When Herding and Contrarianism Foster Market Efficiency: A Financial Trading Experiment

Andreas Park and Daniel Sgroi

No 854

WARWICK ECONOMIC RESEARCH PAPERS

DEPARTMENT OF ECONOMICS

When Herding and Contrarianism Foster Market Efficiency: A Financial Trading Experiment

Andreas Park*
University of Toronto

Daniel Sgroi†
University of Warwick

April 29, 2008

Abstract

While herding has long been suspected to play a role in financial market booms and busts, theoretical analyses have struggled to identify conclusive causes for the effect. Recent theoretical work shows that informational herding is possible in a market with efficient asset prices if information is bi-polar, and contrarianism is possible with single-polar information. We present an experimental test for the validity of this theory, contrasting with all existing experiments where rational herding was theoretically impossible and subsequently not observed. Overall we observe that subjects generally behave according to theoretical predictions, yet the fit is lower for types who have the theoretical potential to herd. While herding is often not observed when predicted by theory, herding (sometimes irrational) does occur. Irrational contrarianism in particular leads observed prices to substantially differ from the efficient benchmark. Alternative models of behavior, such as risk aversion, loss aversion or error correction, either perform quite poorly or add little to our understanding.

JEL Classification: G14, G24, G28.
Keywords: Herding, Informational Efficiency, Experiments.

*E-mail: andreas.park@utoronto.ca
†E-mail: daniel.sgroi@warwick.ac.uk

We thank seminar participants at the Universities of Cambridge, Toronto, and the ESRC Conference on World Economy and Finance at Warwick. Financial support from the ESRC, SSHRC and CERF is gratefully acknowledged. Special thanks go to Annie Jekova an and Malena Digiuni for expert research assistance and Kenji Yee and Kelly Hayward for allowing us to use their computer labs.

1 Introduction

Rational herding has become an important tool in analyzing why and how economic agents learn through observation in groups (Banerjee (1992) and Bikhchandani, Hirshleifer, and Welch (1992)). Economic agents constantly learn from others, through chat, 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 it and then, seemingly disregarding their own information, follow the majority action. A central lesson of 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 nowhere near 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 herding analysis. One could argue that a few early, perhaps incorrect, movements by visible traders induce discontinuous price jumps in one direction or the other. 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.

Yet a theoretical proof for the possibility of meaningful financial market herding proved difficult. An early result by Avery and Zemsky (1998) showed that in a standard financial market-trading setting herding is not possible as the market price always separates people with good and bad information.[1] Subsequent experimental studies by Drehmann, Oechssler, and Roider (2005) and Cipriani and Guarino (2005), which we will comment on in detail below, confirmed this finding.[2]

However, financial commentators and highly respected economists continue to con-

[1]To put this in context, Avery and Zemsky (1998) *did* present 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 herding example, however, 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 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). Cipriani and Guarino (2007) is a recent experimental study of the kind of herd behavior that Avery and Zemsky identify.

[2]Subsequent theoretical work that identified some form of herding relied on market frictions that mitigated the separating power of the market price; see e.g. Lee (1998), Cipriani and Guarino (2003), Chari and Kehoe (2004), or Dasgupta and Prat (2005).

sider rational herding to be a possible explanation for financial crises and shocks,[3] and contrarianism in financial markets is not only prevalent in experimental settings but also has a well-documented presence in market transaction data.[4] Our experiment is the first to be based on a theoretical framework in which herding and contrarianism are admissible rational outcomes, and as such, we are the first to address the question of why this contrarianism occurs, and to attempt to disentangle rational and irrational forms of herding and contrarianism. We also assess alternatives, such as risk- or loss-aversion, slow updating, and numerous behavioral theories, but find them to offer little explanatory power over and above the benchmark rational model. We find that: (a) rational herding does arise in accordance with the theory; (b) rational contrarianism also arises though less often than predicted by theory; and (c) irrational contrarian behavior in particular leads prices to substantially differ from the rational, efficient benchmark. While our results are generally supportive of the theory, to fully reconcile the behavior of our subjects with a rational model, even one that admits herding and contrarianism, one would have to assume dramatically more noise than was actually present.

The experimental design we employ is based on a recent theoretical paper by Park and Sabourian (2008), which describes the conditions on the underlying information structure that are necessary and sufficient for informational herding, contrarianism, and for no-herding and no-contrarianism. To summarize their findings, they show that with more than two states and signals informational herding is possible if the herding candidate receives a 'U-shaped' signal that causes him to redistribute probability weight towards extreme values. Alternatively, for contrarian behavior the trader must receive a 'hill-shaped' signal that causes him to distribute probability weight towards the center. A full summary of their model and results is provided in section 3. Contrasting with the conclusions drawn from the early results on herding in financial markets, their work suggests that rational herding and contrarian behavior may be pervasive and important. They show that prices become more volatile than if there were no herding; liquidity, measured by the bid-ask-spread, is reduced; herding is persistent; and herding affects the process of learning.[5] Using this framework we can analyze situations where rational herding or rational contrarianism are possible as well as situations where neither is mandated by theory.

[3]See for example Shiller (2008) in which Robert Shiller eloquently highlights the scope for herding theory as a potential explanation for the current housing market difficulties faced by the USA and elsewhere.

[4]See for example Chordia, Roll, and Subrahmanyam (2002). Interestingly, Chordia, Roll, and Subrahmanyam (2002) attribute contrarianism to rational traders who counter irrational, herding traders. However, experimental observations so far give the opposite implication: for instance, if prices should rationally rise, they are forced to mean revert by irrational contrarians.

[5]Despite this, Park and Sabourian employ the same model as Avery and Zemsky (1998) and thus do not impose assumptions to mitigate the separating power of the market price.

The theoretical model underlying both our experiment and Park and Sabourian is a sequential trading setup *a la* Glosten and Milgrom (1985) in which risk-neutral subjects trade single units of a financial security with a competitive market maker, arriving at the market in a predetermined, exogenous sequence. 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. In our specification, there are three possible values and three possible signals (high, middle, or low); each subject privately receives a realization of one of these signals. There are three possible actions (buy, pass, or sell). Rational subjects should then buy if their expectation, conditional on all private and public expectation, is above the price and sell if it is below.

This fully rational behavior is the first yardstick for our experimental results.[6] The overall fit of the model is roughly 75%, which corresponds well with other experimental work such as Drehmann et al (2005) and Cipriani and Guarino (2005, 2007). Broken down by player type, these numbers are 82%, 61% and 86% for the lowest, the middle and the highest signal type respectively. Thus while the fit is very high for the highest and lowest signal types, it is low for middle signal types. One obvious reason is that the decision with the middle signal is often much more difficult: high and low signal types would make no mistakes if they simply stick to their initial decision, whereas middle signal types have to follow the development of prices very carefully.

The implicit hypothesis in the spirit of Park and Sabourian is that middle signal types are more prone to herd (with U-shaped signals) or contrarianism (with hill-shaped signals). Indeed, this is what we observe: while the fit to the rationally expected herds is poor (only 31% of the herds that rationality prescribes actually take place), the middle signal types do exhibit a strong tendency to herd when they can. Herding by middle types occurs in 26% of the cases where it is possible whereas it only occurs in 6% and 16% for the highest and lowest types, respectively. Logit, OLS and fixed effects regressions all formally confirm the significance of the shape (U or hill) of the signal.

This is a rather powerful message: people who receive information that may be conflicting, weaker, or 'somewhere in the middle' will shift around with market forces and may exhibit herding or contrarian characteristics.

Individual behavior is of course interesting, but also of interest, especially from a policy perspective, is the behavior of the asset price and overall market efficiency. Using the price that would result if all traders in the model act as theory predicts, we find that about 16% of observed prices coincide with this benchmark. In most treatments

[6] This rational benchmark is arguably quite restrictive since prices are computed on the basis of the rational model which assumes no mistakes. However, if there are mistakes then prices are misleading and as such subjects, knowing that human nature is prone to errors, may actually take the correct decision.

the trading histories are such that the theoretical model would predict rising prices, and indeed this is what happens. Yet on average observed prices are 7% below the benchmark, with a standard deviation of 9.9%. For instance, prices would rise above the benchmark if there is irrational buy-herding, or fall below the benchmark if there is irrational sell-contrarianism. Our numbers indicate that contrarianism (both rational and irrational) is more prevalent than herding.

The consequence of irrational behavior of any sort is, of course, that prices deviate from the efficient benchmark. We provide an indication of the scale of the inefficiency: supposing that these deviations are non-systematic, one could argue that observed prices could be reconciled with the rational benchmark if the latter would account for a higher level of noise. We then find that the required noise level would have to be roughly 3.6 times larger than the actual level of noise trading.

While we believe that our data indicates that people act in the spirit of the rational model, we also examine several alternative hypotheses to see if they capture behavior better. First on the list are risk and loss aversion. Loss aversion can be motivated by observing that as prices approach, say, the highest value, purchasing means exposing oneself to potentially large losses. Both of these models, however, perform quite poorly. Loss aversion matches at most 50% of the data in total. Risk-aversion, as measured by the CARA coefficient, does equally poorly when we assume standard values for the risk-aversion coefficient (2 or higher) and the fit improves as we approach risk-neutrality.

Noting that people do indeed make mistakes, we also examine several specifications in which people dampen the effect of observed trades in computing their expectations, for instance, by re-scaling probabilities or over-estimating noise trading. While these models can generate a higher fit with the data, one needs to assume very high levels of dampening to achieve a measurable effect. Moreover, when assuming these dampening effects to be present, herding behavior cannot be explained. Since formal regressions confirmed the significance of U-shaped signals, these dampening approaches that eliminate the possibility of herding seem inapplicable.

Our experimental method differs slightly from the theoretical setup in that people are not only allowed to pass but that this decision also has a potential interpretation: if people were to act as described by any rational model with risk-neutral agents, then people buy if their private expectation is above the price and sell if it is below — passes would be irrational. However, in our experiment people are initially endowed with a share and their decision is to sell this share, to hang on to it, or to buy an additional share. Holding is therefore a sign of an agent believing that the share has some value but not enough to warrant doubling up the position. [7]

[7]Put differently, the decision to pass in our model can be interpreted as a 'weak buy'. In the main

Overview. The rest of the paper is structured as follows. The next section reviews extant work on experiments that test financial market herding and the relation to our work. 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 5 examines the design of the experiment, including a discussion of the nature of the software, the different information structures embodied in the alternative treatments, and the information provided to subjects. Section 6 examines which hypotheses are implied by the theory and considers a variety of alternative models of behavior which might explain the observed data. Section 7 carries out an econometric analysis. Section 8 outlines results from an examination of alternative behavioral explanations. Section 9 examines the efficiency of prices observed in the experiment. Section 10 summarizes the key findings and concludes with some suggestions for future work. The subject instructions and other supporting materials are provided in the appendices at the end of the paper.

2 Experimental Work on Financial Market Herding

Given the differing views concerning the scope of informational herding to play a strong role in explaining clustered behavior in financial markets, it is not surprising that there is a market for experimental work in the area. Experimental work can act not just as a direct test of existing theories, but as a means to assess which assumptions are most plausible, and how small deviations from the theory, perhaps in the direction of reality, might complicate matters. The need to know exactly what information each trader has in each time period makes traditional empirical work difficult, and makes experimental testing, where information is closely controlled, an ideal setting. Given this, there have been numerous papers examining herding behavior in experimental settings, albeit not necessarily with efficient prices.

The first published experiment to test herding was Anderson and Holt (1997), but without moving prices. They found that herding did indeed occur, but at lower levels than predicted by the theory (73%), which they justified in terms of assumed errors made by predecessor decisions. Like Anderson and Holt, most experiments since have not specifically considered a financial market setting and have therefore implicitly held prices fixed. Notable exceptions are Drehmann, Oechssler, and Roider (2005), Cipriani and Guarino (2005) and Cipriani and Guarino (2007).

Drehmann, Oechssler, and Roider (2005) and also Cipriani and Guarino (2005) aimed

text we will also present results which have this interpretation for passes, thus ruling some passes (not all) as a rational decision akin to a buy; in this introduction, for simplicity of the exposition holds have been counted as 'weak buys'.

to test the no-herding result from Avery and Zemsky (1998) that we alluded to above. We focus here on Drehman et al; Cipriani and Guarino obtain very similar results. In their specification, people had the choice of two actions, A and B,[8] with a moving price that should exhibit the usual separating property of efficient prices. In their study, they found that the distribution of final prices is much more concentrated around the prior, so there was lower volatility than predicted, and prices showed a tendency to revert to the mean. Supported by a simple error model or a Quantal Response Equilibrium (QRE) model[9] they thus concluded that regression to the mean and contrarianism better typified their findings than rational herding or mimicking. They found little evidence of rational herding as expected, and frequent contrarianism.

The paper most closely related to ours is independent work by Cipriani and Guarino (2007), who tested Avery and Zemsky's model with so-called 'event uncertainty' in the lab. As in ours, their experiment allows for rational herding behavior. The first major differences in the design between our two papers are that our subjects were students, theirs were financial professionals. This lends more real-world relevance to their findings, although we had the larger subject pool.[10] The second difference is that they tested the 'event uncertainty' setting from Avery and Zemsky whereas we tested the arguably more general Park and Sabourian. This enables us to study *rational* contrarian behavior. There are other differences, for instance no-trades have different payoff implications in our two frameworks. The results on herding from our two approaches are similar: they also found that there was more herding than in no-herding setups but less than theoretically predicated. The numbers are also similar: they found that 23% of the theoretically predicted herds occur; in our setup this number is 31%. They have no notion of rational contrarianism, and so no comparison is possible on that front.

However, both Drehmann et al (2005) and Cipriani and Guarino (2005, 2007) have only two kinds of signals, and so the initial comparison with Drehmann et al and Cipriani and Guarino to our framework should be done with respect to the two types in our framework that receive the analogous signals. For these types, we observe a very close fit to rationality (82% for the lowest signal types and 86% for the highest signal types). Moreover, the low signal types herd 16% of the time, and the high signal types herd 6% of the time, when herding is possible (when prices rise or fall respectively). Our subjects also show a slightly higher tendency to act as contrarians (23% and 14% respectively). All

[8] Two actions are thus akin to 'buy' and 'sell'; some of their treatments included a no-trade option, which, however, has a different interpretation than the passes in our experiment as argued in the introduction.

[9] A QRE model estimates errors recursively for each round taking into account that predecessors react to their predecessors' possible errors, as developed by McKelvey and Palfrey (1995, 1998).

[10] That being said, they employed an experimental design by which they were able to generate a lot of data.

in all we conclude though that our results on this level do not differ dramatically between our setting and Drehmann et al, which provides a useful robustness check that our high and low signal types behave in line with earlier work. Of course, our setup's main thrust is the insights provided by the behavior of the middle types.

From the literature we observe various lessons for our own design. An experiment tests many things, for example: whether a given herding theory makes reasonable demands on the rationality of subjects; the nature of deviations from rationality; and the response of subjects to the observed non-rationality of others. In the literature the seminal herding papers (Bikchandani, Hirshleifer and Welch (1992) and Benerjee (1992)) have some success in the lab since cascades do happen, but less often than theory predicts. Adding flexible prices in a two-state setting, Avery and Zemsky get some support, with a much reduced incidence of herding, and a measure of irrational contrarianism.

We now go in the opposite direction and seek to investigate whether subjects will indeed herd, or act in a contrarian manner, when the design of experiment is based on a model in which herding or contrarian behavior is a rational outcome. To do this, we use the multi-state multi-signal theoretical framework of Park and Sabourian.

3 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.

Definition (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.* It therefore does not signify that everyone acts alike — but in a financial market setting such behavior would actually not be very useful for explaining crises or booms. The reason is that if all traders were to act alike, then actions would be uninformative, so that prices would not react to actions and stay 'constant' for the duration of herding.[11] Moreover, to theoretically ensure that financial markets function, i.e. that for every seller there is a buyer, one usually assumes the presence of a dedicated market maker. Yet this is a modeling trick that is no longer innocuous if all traders act alike.[12] In summary, the most one should expect from herding is a substantial shift in market activity that impacts price volatility and liquidity in the short run. This is precisely what we study here.[13]

4 The Underlying Theory

Traders arrive in a random sequence and 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 inde-

[11] For instance, models with informational cascades such as Cipriani and Guarino (2003) and Dasgupta and Prat (2005) have the feature that prices no longer move once the informational cascade starts. We focus on informational effects and information externalities. Of course, the mass-uniform behavior and the resultant large order imbalance may lead to payoff externalities akin to a bank-run situation. That notwithstanding, such effects are only second-order consequences of informational herding and thus require a different modeling approach altogether.

[12] In fact, the importance of the role of the market maker comes out very cleanly in Lee (1998)'s insightful paper where booms and crashes occur because the market maker cannot react fast enough to mass behavior and instead has to absorb all trades at the same price (the price adjustment after the rush constitutes the crash outcome). Chari and Kehoe (2004) use a different trick: herding in their model refers to mass-uniform behavior *outside* of capital markets and refers to a fixed investment decision. In Cipriani and Guarino (2003) mass-behavior has two sides: one group of people herds, the other acts as contrarians; both groups are, however, of similar size.

[13] In the literature, there are other definitions of herding (and informational cascades). For instance, Smith and Sørensen (2000) and also Cipriani and Guarino (2003) define herding as 'action convergence' — agents of the same 'type' take the same action. They describe an informational cascade as a situation where an agent takes the same decision irrespective of his private signal. Herding in our set-up refers to the actions of a particular signal type, not to all informed agents collectively. In the Park and Sabourian model the market-maker's zero-profit condition together with the assumption of non-identical signals precludes action-convergence of all informed traders — it is not possible that all informed agents trade on the same side of the market. In Cipriani and Guarino (2003) action convergence refers to specific types with type-characteristics other than just signals and in their model different types take different actions. In any case, the definition of herding which we (and also Avery and Zemsky) employ is in spirit of herd-mentality: people switch actions to follow the crowd and act against their initial private information.

pendent 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. Before meeting a trader, the market maker sets a single price at which he is willing to buy or sell one unit of the security.[14]

Every trader is a noise trader with probability 25% and buys or sells with equal chance. Informed traders receive one of three signals, S_1, S_2, S_3. Signal S_1 is generated with higher probability in state V_1 than V_2, and likewise in state V_2 vs. 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. Traders buy if their expectation conditional on their private signal and on the information derived from past trades exceeds the price, and they sell if the price is below the price.

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

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

A signal is called *csd-monotonic* if its csd is either increasing or decreasing. If $\mathsf{Pr}(S|V_1) > \mathsf{Pr}(S|V_3)$ or $\mathsf{Pr}(S|V_1) < \mathsf{Pr}(S|V_3)$ 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. Applied to the experimental design in which bid-ask-spreads are ignored, the analysis in Park-Sabourian implies the following. (Note that 'S herds' is to be read as 'S herds with positive probability'.)

[14]This is a simplification of the Park and Sabourian model, which itself is an adaptation of a Glosten and Milgrom (1985) style sequential trading model. 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. While using bid- and ask-prices would be desirable, their use would generate a host of complications. Most obviously, participants need to understand the difference between the two, and the software needs to keep track of more variables. Moreover, it may lead to people focussing (subconsciously even) on only one side of the market, i.e. people with initially negative information may follow only the movement of the bid-price and thus make their decision between selling and holding, disregarding the possibility of buying completely. It is not clear what insight would be added by distinguishing between trades at bid- and ask-prices.

Theorem (Existence of Herding and Contrarian Behavior)

(*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 negative bias is necessary and sufficient for H1/C1 to be satisfied, i.e. for someone to have a negative/positive initial opinion.

So what would a history with herding look like? We know that herding requires that prices have increased; this occurs if and only if the probability of the lowest state is smaller than that of the highest state. Now suppose that the history of trades is such that state V_1 can be ignored relative to states V_2 and V_3. The basic intuition for the sufficiency of the U-shape is then that someone with a U-shaped signal would put more weight on V_3 than V_2, and thus, once state V_1 is insignificant enough he would start to buy-herd.

Now suppose there is herding: why does this imply that the S_2's signal is U-shaped? Suppose it is negative hill-shaped. Then type S_2 always shifts less weight towards the tails of his posterior and more towards the posterior's center relative to the prior. Moreover, the shift from the tails towards the center is larger for value V_3 than for V_1. This redistribution of probability mass is true for any prior, and thus when the low state is less likely than the highest state, then this trader still shifts weight towards the center, and more so for the low state. Thus his expectation is, intuitively, left of the center — for herding it must be to the right.

Now suppose we have the negative U-shape. Then the trader always shifts *more* weight towards the tails of his posterior and less towards the center. And while the shift into the tails is larger for V_1 than for V_3, when the low state is less likely than the high state, then eventually state V_3 may be important enough to outweigh states V_2 and V_1 so that type S_2's expectation moves to the right of the center.

With contrarian behavior the intuition is the reverse.

5 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 timeline (Appendix A), a full set of instructions and the material given to subjects (Appendix B), as well as a thorough description of the purpose-built software used for this experiment (Appendix C).

5.1 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% noise traders, with a central computer acting as the market maker.

Prior to each treatment each subject i received an informative private signal, described to them as a "broker's tip", $S_i \in S \in \{S1, S2, S3\}$. 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 them with both the signal distribution and the initial posterior distribution, conditional on receiving a signal, and we told the subjects as much. 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 (up to £25 (for the UK) or \$30 (for Canada)) combined with a one-off participation fee (£5 (UK) or \$15 (Canada)).

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.

Trades in each treatment were organized sequentially, with each subject or noise trader allocated a 5 second interval in which they, and only they, could act. This sequence of trading opportunities $t = 1, 2, 3, \ldots$ produced a history of actions and prices, $H_t = \{(a_1, P_1), \ldots, (a_{t-1}, P_{t-1})\}$ with $H_1 = \emptyset$. 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 $P_t = \mathsf{E}[V|H_t]$ with $P_1 = 100$.

Subjects were told that they had three possible actions $a = \{\text{sell, pass, buy}\}$ which they could undertake only during their own 5 second trading opportunity, t^*, and were also told that the precise time of their opportunity was exogenous, randomly determined and unique to them. The trading opportunity was indicated to subjects via a red border on the screen.

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.[15]

5.2 Treatments

Following Section 3, the rational action for $S1$ and $S3$ subject types was to sell or buy respectively, irrespective of H_t, while for the $S2$ types the nature of H_t and the precise information structure determined a unique optimal action. This action might be to herd or to act in a contrarian manner or neither. The treatments were each designed to enable us to examine a specific information structure:[16]

Treatment 1: negative hill-shaped signal structure making buy-contrarianism possible;
Treatment 2: increasing signal structure ruling out herding or contrarianism;
Treatment 3: negative u-shaped signal structure making buy-herding possible;
Treatment 4: decreasing signal structure ruling out herding or contrarianism;
Treatment 5: positive u-shaped signal structure making sell-herding possible;
Treatment 6: negative hill-shaped signal structure making buy-contrarianism possible.

Therefore under each treatment, once we knew the exact ordering of signals and trades, we could exactly calculate the theoretically predicted actions for each subject; in some cases this might be to herd or act in a contrarian way.

There was also an example treatment, during which the subjects could practise using the software, but which did not form part of the payment calculations, or the results.

[15] Appendix D 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.

[16] In Drehmann, Oechssler, and Roider (2005), the inclusion of transaction costs produced the expected outcomes. Thus we ignored transactions costs altogether which enabled us to focus on the information structure as the key differentiating factor between treatments.

5.3 Theoretical Predictions

Given the proximity of the design to the theoretical model outlined in Section 3, several predictions arose immediately:

1. $S1$ types should always sell;
2. $S3$ types should always buy;
3. $S2$ types' behavior should be a function of both the treatment and H_t.

Specifically, from the $S2$ types, we expect to see:

- Possible herding behavior in Treatments 3 and 5;
- Possible contrarianism in Treatments 1 and 6;
- No herding or contrarianism in Treatments 2 or 4.

Buy-herding is possible in Treatment 3 and sell-herding in Treatment 5. Therefore since we know the outcome of the random elements (noise trades, the exogenous ordering 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. This leads immediately to

Hypothesis 1: Subjects will act in accordance with their theoretically prescribed actions. We break this hypothesis down into six further sub-hypotheses as follows:

Hypothesis 1.1: No types should ever pass.
Hypothesis 1.2: $S1$ types should always sell, never herding or act in a contrarian way.
Hypothesis 1.3: $S3$ types should always buy, never herd or act in a contrarian way.
Hypothesis 1.4: $S2$ types will herd when the conditions required are met.
Hypothesis 1.5: $S2$ types will act as contrarians if the conditions required are met.
Hypothesis 1.6: $S2$ types will not herd or act as contrarians if the conditions are not met.

5.4 Behavioral Predictions

If complexity is an issue, it would seem likely that $S2$ 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 $S1$ and $S3$ types than the $S2$ types.

Needless to say 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 fully rational theory.

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 cannot exceed 125 and cannot 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.

In summary, should Hypothesis 1 fail in any of its forms, we have several further, alternative hypotheses that could be investigated. These are purposefully left as general as possible with the aim being to add more detail during the analysis of the results.

Hypothesis 2.1: (Complexity 1) $S2$ types are more likely to fail to act in accordance with their prescribed optimal actions than $S1$ or $S3$ types.

Hypothesis 2.2: (Complexity 2) $S2$ types fail to act in accordance with prescribed optimal actions when the decision is close.

Hypothesis 3: (Risk aversion) Subjects exhibit signs of risk aversion.

Hypothesis 4: (Prospect theory) Subjects exhibit signs of loss aversion.

Hypothesis 5: (Prior Action) Subjects act only on the basis of their ex-ante optimal action and ignore both prices and the history.

Hypothesis 6: (Prior expectation) Subjects do not update their beliefs at all as prices change but act solely on the basis of their prior expectation.

Hypothesis 7: (Probability Shifting) Subjects update their beliefs on the basis of changing prices at a slower rate than they should.

Hypothesis 8: (Noise) Subjects update their beliefs on the basis of an incorrect estimate of noise trades.

Hypothesis 9: (Trend chasing) Subjects tend to chase short-run trends, buying in response to observed buys by others, and selling in response to observed selling by others.

Note that Hypotheses 6-8 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 the subject would then deliberately slow down their own updating process. Attributing more trades to noise might well not be interpreted as subjects not believing the stated probability of noise trading, but rather as a measure of their lack of trust in their predecessors' level of rationality (instead considering some of them to be randomizing like noise traders).

We have also left the hypotheses unconstrained where a specific variable is important, so we can calibrate to the data. For example, in Hypothesis 7 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 the alternative noise evaluation in Hypothesis 8. 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.

6 Analysis of the Rational Benchmark

We will first examine the results directly, while supplementing it with a formal econometric analysis in Section 7.

In the numbers to follow we exclude noise trades, and focus only on trades by human subjects. The total number of trades was 1382 spread over all 6 treatments. We recorded 28 time-outs, leaving 1354 recognized trades. Time-outs will henceforth be omitted from the analysis except where explicitly stated otherwise.[17] The number of trades conditional on type were 394 ($S1$), 553 ($S2$) and 407 ($S3$).

6.1 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 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.[18] 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 145 passes (10.7% of all trades), 31 from $S1$ types (8%); 87 from $S2$ types (16%) and 27 from $S3$ types (7%). 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 small, especially for the $S1$ and $S3$ types. The figure of 16% for $S2$ types indicates that there is cause for some doubt about Hypothesis 1.1 (no passes) especially from those traders with the middle signal.

As with the overall numbers, we can tentatively support Hypothesis 1.1 for the $S1$ and $S3$ types, while being sufficiently uncertain about $S2$ types to warrant further examination of what might be motivating their behavior.

[17] We might wonder whether traders use passes and timeouts interchangeably. There were only 6 cases where a trader used both passes and timeouts. 35 traders used multiple passes (and/or timeouts) and only 18 traders recorded any timeouts. Since the number of timeouts was far greater in the example round this suggests than when motivated traders did not use timeouts as a substitute for passes, but that they were more likely accidental. We therefore elected to remove them from the analysis, though the tiny number means that the results are neutral to their inclusion.

[18] Our rationale for this specification was to avoid explaining short-selling to subjects.

6.2 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.

Omitting time-outs, the number of trades contradicting the theoretical model was 30.5% when counting passes as categorically incorrect. Now suppose we admit that passes may be 'weak buys'. Then all passes by $S1$ types are still irrational, whereas all passes by $S3$ types are admitted as rational. For the $S2$ types, passes are admitted as rational whenever the rational action was to 'buy'.

With this specification, the overall model fit increases from 69.5% to 74.6%. In total, broken up per trader, our model fit is

	$S1$	$S2$	$S3$	Total
counting passes as wrong	82%	53%	80%	69.5%
excluding passes from wrong	82%	61%	86%	74.6%

Table 1: **Fit of the data to the rational model.**

Examining Hypothesis 1 ('rationality'), we note that the number of trades that conformed to the theoretical predictions was 69% treating passes as suboptimal (in accordance with Hypothesis 1.1 ('no passes')). 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.

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 $S1$ and $S3$ 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 $S1$ and $S3$ types are acting in accordance with the rational theory.

The $S2$ types, however, often do not act rationally — almost half of their trades were against the rational model. Table 2 illustrates this broken up by treatments. In particular in the herding Treatments 3 and 5, the $S2$ do quite poorly, even when admitting passes as weak buys (22% and 37% fit). Had they taken each action at random they would have done better.[19]

At the same time, the $S2$ types face a more difficult decision problem than the $S1$ and $S3$ types. Theoretically, the decisions of the $S1$ and $S3$ types never change, so they can take the correct decision even without following the history carefully. The $S2$ types

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

on the other hand, have to follow the history carefully and small mis-computations can cause them to be categorized as irrational. However, below in Section 8.1 we show that allowing some leeway with respect to errors does not improve the fit.

	Wrong trades (all passes are wrong) S1	S2	S3	Wrong trades (some passes are ok) S1	S2	S3
Treatment 1 negative hill-shape	7	36	23	7	36	14
	53	98	71	53	98	71
	87%	63%	68%	87%	63%	80%
Treatment 2 increasing	10	40	17	10	28	10
	70	84	67	70	84	67
	86%	52%	75%	86%	67%	85%
Treatment 3 negative U-shape	10	73	19	10	59	12
	59	94	74	59	94	74
	83%	22%	74%	83%	37%	84%
Treatment 4 decreasing	13	26	9	13	26	6
	57	99	74	57	99	74
	77%	74%	88%	77%	74%	92%
Treatment 5 positive U-shape	9	56	8	9	44	7
	68	95	63	68	95	63
	87%	41%	87%	87%	54%	89%
Treatment 6 negative hill-shape	20	31	6	20	27	6
	87	83	58	87	83	58
	77%	63%	90%	77%	67%	90%

Table 2: **Total fit of the rational model by treatment.**

6.3 Herding and Contrarianism

In the Section 3 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. Herding and contrarian behavior can be classified either excluding passes or including passes as 'weak buys'; we look at both cases separately.

Herding and Contrarianism at the Aggregate Level

We begin by aggregating behavior across all treatments.

Herding in the strict sense: only rational herding is allowed. Looking only at the cases when herding is theoretically mandated, then we see that, when including passes as weak buys, rational herding occurs only in 31% of the cases when it should. If we count passes as irrational, the fit is even worse with only 18% of herds occurring when they should.

The fit for contrarian behavior is better (62% when passes are wrong, 77% when passes are weak buys), although the total number of contrarian trades as required by rationality is rather small compared to the required herding trades.

	Herding		Contrarianism	
	excl. passes	incl. passes	excl. passes	incl. passes
rational herds/contra	20	34	16	20
required by rational	109	109	26	26
in %	18%	31%	62%	77%

Table 3: **Occurrence of Rational Herding/Contrarianism**

Herding in the loose sense: irrational herding allowed. As argued above, the $S2$ types face a computationally difficult decision. We thus ask now if the $S2$ 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 selling to buying if their signals prices fall?

The following table gives the raw numbers for these scenarios. When including passes, herding arises in about 20% of the cases where it is possible, contrarian behavior arises in 25% of the possible cases.

	Possible herds/contra	realized herd/contra	percentage
total herd without passes	745	89	12%
total herd with passes	745	148	20%
total contra without passes	578	129	22%
total contra with passes	578	146	25%

Table 4: **Possible Herds**

Breaking these up by trader types, one can see that $S2$ types have the highest propensity to herd —26% including passes as weak buys— and to act as contrarians —46% including passes.

	Herding Trades			Contrarian Trades		
	$S1$	$S2$	$S3$	$S1$	$S2$	$S3$
possible	300	363	82	87	179	312
actual trades excl. passes	29	55	5	9	76	44
in %	10%	15%	6%	10%	42%	14%
actual trades incl. passes	49	94	5	20	82	44
in %	16%	26%	6%	23%	46%	14%

Table 5: **Possible Herds by types**

While it is never optimal for the $S1$ types or the $S3$ types to herd or act as contrarians, the behavior of the $S2$ types is at least loosely in line with the spirit of the model as motivated in the introduction. The regression analysis in Section 7 will stress further that herding-types are more likely to herd than any other type, and contrarian types are more likely to act as contrarians than any other type.

Herding and Contrarianism by Treatments

Herding in the strict sense: only rational herding is allowed. The results indicate that both herding and contrarian behavior occurred less often than they should, though the percentages vary considerably by treatment type.

The performance of herding candidates in the positive U-shaped Treatment 5 is rather poor: 90% of the required herds did not occur. However, the number of observations is also rather small. This is due to sequences of trader arrivals that hardly ever led to falling prices — as would have been necessary to trigger sell-herding.

The performance is somewhat better with the negative U-shaped Treatment 3: counting passes as 'weak buys' the fraction of missing herds is 'only' 64%. Notably, this is indeed a much larger fraction of herds than observed in Drehmann et al (2005) or Cipriani and Guarino (2005) (where herding behavior is linked to irrationality and rarely accounts for more than 10-20% of trades, and usually much less), lending some support to the hypothesis that the U-shaped signal structure matters. In the next section we will back this observation by a formal econometric analysis, showing that U-shaped signals indeed explain herding best.

Contrarian behavior has a better performance than herding, although the number of cases with theoretically mandated contrarianism is rather small. One obvious explanation for the better fit is that the hill-shaped contrarian signal is much easier to interpret since it indicates that the true value is the middle one. Consequently, it is comparatively simple for subjects to pick an action that moves prices in the direction of this middle value.

		Required	missing herds/contra w/o pass	rational herds/ contra	% not rational	missing herds/ contra w/ pass	rational pass	% not rational
T1	−ve hill	6	2	4	33%	2	0	33%
T6	−ve hill	20	8	12	40%	4	4	20%
T3	−ve U	89	71	18	80%	57	14	64%
T5	+ve U	20	18	2	90%	18	0	90%

Table 6: **Herding trades split up by treatment.**

Herding in the loose sense: irrational herding allowed. Applying the herding definition at face value (i.e. without insisting that the expectations exceed the price for buy-herding) the following tables detail when herding or contrarian behavior could have occurred and when it did.

Whether or not herd behavior can potentially occur depends on which type is called to trade and on whether prices had gone up beforehand. For instance, $S3$ types cannot herd when prices go up (but could act as contrarians). Hence, for every trade there is either the potential for herd or contrarian behavior.

Looking at both Tables 7 and 8, we observe that the $S2$ type is the trader who is most prone to herding and to act as a contrarian.[20] This indicates that traders with an intermediate signal are most likely to switch their behavior.

Some additional understanding can be obtained from studying the raw data. It appears that the relatively low conversion rate of theoretical to actual herd or contrarian behavior is partly caused by the price level at the point when potential herd or contrarian trading decisions have to be made. Prices were often very high, i.e. close to 125, and hence the scope for profit was relatively low when following the crowd, leaving traders tempted to revert to the mean, rather than to follow their optimal predicted action.

Revisiting the Herding/Contrarian Results by Trader Type. There was almost no herding from $S3$ types which is as predicted by the theory, however there was persistent contrarian behavior. Notably, across all treatments prices typically rose over time, which results in public behavior confirming $S3$ signals. Also, when prices rise, $S3$ types cannot

[20]The exception is Treatment 5 (decreasing) where, however, the numbers of realized herds are very small to begin with.

	Number of Herds/Possible Herds					
	excluding passes			including passes		
	$S1$	$S2$	$S3$	$S1$	$S2$	$S3$
Treatment 1 negative hill-shape	4	11	1	7	34	1
	52	91	6	52	91	6
	8%	12%	17%	13%	37%	17%
Treatment 2 increasing	5	5	1	8	5	1
	59	20	7	59	20	7
	8%	25%	14%	14%	25%	14%
Treatment 3 negative U-shape	3	20	0	10	21	0
	59	94	1	59	94	1
	5%	21%	0%	17%	22%	0%
Treatment 4 decreasing	5	10	0	7	13	0
	34	63	24	34	63	24
	15%	16%	0%	21%	21%	0%
Treatment 5 positive U-shape	5	2	2	7	2	2
	48	36	25	48	36	25
	10%	6%	8%	15%	6%	8%
Treatment 6 negative hill-shape	7	7	1	10	19	1
	48	59	19	48	59	19
	15%	12%	5%	21%	32%	5%

Table 7: **Herding trades split up by treatment.** For each treatment, entries in the top row indicate herding that did occur, entries in the middle row refer to the number of times that herding could have occurred. The third row entries indicate the percentage of realized herds.

herd by definition. Indeed, the total number of times that $S3$ types could herd was very small, as can be seen from Table 7, herding is possible for 82 (or 20%) of the 407 $S3$ trades. However, with prices rising, potentially getting very close to 125, an $S3$ type might reasonably think that there is little to gain when buying at close to the theoretical maximum price, and so contrarianism might be more defensible than herding behavior.

This is tentatively supported by the behavior of the $S1$ types who exhibited contrarian behavior especially in Treatment 6 and only when prices fell (the negative hill-shape case). There was also a degree of herding behavior by $S1$ types across all treatments.

For $S2$ types, who are the only type who should theoretically herd or act in a contrarian manner there is some evidence that herding occurs when it should (in Treatment 3 in particular), but there was little to no observed herding in Treatment 5 when prices fell. Similarly there was some evidence of $S2$ types exhibiting contrarian behavior when they should in Treatment 6, with falling prices (the numbers for Treatment 1 are very small). For example, in Treatment 2 prices were typically increasing, and $S2$ types tended to herd

	Number of contrarians					
	excluding passes			including passes		
	$S1$	$S2$	$S3$	$S1$	$S2$	$S3$
Treatment 1 negative hill-shape	0	4	13	0	4	13
	1	7	65	1	7	65
	0%	57%	20%	0%	57%	20%
Treatment 2 increasing	1	23	9	2	23	9
	9	64	47	9	64	47
	11%	36%	19%	22%	36%	19%
Treatment 3 negative U-shape	0	0	6	0	0	6
	0	0	73	0	0	73
			8%			8%
Treatment 4 decreasing	2	6	6	6	10	6
	19	25	50	19	25	50
	11%	24%	12%	32%	40%	12%
Treatment 5 positive U-shape	1	27	5	2	28	5
	19	59	38	19	59	38
	5%	46%	13%	11%	47%	13%
Treatment 6 negative hill-shape	5	16	5	10	17	5
	39	24	39	39	24	39
	13%	67%	13%	26%	71%	13%

Table 8: **Contrarian trades split up by treatment.**

more than act in a contrarian way, which is counter to the theory. In fact, in general, $S2$ types did occasionally exhibit contrarian behavior across all of the treatments if and when prices rose.

7 The Importance of the Price: Regression Analysis

The summary statistics from the last section gave a rough idea of the determinants of behavior: first, the 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 more frequently than herding. Yet prior studies suggest that people irrationally tend to act as contrarians; thus one might expect that rational contrarianism works in the same direction and might thus be observed even more accurately than any other behavior. Thus far, our numbers indicate that contrarianism does not arise quite as often.

We now take a closer look and run several regressions to test the direct impact of herding and contrarian signals relative to occurrences 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?

We tackle these questions from several angles: first, we perform an analysis for the entire data set of trades restricted to the cases where respectively herding and contrarian behavior is possible. (For instance, prices must have risen and the respective trader must have a signal that would allow him to act as a contrarian (e.g. a positive hill-shaped $S2$ or signal $S3$).) In a second step, we restrict attention to only the actions of $S2$ types. In each scenario we first run plain logit and OLS regressions and then control for trader fixed effects to ensure that our results are not contaminated by traders who, for instance, persistently err.

Overall, we find that while the effect of a herding signal is insignificant in the sample of $S2$ traders, it is significant and positive in the total sample (with and without controlling for trader fixed effects). This confirms the intuition derived from Section 6 that recipients of $S2$ herding-type signals are generally more likely to herd.

Next, the effect of a contrarian signal is significant both in the total sample and in the $S2$ trader sub-sample (with and without controlling for fixed effects).[21] Consequently, the regression analysis in the total sample clearly indicates that an $S2$ contrarian signal is a significant cause of contrarianism relative to all other kinds of signals.

We ran some further, unreported regressions, in an attempt to understand if herding signals have a significant effect on the probability of observing, say, buys. Effectively, one can think about herding as an upward-sloping demand function, that is, the recipient of a herding signal is more likely to buy if prices increase. However, we did not find a significant effect. Considering the findings of past studies, it is not surprising that we do not find evidence of an upward-sloping demand function: past studies identify contrarian behavior[22] even by $S1$ and $S3$ types. So as prices move towards extreme values, herding $S2$ types should be expected to display a similar kind of contrarianism that is observed in other types. In other words, one could imagine that the occurrence of herding in our model relative to the current price level indicates a hill-shaped demand function.

Turning to the behavior of the $S1$ and $S3$ types, we have already observed that even

[21]The insignificant coefficient in the fixed effects regression is likely due to the small number of $S2$ contrarian trading-situations, i.e. there are only few situations where contrarian trades by hill-shaped $S2$ traders are actually possible; see also Table 8 for the total number.

[22]Our findings on the behavior of $S3$ types is in line with the findings in Drehmann, Oechssler and Roider (2005) and Cipriano and Guarino (2005), both of which find those with high signals selling at sufficiently high prices.

	All candidate herding trades			S2 candidate herding trades		
	Logit	OLS	Fixed effects	Logit	OLS	Fixed effects
herding signal	0.059*	0.069*	0.076*	0.04	0.041	0.065
	(0.026)	(0.030)	(0.033)	(0.037)	(0.038)	(0.049)
Constant	-0.215**	0.100**	0.099**	-0.233**	0.128**	0.120**
	(0.009)	(0.013)	(0.012)	(0.014)	(0.022)	(0.025)
Observations	761	761	761	372	372	372
R^2		0.01	0.4		0	0.6

Table 9: **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. Fixed effects regressions control for trader-fixed effects. The data is restricted to include only trades that could be herding trades.

	All candidate contrarian trades			S2 candidate contrarian trades		
	Logit	OLS	Fixed effects	Logit	OLS	Fixed effects
contrarian signal	0.336**	0.458**	0.355**	0.267**	0.267**	0.128
	(0.065)	(0.073)	(0.088)	(0.101)	(0.096)	(0.202)
Constant	-0.239**	0.187**	0.193**	-0.121**	0.378**	0.402**
	(0.004)	(0.017)	(0.016)	(0.039)	(0.040)	(0.052)
Observations	591	591	591	179	179	179
R^2		0.06	0.5		0.04	0.69

Table 10: **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 contrarian trades.

$S3$ types engage in selling and that $S1$ types engage in buying. To this end, we ran multinomial logit regressions of the decisions (buy/hold/sell) on price, controlling for these two signals.[23] Theory (Hypothesis 1.2 & 1.3) predicts that the price should have no impact on whether an $S3$ or $S1$ type buys or sells. However, for the $S3$ types we find that the signs of the first order effect of a price increase is significant and negative, implying that as prices increase, the $S3$ types become more likely to sell. For the $S1$ types, the coefficient for price changes is insignificant (albeit it goes in the right direction so that there is not even the slightest indication of herding), as predicted by theory. That being said, we only have comparatively few sessions with low prices, so that we cannot assert with confidence whether the $S1$ types might act as contrarians if the prices were close to the lowest value.

[23]Given the well-known independence problems when running multinomial logit we focused mainly on standard logit, and also note that when using multinomial logit, the results should be seen as *at best* indicative.

	$S3$		$S1$	
	passes are buys	passes are omitted	passes are buys	passes are omitted
% price change	-0.549**	-0.692**	-0.231	-0.075
	(0.167)	(0.176)	(0.209)	(0.177)
constant	0.246**	0.255**	-0.206**	-0.195**
	(0.016)	(0.016)	(0.016)	(0.018)
number of observations	407	380	394	363

Table 11: **The decision to buy for $S3$ and $S1$.** The table displays the results from a logit regressing the decision to buy on the price change, as measured relative to the prior, for all $S3$ and $S1$ traders respectively. For the $S3$, the probability of a buy is decreasing as the price rises; the hypothesis is that the coefficient is insignificant ($S3$ always buy). For the $S1$, the coefficients are insignificant, consistent with theoretical predictions.

8 Alternative Explanations for Trading Behavior

We have seen in Section 6 that some results are supportive of the theory. Moreover, the formal regression analysis in Section 7 is in favor of the qualitative predictions provided by the Park and Sabourian's analysis. 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.

It is therefore prudent to explore alternative explanations, ranging from omitting decisions that are within an ϵ error-region, to risk- and loss-aversion, to various forms of alternative information updating. These alternative hypotheses 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. Such an approach is, of course, a maximum likelihood technique, albeit a very coarse one.

8.1 Robustness to Small Errors

Hypotheses 2.1 and 2.2 consider the scope for computational complexity to trigger suboptimal behavior. We have already argued that since the $S2$ types have to make a decision that is conditional on H_{t-1} they have a more difficult decision to make (Hypothesis 2.1). They are therefore more likely to make errors than the $S1$ types or $S3$ types. And indeed, the error rates of $S2$ types, as seen above, are highest and so we have some supportive evidence for this view.

Secondly, we might consider that when prices and expectations are close, traders are more likely to make incorrect choices (Hypothesis 2.2). To test this idea, we omit all trades that occur when prices and expectations are within ϵ of each other, for small values of ϵ.

This variation typically worsens the fit of the model. For instance if we set $\epsilon = .01$, then the total model fit is reduced to 67% (for $S1$, $S2$ and $S3$, 78%, 45%, and 80% respectively). We repeated the analysis with different values for ϵ but could not generate a higher fit for the data.

8.2 Risk and Loss Aversion

Risk Aversion. One persistent finding from the last section is that traders exhibit a general tendency to act as contrarians. One might thus also entertain the idea that traders act as contrarians because of risk-aversion (Hypothesis 4). 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 because 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 traverses 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 payoff tomorrow. 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 below 50%; for CARA it is 48% for $\rho = 1$ and 64.5% for $\rho = .01$, rising 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.

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

	Total Number of wrong decisions CRRA utility, $\gamma = 1$ (log-utility)			Total Number of wrong decisions CARA utility, $\rho = 1$			Total Number of wrong decisions CARA utility, $\rho = 0.01$		
	S1	S2	S3	S1	S2	S3	S1	S2	S3
Treatment 1 negative hill-shape	18 33%	58 58%	37 51%	46 84%	60 60%	14 19%	46 84%	58 58%	48 67%
Treatment 2 increasing	31 42%	60 67%	36 53%	60 81%	28 31%	10 15%	60 81%	30 33%	50 74%
Treatment 3 negative U-shape	13 22%	69 73%	37 49%	49 82%	59 63%	6 8%	49 82%	18 19%	61 80%
Treatment 4 decreasing	33 57%	41 41%	33 45%	44 76%	74 75%	6 8%	44 76%	75 76%	65 88%
Treatment 5 positive U-shape	30 43%	66 67%	32 49%	59 86%	29 30%	7 11%	59 86%	25 26%	55 85%
Treatment 6 negative hill-shape	41 47%	51 61%	22 38%	67 76%	43 51%	6 10%	67 76%	48 57%	52 90%
Total number wrong percentage wrong	**166** **41%**	**345** **61%**	**197** **48%**	**369** **80%**	**344** **52%**	**57** **12%**	**369** **80%**	**279** **42%**	**390** **80%**
Total model fit	48.8%			47.9%			64.5%		

Table 12: **Risk-Aversion Analysis.**

S-shaped. Kahnemann and Tversky suggested the following function 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 38%. Table 13 illustrates this observation for the above parameters as well as for one other configuration.[24]

8.3 Decision Rule: Prior Actions

One alternative decision rule formulation is that of naïve traders who ignore the history and who simply stick to their prior action. As such, $S1$ types always sell, $S3$ types always buy and $S2$ types pick the actions that is prescribed at the initial history. For instance, with negative U-shape, $S2$ traders always sell.

[24]Arguably, 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 13: **Loss-Aversion Analysis.**

Initially, this specification appears to do very well: it fits 73.6% of the data which is higher than the rational model (with passes as wrong trades); broken up by type the fit is 82% for $S1$, 63% for $S2$ and 80% for the $S3$, which is again higher than the rational fit without passes but lower than the rational model with passes.

Of course 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 $S1$ and $S3$ types. Since the econometric analysis has already revealed that the $S3$ are sensitive to the price, this decision rule is rather weak.

Indeed, looking only at those actions that are at odds with rationality (and counting passes as wrong as would befit the hypothesis here), only 30% (79 decisions) of the irrational actions of the $S2$ are in line with this hypothesis. This further reveals that the remaining 183 decisions are due to a change of actions, which constitute 33% of the $S2$'s actions.

In summary, while the fit seems high, the model does not help explain any of the observed changes of actions.

8.4 Decision Rule: Prior Expectations without updating

The Decision Rule "Prior Actions" asserts that people ignore prices altogether; a weaker variation of this theme has people ignore the history altogether but are mindful of the price. People thus act based only their prior expectation: if the price exceeds it, they sell, if the price is below it, they buy (Hypothesis 6).

And indeed about 72.2% of people take an action that is in accordance with their prior expectation. Counting passes as wrong trades, trades can be classified as 'following the prior' for 1% of $S1$-, 45% of $S2$- and 23% of $S3$-trades that were classified as wrong.

For instance, for the $S3$ types this means that they do not buy when they should be buying, or for the $S2$ types that they do not herd when they should be herding. Most treatments resulted in increasing prices and so it would be difficult to have a specification of prices and signals in which the $S1$ types were induced to follow the prior, which might explain why the numbers for $S1$ types are so much lower than for $S3$ types (the single rationally 'wrong' trade that is being explained by this decision rule is a contrarian buy by the $S1$ types).

8.5 Probability Scaling and Shifting

We might also consider probability shifting (Hypothesis 7), whereby traders underplay (overplay) low (high) probabilities coming from the observed history H_{t-1}. Alternatively, traders may overstate the probabilities of their prior expectations. The usual symmetric treatment of this under- or overstating of probabilities is to transform probability p into $f(p)$[25] 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 prior signal we would apply an $\alpha > 1$ re-scaling on the initial probabilities. Alternatively,

[25] There 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 ${p_j}^\alpha$ for all states, $j = 1, \ldots, 3$.

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 not matter as much as under the rational model. We tried both specifications and the results are similar.

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

	With $\alpha = 5$			With $\alpha = 10$			With $\alpha = 25$		
	S1	S2	S3	S1	S2	S3	S1	S2	S3
Treatment 1 negative hill-shape	46	61	57	46	61	57	46	61	57
	7	37	14	7	37	14	7	37	14
	87%	62%	80%	87%	62%	80%	87%	62%	80%
Treatment 2 increasing	60	56	57	60	56	57	60	56	57
	10	28	10	10	28	10	10	28	10
	86%	67%	85%	86%	67%	85%	86%	67%	85%
Treatment 3 negative U-shape	49	31	68	49	41	68	49	57	68
	10	62	6	10	52	6	10	36	6
	83%	33%	92%	83%	44%	92%	83%	61%	92%
Treatment 4 decreasing	44	73	65	44	73	65	44	73	65
	13	25	6	13	25	6	13	25	6
	77%	74%	92%	77%	74%	92%	77%	74%	92%
Treatment 5 positive U-shape	59	66	56	59	66	56	59	66	56
	9	29	7	9	29	7	9	29	7
	87%	69%	89%	87%	69%	89%	87%	69%	89%
Treatment 6 negative hill-shape	67	61	52	67	61	52	67	61	52
	20	22	6	20	22	6	20	22	6
	77%	73%	90%	77%	73%	90%	77%	73%	90%
Total trades in line with probability scaling	**325**	**348**	**355**	**325**	**358**	**355**	**325**	**374**	**355**
Percentage explained	84%	64%	88%	84%	66%	88%	84%	69%	88%
total fit		77.1%			77.9%			79.1%	

Table 14: **Probability Scaling.**

Comparing the results here to those in Table 2, one can see that the fit of probability scaling hardly improves for the $S1$ and $S3$ 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 will do a poor job explaining herd-behavior of any sort.

8.6 Overestimating Noise Trading

An alternative specification for hampered updating is to assert that agents overestimate the amount of noise (Hypothesis 8); conceptually this is similar to a Quantal Response Model (QRE) in which agents recursively take their predecessor's decisions as prone to error.[26] The idea for this specification is in spirit similar to probability shifting but the fit obtained with the overweighing of one's own signal provides a better fit with the data. Compared to the rational model the improvement of the fit is minor again as Table 15 illustrates.

	noise .5			noise .75		
	S1	S2	S3	S1	S2	S3
Treatment 1 negative hill-shape	46	62	46	46	61	47
	7	36	25	7	37	24
	87%	63%	65%	87%	62%	66%
Treatment 2 increasing	60	35	50	60	38	49
	10	49	17	10	46	18
	86%	42%	75%	86%	45%	73%
Treatment 3 negative U-shape	49	35	57	49	59	55
	10	59	17	10	35	19
	83%	37%	77%	83%	63%	74%
Treatment 4 decreasing	43	74	64	43	74	65
	14	25	10	14	25	9
	75%	75%	86%	75%	75%	88%
Treatment 5 positive U-shape	59	38	55	59	44	55
	9	57	8	9	51	8
	87%	40%	87%	87%	46%	87%
Treatment 6 negative hill-shape	67	50	52	67	49	49
	20	33	6	20	34	9
	77%	60%	90%	77%	59%	84%
Total identified	**324**	**294**	**324**	**324**	**325**	**320**
Total fit	82%	53%	80%	82%	59%	79%
Total model fit	69.6%			71.6%		

Table 15: **Variations in the Perception of Noise Trading.**

8.7 Short-Term Trend-Chasing

We also consider the simple hypothesis that people chase trends, and so are deferential towards the price movements that they observe. To examine this we computed the trend

[26]There is a subtle difference to the way that Quantal Response Models can be implemented in models with and without prices: in Anderson and Holt, who originally established this notion, agents were in an informational cascade and thus a deviation from the cascading action was, 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 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.

in prices over the last 3 or 5 periods. If the price is above the trend, then we considered the suggested strategy to be sell, if below, then buy. This model performed poorly as a predictor of behavior, with the average fit of a trend-trading model below 50%, and the model seems especially poor at explaining the behavior of the $S2$ types. We thus omit the details of the data analysis.

8.8 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 and overestimating noise trading are essentially very similar, and also have strong similarities to a strategy of following the prior (which is effectively a policy of zero updating).

Several studies (Drehmann, Oechssler and Roider (2005) and Cipriani and Guarino (2005)) have already identified that when prices grow, 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 $S3$ types. Symmetrically, the $S1$ 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 endpoint effect should also influence the $S2$ types, because whatever mechanism or cognitive bias leads $S3$ types to sell for high prices should apply in the same manner to $S2$ 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. If indeed people do update slower then we should observe three things:

- In treatment 2, with an increasing information structure, when prices rise, the $S2$ should start (contrarian) selling at prices before the $S3$ types. In the data we do in fact observe that a much larger fraction of $S2$ act as contrarians than $S3$.

- We should see more irrational herding by $S2$ types than $S1$ types in treatment 4 with a decreasing information structure. In the data, while we do observe less irrational herding by $S2$ types than $S1$ types, the difference is minor.

- In the hill-shaped information structure treatments (1 and 6), herding should not arise. In our data we observe that 19% or 13% (respectively) of trades are herd trades (compared for example to 14% for treatment 3 under u-shaped negative information).

9 Price Efficiency

By now we have established that rational herding and contrarianism occur, although not as often as prescribed by the theory. To compliment this, some subjects irrationally herd, and to an even greater extent, some types irrationally act as contrarians. This all raises the question: does it matter? Should we care? We now argue that it does matter and that one should care because irrational actions severely and somewhat systematically affect market efficiency.

The basis of our experiment is a Glosten-Milgrom type sequential trading model in which prices are informationally efficient by setup. Thus if people deviate from the theoretically prescribed behavior, then prices will automatically be inefficient. As indicated by the title of this paper we now assess how the *lack of herding* or *lack of contrarianism* affects price efficiency.

In more detail: the empirical literature usually distinguishes the strong, the semi-strong and the weak form of the efficient market hypothesis. The strong form applies when prices reflect all private information; theoretically this implies prices react so as to precisely account for the private signal ($S1$, $S2$ or $S3$) of the current trader (arguably allowing for noise trading). Yet the model underlying our experiment is structured so that in equilibrium there are almost always two signal types that take the same action, e.g. $S1$ and $S2$ types both sell or $S2$ and $S3$ types buy. Prices can therefore inherently not reveal individual signals and so even theoretically the strong form cannot apply.

Next, under the semi-strong form, prices include all public information; Glosten-Milgrom style models usually satisfy this form of market efficiency, in particular because the bid- and ask-prices anticipate the information content of a potentially upcoming sale and buy. Since we dispense with bid- and ask-prices, the most one can hope for is that prices are weak-form efficient, i.e. prices include all information contained in past trading.

Of course, whenever people deviate from the rational model, prices are not as prescribed by the efficient benchmark. This rational benchmark is the price that would result if all traders in the model act as theoretically predicted. In the case of our experiment we can then compute that while about 16% of the observed prices coincide with the efficient one, on average observed prices are 7% below the efficient one, with a standard deviation of 9.9%. The left panel in Figure 1 plots the distributions of the deviations of prices from the efficient benchmark. This observation is, of course, in line with our prior findings: we observed that $S3$ types tend to act as contrarian-sellers as prices rise and that $S2$ types tend to not buy-herd when they should. Combined with generally increasing prices this intuitively leads one to predict that observed prices are often not as high as they should be.

To measure the extent of the inefficiency, we start with the argument that less-than rational behavior yields prices that are noisier than predicted by the theory. We can therefore determine the level of noise that is implied by the observed price. Our procedure is as follows: the theoretical level of noise trading is $\lambda = 25\%$. First, we can search for the level of noise trading $\lambda^\star$ for which the distance between this new, pseudo-rational price and the observed price is smallest (measured by minimizing the cumulative squared difference). Secondly, we will also look for the level of noise so that, the average difference is closest to zero.

The outcome of the first exercise is that the minimal value is attained for $\lambda^\star \approx 0.895$. In other words, the observed prices 'correspond' to a noise level that is roughly 3.6 times larger than the one that was modeled. Means for the price-difference for this level of noise are -3.5% and the standard deviation is 9.8%. The middle panel of Figure 1 plots the distribution of price differences (ideally, of course we would like to see a distribution that is tighly centered around 0%).

Next, in order to attain an average price difference of zero for $\lambda^\star \approx 0.75$ with standard deviation 10.9%. The right panel of Figure 1 plots the distribution of price differences for this noise level.

10 Conclusion

The experimental study presented in this paper focuses on a multi-state multi-signal model which admits both herding and contrarian behavior as rational outcomes given the right information structure. We observe that subjects generally behave according to theoretical predictions, yet the fit is lower for types who have the theoretical potential to herd or act in a contrarian manner. Herding (sometimes irrational) does occur, and the types identified by the theory as potential herders, are indeed by far the most likely to herd, but also the most likely to behave at odds with theory.

Examining alternative models of behavior (such as risk aversion, loss aversion or error correction) leaves us with the impression that these specifications perform either quite poorly or add little to our understanding over and above the rational model. Examining price efficiency we see that the observed prices are some way from efficient, and are captured by assuming a theoretical level of noise trading significantly above the actual

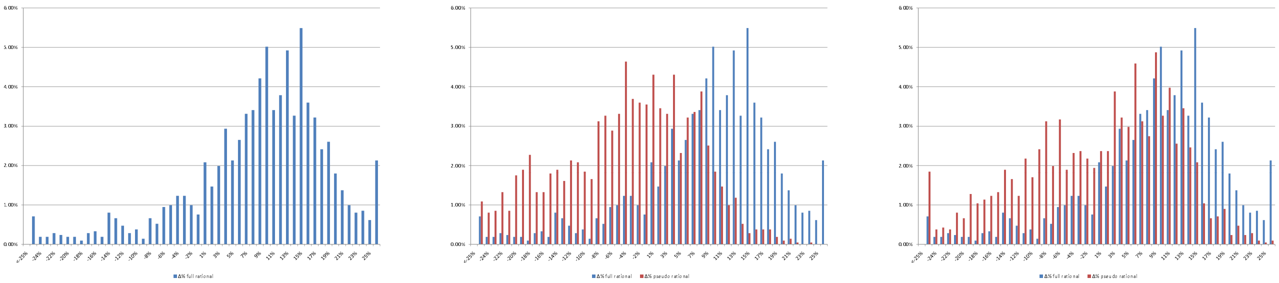

Figure 1: **Histograms of Deviations from the Efficient Benchmark Price.** The left panel plots the distributions of the deviations of prices from the efficient benchmark. The middle panel plots the distribution of price differences assuming noise is set to minimize the cumulative squared differences of price deviations; this noise level is $\lambda^\star \approx .895$. The right panel plots the distribution of price differences that yields an average price difference of zero; this occurs for $\lambda^\star = .75$. The distribution from the left panel is repeated in the middle and in the left panel for illustrative purposes. The percentage deviation is measured as the difference of efficient and observed price, scaled by the efficient price. All plots omit the value for zero deviations for visual clarity; for the left panel this number is about 16%, for the price sequences depicted in the and middle panels, this number would be about 6.7-6.8%.

proportion of noise traders used in the experiment.

Perhaps the most straightforward lessons to be learned are that herding and contrarian behavior happen in the lab, and seem to be as well described by rational behavior as other alternatives, so long as that theory allows herding and contrarianism as rational phenomena. By examining a heterogenous set of traders, differentiated by their signal, we also provide confirming evidence for the theory that while traders with conflicting information are most prone to rational herding, they are also prone to irrational behavior. Finally, irrational contrarianism in particular leads observed prices to substantially differ from the rational efficient benchmark, to the extent that the two can only be reconciled by assuming substantially more noise trading than was actually present.

Perhaps the most significant addition to the experiment examined in this paper (and indeed to those examined in the literature review) would be to allow subjects to make decisions as and when they wish. Incorporating endogenous timing into the design would much more closely mirror real-world financial markets and add an extra dimension of interest: the ability to delay to potentially learn more about the state. Work on this at both a theoretical and experimental level would be a fruitful future avenue for research on herding in financial markets.

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.

A 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.

B 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 your 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 a 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 $S3$), "bad tip" ($S1$) or "middle tip" ($S2$). A good $S3$ tip indicates that he believes the event will be good and the share price will be 125 vcu after it is realized, a bad $S1$ tip that something bad will happen indicates 75 after the event is realized. A middle $S2$ 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: S1, S2 or S3, 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 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 I use the central computer to calculate your final payments. I ask that you close the trading software and any other open software and shut down your computer. I 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.

C 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.

D 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)
6. If you would like to be added to the database of potential participants, and contacted if/when future sessions are organized please write your email address here:

About your decisions today

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

About the experiment

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

E 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 2 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).

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 3 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 S1 (low) and the subject had a single share to sell and a large cash balance to enable him to buy a further share. He could also pass (declining to buy or sell) when he was 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.[27]

[27]Further details about the software are available on request from the authors.

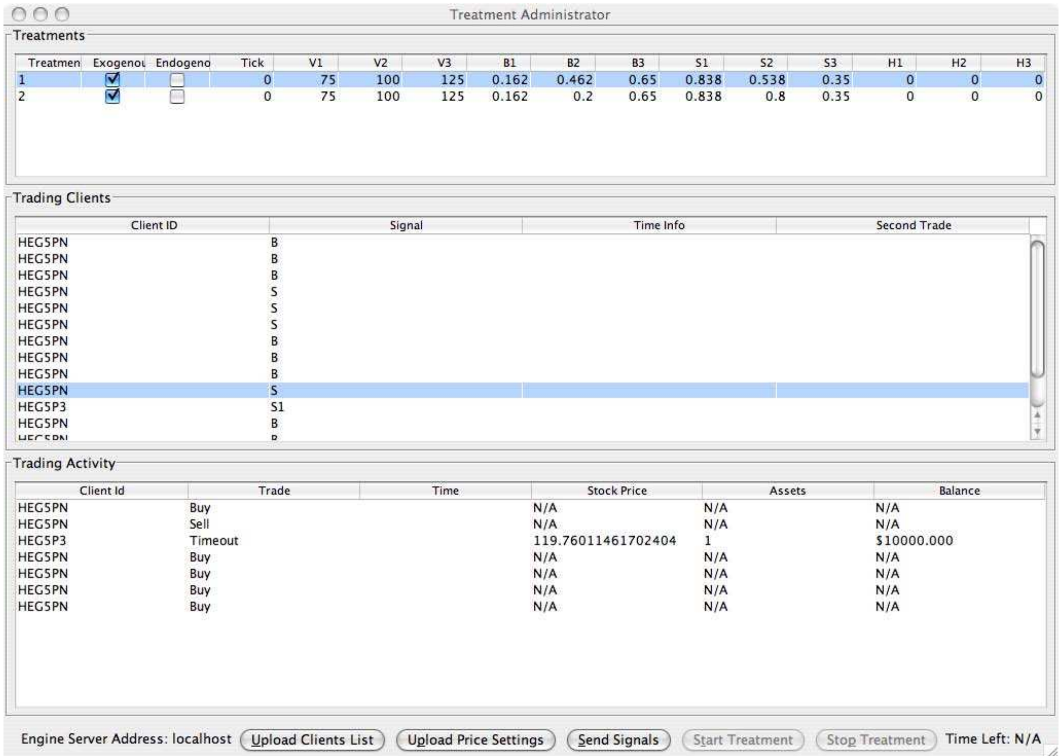

Figure 2: The Administrative Interface

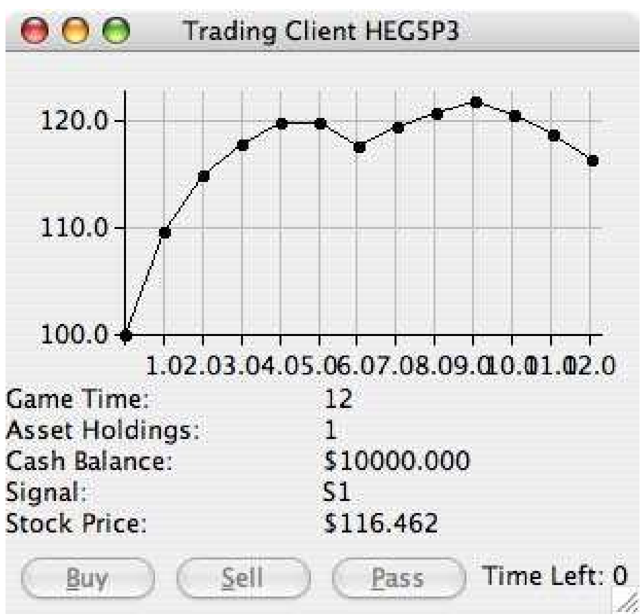

Figure 3: The Trading Client

References

Alevy, J. E., M. S. Haigh, and J. A. List (2007): "Information Cascades: Evidence from a Field Experiment with Financial Market Professionals," *Journal of Finance*, LXII(1), 151–180.

Anderson, L., and C. Holt (1997): "Informational Cascades in the Laboratory," *American Economic Review*, 87, 847–862.

Avery, C., and P. Zemsky (1998): "Multi-Dimensional Uncertainty and Herd Behavior in Financial Markets," *American Economic Review*, 88, 724–748.

Banerjee, A. (1992): "A Simple Model of Herd Behavior," *Quarterly Journal of Economics*, 107, 797–817.

Bikhchandani, S., D. Hirshleifer, and I. Welch (1992): "Theory of Fads, Fashion, Custom, and Structural Change as Informational Cascades," *Journal of Political Economy*, 100, 992–1026.

Bikhchandani, S., and S. Sunil (2000): "Herd Behavior in Financial Markets: A Review," Staff Paper WP/00/48, IMF.

Brunnermeier, M. K. (2001): *Asset Pricing under Asymmetric Information – Bubbles, Crashes, Technical Analysis, and Herding*. Oxford University Press, Oxford, England.

Chari, V., and P. J. Kehoe (2004): "Financial Crises as Herds: Overturning the Critiques," *Journal of Economic Theory*, 119(1), 128–150.

Chordia, T., R. Roll, and A. Subrahmanyam (2002): "Order imbalance, liquidity, and market returns," *Journal of Financial Economics*, 65(1), 111–130.

Cipriani, M., and A. Guarino (2003): "Herd Behavior and Contagion in Financial Markets," mimeo, New York University.

——— (2005): "Herd Behavior in a Laboratory Financial Market," *American Economic Review*, 95(5), 1427–1443.

——— (2007): "Herd Behavior in Financial Markets: A Field Experiment with Financial Market Professionals," mineo, UCL and GWU.

Dasgupta, A., and A. Prat (2005): "Asset Price Dynamics When Traders Care About Reputation," Working paper, London School of Economics and CEPR.

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

Glosten, L., and P. Milgrom (1985): "Bid, Ask and Transaction Prices in a Specialist Market with Heterogenously Informed Traders," *Journal of Financial Economics*, 14, 71–100.

Kahneman, D., and A. Tversky (1979): "Prospect Theory: An Analysis of Decision under Risk," *Econometrica*, 47(2), 263–292.

Lee, I. H. (1998): "Market Crashes and Informational Avalanches," *Review of Economic Studies*, 65(4), 741–759.

McKelvey, R., and T. Palfrey (1995): "Quantal Response Equilibria for Normal Form Games," *Games and Economic Behavior*, 10, 6–38.

——— (1998): "Quantal Response Equilibria for Extensive Form Games," *Experimental Economics*, 1, 9–41.

Park, A., and H. Sabourian (2008): "Herding and Contrarian Behavior in Financial Markets," Working paper, Universities of Toronto and Cambridge.

Shiller, R. (2008): "How a Bubble Stayed Under the Radar," *New York Times*, March 2.

Smith, L., and P. Sørensen (2000): "Pathological Outcomes of Observational Learning," *Econometrica*, 68, 371–398.

Tversky, A., and D. Kahneman (1992): "Advances in Prospect Theory: Cumulative Representation of Uncertainty," *Journal of Risk and Uncertainty*, 5, 297–323.