



The World's Largest Open Access Agricultural & Applied Economics Digital Library

This document is discoverable and free to researchers across the globe due to the work of AgEcon Search.

Help ensure our sustainability.

Give to AgEcon Search

AgEcon Search

<http://ageconsearch.umn.edu>

aesearch@umn.edu

*Papers downloaded from **AgEcon Search** may be used for non-commercial purposes and personal study only. No other use, including posting to another Internet site, is permitted without permission from the copyright owner (not AgEcon Search), or as allowed under the provisions of Fair Use, U.S. Copyright Act, Title 17 U.S.C.*

No endorsement of AgEcon Search or its fundraising activities by the author(s) of the following work or their employer(s) is intended or implied.

Warwick Economics Research Paper Series

Herding and Contrarian Behavior in Financial Markets: An Experimental Analysis

Andreas Park and Daniel Sgroi

January, 2016

Series Number: 1109

ISSN 2059-4283 (online)

ISSN 0083-7350 (print)

This paper also appears as [*CRETA Discussion Paper No: 17*](#)

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

Andreas Park[†]

University of Toronto

Daniel SgROI[‡]

University of Warwick

January 2016

Abstract

We analyze and confirm the existence and extent of rational informational herding *and* rational informational contrarianism in a financial market experiment, and compare and contrast these with 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 sources of both within the data. Correcting for subjects who can be identified as less rational increases our ability to predict herding or contrarian behavior considerably.

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

*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. Special thanks to Antonio Guarino, Jim Peck and Gustavo Borbonis for invaluable advice and to Annie Jekova, Malena Digini, Kenji Yee and Kelly Hayward for expert research assistance.

[†]E-mail: andreas.park@utoronto.ca

[‡]E-mail: daniel.sgROI@warwick.ac.uk

During the 2008 Financial Crisis stock markets displayed extraordinary fluctuations. From September to mid-November 2008, there were eight days when the Dow Jones Industrial Average changed by more than 5% in absolute terms (from close to close). From World War II to 2008 there were only sixteen other days when the day-to-day change exceeded 5% in absolute value. Moreover, although we perceive the time of the 2008 crisis as a time of market decline, there were two days when the Dow rose by more than 10%. Intra-day fluctuations were even more pronounced: on fourteen days the maximum and minimum prices levels between two days were more than 10% apart.

Such extreme price fluctuations are possible only if there are dramatic changes in behavior (from buying to selling or the reverse). Such behavior and the resulting price volatility is often claimed to be inconsistent with rational trading and informationally efficient prices. Commentators invariably attribute dramatic swings to investors “animal instincts”, which to most economists, is deeply unsatisfying. Work by Park and Sabourian (2011), however, provides a theoretical insight showing that switching behavior can be fully rational even with efficient prices and also provides a framework that admits both rational herding and rational contrarian behavior. The nature of information in real financial markets is often complex: a nice feature of Park and Sabourian (2011) is to allow traders to receive signals that are non-monotonic. For instance, a trader may receive a signal suggesting that an upward move in prices or a downward move in prices are more likely than prices remaining the same (what we define in the theory section below as a “hill-shaped signal”) or alternatively a signal may indicate that prices are likely to stay the same but might fall or rise with positive probability (a “U-shaped signal”). The theoretical results vary according to the private signals received by traders. The key hidden parameter that defines informational herding theory is the private information held by traders. Such information is unobservable in the real world which makes it difficult to impossible to test any herding theory with, say, transaction level financial data. In a laboratory setting, however, one can control private information and examine directly how behavior changes with private information. Since our primary objective is to provide an empirical test of this theory, we therefore base our work upon a series of lab experiments, which is also one of the largest yet to examine rational herding with around 1350 trades across six separate treatments.

To appreciate our results and our contribution it is important to understand how our findings relate to the literature. Avery and Zemsky (1998) show theoretically that herding cannot arise if there are only two liquidation values. Experimental research by Drehmann, Oechssler, and Roider (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. Instead, they

also found that subjects tended to (irrationally) act against the crowd, i.e. to behave as contrarians. This natural tendency to be a contrarian may bias behavior against *rational* herding and may also overwhelm rational motivations to be a contrarian: disentangling these two possibilities requires a notion of *rational* contrarianism which is an important feature of our model.

In our work, the overall fit of the data to the theoretical model is roughly 75%. This figure is in line with the results in Drehmann, Oechssler, and Roeder (2005) and Cipriani and Guarino (2005), both of which test simpler models of (no-)herd behavior.¹ Broken down by player type, these numbers are 79% for the recipients of signals with monotonic likelihood functions,² 65% for recipients of “hill-shaped signals”, and 45% for recipients of “U-shaped signals”.

Considering the results on irrational contrarianism in the literature, it is not too surprising that our results 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.

Given these findings, it is important to understand whether the underlying information model has predictive power. To this end we analyze whether receiving a U or hill-shaped signal affects the chance of acting as a herder and contrarian significantly relative to any other signal. We find that this is indeed the case. Compared to receiving any other signal, recipients of U-shaped signals are significantly more likely to herd (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. Next, 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.

The decision problem of herding candidates is most difficult. We thus ask whether those that act more rationally are more likely to act according to the theory. Being less rational is determined by a simple test: the decision not to trade is never rational in our model. Thus, 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 the marginal effect of receiving a U-shaped

¹To 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.

²Two treatments had only monotonic likelihood functions, with one signal being “weaker” than another. Recipients of the weaker signals generally act less frequently in accordance with the theory; see Table 1.

signal on the chance of herding is significantly stronger for the more rational subjects.

OVERVIEW. The rest of the paper is structured as follows. The next section reviews the experimental literature on financial market herding. Section 2 examines the theoretical framework, discusses the modifications undertaken to better fit a laboratory experiment, and outlines the hypotheses we seek to evaluate. Section 3 describes the experimental design, the subject pool, and the various treatments. 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 and also looks at the relationship between prices and trading decisions. Section 7 studies whether there are differences in behavior with respect to herding and contrarianism between more and less rational types. Section 8 provides a brief discussion of some alternative theories, the examination of which is confined to the Appendix. Section 9 summarizes the key findings and concludes. The Appendix contains proofs, the detailed discussion of alternative behavioral explanations for our findings, as well as the subject instructions, an experimental timeline and screenshots.

1 Experimental Work on Financial Market Herding

Several papers have examined informational herding behavior in experimental settings; only a few employ efficient prices and none studies rational contrarianism.³ 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, many experiments since have not specifically considered a financial market setting and have implicitly held prices fixed.⁴ There are however a number of important studies that have considered moving prices such as Drehmann, Oechssler, and Roeder (2005) and Cipriani and Guarino (2005), discussed earlier, and Cipriani and Guarino (2009). The latter tested the model presented in Avery and Zemsky (1998) with so-called ‘event uncertainty’ in the lab, enlisting financial professionals.⁵ They find a degree of rational herding in line with our own results, which

³An interesting alternative is provided by Boortz, Jurkatis, Kremer, and Nautz (2013) who analyze a computer simulation of a model based on Park and Sabourian (2011). Their findings generally support the model though they argue that herd behavior during the 2008 financial crisis was not entirely explainable by the model alone.

⁴A more recent contribution using a setting without moving prices is Alevy, Haigh, and List (2007), who used professional Chicago Board of Trade traders as their subjects.

⁵More recently Cipriani and Guarino (2014) provide estimation tools for real transaction level data in a setting based on the theoretical model of Avery and Zemsky (1998). For the one particular NYSE stock in 1995 that they use as an example, they argue that herd activity accounts for about 4% of the price movements.

suggests that our student subject pool does not distort 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, Roll, and Subrahmanyam (2002)). Prior experimental research may have given the impression that these contrarian tendencies are irrational; our findings here suggest instead that this not necessarily 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. In our framework, such overestimating tendencies would weaken herding instincts and reinforce contrarianism.

Finally, we also carried out a similar experimental analysis in Park and SgROI (2012) in which we allow subjects to decide when to trade rather than limiting them to an exogenous sequence. This has the advantage of adding realism, but the significant disadvantage of not having a theoretical basis on which to judge the rationality of individual subjects.⁶

2 Definitions, Theoretical Results, and Predictions

This section provides a simplified variation of the theory in Park and Sabourian (2011) to provide theoretical predictions and thus guide our experimental design.

2.1 An Illustrative Example

MC, the main competitor of a financial institution, FI, has just declared bankruptcy. This may be good for FI because they may be able to attract the MC's customers. If this situation materializes, a share of FI is worth V^h . MC's failure may also be bad since FI may have made the same mistakes as MC; a share of FI is then worth $V^l < V^h$. We are interested in the behavior of a privately informed investor, who received a (noisy) signal S and who observes several sales. Sales, loosely, convey that the sellers had negative

⁶Shachat and Srinivasan (2011) also allow experimental subjects to trade when they wish, and moreover allow greater flexibility in the number of trades (Park and SgROI (2012) allow at most two trades per subject).

information. We ask: is it possible that a trader sells, even though his private (noisy) signal alone tells him that FI is worth V^h ?

Suppose that the price $p \in (V^l, V^h)$ is *fixed*. Then the answer is yes: for sufficiently many sales, the private information of any trader is swamped by the negative information derived from observing that early sells, $E[V|S, \text{many sales}] < p$, even if the private signal favored state V^h , $\Pr(S|V^h) > \Pr(S|V^l)$.

There is, of course, a major shortcoming to this argument: the price of a financial asset is not fixed. If, as is common in financial market models, the price would be such that $p = E[V|\text{information contained in all past trades}]$, then we would predict that for all past trading activity $E[V|S] \geq p$ if and only if $\Pr(S|V^h) \geq \Pr(S|V^l)$. In other words, someone with favorable information would never sell; see Avery and Zemsky (1998). Experimental evidence has confirmed this finding; see Drehmann, Oechssler, and Roeder (2005) and Cipriani and Guarino (2005).

More recent theoretical research, Park and Sabourian (2011), has found, however, that as we admit more complexity in the form of additional states, it is possible to identify situations when traders act against their information. The ideas behind Park and Sabourian are best explained by extending upon the above example. Namely, suppose that there is a third outcome, one in which FI is unaffected by MC's failure, associated with value V^m with $V^l < V^m < V^h$. Assume all outcomes are equally likely. We are now interested in the behavior of a privately informed investor, who received a signal S , after a good public announcement G and a bad public announcement B . The information content of G is such that the worst state can be ruled out, $\Pr(V^l|G) = 0$, that of B is that the best state can be ruled out, $\Pr(V^h|B) = 0$. Compared to the situation when all outcomes are equally likely, prices will be higher after G and lower after B .

If an investors buys, his expectation must exceed the expectation of the market, and if he sells, his expectation must be below that of the market. Both G and B eliminate one state, so that, after each announcement arguments from two-state setups apply. For instance, after good news G , an investors buys if he thinks, relative to the market, that it is more likely that FI will thrive than be unaffected. Mathematically, $E[V|G] \leq E[V|S, G]$ is equivalent to $\Pr(S|V^m) \leq \Pr(S|V^h)$ and $E[V|S, B] \leq E[V|B]$ is equivalent to $\Pr(S|V^m) \leq \Pr(S|V^l)$. Thus a privately informed investor buying after G and selling after B is equivalent to $\Pr(S|V^l) > \Pr(S|V^m)$ and $\Pr(S|V^h) > \Pr(S|V^m)$. Conversely, a privately informed investor selling after G and buying after B is equivalent to $\Pr(S|V^l) < \Pr(S|V^m)$ and $\Pr(S|V^h) < \Pr(S|V^m)$.

In words, an investor sells after B and buys after G , if and only if, compared to the market, his private information is such that he puts more weight on the extreme outcomes, failing and thriving, and less on the middle one, being unaffected. Such a type has U -

shaped information and, loosely, herds in the sense that he acts like a momentum trader, buying with rising and selling with falling prices. Similarly, an investor buys after B and sells after G , if and only if, he puts less weight on the extreme outcomes and more on the middle one compared to the market. Such a type has *hill-shaped* information and, loosely, trades contrary to the general movement of prices.⁷

There are several points to note about this example. First, the public announcements G and B are degenerate as they each exclude one of the extreme states. Yet the same kind of reasoning holds if we replace G by an announcement that attaches arbitrarily small probability to the worst outcome, V^l , and if we replace B by an announcement that attaches arbitrarily small probability to the best outcome, V^h . Second, in the above illustration, G and B are *exogenous* public signals. In the security model underlying our experiment, on the other hand, public information is created *endogenously* by the history of publicly observable transactions. Yet the intuition behind the results is similar to the above illustration in that, for instance, after a history of many buys, state V^l can be almost ruled out (as after announcement G). The underlying model is thus self-contained in that public signals are generated endogenously by trading. Switching behavior is triggered by observational learning, which is a core feature of informational herding, and from this perspective it is thus justified to label switching with the crowd as “herding”.

The theoretical model underlying our experiment mimics the above example with three states, three signals, and three possible trading actions (buy, no trade, sell). It is based upon a standard 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.

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 in Park and Sabourian (2011), with three states, there are three key types of likelihood functions (henceforth LF) for signals: monotonic, hill-shaped and U-shaped. Recipients of a signal with an increasing LF will always buy, those with a decreasing LF will always sell. As in the above example, loosely, recipients of a signal with a hill-shaped LF will buy if prices fall a lot and sell if prices rise a lot (they “buck the trend” and act as *contrarians*). Finally, recipients of U-shaped LFs will buy if prices rise a lot but sell if prices fall a lot (they “follow the trend”).

⁷In our formal definition of herding and contrarianism, we benchmark behavior against the decision that the trader would take at the *initial* history, but the switching mechanism is akin to what we describe here.

2.2 Formal Definition of Herding

The illustrative example examined above is perhaps enough to motivate the experiment performed in this paper, but it may also be useful to provide a more formal description of the theory which must begin with a 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) his expectation is below the market expectation at the initial history, (ii) he 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) his expectation is below the market expectation at the initial history, (ii) he buys at history H_t , and (iii-c) prices at H_t are lower than at H_1 . Contrarian selling is analogous.*

Both with buy herding and buy contrarianism, a trader prefers to sell 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.

These definitions also capture well-documented financial market trading behavior. The herding definition is a formalization of the idea of rational momentum trading. Contrarianism has a mean-reversion flavor. Both momentum and contrarian trading have been analyzed extensively in the empirical literature and have been found to generate abnormal returns over some time horizon.⁸ Our analysis provides evidence for momentum and mean reversion behavior in an environment where such behavior can be rational.

⁸In the empirical literature, contrarian behavior is found to be profitable in the very short run (1

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 (2011) (which we implement), and capture the social learning (learning from others), imitative aspect of behavior for individual traders that is implied by the notion of herding from the earlier literature.⁹

2.3 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. Noise traders play an important role in generating uncertainty: for example the second trader to act is never certain of the motivation of the first trader. This exogenously generated uncertainty can also play a role in making herding or contrarianism more likely through an unusual sequence of noise trades and so captures the idea that exogenous noise can kick-start a herd. The remaining traders are informed and receive one of three signals, S_1, S_2, S_3 .

All past 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. Since the entry order is exogenous and since traders can act only once, they face an inherently static decision. They thus 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. In the setting used in this experiment, a buy increases the price, a sale decreases the price. Traders can thus perfectly infer past actions from prices and they can compute every type's optimal action at any point in the past.¹⁰

week and 1 month, Jegadeesh (1990) and Lehmann (1990)) and in the very long run (3-5 years, Bondt and Thaler (1985))). Momentum trading is found to be profitable over the medium term (3-12 months, Jegadeesh and Titman (1993)) and exceptionally unprofitable beyond that (the 24 months following the first 12, Jegadeesh and Titman (2001)). Sadka (2006) studies the effect of systematic liquidity risk on returns to momentum trading and links one of the liquidity risk components to informed trading.

⁹Specifically, we follow Brunnermeier (2001)'s (Ch. 5) description of herding as a situation in which "an agent imitates the decision of his predecessor even though his own signal might advise him to take a different action" and we consider the behavior of a *particular* signal type by looking at how the history of past trading can induce a trader to change behavior and trade against his private signal.

¹⁰To assess traders' actions as rational or irrational, it is thus unimportant how exactly prices are

The following are the possible shapes of likelihood functions (LF):

increasing : $\Pr(S|V_1) < \Pr(S|V_2) < \Pr(S|V_3)$; *decreasing* : $\Pr(S|V_1) > \Pr(S|V_2) > \Pr(S|V_3)$;
U-shaped : $\Pr(S|V_i) > \Pr(S|V_2)$ for $i = 1, 3$; *Hill-shaped* : $\Pr(S|V_i) < \Pr(S|V_2)$ for $i = 1, 3$.

For the results in our paper it is also important whether the likelihood of a signal is higher in one of the extreme states V_1 or V_3 relative to the other extreme state. We thus define the *bias* of a signal S as $\Pr(S|V_3) - \Pr(S|V_1)$. A U-shaped LF with a negative bias, $\Pr(S|V_3) - \Pr(S|V_1) < 0$, will be labeled as an nU-shaped LF and a pU-shaped LF with a positive bias, $\Pr(S|V_3) - \Pr(S|V_1) > 0$, will be labeled as a pU-shaped LF. Similarly, we use nHill (pHill) to describe a hill-shaped LF with a negative (positive) bias.¹¹ A signal is called *monotonic* if its LF is either increasing or decreasing and *non-monotonic* if its LF is hill or U-shaped.

In all treatments, signal S_1 is decreasing, signal S_3 is increasing. 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 varies by treatment and we will study settings with all possible shapes.

The following is a corollary to Theorem 1 in Park and Sabourian (2011) and it formalizes the ideas conferred in the illustrative example above. The proof is in appendix. (“ S herds” is to be read as “ S herds with positive probability”).

Theorem (Herding and Contrarian Behavior (Park and Sabourian (2011)))

- (a) A signal type with a decreasing LF always sells.
- (b) A signal type with an increasing LF always buys.
- (c) A signal type buy- (sell-) herds if and only if his LF is negative (positive) U-shaped.
- (d) A signal type acts as a buy-(sell-) contrarian if and only if his LF is negative (positive) hill-shaped.

Adjustments for the Experiment. The full underlying theory in Avery and Zemsky (1998) and Park and Sabourian (2011) has two prices: one at which the market maker buys (the bid) and one at which the market maker 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 necessary conditions from Park and Sabourian (2011) for herding and contrarianism are also sufficient. 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

computed: all that is needed is that traders can correctly infer past actions. Thus it suffices if prices increase for a buy and decrease for a sale.

¹¹The more compact nU/nHill notation is used in the proof of Theorem 1 found in the appendix. In the main text we will stick with the term “negative U-shape” or similar.

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, Oechssler, and Roeder (2005) highlight, bid-ask treatments usually offer no additional insights.

2.4 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 decreasing LF (such as S_1) always sell; (2) signal types with a increasing LF (such as S_3) always buy; (3) the behavior of signal types with a non-monotonic LF (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 an informed trader with the positive U-shaped signal S_2 .¹² No-trades in our setting do not affect the

¹²Although 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”

price, which is synonymous to a no-trade not revealing information. From the lack of a price-movement, subjects can thus infer that there was a no-trade. 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.

Hypothesis 3 (Rational Herding and Contrarianism)

CONTRARIANISM: *Subjects who receive a signal with a negative (positive) hill-shaped LF 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 LF 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, traders with monotonic and hill-shaped signals can herd in the sense of our definition, and those with monotonic and U-shaped signals 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.

In 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, Oechssler, and Roider (2005), or Cipriani and Guarino (2005)).

The decision problem in particular for U and hill-shaped types is much more difficult than for any other type as they must change behavior, and it would almost be surprising if they always do this correctly. A weaker but equally important test of our theory is thus whether receiving a U or hill-shaped signal increases the chance that one acts as a herder and contrarian respectively. For if so, we have evidence that the theory has relevance in

(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 and it ensures that rational herding and contrarianism can only occur for U and hill-shaped types respectively. As explained above, the precise price-setting rule is not essential for assessing rationality: all that is needed is that the price increases after buys and decreases after sales so that subjects can infer past actions (which is important so that they can update correctly). In section 8 and in the appendix we also explore whether alternative updating rules, which may include accounting for others’ irrationality, may do a better job at explaining the data.

the sense that the right signals affect behavior in a significant manner. Hypothesis 4 is thus the most important test as it tests whether the role attributed to underlying signals by the theory is justified.

Hypothesis 4 (The Impact of Signals on Herding and Contrarian Behavior)

Receiving a U-shaped signal significantly increases the chance of acting as a herder compared to receiving any other type of information. Receiving a hill-shaped signal significantly increases the chance of acting as a contrarian compared to receiving any other type of information.

Prior experimental work by Drehmann, Oechssler, and Roider (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 LF sell, and, likewise, as prices fall traders with a decreasing LF buy. If instead they tend to act as herders, then as prices rise, traders with an decreasing LF buy, and, likewise, as prices fall traders with a increasing LF sell. Either type of behavior contradicts the theory.

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

3 Experimental Design

In addition to the information given in what follows, the 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.

3.1 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. 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 in the appendix we analyze how a model of errors as an incorrect assessment of the degree of noise trading performs given our data. 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.¹³ 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 the current exchange rate).

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 prior probabilities that the asset value was high, medium or low.¹⁴ 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.¹⁵ In practice, subjects tended to act before the allotted time was over with very few timeouts. A sequence of trading opportunities $t = 1, 2, 3, \dots$ produced a history of actions and prices, $H_t = \{(a_1, P_1), \dots, (a_{t-1}, P_{t-1})\}$ 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

¹³For examples of such sheets please see the appendix.

¹⁴The 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.

¹⁵Subjects 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.

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).¹⁶ 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.

3.2 Treatments

Following Section 2, 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 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:¹⁷

- 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.

¹⁶In the tabulated regressions we correct for possible within-session correlations by clustering the standard errors at the session level. In untabulated regressions we further checked for the effect of different number of subjects in the sessions by (a) using fixed effects for the sessions and (b) using the number of participants as a control. The results were robust with either approach.

¹⁷Note 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.

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.¹⁸ Across all sessions 390 traders were allocated the S_1 signal, 550 the S_2 signal and 407 the S_3 signal.¹⁹

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.²⁰ 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,

¹⁸One 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.

¹⁹The 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. 24% claimed to have carried out numerical calculations.

²⁰Without 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

presumably in hope of making a profit on that one share. In this sense, a hold is a positive signal albeit weaker than a buy. Therefore, a pass can be counted as a “weak buy”.

As outlined in Section 2 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. 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 section 8 and in the appendix: in short, we found that including risk aversion did 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 1 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.²¹ 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.

²¹Anderson 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 1 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.²²

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).²³

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 traders with U-shaped signals 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 traders with hill-shaped signals can act rationally as contrarians. Yet U-shaped types can also *irrationally* herd in the sense of our definition. Such a situation arises when a trader with a negative U-shaped signal 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

²²In section 5 we will highlight the persistency in their behavior.

²³Future 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 signal traders 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 2 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, Oechssler, and Roider (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 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: Regression Analysis

The analysis thus far revealed that herding and contrarianism often does not arise when it is theoretically mandated. A second, important implication of the theory (by virtue of its economic intuition) is that information matters in the sense that recipients of U and hill-shaped signals should be more likely to act as herders and contrarians than any other signal type. We thus 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:

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

where the dependent variables $\text{herd}_{i,t}$ and $\text{contra}_{i,t}$ 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 $\text{u-shape}_{i,t}$ and $\text{hill shape}_{i,t}$ are signal dummies that are 1 if the individual received a U-shaped or hill-shaped signal (for herding or contrarianism respectively). Given the random assignment of signals and time slots, we can assume that $E[\text{u-shape}_{i,t} \cdot \epsilon_{i,t}] = 0$ and $E[\text{hill shape}_{i,t} \cdot \epsilon_{i,t}] = 0$, the main identifying assumption.

We estimated the model by Logit and report the marginal effects at the mean.²⁴ Standard errors are clustered by sessions to control for group-specific correlations.²⁵

²⁴We 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. Marginal effects were computed with Stata's built-in tools.

²⁵In 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 observable), 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.

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.

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 4 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 of 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 4 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

$$\text{trade}_{i,t} = \alpha + \beta \Delta \text{price}_{i,t} + \epsilon_t, \quad (2)$$

where $\text{trade}_{i,t}$ is a dummy that is 1 if there is a buy or pass, and 0 when there is a sale,²⁶ and the independent variable $\Delta \text{price}_{i,t}$ is the percentage change of the price from 100, i.e. the price at the time of the trade divided by 100 and subtracting 1. Since the time slot in which people are allowed to trade is assigned at random, we can assume that $E[\Delta \text{price}_{i,t} \cdot \epsilon_{i,t}] = 0$.

We estimated the model by Logit²⁷ 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 5 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

²⁶We 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.

²⁷Ordinary least squares regressions with and without trader fixed effects also yield the same insights.

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.²⁸ We thus conclude that there is enough evidence for contrarian behavior of the S_3 types and we have no evidence for herding.

Finding 5 (Price Changes and the Decisions of Types with Monotonic LFs)

Theory predicts that the price does not affect the decision of signal types with monotonic LFs. 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 repeated 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 4 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

$$\text{herd}_i = \alpha + \beta_1 \text{U shape}_i + \beta_2 \text{passer}_i + \beta_3 \text{U shape}_i \times \text{passer}_i + \epsilon_i, \quad (3)$$

$$\text{contra}_i = \alpha + \beta_1 \text{hill shape}_i + \beta_2 \text{passer}_i + \beta_3 \text{hill shape}_i \times \text{passer}_i + \epsilon_i. \quad (4)$$

²⁸In 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.

The dependent variables herd_i and contra_i are the herding and contrarian dummies from the equations in (1), U shape_i and 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 $\text{U shape}_i \times \text{passer}_i$ and $\text{hill shape}_i \times \text{passer}_i$ 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 inter-group 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-shaped nor the hill-shaped 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 6 summarizes the results of our estimation.²⁹

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 or 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 Alternative Theories of Behavior

We argued in Sections 4 and 5 that our findings are generally supportive of the theory. Yet there are still challenges, for instance, regarding the apparent contrarianism by S_3

²⁹Non-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 of these findings indicate that future research into behavioral heterogeneity between traders, along the lines of Ivanov, Levin, and Peck (2009), is warranted.

types for high prices or the lack of herding and (to a lesser extent) contrarianism by S_2 types 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.³⁰ In the appendix we consider a range of alternative, behavioral theories to see if we can find a better fit with our data. In turn we consider: omitting decisions that are within an ϵ error-region, risk aversion, loss aversion, non-Bayesian updating and models of error correction (in the spirit 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)). These alternative hypotheses usually depend on some parameter(s). Our general 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.

Finding 7 (Alternative Theories of Behaviour) *Although some of the alternative theories that we consider in the appendix provide reasonable approximations of behavior, none perform significantly better than the rational theory, and some perform worse.*

9 Conclusion

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 (2011). In the second part, we determine whether people behave in the spirit of Avery and Zemsky (1998) and Park and Sabourian (2011) 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 part of our analysis we focus on understanding the link between information and observed herding and contrarian behavior. Prior experimental work on

³⁰I., Levin, and Peck (2012) provide a nice examination of deviations from Nash equilibrium and discuss the role of noise and cognitive biases. Their setting with endogenous timing of actions and fixed prices complements ours which has exogenous timing of actions and moving prices.

sequential trading outlined in Section 1 has shown that subjects have a “natural” tendency to act as contrarians.³¹ Notwithstanding this finding, 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.

Our findings have 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-shaped and hill-shaped signals.³² We can envisage situations where this is both likely and predictable. For instance, consider a situation where the great majority of traders believe that prices will change, either rising or falling: perhaps a regulatory decision is about to be made that will either greatly benefit or greatly harm a firm, or an electoral decision is imminent and this will advantage or disadvantage the firm. In either case the fraction of traders with U-shaped signals might be arbitrarily close to 1. If such a situation arises in markets, then even the low probability of rational herding may lead to economically meaningful price swings.

References

- ALEVY, J. E., M. S. HAIGH, AND J. A. LIST (2007): “Information Cascades: Evidence from a Field Experiment with Financial Market Professionals,” *Journal of Finance*, LXII(1), 151–180.
- ANDERSON, L., AND C. HOLT (1997): “Informational Cascades in the Laboratory,” *American Economic Review*, 87, 847–862.
- AVERY, C., AND P. ZEMSKY (1998): “Multi-Dimensional Uncertainty and Herd Behavior in Financial Markets,” *American Economic Review*, 88, 724–748.
- BIAIS, B., D. HILTON, K. MAZURIER, AND S. POUGET (2005): “Judgemental Overconfidence, Self-Monitoring, and Trading Performance in an Experimental Financial Market,” *The Review of Economic Studies*, 72(2), 287–312.
- BONDT, W. F. M. D., AND R. THALER (1985): “Does the Stock Market Overreact?,” *The Journal of Finance*, 40(3), 793–805.

³¹It remains to be explained why this happens, for example are there components of the experimental design, such as bounded asset values, that affect this phenomenon?

³²Formally, their fractions in the population depend on the specific details of the signal likelihood function.

- BOORTZ, C., S. JURKATIS, S. KREMER, AND D. NAUTZ (2013): “Herding in financial markets: Bridging the gap between theory and evidence,” SFB 649 Discussion Paper 2013-036, Humboldt University Berlin, Berlin.
- BRUNNERMEIER, M. K. (2001): *Asset Pricing under Asymmetric Information – Bubbles, Crashes, Technical Analysis, and Herding*. Oxford University Press, Oxford, England.
- CHORDIA, T., R. ROLL, AND A. SUBRAHMANYAM (2002): “Order imbalance, liquidity, and market returns,” *Journal of Financial Economics*, 65(1), 111–130.
- CIPRIANI, M., AND A. GUARINO (2005): “Herd Behavior in a Laboratory Financial Market,” *American Economic Review*, 95(5), 1427–1443.
- (2009): “Herd Behavior in Financial Markets: A Field Experiment with Financial Market Professionals,” *Journal of the European Economic Association*, 7(1), 206–233.
- (2014): “Estimating a Structural Model of Herd Behavior in Financial Markets,” *The American Economic Review*, 104(1), 224–251.
- COSTA-GOMES, M., V. CRAWFORD, AND B. BROSETA (2001): “Cognition and Behavior in Normal-Form Games: An Experimental Study,” *Econometrica*, 69, 1193–1235.
- DREHMANN, M., J. OECHSSLER, AND A. ROIDER (2005): “Herding and Contrarian Behavior in Financial Markets: An Internet Experiment,” *American Economic Review*, 95(5), 1403–1426.
- GLOSTEN, L., AND P. MILGROM (1985): “Bid, Ask and Transaction Prices in a Specialist Market with Heterogenously Informed Traders,” *Journal of Financial Economics*, 14, 71–100.
- I., A., D. LEVIN, AND J. PECK (2012): “Behavioral Biases in Endogenous-Timing Herding Games: an Experimental Study,” *Journal of Economic Behavior and Organization*, 87, 25–34.
- IVANOV, A., D. LEVIN, AND J. PECK (2009): “Hindsight, Foresight, and Insight: An Experimental Study of a Small-Market Investment Game with Common and Private Values,” *American Economic Review*, 99(4), 1484–1507.
- JEGADEESH, N. (1990): “Evidence of Predictable Behavior of Security Returns,” *The Journal of Finance*, 45(3), 881–898.

- JEGADEESH, N., AND S. TITMAN (1993): “Returns to Buying Winners and Selling Losers: Implications for Stock Market Efficiency,” *The Journal of Finance*, 48(1), 65–91.
- (2001): “Profitability of Momentum Strategies: An Evaluation of Alternative Explanations,” *The Journal of Finance*, 56(2), 699–720.
- KAHNEMAN, D., AND A. TVERSKY (1979): “Prospect Theory: An Analysis of Decision under Risk,” *Econometrica*, 47(2), 263–292.
- LEHMANN, B. N. (1990): “Fads, Martingales, and Market Efficiency,” *The Quarterly Journal of Economics*, 105(1), 1–28.
- McKELVEY, R., AND T. PALFREY (1995): “Quantal Response Equilibria for Normal Form Games,” *Games and Economic Behavior*, 10, 6–38.
- (1998): “Quantal Response Equilibria for Extensive Form Games,” *Experimental Economics*, 1, 9–41.
- PARK, A., AND H. SABOURIAN (2011): “Herding and Contrarian Behavior in Financial Markets,” *Econometrica*, 79, 973–1026.
- PARK, A., AND D. SGROI (2012): “Herding, Contrarianism and Delay in Financial Market Trading,” *European Economic Review*, 56, 1020–1037.
- PLOTT, C. R., AND S. SUNDER (1988): “Rational Expectations and the Aggregation of Diverse Information in Laboratory Security Markets,” *Econometrica*, 56(5), 1085–1118.
- SADKA, R. (2006): “Momentum and post-earnings-announcement drift anomalies: The role of liquidity risk,” *Journal of Financial Economics*, 80(2), 309 – 349.
- SHACHAT, J., AND A. SRINIVASAN (2011): “Informational Price Cascades and Non-Aggregation of Asymmetric Information in Experimental Asset Markets,” Working paper, National University of Singapore.
- TVERSKY, A., AND D. KAHNEMAN (1992): “Advances in Prospect Theory: Cumulative Representation of Uncertainty,” *Journal of Risk and Uncertainty*, 5, 297–323.

Table 1
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 2
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 theoretically mandated		20	34	16	20
		109	109	26	26
percent that arises rationally		18%	31%	62%	77%

Table 3
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 4
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. For all tables that follow, standard errors are in parentheses, ** indicates significance at the 5% level, *** at the 1% level. Constants were included in the regression but are not reported for brevity.

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

Table 5
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 4. Constants were included in the regression but are not reported for brevity.

Dependent Variable: Buy Indicator						
Independent Variable: price change in %						
	Increasing LFs			Decreasing LFs		
	S_3	$S_2 \nearrow$	all $\nearrow$	S_1	$S_2 \searrow$	all $\searrow$
<i>Panel A: Passes are buys</i>						
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
<i>Panel B: Passes are omitted</i>						
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 6
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 4 and we omit constants from the report.

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

APPENDIX

Appendix A contains the proof of the main theorem in the paper which is adapted from Park and Sabourian (2011). Appendix B is a thorough discussion of the alternative theories mentioned in the main text in section 8. Appendix C provides an experimental timeline. Appendices D, E and F (respectively) provide the full set of subject instructions, the questionnaire and some details about the software. Tables from Appendix B are provided at the end. 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.

A Proof of the Theorem

The theorem follows directly from the results in Park and Sabourian (2011); we repeat the arguments here for the reader's convenience. We first show the following three properties.

Lemma 1 (i) For any S , t and history H^t , $E[V|S, H^t] - E[V|H^t]$ has the same sign as

$$q_1^t q_2^t [\Pr(S|V_2) - \Pr(S|V_1)] + q_2^t q_3^t [\Pr(S|V_3) - \Pr(S|V_2)] + 2q_1^t q_3^t [\Pr(S|V_3) - \Pr(S|V_1)]. \quad (5)$$

(ii) For any signal S , $E[V|S]$ is less than $E[V]$ if and only if S has a negative bias, and $E[V|S]$ is greater than $E[V]$ if and only if S has a positive bias.

(iii) If $E[V|H^t] > E[V]$ then $q_3^t > q_1^t$ and if $E[V] > E[V|H^t]$ then $q_1^t > q_3^t$.

Proof: (i) Observe first that

$$E[V|S, H^t] - E[V|H^t] = \mathcal{V} q_2^t \left(\frac{\Pr(S|V_2)}{\Pr(S)} - 1 \right) + 2\mathcal{V} q_3^t \left(\frac{\Pr(S|V_3)}{\Pr(S)} - 1 \right).$$

The RHS of the the above equality has the same sign as

$$\begin{aligned} & q_2^t \left(\Pr(S|V_2) \sum_j q_j^t - \sum_j \Pr(S|V_j) q_j^t \right) + 2 q_3^t \left(\Pr(S|V_3) \sum_j q_j^t - \sum_j \Pr(S|V_j) q_j^t \right) \\ &= q_1^t q_2^t (\Pr(S|V_2) - \Pr(S|V_1)) + q_2^t q_3^t (\Pr(S|V_2) - \Pr(S|V_3)) \\ & \quad + 2 q_3^t (q_1^t (\Pr(S|V_3) - \Pr(S|V_1)) + q_2^t (\Pr(S|V_3) - \Pr(S|V_2))). \end{aligned}$$

(ii) By the symmetry assumption on the priors ($q_1^1 = q_3^1$), the (5) is negative (positive) at $t = 1$ if and only if $(\Pr(S|V_3) - \Pr(S|V_1))(q_2^1 + 2q_1^1)q_3^1$ is less (greater) than 0; the latter is equivalent to S having a negative (positive) bias.

(iii) The claim follows from $E[V|H^t] - E[V] = \mathcal{V}[(1 - q_1^t - q_3^t) + 2q_3^t] - \mathcal{V} = \mathcal{V}(q_3^t - q_1^t)$.

Proofs of the Theorem (a) and (b): These follow immediately from part (i) of the lemma because for an increasing LF, all terms in (5) are always positive, for a decreasing LF they are always negative.

Proof of the Theorem (c) and (d): “ $\Leftarrow$ ”: We will prove the “only if” parts only for the case of buy herding and buy contrarian; the proof for the sell cases are analogous. Thus suppose that S buy herds or acts as a buy contrarian at some H^t . The proof proceeds in several steps.

Step 1: S has a negative bias: Buy herding and buy contrarian imply $E[V|S] < \text{bid}_1$. Since $\text{bid}_1 < E[V]$ we must have $E[V|S] < E[V]$. Then by part (ii) of the above lemma, S has a negative bias.

Step 2: $(\Pr(S|V_1) - \Pr(S|V_2))(q_3^t - q_1^t) > 0$: It follows from the definition of buy herding and buy contrarian that $E[V|S, H^t] > \text{ask}^t$. Since $E[V|H^t] < \text{ask}^t$ we must have $E[V|S, H^t] > E[V|H^t]$. By part (i) of the above lemma, this implies that (5) is positive at H^t . Also, by the negative bias (Step 1), the third term in (5) is negative. Therefore, the sum of the first two terms in (5) is positive: $q_3^t(\Pr(S|V_3) - \Pr(S|V_2)) + q_1^t(\Pr(S|V_2) - \Pr(S|V_1)) > 0$. But this means, by negative bias, that $(\Pr(S|V_1) - \Pr(S|V_2))(q_3^t - q_1^t) > 0$.

Step 3a: If S buy herds at H^t then S is nU-shaped: It follows from the definition of buy herding that $E[V|H^t] > E[V]$. By part (iii) of the above lemma, this implies that $q_3^t > q_1^t$. Then it follows from Step 2 that $\Pr(S|V_1) > \Pr(S|V_2)$. Also, since S buy-herds, by parts (a), S cannot have a decreasing csd and we must have $\Pr(S|V_2) < \Pr(S|V_3)$. Thus, S is nU-shaped.

Step 3b: If S acts as a buy contrarian at H^t then S is nHill shaped. It follows from the definition of buy contrarian that $E[V|H^t] < E[V]$. By part (iii) of the above Lemma, this implies that $q_3^t < q_1^t$. But then it follows from Step 2 that $\Pr(S|V_1) < \Pr(S|V_2)$. Since by Step 1 S has a negative bias, we have $\Pr(S|V_2) > \Pr(S|V_1) > \Pr(S|V_3)$. Thus S is nHill-shaped.

“ $\Rightarrow$ ”: To see the “if” part notice first that since by assumption S has a negative bias, it follows from part (ii) of the lemma that S sells at the initial history. Also, since S is U-shaped we have $\Pr(S|V_3) > \Pr(S|V_2)$. Therefore, by part (i) of the lemma, there exists some $\eta > 0$ such that the second term in (5) always exceeds η .

Herding requires that the negative terms in (5) are sufficiently small. Suppose that there exists a history H^t such that $q_1^t/q_3^t < 1$ and $\frac{q_1^t}{q_3^t} + \frac{2q_1^t}{q_2^t} < \eta$. Then by the former inequality and (iii) of the Lemma we have $E[V|H^t] > E[V]$. Also, since the sum of the first and the third term in (5) is greater than $-q_2q_3(\frac{q_1^t}{q_3^t} + \frac{2q_1^t}{q_2^t})$, it follows from $\frac{q_1^t}{q_3^t} + \frac{2q_1^t}{q_2^t} < \eta$ that the sum must also be greater than $-\eta$. This, together with the second term in (5) exceeding η , implies that (5) is greater than zero, and hence S must be buying at H^t .

The proof for buy contrarianism is analogous: by the same reasoning as above S sells at H^1 . Also, since S has a hill shape we have $\Pr(S|V_1) > \Pr(S|V_2)$. Therefore, by part (i) of the Lemma, there exists some $\eta > 0$ such that the first term in (5) always exceeds η .

Suppose that there exists a history H^t such that $q_3^t/q_1^t < 1$ and $\frac{q_3^t}{q_1^t} + \frac{2q_3^t}{q_2^t} < \eta$. Then by the former inequality and (iii) of the Lemma we have $E[V|H^t] < E[V]$. Also, since the sum of the second and the third term in (5) is greater than $-q_1q_2(\frac{q_3^t}{q_1^t} + \frac{2q_3^t}{q_2^t})$, it follows

from $\frac{q_3^t}{q_1^t} + \frac{2q_3^t}{q_1^t} < \eta$ that the sum must also be greater than $-\eta$. Since the first term in (5) exceeds η , this implies that (5) is greater than zero, and hence S must be buying at H^t .

What remains to be shown is that the above mentioned history exists. Consider any arbitrary history H^t and any two values $V_l < V_h$. By (a) type S_1 always sells, by (b) type S_3 always buys. There are thus two cases for a buy at H^t : either only S_3 types buy or S_2 and S_3 types buy. In the former case, $\beta_i^t = \gamma + \mu \Pr(S_3|V_i)$. As S_3 is strictly increasing, there exists $\epsilon > 0$ such that $\beta_h^t - \beta_l^t > \epsilon$. In the latter case,

$$\begin{aligned}\beta_h^t - \beta_l^t &= \mu (\Pr(S_3|V_h) + \Pr(S_2|V_h) - \Pr(S_3|V_l) - \Pr(S_2|V_l)) \\ &= \mu (1 - \Pr(S_1|V_h) - (1 - \Pr(S_1|V_l))) = \mu (\Pr(S_1|V_l) - \Pr(S_1|V_h)).\end{aligned}$$

Since S_1 is strictly decreasing, there exists an $\epsilon > 0$ such that $\beta_h^t - \beta_l^t > \epsilon$.

By a similar reasoning it can be shown that there must exist $\epsilon > 0$ so that $\sigma_l^t - \sigma_h^t > \epsilon$.

Consequently, the probability of a buy is uniformly increasing in the liquidation value at any date and history, i.e. for some $\epsilon > 0$, $\beta_j^t > \beta_i^t + \epsilon$ for any $j > i$ and any t . Similarly, the probability of a sale is uniformly decreasing in the liquidation value at any H^t , i.e. for some $\epsilon > 0$, $\sigma_i^t > \sigma_j^t + \epsilon$ for any $j > i$ and any t .

Since $q_i^{t+1}/q_j^{t+1} = (q_i^t \beta_i^t)/(q_j^t \beta_j^t)$ when there is a buy at date t , it follows that a buy reduces *both* q_1^t/q_2^t and q_1^t/q_3^t uniformly. Thus a sufficiently large number of buys induces the described histories that allow buy herding.

Similarly, with a sale $q_i^{t+1}/q_j^{t+1} = (q_i^t \sigma_i^t)/(q_j^t \sigma_j^t)$. Thus, a sale reduces *both* q_3^t/q_1^t and q_3^t/q_2^t uniformly. Thus, a sufficiently large number of sales induces the described histories needed for buy contrarianism. This argument completes the proof. $\square$

B Alternative Explanations for Trading Behavior

B.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 6 (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 7 (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 8 (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 9 (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.³³ We will thus examine whether error correction formulations will help us understand the data better.

Hypothesis 10 (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 6 to 10 will be analyzed in the following series of subsections.

³³Notions 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.

B.1.1 Robustness to Small Errors

Hypothesis 6 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.

Finding 8 (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 6.*

B.1.2 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 7).³⁴ 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

³⁴With 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.

the behavior by S_2 types. Note that as ρ falls, we move closer to risk neutrality. Table 7 displays the results for some select parameter values that are indicative of the general tendencies in the data.

Finding 9 (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 7.*

B.1.3 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 8 illustrates this observation for the above parameters as well as for one other configuration.³⁵

Finding 10 (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 8 and conclude that loss aversion is not an important influence for behavior in our experiment.*

B.1.4 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

³⁵Arguably, 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.

prescribed at the initial history (e.g. with a negative U-shaped signal, 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); and 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 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 types 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 types 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 types’ 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. 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 $f(p)$ ³⁶ 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 employed properly in the conditional posterior expectation the transformation achieves

³⁶There are various other forms for these switches, e.g. non-symmetric switches where the effects are stronger (or weaker) for larger probabilities. The interpretation and implementation of such asymmetric shifts does, however, become difficult if not impossible with three states. Of the various possible specifications we only pick a few as the spirit of all re-scalings is similar: updating is slowed.

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

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

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 11 (Non-Bayesian Updating) *Having considered several forms of Hypothesis 9 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.*

B.1.5 Error Correction Provisions

To investigate Hypothesis 10, and inspired by level K reasoning (see Costa-Gomes, Crawford, and Broseta (2001)) and Quantal Response Equilibria (see McKelvey and Palfrey (1995) and McKelvey and Palfrey (1998)), we will contemplate an alternative specification for hampered updating in which agents do not trust 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 $\beta_i = .25/2 + .75 \cdot \Pr(S_3|V_i)$ (recalling that $.25/2$ is the probability of a noise buy). Now imagine that instead subjects believe that only fraction δ of the informed buyers act rationally and that the remaining $1 - \delta$ take a decision at random. Then the probability of a buy in state V_i becomes:

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

The task is 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 10): 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.³⁷ As with the simple error correction, we do not obtain a substantially better fit with the data, as can be seen in Table 10: 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 12 (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 10.*

B.1.6 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

³⁷Rather 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 (2009).

since the results did not differ, we choose not to report them here in detail.³⁸

Several studies (Drehmann, Oechssler and Roeder (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 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 6 to 9. 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.

³⁸For example, we considered whether traders chase short-term trends, but found that they do not.

C 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.

D Instructions

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

Instructions 1 (Welcome)

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

Instructions 2 (After Seated)

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

Instructions 3 (The Experimental Setting)

Next I will describe the experiment itself. You will be participating in a series of financial market trading exercises. There will be 7 trading rounds, and each round will last 3-4 minutes. There are [number of participants] participants in the room and everyone is involved in the same trading exercise. Your objective should be to take the most thorough decision possible in order to maximize the money you will make today. The general situation is the following: you are the stockholder of a company and have some cash in hand. Some event may happen to your company that affects the value of the company (for better or worse). You have a broker who provides you with his best guess. You then have to decide whether you want to buy an additional share of the company, whether you want to sell your share, or whether you want to do nothing. We will look at a variety of similar situations: each situation concerns a *different* company, and we will vary the information and the trading rules in each situation. Please note that the situation

described to you in each round is independent of that in any other round. *In other words, what you learned in round 1 tells you nothing about round 2, etc.* In the process of this session you may or may not generate virtual profits. Your trading activities will be recorded automatically; these activities determine your trading profits.

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

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

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

Instructions 4 (The Software)

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

1. It tells you about all the trades that occur during the round; you also see when a trade occurs and whether or not someone bought or sold a share. For your convenience, there is a graph that plots the sequence of prices.
2. Your screen also lists the current market price; people can either buy a share at this price or they can sell their share at this price.
3. In the case where we restrict the time when you can make a trade, a red bar will appear on the bottom of the screen to highlight the fact that you can trade. During this time the buy, sell and pass buttons will be available for your use, typically only once per round, though twice in the final 3 rounds.
4. There is also a box in which you receive some information from your "broker" which I will explain in a few moments.
5. The screen includes a timer which indicates how many seconds have gone past during the round.
6. Finally, the screen updates itself whenever a trade is made.

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

Instructions 5 (Compensation)

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

I will next explain virtual profits. When you trade you will do so at 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 you 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.

E Questionnaire

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

About you

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

About your decisions today

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

11. Did you make any calculations? If so, which ones?

About the experiment

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

F The Software

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

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

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.³⁹

³⁹Further details about the software are available on request from the authors.

Table 7
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. The total number of decisions is in Table 4 from the main text.

	Total Number of wrong decisions CRRRA utility, $\gamma = 1$ (log-utility)			Total Number of wrong decisions CARA utility, $\rho = 1$			Total Number of wrong decisions CARA utility, $\rho = 0.01$		
	S_1	S_2	S_3	S_1	S_2	S_3	S_1	S_2	S_3
Treatment 1 negative hill shape	18 33%	58 58%	37 51%	46 84%	60 60%	14 19%	46 84%	58 58%	48 67%
Treatment 2 increasing	31 42%	60 67%	36 53%	60 81%	28 31%	10 15%	60 81%	30 33%	50 74%
Treatment 3 negative U shape	13 22%	69 73%	37 49%	48 82%	59 63%	6 8%	48 82%	18 19%	61 80%
Treatment 4 decreasing	32 57%	41 41%	33 45%	44 76%	74 75%	6 8%	44 76%	75 76%	65 88%
Treatment 5 positive U shape	29 43%	66 67%	32 49%	58 86%	29 30%	7 11%	58 86%	25 26%	55 85%
Treatment 6 negative hill shape	41 47%	51 61%	22 38%	66 76%	42 51%	6 10%	66 76%	47 57%	52 90%
Total number wrong percentage wrong	164 41%	345 61%	197 48%	322 70%	292 44%	49 10%	322 70%	253 39%	331 68%
Total model fit	48.8%			41.3%			56.5%		

Table 8
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 B.1.3. 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 7.

	Total Number of wrong decisions prospect theory, $\alpha = \beta = 0.8, \gamma = 1$			Total Number of wrong decisions prospect theory, $\alpha = \beta = 0.8, \gamma = 2.25$		
	S_1	S_2	S_3	S_1	S_2	S_3
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 9
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 B.1.4. 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 $\alpha = 25$		
	S_1	S_2	S_3	S_1	S_2	S_3	S_1	S_2	S_3
Treatment 1 negative hill shape	46 7 87%	61 37 62%	57 14 80%	46 7 87%	61 37 62%	57 14 80%	46 7 87%	61 37 62%	57 14 80%
Treatment 2 increasing	60 10 86%	56 28 67%	57 10 85%	60 10 86%	56 28 67%	57 10 85%	60 10 86%	56 28 67%	57 10 85%
Treatment 3 negative U shape	48 10 83%	31 62 33%	68 6 92%	48 10 83%	41 52 44%	68 6 92%	48 10 83%	57 36 61%	68 6 92%
Treatment 4 decreasing	44 12 77%	73 25 74%	65 6 92%	44 12 77%	73 25 74%	65 6 92%	44 12 77%	73 25 74%	65 6 92%
Treatment 5 positive U shape	58 9 87%	66 29 69%	56 7 89%	58 9 87%	66 29 69%	56 7 89%	58 9 87%	66 29 69%	56 7 89%
Treatment 6 negative hill shape	66 20 77%	60 22 73%	52 6 90%	66 20 77%	60 22 73%	52 6 90%	66 20 77%	60 22 73%	52 6 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 model fit	77.1%			77.9%			79.1%		

Table 10
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 B.1.5. 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 $\delta = .24$		
	S_1	S_2	S_3	S_1	S_2	S_3	S_1	S_2	S_3
Treatment 1 negative hill shape	46 7 87%	62 36 63%	46 25 65%	46 7 87%	61 37 62%	47 24 66%	46 7 84%	63 35 63%	46 25 64%
Treatment 2 increasing	60 10 86%	35 49 42%	50 17 75%	60 10 86%	38 46 45%	49 18 73%	60 10 81%	36 48 40%	49 18 72%
Treatment 3 negative U shape	48 10 83%	35 59 37%	57 17 77%	48 10 83%	59 35 63%	55 19 74%	48 10 82%	59 35 63%	58 16 76%
Treatment 4 decreasing	43 13 75%	74 25 75%	64 10 86%	43 13 75%	74 25 75%	65 9 88%	44 12 76%	74 25 75%	65 9 88%
Treatment 5 positive U shape	58 9 87%	38 57 40%	55 8 87%	58 9 87%	44 51 46%	55 8 87%	58 9 86%	42 53 43%	55 8 85%
Treatment 6 negative hill shape	66 20 77%	49 33 60%	52 6 90%	66 20 77%	48 34 59%	49 9 84%	66 20 76%	49 33 60%	49 9 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%





Figure 2
Information Sheet for negative U Shape

Round

Signals: **Case 2**

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

If the true effect will be POSITIVE then you receive

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

If the true effect will be NEGATIVE then you receive

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

If indeed the effect will be NO EFFECT then you receive

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





Figure 3
Information Sheet for negative hill shape

Round

Signals: **Case 1**

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

If the true effect will be POSITIVE then you receive

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

If the true effect will be NEGATIVE then you receive

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

If indeed the effect will be NO EFFECT then you receive

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





Figure 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