Herding and Contrarian Behavior in Financial Markets: An Experimental Analysis[*]

Andreas Park[†]
University of Toronto

Daniel Sgroi[‡]
University of Warwick

September 30, 2009

Abstract

We are the first paper to analyze and confirm the existence and extent of rational informational herding *and* rational informational contrarianism in a financial market experiment, and to compare and contrast these with the equivalent irrational phenomena. In our study, subjects generally behave according to benchmark rationality. Traders who should herd or be contrarian in theory are the significant source of both within the data. Correcting for subjects who chose not to trade at least once (an irrational action in itself), increases our ability to predict herding or contrarian behavior considerably.

JEL Classification: C91, D82, G14.
Keywords: Herding, Contrarianism, Informational Efficiency, Experiments.

[*]We thank seminar participants at the 3rd World Congress of the Game Theory Society at Evanston, the 42nd annual meeting of the Canadian Economic Association at Vancouver and the ESRC Conference on World Economy and Finance at Warwick. We would also like to thank attendees at seminars in Cambridge, Toronto, Southampton and Birmingham. Special thanks to Antonio Guarino, Jim Peck and Gustavo Borbonis for invaluable advice and to Annie Jekova, Malena Digiuni, Kenji Yee and Kelly Hayward for expert research assistance. Financial support from the ESRC (grant number RES-156-25-0023), the SSHRC and the Cambridge Endowment for Research in Finance (CERF) is gratefully acknowledged.

[†]E-mail: andreas.park@utoronto.ca

[‡]E-mail: daniel.sgroi@warwick.ac.uk

Rational informational herding has become an important tool in analysing why and how observational learning by economic agents can affect economic outcomes (Banerjee (1992) and Bikhchandani, Hirshleifer and Welch (1992)). People constantly learn from others through discussion, newspapers, observing actions, and typically in financial markets through observing price movements or the buy and sell decisions of others. Informational herding arises in situations where people observe the actions of others, derive information from them and then, seemingly disregarding their own information, follow the majority action. A central lesson of herding theory is that the actions that people take are not necessarily a direct indication of their information. Moreover, a few early incorrect decisions, through a process of rational observation and inference, can have serious ramifications for all who follow.

A loose application of herding theory to financial market trading might suggest that early movements by visible traders can induce discontinuous price jumps in one direction or the other, potentially leaving 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 avoiding financial crises. To this end, financial commentators and economists, e.g. Shiller (2008), continue to consider rational herding a possible explanation for financial crises and shocks; contrarianism, the counterpart to herding, is also prevalent in market transaction data (see Chordia, Roll and Subrahmanyam (2002)) and in experimental settings (see, for instance, Alevy, Haigh and List (2007)).

Yet, the early work on rational herding was not designed to be directly applied to security market trading. This early work did not consider prices that react to actions, while in financial markets (efficient) market prices drop after sales and rise after buys. A path-breaking paper by Avery and Zemsky (1998) introduced efficient prices to a sequential herding context. They showed that in a simple financial market-trading setting with two values herding is not possible, since an efficient market price always separates people with favourable and unfavourable information about the security so that the former always buy and the latter always sell.

Introducing more states, however, Avery and Zemsky (1998) and, more generally, Park and Sabourian (2008), derive sharp predictions as to when herding should and should not happen depending upon the kind of information that traders receive. We will employ such a general setting in our experimental setup and test whether we find experimental evidence for rational herding. 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 it is often cited as an important phenomenon in real-world markets.

To focus on our contribution, our paper is the first to analyze and confirm the existence and extent of *rational* informational herding and *rational* informational contrarianism in a financial market experiment, and to compare and contrast these with the equivalent *irrational* phenomena. Our experiment is also one of the largest yet to examine rational herding in a laboratory setting, with around 1350 trades across 6 separate treatments.

Theoretical Background. The 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, arriving at the market in a predetermined, exogenous sequence. Past trades and prices are public information, and the market maker adjusts the price after each transaction to include the new information revealed by this trade.

In our experimental setting, there is a single security with three possible liquidation values (high, middle, or low). Subjects can choose between three actions (buy, sell, or no trade). Each subject receives a private realization of one of three possible signals (high, middle, or low) that are informative 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. We refer to these as 'monotonic' signals.

We study four different specifications of the middle signal. In two of these, signal recipients obtain 'weaker' versions of the high and low signal respectively and thus shift weight from high to low values systematically but in a less pronounced manner. In the remaining two specifications, they shift weight either towards the middle value or away from it towards *both* extreme values. We refer to such signals as 'non-monotonic'.

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. As Park and Sabourian (2008) show, recipients of the "low" signal should always sell, recipients of the "high" signal should always buy. Herding occurs if someone switches from selling to buying as prices rise or from buying to selling as prices fall. To do so, a signal recipient must update his or her 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 signal recipients have information 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".

In the natural counterpart situation, 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 upwards slower than the price rises. This slower speed of updating is possible if and only if signal recipients receive information that causes them to shift weight towards the middle outcome, or, in technical terms, if their conditional signal distribution is "hill shaped".

Including subjects who receive non-monotonic signals, our experimental setup is the first that admits herding *and* contrarianism as rational outcomes. We can thus disentangle rational and irrational forms of herding and contrarianism. Our analysis consists of two parts. In the first part, we test whether our experimental subjects abide by the theoretical model put forward by Avery and Zemsky (1998) and Park and Sabourian (2008). In the second part, we ascertain whether signals have the effects on herding and contrarianism as predicted by Avery and Zemsky and Park and Sabourian.

Findings. The overall fit of the data to the theoretical model is roughly 75%. Broken down by player type, these numbers are 82%, 61% and 86% for the low, the middle and the high signal type respectively. These figures are very much in line with the results in Drehmann, Oechssler and Roider (2005) and Cipriani and Guarino (2005).[1] Yet although the fit is very high for the high and low signal types, it is low for middle signal types.

To appreciate our results and our contribution it is important to understand how it relates to the current literature. Avery and Zemsky (1998) show theoretically that herding cannot arise if there are only two liquidation values. Experimental research by Drehmann et al. (2005) and Cipriani and Guarino (2005) confirmed this result: herding does not arise with two liquidation values. One can interpret their findings as showing that people do not exhibit a natural tendency to herd. This line of experimental research also showed that people tend to act as contrarians, or, in other words, that people have a natural (albeit irrational) tendency to act against the crowd.

This natural tendency to be a contrarian may bias behavior against *rational* herding and may also overwhelm rational motivations to be a contrarian. Each effect has implications for price efficiency and volatility. The irrational tendency for contrarianism may prevent herding and the implied excessive price volatility. Yet it also hinders justified movements because when there is irrational contrarianism or a lack of rational herding then prices adapt slower towards extreme values than they should. Furthermore, prices in the theoretical model are a martingale so that "price charting" or "technical analysis" has no merit. If, however, there is an irrational tendency for contrarianism, then there may be persistent mean reversion so that "price charting" can be beneficial.

Given the results on irrational contrarianism in the literature, not surprisingly, our re-

[1] To qualify this, both of these papers consider only two types of agent which are equivalent to the low and high signal type in our experiment.

sults on herding are mixed: we observe that herding arises less often than predicted by the theory (in only about 30% of the predicted cases). Our results on rational contrarianism are stronger: it arises in about 77% of the predicted cases.

In the second part of our analysis we show that non-monotonic signals as implied by the theory are the significant cause for herding and contrarianism. Specifically, when herding arises, it is most likely that this behavior is by someone who has the theoretical potential to herd: obtaining a U shaped signal significantly increases the chance of acting as a herder, compared to receiving any other signal, by 8%. This indicates that while the contrarian tendency of traders is strong enough to prevent some of the herding that should arise, U shaped signals are still the relevant source of herding. Moreover, if contrarianism arises, it is most likely to be caused by someone with contrarian information; receiving a hill shaped signal increases the probability of acting as a contrarian by 50% relative to receiving any other signal.

Next, the types who received weaker versions of the extreme signals behaved significantly less rationally than the extreme types. For instance, take the type that received a weaker version of the extremely good signal. When omitting the non-monotonic signals from the data, receiving this type of information significantly increases the chance of acting as both a sell contrarian *and* as a sell herder.

Finally, the decision not to trade is never rational in our model. Consequently, one can argue that the subjects who decide not to trade at some point in the experiment are less rational than those who choose to always trade. When splitting the subjects into two pools according to this simple rationality test (the first pool contains all subjects who never choose to pass, the second contains all subjects who choose to pass in at least one treatment), we observe that the behavior of the more rational types is much more in line with the theory with respect to trading on the basis of non-monotonic signals.

Experimental Work on Financial Market Herding. To match herding theory with data, one needs to know exactly what information each trader has at each time; this makes traditional empirical work difficult and experimental testing, where information is closely controlled, an ideal setting. Numerous papers have thus examined herding behavior in experimental settings, but only a few employ efficient prices.

The first published experiment to test herding was Anderson and Holt (1997), albeit in a setting without moving prices. They found that herding did occur, but at lower levels than predicted by the theory (73%), which they justified in terms of assumed errors made by predecessor decisions. Following Anderson and Holt, most experiments since have not specifically considered a financial market setting and have implicitly held prices fixed.[2]

[2]A more recent contribution using a setting without moving prices is Alevy et al. (2007), who used professional Chicago Board of Trade traders as their subjects.

Exceptions are Drehmann et al. (2005) and Cipriani and Guarino (2005), discussed earlier, and Cipriani and Guarino (2007). The latter tested Avery and Zemsky's model with so-called 'event uncertainty' in the lab, enlisting financial professionals. They find a degree of rational herding in line with our own results, which suggests that our student subject pool does not drive the results. Using a larger subject pool, we complement their work by studying rational contrarianism and by obtaining results on the impact of information structures on herding and contrarianism. There is a wealth of research in the empirical finance literature that stresses the role of contrarian tendencies in real-world data (e.g. Chordia et al. (2002)). Prior experimental research may have given the impression that these contrarian tendencies are irrational; our findings here suggest instead that this need not be the case.

There also exists an older experimental literature on information aggregation in financial markets (for instance, Plott and Sunder (1988)). Work in this area studies the capacity of prices to aggregate information and to reveal the true state. One important component of these studies is the speed at which traders with either overlapping or hierarchical information learn from each other. We complement this line of work by studying situations without overlapping or hierarchical information and without fully informed parties; in our setup, in finite time, no trader in the market will be able to learn the value of the state. Biais, Hilton, Mazurier and Pouget (2005) identify the large extent to which subjects tend to overestimate the precision of their information. Applied to our framework, such overestimating tendencies would weaken herding instincts and reinforce contrarian behavior.

Overview. The rest of the paper is structured as follows. The next section examines the theoretical framework, discusses the modifications undertaken to better fit a laboratory experiment, and outlines the hypotheses we seek to evaluate. Section 2 describes the experimental design, the subject pool, and the various treatments. Section 3 provides a preliminary, non-structural exploration of the data and its relation to key observable variables. Section 4 describes the overall fit of our data to the theoretical model. Section 5 looks in detail at herding and contrarianism. Section 6 studies the behavior of monotonic signal types. Section 7 studies whether there are differences in behavior with respect to herding and contrarianism between more and less rational types. Section 8 summarizes the key findings and concludes. The subject instructions and other materials are in the supporting appendix. We have also examined numerous behavioral explanations for our subjects' behavior. However, as these behavioral modifications did not add significantly to performance and insight, we relegated them to the supporting appendix.

1 Definitions, Theoretical Results, and Predictions

1.1 Formal Definition of Herding

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 will define herding and contrarianism against this yardstick. Moreover, the benchmark decision for a herder or contrarian is the action that they would take without observing any of the prices.

We thus say that a 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. For the formal definition we use H_t for the trading history at time t; this history includes all past actions, their timing, and the transaction prices; H_1 is the initial history.

Definition (Herding and Contrarianism)
HERDING. *A trader engages in herd-buying in period t after history H_t if and only if (i) he sells at the initial history, (ii) buys at history H_t, and $(iii-h)$ prices at H_t are higher than at H_1. Sell herding is defined analogously.*
CONTRARIANISM. *A trader is a buy contrarian in period t after history H_t if and only if (i) he sells at the initial history, (ii) buys at history H_t, and $(iii-c)$ prices at H_t are lower than at H_1. Contrarian-selling is defined analogously.*

Both with buy herding and buy contrarianism, a trader prefers to sell at the initial history, before observing other traders' actions (condition (i)), but prefers to buy after observing the history H_t (condition (ii)). The key differences between herding and contrarianism are conditions (*iii-h*) and (*iii-c*): the former ensures that the change of action from selling to buying is *with* the general movement of the crowd. The latter condition, requires the price to have dropped so that a buyer acts *against* the movement of prices. Thus there is a symmetry in the definitions, making herding the intuitive counterpart to contrarianism.

In the literature, there are other definitions of herding (and informational cascades). The definitions of herding and contrarianism that we adopt here are analogous to those in Avery and Zemsky (1998) and Park and Sabourian (2008) (which we implement) and capture the history-dependent (or social learning) element of behavior in an informationally efficient financial market.

1.2 The Underlying Theory

The model that underlies our experiment is an adaptation of a Glosten and Milgrom (1985) style sequential trading model in the following way. There is a single security that takes one of three possible liquidation values, $V_1 < V_2 < V_3$, each equally likely. Traders arrive in a random sequence and trade a security with an uninformed market maker. Before meeting a trader, the market maker sets a single price at which he is willing to buy or sell one unit of the security. Every trading slot is designated to a noise trader with a fixed probability (25% in our experimental setting) who buys or sells with equal chance. The remaining traders are informed and receive one of three signals, S_1, S_2, S_3. Signal S_1 is generated with higher probability in state V_1 than V_2, and likewise in state V_2 vs. state V_3. The reverse holds for signal S_3. This implies that the recipient of signal S_1 shifts probability weight towards the lowest state (S_1 is 'bad news'), whereas the recipient of S_3 shifts weight towards the highest state (S_3 is 'good news'). Signal S_2 can take several different shapes which we will outline shortly.

All past trades and prices are public information. The market maker follows a simple pricing rule by setting the unique trading price as the expectation of the true value of the security, conditional on all publicly available past information. Traders buy if their expectation conditional on their private signal and on the information derived from past trades exceeds the price, and they sell if this expectation is below the price.

The following are the possible shapes of conditional signal distributions (csd):

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

A signal is called *csd-monotonic* if its csd is either increasing or decreasing. If $\mathsf{Pr}(S|V_1) > \mathsf{Pr}(S|V_3)$ or $\mathsf{Pr}(S|V_1) < \mathsf{Pr}(S|V_3)$ holds then signals are negatively and positively biased respectively. A negatively biased U shaped csd is referred to as a negative U shape, likewise for hill shape and positive biases. More generally we refer to hill and U shaped csds as *non-monotonic*. The following is a corollary to the combination of Theorems 1 and 3 in Park-Sabourian. ('S herds' is to be read as 'S herds with positive probability'.)

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

(*a*) *A signal type with a decreasing csd always sells.*

(*b*) *A signal type with an increasing csd always buys.*

(c) *A signal type buy- (sell-) herds if and only if his csd is negative (positive) U shaped.*

(d) *A signal type acts as a buy-(sell-) contrarian if and only if his csd is negative (positive) hill shaped.*

A buy-herding trader would elect to sell at the initial history (i). Since the prior is uniform, this implies that sellers attach more weight to the lowest than the highest state, i.e. $\mathsf{Pr}(S|V_1) > \mathsf{Pr}(S|V_3)$. Buy-herding requires that prices have increased ($iii - h$); this occurs if and only if the probability of the lowest is smaller than that of the highest state.

Sufficiency can most easily be seen 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 a type with an increasing signal but such a 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. The argument for necessity is more involved and we refer the reader to the original Park and Sabourian (2008).

Adjustments for the Experiment. The full underlying theory in Avery and Zemsky and Park and Sabourian has two prices: one at which the market maker buys (the bid) and one at which the market makers sells (the ask). We dispense with bid- and ask-prices and focus instead on a single trading price, as is standard in the related experimental literature. With only a single price, the sufficient conditions from Park and Sabourian for herding and contrarianism are also necessary. While using bid- and ask-prices may seem to provide a better fit with the underlying theory, their use would generate a host of complications. Most obviously, participants need to understand the difference between the two prices. Moreover, it may lead to people focusing (subconsciously even) on only one side of the market. Thus people with initially negative information may follow only the movement of the bid-price, disregarding the possibility of buying completely. As Cipriani and Guarino (2005) and Drehmann et al. (2005) highlight, bid-ask treatments often do not offer additional insights.

1.3 Theoretical Predictions for the Experiment

Under rationality, traders should buy if their conditional expectation exceeds the price and sell otherwise. Traders can also choose not to trade. Such a 'pass' could only be optimal when the price to buy (the 'ask') differs from the price to sell (the 'bid') and when the expectation of a subject is between these two prices. In a setup without a bid-ask-spread as in our experiment, however, such a situation cannot arise (but for degenerate cases when the expectation coincides with the price). Thus we have

Hypothesis 1 (No passes) *Subjects will never pass.*

Given the proximity of the design to the theoretical model, further signal-specific predictions arise immediately: (1) signal types with a monotonically decreasing csd (such as S_1) should always sell; (2) signal types with a monotonically increasing csd (such as S_3) should always buy; (3) the behavior of signal types with a non-monotonic csd (hill or U shape) will depend on the trading history.

In the experiment we know the outcomes of the random elements (noise trades, the exogenous ordering and the signals for each subject). Thus, conditional on all other subjects behaving in accordance with the theory, we can calculate the action that each subject should undertake given history H_t and signal S.

Hypothesis 2 (Adherence to the Theory) *Subjects will act as prescribed by the theory: they buy if their theoretic expectation, conditional on their signal and the trading history, exceeds the price, and they sell otherwise.*

The experiment is implemented using a computerized price setting rule that assumes that subjects act in accordance with Hypothesis 2. 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).[3] Notably, no-trades in our setting do not affect the price, which is synonymous to a no-trade not revealing information. Using any other updating rule would have required us to speculate ex ante that a particular type is more prone to not trading than another.

One particular focus of this study is to test for and understand herding and contrarian behavior. The theory here predicts that types with U and hill shaped signals may exhibit such behavior.

[3]Although even ex ante it is foreseeable that some subjects do not act according to the theory, any other pricing rule would require that the underlying algorithm either uses "inside, non-public information" (such as the signal) or that it speculates on a particular bias that subjects may exhibit. Our mechanical price setting rule should thus be seen as the unique "rational" yardstick against which we measure behavior. In the supplementary appendix we also explore whether alternative updating rules, which may include accounting for others' irrationality, may do a better job at explaining the data.

Hypothesis 3 (Rational Herding and Contrarianism)
CONTRARIANISM: *Subjects who receive a signal with a negative (positive) hill shaped csd will act as buy (sell) contrarians when the theory predicts and sell (buy) otherwise. Specifically, they may act as buy (sell) contrarians when prices fall (rise), and they sell (buy) whenever prices rise (fall).*
HERDING: *Subjects who receive a signal with a negative (positive) U shaped csd will act as buy (sell) herders when the theory predicts and sell (buy) otherwise; specifically, they may buy (sell) herd when the price rises (falls), but sell (buy) whenever the price falls (rises).*

Monotonic types cannot *rationally* herd because their expectation will always be either below or above the price and cannot switch. However, monotonic and hill shaped types can herd in the sense of our definition, and monotonic and U shaped types can act as contrarians. What is required is that prices fall for buy contrarian behavior and sell herding and that prices rise for buy herding and sell contrarian behavior.

Being an experiment we do not anticipate a precise mapping of the theory to observed behavior and we shall compare the fit of the theory with that found in the literature (e.g. Anderson and Holt (1997), Drehmann et al. (2005), or Cipriani and Guarino (2005)).

A weaker test of our theory is whether herding and contrarianism is most prevalent for types with U and hill shaped signals respectively. One has to be careful to use the right metric to properly account for differences in the number of signal types and in the number of opportunities that a class of signal types has to exhibit herding or contrarian behavior. We will thus examine whether the role attributed to underlying signals by the theory seems justified. We do this by estimating the link between private signals and occurrences of herding and contrarianism.

Hypothesis 4 (The Impact of Signals on Herding and Contrarian Behavior)
When herding arises, it is most likely to stem from someone with a U shaped signal. When contrarianism arises, it is most likely to stem from someone with a hill shaped signal.

Prior experimental work by Drehmann et al. (2005) and Cipriani and Guarino (2005) has identified that traders exhibit a "natural" tendency to act as contrarians. Traders in these papers had the equivalent of monotonic signals. Anticipating that we may observe some irrational decisions among the monotonic signal types, we shall strive to understand their behavior better. If they tend to act as contrarians, then as prices rise, traders with an increasing csd sell, and, likewise, as prices fall traders with a decreasing csd buy. If instead they tend to act as herders, then as prices rise, traders with an decreasing csd buy, and, likewise, as prices fall traders with a increasing csd sell. Either type of behavior contradicts the theory.

Hypothesis 5 (Price Impact) *Price changes do not affect the decisions of monotonic signal types.*

2 Experimental Design

In addition to the information given in what follows, the supplementary appendix contains a time-line of events, the instructions and materials given to subjects, and a description of the purpose-built software used in this experiment.

2.1 Overview

The financial asset in every treatment can take one of three possible liquidation values $V \in \{75, 100, 125\}$ which correspond to the true value of the asset. The traders were typically made up of 15-25 experimental subjects, plus a further 25% noise traders, with a central computer acting as the market maker.[4] The existence and proportion of noise traders was made known to the experimental subjects in advance. It was also mentioned that noise traders randomized 50:50 between buying and selling.

Prior to each treatment subjects were provided with an information sheet detailing the prior probability of each state, 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. We thus provided subjects with both the signal distribution and the initial posterior distribution for each signal. The information on the sheets was common knowledge to all subjects. After being given the opportunity to study this information, each subject i received an informative private signal, described to them as a "broker's tip", $S_i \in S \equiv \{S_1, S_2, S_3\}$. The subjects were *not* told anything about the implications of non-monotonic or monotonic information structures or the predictions of the theory.

The nature of the compensation system was also made clear in advance, and in particular that it directly implied that they should attempt to make the highest possible virtual profit in each round, since the final compensation was based on overall performance (up to C$30) combined with a one-off participation fee (C$15; equivalent amounts were paid in the UK at 2006 exchange rates).

Prior to the start of trading, subjects were allocated time with pen and paper to contemplate their own signal and their information about the signal distribution and the

[4] Noise traders play an important role in the theory and add only a mild degree of extra complexity to the experimental design. They also have a useful practical role in the experiment, simulating a degree of uncertainty about the usefulness of any observed actions. Generally, noise traders reduce the informativeness of any observed action and we analyzed (details are in the supplementary appendix) how a model of errors as an incorrect assessment of the degree of noise trading performs with the data.

prior probabilities that the asset value was high, medium or low.[5] When trading began, trades in each treatment were organized sequentially: each subject or noise trader was assigned a time interval in which they, and only they, could act.[6] In practice, subjects tended to act before the allotted time was over with very few timeouts. A sequence of trading opportunities $t = 1, 2, 3, \ldots$ produced a history of actions and prices, $H_t = \{(a_1, P_1), \ldots, (a_{t-1}, P_{t-1})\}$ with $H_1 = \emptyset$. During the experiment itself, trading was anonymous and all price movements were clearly visible in real time on the computer screen. Specifically, subjects were shown the history in the form of a continuously updating price chart, and the screen also listed the current price, P_t, with $P_1 = 100$.

Subjects were told that they had three possible actions $a \in \{\text{sell, pass, buy}\}$ that they could undertake in their trading opportunity time window, t^*. They were also told that the time of their opportunity was exogenous, randomly determined and unique to them.

It was stressed to the subjects that their virtual profits per treatment were generated based on the difference between the price at which they traded, P_{t^*} where t^* is the time of their personal trading opportunity, and the true value of the share, V. It was also emphasized that the price that prevails at the end of trading does not affect their payoff (unless this was the price at which they traded).

The subjects were recruited from the Universities of Toronto, Cambridge and Warwick. No one was allowed to participate twice. We ran 13 sessions: 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 for the Warwick sessions: of the subjects there, around 49% were female, around 73% were studying (or had taken) degrees in Economics, Finance, Business, Statistics, Management or Mathematics. 53% claimed to have some prior experience with financial markets, with some 23% owning shares at some point in the past.

2.2 Treatments

Following Section 1, the rational action for recipients of signals S_1 and decreasing S_2 is to sell and for S_3 and increasing S_2 is to buy, irrespective of H_t. For the recipients of non-monotonic S_2 signals the nature of H_t and the precise information structure determined a

[5]The subjects could use the pen and paper as they wished, for example, to make notes or calculations. In the post-experimental questionnaire detailed below around 24% of subjects reported that they carried out numerical calculations.

[6]Subjects were called upon to act by a clearly visible notification on their computer screen. The combination of time in advance of trading (after signals were revealed) and a limited time window of action during trading was designed to match the theoretical model and to capture the spirit of real-world time-limited trading opportunities. It also dealt with the practical time limitations of the experimental sessions.

unique optimal action. This action might be to herd or to act in a contrarian manner. The treatments were each designed to enable us to examine a specific information structure with respect to signal S_2, specified as follows:[7]

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

There was also a training treatment, during which the subjects could practise using the software. This round was not part of the payment calculations, or the results.

3 Exploratory Analysis

Before we begin with a formal analysis of the validity of the theoretical model, we will first provide an exploratory, non-structural analysis of the data. Our aim is to get a sense for the main factors that affect behavior. We begin with a regression in which we seek to determine the likelihood of a buy given certain stylized model variables. The dependent variable, T_j is a dummy which takes value 0 if trade j is a sale and 1 if it is a buy or a pass.[8] The independent variables $X_{i,j}$, $i = 1, \ldots, 15$, and their construction are outlined in the first and second column of Table 1. We estimated a simple linear model

$$T_j = \alpha_j + \beta_1 X_{1,j} + \ldots + \beta_{15} X_{15,j} + \epsilon_j.$$

Several of the dependent variables are, of course, correlated. We shall thus refrain from over-emphasizing the findings in this section but instead treat them as merely indicative. Overall, we want to obtain a first glimpse of the factors that affect behavior and are thus interested merely in the sign and significance of variables.

We estimated the model by Logit and linear regression, and also controlled for subject fixed effects, but since the parameter estimates have the same significance and signs across all methods, we report only the marginal effects at the mean from the Logit regression. We

[7]Note that we employed two negative hill shaped treatments. In the first, Treatment 1, by design there were more candidates for buying, in the second, Treatment 6, there were more candidates for selling. The purpose of this design was to see if hill shaped types might be irrationally prone to herd behavior.

[8]Our choice of handling passes as buys will become clear in the next section. Yet we also ran two other types of regressions: in the first, all passes were ignored; in the second, the dummy took a value of 2 if there was a buy, and a value of 1 if there was a pass. The results are qualitatively similar and thus omitted.

ran the regressions on the entire data set, separate for each signal type (S_1, S_2 and S_3), and separately by the shape of the distribution (monotonic increasing, monotonic decreasing and the two types of non-monotonic hill and U shape).

Exploratory Observation 1: As Table 1 shows, on the entire data set the significant variables that determine the probability of a buy are (1) the price change, (2) the decision taken on the basis of the initial expectation, (3) the decision taken on the basis of the initial history, (4) signal S_3, (5) a hill shaped signal and (6) a negative U shaped signal.

The first noteworthy implication is that the *expected payoff from buying* plays no role in the decision to buy. For monotonic signal types this payoff, assuming rationality, is either always positive or always negative. So for these types the payoff should indeed not matter. For the non-monotonic signal types, however, the sign of the payoff may change and thus the variable should have an impact on their decision to buy or sell. Yet Table 1 shows that it does not, which indicates that we may find less herding and contrarianism than predicted by the theory.

Next, we also observe that the probability to buy decreases as the *price* increases. While this is a most intuitive outcome in a standard economic model, here there are many signal types who should not change their behavior when prices change. Moreover, since we did not include a positive hill shaped treatment, there is no type that should rationally switch from buying to selling as prices rise. Thus this observation indicates that we may find a substantial number of incidences of (irrational) contrarianism. One general observation in our experiments is that prices increased more often than they decreased (about 73% of trading prices exceeded 100, the initial price).[9] For increasing prices, negative hill shaped types should never be buying. Thus it is puzzling that this signal stands out as a significant variable to indicate that someone is buying.

Table 1 also indicates that the *initial expectation* plays an important role in the subjects' decisions. Combined with our observation on the possibility of contrarianism, this hints at subjects updating their expectations slower than mandated by the theory.

Exploratory Observation 2: Turning to the regressions restricted by signal types, we observe that for the S_1 type or, more generally, for the decreasing and increasing signal types none of the variables are significant that were significant on the entire data set. This is good news because these types should never buy. For the S_3 type on the other hand, we observe that as the price rises this type becomes less likely to buy. The same does

[9] In contrast to two state trading models, a buy changes the price usually by a different amount than a sale. Our setup included more traders that were leaning towards selling rather than buying. Ignoring history induced switches, there were 819 traders who were leaning towards buying (233 noise buyers, 407 S_3 signals, 84 with an increasing S_2, and 95 with a positive U shaped signal S_2), compared to 1021 traders who were more inclined to sell (290 noise sellers, 390 S_1 signals, 97 decreasing S_2 signals, 180 negative hill shaped S_2 signals, and 94 negative U shaped signals).

not hold, however, when we combine the S_3 with the increasing S_2. For the S_2 types as a whole, we observe significant impacts of the same variables that were relevant for the entire data set. For the non-monotonic types alone, however, we observe that the decision based on the initial expectation is the only significant variable.

Thus the S_3 types are quite prone to act as contrarians.[10] When looking at the more focused regressions that consider only the trading variables as displayed in Table 3, it turns out that S_3 type's decision to buy is influenced significantly by the price and by the initial expectation. Yet it is puzzling that the significance of the price change vanishes if we look at monotonically increasing csd types jointly. The S_3 type's unwillingness to buy combined with the significance of the variable *prior expectation* $-$ *price*$_t$ indicates the tendency to update slowly. Yet the same should apply to monotonic S_2 types. Moreover, the latter's expectation is below that of the S_3 type, so if anything then the monotonic S_2 types should display more contrarian behavior, and not less.

Exploratory Observation 3: Looking at the more focused Table 3, we observe that prices negatively affect the probability of buying for S_2 types. Moreover, the relation of the current price to the initial expectation is a good predictor for the decision taken by an S_2 type. To single out which variable is more important for the S_2 type, we ran regressions excluding either the initial decision dummy and the initial expectation-based dummy. In turns out that the total price change is unimportant but instead that the level of the price compared to the initial expectation matters. (Any significant effect of the price change is captured in magnitude by the other two variables.) For non-monotonic signal types, the only highly significant variable is *prior expectation* $-$ *price*$_t$. For monotonic types, no variable is significant, and as mentioned before for the S_3 types alone the price level has a significant, negative effect.

Summary of Exploratory Observations: We observe (1) the sign of expected payoffs plays no role in the decision to buy, (2) S_2 and S_3 types are price sensitive and less likely to purchase when prices rise, (3) we expect (irrational) contrarianism, in particular by S_3 types, and we expect S_2 types to act as contrarians for lower prices than S_3 types.

[10] For this type the location dummy is also significant, which indicates that the S_3 types in Warwick and Cambridge were more likely to buy than those in Toronto. This hints at the possibility that contrarianism of S_3 types may be more pronounced for the Torontonian participants. Ran as a separate test, we indeed observed that S_3 types in Toronto are 12% more likely to act as contrarians relative to participants at the other locations.

4 Analysis of the Rational Benchmark

In the numbers to follow we report only trades by human subjects and exclude noise trades. The total number of trades was 1375, spread over all 6 treatments. We recorded 28 time-outs, leaving 1347 recognized trades. Time-outs will henceforth be omitted from the analysis.[11] The number of trades were 390 S_1, 550 S_2 and 407 S_3.[12]

4.1 The Decision to Pass

We admitted the option to pass for several reasons. First, as subjects had to take their decision in a limited time frame, we had to include either the explicit decision to pass or the decision to allow the clock to run out. Including passing as an option allowed us to count a pass as a deliberate action and to distinguish them from accidental timeouts.

Second, the structure of our setup lends meaning to passes. Traders are owners of a share and they have the choice to buy an extra share, or to sell the share that they already own. Our rationale for providing them with an endowment was twofold. First, it allowed us to avoid explaining 'short-selling'. Second, as extant owners any decision has direct payoff consequences. Namely, by electing to pass and thus to retain their share subjects have de facto decided that the share is worth more to them than the current price. By allocating a share endowment and giving subjects the ability to pass we thus made sure that any action had payoff consequences.[13] This contrasts with situations without endowments where a passing decision has no tangible cost and not trading is a risk-free action. In this sense, decisions in our experimental setup mimic the choices of investors, who usually hold positions, as opposed to speculators, who hold no position and just go long or short for a limited time.

In summary, when traders are owners, passing implies that they hold on to that share, presumably in hope of making a profit on that one share. In this sense, a hold is a positive

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

[12] The subjects were asked to comment on their own actions in a questionnaire (provided in the appendix) at the end of each session. 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.

[13] Without the ability to pass subjects could still timeout. This might appear synonymous with passing, but it lacks payoff consequences in the absence of share endowments. By enabling a "pass" decision we focus the mind on payoff relevance as well as differentiating between active passes and passive timeouts.

signal albeit weaker than a buy. Therefore, a pass can be counted as a "weak buy".

As outlined in Section 1 passes contradict the theoretical model and thus Hypothesis 2. We will use the passing decision in three ways: first, we count any pass as an irrational action. Second, we count the decision to pass as a weak buy and thus classify a decision to pass as rational when the theory predicts that the trader should buy. Third, we use the decision to pass as a classification tool of more vs. less rational traders and check if the behavior of these two groups differs.

Overall there were 145 passes (10.7% of all trades), 31 from S_1 types (8%); 87 from S_2 types (16%) and 27 from S_3 types (7%); 56% of the subjects (128 out of 230) never pass. While Hypothesis 1, the strongest interpretation of the theoretical model predicts that we should see no passes at all, we do see some. One explanation for the presence of passes could be risk aversion. We discuss this interpretation in the supplementary appendix: in short, including risk aversion does not explain the data better at all.

The total number of passes is small, especially for the S_1 and S_3 types. However, the figure of 16% for S_2 types indicates that there is cause for some doubt about Hypothesis 1 (no passes) from those traders with the middle signal.

Finding 1 (Passes) *About 10% of trades were passes, contradicting Hypothesis 1. About 44% of the subjects pass at least once.*

4.2 Fit of the data to the rational model

We start with a rough overview of decisions aggregated over all treatments that are in line with rationality; Table 4 displays the data.

About 69.5% of trades are in accordance with the theoretical model when counting passes as categorically incorrect. If we admit passes as "weak buys", as outlined in the last subsection, then all passes by S_1 types remain irrational, whereas all passes by S_3 types are admitted as rational. For the S_2 types, passes are admitted as rational whenever the rational action was to 'buy'. With this specification, the overall model fit is 74.6%.

These numbers are similar to those in Cipriani and Guarino (2005) who obtain 73% rationality.[14] This similarity is noteworthy because our setting is more complex, particularly for the S_2 types. Moreover, Cipriani and Guarino's experiment effectively considers only types that are equivalent to our S_1 and S_3 types, and these types actually performed better in our setup, with rationality in excess of 80%. We might thus reasonably argue that the S_1 and S_3 types are acting in accordance with the rational theory.

[14]Anderson and Holt (1997) have 70% rationality, albeit with a fixed-price setting.

The S_2 types, on the other hand, often do not act rationally. As Table 4 illustrates, almost half of their trades were against the rational model. This holds for both monotonic and non-monotonic signals. In the herding Treatments 3 and 5, the S_2 types perform quite poorly, even when admitting passes as weak buys (22% and 37% fit). Had they taken each action at random they would have done better.[15]

At the same time, the non-monotonic types face a more difficult decision problem than the monotonic types. Theoretically, the decisions of monotonic S_1 and S_3 types never change, so they can take the correct decision even without following the history. The non-monotonic types on the other hand, have to follow the history carefully and small mis-computations can cause them to be categorized as irrational. Yet the fit is also low for increasing S_2 types, even though their decision problem is similar to that of S_3 types (they should always buy).[16]

Finding 2 (Rationality) *69% of trades conformed to the rational choice. This is in line with other experimental studies, thus overall we do not reject Hypothesis 2. In particular, for S_1 and S_3 types this figure exceeds 80% so that for these types Hypothesis 2 should not be rejected. Similarly, for the decreasing S_2 and hill shaped types, rationality is large (combined: 68%). Hypothesis 2 should also not be rejected for increasing S_2 when passes count as weak buys. But Hypothesis 2 is not warranted for U shaped types as well as increasing S_2 types when passes are wrong (combined fit 38%) as well as for U shaped types when passes are weak buys (fit 46%).*

5 Herding and Contrarianism

All signal types can herd or act as contrarians in the sense of our definition. For instance, an S_1 type who buys after prices have risen would engage in buy herding. Yet only U shaped types can *rationally* herd as only their initial expectation can be below and their time t expectation be above the price after prices have risen. Similarly, only hill shaped types can rationally act as contrarians. Yet U shaped types can also *irrationally* herd in the sense of our definition. Such a situation arises when a negative U shaped type buys, prices have risen, but the type's theoretical expectation is still below the price. In other words, the requirements for rational herding and contrarianism are rather restrictive.

In what follows we will first focus on the rational case. Namely, since we know each subject's theoretical expectations at any stage, we know when herding or contrarian behavior is theoretically mandated. As before we can exclude or include passes as 'weak

[15] In the econometric analysis of the next section we will highlight the persistency in their behavior.

[16] Future work may seek to analyze if herding is more prevalent with multiple monotonic types.

buys'; we look at both cases separately. In a second step, we look at herding and contrarianism by all types and determine whether the hill and U shaped types stand out.

5.1 Rational Herding and Contrarianism

Rational herding can arise only in treatment 3 by S_2 types who have a negative U shaped signal and in treatment 5 by S_2 types who have a positive U shaped signal. Rational buy contrarianism can arise only in treatments 1 and 6 from recipients of the negative hill shaped signal.

When passes are considered irrational, only 18% of the herding trades that should occur do occur. When passes are weak buys, rational herding occurs in 31% of the mandated cases. Contrarian behavior arises in 62% of the cases when passes are wrong, 77% when passes are weak buys. The total number of possible rational contrarian trades is, however, rather small compared to the herding trades. Table 5 summarizes the data.

Broken up by treatments, the performance of herding candidates in the positive U shaped Treatment 5 is rather poor: 90% of the required herds did not occur. However, the number of observations is also rather small. The main reason for the small number of observations is that (unintendedly) the sequences of trader arrivals were such that there were rarely falling prices, which is the prerequisite for sell-herding.

The performance is somewhat better in the negative U shaped Treatment 3: counting passes as 'weak buys' the fraction of missing herds is 'only' 64%. However, this is a notably larger fraction of herds than observed in Drehmann et al. (2005) or Cipriani and Guarino (2005) (where herding behavior is irrational and rarely accounts for more than 10-20% of trades, and usually much less). This lends some support to the hypothesis that the U shaped signal structure matters.

Rational contrarian behavior has a better performance than herding, although the number of cases with theoretically mandated contrarianism is rather small, in particular for treatment 1. One obvious explanation for the better fit is that the hill shaped contrarian signal is much easier to interpret since it indicates that the true value is the middle one. Consequently, it is comparatively simple for subjects to pick an action that moves prices in the direction of this middle value.

Finding 3 (Rational Herding and Contrarianism) *Rational herding arises less frequently than predicted by the theory. Contrarianism arises in about 2/3 of the cases as predicted by the theory, yet the rationality of the S_2 contrarian types lags that of the S_1 and S_3 types with respect to rational contrarian actions.*

5.2 General Herding and Contrarianism: Summary Statistics

The rational model is strict and small computational errors by subjects may lead to decisions that are classified as incorrect. In what follows we will no longer restrict attention to *rational* herding and contrarian but instead consider all herding and contrarianism decisions as classified by our definition.

As outlined before, all signal types may engage in herding or contrarian behavior. Moreover, for every trade there is either the potential for herding or for contrarian behavior. For instance, an S_3 type engages in herd selling if he sells when prices fell below 100. Similarly, while S_3 types cannot herd when prices go up, they can act as sell contrarians.

Our main concerns are two-fold. Firstly, we want to determine whether the non-monotonic S_2 types exhibit herding and contrarian behavior "in the right direction", that is, do they switch from selling to buying if prices rise (for buy herding) and do they switch from selling to buying if prices fall (for buy contrarian behavior)? Second, we want to find out if their herding and contrarian decisions by U and hill shaped signal types stand out relative to those by the other signal types.

Revisiting Herding and Contrarianism by Trader Type. There was almost no herding from S_3 types which is as predicted by the theory. However these types exhibited persistent contrarian behavior. Across all treatments prices typically rose over time, which results in public behavior confirming S_3 signals. Also, when prices rise, S_3 types cannot herd by definition. The total number of times that S_3 types could herd was very small. As can be seen from Table 6, herding is possible for only 82 (or 20%) of the 407 S_3 trades.

The S_1 types tentatively exhibited contrarian behavior, especially in Treatment 6 when prices fell. There was also some herding behavior by S_1 types across all treatments.

For U shaped S_2 types (who can theoretically herd or act in a contrarian manner) there is evidence that herding occurs (in Treatment 3 in particular), but there was little to no observed herding in Treatment 5 when prices fell.

Similarly, for hill shaped S_2 there is evidence of contrarian behavior in Treatment 6. The numbers for Treatment 1 are too small to allow a definitive judgement.

Interestingly, monotonic S_2 types, who were present in Treatments 2 and 4, also showed a tendency to engage in both herding and contrarianism, even though theoretically they should not engage in either. These monotonic signals were weaker versions of the S_1 and S_3 signals. We will study these types' behavior in more details later on.

Remark (Herding and Contrarianism by all Types) *In the summary data, the S_2 types are much more likely to herd or act as contrarians than the S_1 types or the S_3 types.*

5.3 General Herding and Contrarianism: Regression Analysis

We now directly test the link between herding (U shaped) and contrarian (hill shaped) signals and incidences of herding and contrarian trades. In particular, we ask:

(1) If someone has a herding (U shaped) signal, is this person more likely to herd than someone who has any other type of signal?

(2) If someone has a contrarian (hill shaped) signal, is this person more likely to act as a contrarian than someone who has any other type of signal?

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

$$\mathsf{herd}_{i,t} = \alpha + \beta \mathsf{u\text{-}shape}_{i,t} + \epsilon_{i,t}, \quad \mathsf{contra}_{i,t} = \alpha + \beta \mathsf{hill\ shape}_{i,t} + \epsilon_{i,t} \tag{1}$$

where the dependent variables $\mathsf{herd}_{i,t}$ and $\mathsf{contra}_{i,t}$ are dummies that apply our definition in that they are 1 if individual i herds or acts as a contrarian respectively at trade t and 0 otherwise, α is a constant, and $\mathsf{u\text{-}shape}_{i,t}$ and $\mathsf{hill\ shape}_{i,t}$ are signal dummies that are 1 if the individual received a U shaped (for herding) or hill shaped (for contrarianism) signal. Given the random assignment of signals and time slots, we can assume that $\mathsf{E}[\mathsf{u\text{-}shape}_{i,t} \cdot \epsilon_{i,t}] = 0$ and $\mathsf{E}[\mathsf{hill\ shape}_{i,t} \cdot \epsilon_{i,t}] = 0$, the main identifying assumption.

We estimated the model by Logit and report the marginal effects at the mean.[17] Standard errors are clustered by subjects' respective sessions to control for inter-group correlations.[18]

Herding and U shaped signals. In this specification, β represents the impact of the signal on a subject's choice of whether or not to herd. It is the main coefficient of interest and should be positive because, according to the theory, the U shaped signal should increase the probability of acting as a herder.

[17] We also ran several alternative unreported specifications: a linear regression, a linear regression controlling for trader fixed effects, and a linear regression controlling for group-level fixed effects (a group is the collection of subjects in one of our 13 sessions). The estimation was less precise for the fixed effects regressions, largely because the group-level clustering of standard errors caused a significant reduction in degrees of freedom. Yet the estimates were qualitatively similar irrespective of the estimation method and we thus only report the Logit results.

[18] In principle, actions are not independent because a trader at time t can only buy herd if sufficiently many other traders before have bought shares so that prices have risen. However, in our estimation we condition on the possibility of herding or contrarianism. Since a herding or contrarian action is not observable (only buys, sales and no trades are), herding decisions should objectively be conditionally independent. Thus our approach is possibly overly cautious and if anything our use of clustered standard errors may slightly underplay the significance of our estimates.

In line with our exposition thus far, we distinguish the case where passes are categorically irrational from the case where passes count as weak buys. The estimation is restricted to the cases where herding is possible. This restriction is reasonable because, 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.

Table 7 summarizes the results from our regression. Overall, for the case where passes are categorically wrong, obtaining a U shaped signal increases the probability of herding by about 5.8% relative to any other signal and it is significant at the 1.7% level. When counting passes as weak buys, the marginal impact of the signal value increases to 8.1%. Overall the estimation confirms the hypothesis that recipients of U shaped, herding signals are generally more likely to herd, providing support for Hypothesis 4.

Contrarianism and Hill Shaped Signals. Next, we estimate equation (1) to test the hypothesis that a hill shaped signal is a significant cause for contrarian behavior. Our theory predicts that the coefficient β is positive so that a hill shaped signal has a larger impact on the occurrence of contrarianism relative to other kinds of signals. As with herding, we restrict to the cases where it is possible that traders act as contrarians.

Table 7 summarizes the results from our regression. Obtaining a hill shaped signal increases the chance of acting as a contrarian by about 33.7% relative to any other kind of signal. When admitting passes as weak buys, the marginal effect increases to about 50%. These coefficients are significantly different from zero at all conventional levels. Overall we confirm the insights from the preceding sections that the impact of the contrarian signal is stronger and the theory is more reliable in yielding predictions. Again, Hypothesis 4 gains support.

Finding 4 (Impact of the Information Structure) *The regression analysis indicates that receiving a U shaped signal significantly raises the probability of acting as a herder compared to obtaining any other signal (by 5.8%). Similarly, receiving a hill shaped signal significantly raises the probability of acting as a contrarian compared to obtaining any other signal (by 33.7%). Thus there is support for Hypothesis 4.*

6 Monotonic Types

While the general behavior of the S_1 and S_3 types is in line with the theoretical model (about 80% of their traders are 'rational'), we do observe that S_3 types engage in selling and that S_1 types engage in buying; similarly for monotonic S_2 types. We now want to assess whether this behavior is systematic. Specifically, we test whether an increase in the price changes the probability of a specific trade. Theoretically, the price should have no impact on the decision because S_1 traders should always sell, S_3 traders should always buy. We thus estimated the following regression

$$\mathsf{trade}_{i,t} = \alpha + \beta \Delta \mathsf{price}_{i,t} + \epsilon_t, \tag{2}$$

where is $\mathsf{trade}_{i,t}$ a dummy that is 1 if there is a buy or pass, and 0 when there is a sale,[19] and the independent variable $\Delta\mathsf{price}_{i,t}$ is the percentage change of the price from 100, i.e. the price at the time of the trade divided by 100 and subtracting 1. Since the time slot in which people are allowed to trade is assigned at random, we can assume that $\mathsf{E}[\Delta\mathsf{price}_{i,t} \cdot \epsilon_{i,t}] = 0$.

We estimated the model by Logit[20] separately for the signals S_1, S_3, increasing S_2, decreasing S_2, all increasing together, and all decreasing together. The main variable of interest in (2) is β which measures whether a rising price affects the probability of a trader buying or selling. Our theory (formalized in Hypothesis 2) predicts that the price should have no impact on whether any of the types under consideration buys or sells. Consequently, parameter β should be insignificantly different from zero. In contrast, if it is not zero, then we gain insights about systematic herding or contrarian behavior. For instance, consider type S_3. If the sign of β is negative, then this type becomes less likely to buy as prices increase. Such behavior tentatively indicates systematic contrarian behavior. Likewise, if β were positive for the S_1 types, then this implies that the S_1 types are more likely to buy when prices rise; this is a tentative herding effect.

Table 8 summarizes the results of our estimation. We find that for the S_3 types a 1% increase in the price from the original level lowers the probability of a buy by 0.5%. This estimate is significantly different from zero at the 1% level. Thus as prices increase, the S_3 types become more likely to sell — which confirms their tendency to act as contrarians. The coefficients for all other types individually are insignificantly different from zero. Increasing types together have a significant negative coefficient, which is driven by the S_3 types (the coefficient has almost the same magnitude as for the S_3 alone). Decreasing types together have a weakly negative coefficient of similar magnitude as the increasing types. While it appears that we should get symmetric results for the S_1 and the S_3 types, for this we must have sufficiently many incidences of falling prices — which we do not have.[21] We thus conclude that there is enough evidence for contrarian behavior of the S_3 types and we have no evidence for herding.

[19] We also estimated specifications in which the dummy takes value 2 for a buy and 1 for a pass, and one where passes are ignored altogether. The findings coincide qualitatively for all specifications.

[20] Ordinary least squares regressions with and without trader fixed effects also yield the same insights.

[21] In some unreported regressions, in which we restrict attention only to those cases where prices fell, there is some weak evidence (significant at the 10% level) that the coefficient sign is negative. This, again, is evidence in favor of contrarianism.

Finding 5 (Price Changes and the Decisions of Types with Monotonic Csds) *Theory predicts that the price does not affect the decision of signal types with monotonic csds. Our regressions show that as prices increase, the* S_3 *types become less likely to buy. The sign and magnitude of the effect is similar for all other monotonic signal types. Statistically significance obtains for signal* S_3*, all increasing types taken together, and all decreasing types taken together.*

S_2 types who receive a monotonically increasing signal arguably receive a weaker version of the S_3 signal; similarly for a decreasing S_2, which is a weaker negative signal than S_1. It is thus curious that S_3 types display a reaction towards the price whereas increasing S_2 types do not. That being said, the coefficient estimate has the same magnitude but is insignificant due to the large standard error.

To complete the picture we thus re-ran the regressions in (1), but omitted all incidences of non-monotonic signals from the data. We observe that, when omitting non-monotonic signals, receiving an increasing S_2 signal significantly increases the probability of both acting as a herder (by about 10.4%) and as a contrarian (by about 17%) over all other (monotonic) signals. This bi-directional behavior thus explains the large standard error in the estimation of equation (2).

7 More vs. Less Rational Types

According to the rational theory the decision to pass is never optimal. Therefore, someone who passes can be considered to be less rational than somebody who does not. About 56% of the subjects (128 out of 230) never pass; the remaining subjects pass at least once.

We now want to analyze to what extent our estimates in Table 7 are affected by the less rational, "passing" types. We thus ask the following question: what is the probability that a subject herds/acts as a contrarian conditional on being a less rational type relative to the more rational types? To answer this question, we ran the following regressions

$$\begin{aligned} \mathsf{herd}_i &= \alpha + \beta_1 \mathsf{U\ shape}_i + \beta_2 \mathsf{passer}_i + \beta_3 \mathsf{U\ shape}_i \times \mathsf{passer}_i + \epsilon_i, \qquad (3) \\ \mathsf{contra}_i &= \alpha + \beta_1 \mathsf{hill\ shape}_i + \beta_2 \mathsf{passer}_i + \beta_3 \mathsf{hill\ shape}_i \times \mathsf{passer}_i + \epsilon_i. \qquad (4) \end{aligned}$$

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, passer_i is a dummy that takes value 1 if the trading subject has passed at least once and 0 otherwise, and $\mathsf{U\ shape}_i \times \mathsf{passer}_i$ and $\mathsf{hill\ shape}_i \times \mathsf{passer}_i$ 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 at the mean. As before, standard errors were clustered at the group level to correct for possible intergroup correlations. The coefficient β_1 allows us to estimate the marginal effect among more rational traders, coefficient β_3 allows us to estimate the differential marginal effect among less rational traders, so that $\beta_1 + \beta_3$ allows us to determine the effect of a signal among less rational traders.

Naturally, the regression is run only for the case where passes are *not* weak buys. We find that for those who pass at least once, neither receiving the U shape nor the hill shape signal affects their probability of engaging in herding and contrarian behavior respectively. For the types who never pass ("the more rational types"), on the other hand, the marginal effects of the signals are (much) stronger and significant. Table 9 summarizes the results of our estimation.[22]

Note however, that, omitting the cases where a pass occurs, subjects who pass do not make more irrational decisions than their counterparts (78% of both types' decisions are correct when ignoring the passing decisions themselves). The results here are thus not driven by the general (ir-)rationality of the passing types' actions.

Finding 6 (More vs. Fewer Rational Types) *Although the behavior of passers and non-passers is overall similar (ignoring the passes themselves), when looking at the decisions that may involve contrarianism or herding, the behavior of passers vs. non-passers is different: U shaped and hill shaped signals do not affect the passers' probabilities of engaging in herding an contrarianism whereas for non-passers the effect of these signals is significant and 75% stronger than in the general population of subjects (10.2% vs. 5.8%).*

8 Conclusion

The experimental study presented in this paper focuses on a multi-state multi-signal model that admits both herding and contrarian behavior as rational outcomes. It is the first study to allow both of these important phenomena to be analyzed together in the context of a rational model of financial market trading.

Our analysis can be loosely separated into two parts. In the first part, we directly test the theoretical findings on herding and contrarianism put forward by Avery and Zemsky (1998) and Park and Sabourian (2008). In the second part, we determine whether people

[22] Non-passers are also more likely to act in accordance with the theory when it comes to herding: they herd when they should in 24.2% compared to the passers who herd only in 12.8% of the relevant cases. All these findings indicate that future research into behavioral heterogeneity between traders, along the lines of Ivanov, Levin and Peck (2007), is warranted.

behave in the spirit of Avery and Zemsky and Park and Sabourian with respect to the effect of their information on their tendency to engage in herding or contrarian behavior.

Since we know all actions and signals, we can compute the theoretically optimal decision for each subject at any time, and we find that about 70% of all decisions are explained by the rational model. This figure is in line with the literature, even though subjects face a more difficult decision than in previous studies. However, herding often does not arise when it is theoretically predicted. Contrarianism is more prolific than herding, but arises both rationally and irrationally.

In the second step of our analysis we focus on understanding the link between information and observed herding and contrarian behavior. In principle, all signal types can irrationally herd or be contrarian. Therefore it is important to understand whether herding and contrarianism are observed equally across types, or whether they are more prevalent among certain types. Our results are in line with the theory in that herding is more commonly observed among types who have the theoretical potential to herd; similarly for contrarianism.

This finding has implications for the economic relevance and significance of social learning in financial markets. Ultimately, the economic importance of herding and contrarianism depends on the number of people who receive U and hill shaped signals.[23] If only .2% of traders obtain U shaped signals, then herding will likely have no impact, even if all recipients of such signals herd. If, on the other hand, 98% of traders receive U or hill shaped signals, then even the low probability of rational herding and the (somewhat higher) probability of contrarianism may lead to economically meaningful price swings.

From a policy perspective, we thus argue that herd (or contrarian) behavior that causes unusual, sudden or very rapid price movements, may well be driven by rational agents. Buy herding in the sense depicted by our definition would be detected by an apparent serial correlation or momentum in buy orders; contrarianism would manifest itself as mean reversion. Our analysis has shown that, for instance, strong behavior momentum such as a herd-selling frenzy can be attributed to *rational* behavior. Policy-makers should thus be careful not to dismiss herd-like behavior as irrational, nor to rule out the degree to which improved information-provision or clearer policies (that eliminate U shaped signals) might help to ameliorate the herding phenomenon.[24]

In the supplementary appendix we examined a host of alternative models of behavior

[23]Formally, their fractions in the population depend on the specific details of the signal distribution.

[24]The efficiency of the market price might also be of interest from a policy perspective. Using the price that would result if all traders in the model act as the theory predicts, we find that about 16% of observed prices coincide with this benchmark. In most treatments the trading histories are such that the theoretical model would predict rising prices, and indeed this is what happens. Yet on average observed prices are 7% below the rational benchmark, with a standard deviation of 9.9%.

(such as risk aversion, loss aversion or error correction). Yet the findings there left us with the impression that these specifications do not add to our understanding over and above the considerable success afforded by the rational model.

References

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

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

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

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

Biais, Bruno, Denis Hilton, Karine Mazurier, and Sbastien Pouget, "Judgemental Overconfidence, Self-Monitoring, and Trading Performance in an Experimental Financial Market," *The Review of Economic Studies*, 2005, *72* (2), 287–312.

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

Chordia, Tarun, Richard Roll, and Avanidhar Subrahmanyam, "Order imbalance, liquidity, and market returns," *Journal of Financial Economics*, 2002, *65* (1), 111–130.

Cipriani, Marco and Antonio Guarino, "Herd Behavior in a Laboratory Financial Market," *American Economic Review*, 2005, *95* (5), 1427–1443.

—— **and** —— , "Herd Behavior in Financial Markets: A Field Experiment with Financial Market Professionals," Mimeo, UCL and GWU June 2007.

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

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

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," Mimeo, forthcoming American Economic Review, The Ohio State University 2007.

Park, Andreas and Hamid Sabourian, "Herding and Contrarian Behavior in Financial Markets," Working Paper, Universities of Toronto and Cambridge 2008.

Plott, Charles R. and Shyam Sunder, "Rational Expectations and the Aggregation of Diverse Information in Laboratory Security Markets," *Econometrica*, 1988, *56* (5), 1085–1118.

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

Table 1
Exploratory Regressions.

The table represents the results from a logit regression of the occurrence of a buy or pass on all the explanatory data that was available to us. For all tables that follow, standard errors are in parentheses, * indicates significance at the 5% level, ** at the 1% level.

dependent variable	details	all traders	only S_1	only S_3	only S_2
profit from buying	$\mathsf{E}_t[V\|S] - p_t$	-0.004 (0.01)	0.004 (0.00)	0.004 (0.00)	-0.014 (0.01)
% price change	$(p_t - 100)/100$	-1.150** (0.34)	-0.326 (0.31)	-0.800** (0.27)	-1.148* (0.53)
decision taken on the basis of initial expectation − price	buy if $\mathsf{E}_1[V\|S] > p_t$, else sell	0.128* (0.06)	-0.055 (0.17)	-0.047 (0.04)	0.148 (0.09)
Initial decision	buy if $\mathsf{E}_1[V\|S] > p_1$, else sell	0.473** (0.10)			0.533** (0.12)
Net order flow scaled by volume	$(b_t - s_t)/(b_t + s_t)$	0.038 (0.07)	0.06 (0.07)	0.09 (0.09)	-0.028 (0.12)
Scaled Volume	$(b_t + s_t)/(\#\text{traders}/.75)$	-0.067 (0.07)	-0.111 (0.07)	0.019 (0.06)	-0.069 (0.10)
change from last price	$(p_t - p_{t-1})/p_{t-1}$	0.173 (0.61)	-0.209 (0.68)	0.384 (0.55)	0.308 (0.86)
Decision of predecessor	buy if predecessor bought, sell if predecessor sold	0.006 (0.06)	-0.01 (0.06)	-0.01 (0.05)	-0.003 (0.08)
Signal S_1		-0.088 (0.08)			
Signal S_3		0.345** (0.08)			
positive U shaped signal		-0.016 (0.08)			-0.015 (0.09)
negative U shaped signal		0.289** (0.09)			0.384** (0.11)
hill shaped signal		0.326** (0.08)			0.337** (0.09)
number of people in session		-0.005 (0.01)	-0.005 (0.01)	0	-0.004 (0.01)
location dummy	0 for Toronto, 1 for Cambridge, 2 for Warwick	0.01 (0.03)	-0.015 (0.03)	0.063** (0.02)	-0.044 (0.04)
Constant		-0.159 (0.13)	0.016 (0.15)	0.175 (0.10)	-0.148 (0.17)
Observations		1315	384	391	540

Table 2

Exploratory Regressions by Signal Shape.

The table represents the results from a logit regression of the occurrence of a buy or pass on all the explanatory data that was available to us. Non-monotonic signals are U- and hill shaped S_2, increasing signals are monotonically increasing S_2 (as in treatment 2) and signal S_3, decreasing signals are monotonically decreasing S_2 (as in treatment 4 and signal S_1. Compared to Table 1, some variables were dropped due to multi-collinearity. Standard errors and significance levels are denoted as in Table 1.

dependent variable	details	non-monotonic	increasing	decreasing
profit from buying	$\mathsf{E}_t[V\|S] - p_t$	-0.015 (0.01)	0.004 (0.01)	0.003 (0.00)
% price change	$(p_t - 100)/100$	-0.7 (0.65)	-0.935** (0.34)	-0.447 (0.30)
decision taken on the basis of initial expectation − price	buy if $\mathsf{E}_1[V\|S] > p_t$, else sell	0.360** (0.13)	-0.041 (0.04)	0.057 (0.09)
Initial decison	buy if $\mathsf{E}_1[V\|S] > p_1$, else sell	0.061 (0.10)		
Net order flow scaled by volume	$(b_t - s_t)/(b_t + s_t)$	-0.162 (0.16)	0.107 (0.11)	0.057 (0.07)
Scaled Volume	$(b_t + s_t)/(\#\text{traders}/.75)$	-0.086 (0.13)	0.029 (0.07)	-0.104 (0.07)
change from last price	$(p_t - p_{t-1})/p_{t-1}$	-0.298 (1.02)	0.594 (0.64)	0.003 (0.63)
Decision of predecessor	buy if predecessor bought, sell if predecessor sold	0.063 (0.10)	-0.025 (0.05)	-0.015 (0.06)
hill shaped signal		-0.105 (0.10)		
number of people in session		0 (0.01)	-0.003 (0.01)	-0.003 (0.01)
location dummy	0 for Toronto, 1 for Cambridge, 2 for Warwick	-0.045 (0.05)	0.068** (0.02)	-0.033 (0.03)
Signal S_3		0.171** (0.04)		
Signal S_1			-0.006 (0.06)	
Constant		0.117 (0.20)	0.117 (0.11)	0.004 (0.12)
Observations		369	475	471

Table 3

Exploratory Regressions with Dynamic Variables.

The table represents the results from a logit regression of the occurrence of a buy or pass on all the explanatory data that changes with trading, except the initial decision (for S_1 and S_3 this variable is always 0 and 1 respectively and was thus dropped). Non-monotonic signals are U- and hill shaped S_2, increasing signals are monotonically increasing S_2 (as in treatment 2) and signal S_3, decreasing signals are monotonically decreasing S_2 (as in treatment 4 and signal S_1. Standard errors and significance levels are denoted as in Table 1.

dependent variable	details	all traders	only S_1	only S_3		only S_2		non-monotonic	increasing	decreasing
profit from buying	$\mathsf{E}_t[V\|S] - p_t$	-0.002 (0.00)	0 (0.00)	-0.003 (0.00)	-0.001 (0.01)	-0.002 (0.01)	0.011 (0.01)	-0.009 (0.01)	0.006 (0.00)	0.001 (0.00)
% price change	$(p_t - 100)/100$	-0.640** (0.23)	-0.293 (0.22)	-0.686** (0.24)	-0.378 (0.34)	-1.114** (0.26)	-0.531 (0.33)	-1.090* (0.49)	-0.073 (0.24)	-0.322 (0.21)
decision taken on the basis of initial expectation − price	buy if $\mathsf{E}_1[V\|S] > p_t$, else sell	0.164** (0.06)	-0.034 (0.17)	-0.017 (0.04)	0.268** (0.08)		0.318** (0.08)	0.363** (0.12)	0.058 (0.04)	0.089 (0.08)
initial decision	buy if $\mathsf{E}_1[V\|S] > p_1$, else sell	0.235** (0.05)			0.218** (0.06)	0.252** (0.06)		0.11 (0.08)		
Signal S_1		-0.283** (0.06)								
Signal S_3		0.302** (0.06)								
Constant		-0.090* (0.04)	-0.206** (0.04)	0.286** (0.06)	-0.127* (0.05)	-0.021 (0.04)	-0.058 (0.05)	0.017 (0.08)	0.147** (0.05)	-0.199** (0.03)
Observations		1347	390	407	550	550	550	369	491	487

Table 4
Total Fit of the Rational Model by Treatments.

In each box, the first entry signifies the number of choices that are not according to the rational model, the second number indicates the total number of choice by a type for that treatment, and the third number is the proportion of rational decisions.

		all passes are wrong			some passes are ok		
		S_1	S_2	S_3	S_1	S_2	S_3
Treatment 1, negative hill shape	wrong	7	36	23	7	36	14
	total	53	98	71	53	98	71
	% correct	87%	63%	68%	87%	63%	80%
Treatment 2, increasing		10	40	17	10	28	10
		70	84	67	70	84	67
		86%	52%	75%	86%	67%	85%
Treatment 3, negative U shape		10	73	19	10	59	12
		58	94	74	58	94	74
		83%	22%	74%	83%	37%	84%
Treatment 4, decreasing		12	26	9	12	26	6
		56	99	74	56	99	74
		77%	74%	88%	77%	74%	92%
Treatment 5, positive U shape		9	56	8	9	44	7
		67	95	63	67	95	63
		87%	41%	87%	87%	54%	89%
Treatment 6, negative hill shape		20	31	6	20	27	6
		86	82	58	86	82	58
		77%	63%	90%	77%	67%	90%
Total wrong		68	262	82	68	220	55
Total Trades		322	290	325	322	332	352
Total correct %		82%	53%	80%	82%	61%	86%
Overall			69.5%			74.6%	

Table 5

Occurrence of Rational Herding/Contrarianism.

The first entry in each box denote the number of herding or contrarian actions that occur, the second entry denotes the number of herding/contrarian actions that were theoretically mandated, and the third entry is the fraction of rational herding or contrarian actions that is observed.

		Herding		Contrarianism	
		passes are irrational	passes are rational	passes are irrational	passes are rational
Treatment 1, negative hill shape	does occur			4	4
	should occur			6	6
	% as expected			67%	67%
Treatment 3, negative U shape		18	32		
		89	89		
		20%	36%		
Treatment 5, positive U shape		2	2		
		18	18		
		11%	11%		
Treatment 6, negative hill shape				12	16
				20	20
				67%	80%
herding/contrarianism occurring		20	34	16	20
theoretically mandated		109	109	26	26
percent that arises rationally		18%	31%	62%	77%

Table 6
Herding and Contrarian Trades split up by Treatment.

For each treatment, entries in the top row indicate herding/contrarianism that did occur, entries in the middle row refer to the number of times that herding/contrarianism could have occurred. The third row entries indicate the percentage of realized herding/contrarian trades.

		Herding Trades						Contrarian trades					
		pass=pass			pass=weak buy			pass=pass			pass=weak buy		
		S_1	S_2	S_3	S_1	S_2	S_3	S_1	S_2	S_3	S_1	S_2	S_3
Treatment 1 negative hill shape	occurs	4	11	1	7	34	1	0	4	13	0	4	13
	possible	52	91	6	52	91	6	1	7	65	1	7	65
	% occurs	8%	12%	17%	13%	37%	17%	0%	57%	20%	0%	57%	20%
Treatment 2 increasing		5	5	1	8	5	1	1	23	9	2	23	9
		59	20	7	59	20	7	9	64	47	9	64	47
		8%	25%	14%	14%	25%	14%	11%	36%	19%	22%	36%	19%
Treatment 3 negative U shape		3	20	0	10	21	0	0	0	6	0	0	6
		58	94	1	58	94	1	0	0	73	0	0	73
		5%	21%	0%	17%	22%	0%			8%			8%
Treatment 4 decreasing		4	10	0	6	13	0	2	6	6	6	10	6
		33	63	24	33	63	24	19	25	50	19	25	50
		12%	16%	0%	18%	21%	0%	11%	24%	12%	32%	40%	12%
Treatment 5 positive U shape		5	2	2	7	2	2	1	27	5	2	28	5
		48	36	25	47	36	25	19	59	38	19	59	38
		10%	6%	8%	15%	6%	8%	5%	46%	13%	11%	47%	13%
Treatment 6 negative hill shape		7	7	1	10	19	1	5	16	5	10	17	5
		48	58	19	48	58	19	38	24	39	38	24	39
		15%	12%	5%	21%	32%	5%	13%	67%	13%	26%	71%	13%
actual herding/contrarian trades		28	55	5	48	94	5	9	76	44	20	82	44
possible herding/contrarian trades		300	363	82	300	363	82	86	179	312	86	179	312
in percent		10%	15%	6%	16%	26%	6%	10%	42%	14%	6%	23%	46%

Table 7

U-Shaped and Hill Shaped Signals vs. Herding and Contrarianism.

The table combines the situations where passes never count as herding or contrarian trades and where they do count as weak buys. Panel A in the table represents regressions of the occurrence of a herding trade on the trader receiving a U shaped signal, as in the left equation in (1). Panel B in the table represents regressions of the occurrence of a contrarian trade on the trader receiving a hill shaped signal, as in the right equation in (1). We present the marginal effects obtained by logit regressions. Standard errors were clustered by sessions. The data is restricted to include only trades that could be herding and contrarian trades respectively. Standard errors and significance levels are denoted as in Table 1. Constants were included in the regression but are not reported for brevity.

	Dependent Variable: Herding/Contrarianism Indicator			
	Passes do not count		*Passes count as weak buys*	
	U shaped signal	hill shaped signal	U shaped signal	hill shaped signal
Panel A:				
herding	0.058*			
	(0.024)			
			0.081*	
			(0.035)	
Panel B:				
contrarianism		0.337**		
		(0.051)		
				0.499**
				(0.097)
Observations	741	577	741	577

Table 8
The Decision to Buy for Monotonic Signal Types.

The table displays the results from a logit regression of equation (2), i.e. the decision to buy on the price change; standard errors were robustly clustered by sessions. Symbol $\nearrow$ stands for 'increasing', $\searrow$ for 'decreasing'. For all types, the probability of buying declines as the price increases. The hypothesis is that the coefficient is insignificant (monotonic types always either buy or sell, irrespective of the price). Coefficients are significant for the S_3 types, contrary to theoretical predictions, but they are insignificant for all other types. The numbers can be interpreted as follows. For instance, for S_3 types, a 10% increase in the price will lead to a 5% reduction in the probability that this type buys. Standard errors and significance levels are denoted as in Table 1. Constants were included in the regression but are not reported for brevity.

	Dependent Variable: Buy Indicator **Independent Variable: price change in %**					
	Increasing csds			Decreasing csds		
	S_3	$S_2 \nearrow$	all $\nearrow$	S_1	$S_2 \searrow$	all $\searrow$
Panel A: Passes are buys						
Δp	-0.549**					
	(0.149)					
		-0.457				
		(0.589)				
			-0.420*			
			(0.137)			
				-0.276		
				(0.216)		
					-0.866	
					(0.386)	
						-0.419*
						(0.221)
Observations	407	84	491	390	97	487
Panel B: Passes are omitted						
Δp	-0.692**					
	(0.168)					
		-0.935				
		(0.558)				
			-0.585**			
			(0.165)			
				-0.122		
				(0.234)		
					-0.378	
					(0.376)	
						-0.201
						(0.234)
Observations	380	72	452	359	90	449

Table 9

Do People who Miss the Rationality Test Act Differently?

The table displays the results from a logit regression of equation (3) which assesses whether the occurrence of a herding/contrarian trade is more likely to occur when a U shaped signal recipient passed the first rationality test. According to this test, a subject is less rational if s/he chose the pass decision (which is categorically irrational) in any treatment. For this regression the decision to pass is, of course, *not* considered to be a weak buy. The table reports the marginal effects. The data is restricted to include only trades that could be herding or contrarian trades respectively. Standard errors were clustered by sessions. Significance levels are denoted as in Table 1 and we omit constants from the report.

	U shape	U shape × passer	hill shape	hill shape × passer	passer	Observations
Panel A: Herding						
Herding Indicator	0.102**	-0.114**			0.025	741
	(0.029)	(0.042)			(0.036)	
Panel B: Contrarianism						
Contrarianism Indicator			0.475**	-0.306	-0.084*	577
			(0.096)	(0.161)	(0.033)	

Supplementary Appendix for *Herding and Contrarian Behavior in Financial Markets: An Experimental Analysis* —not for publication—

Andreas Park
University of Toronto

Daniel Sgroi
University of Warwick

September 30, 2009

Abstract

Appendix A analyzes alternative, behavioral explanations for trading behavior. The time-line in Appendix B gives a complete run down of the structure of a typical experimental setting and the text (not the section headings) of the instructions detailed in Appendix C are a reproduction of what is read to the subjects during the experiment. Answers to questions are of course unscripted, but all other communication with subjects was kept to a minimum. The questionnaire is reproduced in Appendix D. Practical details about our software are provided in Appendix E. We also attach the information pages that were given subjects before each treatment, Figures 1 to 5. These pages include the parameter values that were used in the experiment. 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

A.1 Development of Alternative Hypotheses

While we believe that our data indicates that people act in the spirit of the rational model with risk neutrality, we examine numerous behavioral explanations for any observed irrationality. In particular we study specifications in which people dampen the effect of observed trades in computing their expectations, namely, by re-scaling probabilities, or by correcting for their predecessors' presumed mistakes. However, while some of these models can generate a higher fit with the data, very high levels of dampening need to be assumed to achieve a measurable effect and these same dampening effects remove the scope to explain rational herding behavior. We also consider the potential for risk and loss aversion to explain the data but find no effect.

Specifically, the theoretical model implicitly imposes strong requirements on subjects being able to compute history-dependent expectations correctly. We will now discuss ways to capture possible departures of the theory. If subjects do not act in accordance with the theory, we have to check if giving them leeway when the decision is close allows for a better fit of the data to the model. We will thus examine

Hypothesis A (Small Errors) *Subjects generally act according the theory as in Hypothesis 2 if the decision is not too close, that is, if prices and private expectations are sufficiently different.*

Standard economic theory suggests that people are risk averse. If they are, then this may affect the performance of our model and we thus have to examine

Hypothesis B (Risk Aversion) *Subjects generally act according the theory as in Hypothesis 2 subject to a risk averse utility function.*

Similarly, a large body of experimental research has found that people tend to react differently to relative gains and losses. We will thus check whether this kind of behavior would affect the performance of our model.

Hypothesis C (Loss Aversion) *Subjects generally act according the theory as in Hypothesis 2 subject to a loss averse valuation function.*

Next, for subjects to act in accordance with the theory, it is imperative that they perform Bayesian updating correctly. Yet there are various behavioral theories which contradict Bayesian updating. We aim to examine whether, when and how these might explain departures from the standard fully rational theory.

Hypothesis D (Non-Bayesian Updating) *Subjects generally act according the theory as in Hypothesis 2 subject to updating their beliefs in a non-Bayesian fashion.*

Finally, the underlying decision problem is not simple, and this is common knowledge among subjects. They may thus believe that at least some of their counterparts may persistently err. Moreover, they may also think that some of their counterparts think the same way and react to this irrationality.[1] We will thus examine whether error correction formulations will help us understand the data better.

Hypothesis E (Error Correction) *Subjects generally act according the theory as in Hypothesis 2 subject to correcting for possible errors and rational reactions to errors that their fellow experiment participants may make.*

Each of the alternative Hypotheses A to E will be analyzed within a parameterized model which will be made more explicit in Section A.

A.2 Results for the Alternative Explanations

We argued in the main text that our findings are generally supportive of the theory, yet there are still problems explaining for instance the apparent contrarianism displayed by S_3 types for high prices or the lack of herding and (to a lesser extent) contrarian trades relative to theoretical predictions. It is well-established in experimental work, that models with Bayesian rationality and risk-neutral agents may not provide the best fit for the data, and so this section we look at alternative models to see if we can explain findings which the rational theory cannot.

We explore a range of alternative explanations, ranging from omitting decisions that are within an ϵ error-region, to risk- and loss-aversion, to various forms of alternative information updating. These alternative hypotheses usually depend on some parameter(s). Our approach is to vary this parameter and see how the variation improves the overall fit of the alternative model to the data. Such an approach is, of course, a maximum likelihood technique, albeit a coarse one.

A.2.1 Robustness to Small Errors

Hypothesis A considers the scope for subjects to make errors when the decision they have to make is a close one. To test this idea, we omit all trades that occur when prices and expectations are within ϵ of each other, for small values of ϵ. This variation typically worsens the fit of the model. The reason is that while it does capture some errors made when the decision is close, it also rules out some correct decisions that were made when the decision is close. For instance if we set $\epsilon = .01$, then the total model fit is reduced to 67% (for S_1, S_2 and S_3, 78%, 45%, and 80% respectively). We repeated the analysis with different values for ϵ but could not generate a higher fit for the data.

[1]Notions of level K reasoning (see Costa-Gomes, Crawford and Broseta (2001)) and Quantal Response Equilibria (see McKelvey and Palfrey (1995) and McKelvey and Palfrey (1998)) describe this type of behavior.

Finding A (Errors when the decision is close) *Subjects do not seem to be more likely to act in accordance with prescribed optimal actions when the decision is not too close and hence we do not support Hypothesis A.*

A.2.2 Risk Aversion

One persistent finding from the last sections is that traders exhibit a general tendency to act as contrarians. One might thus also entertain the idea that traders act as contrarians because of risk-aversion (Hypothesis B).[2] We can go about examining this by computing the optimal action when people have a concave utility function. We checked this employing both CARA and CRRA utility functions:

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

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

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

Computing the expected utilities we find, however, that the performance of a model with risk aversion is worse for all reasonable levels of risk aversion. For CRRA with log-utility ($\gamma = 1$), it is below 50%; for CARA it is 48% for $\rho = 1$ and 64.5% for $\rho = .01$, rising as ρ declines. As ρ declines, we capture more of the behavior by S_3 types but less of the behavior by S_2 types. Note that as ρ falls, we move closer to risk neutrality. Table I displays the results for some select parameter values that are indicative of the general tendencies in the data.

Finding B (Risk Aversion) *The performance of a model with risk aversion is worse for all reasonable levels of risk aversion and so we do not support Hypotheses B.*

[2] With typically high and rising prices, acting in a contrarianism entails selling. Selling gives an immediate payoff equal to price, whereas holding or buying entails a wait for a risky return. Contrarian behavior may therefore be justifiable in terms of risk-aversion.

A.2.3 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. We tried various alternative parameter configurations, searching for the best fit possible, but none performed significantly better than $\alpha = \beta = 0.8$ and $\gamma = 2.25$, which is a common specification for the parameters stemming from experimental observations (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 benchmark rational model. For parameters as estimated by Tversky and Kahneman (1992), the fit is below 38%. Table II illustrates this observation for the above parameters as well as for one other configuration.[3]

Finding C (Loss Aversion) *Using a variety of parameterizations we could not achieve a better fit than under the benchmark rational model. Thus we do not support Hypothesis C and conclude that loss aversion is not an important influence for behavior in our experiment.*

A.2.4 Non-Bayesian Forms of Updating

We consider various forms of non-Bayesian updating to assess whether they offer insights over and above the benchmark rational model. One extreme 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 actions that is prescribed at the initial history (e.g. with negative U shape, S_2 traders always sell).

Initially, this specification appears to do well: it fits 73.6% of the data which is higher than the rational model (with passes as wrong trades); broken up by type the fit is 82% for S_1, 63% for S_2 and 80% for the S_3, which is again higher than the rational fit without

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

passes but lower than the rational model with passes. Of course with this alternative model, we cannot accommodate passes as 'weak buys' because this would be contrary to the spirit of 'no changes of the action'. Indeed this illustrates the first weakness: a model based on people choosing their prior action will not help us to understand any changes in behavior that might have occurred, in particular not for S_1 and S_3 types. Since the econometric analysis has already revealed that the S_3 are sensitive to the price, this decision rule is rather weak. Looking only at those actions that are at odds with rationality (and counting passes as wrong as would befit the hypothesis here), only 30% (79 decisions) of the irrational actions of the S_2 are in line with this hypothesis. This further reveals that the remaining 183 decisions are due to a change of actions, which constitute a total of 33% of the S_2's actions.

While the fit of a model which emphasizes the ex-ante optimal action at the expense of prices and history seems high, the model does not help explain any of the observed changes of actions. We therefore argue that given the extra complexity of the model, and the apparent only slight improvement in fit over and above the rational model, a model focusing on prior actions is of limited use.

We also investigated the possibility that subjects do not update their beliefs at all as prices change but act solely on the basis of their prior expectation. In fact, the results from the regression in Section 3 already touch upon this topic and there we argued that this variable is insignificant for most signal types.

Finally, we considered probability shifting, whereby traders underplay (overplay) low (high) probabilities coming from the observed history, which is equivalent to traders overstating the probabilities of their prior expectations. The usual symmetric treatment of this under- or overstating of probabilities is to transform probability p into $f(p)$[4] as follows

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

Parameter values $\alpha > 1$ are associated with S-shaped re-valuations (extreme probabilities (those close to 0 and 1) get overstated, moderate probabilities (those close to 1/2) understated), $\alpha < 1$ with reverse S-shaped valuations (extreme probabilities get understated, moderate probabilities overstated). Note that the transformation $f(p)$ applied to probabilities of all three states do not yield a probability distribution. However, when

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

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 not matter as much as under the rational model. We tried both specifications and the results are similar. Here we report the results where $\mathsf{Pr}(V|H_1) \times \mathsf{Pr}(S|V)$ has been re-scaled with an $\alpha > 1$; downward scaled probabilities of the history $\mathsf{Pr}(V|H_t)$ yield similar insights.

Comparing the results here to those in Table 4, one can see that the fit of probability scaling 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. However, re-scaling does a poor job at explaining herd-behavior of any sort.

Finding D (Non-Bayesian Updating) *Having considered several forms of Hypothesis D including emphasizing the ex-ante optimal action at the expense of prices and history, acting solely on the basis of price and not history, and the underplaying of posteriors derived from observed history or the overplaying of the signal-prior or we have discovered no insights over and above the benchmark rational model.*

A.2.5 Error Correction Provisions

To investigate Hypothesis E, and inspired by level K reasoning (see Costa-Gomes et al. (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 their peers and instead believe that their actions are random. 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 to find the δ for which this specification yields the best data fit; we obtained the best fit for $\delta = 1/3$. However, compared to the rational model the improvement of the fit is minor (see Table IV): the rational fit is 69.5% vs. 71.6% with the error correction.

An alternative interpretation for this error correction is that the level of noise trading is perceived higher than it actually is because other subjects act randomly: a δ of 33.33% translates into a factual noise level of 75%. 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.[5] Alas, as with the simple error correction, we do not obtain a substantially better fit with the data, as can be gleaned from Table IV: we obtained the best fit for $\delta = .24$ but the improvement is merely from a 71.6% to a 71.7% fit and is thus negligible.

Finding E (Error Correction) *A model in which agents recursively take their predecessor's decisions as prone to error offers no noteworthy improvement in fit over the benchmark rational model and so we do not support Hypothesis E.*

[5]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 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 (2007).

A.2.6 Summary of Alternative Behavioral Explanations

While forms of slow updating improve the fit of the data slightly, no alternative model is capable of providing a convincing explanation for the results. Slow updating and error correction specifications (which imply overestimation of noise trading) are essentially very similar, and also have strong similarities to a strategy of following the prior (which is effectively a policy of zero updating). We also tested several other, related models but since the results did not differ, we choose not to report them here in detail.[6]

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

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

(1) In treatment 2, with an increasing information structure, when prices rise, the S_2 should start (contrarian) selling at prices before the S_3 types. In the data we do in fact observe that a much larger fraction of S_2 act as contrarians than S_3.

(2) We should see more irrational herding by S_2 types than S_1 types in treatment 4 with a decreasing information structure. In the data, while we do observe less irrational herding by S_2 types than S_1 types, the difference is minor.

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

In general therefore, there is not enough evidence of slow updating over and above the

[6]For example, we considered whether traders chase short-term trends, but found that they do not.

benchmark rational theory to allow us to support it as a general description of subject behavior. In conclusion and to summarize the various findings in this section, we feel that the variations and behavioral alternatives to the benchmark model of Bayesian rationality which we have considered do not provide sufficient improvement in fit to allow us to support Hypotheses A to D. We should emphasize the positive nature of these findings, since it reinforces the relative success of the rational model in describing the data, which was the key insight of rational herding theory from its earliest incarnations: to provide a rational explanation of apparently irrational phenomena.

B Time-line

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

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

C Instructions

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

Instructions 1 (Welcome)

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

Instructions 2 (After Seated)

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

Instructions 3 (The Experimental Setting)

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

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

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

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

Instructions 4 (The Software)

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

1. It tells you about all the trades that occur during the round; you also see when a trade occurs and whether or not someone bought or sold a share. For your convenience, there is a graph that plots the sequence of prices.

2. Your screen also lists the current market price; people can either buy a share at this price or they can sell their share at this price.

3. In the case where we restrict the time when you can make a trade, a red bar will appear on the bottom of the screen to highlight the fact that you can trade. During this time the buy, sell and pass buttons will be available for your use, typically only once per round, though twice in the final 3 rounds.

4. There is also a box in which you receive some information from your "broker" which I will explain in a few moments.

5. The screen includes a timer which indicates how many seconds have gone past during the round.

6. Finally, the screen updates itself whenever a trade is made.

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

Instructions 5 (Compensation)

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

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

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

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

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

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

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

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

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

Instructions 6 (The Information Setting)

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

Your broker will give you a tip via your computer screen that indicates his view about what sort of event will occur. He might give you a "good tip" (which we call $S3$), "bad tip" ($S1$) or "middle tip" ($S2$). A good $S3$ tip indicates that he believes the event will be good and the share price will be 125 vcu after it is realized, a bad $S1$ tip that something bad will happen indicates 75 after the event is realized. A middle $S2$ tip is a bit more complex but indicates he feels 100 vcu is his best guess:

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

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

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

You will be given 2 minutes to examine the relevant sheet before each round. You will then receive notification on your computer screen of the actual tip sent to you from the broker: S1, S2 or S3, and will have another minute to consider this. The beginning of the round will then be announced and trading will begin. Remember you can only trade during the 5 second window indicated by a red bar on your screen. The buttons on the screen (buy, sell or pass) can only be pressed during this time and only once per round.

Instructions 7 (Summary)

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

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

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

Instructions 8 (Experiment End)

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

D Questionnaire

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

About you

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

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:

E The Software

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

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

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

Here you can see that the price initially rose from a level of 100, indicating buying at the early stages, but then price started to fall back, it rallied and then fell back further to a value of around 116. This subject's private signal was S1 (low) and the subject had a single share to sell and a large cash balance to enable him to buy a further share. He could also pass (declining to buy or sell) when he was given the opportunity to trade.

The software was custom-programmed for the experiment, as existing software mimics order driven markets in which traders submit both limit and market orders.[7]

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

References

Costa-Gomes, M., V. Crawford, and B. 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," Mimeo, forthcoming American Economic Review, The Ohio State University 2007.

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

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

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

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

Table I
Risk-Aversion Analysis.

The table classifies trades as 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 $\rho = 1$ 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. We omit the total number of decisions as they can be straightforwardly Table 4.

	Total Number of wrong decisions CRRA utility, $\gamma = 1$ (log-utility)			Total Number of wrong decisions CARA utility, $\rho = 1$			Total Number of wrong decisions CARA utility, $\rho = 0.01$		
	S1	S2	S3	S1	S2	S3	S1	S2	S3
Treatment 1 negative hill shape	18	58	37	46	60	14	46	58	48
	33%	58%	51%	84%	60%	19%	84%	58%	67%
Treatment 2 increasing	31	60	36	60	28	10	60	30	50
	42%	67%	53%	81%	31%	15%	81%	33%	74%
Treatment 3 negative U shape	13	69	37	48	59	6	48	18	61
	22%	73%	49%	82%	63%	8%	82%	19%	80%
Treatment 4 decreasing	32	41	33	44	74	6	44	75	65
	57%	41%	45%	76%	75%	8%	76%	76%	88%
Treatment 5 positive U shape	29	66	32	58	29	7	58	25	55
	43%	67%	49%	86%	30%	11%	86%	26%	85%
Treatment 6 negative hill shape	41	51	22	66	42	6	66	47	52
	47%	61%	38%	76%	51%	10%	76%	57%	90%
Total number wrong	**164**	**345**	**197**	**322**	**292**	**49**	**322**	**253**	**331**
percentage wrong	**41%**	**61%**	**48%**	**70%**	**44%**	**10%**	**70%**	**39%**	**68%**
Total model fit	48.8%			41.3%			56.5%		

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 Section A.2.3. 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. We also tried many different parametric configurations but could not provide a higher fit. The structure of the table is similar to that of Table I.

	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%	40 71%	55 56%	48 86%
Treatment 5 positive U shape	32 48%	70 71%	32 49%	32 48%	73 74%	46 71%
Treatment 6 negative hill shape	41 47%	59 71%	22 38%	41 47%	59 71%	22 38%
Total number wrong wrong percentage	**185** **46%**	**391** **69%**	**197** **48%**	**187** **47%**	**408** **72%**	**277** **67%**
Total model fit	43.8%			36.7%		

Table III
Probability Scaling.

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.2.4. The structure of the table is similar to that in Table 4 with correct and wrong actions listed alongside one another.

	With $\alpha = 5$			With $\alpha = 10$			With $\alpha = 25$		
	S1	S2	S3	S1	S2	S3	S1	S2	S3
Treatment 1 negative hill shape	46	61	57	46	61	57	46	61	57
	7	37	14	7	37	14	7	37	14
	87%	62%	80%	87%	62%	80%	87%	62%	80%
Treatment 2 increasing	60	56	57	60	56	57	60	56	57
	10	28	10	10	28	10	10	28	10
	86%	67%	85%	86%	67%	85%	86%	67%	85%
Treatment 3 negative U shape	48	31	68	48	41	68	48	57	68
	10	62	6	10	52	6	10	36	6
	83%	33%	92%	83%	44%	92%	83%	61%	92%
Treatment 4 decreasing	44	73	65	44	73	65	44	73	65
	12	25	6	12	25	6	12	25	6
	77%	74%	92%	77%	74%	92%	77%	74%	92%
Treatment 5 positive U shape	58	66	56	58	66	56	58	66	56
	9	29	7	9	29	7	9	29	7
	87%	69%	89%	87%	69%	89%	87%	69%	89%
Treatment 6 negative hill shape	66	60	52	66	60	52	66	60	52
	20	22	6	20	22	6	20	22	6
	77%	73%	90%	77%	73%	90%	77%	73%	90%
Total trades in line with probability scaling	**322**	**347**	**355**	**322**	**357**	**355**	**322**	**373**	**355**
Percentage explained	84%	64%	88%	84%	66%	88%	84%	69%	88%
total fit		77.1%			77.9%			79.1%	

Table IV
Error Correction Provisions.

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.2.5. 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 4 with correct and wrong actions listed alongside one another.

	simple noise shift $\delta = 2/3$			simple noise shift $\delta = 1/3$			Level 2 noise shift $\delta = .24$		
	S1	S2	S3	S1	S2	S3	S1	S2	S3
Treatment 1 negative hill shape	46	62	46	46	61	47	46	63	46
	7	36	25	7	37	24	7	35	25
	87%	63%	65%	87%	62%	66%	84%	63%	64%
Treatment 2 increasing	60	35	50	60	38	49	60	36	49
	10	49	17	10	46	18	10	48	18
	86%	42%	75%	86%	45%	73%	81%	40%	72%
Treatment 3 negative U shape	48	35	57	48	59	55	48	59	58
	10	59	17	10	35	19	10	35	16
	83%	37%	77%	83%	63%	74%	82%	63%	76%
Treatment 4 decreasing	43	74	64	43	74	65	44	74	65
	13	25	10	13	25	9	12	25	9
	75%	75%	86%	75%	75%	88%	76%	75%	88%
Treatment 5 positive U shape	58	38	55	58	44	55	58	42	55
	9	57	8	9	51	8	9	53	8
	87%	40%	87%	87%	46%	87%	86%	43%	85%
Treatment 6 negative hill shape	66	49	52	66	48	49	66	49	49
	20	33	6	20	34	9	20	33	9
	77%	60%	90%	77%	59%	84%	76%	60%	84%
Total identified	**321**	**293**	**324**	**321**	**324**	**320**	**325**	**324**	**322**
Total fit	84%	54%	81%	84%	59%	80%	82%	59%	79%
Total model fit	69.5%			71.5%			71.7%		

Figure 1
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%

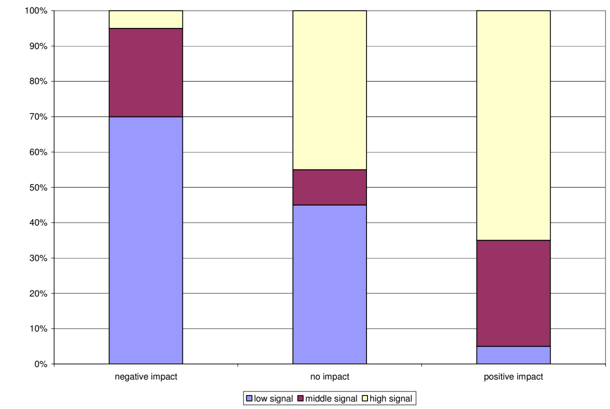

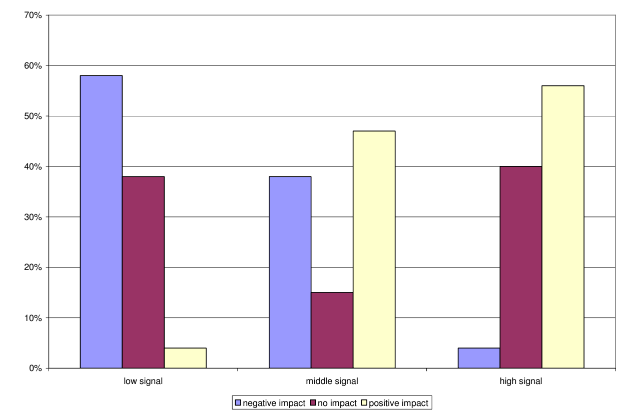

Round

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%

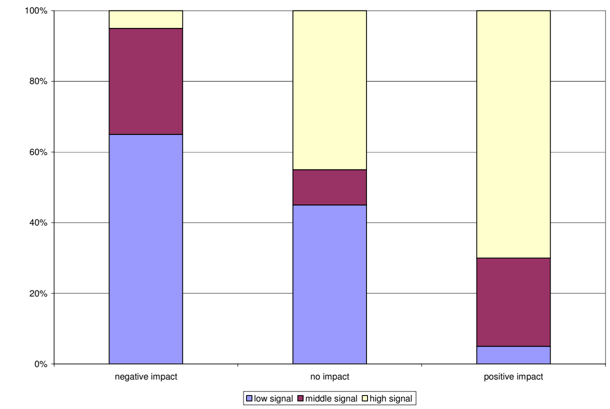

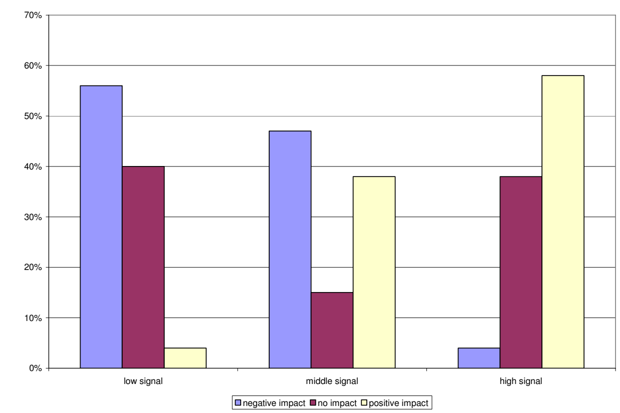

Figure 2
Information Sheet for negative U Shape

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%

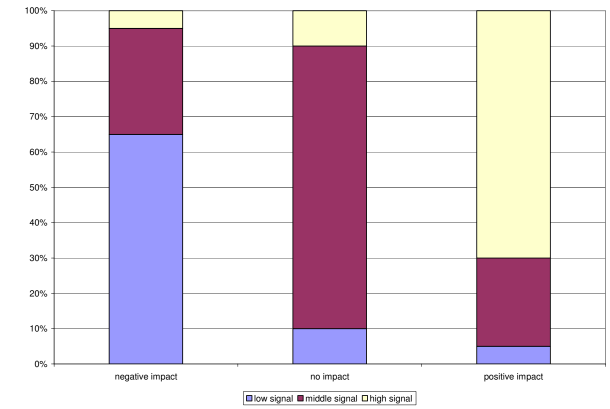

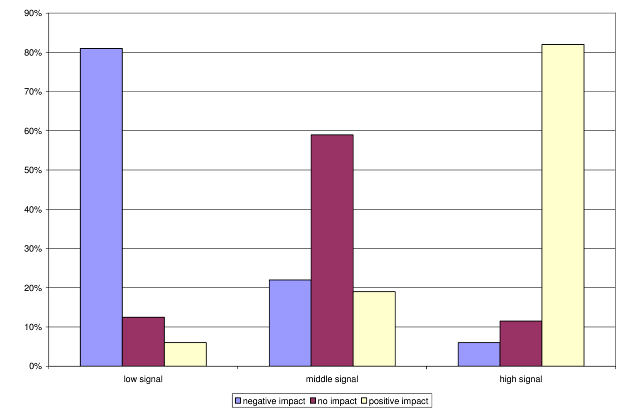

Figure 3
Information Sheet for negative hill shape

Figure 4
Information Sheet for increasing csd

Round

Signals: **Case 3**

- 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 a *positive* effect but he is not confident enough to give the good signal.

If the true effect will be POSITIVE then you receive

- Signal S1 (bad) with chance 5%
- Signal S2 (no effect) with chance 35%
- Signal S3 (good) with chance 60%

If the true effect will be NEGATIVE then you receive

- Signal S1 (bad) with chance 80%
- Signal S2 (no effect) with chance 15%
- Signal S3 (good) with chance 5%

If indeed the effect will be NO EFFECT then you receive

- Signal S1 (bad) with chance 35%
- Signal S2 (no effect) with chance 25%
- Signal S3 (good) with chance 40%

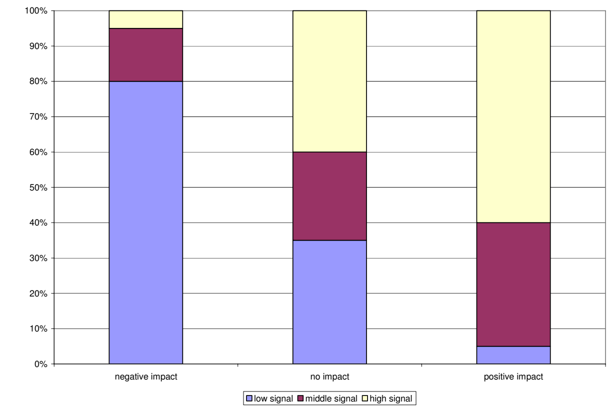

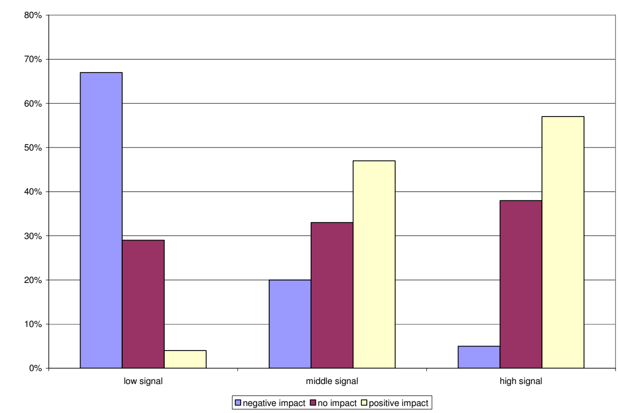

Figure 5
Information Sheet for decreasing csd

Round

Signals: **Case 3**

- 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 a *negative* effect but he is not confident enough to give the good signal.

If the true effect will be POSITIVE then you receive

- Signal S1 (bad) with chance 5%
- Signal S2 (no effect) with chance 15%
- Signal S3 (good) with chance 80%

If the true effect will be NEGATIVE then you receive

- Signal S1 (bad) with chance 60%
- Signal S2 (no effect) with chance 35%
- Signal S3 (good) with chance 5%

If indeed the effect will be NO EFFECT then you receive

- Signal S1 (bad) with chance 40%
- Signal S2 (no effect) with chance 25%
- Signal S3 (good) with chance 35%

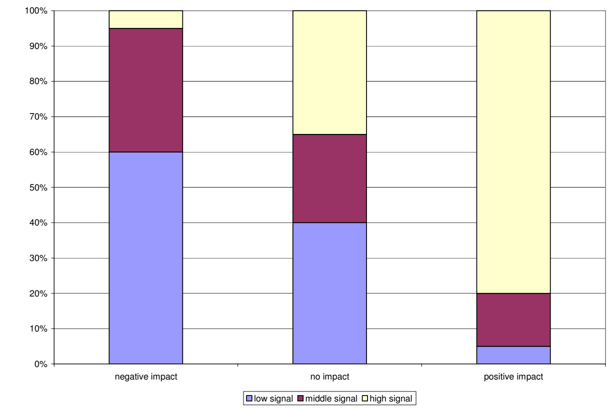

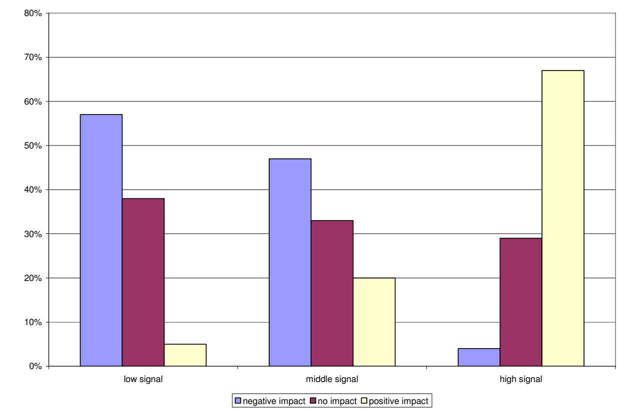

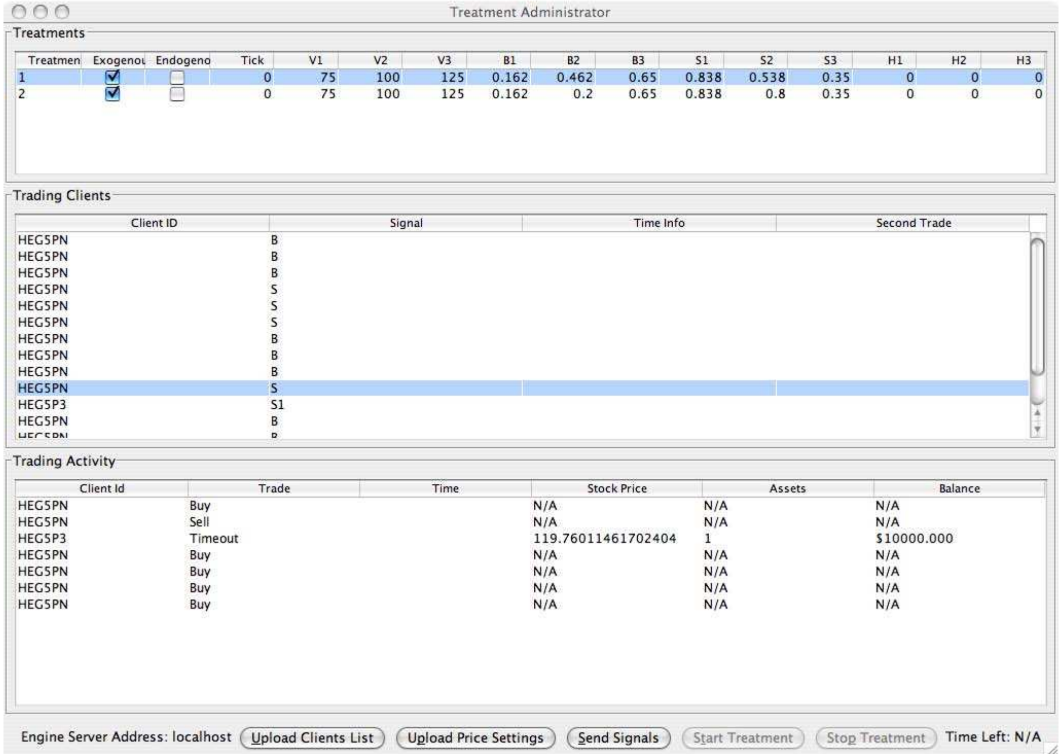

Figure 6
The Administrative Interface

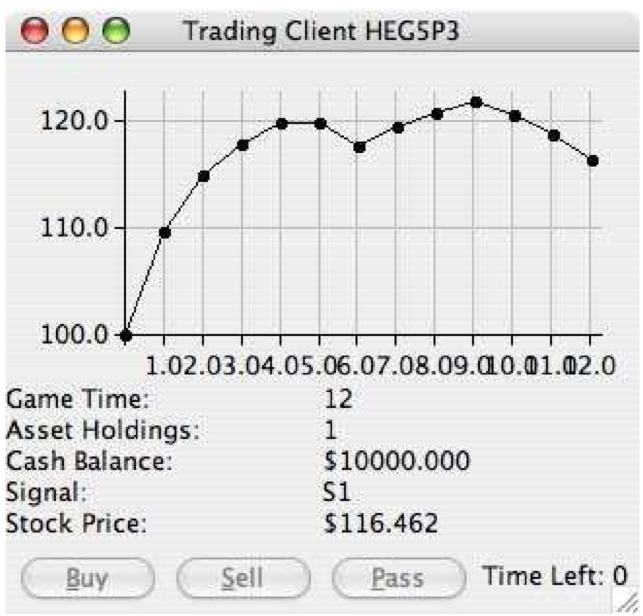

Figure 7
The Trading Client