Herding, Contrarianism and Delay in Financial Market Trading[*]

Andreas Park[†]
University of Toronto

Daniel Sgroi[‡]
University of Warwick

September 30, 2009

Abstract

Herding and contrarian behavior are often-cited features of real-world financial markets. Theoretical models of continuous trading that study herding and contrarianism, however, usually do not allow traders to choose when to trade or to trade more than once. We present a large-scale experiment to explore these features within a tightly controlled laboratory environment. Herding and contrarianism are significantly more pronounced than in comparable studies that do not allow traders to time their decisions. Traders with extreme information tend to trade earliest, followed by those with information conducive to contrarianism, while those with the theoretical potential to herd delay the most.

JEL Classification: C91, D82, G14.
Keywords: Herding, Contrarianism, Endogenous-time Trading, Experiments.

[*]Financial support from the ESRC (grant number RES-156-25-0023), the SSHRC, the Leverhulme Trust and the Cambridge Endowment for Research in Finance (CERF) is gratefully acknowledged. Special thanks go to Annie Jekova and Malena Digiuni for expert research assistance. We thank seminar participants at the Universities of Cambridge, Toronto, Southampton, Warwick, and Zurich, and the Midwest Economic Theory Meetings 2008. We are grateful to Gustavo Bobonis, Ernst Fehr, Jordi Mondria, Jim Peck and Xiadong Zhu for extensive discussions.

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

1 Introduction

> "Were all these people stupid? It can't be. We have to consider the possibility that perfectly rational people can get caught up in a bubble. In this connection, it is helpful to refer to an important bit of economic theory about herd behavior." Robert Shiller, New York Times, March 2$^{\text{nd}}$, 2008.

Robert Shiller believes that rational herding can and does explain the current housing market difficulties and, by implication, the financial problems faced all over the world. The theory to which he explicitly refers was pioneered by Welch (1992), Bikhchandani, Hirshleifer, and Welch (1992) and Banerjee (1992) who highlight that rationality is no defence against the randomness of herd behavior.[1] Put simply, a few early incorrect decisions, through a process of rational observation and inference, can have serious ramifications for all who follow.[2]

A loose application of herding theory to financial market trading might suggest that early movements by visible traders can provide a catalyst for momentum trading, induce discontinuous price jumps in one direction or the other, and potentially leave share prices far from their fundamental value. If we could describe cases of wild, short-term movements or dubious run-ups of asset prices 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 controlling or averting financial crises.

Yet, the early work on rational herding was not designed to be directly applied to security market trading. First, this early work on herding does not consider prices that react to actions, while in financial markets (efficient) market prices drop after sales and rise after buys. Second, agents must act in a strict, exogenous sequence — they cannot decide when to trade. Arguably, in a market trading environment where learning from others is important, the timing of actions should matter a great deal. In fact, one of the key features of real-world financial frenzies is the clustering of actions in time, a phenomenon that cannot be examined when timing is not considered.

A path-breaking paper by Avery and Zemsky (1998) introduced efficient prices to a

[1] The first published paper on the breakdown of informational learning by rational agents is Welch (1992); it is also the first application of herding theory to a financial market setting.

[2] Consider a traditional herding setting in which agents receive an informative but noisy signal about which of two states is better. Suppose state A is truly worse than state B. Then it is possible that the first two agents happen to draw incorrect signals, and thereby opt for A. For agent 3, under a natural indifference condition, this means disregarding whatever signal she possesses and following the actions of the first two agents. All later agents find themselves in the same position as the third agent and will follow in the same manner even though they realize that it is only the information conveyed in the first two actions that determines behavior. As the direction of the herd disproportionately depends on the first movers, the ultimate outcome is exposed to a degree of randomness that is not warranted by fundamentals.

sequential herding context, but showed that in a simple financial market-trading setting with two values herding is not possible because the market price always separates people with good and bad information so that the former always buy and the latter always sell. However, with multiple states herding can arise when traders receive a particular kind of information; we will employ such a setting in our experimental setup.[3]

Moreover, we will also analyze the natural and often cited counterpart situation to herding in which people rationally act against the majority action. Such *rational contrarianism* has not been studied explicitly in experimental markets before, yet in real-world markets it is often cited as an important force for the mean-reversion of asset prices.[4]

Despite the considerable advances in the market microstructure literature there has been little work on how heterogeneously informed traders act in endogenous time, two prominent features of real-world markets. We focus here on Smith (2000), which provides sharp predictions for the optimal timing behavior, suggesting that traders with stronger private information will trade quickly, while those with weaker information will delay. Our work is also in the tradition of Bloomfield, O'Hara, and Saar (2005) who tackle a related issue with controlled laboratory experiments.[5] In effect we jointly test if the direction of trading is in line with rational herding theory, and if the timing of trades follows the predictions of Smith (2000).

When the timing of actions is a choice, the existence and the effect of herding and contrarianism are not immediately obvious. One can argue that eliminating exogenous timing removes an artificial friction, in the absence of which herding and contrarianism might vanish. Alternatively, behavior may turn out to be much more pronounced than with exogenous timing. For instance, traders prone to herd are likely less confident about the asset value may delay to learn from observing others. Once they do trade, their herding behavior may lead to stronger price distortions than under exogenous-timing. Furthermore, if traders with stronger signals act as early as predicted by Smith, then they will avoid herding altogether.

To summarize our results, in one of the largest laboratory experimental studies of its

[3]See, for example, Park and Sabourian (2008) which gives the conditions on information that must be satisfied to admit herding, contrarianism, or neither of the two.

[4]For example, see Chordia, Roll, and Subrahmanyam (2002) for a discussion of the importance of contrarianism in financial market trading. The exception is our companion paper, Park and Sgroi (2009) which studies contrarianism in a setting with exogenous timing.

[5]Bloom et al. also provide an insightful discussion on the complexity of the trade-timing decision and how controlled experiments can be used to glean insights into theoretically inaccessible problems. Moreover, experimental work can play an important complementary role to empirical analyses: transaction level data from standard databases such as TAQ does not reveal the party behind a transaction. This makes it difficult to test information based motives for trading. Experiments can bridge the gap between theory and empirics as the experimenter can control the information of each experimental subject and can explicitly observe the behavior of individuals.

type (with around 2000 trades spread over 6 treatments) we find scope for rational herding and contrarian behavior, even though prices are efficient. Moreover, comparing our findings with earlier studies, the nature of herding and contrarianism is more pronounced when traders have the opportunity to delay. Importantly, recipients of the signals that theoretically admit herding and contrarianism systematically delay more than those with high or low signals. There is also significant clustering of decision-making in time. Our findings highlight a key linkage between financial economics and informational herding theory: price momentum can be associated with rational herding, price mean-reversion can be associated with rational contrarianism, and both effects are bolstered by endogenous timing.

We will next go through the relevant theory and our pertinent findings in more detail, examining first the trading direction, then the timing decision and finally the issue of multiple trades. In what follows, we refer to the choice of trading direction as the static decision problem, and the timing decision as the dynamic problem, all the while emphasizing that these two belong together in the full equilibrium problem.

The Static Decision Problem. The 'static' theoretical model underlying our experiment is a sequential trading setup in the tradition of Glosten and Milgrom (1985), in which risk-neutral subjects trade single units of a financial security with a competitive market maker. Past trades and prices are public information, and the market maker adjusts the price after each transaction to include the new information revealed by this trade. We differ from Glosten and Milgrom by admitting endogenous timing.

In our specification, there is a single security with three possible liquidation values (high, middle, low). There are three possible actions (buy, sell, or no trade). Each subject receives a private realization of one of three possible signals (high, middle, or low); these signals are a piece of information about the true liquidation value. Recipients of the low signal systematically shift probability weight towards low values, recipients of the high signal systematically shift probability weight towards high values. Recipients of the middle signal value shift weight either towards the middle value or towards *both* the high *and* the low value. From the static perspective, rational subjects should buy if their expectation, conditional on their private signal and all public information, is above the price and sell if it is below the price. It turns out that, theoretically, recipients of the "low" signal should always sell and recipients of the "high" signal should always buy.

A key result of herding theory for financial market trading is that the existence of rational herding crucially depends upon the shape of the underlying conditional signal distribution for the middle value. Herders switch from selling to buying as prices rise or from buying to selling as prices fall. To do so, they must update their private expectations upwards faster than the price rises or downwards faster than the price falls. This faster speed of updating is possible if and only if traders have signals that make them shift weight towards *both* extreme outcomes (see Avery and Zemsky (1998) or Park and Sabourian (2008)), or, in technical terms, if their conditional signal distribution is "U shaped". Contrarians switch from selling to buying as prices fall and from buying to selling as prices rise. To do so, they must adjust their private expectations downwards slower than the price falls and vice versa when the price rises. This slower speed of updating is possible if and only if traders receive signals that cause them to shift weight towards the middle outcome, or, in technical terms, if their conditional signal distribution is "hill shaped".

As an example of when U shaped signals might be prevalent, take the situation that markets found themselves in on September 29, 2008, the day when the United States Congress rejected the first version of the TARP bill. At least three future scenarios were imaginable: the bill might be re-introduced as is, an arguably worse bill could go to the floor, or there will be no bill at all (causing a widespread collapse of the banking system). In this environment it is imaginable that investors were pulled between two opposing possibilities, either thinking that Congress was merely flexing its muscles but with every intention of eventually going with Treasury Secretary Paulson's recommendations or thinking that Congress would block any attempted bailout thus triggering the banking doomsday. Theory here predicts the potential for herd behavior.

Turning to our results, overall about 70% of the decisions are in line with the predictions of the static model. This fit compares well with other experimental work that focuses on trading with exogenous timing such as Drehmann et al (2005), Cipriani and Guarino (2005, 2007) and Park and Sgroi (2009). Broken down by signal type, these numbers are 87% and 83% for recipients of the low and high signals respectively. About 67% of recipients of signals that admit contrarianism act in accordance with the theory. For recipients of signals that admit herding, however, the fit is only 37%. One reason for this lower fit is that the decision with this signal is most difficult. Recipients of high and low signals act according to the static theory by simply sticking to their initial best action. Recipients of the U shaped (herding) and hill shaped (contrarian) signals, on the other hand, have to follow the development of prices very carefully. If they make small mistakes in their computations then their resulting decisions can violate the predictions of the static model.

We thus analyze whether the behavior of traders who have U and hill shaped signals is in the spirit of the predictions of herding theory. Namely, the implicit hypothesis in the static rational model is that traders who have U and hill shaped signals are more prone to engage in herd behavior or contrarianism respectively than recipients of high and low signals. This is what we observe: receiving a U shaped signal increases the probability of engaging in herding behavior by about 12%. Similarly, receiving the hill shaped signal increases the probability of engaging in contrarian behavior by 35% while receiving any other signal reduces the chance of acting as a herder or contrarian. The effect of the U shaped (herding) signal is notably stronger than that in Park and Sgroi (2009) who study a similar setup albeit with exogenous timing. They found that a herding signal increases the chance of acting as a herder by less than 6%. Our stronger effect highlights the importance of endogenous timing for herding.

The Dynamic Decision Problem. Smith (2000) presents a single-trade setup with a single informed trader who can decide whether to act early or late. Smith shows that someone who receives a signal that is either good or bad news (which we model with high or low signals) should trade 'early', and he also presents an example of a U shaped signal structure under which people optimally delay. Although this setup is not an exact match for the situation that we analyze, it provides important insights for us to draw upon. Namely, recipients of the high and low signals gain nothing from waiting as their optimal action does not depend on the history and they should thus act soonest.[6]

In our data, we find evidence that people with the high and low signals trade systematically earlier than people with the U shaped (herding) or hill shaped (contrarian) signal. Moreover, recipients of the hill shaped signal make most of their trades after the first few trades. In particular, they trade earlier than those who receive the U shaped signals. Thus contrarians act comparatively earlier than herders.

Our findings on timing complete the picture. As in any herding model, early movers trigger the herd behavior of the later movers. But we identify a systematic difference in who moves early and late: early movers are the ones who receive signals without the theoretically potential to herd or act as a contrarian, and latest movers are the ones who have the signals for which rational herding is possible. Thus recipients of herding signals wait to observe others, and when they have seen sufficient activity they join the majority. Traders who receive other signals act too early to be counted as joining the majority. Taken together, the behavior causes herding to be more pronounced.

We also find evidence of intertemporal clustering of trades in the sense of leader-follower patterns: a significant percentage of all trades (about 50-55%) occur within an interval of 1.5 seconds of their predecessors. Interestingly, private signals play no role in these patterns.

Multiple Trades. We might also ask what changes when traders are allowed to trade more than once? To address this we considered a set of treatments in which subjects are allowed to trade at most twice. With multiple trades the problem becomes extremely

[6]The seminal paper which studies investment timing with a single irreversible action, but without moving prices, is Chamley and Gale (1994). Similarly to Smith, their key message is that decision-makers will move very quickly, since waiting only makes sense when new information arises.

complex, for a trader's early transaction may affect the price and thus other traders' actions in such a way as to change his or her own optimal action later. For this reason there is no existing experimental work or theoretical paper close enough to our setup to provide clear predictions or comparisons. Therefore, our results under this treatment will be largely speculative and our aim will be to provide a set of experimental findings which will need to be addressed by future theoretical work.

We might conjecture that with two trades, subjects may attempt to "buy low and sell high" or "sell high and buy low". As this behavior may be governed by considerations other than the signals, such "round-trip" transactions may bias our estimate of the impact of signals. We show, however, that this impact is either insignificant or, if anything, leads to a slight underestimation of the impact of herding and contrarian signals.

Overview. The remainder of the paper is structured as follows. The next section discusses the related experimental literature. Section 3 provides a formal definition of herding and contrarianism. Section 4 examines the theoretical framework in more detail, discusses the modifications undertaken to better fit a laboratory experiment, and develops the hypotheses that are implied by the theory. 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 presents the summary results of the experiments and their fit to the rational model. Section 7 carries out a formal econometric analysis of the static decision. Section 8 analyzes the timing behavior. Section 9 studies the relation of the first and second trades analyzing the impact of "buy-low, sell-high"-trading. Section 10 summarizes the key findings and concludes. The supplementary appendix outlines results from an examination of several alternative behavioral explanations for the observations.[7] The appendix also includes a discussion of the role of information theory in the timing of decisions, as well as the subject instructions and other supporting materials, including the explicit parameter values.

2 Related Experiments

In what follows we discuss several related experiments that share some the features with ours. However, none of these papers incorporate both moving prices and endogenous timing. The first published paper to consider herding in endogenous-time was Sgroi (2003),

[7] Although we believe that our data indicates that people act in the spirit of the static rational model, we also examined several alternative hypotheses, including different assumptions on risk preferences, to see if we could further our understanding. While error adjustment models or probability over- or under-weighting schemes improve the fit of behavior to the data to some degree, they come at the cost of implicitly ruling out some interesting behavior (such as herding).

a close implementation of the Chamley and Gale (1994) fixed-capital investment framework. This framework was also examined experimentally in Ziegelmeyer, My, Vergnaud, and Willinger (2005). We complement this line of work that does not admit moving prices by explicitly allowing prices that adjust after actions.

There are several other papers that study herding behavior in a financial market environment: Drehmann, Oechssler, and Roider (2005), Cipriani and Guarino (2005, 2007), Alevy, Haigh, and List (2007), and Park and Sgroi (2008). The first two provide insights into trading behavior with exogenous timing in a setting where herding is theoretically impossible (as depicted in the underlying theory by Avery and Zemsky (1998)). Both papers confirm that herding does not commonly arise, thus showing that there is no 'natural tendency' of people to herd. Cipriani and Guarino (2009) and Park and Sgroi (2009) consider settings that do allow herding theoretically and both show that herding does arise, though not as often as predicted by the theory.

Alevy, Haigh, and List (2007) provide an innovative bridge between theory and empirical work by observing professional Chicago Board of Trade traders (who effectively generate the empirical data that one gathers from databases) in a controlled environment. They observe that the group of professional traders behaves quite differently than student control groups. For instance, professional traders are able to assess public information better than the student group.[8] This suggests that market data may be less susceptible to being contaminated by misleading herds. Our finding that contrarianism is more pronounced than herding reinforces their results. As Alevy, Haigh, and List (2007) employ a model with fixed prices, an interesting alternative would be to observe the behavior of professional traders in the environment depicted here where prices adjust after each trade.

The major difference between our setup and the five studies described above is that they all analyze settings with exogenous timing of trading decisions and all except Park and Sgroi (2009) do not allow for rational contrarianism. We build upon these studies by providing insights into the impact of endogenous timing on herding and contrarianism.

There are two related experimental papers that do analyze people's timing behavior in a financial market environment. Bloomfield, O'Hara, and Saar (2005) study a financial market in which people can trade repeatedly throughout a trading day. The focus of their study is on the timing behavior of informed traders and on their choice of limit- or market-orders depending on the passage of time, and they do not employ information that could (theoretically) trigger herding or contrarianism.

Ivanov, Levin, and Peck (2008) implement Levin and Peck (2008), which is a model of fixed capital (green-field), non-financial investments, and they develop important insights

[8]It is noteworthy, however, that the findings of Cipriani and Guarino (2009), who use professional traders, and Park and Sgroi (2009), who use students, differ only slightly.

into the timing behavior of people's investment choices. They find that behavior can be classified into three categories: self-contained (ignoring information from observing others), myopic (acting upon the current best decision, ignoring the option to learn from others in the future), and foresighted (the perfectly rational decision). Focusing on non-financial investments, the setting does not consider moving prices.

3 Definition of Herding and Contrarianism

The general movement of prices captures the majority or 'crowd' action: rising prices indicate that there are more buyers than sellers, falling prices indicate that there are more sellers than buyers. We define herding and contrarianism against this yardstick. Start with a trader who before any observable price changes has a trade (buy or sell) in mind. We say that this trader engages in herding behavior if he switches from selling to buying in the face of rising prices, or if he switches from buying to selling in the face of falling prices. The counterpart situation, contrarianism, arises when a trader switches from selling to buying in the face of falling prices or if he switches from buying to selling in the face of rising prices.[9] Thus, herding represents a history-dependent (social learning induced) switch of opinion *in the direction of the crowd*. Contrarianism describes a switch *against the direction of the crowd*.[10]

For the formal definition we use notation H_t for the trading history at time t; this history includes all past actions, their timing, and the transaction prices; H_1 is the initial history. Finally, we classify a person as a herder or contrarian at the time when the trade is made, that is, we do not speculate if the action may be driven by other motives such as price manipulation.

Definition 1 (Herd- and Contrarian-Behavior)

(*a*) *A trader engages in herd-buying after a history of trade H_t if and only if*

(H1) *he would sell at the initial history H_1,*

(H2) *he buys at history H_t, and*

(H3) *prices at H_t are higher than at H_1.*

Sell herding is defined analogously.

(*b*) *A trader engages in buy-contrarianism after a history of trade H_t if and only if*

(C1) *he would sell at the initial history H_1,*

[9] The definition of herding is in the spirit of Avery and Zemsky (1998) who were the first to provide a formal definition of herding in financial markets. The definition of contrarianism is in the natural counterpart of the definition of herding. A very loose intuition for herding is that herding types have increasing demand functions: they sell when prices are low and buy when they are high.

[10] Notice that this is a strict definition as it specifically excludes those who might have "herded" if they had the right signal but who happen to have a private signal that confirms the decision of the herd.

(C2) *he buys at history* H_t*, and*

(C3) *prices at* H_t *are lower than at* H_1*.*

Sell-contrarianism is defined analogously.

Herding here 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. If all traders were to act alike, then actions would be uninformative, so that an informational cascade arises and prices would not react to actions, staying "constant" for the duration of herding.[11] Instead, our definition of herding and contrarianism captures social learning, the key element of herding theory, and it also admits outcomes where prices move (often quite dramatically), which is an important real-world concern.

4 The Underlying Theory and Testable Predictions

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

4.1 The Static Decision of the Trading Direction

All traders trade a security with an uninformed market maker. The security takes one of three possible liquidation values, $V_1 < V_2 < V_3$, each equally likely. Traders can be informed, in which case they receive a conditionally independent signal about the true value of the security, or they can be noise traders in which case they trade for reasons outside the model. The market maker sets a single price at which he is willing to buy or sell one unit of the security.[12]

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

[11]For instance, models with informational cascades such as Cipriani and Guarino (2008) and Dasgupta and Prat (2008) have the feature that prices no longer move once the informational cascade starts.

[12]This is a simplification of a sequential trading model with three signals and three states in the tradition of Glosten and Milgrom (1985). In these models, a competitive market maker sets a zero-profit bid and offer price. In our experiments, we dispense with bid- and ask-prices and focus instead on a single trading price to minimize complexity. This is standard practice in the related experimental literature. E.g. Cipriani and Guarino (2005) report that that the spread has little to no effect.

the recipient of S_3 shifts weight towards the highest state (S_3 is 'good news'). Both of these signals have the feature that they are *monotonic*.

The signal S_2 that we employ here can take several different shapes. If it is hill shaped or single-polar, then it is generated with highest probability in the middle state V_2. The recipient of this signal will put most weight on the middle value. If signal S_2 is U shaped or bi-polar, then it generated with high probability in both the good state V_3 and the bad state V_1. The recipient of this signal puts most weight on both of the extreme outcomes. These types of signals are *non-monotonic*.

To determine whether someone herds or acts as a contrarian, we must determine which decision the person would take without observing trading activities of others. This benchmark decision is determined simply by the probability weight that the trader puts on the high relative to the low state. If the trader puts more weight on the high state, he would buy and we speak of a positively biased signal, if he puts more weight on the low state, the trader would sell and we speak of a negatively biased signal. We then refer to a negatively biased U shaped signal distribution as a negative U shape, and likewise for hill-shape and positive biases.

All past trades and prices are public information. We now have to distinguish between the settings in which traders can make one and two trades. For the single unit trade, for a given transaction price, once a trader decides to act, he buys if his expectation conditional on his private signal and on the information derived from past trades exceeds the price, and he sells otherwise. We will study a simple pricing rule for the market maker according to which the current price coincides with the public expectation conditional on all publicly available past information. This rule does not account for the possible timing decisions of traders but instead is purely backward looking.[13]

With two trades, matter are more complex. Take, for instance, a situation when a trader buys twice in quick succession. If his identify were known, then the price at which he buys the first and second unit should be the same because the second trade reveals no new information. Yet trading in our experiment is anonymous, and thus such a rule cannot be applied. We implement the price setting in the experiment by following the same price-changing rule as in the single-unit case, i.e. the market maker changes prices as if every trade has new information content. This ensures consistency from the perspective of the subjects.

Yet as a consequence, we cannot (and will not) make statements about whether or not behavior is "rational" in the two-trade setting. Nevertheless, what we can do is provide insights as to how people behave when they have more trading options available

[13] A backward looking market maker should theoretically make losses, but from an experimental perspective keeping matters simple for subjects if a greater concern.

for a fixed price-setting rule. Importantly, however, the problem of the right pricing rule combined with the rationality of the trading decision is intrinsic to *any* experimental study of endogenous timing.

The underlying decisions that we categorize for the single trade setups are assumed to be static, i.e. we ignore possible dynamic considerations that may govern a subject's trading decision. Moreover, we also assume that traders interpret every past decision as taken on the basis of only static considerations. Applied to the experimental design without bid-ask-spreads,[14] a monotonic type would never herd or act as contrarians, an S_2 type buy-herds (sell-herd) if and only if his signal distribution are negative (positive) U shaped, and an S_2 type acts as buy (sell) contrarian if and only if his signal distribution is negative (positive) hill shaped. In this spirit, we can refer to U shaped signals as *herding signals* and to hill shaped signals as *contrarian signals*. We shall also refer to signal S_1 as *bad news* and to S_3 as *good news*.

To understand the intuition behind the above statement, observe first, that a buy-herding trader would be selling at the initial history (H1/C1 from Definition 1). Since the prior is uniform, this implies that sellers attach more weight to the lowest than the highest state, i.e. $\mathsf{Pr}(S|V_1) > \mathsf{Pr}(S|V_3)$. Next, buy-herding also requires that prices have increased (H3 from Definition 1); this occurs if and only if the probability of the lowest state is smaller than that of the highest state.

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

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

In our experimental setup, we know the outcome of the random elements (noise trades and the signals for each subject). Conditional on all other subjects behaving optimally, we can calculate which action each subject should have undertaken given H_t and S, irrespective of the timing. Specifically,

Hypothesis 1 *When traders can trade once, S_1 types sell, and S_3 types buy. Moreover,*

[14] Note that 'S herds' is to be read as 'S herds with positive probability'. For a formal theory backing these assertions see Park and Sabourian (2008).

if S_2 types buy then their static expectation exceeds the price, and if they sell, then their static expectation is below the price.

Hypothesis 1 imposes strict requirements in that we posit a precise mapping from expected payoffs into decisions. Specifically, recipients of non-monotonic signals are expect to be able to compute their expectations precisely, and small errors can cause a trade to be classified as wrong. We will thus additionally consider a more generous interpretation. We expect the U shaped types to be the source of herding behavior and the hill shaped types to be the source of contrarian behavior.

Hypothesis 2 *If we observe herding then this is most likely to be caused by U shaped types. If we observe contrarianism, then this is most likely caused by hill shaped types.*

4.2 The Dynamic Decision of the Trading Time

Most work on financial market microstructure constructs the trading decision to be either stationary or static.[15] Those that do analyze a dynamic problem usually allow at most one informed insider.[16] Yet the essence of traders rationally herding or acting as contrarians is that (a) behavior is not stationary and (b) behavior is driven by social learning, i.e. by observing the actions of other, possibly informed agents.

The theoretical paper closest to our experimental setup is Smith (2000). Although he models a single trader who can make a single trade at one of two points in time (early or late), his results intuitively extend to the case of multiple traders.[17] Smith shows that a trader with a good or bad news signal will trade early. Moreover, Smith also presents an example with a U shaped signal and shows that within his framework this trader would rationally delay.[18] Applying Smith to our framework, the monotonic S_1 and S_3 types have bad-news and good-news signals respectively and thus they should always act immediately at the beginning of a trading session.[19] Thus we have

[15] Hasbrouck (2007) and Vives (2008) survey the recent theoretical market microstructure literature.

[16] See, for instance, Kyle (1985), Back and Baruch (2004, 2007), or Chakraborty and Yilmaz (2004). Also, see Bloomfield, O'Hara, and Saar (2005) for a discussion of trade-timing models.

[17] In Smith (2000) the trader obtains a piece of information about a public signal that will be released. After the release of this public signal, prices will adjust instantaneously to the fundamental value implied by the signal. To see the equivalence to our setting assume that people trade as described by the static problem. Then their actions will affect the price and (noisily) reveal traders' information. Thus the price at the end of the trading day is, loosely, a sufficient statistic for all traders' private information. Moreover, the price is public information. Thus the price at the end of the trading day is the same as the price in Smith (2000) after the release of the public signal. A trader's information in our model can thus be understood as a signal about the information that is to be revealed through trading.

[18] Smith's focus is to derive sufficient conditions for early trading, not to provide conditions for delay. So Smith's work, while insightful in the early trading dimension, cannot answer whether U or hill shaped types would generally trade early or late.

[19] An alternative approach to Smith's finding is that the expectations that the S_1 and S_3 types form

Hypothesis 3 *In the single trade case, monotonic types trade once the treatment begins.*

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

While the theory predicts that monotonic types should act immediately at the beginning of trading, we would expect that their actions are, to some extent, spaced out. In analyzing their timing behavior we will thus look at the *distribution* of trading times and we are thus interested in the relative ordering of the distributions of trading times for the different signal types. This yields the following three hypotheses.

Hypothesis 4 *The U shaped types will act later than the monotonic types.*

Hypothesis 5 *Hill shaped types will act before U shaped types.*

Hypothesis 6 *Hill shaped types act after the first few trades have occurred.*

The statements so far cover the case of a single trade, and, as outlined in the introduction, there is no theoretical work in the literature that provides guidance as to how people would optimally act when there are multiple heterogeneously informed traders who can trade repeatedly. Despite the lack of a theory to guide our analysis, we can conjecture possible behaviors. With two trades, people can either trade twice in the same direction or they can trade in opposite directions (sell-buy or buy-sell); the latter is referred to as a 'return'-trade (or 'round-trip' transaction). With return trades possible, we must revisit the decision about the trading direction.

Now consider decision of the bad-news (S_1) types. Suppose that they perform a return-trade, for instance, by buying early and selling late. Then after these two transactions, they still hold a share, and on this single share they make an expected loss. While they

over the future price are super- and sub-martingales respectively. For instance, an S_3 type would always expect that there are more people who have the same signal and thus expects prices to rise. For their own expectation, however, the law of iterated expectations applies.

may gain on the return transaction, it seems most intuitive that the bad news types should sell twice. A similar argument applies to the good news (S_3) types: if they sell early and buy late (or vice-versa) they would hold only one share while they could hold three. Thus

Hypothesis 7 *When traders can trade twice, S_1 types sell twice, S_3 types buy twice.*

Hypothesis 8 *In the two trade case, monotonic types trade at the start of the treatment.*

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

Hypothesis 9 *The first trade of U shaped types occurs later than the first trade of hill shaped and monotonic types.*

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

As prices move, traders' information rents are reduced in expectation, at least for the S_1 and S_3 types (by the same arguments as described in Smith). Thus the greater the number of trades, the more intense is the competition for information rents. A very straightforward assertion is thus that trades should occur earlier when people can trade more often. When two trades are allowed, as subjects know that people can trade twice, any delay motive that they might possess will be diminished. Consequently, we expect

Hypothesis 10 *First trades in the two trade cases are earlier than in the single trade cases.*

One conceptual difficulty that arises in the experimental implementation is the manner in which the price is updated. In principle, this should be dictated by a theory that determines who trades when and in what direction. For lack of such a theory we used the reasonable updating rule that would account for the statically optimal actions. For instance, in a setting with a negative U shaped signal S_2 and absent herding, a buy would have been assumed to come from either a noise trader or an informed trader with good news (S_3). Likewise, in a setting with a positive U shaped signal S_2 and absent herding, a buy would have been assumed to come from either a noise trader, an informed trader with good news (S_3), or and informed trader with the positive U shaped signal (S_2). As we argue below, this is ex post justified as the behavior is largely in line with the presumed statically optimal behavior. Moreover, whether a price change is based on the assertion that an S_1 type sold relative to a situation where both the S_1 and S_2 types sold is numerically so small that it does not affect the statically optimal action of traders in the situations that we study. In other words, although the observed timing behavior may render the pricing rule ex post imprecise, the numerical imprecision is too small to negatively affect our analysis, at least for the single trade case.

Furthermore, the details of the pricing rule and the relation to the theoretical behavior matters only for our analysis in the single trade case where we do assess traders' rationality. Moreover, for the general analysis of herding and contrarian behavior we note that their definitions apply to any pricing rule and do not depend on a specific, theoretically correct pricing rule. It is thus sufficient for most of our analysis to study a mechanical market maker, to treat subjects as if they know the pricing rule, and to assess their behavior relative to the pricing rule.

5 Experimental Design

Here we discuss the experimental design, focusing on the information provided to the subjects, the differences between treatments, and our predictions. The supplementary appendix contains further information including a time-line (Appendix C), a full set of instructions and the material given to subjects (Appendices D-F), as well as a description of the custom software used in this experiment (Appendix G).

5.1 Overview

The design focuses 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 group of traders was made up of 13-25 experimental subjects plus a further 25% of computerized noise traders, with a central computer acting as the market maker. Subjects were informed that they would not interact with each other directly but rather that the actions of all of the traders would effect the current price. They were told that this price is set by the central computer that is operating at the front of the experimental laboratory and that a decision to purchase by a trader raises the price while a decision to sell lowers it.

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

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

The existence and proportion of noise traders was made known to the experimental subjects in advance, who were also aware that noise traders randomized 50:50 between buying and selling and that they trade at random times.[20]

The subjects were informed that the sessions would last 3 minutes and that they would receive announcements about the remaining time after 2:30 minutes, and 2:50 minutes. We considered two classes of treatments: in the first people were allowed to trade once, in the second they could trade twice. The software allowed subjects to trade at most this specific number of times. The sequence of transactions produced a history of actions and prices, H_t with $t \in (0, 180)$, that recorded the timing (in seconds), price, and direction of each transaction. Subjects were shown the history in the form of a continuously updating price chart during each treatment, and they were also given the current price, P_t. This price was calculated by the computer as $P_t = \mathsf{E}[V|H_t]$ with $P_1 = 100$.

Subjects were told that they had three possible actions $a = \{\text{sell, pass, buy}\}$ one (or two) of which they could undertake during the 3 minutes of trading time. For the treatments in which two trades were allowed, subjects were additionally informed that they could trade twice, so they could "buy and buy", "sell and sell", "sell and buy", etc. as and when they wished during the three minute period. They were instructed that

[20]Noise traders play an important role in the static theory and add only a mild degree of extra complexity to the experimental design. They also play a useful practical role in the experiment by simulating a degree of uncertainty about the usefulness of any observed actions. Generally, noise traders reduce the informativeness of any observed action and we will see later that this allows us to model errors as an incorrect assessment of the degree of noise trading (see the Supplementary Appendix). Moreover, they also reduce the ability of subjects to accurately predict when other informed traders may have acted. For instance, in a room with 25 subjects, and no noise traders, observing 24 actions immediately tells the final subject that there are no more informative actions to come, and so no further reason to delay.

pressing the "pass"-button would count as one of the actions that they were allowed. It was stressed to the subjects that their virtual profits per treatment were generated based on the difference between the price at which they traded, P_t, and the true value of the share, V. It was emphasized that the price at the end of the trading session would not be relevant for their payoffs.

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 (13 subjects each), 6 at the University of Warwick (18, 19, 22, 22, 22, and 25 subjects) and 4 at the University of Toronto (17, 18, 13, and 13 subjects). 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 Mathematics. 53% claimed to have some prior experience of financial markets, with some 23% owning shares at some point.[21]

5.2 Treatments

In the first three treatments, subjects were allowed to trade at most once, in the last three treatments subjects were allowed to trade at most twice. The treatments were each designed to enable us to examine behavior under a specific information structure:[22]

Treatment 1: negative U shaped signal structure making buy herding possible;
Treatment 2: negative hill-shape making buy-contrarianism possible;
Treatment 3: positive U shaped signal structure making sell-herding possible;
Treatment 4: as Treatment 2 but with two trades;
Treatment 5: as Treatment 3 but with two trades;
Treatment 6: as Treatment 1 but with two trades.

The underlying parameters are listed in the supplementary appendix together with the instructions given to subjects.

5.3 Behavioral, non-rational predictions for the static decision

To complete the analysis, we considered the possible impact of risk aversion and loss aversion on decision making, and various behavioral alternatives to Bayesian updating.

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

[22] Since in Drehmann, Oechssler, and Roider (2005) the inclusion of transaction costs produced the expected outcomes, we ignored transactions costs and instead focused on the information structure as the key differentiating factor between treatments.

First, we considered a model in which subjects do not update their beliefs as prices change but act solely on the basis of their prior expectation. Second, we considered one setting in which subjects update their beliefs on the basis of changing prices at a slower rate than they should and one setting in which people overweigh their own private information. Finally, we developed error correction models in which subjects account for errors made by their peers and react rationally to these errors; these models are in the spirit of level-k beliefs (see Costa-Gomes, Crawford, and Broseta (2001)) and the Quantal Response Equilibrium (see McKelvey and Palfrey (1995) and McKelvey and Palfrey (1998)).

Since none of the alternative specifications provided significant additional insights over and above the standard model of Bayesian rationality and since risk and loss aversion did not provide an improved fit, we delegated the details of the specifications and testing to Section A of the supplementary appendix.

6 Analysis of the Static Trade-Direction Decision

We will first examine the results by summary statistics and then expand on them with a formal econometric analysis in Section 7.

In the numbers to follow we exclude noise trades, and focus only on trades by human subjects. The total number of trades was 1991 spread over all 6 treatments; broken up by trader type we have 623 (S_1), 786 (S_2) and 584 (S_3). For treatments 1 to 3 we had 683 trades (197 S_1, 276 S_2 and 210 S_3), for treatments 4-6 there were 1308 (425 S_1, 510 S_2 and 373 S_3) trades.

6.1 Overview of the fit of the data to the static model

Examining Hypotheses 1 and 2, let us turn to the findings detailed in Table 1. We note that the number of trades that conformed to the theoretical static predictions was 70%. 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. Park and Sgroi (2009) employed a similar trading framework but with trading in an exogenous sequence. The fit there was also 70%. The similarity to the results in the literature is noteworthy because the endogenous timing setting in our experiment is much more complex. Moreover, the experiment in Cipriani and Guarino (2005) effectively considers only types that are equivalent to our monotonic types. With rationality in excess of 80%, these types actually performed better in our setting than those in Cipriani and Guarino. Therefore, we find it reasonable to argue that the static theory is a good description for the behavior of the monotonic types.

The non-monotonic S_2 types, however, often do not act according to the static theory — almost half of their trades were against the static model. In particular in the buy-herding Treatment 1, the S_2 do deviate from the static predictions, even when admitting passes as weak buys. Had they taken each action at random they would have been more likely to trade according to the static theory.

At the same time, the non-monotonic types face a more difficult decision problem than the monotonic types. Theoretically, the decisions of the monotonic types never change, so they can take the correct decision even without following the history carefully. The non-monotonic types on the other hand, have to follow the history of prices carefully and small mis-computations can cause them to be categorized as not fitting the static model.

A further complication that we discuss in Section 8 is that trades are often clustered. If subjects observe a cluster, they may deduce that many trades in the cluster are by other, informed subjects and not by noise traders. As a consequence, they may (rightfully) consider the pricing mechanism to be flawed in this cluster and react to it. Finally, the benchmark against which we measure the decisions assumes that all traders act rationally and it assumes that the subjects themselves trust their peers to act rationally. If subjects suspect that this assumption is not warranted, then they may rationally choose an action that is against the static model.

We explore these possibilities in more detail in the supplementary appendix. More generally, however, we are interested in the decisions that people make in relation to their signals once prices move. Specifically, we are interested whether specific signals are more likely to trigger a change in behavior (contrarianism or herding) than others. In the formal econometric analysis of the next section we will see that such an effect is indeed prevalent in subjects' behavior as a function of their signals.

Finding 1 (Hypothesis 1) *About 85% of monotonic types, 67% of hill shaped types, but only 37% of U shaped types act in accordance with the static theory.*

6.2 Herding and Contrarianism

Based on the definitions from Section 3 of herding and contrarianism we now describe how often such behavior is observed. Specifically, we ask whether the herding candidate S_2-types switch from selling to buying if prices rise and whether they switch from buying to selling if prices fall. While according to the static theory only S_2 types can rationally herd, herding and contrarianism according to Definition 1 can be observed for all types.

Table 2 gives the raw numbers. Herding arises in about 16% of the cases when it is possible, contrarian behavior arises in 27% of the possible cases. Non-monotonic types have the highest propensity to herd (24%). Similarly, non-monotonic types have the highest propensity to act as contrarians (43%), but S_1 and positive U shaped types also show a tendency to act as contrarians (31% and 37%). Our formal regression analysis will later show that as predicted the hill shaped signal S_2 is the major cause for contrarianism.

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

Finding 2 (Hypothesis 2, based on summary statistics) *Non- monotonic types are more likely to cause herding than non-monotonic types. Most instances of contrarianism stem from non-monotonic types, but when they can, S_1 and positive U shaped types tend to act as contrarians. (31% and 37% respectively).*

7 Regression Analysis of the Static Decision

The summary statistics from the last section gave a good idea of the determinants of behavior: first, recipients of U shaped signals are more likely to herd than recipients of the extreme signals. Second, contrarian behavior occurs more frequently and is prevalent both among the S_1 and the S_2 types.

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

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

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

The random assignment of signals to traders allows us to interpret mean differences in signal-specific effects as the average causal effect of the signal. It is important to note that the only exogenous variation for traders is in their signal. The timing of their action and also the price at which a trade occurs are endogenous and we will thus not use such measures in our formal regression analysis.

Formally, we estimate the following equation to test the hypothesis that a type of signal is a significant cause for herding or contrarian behavior respectively:

$$\mathsf{herd}_i = \alpha + \beta \mathsf{u\text{-}shape}_i + \mathsf{fixed}_i + \epsilon_i, \quad \mathsf{contra}_i = \alpha + \beta \mathsf{hill\text{-}shape}_i + \mathsf{fixed}_i + \epsilon_i \qquad (1)$$

where the dependent variables herd_i and contra_i are dummies that apply Definition 1 in the sense that they are set equal to 1 if individual i herds or acts as a contrarian respectively at a given trade and 0 otherwise, α is a constant, and $\mathsf{u\text{-}shape}_i$ and $\mathsf{hill\text{-}shape}_i$ are signal dummies that are set equal to 1 if the individual received a U shaped (for the herding estimation) or hill shaped (for the contrarian estimation) signal. Parameter fixed_i is an individual fixed effect that controls for specific traders who persistently err.[23] Given the random assignment of signals, we can assume that $\mathsf{E}[\mathsf{u\text{-}shape}_i \cdot \epsilon_i] = 0$ and $\mathsf{E}[\mathsf{hill\text{-}shape}_i \cdot \epsilon_i] = 0$, which are the main identifying assumptions.

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

We ran these regressions on a total of 5 different subsets of the data: for all treatments, for treatments 1-3 (one trade), for treatments 4-6 (two trades), for the first trade in treatments 4-6, and for the second trade in treatments 4-6.

In each scenario we estimated the model by logit and linear probability regressions without fixed effects (i.e. γ_i is omitted from (1)) and then ran the linear probability regression controlling for trader fixed effects. All regressions included a constant which is significant at all conventional levels but omitted from the results tables. As a general convention we report standard errors in parentheses below the estimates and a $*$ indicates significance at the 5% level, $**$ signifies significance at the 1% level.

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

Table 3 summarizes the result from our regression. Overall, obtaining a U shaped signal increases the probability of herding by about 12% relative to any other signal and it is significant with and without controlling for trader fixed effects; moreover, the coefficients from the logit and the linear regression as well as he linear fixed effects regressions are similar for all subsamples studied. The estimates are significantly different from zero at conventional levels except for some of the the fixed effects regressions, where the standard error increases. Including trader-fixed effects decreases the linear coefficients only slightly; this indicates that the estimates are not driven by some specific traders who persistently take no-herding actions.

Overall the regression confirms the hypothesis that recipients of U shaped herding-type signals are generally more likely to herd. The effect of the signal is also double that

[23] In unreported regressions we also controlled for treatment, session and treatment-session fixed effects as well as random effects. The results remain unaffected.

found in related work in the literature, namely in Park and Sgroi (2009), where we find a marginal effect for the herding signal of about 6% in an exogenous timing setting.

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

Table 4 summarizes the results from our regression. Receiving the hill shaped S_2 signal increases the chance of acting as a contrarian by about 36% relative to any other kind of signal. As with herding, the linear probability model's coefficient changes only slightly when we include trader fixed effects. All coefficients are significant at the 1% level except when we focus on treatments 1-3 and include fixed effects, where it is significant at the 5% level due to the large standard error.

Overall we confirm that the hill shaped signal is the significant source of contrarianism relative to all other signals. The marginal effect is comparable to that found in the literature, namely in Park and Sgroi (2009), where it is estimated to be about 34%.

Finally, we also ran the regressions in this section separating subjects by their locations (Cambridge, Toronto, and Warwick). While the estimates vary slightly, we found no substantial differences in the results and thus omit the corresponding tables.

Finding 3 (Hypothesis 2, based on formal regression analysis) *The U shaped and hill shaped signals are the significant sources for herding and contrarianism respectively.*

Our findings here are noteworthy for three reasons. First, in the last section it appeared as if an S_1 signal could also be a significant source of contrarianism. As we show here, this casual observation missed the subtlety of a signal's impact. Second, (irrational) contrarianism has been observed in other experiments before (e.g. Drehmann, Oechssler, and Roider (2005), Cipriani and Guarino (2005), Alevy, Haigh, and List (2007)). Thus arguably, people exhibit a tendency to act against the crowd. Here we show that despite this tendency, contrarianism is still most likely caused by recipients of signals that admit contrarianism theoretically.

Third, the marginal effect of a U shaped signal on the probability of herding is much stronger than that in our exogenous timing paper, Park and Sgroi (2009). It is not as strong as the effect of the hill shaped signal on the probability of contrarianism. But this is expected, given the aforementioned natural tendency to of people to act as contrarians. Combined with the fact that U shaped types do not herd as much as they should theoretically, this implies that *due to the timing of actions* (a) non-U shaped types herd proportionately much less than with exogenous timing and (b) types with herding signals are more strongly singled out as herders.

8 Analysis of Timing

The strongest interpretation of the theory is that monotonic types should trade immediately when the session starts. Consequently, according to this view we should observe that all these types trade within the first few seconds of the game. And indeed, we do observe a very large number of trades at the very beginning of trading. Table 5 displays how many trades we observe in the first five seconds of trading.

Overall, 355 out of the 1206 trades made by monotonic types occur within the first five seconds. Although this number is substantial, it is less than 30% of the trades. These numbers differ only slightly when they are broken up S_1 and S_3 types or by the single trade and the two trade cases. The exception is the first trade in the two-trade cases: of these, about 49% of the first trades occur within in first five seconds. However, implicit to Hypothesis 8 is that both trades occur at the beginning of trading. For the second trade we find, however, that only a small fraction of these occurs in the first five seconds.[24]

Finding 4 (Hypotheses 3) *Monotonic types perform about 28% of the their trades within the first five seconds. Thus we cannot support the hypothesis that in the one trade case, all monotonic types trade immediately.*

Finding 5 (Hypotheses 8) *Monotonic types perform about 30% of all their trades within the first five seconds. They perform 49% of their first trades and only 11% of their second trades within this time span. Thus we cannot support the hypothesis that in the two trades case, all monotonic types trade immediately.*

We now attempt to identify systematic differences in the timing behavior for the various signal types and treatment settings. Specifically, we will compare the cumulative distributions of the trade-times for different categories of types. The strongest result that one can hope for in this context is that one cumulative distribution function (henceforth, cdf) of trade-times stochastically dominates another: distribution F first order stochastically dominates distribution G if G is larger than F for all entry times. If we indeed observe that F first order stochastically dominates G, then we can say that the entry times under F are systematically later than under G.

We computed the cdfs for a large variety of subsamples, such as treatments 1-3, 4-5, 1-3 and 4-5 (first trades), and so on. Since the insights from the different subsamples differ only slightly, we present the results for case where (1) we combine all treatments, (2) S_1 and S_3 are combined as *monotonic* signals (their individual cdfs show no pattern)

[24]The numbers increase to only 35 (17%) and 28 (15%) trades when we expand the range to the first 10 seconds.

and (3) positive and negative U shape are combined as *U shape*. Figure 1 provides plots of the relevant differences of cdfs.[25]

First, Hypothesis 5 (and also 9) state that monotonic types trade before U shaped types (bottom left panel). This can be extended to say that monotonic types trade before non-monotonic types (top left panel). We observe that

Finding 6 (Hypothesis 4) *U shaped types trade later than monotonic types.*

Consequently, our findings comply with Smith (2000)'s prediction that people with good-news or bad-news-signals trade early, and that people who receive mixed information delay. The bottom right panel displays the relation of the hill shaped types' timing to the U shaped types. We observe

Finding 7 (Hypothesis 5) *On almost the entire domain (apart from the first few seconds) the hill shaped type trades systematically earlier than the U shaped types.*

To consider Hypothesis 6, one can compare the top right and bottom right panels. As can be seen, for the first few seconds, more trades stem from non-hill shaped types. Yet after these first few trades, the hill shaped types trade strongly (this follows as their cdf rises strongly relative to the other cdfs).

Finding 8 (Hypothesis 6) *The trades by hill shaped types are concentrated after the first few transactions have occurred.*

Hypothesis 10 states that people trade faster when they can trade more often. The intuition here is that a large number of trades dissipates information rents and thus trading early is most valuable. We aggregated all trades in treatments 1-3 and 4-6 as well as first and second trades in treatments 4-6. The panels in Figure 2 plot the corresponding cdfs and paint a clear picture.

Finding 9 (Hypothesis 10) *Allowing people to trade more often speeds up their trade-times in that the first trade in treatments 4-6 occurs earlier than the single transaction in treatments 1-3.*

Finally, as outlined in Section 3, we find no particular ordering for the timing of the second trade, so we omit the graph. This non-finding suggests that settings in which people can trade more than once may not yield clear insights with respect to people's timing behavior. This implicitly supports the choice of our stylized setup because because

[25] Tests of stochastic dominance have low power. The plots of cdfs that we show here, however, paint a very clear picture in that for almost the entire domain we observe a clear ordering of the distributions of trading times.

in a more complicated setup with a larger choice set, it would be much more difficult to track the impact of signals on traders' decisions.

The supplementary appendix discusses an alternative analysis, examining the scope for pure information theory to explain the timing of decisions. As is common in the literature on information theory we use the entropy of posteriors to measure the informativeness of a signal. We observe, however, that information theory does not seem sufficient to explain the timing of decisions, and that instead contrarianism and herding motives play an important role.

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

To get a sense of systematic behavior, we first determined the number of seconds between one trade and its predecessor. We then looked at the frequencies for which trades occur within 1.5 seconds of one another, split up by signal types. Table 6 lists the relevant percentages, Figure 3 plots the frequencies for all signal types taken together.

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

Note that there is at least one additional reason to believe that clustering is rational in the context of endogenous-timing. Since noise traders trade at random times spread uniformly across the treatment, observing several trades in close temporal proximity suggests that they are being made by a disproportionate number of informed traders. Subjects might reason that the market maker, who updates prices without taking timing into account, will fail to account for this behavior so that the price is updated too slowly during a cluster. This creates an opportunity for gains if a trade is made soon after a cluster is first observed, and produces a subtle feedback loop which might work to increase the scale of clustering behavior.

9 The Relation Between the First and Second Trade

Subjects have the opportunity to make so-called 'return' (or 'round-trip') trades by selling first and then buying later or vice versa. This way, they can realize a trading profit in the process. At the same time, if they make such a 'return' trade then by the end of the treatment they still own one unit of the security. They thus forego a trading profit that they could make by selling this unit or by not buying an additional unit.

Return trades can arise only in treatments 4-6 and thus this section deals only with this subset of the data. Table 7 provides summary statistics for the second trade with emphasis on the return trades.

Not all traders act twice — in about 6% of cases they forego the second trading opportunity. About 77% of all second trades go in the same direction as the first. Hypothesis 7 asserts that monotonic types act in the same manner irrespective of the number of trades. Here we find

Finding 10 (Hypothesis 7) *66% of S_1 types sell twice,* 68% *of S_3 types buys twice. This fraction is lower than sell/buy fraction for the S_1/S_3 in the single trade case, but remains supportive that signals are paramount in driving behavior.*

About 23% of second trades are part of a round trip transaction. Most of these are performed by the non-monotonic types. These types also account for the largest fraction of all type-specific return trades. About 76% of the return trades yielded a trading profit which suggests that return-trades were performed on the basis of "buy low, sell high" (or "sell high, buy low").

This raises the question of whether there are systematic features of the return trades.[26] Return trades are always either buy-sell or sell-buy. So we ask if is there a particular fashion in which the different signal types perform their return trades.

There is a specific interpretation that can be attached to return trades: for the S_1 types a buy-sell return trade has a manipulative connotation in that this type may try to drive up the price in order to sell high later; similarly for negative hill and negative U shaped types and for sell-buy return trades by the S_3 or positive U shaped types. Yet, combining the negative and positive signals respectively, we observe this tendency only in 40% and 31% of cases.

[26] In some unreported regressions we analyzed whether the payoffs from return trades are smaller when they occur late. For this we regressed factual payoffs from return trades and also expected payoffs from return trades on the trading time. Yet we found no significant relation and thus concluded that time has no impact on the payoff of return trades. Moreover, we could not find a general relationship between return trades and time for the S_2 and S_3 types either. However, for the S_1 types we found a relationship that is significant at the 1%-level. This finding is intuitive in the context of a different finding (which we report below) in that the S_1 types tend to act as herders when they perform a return trade. This implies that the S_1 types first sold, then waited and observed that prices were increasing. They then bought back the share that they sold.

The next question is whether the return trades affects our estimates of herding and contrarianism. From the above analysis it is conceivable that traders who engage in round-trip transactions have different motives for trading than those who do not engage in such behavior. A trader who merely aims to buy low and sell high may thus act for reasons that have little to do with his information. Yet in our analysis, such a trade may be classified as herding or contrarianism and we would thus obtain spurious estimates.

We want to analyze to what extent our estimates in Tables 3 and 4 change when we take account of this possible misclassification. We thus ask the following question: what is the probability that a first/second trade is a herding trade conditional on the trade being a return trade (by type S when herding is possible) relative to the case where it is not a return trade? To answer this question, we ran the following regressions

$$\mathsf{herd}_i = \alpha + \beta_1 \mathsf{U\ shape}_i + \beta_2 \mathsf{return\ trade}_i + \beta_3 \mathsf{U\ shape}_i \times \mathsf{return\ trade}_i + \epsilon_i, \quad (2)$$

$$\mathsf{contra}_i = \alpha + \beta_1 \mathsf{hill\ shape}_i + \beta_2 \mathsf{return\ trade}_i + \beta_3 \mathsf{hill\ shape}_i \times \mathsf{return\ trade}_i + \epsilon_i. \quad (3)$$

The dependent variables herd_i and contra_i are the herding and contrarian dummies from the equations in (1), $\mathsf{U\ shape}_i$ and $\mathsf{hill\ shape}_i$ are the signal dummies, α is a constant, $\mathsf{return\ trade}_i$ is a dummy for the incidence of a return trade (both the first and second transaction of a return trade have value 1), and $\mathsf{U\ shape}_i \times \mathsf{return\ trade}_i$ and $\mathsf{hill\ shape}_i \times \mathsf{return\ trade}_i$ are products of the two dummies.

For each case we estimated the model by logit, restricted to incidences where herding and contrarianism respectively can occur; we report the marginal effects. The coefficient β_1 allows us to estimate the marginal effect among non-return traders and the coefficient β_3 allows us to estimate the differential marginal effect among return traders, so that $\beta_1 + \beta_3$ allows us to determine the effect of a signal among return traders.

Table 8 summarizes our findings and indicates that our herding estimates from Section 7 are biased *downwards* by return trades (the coefficients on the product term are negative and significant) and that our contrarian estimates are unaffected. This is good news for our analysis thus far as it indicates that, if anything, the effect of a herding signal as a source for herd behavior is underestimated by the possibility of round trip transactions.

Finding 11 (Impact of Return Trades on Estimates for Hypothesis 2) *The estimates underlying Finding 3 for herding become stronger and those for contrarianism remain unaffected when we correct for misclassified round trip transactions.*

10 Conclusion

Herding has long been suspected to play a role in financial market booms and busts. Recent theoretical work shows that informational herding (or contrarianism) is possible in a market with efficient asset prices if the conditional signal distribution for traders has a specific shape. Other work shows that when timing is endogenous to the decision, traders with good or bad news should trade earlier than those with less informative signals. In the context of financial market herding, giving traders a choice of when to act is not only natural, but there are also important insights that can be gleaned from such an analysis. For instance, if herding-prone types delay their actions systematically, then herd behavior can become more pronounced and significant compared to exogenous timing settings.

We collected almost 2000 trades in the first experimental test which allows traders to chose both the trade direction and time. By contrast with almost all existing experiments, we employed a theoretical framework that allowed rational herding and contrarianism in a financial market environment with moving prices. We found that subjects decisions were generally in line with the predictions of static informational learning theory when that theory admits rational herding and contrarianism. For example, types theoretically prone to herd or be contrarian are the significant and important source of this kind of behavior when it does arise. Types with extreme information (good or bad) trade systematically earliest, and those with signals conducive to contrarianism trade earlier than those with information conducive to herding. We thus find strong evidence for the impact of the type of information both with respect to the direction and the timing of trades.

The work presented here has five key messages. First, we find additional and qualitatively novel support for (sequential) rational herding theory in the laboratory. Second, adding endogenous-timing leaves the key predictions of sequential herding theory unchanged as concerns the direction of trade. Therefore, our results suggest that earlier work which forces subjects to act in a strict sequence remains valid even though the timing assumptions are unrealistic. Third, extant laboratory experiments with sequential trading are not designed to capture timing decisions. We thus provide the first experimental test of Smith (2000), confirming his theoretical predictions. Fourth, we join two literatures by linking information-based trade directions and timing. Signals that push subjects towards herd or contrarian behavior also push them towards delay, relative to the sorts of signals which guide subjects towards clear buy or sell decisions. This point is a potentially very important avenue for future research as the combination of herding/contrarianism in decision-making and clustering in time can work together to potentially exacerbate/counter prices movements which drift away from fundamentals. Finally, we also provide an exploration of the impact of allowing traders to trade more than once.

Our findings are more speculative in the two-trade treatments, though we find evidence that if anything the ability to trade more than once may lead to even more pronounced rational herding and contrarian behavior. Following on from this last point we believe that one of the most profitable areas for future work is in the theoretical examination of endogenous-time financial herding with multiple trades. Such a model would bring theoretical and experimental findings into closer alignment with real-world markets.

References

Alevy, Jonathan E., Michael S. Haigh, and John A. List, 2007, Information cascades: Evidence from a field experiment with financial market professionals, *Journal of Finance* LXII, 151–180.

Anderson, Lisa R., and Charles A. Holt, 1997, Information cascades in the laboratory, *American Economic Review* 87, 847–862.

Avery, Christoper, and Peter Zemsky, 1998, Multi-dimensional uncertainty and herd behavior in financial markets, *American Economic Review* 88, 724–748.

Back, Kerry, and Shmuel Baruch, 2004, Information in securities markets: Kyle meets Glosten and Milgrom, *Econometrica* 72, 433–466.

———, 2007, Working orders in limit order markets and floor exchanges, *Journal of Finance* 62, 1589–1621.

Banerjee, Abhijit V., 1992, A simple model of herd behavior, *Quarterly Journal of Economics* 107, 797–817.

Bikhchandani, Sushil, David Hirshleifer, and Ivo Welch, 1992, Theory of fads, fashion, custom, and structural change as informational cascades, *Journal of Political Economy* 100, 992–1026.

Bloomfield, Robert, Maureen O'Hara, and Gideon Saar, 2005, The 'make or take' decision in an electronic market: Evidence on the evolution of liquidity, *Journal of Financial Economics* 75, 165–199.

Chakraborty, Archishman, and Bilge Yilmaz, 2004, Informed manipulation, *Journal of Economic Theory* 114, 132–152.

Chamley, Christophe, and Douglas Gale, 1994, Information revelation and strategic delay in a model of investment, *Econometrica* 62, 1065–1085.

Chordia, Tarun, Richard Roll, and Avanidhar Subrahmanyam, 2002, Order imbalance, liquidity, and market returns, *Journal of Financial Economics* 65, 111–130.

Cipriani, Marco, and Antonio Guarino, 2005, Herd behavior in a laboratory financial market, *American Economic Review* 95, 1427–1443.

———, 2008, Herd behavior and contagion in financial markets, *The B.E. Journal of Theoretical Economics (Contributions)* 8, Article 24.

———, 2009, Herd behavior in financial markets: A field experiment with financial market professionals, *Journal of the European Economic Association* 7, 206–233.

Costa-Gomes, Miguel, Vincent P. Crawford, and Bruno Broseta, 2001, Cognition and behavior in normal-form games: An experimental study, *Econometrica* 69, 1193–1235.

Dasgupta, Amil, and Andrea Prat, 2008, Information aggregation in financial markets with career concerns, *Journal of Economic Theory* 143, 83 – 113.

Drehmann, Mathias, Jörg Oechssler, and Andreas Roider, 2005, Herding and contrarian behavior in financial markets: An internet experiment, *American Economic Review* 95, 1403–1426.

Glosten, Lawrence R., and Paul R. Milgrom, 1985, Bid, ask and transaction prices in a specialist market with heterogenously informed traders, *Journal of Financial Economics* 14, 71–100.

Hasbrouck, Joel, 2007, *Empirical Market Microstructure* (Oxford University Press).

Ivanov, Asen, Dan Levin, and James Peck, 2008, Hindsight, foresight, and insight: An experimental study of a small-market investment game with common and private values, *American Economic Review* forthcoming.

Kyle, Albert S., 1985, Continuous auctions and insider trading, *Econometrica* 53, 1315–1336.

Levin, Dan, and James Peck, 2008, Investment dynamics with common and private values, *Journal of Economic Theory* 143, 114 – 139.

McKelvey, Richard D., and Thomas R. 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, Andreas, and Hamid Sabourian, 2008, Herding and contrarian behavior in financial markets, Working paper Universities of Toronto and Cambridge.

Park, Andreas, and Daniel Sgroi, 2009, Herding and contrarian behaviour in financial markets: An experimental analysis, Working paper Universities of Toronto and Warwick.

Sgroi, Daniel, 2003, The right choice at the right time: A herding experiment in endogenous time, *Experimental Economics* 6, 159–180.

Smith, Lones, 2000, Private information and trade timing, *American Economic Review* 90, 1012–1018.

Vives, Xavier, 2008, *Information and Learning in Markets* (Princeton University Press).

Welch, Ivo, 1992, Sequential sales, learning, and cascades, *The Journal of Finance* 47, 695–732.

Ziegelmeyer, Anthony, Kene Boun My, Jean-Christophe Vergnaud, and Marc Willinger, 2005, Strategic delay and rational imitation in the laboratory, Working paper Max Planck Insitute of Economics.

Table 1
Fit of the data to the static model.

The table breaks up the decisions of traders by type and treatment into correct and incorrect. A decision is correct if, given the trading history, a trader buys if his (theoretical static) expectation exceeds the price and if he sells if his expectation is below the price. Details of the underlying treatments are in Section 5.

		static model S_1	S_2	S_3
Treatment 1, negative U shape	correct	58	22	64
	incorrect	3	63	16
	% correct	95%	26%	80%
Treatment 2, negative hill-shape		47	66	61
		13	33	9
		78%	67%	87%
Treatment 3, positive U shape		66	44	49
		10	48	11
		87%	48%	82%
Total		171	132	174
%		87%	48%	83%
Overall		**70%**		

Table 2
Herding and contrarian trades for all traders by treatment.

The first row in each treatment grouping lists how many herding trades were observed, the second row entries list the number of possible herding trades. An S_1 type cannot herd-sell and can herd-buy only if the price has risen. An S_3 type cannot buy-herd and can sell-herd only if the price has fallen. Similarly, an S_1 type cannot be a sell-contrarian and acts as a buy-contrarian only when buying after prices have fallen; conversely for the S_3 types. The description for the herding and contrarian actions for the S_2 types are more involved, but they are described in detail in Section 4.

		Herding						Contrarianism					
		S_1	S_3	S_2	hill	negative U shape	positive U shape	S_1	S_3	S_2	hill	negative U shape	positive U shape
Treatment 1	occurred	2		22		22		0	14	0		0	
U negative	possible	61		83		83		0	67	0		0	
	% occurred	3%		27%		27%			21%				
Treatment 2		4	3	14	14			6	6	18	18		
hill shape		40	30	65	65			17	25	30	30		
		10%	10%	22%	22%			35%	24%	60%	60%		
Treatment 3		6	0	1			1	2	6	31			31
U positive		70	17	12			12	3	38	74			74
		9%	0%	8%			8%	67%	16%	42%			42%
Treatment 4		14	2	24	24			4	15	22	22		
hill shape		127	47	114	114			20	62	30	30		
		11%	4%	21%	21%			20%	24%	73%	73%		
Treatment 5		15	0	4			4	3	5	54			54
U positive		139	15	20			20	8	84	156			156
		11%	0%	20%			20%	38%	6%	35%			35%
Treatment 6		14	0	48		48		1	28	0		0	
U negative		116	10	176		176		3	124	2		2	
		12%	0%	27%		27%		33%	23%	0%		0%	
Total possible		55	5	113	38	70	5	16	74	125	40	0	85
Total occurred		553	119	470	179	259	32	51	400	292	60	2	230
Total % occurred		10%	4%	24%	21%	27%	16%	31%	19%	43%	67%	0%	37%
single trade treatments		7%	6%	23%	22%	27%	8%	40%	20%	47%	60%		42%
two trades treatments		11%	3%	25%	21%	27%	20%	26%	18%	40%	73%	0%	35%

Table 3
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 as expressed in equation (1). Logit regressions report the marginal effects. Linear probability fixed effects regressions control for trader-fixed effects. The data is restricted to include only trades that could potentially be considered as herding trades. For all tables that follow, standard errors are in parentheses, * indicates significance at the 5% level, ** at the 1% level. We also omit the constants from the reports.

	Dependent Variable: Herding trade indicator				
Sample	all treatments	single trade treatments	two trade treatments both trades	two trade treatments first trade	two trade treatments second trade
Panel A: logit					
U shaped signal	0.120**	0.115**	0.123**	0.151**	0.094**
	(0.020)	(0.033)	(0.026)	(0.037)	(0.035)
Panel B: linear					
U shaped signal	0.143**	0.140**	0.144**	0.179**	0.109**
	(0.024)	(0.040)	(0.030)	(0.044)	(0.040)
Panel C: linear, fixed effects					
U shaped signal	0.128**	0.086*	0.114**	0.124*	0.103
	(0.024)	(0.043)	(0.035)	(0.051)	(0.052)
Observations	1142	378	764	398	366
R^2 linear	0.03	0.03	0.03	0.04	0.02
R^2 fixed effects	0.41	0.69	0.45	0.67	0.65

Table 4
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, as expressed in equation (1). Logit regressions report the marginal effects. Fixed effects regressions control for trader-fixed effects. The data is restricted to include only trades that could potentially be considered as contrarian trades. Standard errors and significance levels are denoted as in Table 3.

	Dependent Variable: Contrarian trade indicator				
Sample	all treatments	single trade treatments	two trade treatments both trades	two trade treatments first trade	two trade treatments second trade
Panel A: logit					
hill shaped signal	0.358**	0.284**	0.422**	0.479**	0.372**
	(0.059)	(0.088)	(0.085)	(0.132)	(0.114)
Panel B: linear					
hill shaped signal	0.410**	0.310**	0.494**	0.566**	0.431**
	(0.059)	(0.089)	(0.081)	(0.115)	(0.113)
Panel C: linear, fixed effects					
hill shaped signal	0.402**	0.274*	0.482**	0.458*	0.655**
	(0.065)	(0.124)	(0.111)	(0.204)	(0.172)
Observations	743	254	489	223	266
R^2 linear	0.06	0.05	0.07	0.10	0.05
R^2 fixed effects	0.44	0.72	0.55	0.78	0.8

Table 5
The Number of Trades in the Initial Five Seconds of Trading.

The table shows the number of trades in the first 5 seconds of trading. Note that the smallest time unit that our software supports is a second.

			S_1					S_3		
Treatments Trade	all	T1-T3	T4-T6 both	T4-T6 first	T4-T6 second	all	T1-T3	T4-T6 both	T4-T6 first	T4-T6 second
<1 sec	28	10	18	18	0	72	33	39	39	0
1 sec	46	13	33	31	2	30	12	18	14	4
2 sec	30	8	22	19	3	30	8	22	19	3
3 sec	37	9	28	16	12	13	2	11	8	3
4 sec	21	6	15	11	4	12	1	11	7	4
5 sec	16	5	11	10	1	20	6	14	9	5
$\leq$ 5 sec	178	51	127	105	22	177	62	115	96	19
all	623	198	425	222	203	583	210	373	190	183
% $\leq$ 5 sec	**29%**	26%	30%	47%	11%	**30%**	30%	31%	51%	10%

Table 6
Clustering of Trades.

The table lists the proportion of trades that were made within 1.5 seconds of their predecessor. Columns divide by signal types, rows successively exclude trades made in the first 5, 10, 20 and 30 seconds. As can be seen, there are no substantial differences by signal type.

	All	S_1	S_2	S_3	hill	$-$ve U	+ve U
All times	67%	66%	63%	71%	64%	66%	61%
total time >5 sec	58%	56%	57%	62%	57%	60%	55%
total time >10 sec	54%	52%	53%	58%	55%	55%	50%
total time >20 sec	51%	48%	51%	55%	56%	50%	49%
total time >30 sec	50%	44%	50%	54%	56%	49%	46%

Table 7
Return Trades.

The table lists summary statistics for return (or round-trip) transactions. Row 1 lists the total trades by types in treatments 4-6 (where two trades are possible). Row 2 lists the number of trades that were first trades. Row 3 lists the number of second trades. A discrepancy between Row 2 and 3 indicates that some people choose not to trade twice (Row 4). Row 5 lists how many of the second trades were classified as return trades (buy-sell or sell-buy). Row 6 lists how many of the return trades lead to an immediate trading profit. Row 7 lists the extend of buy-sell transactions (the remainder are sell-buy). Row 8 lists whether the first trade was in the same direction as prices thus far (i.e. did prices rise and was the first trade a buy or did prices fall and the first trade was a sale). Row 9 computes the same as Row 8 for the second trade. Some "trades" were "passes". For this table we count only the transactions; percentages in rows 5-7 do not add to one as there may be passes. (This affected 55 "trades". Specifically, there were 16 buy-holds, 22 sell-hold, 7 hold-sells, and 7 hold-holds. Most buy-holds (9) stemmed from S_3 types, most sell-holds (9) stemmed from S_1 types.)

	decreasing	increasing	hill-shape	pU shape	nU shape	All
Total trades	425	373	146	183	181	1308
First trades	222	190	76	94	92	674
Second trades	203	183	70	89	89	634
Percent foregone	9%	4%	8%	3%	5%	6%
buy-buy	7 3%	130 68%	9 12%	37 39%	14 15%	197 31%
sell-sell	146 66%	8 4%	28 37%	11 12%	46 50%	239 38%
Return trades	36 18%	29 16%	25 36%	33 37%	20 22%	143 23%
Profitable return	29 81%	19 66%	20 80%	25 76%	15 75%	108 76%
buy-sell	17 47%	21 72%	14 56%	22 67%	15 75%	54 72%
1st trade with price ($\mathsf{p} \nearrow \Rightarrow$ buy)	14 39%	20 69%	7 28%	22 67%	15 75%	78 55%
2nd trade with price ($\mathsf{p} \nearrow \Rightarrow$ buy)	15 42%	7 24%	5 20%	8 24%	5 25%	40 28%

Table 8
Impact of Return Trades on Herding and Contrarianism.

The table condenses six regressions of the equations in lines (2) and (3) (by signal type and then with respect to herding and contrarian behavior separately). When cells are empty, there was insufficient data or the variable was dropped. Constants were omitted from the report. Standard errors and significance levels are denoted as in Table 3.

	Herding			Contrarianism		
	all trades	first trade	second trade	all trades	first trade	second trade
U shaped signal	0.139**	0.133**	0.143**			
	(0.028)	(0.042)	(0.035)			
return × U shaped signal	-0.105*	0.024	-0.195**			
	(0.048)	(0.082)	(0.063)			
hill shaped signal				0.368**	0.324*	0.382**
				(0.102)	(0.136)	(0.143)
return × hill shaped				-0.072		-0.348
				(0.171)		(0.202)
return trade	0.233**	0.253**	0.212**	0.401**	0.189**	0.583**
	(0.028)	(0.043)	(0.037)	(0.044)	(0.059)	(0.077)
Observations	764	398	366	489	216	266

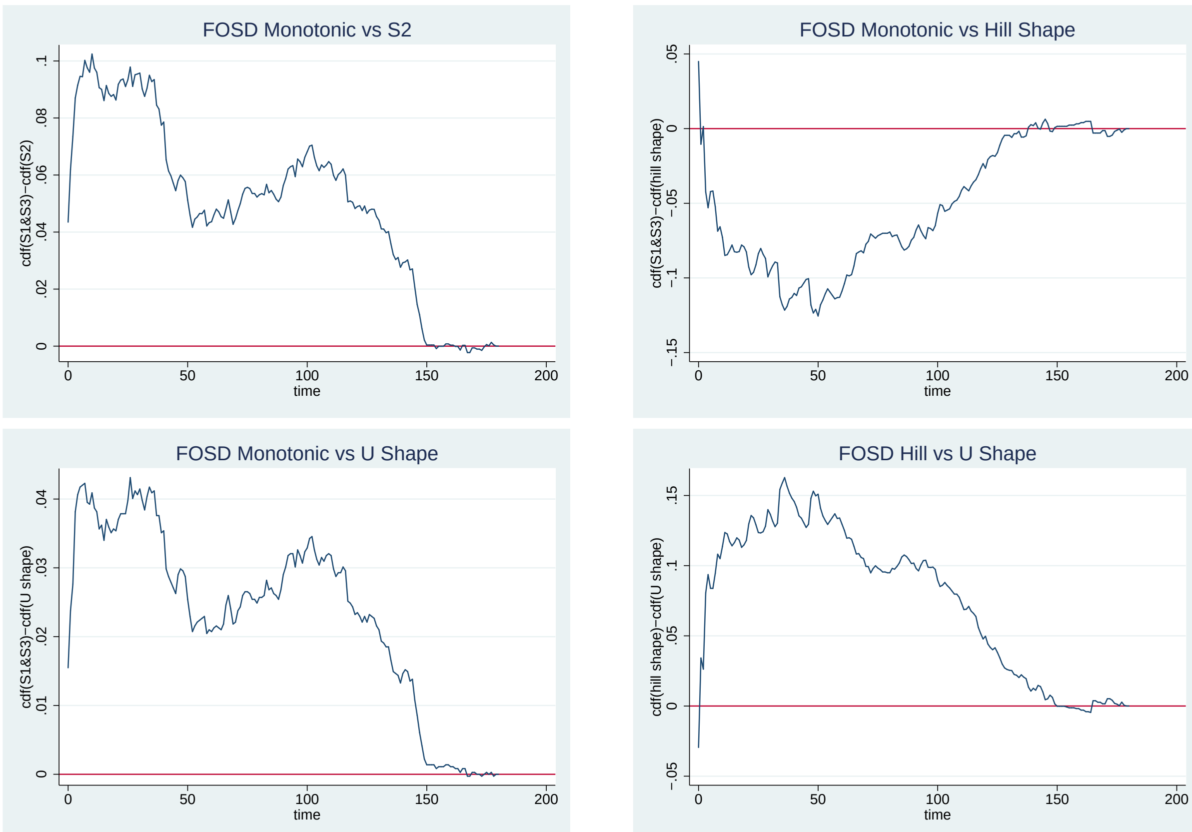

Figure 1

Plots for the differences of timing cdfs by signal types for treatments 1-6.

The four panels plot the differences of the distributions of the trading times, split up by signal types. Time is always on the horizontal axis, with 180 seconds signifying the end of trading. Differences of cumulative probabilities are on the vertical axes. The panels are labeled to signify the difference of distributions that was computed.

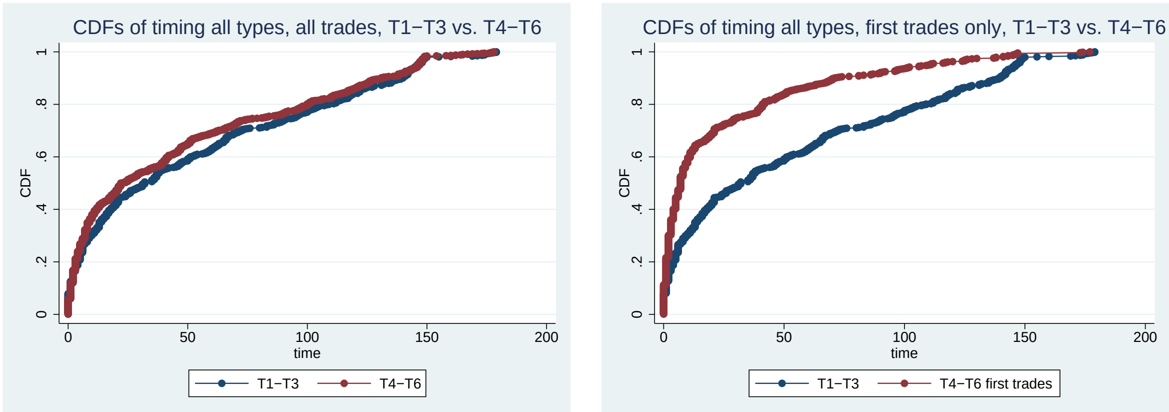

Figure 2
Plots for the timing cdfs by number of trades.

The three panels plot distributions of the trading times, split up by treatments with one and two trades. Axes are as in Figure 1. The left panel plots the cumulative probabilities for all trades; trades in treatments 4-6 occur weakly before those in treatments 1-3. The right panel looks only the distribution of the first trades in treatments 4-6 and all the trades in treatments 1-3. Here, clearly, trading occurs much earlier in treatments 4-6.

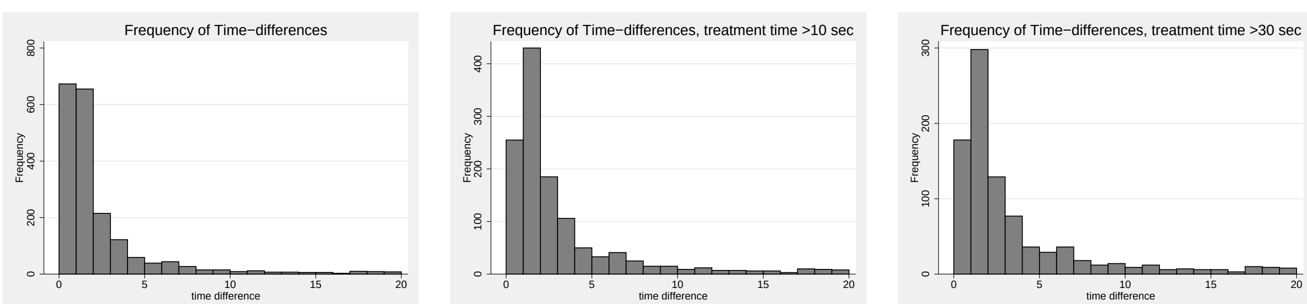

Figure 3
Plots for the time-difference pdfs.

The three panels complement Table 6 by plotting the frequencies of the time-difference of each trade to its immediate predecessor. The left panel plots the time-difference for all trades, the middle panel from all trades excluding those that happen in the first 10 seconds, and the right panel excludes trades in the first 30 seconds.

– Not for Publication –
Supplementary Appendix for
Herding, Contrarianism and Delay in Financial Market Trading

Andreas Park
University of Toronto

Daniel Sgroi
University of Warwick

September 30, 2009

Abstract

Part A of this document examines the performance of alternative behavioral models discussed in Section 4.4 of the main paper. Part B discusses whether information theory may help in understanding the timing decisions. Parts C through G detail the experimental time-line, instructions, an example information sheet, the software and the questionnaire. References to sections are with regards to those in the main paper. This appendix is not intended to be published with the paper, but rather provides additional information for the benefit of an interested audience.

A Alternative Explanations for Trading Behavior

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

The general assumption of our model is that people are risk-neutral. As a first check we will see if this assumption is warranted. Next, we will analyze if loss-aversion may play a role in people's behavior. We present the results for "standard parameters" but emphasize that we have also tried other specifications without being able to improve the fit. Finally we will check if various forms of alternative information updating provide a better fit with the data. These approaches usually depend on some parameter(s). Our approach is to vary this parameter and see how the variation improves the overall fit of the alternative model to the data. In this appendix we focus on the static decision only.

A.1 Risk and Loss Aversion

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

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

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

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

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

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

Loss-Aversion — S-Shaped Valuation Functions. A host of experimental work in prospect theory following Kahneman and Tversky (1979) has indicated that people pick choices based on change in their wealth rather than on levels of utilities. These costs and benefits of changes in wealth are usually assessed with valuation functions that are S-shaped. Kahnemann and Tversky suggested the following functional form

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

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

As with risk aversion, the performance of this model applied to our setup is much worse than the performance of the rational model. For parameters as estimated by Tversky and Kahneman (1992), the fit is below 49%. Table II. illustrates this observation for the above parameters as well as for one other configuration.[1]

A.2 Decision Rule: Prior Actions or No Updating

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

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

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

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

Table III. contains the details of the fit that is obtained under the two specifications outlined here.

[1] 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 relate to re-scaled probabilities which we analyze separately so as to be able to distinguish the effects of the two components.

A.3 Probability Scaling and Shifting

A yet weaker version of the no-updating alternative rule is probability shifting, whereby traders underplay (overplay) low (high) probabilities coming from the observed history H_{t-1}. Alternatively, traders may overstate the probabilities of their prior expectations; we present results from the latter but point out, that the former yields similar insights. The usual symmetric treatment of this under- or overstating of probabilities is to transform probability p into $f(p)$ as follows[2]

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

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

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

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

[2]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, $j = 1, \ldots, 3$.

A.4 Error Correction Provisions

Inspired by level-k reasoning (see Costa-Gomes, Crawford and Broseta (2001)) and Quantal Response Equilibria (see McKelvey and Palfrey (1995) and McKelvey and Palfrey (1998)), we will contemplate an alternative specification for hampered updating in which agents do not trust that their peers act fully rationally. In the rational model, consider a buy without herding in state V_i: this event occurs with probability $\beta_i = .25/2 + .75 \cdot \mathsf{Pr}(S_3|V_i)$ (recalling that .25/2 is the probability of a noise buy). Now imagine that instead subjects believe that only fraction δ of the informed buyers act rationally and that the remaining $1-\delta$ take a decision at random. Then the probability of a buy in state V_i becomes

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

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

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

A variation on this error correction theme is a specification in which a subject believes that fraction $1-\delta$ act randomly but the subject assumes that the remaining fraction δ takes this irrationality into account and reacts rationally to it. The difference to the first specification is that in the first, the subject not only assumes irrationality on the part of informed traders but also considers himself to be the only informed trader to take this into consideration. Now we instead allow a later subject to believe that his predecessors are also aware of the possible irrationality on the part of informed traders and employ this knowledge in their decision-making. Consequently, in the first specification, S_3 traders would never have been presumed to rationally sell, whereas in the second specification such behavior is admitted as rational.[3] Alas, as with the simple error correction, we do not obtain a substantially better

[3] Rather than directly implementing level-k reasoning or Quantal Response Equilibria, we choose our alternative specification because it is an unusually complex task for the subjects to calculate these more general measures of naive reasoning with 4 different known types of traders (noise traders and three types of informed trader). Moreover, there is a subtle difference of our approach to the way that Quantal Response

fit with the data, as can be gleaned from Table V.: we obtained the best fit for $\delta = 0$ in which case people act only on the basis of their prior expectation and do not update. For $\delta = .22$ (presented in the table; the figures for $\delta = 0$ coincide with those of the no-updating case), the fit is best for treatments 1-3 (treatments 4-6 have the best fit for $\delta = 0$). In the latter case, the improvement for treatments 1-3 only is from 69.8% to 76.1%.

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

A.5 Summary of Alternative Behavioral Explanations

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

Several studies (Drehmann, Oechssler and Roider (2005) and Cipriani and Guarino (2005)) have already identified that when prices rise, people with high signals tend to act as contrarians, i.e. they sell. There are multiple possible explanations, ranging from risk aversion (which we refute) to slow or no updating. We observe the same kind of end-point behavior by the S_3 types. Symmetrically, the S_1 types should exhibit similar behavior when prices approach the lower bound. However our data rarely involves prices that fall to a sufficient extent to examine the symmetric claim, since in general across all treatments, prices tend to tentatively rise. Note that the end-point effect should also influence the S_2 types, because whatever mechanism or cognitive bias leads S_3 types to sell for high prices should apply in the same manner to S_2 types.

Irrespective of which hypothesis is correct, if the end result is observationally equivalent to slow updating then this has a profound effect on how much herding or contrarian behavior one might expect to see: when people update slowly, it takes longer for them to reach a

Models can be implemented in models with and without prices. In an informational cascade without prices a deviation from the cascading action is, in principle, a deviation from rationality. With moving prices, such a simple observation can no longer be made, neither is it possible for subjects to determine if there is a genuine error. Our notion of overweighing noise is therefore a simple means for subjects to model the lack of trust in predecessors' actions, without implying a definitive or systematic direction of the error. Traders thus act as if the proportion of noise traders were higher than 25% by downgrading the quality of information extracted from the history of actions embodied in H_{t-1} or q_t. Finally, since we already have noise traders built into the experiment, by opting to allow traders to increase their estimates of the percentage of expected noise trades above 25% our method is arguably an especially simple and intuitive rule of thumb which enables subjects to incorporate naive reasoning on the part of their peers. For more on rules of thumb by laboratory subjects in a herding context see Ivanov, Levin and Peck (2008).

(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.

B Information Theory and Timing

Our analysis in the main text shows that the timing behavior of individuals depends strongly on the type of their signal. For instance, we argue that subjects with good-news–bad-news information act systematically earlier than those with bi-polar and single-polar information. We now take a second look at the signals' information content, trying to assert if the timing behavior is consistent with information theory.[4] Specifically, one of the standard measures of signal informativeness is *entropy.* If $p|S = (\mathsf{Pr}(V_1|S), \mathsf{Pr}(V_2|S), \mathsf{Pr}(V_3|S))$ is a conditional probability distribution for the three states given signal S, then the entropy of this distribution is

$$H(p|S) = -\sum_{i=1}^{3} \mathsf{Pr}(V_i|S) \log_2(\mathsf{Pr}(V_i|S)).$$

The larger H, the smaller the information content; its minimum is attained for a uniform distribution. The subjects were given the following signal distributions:

Signal Distribution

	S_1			S_2			S_3		
Type	V_1	V_2	V_3	V_1	V_2	V_3	V_1	V_2	V_3
U-negative	0.65	0.45	0.05	0.3	0.1	0.25	0.05	0.45	0.7
hill	0.65	0.1	0.05	0.3	0.8	0.25	0.05	0.1	0.7
U-positive	0.7	0.45	0.05	0.25	0.1	0.3	0.05	0.45	0.65
Posterior Distribution on values									
U-negative	0.565	0.391	0.043	0.462	0.154	0.385	0.042	0.375	0.583
hill	0.813	0.125	0.063	0.222	0.593	0.185	0.059	0.118	0.824
U-positive	0.583	0.375	0.042	0.385	0.154	0.462	0.043	0.391	0.565

Applied to the posteriors generated by these signals, we can then compute the following entropies

[4]For comprehensive overviews see Khinchin (1957) or Reza (1994).

Type	entropy $H(p|S)$ S_1	S_2	S_3
U-negative	1.192	1.460	1.175
hill	0.868	1.380	0.834
U-positive	1.175	1.460	1.192

This table yields an information-ranking of the nine signals, specifically, 1. Hill S_3, 2. Hill S_1, 3. U-negative S_3 and U-positive S_1, 4. U-positive S_3 and U-negative S_1, 5. Hill S_2, and 6. U-positive and U-negative S_2. Of course, we have already seen in the main text that 5. and 6. are dominated by the combination of 1.-4.

Next, the entropy measures for 1. and 2., 3. and 4. and 5. and 6. are very close. The left panel in Figure 1 depicts the cumulative distributions of the combined 'similar' signals. Again, our results thus far clearly indicate that 5. and 6. combined are dominated by the other two combinations. It is however, noteworthy that 1. and 2. and 3. and 4. both depict good-news–bad-news signals. Thus applying Smith (2000), there should be no order — yet there is one.

There are, however, some conceptual objections that one may want to put forward: while a hill-shaped signal S_2 has a bad entropy value, the signal itself is generally a strong endorsement for the middle state and it is intuitively not clear why it should be dominated by cases 3. and 4. We thus split up the distributions by the six entropy values (the graph is for treatments 1-3, but the cdfs look similar for the other combinations that we consider in the main text). The right panel in Figure 1 depicts the respective cdfs. Focussing on the hill-shaped S_2 types, we observe, that the S_1 and S_3 types do trade systematically earlier in the hill-shaped treatment 2 (this is with some reservation for the S_1 types). There is also an order between the U-shaped and the hill-shaped S_2 types. But there is no clear order between the hill-shaped S_2 types and the S_1 and the S_3 types in the U-shaped treatments 1 and 3. This is notable because their entropy values are further from the hill-shaped S_2 types than are the U-shaped S_2 types.

In other words, there must be some other factors driving the timing decision that are not covered by information theory only. With this in mind, we believe that the analysis thus far indicates that herding and contrarian motivations can contribute an important part to understanding the timing behavior.

C Time-line

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

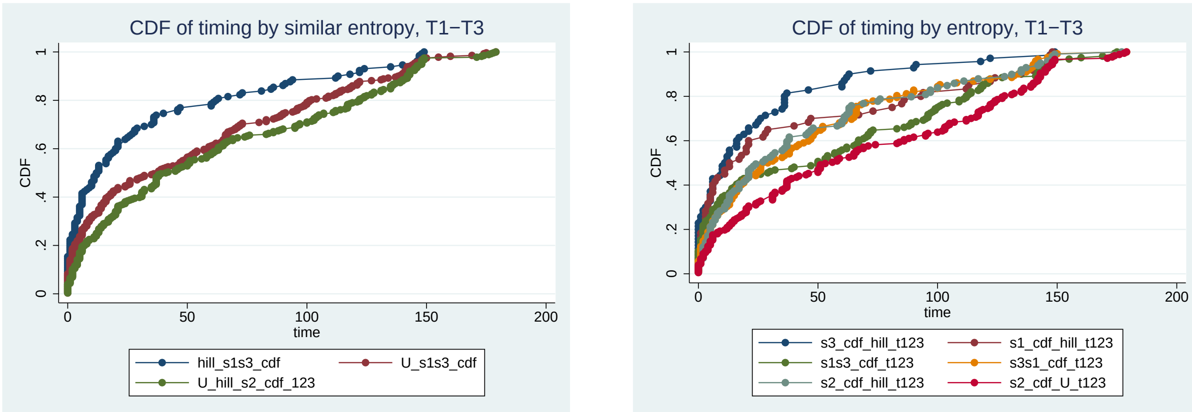

Figure 1
Cumulative distributions ordered by signal entropy.

Analogously to the timing figures employed in the main text, the two panels plot distributions of the trading times, split up by signals types and type of treatment. Time is always on the horizontal axis, with 180 seconds signifying the end of trading. Cumulative probabilities are on the vertical axes. The left panel aggregates and collects trading times for signals with similar entropy values (see Section B for details). The right panel collects the trading times separately for the six different entropy levels that the experiments employed. The trading times are collected only for treatments 1-3, but the graphs look similar for other data-specifications (e.g. for treatments 4-6 etc.).

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.

D 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. The instructions are long, and 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. Note that in the instructions the example round is called "round 1" with the true experiment encompassing rounds 2-7. In the main text of the paper we instead call the rounds "treatments", and ignoring the example round, renumber them to be treatments 1-6.

D.1 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.

D.2 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.

D.3 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 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 or shares in 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 an additional one or two shares or shares. Round 1 will be an example round and your final payment will not reflect how you perform during this round. In rounds 2-4 you will be allowed to trade once (ie to buy, sell or hold one time only), and in rounds 5-7 you will be allowed to trade twice. You will have 3 minutes in which to trade, and we will announce when the time reaches 2 minutes and 30 seconds and 2 minutes and 50 seconds.

During the rounds you may sell your share, you may buy one or in some cases two additional shares or you may do nothing. When you decide to trade (by hitting the buy, sell or pass button) that trade cannot be undone and will be recorded as your first trade.

Depending upon the rules of each round you may be able to trade again. Once you have hit the button it may take the system a fraction of a second to register your trade. You should not double-click or attempt to click more than once, unless of course you wish to record two trades in close succession.

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

D.4 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.

D.5 Instructions 5 (Compensation)

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

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

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

Your profit relates to the difference between the current price that you buy or sell a share at today, and the price revealed after the event takes place. An example of a good event happening to the company might be that it wins a court case or gains a patent. A bad thing might be the opposite, so the firm loses a court case or fails to gain a patent. Note that as already stressed, each round is an independent experiment, so in round 1 it may be that the bad event takes place so the share price becomes 75 after trading finishes, while in round 2 it may be worth 125, etc.

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

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

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

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

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

D.6 Instructions 6 (The Information Setting)

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

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

- It could mean that he believes nothing at all will happen hence he believes the price will revert to the original 100 vcu and we call this *case 1.*
- Or it could mean that he believes an event will happen but he is not sure whether it is either good or bad, and we call this *case 2.*
- Or it could mean that he believes something good or bad will happen and he has a feel for which, but he is not sufficiently sure to indicate the good or bad tip and would prefer to indicate middle and we call this *case 3.*

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

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

You will be given 2 minutes to examine the relevant sheet before each round. You will then receive notification on your computer screen of the actual tip sent to you from the broker: S_1, S_2 or S_3, and will have another minute to consider this. The beginning of the round will then be announced and trading will begin. Remember that each round only lasts for 3 minutes and you will be informed when 2 minutes and 30 seconds and when 2 minutes and 50 seconds have elapsed. The buttons on the screen (buy, sell or pass) can only be pressed during this time and only once per round in rounds 1-4 and twice in rounds 5-7.

D.7 Instructions 7 (Summary)

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

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

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

D.8 Instructions 8 (Experiment End)

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

E 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.

F Questionnaire

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

About you

1. Your age:

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

About your decisions today

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

About the experiment

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

G 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 5 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).

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 6 below shows a screen shot of the software in action.

Here you can see that the price initially rose from a level of 100, indicating buying at the early stages, but then price started to fall back, it rallied and then fell back further to a value of around 116. This subject's private signal was S_1 ("bad") and the subject had a single share to sell and a large cash balance to enable the purchase of a further share. The subject could also pass (declining to buy or sell) when given the opportunity to trade.

The software was purposefully built for the experiment, since existing software was unable to provide the sort of information structure needed in a price-driven (as opposed to order-driven) market.[5]

References

Costa-Gomes, Miguel, Vincent P. Crawford, and Bruno Broseta, "Cognition and Behavior in Normal-Form Games: An Experimental Study," *Econometrica*, 2001, *69*, 1193–1235.

Ivanov, Asen, Dan Levin, and James Peck, "Hindsight, Foresight, and Insight: An Experimental Study of a Small-Market Investment Game with Common and Private Values," *American Economic Review*, 2008, *forthcoming.*

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

Khinchin, A., *Mathematical Foundations of Information Theory*, New York: Dover Publications, Inc., 1957.

McKelvey, Richard D. and Thomas R. Palfrey, "Quantal Response Equilibria for Normal Form Games," *Games and Economic Behavior*, 1995, *10*, 6–38.

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

Reza, F., *An Introduction of Information Theory*, New York: Dover Publications, Inc., 1994.

Smith, Lones, "Private Information and Trade Timing," *American Economic Review*, 2000, *90*, 1012–1018.

Tversky, Amos and Daniel Kahneman, "Advances in Prospect Theory: Cumulative Representation of Uncertainty," *Journal of Risk and Uncertainty*, 1992, *5*, 297–323.

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

		Total Number of wrong decisions CRRA utility, $\gamma = 1$ (log-utility)			Total Number of wrong decisions CARA utility, $\rho = 2$	
	S1	S2	S3	S1	S2	S3
treatment 1 correct	58	61	14	58	22	62
U-negative wrong	3	24	66	61	85	80
% correct	95%	72%	18%	95%	26%	78%
treatment 2	47	53	9	47	62	60
hill	13	46	61	60	99	70
	78%	54%	13%	78%	63%	86%
treatment 3	66	33	6	66	25	49
U-positive	10	59	54	76	92	60
	87%	36%	10%	87%	27%	82%
treatment 4	127	90	17	126	97	98
hill	25	57	104	152	147	121
	84%	61%	14%	83%	66%	81%
treatment 5	123	59	5	123	82	98
U-positive	30	124	101	153	183	106
	80%	32%	5%	80%	45%	92%
treatment 6	103	120	28	103	46	110
U-negative	18	60	119	121	180	147
	85%	67%	19%	85%	26%	75%
Total%	84%	53%	14%	84%	42%	82%
treatments 1-3%	87%	53%	14%	87%	39%	81%
treatments 4-6%	83%	53%	13%	83%	44%	82%
Fit total	51%			67%		
fit treatments 1-3	51%			66%		
fit treatments 4-6	51%			67%		

Table I.
Risk-Aversion Analysis.

The table classifies trades as right or wrong assuming that traders took the decisions according to an underlying model that admitted risk-averse behavior. The first set of columns looks at the case with constant relative risk aversion utility (or power utility; we obtained the best fit for the log-utility function). The second set of columns looks at the case of constant absolute risk aversion (or exponential utility); while the fit for risk aversion parameter $\rho = 2$ is not the best, it is indicative. As ρ decreases so that we approach risk neutrality, the fit improves and it is bounded above by the fit of the risk neutral model.

	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 II.
Loss-Aversion Analysis.

The table classifies trades as right or wrong assuming that traders took the decisions according to an underlying model that admitted a loss-averse valuation function as depicted in Subsection A.1. The two sets of columns depict popular specifications for the Kahneman and Tversky parameters α, β, γ. As can be seen, the fit is much lower than with the rational, risk-neutral model. The structure of the table is similar to that of Table I.; we omit the number of wrong decisions as they can be straightforwardly obtained from the total number of decisions in Table I..

		No updating			prior action		
		S1	S2	S3	S1	S2	S3
treatment 1	correct	60	62	64	58	63	64
U-negative	wrong	1	23	16	3	22	16
	% correct	98%	73%	80%	95%	74%	80%
treatment 2		47	61	63	47	53	61
hill		13	38	7	13	46	9
		78%	62%	90%	78%	54%	87%
treatment 3		66	47	47	66	46	49
U-positive		10	45	13	10	46	11
		87%	51%	78%	87%	50%	82%
treatment 4		134	108	97	127	74	97
hill		18	39	24	25	73	24
		88%	73%	80%	84%	50%	80%
treatment 5		123	81	97	123	80	99
U-positive		30	102	9	30	103	7
		80%	44%	92%	80%	44%	93%
treatment 6		103	115	118	103	89	114
U-negative		18	65	29	18	91	33
		85%	64%	80%	85%	49%	78%
Total%		86%	60%	83%	84%	52%	83%
treatments 1-3%		88%	62%	83%	87%	59%	83%
treatments 4-6%		85%	60%	83%	83%	48%	83%
Fit total		75%			71%		
fit treatments 1-3		76%			74%		
fit treatments 4-6		75%			69%		

Table III.
No Updating and Prior Actions.

The table lists the results from comparing the decisions taken to those that would be optimal if agents do not update (the first set of columns) or simply take the decision that is optimal ignoring the history and all prices (the second set of columns). The structure of the table is similar to that in Table I. with correct and wrong actions listed alongside one another.

	With $\alpha = 25$			With $\alpha = 10$			With $\alpha = 5$		
	S1	S2	S3	S1	S2	S3	S1	S2	S3
treatment 1	58	60	66	58	35	66	58	25	66
U-negative	3	25	14	3	50	14	3	60	14
	95%	71%	83%	95%	41%	83%	95%	29%	83%
treatment 2	47	69	61	47	69	61	47	68	61
hill	13	30	9	13	30	9	13	31	9
	78%	70%	87%	78%	70%	87%	78%	69%	87%
treatment 3	66	59	54	66	59	54	66	59	54
U-positive	10	33	6	10	33	6	10	33	6
	87%	64%	90%	87%	64%	90%	87%	64%	90%
treatment 4	127	110	104	127	110	104	127	110	104
hill	25	37	17	25	37	17	25	37	17
	84%	75%	86%	84%	75%	86%	84%	75%	86%
treatment 5	123	124	101	123	124	101	123	121	101
U-positive	30	59	5	30	59	5	30	62	5
	80%	68%	95%	80%	68%	95%	80%	66%	95%
treatment 6	103	99	119	103	77	119	103	62	119
U-negative	18	81	28	18	103	28	18	118	28
	85%	55%	81%	85%	43%	81%	85%	34%	81%
Total%	84%	66%	86%	84%	60%	86%	84%	57%	86%
treatments 1-3%	87%	68%	86%	87%	59%	86%	87%	55%	86%
treatments 4-6%	83%	65%	87%	83%	61%	87%	83%	57%	87%
Fit total	78%			75%			74%		
fit treatments 1-3	79%			75%			74%		
fit treatments 4-6	77%			75%			74%		

Table IV.
Overweighting of the Prior.

The table lists the results from comparing the decisions taken with those that would be optimal under the hypothesis that traders rescale and overweight their prior as depicted in Subsection A.3. The structure of the table is similar to that in Table I. with correct and wrong actions listed alongside one another.

	simple noise shift $\delta = 2/15$			simple noise shift $\delta = 1/3$			level 2 noise shift $\delta = .22$		
	S1	S2	S3	S1	S2	S3	S1	S2	S3
treatment 1	58	61	60	58	60	65	58	61	66
U-negative	3	24	20	3	25	15	3	24	14
	95%	72%	75%	95%	71%	81%	95%	72%	83%
treatment 2	47	64	63	47	58	63	47	66	63
hill	13	35	7	13	41	7	13	33	7
	78%	65%	90%	78%	59%	90%	78%	67%	90%
treatment 3	66	44	47	66	42	48	66	44	49
U-positive	10	48	13	10	50	12	10	48	11
	87%	48%	78%	87%	46%	80%	87%	48%	82%
treatment 4	127	94	98	127	92	98	127	94	98
hill	25	53	23	25	55	23	25	53	24
	84%	64%	81%	84%	63%	81%	84%	64%	80%
treatment 5	123	69	97	123	64	97	123	67	98
U-positive	30	114	9	30	119	9	30	116	8
	80%	38%	92%	80%	35%	92%	80%	37%	92%
treatment 6	103	116	117	103	97	114	103	108	114
U-negative	18	64	30	18	83	33	18	72	33
	85%	64%	80%	85%	54%	78%	85%	60%	78%
Total%	84%	57%	83%	84%	53%	83%	84%	56%	83%
treatments 1-3%	87%	61%	81%	87%	58%	84%	87%	62%	85%
treatments 4-6%	83%	55%	83%	83%	50%	83%	83%	53%	83%
Fit total	73%			71%			72.9%		
fit treatments 1-3	75%			74%			76%		
fit treatments 4-6	72%			70%			71%		

Table V.
Variations in the Perception of Noise Trading.

The table lists the results from comparing the decisions taken with those that would be optimal under the hypothesis that traders correct for the possibly of random actions by their peers as depicted in Subsection A.4. The first two sets of columns look at the situation in which a certain fraction takes a random action; this can also be understood as an overweighing of the extent of noise trading. The third set of columns considers the possibility that the fraction of traders that does not act irrationally reacts rationally to the irrationality of the remaining players. The structure of the table is similar to that in Table I. with correct and wrong actions listed alongside one another.

Figure 2
Information Sheet for positive U Shape

Round

Signals: **Case 2**

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

If the true effect will be POSITIVE then you receive

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

If the true effect will be NEGATIVE then you receive

• Signal S1 (bad) with chance	70%
• Signal S2 (no effect) with chance	25%
• Signal S3 (good) with chance	5%

If indeed the effect will be NO EFFECT then you receive

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

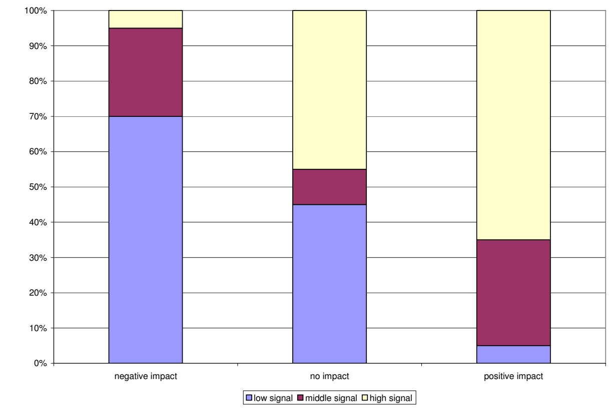

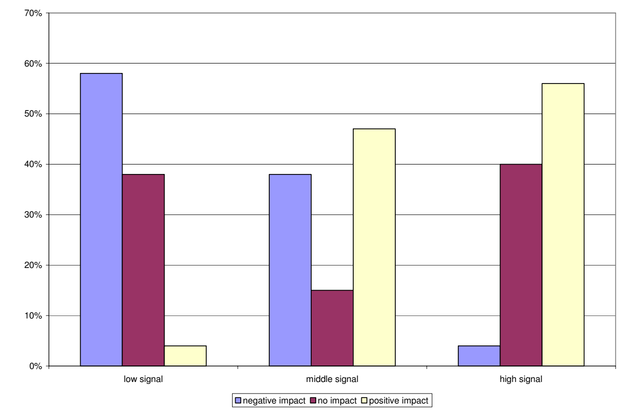

Figure 3
Information Sheet for negative U Shape

Round

Signals: **Case 2**

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

If the true effect will be POSITIVE then you receive

- Signal S1 (bad) with chance 5%
- Signal S2 (no effect) with chance 25%
- Signal S3 (good) with chance 70%

If the true effect will be NEGATIVE then you receive

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

If indeed the effect will be NO EFFECT then you receive

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

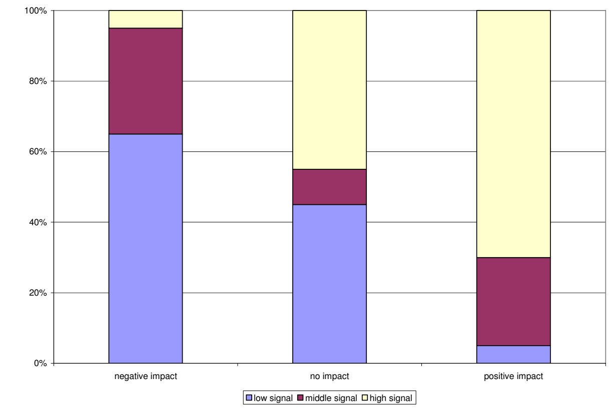

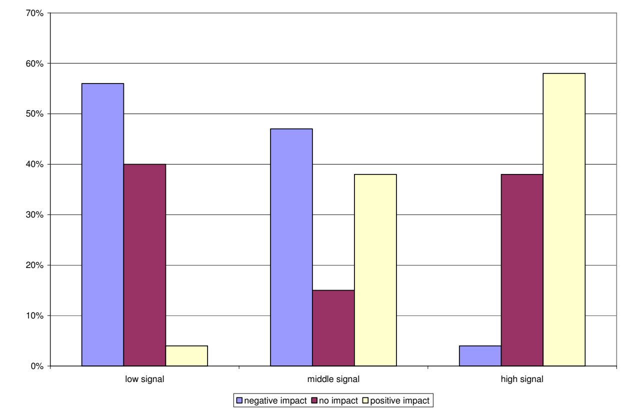

Figure 4
Information Sheet for negative Hill shape

Round

Signals: **Case 1**

- If you receive signal S1 (the “bad” signal), then the broker indicates a *negative* impact.
- If you receive signal S3 (the “good” signal), then the broker indicates a *positive* impact.
- If you receive signal S2 (the “middle”), then the broker indicates that there is *no effect.*

If the true effect will be POSITIVE then you receive

- Signal S1 (bad) with chance 5%
- Signal S2 (no effect) with chance 25%
- Signal S3 (good) with chance 70%

If the true effect will be NEGATIVE then you receive

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

If indeed the effect will be NO EFFECT then you receive

- Signal S1 (bad) with chance 10%
- Signal S2 (no effect) with chance 80%
- Signal S3 (good) with chance 10%

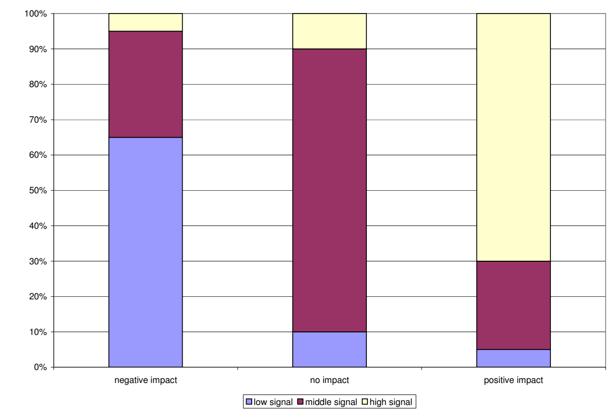

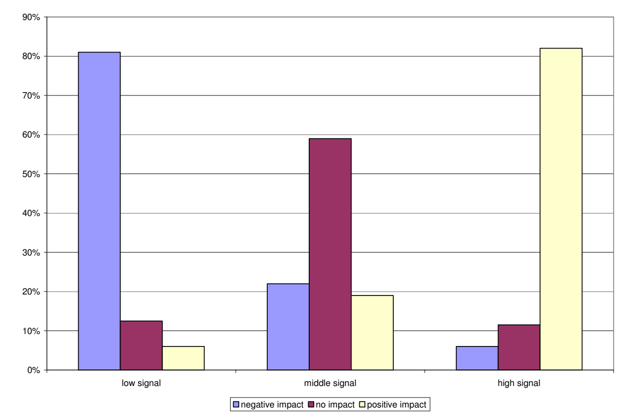

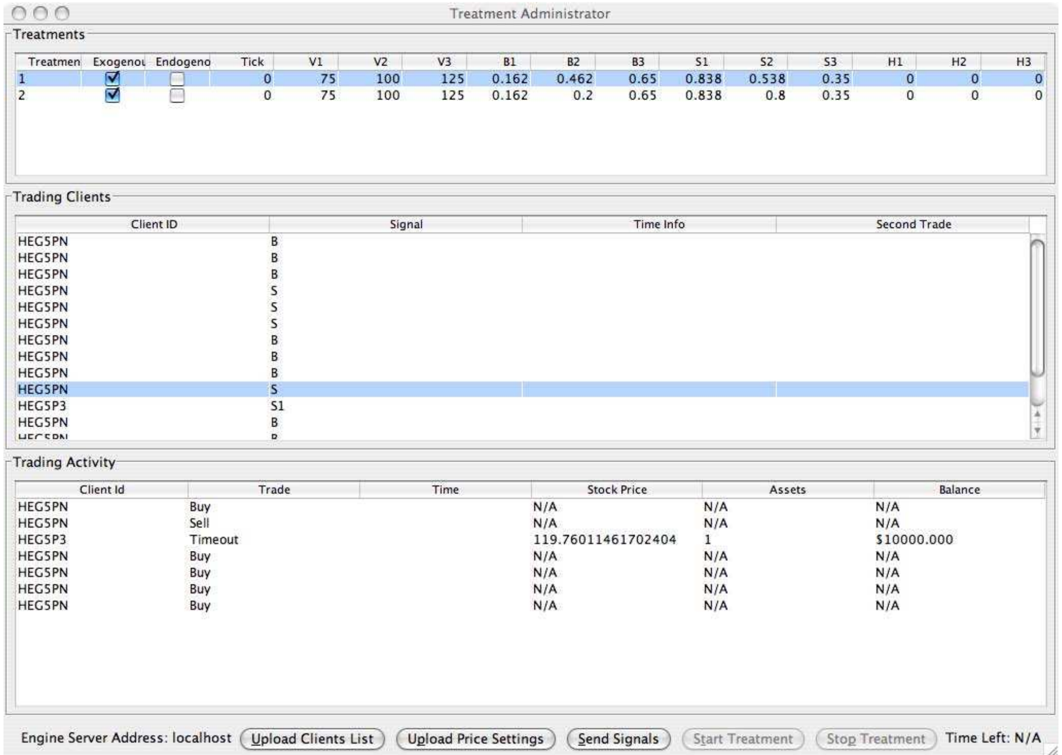

Figure 5
The Administrative Interface

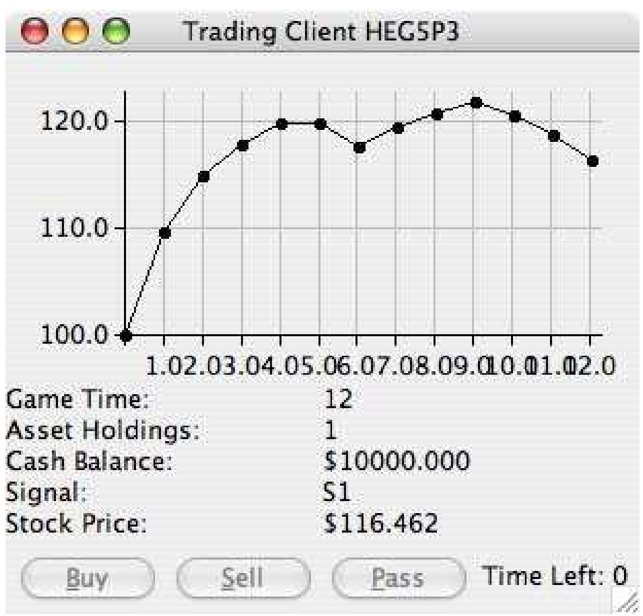

Figure 6
The Trading Client