2}{2^{2+1}} + \frac{1}{3} \frac{13}{3} \frac{2^3}{3} 56 \frac{2^8}{2^{8+1}} \\ & + \frac{1}{3} \frac{11}{3} \frac{2^5}{3} 280 \frac{2^4}{2^{4+1}} + \frac{1}{3} \frac{12}{3} \frac{2^4}{3} 70 \frac{2^6}{2^{6+1}} + \frac{1}{3} \frac{10}{3} \frac{2^6}{3} 280 \frac{2^2}{2^{2+1}} + \frac{1}{3} \frac{11}{3} \frac{2^5}{3} 56 \frac{2^4}{2^{4+1}} + \frac{1}{3} \frac{10}{3} \frac{2^6}{3} 28 \frac{2^2}{2^{2+1}} \end{aligned}$$

All multiplied by \$3:25, yields 0:5p which is below the cost of delay of 20p so the subject acting optimally should choose red if he observes a RR signal. By symmetry, it is also the case that a subject with the signal WW should choose white in the first period.

For a neutral signal we follow a similar procedure except we have a mixed signal of W and R as our initial signal. The subject knows that the distribution is skewed in favour of the correct choice, so anticipates a $\frac{2}{3}$ probability of the correct ball being selected from the urn. We simply assume white is correct for this calculation, since this will give the same result as assuming red is correct by symmetry. For the cost of delay we instead assume no useful information

so any choice (red or white) will result in a $\frac{1}{2}$ probability of the correct action. This yields $\frac{1}{2}(\$3:50 - \$3:25) = 12:5p$.

The sets which lead to regret in this case are slightly different:

f8WWg, f7WW; 1RRg, f6WW; 2RRg, f5WW; 3RRg, f7WW; 1WRg, f6WW; 1RR; 1WRg, f5WW; 2RR; 1WRg, f4WW; 3RR; 1WRg, f6WW; 2WRg, f5WW; 1RR; 2WRg, f4WW; 2RR; 2WRg, f5WW; 3WRg, f4WW; 1RR; 3WRg, f3WW; 2RR; 3WRg, f4WW; 4WRg, f3WW; 1RR; 4WRg, f3WW; 5WRg, f2WW; 1RR; 5WRg, f2WW; 6WRg, f1WW; 7WRg.

Summing each possibility respectively produces:

$$\begin{aligned} & \frac{2^{16}}{3} \frac{2^{16}}{2^{16}+1} + \frac{2^{14}}{3} \frac{1^2}{3} 8 \frac{2^{12}}{2^{12}+1} + \frac{2^{12}}{3} \frac{1^4}{3} 28 \frac{2^8}{2^8+1} + \frac{2^{10}}{3} \frac{1^6}{3} 56 \frac{2^4}{2^4+1} + \frac{2^{15}}{3} \frac{1^1}{3} 8 \frac{2^{14}}{2^{14}+1} \\ & + \frac{2^{13}}{3} \frac{1^3}{3} 56 \frac{2^{10}}{2^{10}+1} + \frac{2^{11}}{3} \frac{1^5}{3} 168 \frac{2^6}{2^6+1} + \frac{2^9}{3} \frac{1^7}{3} 280 \frac{2^2}{2^2+1} + \frac{2^{14}}{3} \frac{1^2}{3} 28 \frac{2^{12}}{2^{12}+1} + \frac{2^{12}}{3} \frac{1^4}{3} 168 \frac{2^8}{2^8+1} \\ & + \frac{2^{10}}{3} \frac{1^6}{3} 420 \frac{2^4}{2^4+1} + \frac{2^{13}}{3} \frac{1^3}{3} 56 \frac{2^{10}}{2^{10}+1} + \frac{2^{11}}{3} \frac{1^5}{3} 280 \frac{2^6}{2^6+1} + \frac{2^9}{3} \frac{1^7}{3} 560 \frac{2^2}{2^2+1} + \frac{2^{12}}{3} \frac{1^4}{3} 70 \frac{2^8}{2^8+1} \\ & + \frac{2^{10}}{3} \frac{1^6}{3} 280 \frac{2^4}{2^4+1} + \frac{2^{11}}{3} \frac{1^5}{3} 56 \frac{2^6}{2^6+1} + \frac{2^9}{3} \frac{1^7}{3} 168 \frac{2^2}{2^2+1} + \frac{2^{10}}{3} \frac{1^6}{3} 28 \frac{2^4}{2^4+1} + \frac{2^9}{3} \frac{1^7}{3} 8 \frac{2^2}{2^2+1} \end{aligned}$$

All multiplied by \$3:25, yields 38:5p to the nearest half-penny. Since this exceeds the cost of delay, which equals 12:5p the subject with a mixed signal of R and W should wait.

Appendix C: Raw Data

This appendix details the raw data taken over the course of the experiment. The subjects are indexed by day and actual payoff. For example, the highest payoff subject on day two is given the index 2.1.

Table C1: Optimal and Actual Actions, Day 1¹

Subject	Signals	Actions ²			Payoffs		
		Ex Ante	Ex Post	Actual	Ex Ante	Ex Post	Actual
1.1	WR, RR, RR, WW	?2, r1, r1, w1	w2, r1, r1, w1	w2, w3, r1, w1	\$18.125	\$19.75	\$16.25
1.2	WW, RW, WR, WW	w1, r2, ?2, w1	w1, w2, ?2, w1	w1, w2, r2, w1	\$17.875	\$14.625	\$16.25
1.3	RR, RR, RR, WW	r1, r1, r1, w1	r1, r1, r1, w1	w2, w3, r1, w1	\$16.50	\$16.50	\$16.25
1.4	RR, RR, WR, RW	r1, r1, ?2, w2	r1, r1, ?2, w2	r1, r1, r2, w2	\$14.625	\$14.375	\$15.00
1.5	WW, WW, RW, WW	w1, w1, ?2, w1	w1, w1, ?2, w1	w1, w1, w2, w1	\$14.625	\$14.625	\$13.00
1.6	WW, WW, RW, WR	w1, w1, ?2, w2	w1, w1, ?2, w2	w1, w1, w2, w2	\$14.375	\$14.375	\$12.75
1.7	RR, RW, WR, WW	r1, r2, ?2, w1	r1, w2, ?2, w1	r1, w2, r2, w1	\$14.375	\$11.125	\$12.75
1.8	RW, RW, WW, WW	?2, r2, w1, w1	w2, w2, w1, w1	w2, w2, w1, w2	\$14.375	\$12.75	\$12.50
1.9	WR, WR, WW, WR	?2, r2, w1, w2	w2, w2, w1, w2	w2, w2, w1, w2	\$14.125	\$12.50	\$12.50

Notes: 1 Correct choices: w,r,r,w; 2 For example, r1 is red in period 1, ?2 is randomize in period 2.

Table C2: Optimal and Actual Actions, Day 2¹

Subject	Signals	Actions ²			Payo@s		
		Ex Ante	Ex Post	Actual	Ex Ante	Ex Post	Actual
2.1	WW, WW, RR, RR	w1, w1, r1, r1	w1, w1, r1, r1	w1, w1, r1, r1	\$20.00	\$20.00	\$20.00
2.2	WW, WW, RR, RR	w1, w1, r1, r1	w1, w1, r1, r1	w1, w1, r2, r1	\$20.00	\$20.00	\$19.75
2.3	WW, WW, WR, RW	w1, w1, r2, r2	w1, w1, r2, r2	w1, w1, r2, r2	\$19.50	\$19.50	\$19.50
2.4	WW, WW, RW, RW	w1, w1, r2, r2	w1, w1, r2, r2	w1, w2, r2, r2	\$19.50	\$19.50	\$19.25
2.5	WR, RW, RW, RR	w2, w2, r2, r1	w2, w2, r2, r1	w2, w2, r2, r1	\$19.25	\$19.25	\$19.25
2.6	RW, RW, RR, RW	w2, w2, r1, r2	w2, w2, r1, r2	w2, w2, r1, r2	\$19.25	\$19.25	\$19.25
2.7	RW, WW, WR, RR	w2, w1, r2, r1	w2, w1, r2, r1	w2, w2, r2, r2	\$19.50	\$19.50	\$19.00
2.8	WW, WW, RR, WW	w1, w1, r1, w1	w1, w1, r1, w1	w1, w2, r1, w1	\$16.50	\$16.50	\$16.25
2.9	RR, WW, RR, RW	r1, w1, r1, r2	r1, w1, r1, r2	r1, w2, r1, r2	\$16.25	\$16.25	\$16.00

Notes: 1 Correct choices: w,w,r,r; 2 Actions are denoted as in Table C1.

Table C3: Optimal and Actual Actions, Day 3¹

Subject	Signals	Actions ²			Payo@s		
		Ex Ante	Ex Post	Actual	Ex Ante	Ex Post	Actual
3.1	WW, WW, RR, RR	w1, w1, r1, r1	w1, w1, r1, r1	w1, w1, r1, r1	\$20.00	\$20.00	\$20.00
3.2	WW, RW, RR, RR	w1, w2, r1, r1	w1, ?2, r1, r1	w1, w2, r1, r1	\$19.75	\$18.125	\$19.75
3.3	WW, WR, WR, RR	w1, w2, r2, r1	w1, ?2, r2, r1	w1, w2, r2, r1	\$19.50	\$17.875	\$19.50
3.4	WW, WR, RW, WR	w1, w2, r2, r2	w1, ?2, r2, r2	w1, r2, r2, r2	\$19.25	\$17.625	\$16.00
3.5	WW, WR, RW, RR	w1, w2, r2, r1	w1, ?2, r2, r1	w2, r2, r2, r2	\$19.50	\$17.875	\$15.75
3.6	WR, WW, RW, RW	w2, w1, r2, r2	w2, w1, r2, r2	w2, r3, r3, r2	\$19.25	\$19.25	\$15.50
3.7	RR, RW, RR, RR	r1, w2, r1, r1	r1, ?2, r1, r1	r1, r2, r1, r1	\$16.25	\$14.625	\$13.00
3.8	RW, RR, WW, WR	w2, r1, w1, r2	w2, r1, w1, r2	w2, r1, w1, r2	\$12.50	\$12.50	\$12.50
3.9	RR, RW, RW, RR	r1, w2, r2, r1	r1, ?2, r2, r1	r1, r2, r2, r1	\$16.00	\$14.375	\$12.50

Notes: 1 Correct choices: w,w,r,r; 2 Actions are denoted as in Table C1.

Table C4: Optimal and Actual Actions, Day 4¹

Subject	Signals	Actions ²			Payo@s		
		Ex Ante	Ex Post	Actual	Ex Ante	Ex Post	Actual
4.1	RR, WW, RR, WW	r1, w1, r1, w1	r1, w1, r1, w1	r1, w1, r1, w1	\$20.00	\$20.00	\$20.00
4.2	RR, WW, RR, WW	r1, w1, r1, w1	r1, w1, r1, w1	r1, w1, r1, w1	\$20.00	\$20.00	\$20.00
4.3	RR, WW, WR, WR	r1, w1, r2, w2	r1, w1, r2, w2	r2, w1, r2, w2	\$19.50	\$19.50	\$19.50
4.4	RW, WW, WR, WR	r2, w1, r2, w2	?2, w1, r2, w2	r3, w1, r2, w2	\$19.25	\$17.625	\$19.00
4.5	RW, WW, WR, WR	r2, w1, r2, w2	?2, w1, r2, w2	r3, w1, r2, w2	\$19.25	\$17.625	\$19.00
4.6	WW, WW, RW, WW	w1, w1, r2, w1	w1, w1, r2, w1	w1, w1, r2, w1	\$16.25	\$16.25	\$16.25
4.7	WW, WR, WR, RW	w1, w2, r2, w2	w1, w2, r2, w2	w1, w2, r2, w2	\$16.00	\$16.00	\$16.00
4.8	RW, RR, RR, RW	r2, r1, r1, w2	?2, r1, r1, w2	r3, r1, r1, w2	\$16.00	\$14.275	\$15.75
4.9	RW, RW, WW, RR	r2, w2, w1, r1	?2, w2, w1, r1	r2, w2, w1, r1	\$12.50	\$10.875	\$12.50

Notes: 1 Correct choices: r,w,r,w; 2 Actions are denoted as in Table C1.

Table C5: Optimal and Actual Actions, Day 5¹

Subject	Signals	Actions ²			Payo@s		
		Ex Ante	Ex Post	Actual	Ex Ante	Ex Post	Actual
5.1	RR, WW, RR, RR	r1, w1, r1, r1	r1, w1, r1, r1	r1, w1, r1, r1	\$20.00	\$20.00	\$20.00
5.2	RR, WW, RW, RR	r1, w1, r2, r1	r1, w1, ?2, r1	r1, w1, r2, r1	\$19.75	\$18.125	\$19.75
5.3	RW, RW, RR, RR	r2, w2, r1, r1	r2, w2, r1, r1	r2, w2, r1, r1	\$19.50	\$19.50	\$19.50
5.4	RR, WR, WR, RR	r1, w2, r2, r1	r1, w2, ?2, r1	r1, w2, r2, r1	\$19.50	\$17.875	\$19.00
5.5	RW, RW, RR, RW	r2, w2, r1, r2	r2, w2, r1, r2	r2, w3, r2, r2	\$19.25	\$19.25	\$18.75
5.6	RR, WW, WR, WR	r1, w1, r2, r2	r1, w1, ?2, r2	r2, w2, w2, r2	\$19.50	\$17.875	\$15.75
5.7	WW, WR, WR, WR	w1, w2, r2, r2	w1, w2, ?2, r2	w1, w2, w2, r2	\$15.75	\$14.125	\$12.50
5.8	RR, RR, WW, WW	r1, r1, w1, w1	r1, r1, w1, w1	r1, r1, w1, w1	\$9.50	\$9.50	\$9.50
5.9	WW, WR, WW, WW	w1, w2, w1, w1	w1, w2, w1, w1	w1, w2, w1, w1	\$9.25	\$9.25	\$9.25

Notes: 1 Correct choices: r,w,r,r; 2 Actions are denoted as in Table C1.

Table C6: Optimal and Actual Actions, Day 6¹

Subject	Signals	Actions ²			Payo@s		
		Ex Ante	Ex Post	Actual	Ex Ante	Ex Post	Actual
6.1	WW, WW, WW, WR	w1, w1, w1, r2	w1, w1, w1, r2	w1, w1, w1, r2	\$19.75	\$19.75	\$19.75
6.2	WW, WR, WW, RR	w1, w2, w1, r1	w1, w2, w1, r1	w1, w2, w1, r1	\$19.75	\$19.75	\$19.75
6.3	WW, WW, WR, RR	w1, w1, ?2, r1	w1, w1, ?2, r1	w2, w1, w2, r1	\$18.125	\$18.125	\$19.50
6.4	WR, WW, RW, WR	w2, w1, ?2, r2	?2, w1, ?2, r2	w3, w1, w2, r2	\$17.625	\$16.00	\$19.00
6.5	RR, WW, WR, RR	r1, w1, ?2, r1	r1, w1, ?2, r1	r1, w1, w2, r1	\$14.625	\$14.625	\$16.25
6.6	WW, WR, RR, RW	w1, w2, r1, r2	w1, w2, r1, r2	w1, w2, r1, r2	\$16.00	\$16.00	\$16.00
6.7	RW, WW, RR, RW	w2, w1, r1, r2	?2, w1, r1, r2	r2, w1, r1, r2	\$16.00	\$14.375	\$12.75
6.8	RR, RW, WR, RR	r1, w2, ?2, r1	r1, w2, ?2, r1	r1, w2, r2, r1	\$14.375	\$14.375	\$12.75
6.9	RR, RR, WR, RW	r1, r1, ?2, r2	r1, r1, ?2, r2	r1, r1, r2, r2	\$10.875	\$10.875	\$9.25

Notes: 1 Correct choices: w,w,w,r; 2 Actions are denoted as in Table C1.

The next two data tables were based on the alternative condition in which the subjects are told what the correct answer is after each game.

Table C7: Optimal and Actual Actions, Day 7¹

Subject	Signals	Actions ²			Payo@s		
		Ex Ante	Ex Post	Actual	Ex Ante	Ex Post	Actual
7.1	WW, RW, RR, RR	w1, w2, r1, r1	w1, w2, r1, r1	w1, w2, r1, r1	\$19.75	\$19.75	\$19.75
7.2	WW, WW, WW, RR	w1, w1, w1, r1	w1, w1, w1, r1	w1, w1, w1, r1	\$16.50	\$16.50	\$16.50
7.3	WR, WW, WR, RW	w2, w1, r2, r2	w2, w1, w2, r2	w2, w1, w2, r2	\$19.25	\$16.00	\$16.00
7.4	WW, RW, WR, RW	w1, w2, r2, r2	w1, w2, w2, r2	w2, w2, w2, r2	\$19.25	\$16.00	\$15.75
7.5	WR, WR, WR, RW	w2, w2, r2, r2	w2, w2, w2, r2	w2, w2, w2, r3	\$19.00	\$15.75	\$15.50
7.6	WR, WR, RR, RW	w2, w2, r1, r2	w2, w2, r1, r2	w2, w2, w3, r3	\$19.25	\$19.25	\$12.50
7.7	WR, WW, RW, WW	w2, w1, r2, w1	w2, w1, w2, w1	w2, w1, w2, w1	\$16.00	\$12.75	\$12.50
7.8	RR, WR, RR, RR	r1, w2, r1, r1	r1, w2, r1, r1	r1, w2, w2, r2	\$16.25	\$16.25	\$12.50
7.9	RW, RR, WW, RW	w2, r1, w1, r2	w2, r1, w1, r2	w2, r1, w1, r3	\$12.75	\$12.75	\$12.50

Notes: 1 Correct choices: w,w,r,r; 2 Actions are denoted as in Table C1.

Table C8: Optimal and Actual Actions, Day 8¹

Subject	Signals	Actions ²			Payoffs		
		Ex Ante	Ex Post	Actual	Ex Ante	Ex Post	Actual
8.1	RR, WW, RR, WR	r1, w1, r1, w2	r1, w1, r1, w2	r1, w1, r1, w2	\$16.25	\$16.25	\$16.25
8.2	RW, RW, WR, WW	r2, r2, r2, w1	r2, w2, r2, w1	r2, w2, r2, w1	\$19.25	\$16.00	\$16.00
8.3	RR, RW, WR, RW	r1, r2, r2, w2	r1, w2, r2, w2	r1, w2, r2, w2	\$19.25	\$16.00	\$16.00
8.4	RR, WR, WR, RW	r1, r2, r2, w2	r1, w2, r2, w2	r1, w2, r2, w2	\$19.25	\$16.00	\$16.00
8.5	WR, WW, RW, WW	r2, w1, r2, w1	r2, w1, r2, w1	r2, w1, r3, w1	\$16.00	\$16.00	\$15.75
8.6	WW, RR, WR, RW	w1, r1, r2, w2	w1, r1, r2, w2	w1, w3, r3, w2	\$16.00	\$16.00	\$12.25
8.7	RR, WR, RR, WR	r1, r2, r1, w2	r1, w2, r1, w2	r2, w2, r2, w3	\$19.50	\$16.25	\$15.50
8.8	WR, RR, RW, WR	r2, r1, r2, w2	r2, r1, r2, w2	r2, w3, r3, w3	\$19.25	\$19.25	\$15.25
8.9	WR, RR, WW, RR	r2, r1, w1, r1	r2, r1, w1, r1	r2, r1, w2, r1	\$12.75	\$12.75	\$12.50

Notes: 1 Correct choices: r,r,r,w; 2 Actions are denoted as in Table C1.

This data was collected over eight days, the first two of which were in June 1999, the final six of which were in November 1999. In all cases subjects had no previous experience of similar experiments and no prior knowledge of herding theory, or game theory in general. Graduate students in economics and related fields were excluded from consideration as subjects. The age range was from late teens to late twenties and there was a roughly even gender mix.

References

- [1] Akerlof, G, and Yellen, J. (1987). 'Rational Models of Irrational Behaviour.' *American Economic Review*, vol. 77, pp. 137-142.
- [2] Anderson, R, and Holt, C. (1997). 'Information Cascades in the Laboratory.' *American Economic Review*, vol. 87, pp. 847-862.
- [3] Banerjee, A. V. (1992). 'A Simple Model of Herd Behaviour.' *Quarterly Journal of Economics*, vol. 107, pp. 797-817.
- [4] Bikhchandani, S, Hishleifer, D and Welch, I. (1992). 'A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades.' *Journal of Political Economy*, vol. 100, pp. 992-1026.
- [5] Chamley, C, and Gale, D. (1994). 'Information Revelation and Strategic Delay in a Model of Investment.' *Econometrica*, vol. 108, pp. 60-74.
- [6] Dixit, A. *The Art of Smooth Pasting* (1993). *Fundamentals of Pure and Applied Mathematics* volume 55, Harwood Academic Publishers.
- [7] Gul, F, and Lundholm, R. (1995). 'Endogenous Timing and the Clustering of Agents' Decisions.' *Journal of Political Economy*, vol. 103, pp. 1039-1066.
- [8] Lee, I. H. (1993). 'On the Convergence of Information Cascades.' *Journal of Economic Theory*, vol. 61, pp. 395-411.

Nuffield College, Oxford OX1 1NF

E-mail address: daniel.sgroi@economics.oxford.ac.uk

URL: <http://www.nuff.ox.ac.uk/users/sgroi>