

Herding and Contrarianism in a Financial Trading Experiment
with Endogenous Timing

Andreas Park and Daniel Sgroi

No 868

WARWICK ECONOMIC RESEARCH PAPERS

DEPARTMENT OF ECONOMICS

THE UNIVERSITY OF
WARWICK

Herding and Contrarianism in a Financial Trading Experiment with Endogenous Timing

Andreas Park*

University of Toronto

Daniel SgROI†

University of Warwick

This version: October 1, 2008

Abstract

We undertook the first market trading experiments that allowed heterogeneously informed subjects to trade in endogenous time, collecting over 2000 observed trades. Subjects' decisions were generally in line with the predictions of exogenous-time financial herding theory when that theory is adjusted to allow rational informational herding and contrarianism. While herding and contrarianism did not arise as frequently as predicted by theory, such behavior occurs in a significantly more pronounced manner than in comparable studies with exogenous timing. Types with extreme information traded earliest. Of those with more moderate information, those with signals conducive to contrarianism traded earlier than those with information conducive to herding.

JEL Classification: C91, D82, G14.

Keywords: Herding, Contrarianism, Endogenous-time, Informational Efficiency, Experiments.

*E-mail: andreas.park@utoronto.ca

†E-mail: daniel.sgROI@warwick.ac.uk

Financial support from the ESRC, SSHRC, the Leverhulme Trust and the Cambridge Endowment for Research in Finance is gratefully acknowledged. Special thanks go to Annie Jekova and Malena Digiuni for expert research assistance.

1 Introduction

Rational herding is an important tool in analyzing how and why economic agents learn through observation in groups (Banerjee (1992) and Bikhchandani, Hirshleifer and Welch (1992)). Economic agents constantly learn from others, through direct conversation, newspapers, observing actions and typically in financial markets, through observing price movements or the buy and sell decisions of others. Informational herding arises in situations where people observe the actions of others, derive information from them and then, seemingly disregarding their own information, follow the majority action. A central lesson of informational herding theory is that what we can observe is not necessarily a direct indication of the information possessed by others. Furthermore, while for each individual following the herd is likely to be a better decision than ignoring the actions of others altogether, the collective choice of the herd is still not necessarily the correct choice, and certainly it is not as reliable a choice as when all individuals pool their information.

It is tempting to argue that this problem must be endemic to finance, where persistent price spikes and crashes, linked movements across national boundaries, and often-cited “crazy” market behavior might be considered prime candidates for a herding analysis. Indeed many highly respected financial commentators and economists have linked herding behavior with financial fragility, including the current period of world-wide financial instability.¹ One could argue that a few early, perhaps incorrect, movements by visible traders induce discontinuous price jumps in one direction or the other, potentially leaving share prices far from their fundamental value. If we could describe such movements using the tools of herding theory, not only would we better understand financial markets, but we would have an intellectual framework to ponder policy suggestions aimed at avoiding painful financial crises. Nevertheless, early work on financial market herding had difficulties identifying such behavior in an environment with informationally efficient prices.²

¹See Shiller (2008) in which Robert Shiller highlights the scope for herding theory as a potential explanation for the current housing and subprime mortgage market difficulties in the USA and elsewhere.

²An early result by Avery and Zemsky (1998) showed that in a simple financial market-trading setting with two values herding is not possible because the market price always separates people with good and bad information. While they also presented an example with moving prices and herding, attributing their finding to ‘multidimensional uncertainty/risk’ (where investors have a finer information structure than the market), in their example prices hardly move, and herding is ‘self-defeating’ as herd-buys eventually stop the herd. Since the underlying ‘multidimensional’ information structure seemed very specific, and since the implications of herding as identified in their paper seemed very short-term, rational herding was typically not thought a good explanation for financial crises. See, for instance, Brunnermeier (2001), Bikhchandani and Sunil (2000), Chamley (2004), or Alevy, Haigh and List (2007). Subsequent experimental studies by Drehmann, Oechssler and Roeder (2005) and Cipriani and Guarino (2005), which we will comment on in detail below, confirmed this finding. Cipriani and Guarino (2007) is a recent experimental study of the kind of herd behavior that Avery and Zemsky *did* identify.

The work presented here is the second part of a project to examine experimentally a theoretical framework that *does* admit herding *and* contrarianism as rational outcomes. The first part of the project employed a setup in which people acted in an exogenously defined sequence, directly examining the set of conditions under which herding and contrarianism can occur as identified in Park and Sabourian (2008). The second part, presented in this paper, is the first experimental work to combine an informational herding story with moving prices and, crucially, *endogenous timing*. We explicitly allow subjects to choose both how *and when* to trade and, in some treatments, to trade twice. This enables us to analyze the impact of information on both the direction of a trade and also on the timing decision.

The importance of analyzing endogenous-time as opposed to exogenous-time is hard to overestimate. First, the extra layer of realism is considerable: in real markets traders do not stand in line waiting their turn to act. In fact, one of the key features of financial frenzies is the clustering of actions in time, which cannot be examined in a sequential trading setting.

Second, allowing people to choose both the direction and the time of their trade greatly complicates the subject's decision problem. A testament to the complexity of the problem is that there is currently no theoretical model that encompasses a dynamic trading decision with multiple heterogeneously informed agents. However, the lack of a full-fledged theory does not inhibit an experimental examination of people's behavior under tightly controlled conditions.

Third, with endogenous time the effect of herding and contrarianism may well be much more pronounced than with exogenous timing. To see this, observe that herding types are, at least by casual intuition, in some way or other less confident about the value of an asset. One may then speculate that they delay their trading decision to gain more information by observing others. So once they do trade, their herding behavior may well lead to stronger price distortions under endogenous-timing than exogenous-timing. Our results indicate that indeed herd behavior becomes much more pronounced when traders can decide when to trade. Thus existing studies that deny traders the opportunity to time their actions may be underplaying likely herding and contrarian behavior.

It is also worth noting that the experiment is also one of the largest yet to examine rational herding in a laboratory setting, with about 2000 trades spread across 6 separate treatments.

To simplify the exposition, we will distinguish and separate the decision of the trading direction from the timing decision; we will refer to the former as the static decision problem, and the latter as the dynamic problem. All the while we emphasize that these two belong together in the full equilibrium problem.

THE STATIC DECISION PROBLEM. The ‘static’ theoretical model underlying our experiment is a sequential trading setup in the tradition of Glosten and Milgrom (1985) in which risk-neutral subjects trade single units of a financial security with a competitive market maker. Past trades and prices are public information, and the market maker adjusts the price after each transaction to include the new information revealed by this trade. We differ from Glosten and Milgrom by admitting endogenous timing, whereas in their model traders act at exogenous times. In our specification, there are three possible values and three possible signals (high, middle, or low); each subject receives a private realization of one of these signals. There are three possible actions (buy, pass, or sell). In this underlying static view, rational subjects should buy if their expectation, conditional on all private and public information, is above the price and sell if it is below.

A key result of herding theory is that much depends upon the shape of the underlying conditional signal distribution. Put simply, subjects with a conditional signal distribution that is bi-polar (“U-shaped”), with weight at the extremes, switch their optimal action from buying to selling or vice versa because they update their expectations faster upwards than the price rises, and so they may rationally herd. Alternatively, with single polar information (“hill-shaped”) they adjust their expectation downwards slower than the price falls and thus switch from selling to buying, acting in a rational contrarian manner.³ In our specification there are also two signal types that never switch their action: one type always sells (the “low” signal type), the other always buys (the “high” signal type). Herding or contrarianism thus stems from “middle” types.

THE DYNAMIC DECISION PROBLEM. We employed two different classes of treatments: in the first, people are allowed to trade at most once, in the second they can trade twice. The seminal paper which studies this problem with a single irreversible action and without moving prices is Chamley and Gale (1994). A key message of their work is that decision-makers will move very quickly, since waiting only makes sense when new information arises. Furthermore, any period of low activity can produce a general lull in behavior as each decision-maker decides to delay further to await more information. Both of these insights are difficult to transfer into a setting where prices can move in response to actions, because new public information will immediately be incorporated into the price. Smith (2000) presents a single-trade setup, in which only a single informed trader is considered. He shows that people who receive signals that are either good or bad news (as is the case for our high and low signals) should trade ‘early’, and he also presents an example of a U-shaped signal structure under which people optimally delay. Alternatively, Ivanov, Levin and Peck (2008) argue that generically only subjects with

³A full description of this process is undertaken in Section 3 below.

sufficient foresight are likely to delay to take advantage of new information, while other more self-contained or myopic subjects are unlikely to do so.

RESULTS. The overall fit of the static theoretical model is roughly 73%, which compares well with other experimental work that allows exogenous timing such as Drehmann et al (2005), Cipriani and Guarino (2005, 2007) and Park and SgROI (2008). Broken down by player type, these numbers are 84%, 54% and 86% for the low, the middle and the high signal type respectively. Thus while the fit is very high for the high and low signal types, it is lower for middle signal types. One obvious reason is that the decision with the middle signal is much more difficult: high and low signal types would make no mistakes if they decide simply to stick to their initial decision, whereas middle signal types have to follow the development of prices very carefully. Finally, the overall fit is similar to the results in our companion paper on exogenous timing, Park and SgROI (2008).

The implicit hypothesis in the static rational model is that middle signal types are more prone to herd behavior or contrarianism conditional on the signal distribution. Indeed, this is what we observe: having a U-shaped signal increases the probability of engaging in herding behavior by about 30%; similarly the hill-shaped signal increases the probability of engaging in contrarian behavior by 35% compared to all other signals. Receiving any other signal, in fact, reduces the chance of herding or contrarianism. The effect of the herding signal in particular is much stronger than in the exogenous time setting, studied in our companion paper. There we found a marginal effect of roughly 6% which is a strong indication of the importance of endogenous timing.

We also find evidence that people with the high signal (good-news) and the low signal (bad-news) trade systematically earlier than people with the middle (hill- or U-shaped) signal. Of the latter, those with single-polar information (hill-shape) trade earlier than those with bi-polar (U-shaped) information. In economic terms this implies that contrarians should act comparatively earlier than herding types. A corollary to this statement is that the herd is triggered by some rational people trading quickly in one direction.

Next, there is evidence of intertemporal clustering of trades in the sense of leader-follower patterns: a significant percentage of all trades (about 50-55%) occurs within an interval of 1.5 seconds. Information, however, seems to play no role in these patterns which instead seem intrinsic to all signal types' behavior.

While we believe that our data indicates that people act in the spirit of the rational model, we also examined several alternative hypotheses, including different assumptions on risk preferences, to see if we could further add to our understanding. In general, alternative models produced no extraordinary insights. An exception are error adjustment models or probability over- or under-weighting schemes which may be able to improve

the fit, but only at the cost of implicitly ruling out some interesting behavior (such as herding).⁴

To summarize the key results, We find that: (a) rational herding and contrarianism do arise in accordance with the theory and are triggered by the people who hold certain signals as predicted by the theory, and (b) timing behavior is systematically influenced by signals in that people with clear signals (good, bad and single-polar) move before people with ambiguous signals.

RELATED EXPERIMENTS. The first published paper to consider herding in endogenous-time was SgROI (2003), a close implementation of the theoretical Chamley and Gale (1994) framework, which was also examined experimentally in ZiegelmeYer, My, Vergnaud and Willinger (2005). Neither of these experiments allowed a moving price and were not designed to examine financial trading environments. There are two related experimental papers that do analyze people’s timing behavior in a financial market environment.⁵ Bloomfield, O’Hara and Saar (2005) study a financial market in which people can trade repeatedly throughout a trading day, but without considering herding or contrarianism. The focus of their study is on the timing behavior of some informed insiders and the choice of limit- or market-orders depending on the time. The key to their setup is not the kind of information so much as the fact that people have information. Perhaps the closest link between our work and Bloomfield et al. (2005) is that they also use a controlled experiment to glean insights into forces that drive trading behavior in the absence of an underlying guiding theory.

Ivanov et al. (2008) implement Levin and Peck (2008) (which relates to Chamley and Gale (1994)). While they also study the timing behavior of people’s investment choices, they do not consider a setup with moving prices; their model is best described as capturing the decision of a non-financial (green-field) investment. They find that first, behavior can be classified into three categories: self-contained (ignoring information from observing others), myopic (acting upon the current best decision, ignoring the option to learn from others in the future), and foresighted (the perfectly rational decision). While we have many subjects and many trades, we feel that there is not sufficient variation in our framework to warrant such a description for our traders; instead, in all of our behavioral analyses we assume that informed traders follow a single behavioral rule. Our focus is also different in that we try to understand how different information structures affect the decision to trade.

⁴For a full explanation and more detail on the alternative models examined see the appendix.

⁵There are at least four other papers that study herding behavior in a financial market environment: Drehmann et al. (2005), Cipriani and Guarino (2005, 2007), and our own prequel, Park and SgROI (2008). None of these studies, however, allows subjects to decide the timing of their trading decisions.

Overview. The rest of the paper is structured as follows. In the next section we define herding and contrarianism formally. We then examine the theoretical framework in more detail in Section 3 and also discuss the modifications undertaken to better fit a laboratory experiment. Section 4 examines the design of the experiment, including a discussion of the nature of the software, the different information structures embodied in the alternative treatments, the information provided to subjects and the hypotheses that are implied by the theory. Section 5 presents the results of the experiments and their fit to the rational model. Section 6 carries out a formal econometric analysis of the static decision. Section 7 considers the impact that prices have on decisions. Section 8 studies the relation of the first and second trades (for instance, is there much “buy-low, sell-high” trading). Section 9 analyzes the timing behavior. Section 10 summarizes the key findings and concludes. Appendix A outlines results from an examination of alternative behavioral explanations for the observed data. The subject instructions and other supporting materials are provided in the final appendices.

2 Formal Definition of Herding

In the spirit of Avery and Zemsky, and Park and Sabourian, we employ a set of definitions for herding (contrarian) behavior in which a trader switches from buying to selling after observing a history with increasing (decreasing) prices; and for our experimental analysis, we shall categorize herding as such. A very loose intuition is that herding types have increasing demand functions: they sell when prices are low and buy when they are high. We use notation H_t for the trading history at time t ; this history includes all past actions, their timing, and the transaction prices; H_1 is the initial history.

Definition 1 (Herd- and Contrarian-Behavior)

- (a) *A trader engages in herd-buying after a history of trade H_t if and only if*
- (H1) *he would sell at the initial history H_1 ,*
 - (H2) *he buys at history H_t , and*
 - (H3) *prices at H_t are higher than at H_1 .*
- Sell herding is defined analogously.*
- (b) *A trader engages in buy-contrarianism after a history of trade H_t if and only if*
- (C1) *he would sell at the initial history H_1 ,*
 - (C2) *he buys at history H_t , and*
 - (C3) *prices at H_t are lower than at H_1 .*
- Sell-contrarianism is defined analogously.*

According to the above definitions, agents with a particular signal engage in herding if, as a result of observing the behavior of others, they take a different action from the one that they would take initially. Thus, herding in our set-up (as well as in Avery and Zemsky) represents any history-induced switch of opinion *in the direction of the crowd*, whereas the logical counterpart, contrarianism, describes a history-induced switch of opinion *against the direction of the crowd*.

3 The Underlying Theory

Subjects face a complex decision problem, having to decide both on the timing and direction of their trade. We split the description here into the trade-direction and the timing component, while noting that a full equilibrium model requires a simultaneous description of both.

3.1 The Static Decision of the Trading Direction

All traders trade a security with an uninformed market maker. The security takes one of three possible liquidation values, $V_1 < V_2 < V_3$, each equally likely. Traders can be informed, in which case they receive a conditionally independent signal about the true value of the security, or they can be a noise trader in which case they trade for reasons outside the model. The market maker sets a single price at which he is willing to buy or sell one unit of the security.⁶

Every trader is a noise trader with a fixed probability (25% in our setting) and buys or sells each with 50% probability. Informed traders receive one of three signals: S_1, S_2 or S_3 . Signal S_1 is generated with higher probability in state V_1 than V_2 , and likewise in state V_2 than state V_3 . The reverse holds for signal S_3 . This implies that the recipient of signal S_1 shifts probability weight towards the lowest state (S_1 is ‘bad news’), whereas the recipient of S_3 shifts weight towards the highest state (S_3 is ‘good news’). Signal S_2 can take several different shapes which we will outline shortly.

All past trades and prices are public information. The market maker follows a simple pricing rule by setting the unique trading price as the expectation of the true value of the security, conditional on all publicly available past information; this price does not account for possible timing decisions of traders but instead is purely backward looking. Once a trader decides to act, he buys if his expectation conditional on his private signal and on

⁶This is a simplification of the Park and Sabourian model, which itself a sequential trading model in the style of Glosten and Milgrom (1985). In these models, a competitive market maker sets a zero-profit bid and offer price. In our experiments, as is standard in the related experimental literature, we dispense with bid- and ask-prices and focus instead on a single trading price to minimize complexity.

the information derived from past trades exceeds the price, and he sells if this expectation is below the price.

The following conditions describe the shape of the herding candidate's conditional signal distribution (henceforth: *csd*):

$$\begin{aligned} \text{increasing} &\Leftrightarrow \Pr(S|V_1) < \Pr(S|V_2) < \Pr(S|V_3) \\ \text{decreasing} &\Leftrightarrow \Pr(S|V_1) > \Pr(S|V_2) > \Pr(S|V_3) \\ \text{U-shaped} &\Leftrightarrow \Pr(S|V_i) > \Pr(S|V_2) \text{ for } i = 1, 3 \\ \text{hill-shaped} &\Leftrightarrow \Pr(S|V_i) < \Pr(S|V_2) \text{ for } i = 1, 3. \end{aligned}$$

A signal is called *csd-monotonic* if its *csd* is either increasing or decreasing. If $\Pr(S|V_1) > \Pr(S|V_3)$ or $\Pr(S|V_1) < \Pr(S|V_3)$ holds then signals are negatively and positively biased respectively. A negatively biased U-shaped *csd* is referred to as a negative U-shape, likewise for hill-shape and positive biases.

The underlying decisions that we categorize are assumed to be static or myopic, i.e. we ignore possible dynamic considerations that may govern a subject's trading decision. Moreover, we also assume that every past decision was taken on the basis of only static considerations. Applied to the experimental design in which bid-ask-spreads are ignored, the analysis in Park-Sabourian implies the following:⁷

Theorem 1 (Herding and Contrarian Behavior, Park and Sabourian (2008))

- (a) *Types S_1 and S_3 never herd.*
- (b) *Type S_2 buy-herds (sell-herds) if and only if his *csd* is negative (positive) U-shaped.*
- (c) *Type S_2 is a buy-(sell-) contrarian if and only if his *csd* is negative (positive) hill-shaped.*

First observe, that a buy-herding trader would be selling at the initial history (H1/C1). Since the prior is uniform, this implies that sellers attach more weight to the lowest than the highest state, i.e. $\Pr(S|V_1) > \Pr(S|V_3)$. Next, buy-herding also requires that prices have increased (H3); this occurs if and only if the probability of the lowest state is smaller than that of the highest state.

Sufficiency can be best explained by imagining that the probability of the lowest state V_1 has dropped to the point where the state can be ignored relative to states V_2 and V_3 . Then a trader who is buying must attach more weight to state V_3 than V_2 , $\Pr(S|V_3) > \Pr(S|V_2)$. This holds for trader S_3 but this type would not be selling at the initial history. Combining the requirements $\Pr(S|V_3) > \Pr(S|V_2)$ and $\Pr(S|V_1) > \Pr(S|V_3)$, we observe that a U-shaped signal allows herding.

⁷Note that 'S herds' is to be read as 'S herds with positive probability'.

A similar idea applies to the occurrence of contrarian behavior where the price falls and so that state V_3 can be ignored relative to V_1 and V_2 . Now a buyer must put more weight on V_2 than V_1 , $\Pr(S|V_2) > \Pr(S|V_1)$ which, together with $\Pr(S|V_1) > \Pr(S|V_3)$ lets us conclude that a hill-shaped signal allows contrarianism.⁸

3.2 The Dynamic Decision of the Trading Time

Most work on financial market microstructure constructs the trading decision to be either stationary or static.⁹ Those that do analyze a dynamic problem usually allow at most one informed insider.¹⁰

The theoretical paper closest to ours is Smith (2000) which models a single trader who can make a single trade early or late. Smith shows that this trader will trade early if his signal is either good or bad news.¹¹ Moreover, Smith also presents an example with a U-shaped signal and shows that within his framework this trader should indeed delay.¹² Applying Smith to our framework, the S_1 and S_3 types have bad-news and good-news signals respectively and thus they should always act immediately.¹³

Matters are more complex for the S_2 types. A U-shaped signal appears to be rather uninformative, giving only an indication that an extreme state has happened. It thus seems reasonable to assert that these types should delay initially to accumulate information. A hill-shaped signal, on the other hand, is a strong indication that the middle state occurred. Consequently, initially there seems to be no sense for hill-shaped S_2 types to trade. If, however, prices move away from the middle state, then it should pay for the hill-shaped S_2 types to trade against this general flow because they should expect prices to revert to the middle value. Comparing traders with hill- and U-shaped signals, it seems reasonable that the U-shaped types delay for longer to get a better sense of the direction in which the market is moving whereas for hill-shaped types, it may make sense to act quickly on deviations of the price from the middle value.

The statements so far cover the case of a single trade, and there is no theoretical work

⁸The argument for necessity is more involved and we refer the reader to Park and Sabourian (2008).

⁹The classic market theoretical market microstructure literature is surveyed in Hasbrouck (2007) or Vives (2008).

¹⁰See, for instance, Kyle (1985), Back and Baruch (2004), or Chakraborty and Yilmaz (2004).

¹¹There is an additional requirement in that the distribution of public information satisfies the monotone likelihood ratio property (see Milgrom (1981)); this condition is to ensure that, effectively, the true state will be revealed in a systematic manner. Our experimental setup will have this feature.

¹²Smith's results are about sufficiency and thus provide no general answer as to how U- or hill-shaped types should trade.

¹³An alternative approach to Smith's finding is that the expectations that the S_1 and S_3 types form over the future price are super- and sub-martingales respectively. For instance, an S_3 type would always expect that there are more people who have the same signal and thus expects prices to rise. For their own expectation, however, the law of iterated expectations applies.

in the literature that provides guidelines as to how people would optimally act when there are multiple heterogeneously informed traders who can trade repeatedly. One reason is the sheer complexity of the issue, having to account for the impact of one's own decision on others' actions, how their reaction affects one's own decision, iterated *ad infinitum*. Despite the lack of a theory to guide our analysis, we can conjecture possible behaviors. With two trades, people can either trade twice in the same direction or they can trade in opposite directions (sell-buy or buy-sell); the latter is referred to as a 'return'-trade. With return trades possible, we must revisit the decision about the trading direction for it is now conceivable that a dynamic trading strategy might involve trading at a static expected loss.

Consider first the S_1 types. Suppose that they perform a return-trade, for instance, by buying early and selling late. Then after these two transactions, they still hold a share, and on this single share they make an expected loss. While they may gain on the return transaction, it seems most intuitive that the S_1 types should sell twice. A similar argument applies to the S_3 types: if they sell early and buy late (or vice-versa) they would hold only one share while they could hold three.

For the S_2 types, however, we have no prediction about the dynamic trading direction, nor do we have an intuitive sense for how they should act in equilibrium with two trades. They could trade in the same direction twice, following the crowd or acting against it. Return trades, however, can also be rationalized since even taking a static view the S_2 types can change their optimal action after observing different histories of trade. In our analysis we will check if there is persistency in their behavior with respect to (static) herding or contrarianism. Our experiment is thus an exploration of the strategic timing of behavior.

Coming back to the timing of a second trade, we note that a second, same-direction trade may be delayed. Since traders are not infinitesimal, their actions have a discrete price impact. A trader who traded early may thus delay, hoping that the price reverts back against the movement that the trader caused. If delay of the second trade is observed, then it is also no longer clear that we should expect to find a significant difference between the timing of second trades for S_1 or S_3 types and S_2 types.

4 Experimental Design

Here we discuss the experimental design, focusing on the information provided to the subjects, the differences between treatments, and the predictions made in advance. The appendix contains further information including a time-line (Appendix B), a full set of instructions and the material given to subjects (Appendix C), as well as a thorough description of the purpose-built software used for this experiment (Appendix D).

4.1 Overview

The design focused on a financial asset that can take one of three possible liquidation values $V \in \{75, 100, 125\}$ which correspond to the true value of the asset. The traders were typically made up of 15-25 experimental subjects plus a further 25% of computerized noise traders, with a central computer acting as the market maker.

Prior to each treatment each subject i received an informative private signal, $S_i \in S \in \{S_1, S_2, S_3\}$. They were also provided with an information sheet detailing the prior probability of each state, and a list of what each possible signal would imply for the probability of each state, and the likelihood of each signal being drawn given the state. In other words, we provided both the signal distribution and the initial posterior distribution, conditional on receiving a signal. The information on the sheet was common knowledge to all subjects. In particular the subjects therefore *should* have realized that the quality of the signal was *ex post* identical for all subjects. The subjects were *not* told anything about the implications of U-shaped, hill-shaped or monotonic information structures or the predictions of the theory.

The nature of the compensation system was also made very clear in advance, and in particular that it directly implied that they should attempt to make the highest possible virtual profit in each round, since the final compensation was based on overall performance (in UK currency up to £25 combined with a one-off participation fee of £5, or the equivalent in Canadian dollars).

The existence and proportion of noise traders was made known to the experimental subjects in advance, who were also aware that noise traders randomized 50 : 50 between buying and selling and that they trade at random times.

The subjects were informed that the sessions would last 3 minutes and that they would receive announcements about the remaining time after 2:30 minutes, and 2:50 minutes. We considered two classes of treatments: in the first people were allowed to trade once, in the second they could trade twice. The software allowed subjects to trade only this specific number of times. The sequence of transactions produced a history of actions and prices, H_t with $t \in (0, 3)$, that recorded the price, timing (in seconds) and direction of each transaction. Subjects were shown the history in the form of a continuously updating price chart during each treatment, and they were also given the current price, P_t . This price was calculated by the computer as $P_t = E[V|H_t]$ with $P_1 = 100$.

Subjects were told that they had three possible actions $a = \{\text{sell, pass, buy}\}$ which they could undertake only during the 3 minutes of trading time. They were instructed that pressing the ‘pass’-button would count as one of the actions that they were allowed.

It was stressed to the subjects that their virtual profits per treatment were generated

based on the difference between the price at which they traded, P_{t^*} where t^* is the time of their personal trading opportunity, and the true value of the share, V .

The subjects themselves were recruited from the Universities of Toronto, Cambridge and Warwick. No one was allowed to take part twice. We ran 13 sessions in all: 3 at the University of Cambridge, 6 at the University of Warwick and 4 at the University of Toronto. We collected demographic data only for the Warwick sessions: of the subjects there, around 49% were female, around 73% were studying (or had already taken) degrees in Economics, Finance, Business, Statistics, Management or Maths. 53% claimed to have some prior experience of financial markets, with some 23% owning shares at some point in the past.¹⁴

4.2 Treatments

Following Section 3, the rational action for S_1 and S_3 subject types was to sell or buy respectively, irrespective of H_t , while for the S_2 types the nature of H_t and the precise information structure determined a unique optimal action. In the first three treatments, subjects were allowed to trade at most once, in the last three treatments subjects were allowed to trade at most twice. The treatments were each designed to enable us to examine a specific information structure.¹⁵

Treatment 1: negative U-shaped signal structure making buy herding possible;

Treatment 2: negative hill-shape making buy-contrarianism possible;

Treatment 3: positive U-shaped signal structure making sell-herding possible;

Treatment 4: as Treatment 2 but with two trades;

Treatment 5: as Treatment 3 but with two trades;

Treatment 6: as Treatment 1 but with two trades.

Therefore under each treatment, once we knew the signals and trades, we could exactly calculate the theoretically predicted static (or myopic) action for each subject; in some cases this might be to herd or act in a contrarian way.

The endogenous-time treatments presented in this paper were run as part of a larger set of experiments that involved the same computer software. Since subjects were exposed

¹⁴Appendix E details the questions asked in the questionnaire. When asked what motivated their decisions (across different sessions) 44% of subjects mentioned a combination of prices and signals, 31% only price, 18% only signal and the remaining 7% had other motivations. 38% thought that in general the current price was more important than the signal, 36% thought the signal was more important than the current price and the remaining 26% felt they were of similar value. Roughly 24% claimed to have carried out numerical calculations.

¹⁵Since in Drehmann et al. (2005) the inclusion of transaction costs produced the expected outcomes, we ignored transactions costs and instead focused on the information structure as the key differentiating factor between treatments.

to an example round, and the endogenous-time treatments were run last, subjects had ample time to familiarize themselves with the software.

4.3 Theoretical Predictions

Single-Trade Case. Given the proximity of the design to the theoretical model outlined in Section 3, several predictions arise immediately:

Hypothesis 1 *S_1 types sell as soon as the treatment starts, and S_3 types buy as soon as the treatment starts.*

While we generally predict that the S_1 and S_3 types should trade immediately, we would expect that the distribution of trades for the S_1 and S_3 types should be strongly tilted towards time 0.

As described in Section 3 the S_2 types' behavior is a function of both the treatment and history H_t . Specifically, from the S_2 types, we expect to see possible herding behavior in Treatments 1 and 3 and possible contrarianism in Treatment 2. More formally,

Hypothesis 2 *S_2 types herd and act as contrarians if only if the conditions for herding and contrarianism are met.*

Buy-herding is possible in Treatment 1 and sell-herding in Treatment 3. Therefore, since we know the outcome of the random elements (noise trades and the signals for each subject) and conditional on all other subjects behaving optimally, we can calculate which action each subject should have undertaken given H_t and S , irrespective of the timing.

As stressed before, we have no theoretical prediction concerning the timing of the S_2 types' transactions. Following the discussion above, however, we conjecture that

Hypothesis 3 *The S_2 types will act later than the S_1 and S_3 .*

Hypothesis 4 *Hill-shaped S_2 -types will act before U-shaped S_2 types.*

Two-Trade Case. For the cases with two transactions, we have no theoretical predictions, even for the S_1 and S_3 types. However, we conjecture that

Hypothesis 5 *S_1 types sell twice and as soon as the treatment starts, and S_3 types buy twice and as soon as the treatment starts.*

The optimal behavior of the S_2 types is even more difficult to determine. In particular, it is not clear whether the second trade should follow the first one immediately. Following the discussion of the last section, we conjecture that

Hypothesis 6 *For the first trade the S_2 types will act later than the S_1 and S_3 types. Likewise, for the first trade hill-shaped S_2 types will act before U-shaped S_2 types.*

If the S_1 and S_3 types do not perform their second trade immediately after their first trade, then we do not have any prediction as to whether or not the S_2 types should make their second trade before or after the other two types.

Finally, it is worth comparing the treatments with one and two trades. When two trades are allowed, since subjects know that people can trade more frequently, any delay motive that they might possess will be diminished. Consequently, we would expect

Hypothesis 7 *Trades in the two-unit trade treatments, T_4 - T_6 , should be earlier than in the single trade treatments, T_1 - T_3 .*

4.4 Behavioral, non-rational predictions for the static decision

If complexity is an issue, it would seem likely that S_2 types are most likely to behave irrationally because they have to take a history and signal dependent decision. So while first hoping that the theory holds, a secondary hypothesis would be that if the theory does not provide a full explanation of behavior then it should at least do a better job at explaining the behavior of the S_1 and S_3 types than the S_2 types.

There are various behavioral theories which contradict Bayesian updating, and we also aim to examine whether, when and how any of these might explain any departures from the standard theory which incorporates Bayesian rationality. We also consider the influence of risk aversion on decision-making. Finally, price movements themselves might influence decision-making through end-point effects, since subjects were told what the possible state-values were in advance. Specifically, they know that values should not exceed 125 or fall below 75. Further to this, they must realize that when prices approach either extreme there is little to gain in buying at a price close to 125 or selling at a price close to 75. Therefore, we might consider how actions will change as prices near their extremes. While we predict that people trade earlier rather than later, when there is delay, prices may affect the decision.

In summary, should Hypotheses 1 and 2 fail, we have several further, alternative hypotheses that could be investigated. These are purposefully left as general as possible with the aim of adding more detail during the analysis of the results.

Hypothesis 8 (Complexity) *S_2 types fail to act in accordance with prescribed optimal actions when the decision is close.*

Hypothesis 9 (Risk aversion) *Subjects exhibit signs of risk aversion.*

Hypothesis 10 (Prospect theory) *Subjects exhibit signs of loss aversion.*

Hypothesis 11 (Prior expectation) *Subjects do not update their beliefs at all as prices change but act solely on the basis of their prior expectation.*

Hypothesis 12 (Probability Shifting) *Subjects update their beliefs on the basis of changing prices at a slower rate than they should.*

Hypothesis 13 (Error Correction) *Subjects account for errors that some of their peers make systematically and they react rationally to these errors.*

Note that Hypotheses 11-13 are similar: in all cases updating is slowed down. For instance, if a subject incorrectly believes that there are more noise traders than specified, then subjects would deliberately slow down their own updating process.

We have also left the hypotheses unconstrained where a specific variable is important, so we can calibrate to the data. For example, in Hypothesis 12 we have purposefully left the degree to which subjects might engage in probability shifting unstated so we can examine this for various different values; similarly for Hypothesis 13. Needless to say, by taking the value that allows each hypothesis to perform best, we can more easily rule out any hypotheses if they still provide an inadequate description of behavior. Furthermore, for reasons of parsimony, should a behavioral theory perform no better than the Bayes rational theory, we thereby have a presumption to favor the later.

5 Analysis of the Static Trade-Direction Decision

We will first examine the results by summary statistics and then expand on them with a formal econometric analysis in Section 6.

In the numbers to follow we exclude noise trades, and focus only on trades by human subjects. The total number of trades was 1993 spread over all 6 treatments; broken up by trader type we have 623 (S_1), 786 (S_2) and 584 (S_3). For T1-T3 we had 683 trades (197 S_1 , 276 S_2 and 210 S_3), for T4-T6 there were 1310 (426 S_1 , 510 S_2 and 374 S_3) trades.¹⁶

5.1 The Decision to Pass

Before we discuss the general fit of behavior towards the theoretical model, we need to consider the decision to ‘pass’. Under rationality, traders should buy if their conditional

¹⁶We had data on 23 additional trades from one treatment in one session that were excluded from the analysis due to computer error.

expectation exceeds the ask price and sell otherwise. Thus passes contradict the theoretical model.

That being said, the structure of our setup lends some additional meaning to passes. Traders are owners of a share and they have the choice to buy an extra share, or to sell the share that they already own.¹⁷ The third possibility, passing, implies that they hold on to that share, presumably in hope of making a profit on that one share. In this sense, a hold is a positive signal albeit weaker than a buy, so that a pass can be counted as a “weak buy”.

Overall there were 99 passes (< 5% of all trades), 26 from S_1 types (4% of S_1 trades); 52 from S_2 types (7%) and 21 from S_3 types (4%). While the theoretical model predicts that we should see no passes at all, we do see some; one explanation for the presence of passes could be risk aversion, and we will comment on this interpretation at length below. The total number of passes is very small; as a percentage of trades, passes occur much less frequently than in our exogenous timing companion project.

5.2 Overview of the fit of the data to the rational model

Let us start with a rough overview of decisions that are in line with rationality, as aggregated over all treatments.

The number of trades contradicting the theoretical model was about 30% when counting passes as categorically incorrect. Now suppose we admit that passes may be ‘weak buys’. Then all passes by S_1 types are still irrational, whereas all passes by S_3 types are admitted as rational. For the S_2 types, passes are admitted as rational whenever the rational action was to ‘buy’.

Examining Hypotheses 1 and 2, we note that the number of trades that conformed to the theoretical predictions was 70% treating passes as suboptimal and 73% if passes are admitted as ‘weak buys’. These numbers are similar to those in Cipriani and Guarino (2005) who obtain 73% rationality, or Anderson and Holt (1997) who have 70% rationality, albeit with a fixed-price setting. Our own companion paper, Park and SgROI (2008), with exogenous timing displayed a 70%/75% fit. This similarity to the results in the literature is noteworthy because the setting in our experiment is much more complex. Moreover, the Cipriani and Guarino (2005) experiment effectively considers only types that are equivalent to our S_1 and S_3 types. Our traders of these types actually performed better than those in Cipriani and Guarino, with rationality in excess of 80%. We might thus reasonably argue that the S_1 and S_3 types are acting in accordance with the rational static theory.

¹⁷Our rationale for this specification was to avoid explaining short-selling to subjects.

		Rational model without passes			Rational model with passes		
		S1	S2	S3	S1	S2	S3
Treatment 1, negative U-shape	correct	58	22	64	58	23	66
	wrong	3	63	16	3	62	14
	% correct	95%	26%	80%	95%	27%	83%
Treatment 2, negative hill-shape		47	66	61	47	70	61
		13	33	9	13	29	9
		78%	67%	87%	78%	71%	87%
Treatment 3, positive U-shape		66	44	49	66	55	54
		10	48	11	10	37	6
		87%	48%	82%	87%	60%	90%
Treatment 4, negative hill-shape		127	104	97	127	107	104
		25	43	24	25	40	17
		84%	71%	80%	84%	73%	86%
Treatment 5, positive U-shape		123	106	99	123	116	101
		30	77	7	30	67	5
		80%	58%	93%	80%	63%	95%
Treatment 6, negative U-shape		103	47	114	103	57	119
		18	133	33	18	123	28
		85%	26%	78%	85%	32%	81%
Total		524	389	484	524	428	505
T1-T3		171	132	174	171	148	181
T4-T6		353	257	310	353	280	324
Total%		84%	49%	83%	84%	54%	86%
T1-T3%		87%	48%	83%	87%	54%	86%
T4-T6%		83%	50%	83%	83%	55%	87%
Overall		70%			73%		

Table 1: Fit of the data to the rational, static model.

	missing herds/ contra w/o pass	missing herds/ contra w pass
T1 (U-negative)	61/82 (74%)	60/82 (73%)
T2 (hill)	10/28 (21%)	6/28 (21%)
T3 (U-positive)	4/4 (100%)	3/4 (75%)
T4 (hill)	8/27 (30%)	5/27 (19%)
T5 (U-positive)	12/16 (75%)	12/16 (75%)
T6 (U-negative)	128/173 (74%)	118/173 (68%)
Total% realized	32%	38%
Herding% realized	25%	30%
Contrarian% realized	67%	80%

Table 2: **Missing herding and contrarian trades split up by treatment.**

The S_2 types, however, often do not act rationally — almost half of their trades were against the rational model. In particular in the herding Treatments 1 and 6, the S_2 do quite poorly, even when admitting passes as weak buys. Had they taken each action at random they would have done better.¹⁸

At the same time, the S_2 types face a more difficult decision problem than the S_1 and S_3 types. Theoretically, the decisions of the S_1 and S_3 types never change, so they can take the correct decision even without following the history carefully. The S_2 types on the other hand, have to follow the history carefully and small mis-computations can cause them to be categorized as irrational.

5.3 Herding and Contrarianism

In Section 2 we have defined herding and contrarian behavior. There is a difference, however, between theoretically mandated herding and experimentally observed herding (and contrarian behavior): according to the theoretical model, a trader buys only if his expectation exceeds the ask price. Thus *theoretically mandated*, rational herd behavior arises whenever a trader’s initial expectation is below the price, his time- t expectation is above the price, and prices have risen.

Since we can compute the theoretical expectations at any stage, we know when herding or contrarian behavior is theoretically mandated, at least from the static optimization perspective (again we note that the dynamic strategies may look different). Herding and contrarian behavior can be classified either excluding passes or including passes as ‘weak buys’; we look at both cases separately.

¹⁸In the formal econometric analysis of the next section we will see that there is some persistency to their behavior.

	Number of herds						Number of contrarians					
	S1		S2		S3		S1		S2		S3	
T1	3	5%	23	28%	0	0%	0	0%	0	0%	14	21%
U-negative	60		83		0		0		0		67	
T2	5	13%	22	27%	3	10%	8	47%	20	67%	6	24%
hill	40		83		30		17		30		25	
T3	8	11%	2	17%	0	0%	2	67%	32	43%	6	16%
U-positive	70		12		17		3		74		38	
T4	17	13%	32	28%	2	4%	8	40%	22	73%	15	24%
hill	128		115		47		20		30		63	
T5	25	18%	5	25%	0	0%	5	63%	54	35%	5	6%
U-positive	139		20		15		8		156		84	
T6	17	15%	50	29%	0	0%	1	33%	0	0%	28	23%
U-negative	116		175		10		3		2		124	
Total	75		134		5		24		128		74	
T1-T3	16		47		3		10		52		26	
T4-T6	59		87		2		14		76		48	
Total%	14%		27%		4%		47%		44%		18%	
T1-T3%	9%		26%		6%		50%		50%		20%	
T4-T6%	15%		28%		3%		45%		40%		18%	

Table 3: **Herding and contrarian trades for all traders by treatment.** The second row entries list the number of possible herding trades. For instance, for the S_3 herding is only possible if prices fell. We report only the figures for which passes are counted as ‘weak buys’.

Herding and Contrarianism in the strict sense: only (static) rational behavior is considered. Table 2 summarizes our findings here. Looking only at the cases when herding is theoretically mandated we see that, if we include passes as weak buys, rational herding occurs only in 38% of the cases when it should. If we count passes as irrational, the fit is even worse with only 32% of herds occurring when they should.

The fit for contrarian behavior is better (67% when passes are wrong, 80% when passes are weak buys). The table also indicates that there are only very few rationally required contrarian trades and there are only very few rationally required herding trades for the case of U-positive signals. The reason is that both cases require that prices fall — and this generally did not happen very often.

Herding and Contrarianism in the loose sense: irrational herding is allowed.

As argued above, the S_2 types face a computationally difficult decision. We thus ask now if the S_2 types at least exhibit herding and contrarian behavior “in the right direction”, we apply the herding/contrarianism definition without looking at private expectations. So we ask, for instance, do they switch from selling to buying if prices rise and do they switch from buying to selling if prices fall? Of course, while theoretically only S_2 types can rationally herd, irrational herding and contrarianism can now be observed for all types.

Table 3 gives the raw numbers for these scenarios. When including passes, herding arises in about 18% of the cases where it is possible, contrarian behavior arises in 33% of the possible cases. Breaking these up by trader types, one can see that S_2 types have the highest propensity to herd — 27% including passes as weak buys. While there are relatively more S_1 types that act as contrarians than S_2 types (47% vs. 44%), the number of possible contrarian trades for the S_1 types is also small. Our formal regression number analysis will later show that as predicted the hill-shaped signal S_2 is the major cause for contrarianism.

Notably, the fraction of S_2 type herders is larger than that observed in Drehmann et al (2005) or Cipriani and Guarino (2005), where herding was generally irrational. Our formal econometric analysis in the next section will confirm that the U-shaped signal is the significant cause for herding behavior relative to all other types of signals.

6 Regression Analysis of the Static Decision

The summary statistics from the last section gave a good idea of the determinants of behavior: first, theoretically predicted herding does not arise as often as mandated by the theory. Second, when applying the looser experimental definition, recipients of middle signals are more likely to herd than recipients of the extreme signals. Third, contrarian behavior both according to the theoretical definition and the experimental definition is observed and occurs more frequently than herding.

We now take a closer look and run several regressions to test the direct impact of herding and contrarian signals relative to incidences of herding and contrarian trades. In particular we ask the following questions:

(1) Given that someone has a herding signal (aka a U-shaped signal), is this person more likely to herd than someone who does not have the U-shaped signal?

(2) Given that someone has a contrarian signal (aka a hill-shaped signal), is this person more likely to act as a contrarian than someone who does not have the U-shaped signal?

The random assignment of signals to traders and time slots allows us to interpret mean differences in signal-specific effects as outcomes as the average causal effects of the signal. Formally, we estimate the following equation to test the hypothesis that a type of signal, specifically U-shaped or hill shaped, is a significant cause for herding or contrarian behavior respectively:

$$\text{herd}_{i,t} = \alpha + \beta \text{u-shape}_{i,t} + \text{fixed}_i + \epsilon_{i,t}, \quad \text{contra}_{i,t} = \alpha + \beta \text{hill-shape}_{i,t} + \text{fixed}_i + \epsilon_{i,t} \quad (1)$$

where the dependent variables $\text{herd}_{i,t}$ and $\text{contra}_{i,t}$ are dummies that apply Definition 1

in the sense that they are set equal to 1 if individual i herds or acts as a contrarian respectively at trade t and 0 otherwise, α is a constant, and $\mathbf{u-shape}_{i,t}$ and $\mathbf{hill-shape}_{i,t}$ are signal dummies that are set equal to 1 if the individual received a U-shaped (for the herding estimation) or hill-shaped (for the contrarian estimation) signal. Parameter \mathbf{fixed}_i is an individual fixed effect that controls for specific traders who persistently err. Given the random assignment of signals, we can assume that $E[\mathbf{u-shape}_{i,t} \cdot \epsilon_{i,t}] = 0$ and $E[\mathbf{hill-shape}_{i,t} \cdot \epsilon_{i,t}] = 0$, the main identifying assumptions.

Overall, we restrict attention to the cases of trades where herding and contrarianism respectively are at all possible. This is reasonable since, for instance, when prices rise and a trader has signal S_3 , then such a trader cannot herd because none of his actions would satisfy the definition of herding.

We ran these regressions on a total of 10 different subsets of the data: we ran the regressions for all treatments, for treatments T1-T3 (one trade), for T4-T6 (two trades), for the first trade in T4-T6, and for the second trade in T4-T6. For these five subsets we looked at the situation where all types are included and those where only the S_2 types were included. In each scenario we estimated the model by logit and ordinary least squares regressions without fixed effects (i.e. γ_i is omitted from (1)) and then ran OLS controlling for trader fixed effects. All regressions included a constant which is significant at all conventional levels but omitted from the results tables. As a general convention we report standard errors in parentheses and a * indicates significance at the 5% level, ** signifies significance at the 1% level.

HERDING. In this specification, β represents the impact of the signal on individuals' choices of whether or not to herd, and should be positive if, as dictated by the theory, the U-shaped signal increases the probability of herding. If the inclusion of the fixed effects parameter \mathbf{fixed}_i alters the coefficients or the significance of estimates, then this indicates that the results are driven by specific individuals.

Tables 4 and 5 summarize the results from our regression. Overall, obtaining a U-shaped signal S_2 increases the probability of herding by about 29% relative to any other signal and it is significant with and without controlling for trader fixed effects; moreover, the coefficients from the logit and OLS regression are similar. Among the S_2 types, this probability remains about the same at 30.5%. This estimate is significantly different from zero at conventional levels for all subsamples with the exception of the T1-T3 trades when focusing on S_2 types trades only. Including trader-fixed effect decreases the OLS coefficients only slightly; this indicates that the estimates are not driven by some specific traders who persistently take no-herding actions.

Overall the regression confirms the hypothesis that recipients of S_2 herding-type signals

are generally more likely to herd. The effect of the signal is also much stronger than that in the comparable exogenous timing setting: in our companion paper Park and SgROI (2008) we find a marginal effect of about 6%.

To obtain a complete picture, we also estimated

$$\text{herd}_{i,t} = \alpha + \beta_2 \text{u-shape}_{i,t} + \beta_1 S_1 \text{ signal}_{i,t} + \beta_3 S_3 \text{ signal}_{i,t} + \epsilon_{i,t} \quad (2)$$

where variables $S_1 \text{ signal}_{i,t}$ and $S_3 \text{ signal}_{i,t}$ are dummies for signals S_1 and S_3 . Estimated by OLS this yields coefficients $\hat{\beta}_2 = .278^{**}(.034)$ on herd-signals, $\hat{\beta}_1 = -.110^{**}(.031)$ on S_1 and $\hat{\beta}_3 = -.163^{**}(.043)$ on S_3 . So obtaining the herding signal strongly increases the chance of acting as a herder, and the probability of herding decreasing for all other signals.

CONTRARIANISM. Next, we estimate equation (1) to test the hypothesis that a hill-shaped signal is a significant explanation of contrarian behavior. Our theory predicts that the coefficient β is positive so that a hill-shaped signal indeed has a larger impact on the occurrence of contrarianism relative to other kinds of signals.

Tables 6 and 7 summarize the results from our regression. Obtaining the hill-shaped S_2 signal increases the chance of acting as a contrarian by about 36.1% relative to any other kind of signal as the marginal effect at the mean. As with herding, the OLS coefficient changes only slightly when we include trader fixed effects. All coefficients are significant at the 1% level. When restricting attention to the sample of S_2 types the marginal effect decreases slightly to 31%, and OLS and marginal logit estimates coincide. All coefficients are significant at the 1% level except for the T1-T3 subsample of S_2 types trades, where we have significance at the 5% level (due to the large standard error). Including fixed effects renders the coefficient insignificant for the second trades, restricted to only S_2 types; this is again due to the small number of relevant data points and the large standard error that is involved.

Overall we confirm that the hill-shaped signal is the significant source of contrarianism relative to all other signals. The marginal effect is comparable to that in the exogenous time setting (Park and SgROI (2008)), which is estimated to be about 34%.

As with herding, to obtain a complete picture, we also estimated

$$\text{contra}_{i,t} = \alpha + \beta_2 \text{hill-shape}_{i,t} + \beta_1 S_1 \text{ signal}_{i,t} + \beta_3 S_3 \text{ signal}_{i,t} + \epsilon_{i,t}. \quad (3)$$

Coefficients here are $\hat{\beta}_2 = .312^{**} (.060)$ for the hill-shaped signal dummy, $\hat{\beta}_1 = -.0414 (.065)$ (insignificant at all conventional levels) for the signal S_1 dummy and $\hat{\beta}_3 = -.191^{**} (.033)$ for the S_3 dummy. So again, the hill-shaped S_2 signal is clearly identified as the major cause of contrarianism.

	all types	T1-T3	T4-T6	first trades T4-T6	second trade T4-T6
Logit	0.292** (-0.022)	0.114** (-0.032)	0.397** (-0.032)	0.228** (-0.025)	0.446** (-0.05)
OLS	0.378** (-0.025)	0.138** (-0.039)	0.495** (-0.031)	0.293** (-0.03)	0.552** (-0.043)
OLS fixed effects	0.352** (-0.027)	0.081 (-0.042)	0.434** (-0.038)	0.276** (-0.032)	0.545** (-0.057)
Observations	1172	391	781	805	367
R^2 OLS	0.16	0.03	0.25	0.11	0.31
R^2 fixed	0.39	0.67	0.52	0.40	0.74

Table 4: **The Effect of U-Shaped Signals on the Probability of Herding.** The table represents regressions of the occurrence of a herding trade on the trader receiving a U-shaped signal. Logit regressions report the marginal effects. OLS fixed effects regressions control for trader-fixed effects. The data is restricted to include only trades that could be herding trades. Standard errors are in parentheses, * indicates significance at the 5% level, ** at the 1% level.

7 The Impact of the Price

While the general behavior of the S_1 and S_3 types is in line with the theoretical static model (over 80% of their traders are ‘rational’), we do observe that S_3 types engage in selling and that S_1 types engage in buying. We now want to assess if this behavior is systematic. Specifically, we test whether an increase in the price changes the probability of a specific trade. Theoretically, the price should have no impact on the decision because S_1 types should always sell and S_3 types should always buy. We therefore estimate the following regression

$$\text{trade}_{i,t} = \alpha + \beta \Delta \text{price}_{i,t} + \epsilon_t, \quad (4)$$

where $\text{trade}_{i,t}$ is a dummy that is 1 if there is a buy or hold, and 0 when there is a sale, and the independent variable $\Delta \text{price}_{i,t}$ is the percentage change of the price from 100, i.e. the price at the time of the trade divided by 100 and subtracting 1.¹⁹ Considering the optimal static decision at the time at which people trade we assert that $E[\Delta \text{price}_{i,t} \cdot \epsilon_{i,t}] = 0$.

We estimated the model by logit separately for the three signals. The main variable of interest in (4) is β which measures whether a rising price affects the probability of a trader buying or selling. Our static theory predicts that the price should have no impact on whether an S_1 or S_3 type buys or sells. Consequently, parameter β should be insignificantly different from zero. In contrast, if it is not zero, then we gain insights

¹⁹We ran two unreported control regressions: in the first, we dropped all holds, and in the second, we estimated a specification in which buys received value 2, holds 1 and sales 0 (this specification was estimated using a multinomial logit). The direction of the results and the significance of the coefficients remained unaffected.

	only S_2	T1-T3	T4-T6	first trades T4-T6	second trade T4-T6
Logit	0.305** (-0.049)	0.035 (-0.066)	0.446** (-0.066)	0.194** (-0.055)	0.565** (-0.104)
OLS	0.284** (-0.044)	0.034 (-0.066)	0.406** (-0.054)	0.186** (-0.052)	0.493** (-0.075)
OLS fixed effects	0.315** (-0.055)	0.031 (-0.069)	0.409** (-0.081)	0.236** (-0.071)	0.591** (-0.123)
Observations	481	166	315	330	151
R-squared OLS	0.08	0	0.15	0.04	0.22
R-squared Fixed	0.47	0.91	0.61	0.55	0.89

Table 5: **The Effect of U-Shaped Signals on the Probability of Herding, only S_2 types.** The table represents regressions of the occurrence of a herding trade on the trader receiving a U-shaped signal. Logit regressions report the marginal effects. OLS fixed effects regressions control for trader-fixed effects. The data is restricted to include only trades that could be herding trades and that are made by S_2 types.

about systematic herding or contrarian behavior. For instance, consider type S_3 . If the sign of β is negative, then this type becomes less likely to buy as prices increase. Such behavior tentatively indicates systematic contrarian behavior. Likewise, if β were positive for the S_1 type, then this implies that the S_1 types are more likely to buy when prices rise; this is a tentative herding effect.

For the S_2 types, however, there should be an effect, the sign depending on the shape of the signal distribution. Specifically, for a hill-shaped signal distribution, the S_2 type should become less likely to buy as the price increases (for prices above 100, the coefficient should be zero), for a negative U-shaped distribution the probability should be increasing in the price, and for a positive U-shaped distribution it should be increasing (for low prices, this type should be sell-herding).

Table 8 summarizes the results of our estimation, Table 9 splits up the S_2 trades by signal type.

We find that for all types of signals, an increase in the price decreases the chance that the trader buys: all estimates are significantly different from zero at the 1% level. This is contrary to the theory for the S_1 and S_3 types who should always be selling and buying respectively, irrespective of the price. While the sign of the estimate for the hill-shaped type has the right sign, the sign is wrong for the U-shaped types. This is, of course, not surprising, given that even the S_3 types (the high types) act in a contrarian manner, so it is reasonable to expect the same behavior from the S_2 types.

	all types	T1-T3	T4-T6	first trades T4-T6	second trade T4-T6
Logit	0.361** (-0.056)	0.304** (-0.081)	0.419** (-0.082)	0.358** (-0.064)	0.371** (-0.114)
OLS	0.434** (-0.057)	0.353** (-0.085)	0.508** (-0.079)	0.439** (-0.066)	0.429** (-0.114)
OLS fixed effects	0.406** (-0.063)	0.300* (-0.117)	0.473** (-0.108)	0.405** (-0.076)	0.655** (-0.177)
Observations	820	293	527	553	267
R-squared OLS	0.07	0.06	0.07	0.07	0.05
R-squared Fixed	0.41	0.67	0.52	0.51	0.79

Table 6: **The Effect of Hill-Shaped Signals on the Probability of Acting as a Contrarian.** The table represents regressions of the occurrence of a contrarian trade on the trader receiving a hill-shaped signal. Logit regressions report the marginal effects. Fixed effects regressions control for trader-fixed effects. The data is restricted to include only trades that could be classified as contrarian.

8 The relation between the first and the second trade

People have the opportunity to make so-called ‘return’ trades by sell first and then buying later or vice versa. This way, they can realize a trading profit in the process. At the same time, if they make such a ‘return’ trade then by the end of the treatment they still own one unit of the security.

Table 10 provides summary statistics for the second trade with emphasis on the return trades (recall that return trades arise only in T4-T6).

Not all traders act twice — in about 6% of cases they forego the second trading opportunity. About 76% of all second trades go in the same direction of the first. The remaining 24% are return trades, most of which are performed by the S_2 types. The S_2 also account for the largest fraction of all type-specific trades (33% of all S_2 second trades, 19% and 16% for the S_1 and S_3 types respectively). About 77% of the return trades yield a trading profit which suggests that return-trades were performed on the basis of “buy low, sell high” (or “sell high, buy low”). Now suppose we account also for expected payoffs, i.e. we employ the expected static profit of a share at the time of each trade. With this measure, only 49% of the return trades are profitable and the S_2 types do particularly poorly.

This raises the question of whether there are systematic features of the return trades.²⁰

²⁰In some unreported regressions we analyzed whether the payoffs from return trades are smaller when they occur late. For this we regressed factual payoffs from return trades and also expected payoffs from return trades on the trading time. Yet we found no significant relation and thus concluded that time has no impact on the payoff of return trades. Moreover, we also could find no general relationship between return trades and time for the S_2 and S_3 types. However, for the S_1 types we found a relationship that is

	only S_2	T1-T3	T4-T6	first trades T4-T6	second trade T4-T6
Logit	0.313** -0.074	0.213* -0.108	0.402** -0.108	0.323** -0.087	0.300* -0.145
OLS	0.312** -0.069	0.213* -0.105	0.394** -0.094	0.324** -0.081	0.292* -0.134
OLS fixed effects	0.227* -0.096	0 -0.113	0.343 -0.19	0.224* -0.111	0.667 -0.43
Observations	305	110	195	208	97
R-squared OLS	0.06	0.04	0.08	0.07	0.05
R-squared Fixed	0.62	0.96	0.61	0.82	0.97

Table 7: **The Effect of Hill-Shaped Signals on the Probability of Acting as a Contrarian.** The table represents regressions of the occurrence of a contrarian trade on the trader receiving a hill-shaped signal. Logit regressions report the marginal effects. Fixed effects regressions control for trader-fixed effects. The data is restricted to include only trades that could be classified as contrarian and that are made by S_2 types.

Return trades are always either buy-sell or sell-buy. So we ask if there is a particular fashion in which the different signal types perform their return trades.

There is a specific interpretation that can be attached to return trades: for the S_1 types a buy-sell return trade has a manipulative connotation in that this type may try to drive up the price so that he can sell high later-on; similarly for sell-buy return trades by the S_3 types. The reverse order is either contrarian or herding, where the difference between contrarianism and herding is determined by the general movement of prices. However, for the S_2 types it is not clear how one would classify manipulative behavior with the order of buys and sells. To get to the bottom of the common direction and the connotation, we ran the following regressions

$$\text{herd}_{i,t} = \alpha + \beta \text{return trade}_{i,t} + \epsilon_{i,t}, \quad \text{contra}_{i,t} = \alpha + \beta \text{return trade}_{i,t} + \epsilon_{i,t}, \quad (5)$$

where the dependent variables $\text{herd}_{i,t}$ and $\text{contra}_{i,t}$ are the herding and contrarian dummies from equations (1), α is a constant, and $\text{return trade}_{i,t}$ is a dummy for the incidence of a return trade.

In each scenario we estimated the model by logit, restricted for the individual signals and the incidences of the second trades; we report the marginal effects. Table 11 summarizes our findings.

significant at the 1%-level. This finding is intuitive in the context of a different finding (which we report below) in that the S_1 types tend to act as herders when they perform a return trade. This implies that the S_1 types first sold, then waited and observed that prices were increasing. They then bought back the share that they sold.

		logit	OLS	OLS fixed
S1	Δprice	-1.287** (-0.193)	-1.317** (-0.204)	-1.303** (-0.232)
	Constant	-0.110** (-0.016)	0.271** (-0.022)	0.270** (-0.023)
	Observations R-squared	623	623 0.06	623 0.5
S2	Δprice	-3.381** (-0.322)	-2.850** (-0.23)	-3.301** (-0.268)
	Constant	0.208** (-0.03)	0.678** (-0.023)	0.711** (-0.024)
	Observations R-squared	786	786 0.16	786 0.5
S3	Δprice	-1.299** (-0.159)	-1.681** (-0.192)	-1.210** (-0.237)
	Constant	0.260** (-0.015)	0.945** (-0.016)	0.922** (-0.017)
	Observations R-squared	583	583 0.12	583 0.5

Table 8: **The decision to buy as a function of the price.** The table displays the results from a logit regression of the decision to buy on the price change, as measured relative to the prior, $(p_t - 100)/100$, for all S_1 , S_2 and S_3 types respectively. For all traders the probability of a buy is decreasing as the price rises.

Overall we observe that S_1 types act as herders and the S_3 types act as contrarians. For the S_2 types we observe a negative herding coefficient (significant at the 5% level) and a positive contrarian coefficient (significant at the 1% level). Thus the likelihood of a S_2 type herding trade is reduced by the incidence of a return trade and instead, an incidence of a return trade increases the probability that the trade is contrarian.

9 Analysis of Timing

As outlined in our discussion in Section 3, there is little theory to guide our analysis of subjects' timing decisions; Hypotheses 1 to 7 formulate those assertions that we feel comfortable making.

The strongest interpretation of the theory is that S_1 and S_3 types should trade immediately when the session starts. Consequently, according to this view we should observe that all these types trade within the first few seconds of the game. And indeed, we do observe a very large number of trades at the very beginning of trading. Overall, the S_1 and S_3 types account for 1206 transactions. Within the first few seconds of trading we

	U-negative	Hill	U-Positive
Δ price	-1.632** (-0.576)	-2.706** (-0.404)	-4.060** (-0.73)
Constant	-0.012 (-0.059)	0.129** (-0.048)	0.330** (-0.035)
Observations	266	245	275

Table 9: **The decision to buy as a function of the price for the S_2 types.** The table displays the results from a logit regressing the decision to buy on the price change, as measured relative to the prior, for the S_2 types split up by signal-type.

	S1	S2	S3	All
Total trades	425	510	373	1308
First trades	222	262	190	674
Second trades	203	248	183	634
Return trades	39	82	29	150
Same direction	164	166	154	484
Forgone second trade	19	14	7	40
Percent return	19%	33%	16%	24%
Percent return of all trades	9%	16%	8%	11%
Percent foregone	9%	5%	4%	6%
Profitable return	32	64	19	115
Percent profitable return	82%	78%	66%	77%
Profitable return including expectations	23	29	21	73
Percent profitable return including expectations	59%	35%	72%	49%

Table 10: **Return Trades.**

observe the following number of transactions (note that the smallest time unit that our software supports is a second).

	<1 sec.	1 sec.	2 sec.	3 sec.	4 sec.	5 sec.
Number of trades	131	111	90	72	51	57
by S_1	28	46	30	37	21	16
by S_3	72	30	30	13	12	20

Overall, 355 out of the 1206 trades made by S_1 and S_3 types occur within the first five seconds. Yet while this number is substantial, it is less than 30% of the trades, and so we cannot support the hypothesis that *all* of the S_1 and S_3 types trade immediately.

We now attempt to identify systematic differences in the timing behavior for the various signal types and treatment settings. Specifically, we will compare the cumulative distributions of the trade-times for different categories of types. The strongest result that one can hope for in this context is that one cumulative distribution function (henceforth,

	S1	S2	S3
herding case	0.085** (-0.028)	-0.129* (-0.066)	
contrarian case		0.181** (-0.039)	0.312** (-0.064)
Constant for herding case	-0.166** (-0.028)	-0.128** (-0.03)	
Constant for contrarian case		-0.276** (-0.017)	-0.252** (-0.027)
Observations	205	248	182

Table 11: **Trade-Direction of Return Trades.** The table condenses six regressions (by type and then with respect to herding and contrarian behavior separately). For the S_1 types there was no simultaneous occurrence of contrarian behavior and return trades; similarly for the S_3 types and herding behavior. Hence the empty cells.

cdf) of trade-times stochastically dominates another: distribution F first order stochastically dominates distribution G if G is larger than F for all entry times. If we indeed observe that F first order stochastically dominates G , then we can say that the entry times under F are systematically later than under G .

Figure 1 provides plots of the relevant cdfs. Specifically, we computed the cdfs for the following subsamples:

- D1 all treatments, separated by types S_1, S_2, S_3 ,
- D2 T1-T3 and first trades in T4-T6, separated by types S_1, S_2, S_3 ,
- D3 T1-T3 separated by types S_1, S_2, S_3 ,
- D4 T4-T6 separated by types S_1, S_2, S_3 ,
- D5 T4-T6 first trades, separated by types S_1, S_2, S_3
- D6 T4-T6 second trades, separated by types S_1, S_2, S_3
- D7 T1-T3 not separated by types
- D8 T4-T6 not separated by types
- D9 T4-T6 first trades only, not separated by types
- D10 all treatments, only for S_2 , separated by hill, U-negative and U-positive
- D11 T1-T3 and first trades in T4-T6, only for S_2 , separated by hill, U-negative, and U-positive

We will now address the systematic timing issues by referring to the above sets of distributions.

Question 1: *Which type acts earliest, which latest?* To answer this, we compared the timing-cdfs split up by types. All panels in Figure 1 save for the bottom row, right panel

are very clear: the S_1 and S_3 are indistinguishable, and the S_2 trade after the S_1 and S_3 because the S_2 type's distribution of entry times is below the entry time distribution for the S_1 and S_3 . Specifically, this holds for all treatments D1, for all first trades D2, for treatments T1-T3, T4-T6, and for the first trades only in T4-T6. Consequently our findings comply with Smith (2000)'s prediction that people with good-news or bad-news-signals trade early, and that people who receive mixed information delay.

Question 2: *Is there a timing order for the second trade?* The bottom row, right panel in Figure 1 indicates that there is no timing relation between the three types when it comes to the second trade. This non-finding suggests that settings in which people can trade more than once would likely not yield clear insights with respect to people's timing behavior.

Question 3: *Will people trade faster when they can trade more often?* To answer this question we compare D7 to D8 and D7 to D9. The left and middle panels in Figure 2 again paint a clear picture. Allowing people to trade more often speeds up their trade-times in particular when looking only at the first trade in T4-T6.

Question 4: *For the S_2 , which type of signal will trade earliest, which latest?* To answer this question we regard the cdfs in D10 and D11; Figure 3 visualizes these cdfs. Again, the picture here is clear: recipients of the hill-shaped signal act earliest, whereas recipients of the U-positive signal act last.

In Section 3 we argued that one may expect to see the hill-shaped types act before the U-shaped type since they would generally act against the general movement of prices. However, note that when prices rise such trades are not contrarian trades as they are in line with the original, slightly negative opinion.

In summary, the S_2 types act systematically after the S_1 and S_3 types, and the hill-shaped types act before the U-shaped types.

Clustering. One significant feature observed during the experiment is that trades are often clustered. We have already explained that many, though not all, trades occur at the very beginning of the treatment. There also appears to be leader-follower trading in the sense that when one trade occurs, others follow in very quick succession. This suggests that people wait for someone to make a move with the intention of immediately reacting with a trade.

To get a sense of systematic behavior, we first determined the number of seconds between one trade and its predecessor. We then looked at the frequencies for which trades occur within 1.5 seconds of one another, split up by signal types. Table 12 lists the relevant percentages, Figure 4 plots the frequencies for the case where all signal types are taken together.

As can be seen, there are a large number of trades that occur within 1.5 seconds of each other; the concentration is strongest at the very beginning of the treatment and changes only very little as the treatment progresses. The concentration at the beginning is no surprise as about 25% of all trades occur within the first 5 seconds (it takes another 25 seconds before a total of 50% of trades are made).²¹ The table also illustrates that there is no measurable variation among types; unreported regressions confirm this insight. We thus confirm that there is a general tendency to trade in a clustered manner, though signals have no impact on the clustering.

	All	S1	S2	S3	hill	-ve U	+ve U
All times	67%	66%	63%	71%	64%	66%	61%
total time >5 sec	58%	56%	57%	62%	57%	60%	55%
total time >10 sec	54%	52%	53%	58%	55%	55%	50%
total time >20 sec	51%	48%	51%	55%	56%	50%	49%
total time >30 sec	50%	44%	50%	54%	56%	49%	46%

Table 12: **Clustering of Trades: The percentage of trades that follow within 1.5 seconds of another.**

10 Conclusion

Herding has long been suspected to play a role in financial market booms and busts. Recent theoretical work shows that informational herding (or contrarianism) is possible in a market with efficient asset prices if the conditional signal distribution for traders has a specific shape. To keep this theory tractable one needs strong assumptions such as an exogenous entry order. Having explored financial market herding with exogenous timing elsewhere, we address how behavior changes when traders can choose the time of their trade. Giving traders a choice of when to act is not only natural, but there are also important insights that can be gleaned from such an analysis. For instance, if herding-prone types delay their actions systematically, then herd behavior can become more pronounced and significant compared to exogenous timing settings.

We thus examined the results of an experimental test involving almost 2000 trades, exploring people's behavior when they can not only choose *how* to trade, but also *when* to trade. By contrast with almost all existing experiments, we employ a theoretical framework that allows rational herding and contrarianism in a financial market environment with moving prices.

²¹The results are similar when looking at the trades in intervals, e.g. those that occur between 20 and 40 seconds etc.




Figure 1: Plots for the timing cdfs by signal types and treatments.




Figure 2: Plots for the timing cdfs by number of trades.




Figure 3: Plots for the timing cdfs of the $S2$ by signal shape.




Figure 4: Plots for the time-difference pdfs.

We found that subjects' decisions were generally in line with the predictions of exogenous-time financial herding theory when that theory admits rational herding and contrarianism. While herding and contrarianism do not arise as frequently as predicted by theory, types theoretically prone to herd or be contrarian are the significant source of this kind of behavior when it does arise. Moreover, compared to an exogenous timing setting, the herding signal has a much larger impact on the likelihood of a herding action.

Types with extreme information (good or bad) trade systematically earliest, and those with signals conducive to contrarianism trade earlier than those with information conducive to herding. We thus find strong evidence for the impact of the type of information both with respect to the direction and the timing of trades.

Finally, examining alternative models of behavior (such as risk aversion, loss aversion or probability shifting) leaves us with the impression that most of these specifications either perform poorly or add little to our understanding over and above the risk-neutral rational model. Moreover since alternative theories often require additional parameters, for reasons of parsimony our presumption remains in favor of the standard model of Bayesian rationality. The exception are models of behavior which essentially result in subjects putting too much weight on their own information or equivalently slowing down their updating, which can partially explain why fewer herding or contrarian trades occur in practice than predicted by the rational theory.

A Alternative Explanations for Trading Behavior

We have seen in Section 5 that some results are supportive of the static theory, confirmed by a formal regression analysis in Section 6. Yet it is also well-established in experimental work, that models with Bayesian rationality and risk-neutral agents may not provide the best fit for the data.

The general assumption of our model is that people are risk-neutral. As a first check we will see if this assumption is warranted. Next, we will analyze if loss-aversion may play a role in people's behavior. Finally we will check if various forms of alternative information updating provide a better fit with the data. These approaches usually depend on some parameter(s). Our approach is to vary this parameter and see how the variation improves the overall fit of the alternative model to the data. In this appendix we focus on the static decision only.

A.1 Risk and Loss Aversion

Risk Aversion. One persistent finding from the Section 7 is that traders exhibit a general tendency to act as contrarians. One might thus entertain the idea that traders act as contrarians because of risk-aversion. We can go about examining this by computing the optimal action when people have a concave utility function. We checked this employing both CARA and CRRA utility functions:

$$\text{utility}_{\text{CARA}}(\text{payoff}|\text{action}) = -e^{\rho \cdot \text{payoff}}, \quad \text{utility}_{\text{CRRA}}(\text{payoff}|\text{action}) = \frac{\text{payoff}^{1-\gamma}}{1-\gamma}.$$

Theoretically, the CARA utility function is the superior choice in the framework since we can ignore income effects.

For each type we determined the optimal action given the respective utility function and compared it to the action taken by the subjects. Within a setup with risk-aversion, a pass is indeed an action that has payoff consequences and may be optimal for some posterior probabilities. Usually, as prices (and thus the probability of a high outcome) rise, the optimal action changes from a buy to a pass to a sell. Risk-aversion biases decisions against buys and holds, because sells yield an immediate cash flow, whereas holding the stock exposes the subject to the risky future payoff. The larger the risk-aversion coefficient, the stronger the bias against buying.

Computing the expected utilities we find, however, that the performance of a model with risk aversion is worse for all reasonable levels of risk aversion. For CRRA with log-utility ($\gamma = 1$), it is 67% , which is below the risk-neutral model (70%) and the fit is only 42% for the S_2 types; for CARA with $\rho = 2$ it is 51% (the fit rises as ρ declines). As ρ declines, we capture more of the behavior by S_3 types but less of the behavior by S_2 types. Note that as ρ falls, we move closer to risk neutrality.

Overall, we conclude that the assumption of risk-neutrality captures behavior quite well, with risk-aversion playing at most a negligible role.

Loss-Aversion — S-Shaped Valuation Functions. A host of experimental work in prospect theory following Kahneman and Tversky (1979) has indicated that people pick

	Total Number of wrong decisions CRRA utility, $\gamma = 1$ (log-utility)			Total Number of wrong decisions CARA utility, $\rho = 2$		
	S1	S2	S3	S1	S2	S3
T1	58	61	14	58	22	62
U-negative	3	24	66	61	85	80
	95%	72%	18%	95%	26%	78%
T2	47	53	9	47	62	60
hill	13	46	61	60	99	70
	78%	54%	13%	78%	63%	86%
T3	66	33	6	66	25	49
U-positive	10	59	54	76	92	60
	87%	36%	10%	87%	27%	82%
T4	127	90	17	126	97	98
hill	25	57	104	152	147	121
	84%	61%	14%	83%	66%	81%
T5	123	59	5	123	82	98
U-positive	30	124	101	153	183	106
	80%	32%	5%	80%	45%	92%
T6	103	120	28	103	46	110
U-negative	18	60	119	121	180	147
	85%	67%	19%	85%	26%	75%
Total%	84%	53%	14%	84%	42%	82%
T1-T3%	87%	53%	14%	87%	39%	81%
T4-T6%	83%	53%	13%	83%	44%	82%
Fit total	51%			67%		
fit T1-T3	51%			66%		
fit T4-T6	51%			67%		

Table 13: **Risk-Aversion Analysis.**

choices based on change in their wealth rather than on levels of utilities. These costs and benefits of changes in wealth are usually assessed with valuation functions that are S-shaped. Kahnemann and Tversky suggested the following functional form

$$V(\Delta\text{wealth}|\text{action}) = \begin{cases} (\Delta\text{wealth})^\alpha & \text{for } \Delta\text{wealth} \geq 0 \\ -\gamma(-\Delta\text{wealth})^\beta & \text{for } \Delta\text{wealth} < 0 \end{cases}$$

where Δwealth is the change in wealth and α, β, γ are parameters. A common specification for the parameters stemming from experimental observations is $\alpha = \beta = 0.8$ and $\gamma = 2.25$ (Tversky and Kahneman (1992)).

As with risk aversion, the performance of this model applied to our setup is much worse than the performance of the rational model. For parameters as estimated by Tversky and Kahneman (1992), the fit is below 49%. Table 14 illustrates this observation for the above parameters as well as for one other configuration.²²

²²Arguably, we are only using one part of the tools developed in prospect theory, S-shaped valuations, and ignore that other component, decision weights. However, the latter have a relation to re-scaled probabilities which we analyze separately.

	Total Number of wrong decisions prospect theory, $\alpha = \beta = 0.8, \gamma = 1$			Total Number of wrong decisions prospect theory, $\alpha = \beta = 0.8, \gamma = 2.25$		
	S1	S2	S3	S1	S2	S3
Treatment 1 negative hill-shape	20 36%	81 81%	37 51%	22 40%	82 82%	37 51%
Treatment 2 increasing	31 42%	57 63%	36 53%	31 42%	71 79%	57 84%
Treatment 3 negative U-shape	21 35%	69 73%	37 49%	21 35%	68 72%	67 88%
Treatment 4 decreasing	41 71%	55 56%	33 45%	41 71%	55 56%	48 65%
Treatment 5 positive U-shape	33 48%	70 71%	32 49%	33 48%	73 74%	46 71%
Treatment 6 negative hill-shape	41 47%	60 71%	22 38%	41 47%	60 71%	22 38%
Total number wrong wrong percentage	187 46%	392 69%	197 48%	189 47%	409 72%	277 67%
Total model fit	43.8%			36.7%		

Table 14: **Loss-Aversion Analysis.**

A.2 Decision Rule: Prior Actions or No Updating

One alternative decision rule formulation is that of naïve traders who ignore the history and who simply stick to their prior action. As such, S_1 types always sell, S_3 types always buy and S_2 types pick the action that is prescribed at the initial history. For instance, with negative U-shape, S_2 traders always sell.

This specification does no better than the rational model, fitting 71% of the data; broken up by type the fit is similar to the rational model. Moreover, with this alternative model, we cannot accommodate passes as ‘weak buys’ because this would be contrary to the spirit of ‘no changes of the action’. Indeed this illustrates the first weakness: a model based on people choosing their prior action will not help us to understand any changes in behavior that might have occurred, in particular not for S_1 and S_3 types. Since the econometric analysis has already revealed that traders are sensitive to the price, this decision rule is rather weak.

A weaker variation of the ‘stick to the prior action’-theme has traders ignore the history altogether but remain mindful of the price. Traders thus act based only their prior expectation: if the price exceeds it, they sell, if the price is below it, they buy.

And indeed about 75% of people take an action that is in accordance with their prior expectation. For instance, for the S_3 types this means that they do not buy when they should be buying, or for the S_2 types that they do not herd when they should be herding.

A.3 Probability Scaling and Shifting

A yet weaker version of the no-updating alternative rule is probability shifting, whereby traders underplay (overplay) low (high) probabilities coming from the observed history

	No updating			prior action		
	S1	S2	S3	S1	S2	S3
T1	60	62	64	58	63	64
U-negative	61	85	80	61	85	80
	98%	73%	80%	95%	74%	80%
T2	47	61	63	47	53	61
hill	60	99	70	60	99	70
	78%	62%	90%	78%	54%	87%
T3	66	47	47	66	46	49
U-positive	76	92	60	76	92	60
	87%	51%	78%	87%	50%	82%
T4	134	108	97	127	74	97
hill	152	147	121	152	147	121
	88%	73%	80%	84%	50%	80%
T5	123	81	97	123	80	99
U-positive	153	183	106	153	183	106
	80%	44%	92%	80%	44%	93%
T6	103	115	118	103	89	114
U-negative	121	180	147	121	180	147
	85%	64%	80%	85%	49%	78%
Total%	86%	60%	83%	84%	52%	83%
T1-T3%	88%	62%	83%	87%	59%	83%
T4-T6%	85%	60%	83%	83%	48%	83%
Fit total	75%			71%		
fit T1-T3	76%			74%		
fit T4-T6	75%			69%		

Table 15: **No Updating and Prior Actions.**

H_{t-1} . Alternatively, traders may overstate the probabilities of their prior expectations; we present results from the latter but point out, that the former yields similar insights. The usual symmetric treatment of this under- or overstating of probabilities is to transform probability p into $f(p)$ ²³ as follows

$$f(p) = \frac{p^\alpha}{p^\alpha + (1-p)^\alpha}.$$

Parameter values $\alpha > 1$ are associated with S-shaped re-valuations (high probabilities get overstated, low probabilities understated), $\alpha < 1$ with reverse S-shaped valuations (high probabilities get understated, low probabilities overstated). Note that transformation $f(p)$ applied to probabilities of all three states do not yield a probability distribution. However, when employed properly in the conditional posterior expectation the transformation achieves the effect of a probability distribution.

Consequently, when modeling an overconfident trader who puts more weight on his

²³There are various other forms for these switches, e.g. non-symmetric switches where the effects are stronger (or weaker) for larger probabilities. The interpretation and implementation of such asymmetric shifts does, however, become difficult if not impossible with three states. Of the various possible specifications we only pick a few as the spirit of all re-scalings is similar: updating is slowed.

In f , one re-scales p^α by itself and the counter-probability; alternatively, if p_i signifies the probability of one state, one could imagine a re-scaling by p_j^α for all states, $j = 1, \dots, 3$.

prior signal we would apply an $\alpha > 1$ re-scaling on the initial probabilities. Alternatively, one can also model slow updating directly by applying an $\alpha < 1$ re-scaling to the posterior probabilities. Of course the effect will be similar: in both cases the histories or updated probabilities would be less important to traders than under the rational model. We considered both specifications.

Here we report the results where $\Pr(V|H_1) \times \Pr(S|V)$ has been re-scaled with an $\alpha > 1$; downward scaled probabilities of the history $\Pr(V|H_t)$ yield similar insights.

	With $\alpha = 25$			With $\alpha = 10$			With $\alpha = 5$		
	S1	S2	S3	S1	S2	S3	S1	S2	S3
T1	58	60	66	58	35	66	58	25	66
U-negative	3	25	14	3	50	14	3	60	14
	95%	71%	83%	95%	41%	83%	95%	29%	83%
T2	47	69	61	47	69	61	47	68	61
hill	13	30	9	13	30	9	13	31	9
	78%	70%	87%	78%	70%	87%	78%	69%	87%
T3	66	59	54	66	59	54	66	59	54
U-positive	10	33	6	10	33	6	10	33	6
	87%	64%	90%	87%	64%	90%	87%	64%	90%
T4	127	110	104	127	110	104	127	110	104
hill	25	37	17	25	37	17	25	37	17
	84%	75%	86%	84%	75%	86%	84%	75%	86%
T5	123	124	101	123	124	101	123	121	101
U-positive	30	59	5	30	59	5	30	62	5
	80%	68%	95%	80%	68%	95%	80%	66%	95%
T6	103	99	119	103	77	119	103	62	119
U-negative	18	81	28	18	103	28	18	118	28
	85%	55%	81%	85%	43%	81%	85%	34%	81%
Total%	84%	66%	86%	84%	60%	86%	84%	57%	86%
T1-T3%	87%	68%	86%	87%	59%	86%	87%	55%	86%
T4-T6%	83%	65%	87%	83%	61%	87%	83%	57%	87%
Fit total	78%			75%			74%		
fit T1-T3	79%			75%			74%		
fit T4-T6	77%			75%			74%		

Table 16: **Overweighting of the Prior.**

Comparing the results here to those in Table 1, one can see that the fit of prior overweighting hardly improves for the S_1 and S_3 types. Moreover, while the total fit does improve relative to the rational model, it does not improve dramatically. Most of the improvement stems from contrarian trades that are now given a rationale. At the same time, re-scaling does a poor job explaining herd-behavior of any sort.

A.4 Error Correction Provisions

Inspired by level K reasoning (see Costa-Gomes, Crawford and Broseta (2001)) and Quantal Response Equilibria (see McKelvey and Palfrey (1995) and McKelvey and Palfrey (1998)), we will contemplate an alternative specification for hampered updating in which agents do not trust that their peers act fully rationally. In the rational model, consider a buy without herding in state V_i : this event occurs with probability $\beta_i =$

$.25/2 + .75 \cdot \Pr(S_3|V_i)$ (recalling that $.25/2$ is the probability of a noise buy). Now imagine that instead subjects believe that only fraction δ of the informed buyers act rationally and that the remaining $1 - \delta$ take a decision at random. Then the probability of a buy in state V_i becomes

$$\beta_i = .25 + .75((1 - \delta)/2 + \delta \cdot \Pr(S_3|V_i)).$$

The task is then to find the δ for which this specification yields the best fit with the data. We obtained the best fit for $\delta = 2/15$. However, compared to the rational model the improvement of the fit is minor (see Table 17): the rational fit is 70% vs. 73% with error correction provisions.

An alternative interpretation for this error correction is that the level of noise trading is perceived higher than it actually is because other subjects act randomly: a δ of $2/15$ translates into a factual noise level of 90%. As the informational impact of each transaction on the subject's beliefs is dampened, after any history the private signal has a larger impact than under the rational model. This specification is thus in spirit similar to probability shifting, but focuses on the idea that subjects believe that others either ignore their signals or are simply unable to interpret it correctly.

A variation on this error correction theme is a specification in which a subject believes that fraction $1 - \delta$ act randomly but the subject assumes that the remaining fraction δ takes this irrationality into account and reacts rationally to it. The difference to the first specification is that in the first, the subject not only assumes irrationality on the part of informed traders but also considers himself to be the only informed trader to take this into consideration. Now we instead allow a later subject to believe that his predecessors are also aware of the possible irrationality on the part of informed traders and employ this knowledge in their decision-making. Consequently, in the first specification, S_3 traders would never have been presumed to rationally sell, whereas in the second specification such behavior is admitted as rational.²⁴ Alas, as with the simple error correction, we do not obtain a substantially better fit with the data, as can be gleaned from Table 17: we obtained the best fit for $\delta = 0$ in which case people act only on the basis of their prior expectation and do not update. For $\delta = .22$ (presented in the table; the figures for $\delta = 0$ coincide with those of the no-updating case), the fit is best for treatments T1-T3 (T4-T6 have the best fit for $\delta = 0$). In the latter case, the improvement for treatments T1-T3

²⁴Rather than directly implementing level K reasoning or Quantal Response Equilibria, we choose our alternative specification because it is an unusually complex task for the subjects to calculate these more general measures of naive reasoning with 4 different known types of traders (noise traders and three types of informed trader). Moreover, there is a subtle difference of our approach to the way that Quantal Response Models can be implemented in models with and without prices. In an informational cascade without prices a deviation from the cascading action is, in principle, a deviation from rationality. With moving prices, such a simple observation can no longer be made, neither is it possible for subjects to determine if there is a genuine error. Our notion of overweighing noise is therefore a simple means for subjects to model the lack of trust in predecessors' actions, without implying a definitive or systematic direction of the error. Traders thus act as if the proportion of noise traders were higher than 25% by downgrading the quality of information extracted from the history of actions embodied in H_{t-1} or q_t . Finally, since we already have noise traders built into the experiment, by opting to allow traders to increase their estimates of the percentage of expected noise trades above 25% our method is arguably an especially simple and intuitive rule of thumb which enables subjects to incorporate naive reasoning on the part of their peers. For more on rules of thumb by laboratory subjects in a herding context see Ivanov et al. (2008).

only is from 69.8% to 76.1%.

In summary, a model specification in which agents recursively take their predecessor's decisions as prone to error provides a worse fit with a data than the overweighing of one's own signal. Compared to the rational model there is an improvement of fit, though it is small.

	simple noise shift $\delta = 2/15$			simple noise shift $\delta = 1/3$			level 2 noise shift $\delta = .22$		
	S1	S2	S3	S1	S2	S3	S1	S2	S3
T1	58	61	60	58	60	65	58	61	66
U-negative	3	24	20	3	25	15	3	24	14
	95%	72%	75%	95%	71%	81%	95%	72%	83%
T2	47	64	63	47	58	63	47	66	63
hill	13	35	7	13	41	7	13	33	7
	78%	65%	90%	78%	59%	90%	78%	67%	90%
T3	66	44	47	66	42	48	66	44	49
U-positive	10	48	13	10	50	12	10	48	11
	87%	48%	78%	87%	46%	80%	87%	48%	82%
T4	127	94	98	127	92	98	127	94	98
hill	25	53	23	25	55	23	25	53	24
	84%	64%	81%	84%	63%	81%	84%	64%	80%
T5	123	69	97	123	64	97	123	67	98
U-positive	30	114	9	30	119	9	30	116	8
	80%	38%	92%	80%	35%	92%	80%	37%	92%
T6	103	116	117	103	97	114	103	108	114
U-negative	18	64	30	18	83	33	18	72	33
	85%	64%	80%	85%	54%	78%	85%	60%	78%
Total%	84%	57%	83%	84%	53%	83%	84%	56%	83%
T1-T3%	87%	61%	81%	87%	58%	84%	87%	62%	85%
T4-T6%	83%	55%	83%	83%	50%	83%	83%	53%	83%
Fit total	73%			71%			72.9%		
fit T1-T3	75%			74%			76%		
fit T4-T6	72%			70%			71%		

Table 17: Variations in the Perception of Noise Trading.

A.5 Summary of Alternative Behavioral Explanations

While forms of slow updating improve the fit of the data slightly, no alternative model is capable of providing a convincing explanation for the results. Slow updating, overweighing of one's own signal, and overestimating noise trading are essentially very similar, and also have strong similarities to a strategy of following the prior (which is a policy of zero updating).

Several studies (Drehmann, Oechssler and Roeder (2005) and Cipriani and Guarino (2005)) have already identified that when prices rise, people with high signals tend to act as contrarians, i.e. they sell. There are multiple possible explanations, ranging from risk aversion (which we refute) to slow or no updating. We observe the same kind of end-point behavior by the S_3 types. Symmetrically, the S_1 types should exhibit similar behavior when prices approach the lower bound. However our data rarely involves prices

that fall to a sufficient extent to examine the symmetric claim, since in general across all treatments, prices tend to tentatively rise. Note that the end-point effect should also influence the S_2 types, because whatever mechanism or cognitive bias leads S_3 types to sell for high prices should apply in the same manner to S_2 types.

Irrespective of which hypothesis is correct, if the end result is observationally equivalent to slow updating then this has a profound effect on how much herding or contrarian behavior one might expect to see: when people update slowly, it takes longer for them to reach a (subjective) expectation for which they would herd. However, with slow updating, they will also be slower to reduce prices and thus it is conceivable that they herd when prices move “against” the herd.

Appendices

The time-line in appendix 1 gives a complete run down of the structure of a typical experimental setting and the text (not the section headings) of the instructions detailed in appendix 2 are a reproduction of what is read to the subjects during the experiment. Answers to questions are of course unscripted, but all other communication with subjects was kept to a minimum. An example information sheet given to the subjects is reproduced in appendix 3 and the questionnaire is reproduced in appendix 4. Practical detail about the software used is provided in appendix 5.

B Time-line

What follows is a precise chronological ordering of events during the experiment.

1. The room is prepared and software pre-loaded into the machines to be used, which are allocated each to one ID number.
2. Read instructions 1 including random distribution of ID cards and seat subjects on the basis of the allocated ID cards.
3. Read instructions 2 including the completion and collection of permission forms.
4. Read instructions 3 which explains the experimental setting.
5. Read instructions 4 which explains the software.
6. Read instructions 5 which explains the compensation.
7. Read instructions 6 which explains the information setting.
8. Read instructions 7 which summarizes the instructions and pause to answer any questions.
9. Run treatment 1 (the example round).
10. Pause to answer final questions.

11. Run treatments 2-7.
12. Read instructions 8, which ends the experiment.
13. Calculate and distribute payments while participants complete receipts and questionnaires.

C Instructions

Note that the parts of the instructions in bold indicate that a name, number or currency be included in the instructions which vary by session. Words in italics are emphasized. and pause to answer any questions. The instructions are long, and the pre-experimental instructions (1-7) took an average of around 25 minutes to deliver including typical questions. Payment calculations typically took around 5 minutes during which subjects were asked to shut down open software and complete a questionnaire.

Instructions 1 (Welcome)

Welcome to everyone participating in today's experiment. My name is [**name**] and my assistants for today will be [**names**]. The experiment should take around one and half to two hours and will mainly involve using a computer. I ask that for the entirety of the experiment you refrain from talking unless you wish to ask a clarifying question or point out a computer error to me or one of my assistants, and you will be told when you can and cannot ask questions. You will be paid a turn up fee of £5 [**equivalent in Canadian dollars**] and can earn anything up to a further £25 [**equivalent in Canadian dollars**] based on your performance, so try to do your best! I will now distribute your ID cards. Please keep these safe as they not only determine where you will sit, but also what your payments will be. Actions during this experiment are anonymous in the sense that we are aware only of your ID number as indicated on your ID card when calculating payments and not your names. Please could you now take a seat in front of the computer indicated by your ID number. The computers are all divided by large screens for a reason, so please do not attempt to examine other people's computers.

Instructions 2 (After Seated)

After taking a seat make sure you are using the computer that is appropriate for your ID number. You will notice that there is a graph displayed on the screen with several on-screen buttons which are currently not highlighted. Next please read and sign the permission form using the pen provided. The permission form confirms that you have given permission for us to use you as willing participants in this experiment. You will also need to complete a receipt which you will be given at the end of the experiment before you receive your payment. My assistant(s) and I will now collect your permission forms.

Instructions 3 (The Experimental Setting)

Next I will describe the experiment itself. You will be participating in a series of financial market trading exercises. There will be 7 trading rounds, and each round will last 3-4 minutes. There are [number of participants] participants in the room and everyone is involved in the same trading exercise. Your objective should be to take the most thorough decision possible in order to maximize the money you will make today. The general situation is the following: you are the stockholder of a company and have some cash in hand. Some event may happen to your company that affects the value of the company (for better or worse). You have a broker who provides you with his best guess. You then have to decide whether you want to buy an additional share of the company, whether you want to sell your share, or whether you want to do nothing. We will look at a variety of similar situations: each situation concerns a *different* company, and we will vary the information and the trading rules in each situation. Please note that the situation described to you in each round is independent of that in any other round. *In other words, what you learned in round 1 tells you nothing about round 2, etc.* In the process of this session you may or may not generate virtual profits. Your trading activities will be recorded automatically; these activities determine your trading profits.

Before each round starts, you are given one share of the company and you have sufficient cash to buy a share. Round 1 will be an example round and your final payment will not reflect how you perform during this round.

During the rounds you may sell your share, you may buy one additional share or you may do nothing. You can only trade within a specific time window indicated by the software a red blinking bar appearing around the trading buttons below the graph. You will receive a notification by the system on your screen and then you have *5 seconds* to make your trade. The frantic blinking will continue for 5 seconds irrespective of whether you trade or not. *Note that you can trade only once*, in other words, you can only buy or sell, you cannot do both. Once you have hit the button it may take the system a second or two to register your trade. You should not double-click or attempt to click more than once.

There will be a pause after round 1, the example round, when you can answer questions. During rounds 2-7 you will be required to remain silent.

Instructions 4 (The Software)

Now please examine your computer screen, without hitting any buttons. Before you is a screen that contains several pieces of information:

1. It tells you about all the trades that occur during the round; you also see when a trade occurs and whether or not someone bought or sold a share. For your convenience, there is a graph that plots the sequence of prices.
2. Your screen also lists the current market price; people can either buy a share at this price or they can sell their share at this price.
3. In the case where we restrict the time when you can make a trade, a red bar will appear on the bottom of the screen to highlight the fact that you can trade. During

this time the buy, sell and pass buttons will be available for your use, typically only once per round, though twice in the final 3 rounds.

4. There is also a box in which you receive some information from your "broker" which I will explain in a few moments.
5. The screen includes a timer which indicates how many seconds have gone past during the round.
6. Finally, the screen updates itself whenever a trade is made.

Note that you are not directly interacting with any of the other participants in the experiment, rather the actions of all of the traders including you and your fellow participants will effect the current price which is set by the central computer being operated at the front of the experimental laboratory such that a decision to purchase by a trader will raise price and to sell will lower it. This central computer will also be producing trades itself which will account for 25% of all the possible trades during each round and will be determined randomly so there is a 50% chance a computer trader will buy and a 50% chance he will sell.

Instructions 5 (Compensation)

Next I will describe the payment you will receive. You will receive £5 [**Canadian equivalent**] in cash for showing up today. You can add to that up to a further £25 [**Canadian equivalent**] as a bonus payment. In this trading experiment, you will be buying or selling a share (with virtual units of a virtual currency), and this trading may or may not lead to virtual profits. Your bonus payment depends on how much profit you generate in total across all of the rounds with the exception of the example round. In general, the more thorough your decisions are, the greater are your chances of making profits, and the higher will be your bonus.

I will next explain virtual profits. When you trade you will do so at the current price appearing on your computer screen. The initial price is 100 virtual currency units (vcu). This price changes based upon the trading that goes on during the round including those by your fellow participants and the random computer traders. While you will trade today during the experiment, we can imagine that after the end of each round of trading there is a second day during which the event (good, bad or neutral) is realized and the price of the share is updated to reflect this: this will be either 75, 100 or 125 vcu. To stress, which price is realized depends upon which event takes place:

- if something good happens to the company, the price will be 125 after the realization of the event;
- if something bad happens, so the price will be 75;
- if neither of these, so the price reverts to the initial value of 100.

Your profit relates to the difference between the current price that you buy or sell a share at today, and the price revealed after the event takes place. An example of a good

event happening to the company might be that it wins a court case or gains a patent. A bad thing might be the opposite, so the firm loses a court case or fails to gain a patent. Note that as already stressed, each round is an independent experiment, so in round 1 it may be that the bad event takes place so the share price becomes 75 after trading finishes, while in round 2 it may be worth 125, etc.

Next I will go through some simple numerical examples of what might happen.

Example 1 *If you buy a share at a price of 90 vcu, and after the event takes place the price of the share is updated to 125 vcu. You have therefore made 35 vcu of virtual profits on your trade. If you instead sold at 90 vcu you would have lost 35 vcu. If you did nothing you would make a profit of 25 vcu since your share was originally worth 100 vcu and is worth 125 vcu after the event is realized.*

Example 2 *If you buy a share at a price of 110 vcu, and after the event takes place the price of the share is updated to 100 vcu you have lost 10 vcu of virtual profits on your trade. If you instead sold at 110 vcu you would have made 10 vcu. If you did nothing you would have neither made a profit or a loss on your trade.*

So note that what matters is the price when you take an action and the true value after the good, bad or neutral event. Which event occurs will not be revealed to you during the experiment though you will receive information about which is more likely before the start of trading. I will explain the nature of this information in a moment.

Please remember that each round represents a completely different situation with a different share and a different firm. In every round you may make or lose virtual profits and by the end the central computer will have a complete record of your performance. On the basis of your overall performance the central computer will calculate your bonus payment.

Instructions 6 (The Information Setting)

I will now explain the *broker's tip* and the information you have before each round begins. Next to your computer is a set of sheets which correspond to each round. For example, the top sheet is called "Example Round 1", and has several pieces of information about the share. For instance the sheet indicates to you the chance that the share price will be 75, 100 or 125 vcu after the event. Next it indicates what sort of broker's tips you might receive. Each participant has identical sheets, the text, numbers and diagrams are literally the same for every participant.

Your broker will give you a tip via your computer screen that indicates his view about what sort of event will occur. He might give you a "good tip" (which we call S_3), "bad tip" (S_1) or "middle tip" (S_2). A good S_3 tip indicates that he believes the event will be good and the share price will be 125 vcu after it is realized, a bad S_1 tip that something bad will happen indicates 75 after the event is realized. A middle S_2 tip is a bit more complex but indicates he feels 100 vcu is his best guess:

- It could mean that he believes nothing at all will happen hence he believes the price will revert to the original 100 vcu and we call this *case 1*.

- Or it could mean that he believes an event will happen but he is not sure whether it is either good or bad, and we call this *case 2*.
- Or it could mean that he believes something good or bad will happen and he has a feel for which, but he is not sufficiently sure to indicate the good or bad tip and would prefer to indicate middle and we call this *case 3*.

Before each round you are told which case would apply if you receive a middle signal together with a background probability that there will be a good, neutral or bad event which will make tomorrow's price 75, 100 or 125 respectively.

Unlike the contents of the information sheet the tip you receive is private to you, and other participants may receive the same or a different tip. In other words it is possible that your broker might believe a good event is going to happen so the price will be 125 after this realization, while other participants might have brokers who agree or disagree with your broker's tip. There are also other pieces of information on the sheet including the probability that the broker is correct when he gives you a tip, and this probability is the same for all participants.

You will be given 2 minutes to examine the relevant sheet before each round. You will then receive notification on your computer screen of the actual tip sent to you from the broker: S_1 , S_2 or S_3 , and will have another minute to consider this. The beginning of the round will then be announced and trading will begin. Remember you can only trade during the 5 second window indicated by a red bar on your screen. The buttons on the screen (buy, sell or pass) can only be pressed during this time and only once per round.

Instructions 7 (Summary)

To summarize, you are in a market experiment with a central computer that both records your actions and produces random trades (which account for 25% of all trades). All other participants will also have the opportunity to trade. You will receive a private signal from a broker and other information pertaining to the price of the share after a possible event occurs, including the likelihood of the broker being correct. The information on your information sheet is common to everyone (for example, everyone's broker is just as likely to be correct as yours), but the broker's signal is private to you while others will receive a signal which may be the same or different from yours. Each market participant, yourself included, has their own different broker in each round. The rounds are all different in the sense that the share is for a different company, the broker is different and earlier actions and prices are not relevant. You will make virtual profits based on the difference between your trading price in vcu and the price after the event which will be 75, 100 or 125 vcu. The total of your virtual profits across all rounds, excluding the example round, will be used to calculate your bonus payment. To maximize your bonus payment you will then have to make high virtual profits and therefore make as thorough a decision as you can.

Please do not talk, signal or make noises to other participants, please do not show anyone your screen or discuss your information, please do not try to look at other people's screens and we would appreciate it if would not leave the room until the experiment is over.

You may ask questions now or just after the example round. Once we begin rounds 2-7 you will not be allowed to ask clarifying questions, though you should inform us if there is a software problem.

Instructions 8 (Experiment End)

Many thanks for participating in today's experiment. Please remain in your seats for a minute or two while we use the central computer to calculate your final payments. We ask that you close the trading software and any other open software and shut down your computer. We also ask that you leave the pen and all sheets on your desks, and keep only the ID card which you will need to bring with you to the front desk in order to receive your payment. When you receive your payment you will also be asked to complete and sign a receipt. It would be useful if you could complete the questionnaire that is on your desk, and hand it in as you leave, though this is not compulsory. After you leave, we ask that you try to avoid any discussion of this experiment with any other potential participants, and once again many thanks for your participation.

D Information Sheets

Here we present an example "information sheet" comprised of some text and two diagrams. The one presented here is taken from the example round, but one of these was provided for each treatment.

E Questionnaire

Many thanks for taking part in today's experiment. The official part of the experiment is now over. Your payments are now being worked out and you will be paid based on your ID number (the computer you are using). Please answer the following questions. In particular this will help us to make future experiments better and may help us understand the results.

About you

1. Your age:
2. Your gender:
3. Your degree subject:
4. Have you ever owned shares?
5. Do you have any experience of financial markets? (if so, what are your experiences)

About your decisions today

6. What made you decide to buy, sell or pass?
7. How important was the current price?
8. How important was the past price data (the graph)?
9. How important was your “broker’s tip”?
10. What else mattered?
11. Did you make any calculations? If so, which ones?

About the experiment

12. Anything else you would like to report, including how to make the experiment better, can be done so here:

F The Software

The trading market was simulated through a software engine, run on a central computer, networked to a number of client machines each running the one version of the client for each subject. The central computer acted to record and analyze results, as well as to distribute signals (through an administrator application) and provide a continuously updated price chart for subjects. The sequence of signals and noise trades was pre-specified and the computer also organized the allocations of time-slots for each trader and noise trades and it provided an indication to traders of when they could trade.

Figure 6 shows the administrator software. The screen shot is not taken from an actual session, but simply shows the layout on screen for a fictional session. It is currently listed as recording the activity of traders in “Treatment 1”. As can be seen in the figure there are more noise traders than would be normal in an actual session (indicated by the final letter N, whereas subjects are indicated by a final ID number). As can be seen here trader HEG5P3 has “timed out” (failed to act in their 5 second window, which is a feature not present in the current experiment but is part of the setup in our prequel paper).

The client software provided a simple to use graphical interface which enabled subjects to observe private information (their signal), and public information (the movement of prices and the current price), as well as indicating to them when they could trade (flashing red and enabling trading buttons) and providing the means of trade (buy, sell and pass buttons). Figure 7 below shows a screen shot of the software in action.

Here you can see that the price initially rose from a level of 100, indicating buying at the early stages, but then price started to fall back, it rallied and then fell back further to a value of around 116. This subject’s private signal was S_1 (“bad”) and the subject had a single share to sell and a large cash balance to enable the purchase of a further

share. The subject could also pass (declining to buy or sell) when given the opportunity to trade.

The software was purposefully built for the experiment, since existing software was unable to provide the sort of information structure needed in a price-driven (as opposed to order-driven) market.²⁵

²⁵Further details about the software are available on request from the authors.

Example Round 1

Signals: **Case 2**

- If you receive signal S1 (the “bad” signal), then the broker indicates a negative impact.
- If you receive signal S3 (the “good” signal), then the broker indicates a positive impact.
- If you receive signal S2 (the “middle”), then the broker indicates that there is an effect but he is not sure which one; he is leaning towards negative.

The broker provides you with his best guess; he may be wrong. He makes his best guess on the basis of data that he analyses. For instance, if the true announcement effect is going to be positive, the broker is likely to find data that makes him conclude that something positive is about to happen. But – he may also read the data the wrong way. In other words, the broker makes his statement with certain probabilities, and these depend on the true effect of the announcement.

If the true effect is **POSITIVE** then you receive


- Signal S1 (bad) with chance 5% (he gets it badly wrong)
- Signal S2 (no effect) with chance 25% (he thinks nothing happens although it does)
- Signal S3 (good) with chance 70% (he gets it right)

If the true effect is **NEGATIVE** then you receive


- Signal S1 (bad) with chance 65%
- Signal S2 (no effect) with chance 30%
- Signal S3 (good) with chance 5%

If indeed the effect is **NO EFFECT** then you receive

- Signal S1 (bad) with chance 45%
- Signal S2 (no effect) with chance 10%
- Signal S3 (good) with chance 45%



Probability of receiving a signal, given the true value



Probability of an announcement, given the signal




Figure 6: The Administrative Interface




Figure 7: The Trading Client

References

- Alevy, Jonathan E., Michael S. Haigh, and John A. List**, “Information Cascades: Evidence from a Field Experiment with Financial Market Professionals,” *Journal of Finance*, February 2007, *LXII* (1), 151–180.
- Anderson, L. and C. Holt**, “Informational Cascades in the Laboratory,” *American Economic Review*, 1997, *87*, 847–862.
- Avery, C. and P. Zemsky**, “Multi-Dimensional Uncertainty and Herd Behavior in Financial Markets,” *American Economic Review*, 1998, *88*, 724–748.
- Back, Kerry and Shmuel Baruch**, “Information in Securities Markets: Kyle meets Glosten and Milgrom,” *Econometrica*, March 2004, *72* (2), 433–466.
- Banerjee, A.V.**, “A Simple Model of Herd Behavior,” *Quarterly Journal of Economics*, 1992, *107*, 797–817.
- Bikhchandani, S. and S. Sunil**, “Herd Behavior in Financial Markets: A Review,” Staff Paper WP/00/48, IMF 2000.
- , **D. Hirshleifer, and I. Welch**, “Theory of Fads, Fashion, Custom, and Structural Change as Informational Cascades,” *Journal of Political Economy*, 1992, *100*, 992–1026.
- Bloomfield, Robert, Maureen O’Hara, and Gideon Saar**, “The ‘make or take’ decision in an electronic market: Evidence on the evolution of liquidity,” *Journal of Financial Economics*, January 2005, *75* (1), 165–199.
- Brunnermeier, Markus K.**, *Asset Pricing under Asymmetric Information – Bubbles, Crashes, Technical Analysis, and Herding*, Oxford, England: Oxford University Press, 2001.
- Chakraborty, Archishman and Bilge Yilmaz**, “Informed Manipulation,” *Journal of Economic Theory*, 2004, *114* (1), 132–152.
- Chamley, C. and D. Gale**, “Information Revelation and Strategic Delay in a Model of Investment,” *Econometrica*, 1994, *62*, 1065–1085.
- Chamley, Christophe**, *Rational Herds*, Cambridge, United Kingdom: Cambridge University Press, 2004.
- Cipriani, Marco and Antonio Guarino**, “Herd Behavior in a Laboratory Financial Market,” *American Economic Review*, 2005, *95* (5), 1427–1443.
- and —, “Herd Behavior in Financial Markets: A Field Experiment with Financial Market Professionals,” mimeo, UCL and GWU June 2007.
- Costa-Gomes, M., V. Crawford, and B. Broseta**, “Cognition and Behavior in Normal-Form Games: An Experimental Study,” *Econometrica*, 2001, *69*, 1193–1235.

- Drehmann, Mathias, Jörg Oechssler, and Andreas Roeder**, “Herding and Contrarian Behavior in Financial Markets: An Internet Experiment,” *American Economic Review*, 2005, *95* (5), 1403–1426.
- Glosten, L.R. and P.R. Milgrom**, “Bid, Ask and Transaction Prices in a Specialist Market with Heterogenously Informed Traders,” *Journal of Financial Economics*, 1985, *14*, 71–100.
- Hasbrouck, Joel**, *Empirical Market Microstructure*, Oxford University Press, 2007.
- Ivanov, Asen, Dan Levin, and James Peck**, “Hindsight, Foresight, and Insight: An Experimental Study of a Small-Market Investment Game with Common and Private Values,” *American Economic Review*, 2008, *forthcoming*.
- Kahneman, Daniel and Amos Tversky**, “Prospect Theory: An Analysis of Decision under Risk,” *Econometrica*, March 1979, *47* (2), 263–292.
- Kyle, A.S.**, “Continuous Auctions and Insider Trading,” *Econometrica*, 1985, *53*, 1315–1336.
- Levin, Dan and James Peck**, “Investment Dynamics with Common and Private Values,” *Journal of Economic Theory*, 2008, *forthcoming*.
- McKelvey, R.D. and T.R. Palfrey**, “Quantal Response Equilibria for Normal Form Games,” *Games and Economic Behavior*, 1995, *10*, 6–38.
- and — , “Quantal Response Equilibria for Extensive Form Games,” *Experimental Economics*, 1998, *1*, 9–41.
- Milgrom, Paul**, “Good news and bad news: representation theorems and applications,” *The Bell Journal of Economics*, 1981, *12*, 380–91.
- Park, Andreas and Daniel SgROI**, “When Herding and Contrarianism Foster Market Efficiency: A Financial Trading Experiment,” Working Paper, Universities of Toronto and Warwick 2008.
- and **Hamid Sabourian**, “Herding and Contrarian Behavior in Financial Markets,” Working Paper, Universities of Toronto and Cambridge 2008.
- SgROI, Daniel**, “The Right Choice at the Right Time: A Herding Experiment in Endogenous Time,” *Experimental Economics*, October 2003, *6* (2), 159–180.
- Shiller, R.**, “How a Bubble Stayed Under the Radar,” *New York Times*, 2008, *March 2*.
- Smith, Lones**, “Private Information and Trade Timing,” *American Economic Review*, 2000, *90*, 1012–1018.
- Tversky, Amos and Daniel Kahneman**, “Advances in Prospect Theory: Cumulative Representation of Uncertainty,” *Journal of Risk and Uncertainty*, 1992, *5*, 297–323.
- Vives, Xavier**, *Information and Learning in Markets*, Princeton University Press, 2008.

Ziegelmeyer, Anthony, Kene Boun My, Jean-Christophe Vergnaud, and Marc Willinger, “Strategic Delay and Rational Imitation in the Laboratory,” Working Paper, Max Planck Institute of Economics 2005.