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

Andreas Park[†]
University of Toronto

Daniel Sgroi[‡]
University of Warwick

December 20, 2011

Abstract

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

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

[*]Financial support from the ESRC (grant number RES-156-25-0023), the SSHRC, the Leverhulme Trust and the Cambridge Endowment for Research in Finance (CERF) is gratefully acknowledged. Special thanks go to Annie Jekova and Malena Digiuni for expert research assistance. We thank seminar participants at the Universities of Cambridge, Toronto, Southampton, Warwick, and Zurich, and the Midwest Economic Theory Meetings 2008 and the Northern Finance Association Meeting 2009. We are grateful to Gustavo Bobonis, Ernst Fehr, Jordi Mondria, Jim Peck, and Xiadong Zhu for extensive discussions, and we thank an Associate Editor and an anonymous referee for helping us improve our work.

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

1 Introduction

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). Since World War II there have been only sixteen other days where 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 substantial changes in behaviour (from buying to selling or the reverse). Such behaviour and the resulting price volatility are often claimed to be inconsistent with rationally motivated trading and informationally efficient prices. Commentators invariably attribute dramatic swings to investors' animal instincts, which to most economists, is a deeply unsatisfying explanation. "Rational herding theory", on the other hand, provides new theoretical insights that show that seemingly erratic, switching back-and-forth behaviour can be driven by rational, information-based motives.

Herding theory was pioneered by Welch (1992), Bikhchandani, Hirshleifer, and Welch (1992) and Banerjee (1992) who highlight that rationality is no defence against the randomness of herd behaviour.[1] Put simply, a few early incorrect decisions, through a process of rational observation and inference, can have serious ramifications for all who follow.[2] A loose application of herding theory to financial market trading might suggest that early movements by visible traders can provide a catalyst for momentum trading, induce discontinuous price jumps in one direction or the other, and potentially leave share prices far from their fundamental value.

The early work on rational herding was not designed, however, for security market trading since it did not admit prices that react to actions, whereas one key feature of financial markets is that (efficient) market prices adjust after trades, with prices dropping after sales and rising after buys. Furthermore, those models that do admit moving prices restrict agents to act in a strict, exogenous sequence — they cannot decide when to

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

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

trade. Finally, the latter models also restrict traders to act only once. In one of the largest laboratory experimental studies of its type (with around 2000 trades spread over 6 treatments) we bring together all of these features: a model of financial trading with asymmetric information across traders, the potential for rational herding and rational contrarianism, the ability to time trades, and the ability to trade more than once.

To understand our contribution it is important to understand the history of the literature. It was first thought that when prices can adjust to actions information based herding was either not possible or economically irrelevant. A path-breaking paper by Avery and Zemsky (1998) introduced efficient prices to a sequential herding context, but showed that in a simple financial market-trading setting with two values herding is not possible because the market price always separates people with good and bad information so that the former always buy and the latter always sell. Experimental work has confirmed these predictions (Drehmann, Oechssler, and Roider (2005), Cipriani and Guarino (2005), Cipriani and Guarino (2009)). More recently, however, Park and Sabourian (2011) showed that with multiple states herding can arise and they gave conditions on information that must be satisfied to admit rational herding; they also described conditions for rational contrarianism.[3] They showed that (economically meaningful) herding can arise by traders who believe that extreme outcomes (big price rises or falls) are more likely than moderate ones, and that contrarianism can arise by traders who believe that moderate outcomes are more likely. The signals that generate these situations are, respectively, *U-shaped* and *hill-shaped*. An experiment by (Park and Sgroi (2009)) showed that this expanded theory has bite. We employ the information-based trading framework developed by Park and Sabourian (2011) in our experiment.

Next, in a market-trading environment where learning from others is important, the timing of actions may affect the possibility and extent of herding.[4] First, one of the key

[3]Rational contrarianism is often cited as an important force for the mean-reversion of asset prices, see Chordia, Roll, and Subrahmanyam (2002).

[4]The seminal paper which studies investment timing with multiple agents and a single irreversible action, but *without moving prices*, is Chamley and Gale (1994) and is also explored in Gale (1996). Their key message is that decision makers will act very quickly in response to their information, since waiting only makes sense when additional, new information arises. The first published experiment to consider herding in endogenous-time was Sgroi (2003), a close implementation of Chamley and Gale (1994). This framework was also examined experimentally in Ziegelmeyer, My, Vergnaud, and Willinger (2005). We complement this line of work by explicitly considering prices that adjust after actions. There are also two other related experimental papers and a theoretical paper. Bloomfield, O'Hara, and Saar (2005) study a financial market in which people can trade repeatedly throughout a trading day. The focus of their study is on the timing behaviour of informed traders and on their choice of limit or market orders depending on the passage of time. They do not employ information that could (theoretically) trigger herding or contrarianism. Ivanov, Levin, and Peck (2009) implement Levin and Peck (2008), which is a model of fixed capital (green-field), non-financial investments, and they develop important insights into the timing behaviour of people's investment choices. Their setting does not, however, consider moving prices. Finally, Smith (2000) studies endogenous timing theoretically in a single trader environment and

features of real-world financial frenzies is the clustering of actions in time, a phenomenon that cannot be examined when timing is not considered. Second, one can imagine that removing yet another friction from sequential trading models may make informational herding a non-issue. Alternatively, one can imagine that herding becomes more pronounced as those with herding signals delay their actions and then rush in eventually. Our experiment can thus shed light on the impact of the endogenous timing of actions on herding. We identify systematic effects caused by information across treatments and participants that are qualitatively in line with theory on the direction of trades, with marginal effects of information that are stronger relative to exogenous timing setups. In particular we see that contrarianism is caused by hill-shaped signals, that herding is caused by U-shaped signals, and that there is a separation of timing of trades across time by which traders with clearly positive and negative information trade systematically before those with U-shaped information. We also identify a new stylized fact in that traders cluster their trading in time, thus complementing clusters in action; there is, however, no evidence that this clustering is information- or herding-driven. In some experimental treatments, we also explore how the ability to trade twice affects behaviour. For lack of a theory these results are much more difficult to interpret. However, behaviour is qualitatively similar to the single trade treatments, except that trading occurs systematically earlier.

Overview. Section 2 provides a formal definition of herding and contrarianism. Section 3 outlines the guiding theoretical framework and develops qualitative hypotheses. Section 4 examines the design of the experiment and lists the different treatments. Section 5 studies the impact of information on the decision of the trade-direction. Section 6 analyzes the impact of information on the absolute and relative timing of actions (in particular, on clusters). Section 7 studies the differential implications of the two vs. one trade settings. Section 8 summarizes the key findings and concludes. The supplementary appendix outlines examinations of alternative explanations, a discussion of the role of information theory on timing, the subject instructions, and the explicit parameter values.

2 Definition of Herding and Contrarian Behavior

In the literature there are several definitions of herding. We follow Avery and Zemsky (1998) and Park and Sabourian (2011) whose definition focuses on the social learning (learning from others) aspect of behaviour for individual traders that is implied by the notion of herding from the earlier literature. Specifically, this definition follows 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

offers some qualitative predictions as to which sort of information induces rapid decision making.

different action" and we consider the behaviour of a *particular* signal type by looking at how the history of past trading can induce a trader to change behaviour and trade against his private signal. This definition has been used in other experimental work on social learning in financial markets (see, for instance, Cipriani and Guarino (2005) or Drehmann, Oechssler, and Roider (2005)) and it further captures the idea of rational momentum trading, a well-documented financial market trading phenomenon.

Definition *Herding.* A trader with signal S buy herds in period t at history H^t if and only if (i) he would sell at the initial history H^1, (ii) he buys at history H^t, and, (iii-h) prices at H^t are higher than at H^1. Sell herding at history H^t is defined analogously. *Contrarianism.* A trader with signal S engages in buy-contrarianism in period t at history H^t if and only if (i) he would sell at the initial history H^1, (ii) he buys at history H^t, and, (iii-c) prices at H^t are lower than at H^1. Sell contrarianism at history H^t is defined analogously.

Both with buy-herding and buy-contrarianism, the trader with signal S prefers to sell at the initial history, before observing other traders' actions (condition (i)), but prefers to buy after observing the history H^t (condition (ii)). The key differences between buy-herding and buy-contrarianism are conditions (iii-h) and (iii-c). The former the price to rise at history H^t so that a change of action from selling to buying at H^t is *with* the general movement of the prices (crowd), whereas the latter condition requires the public expectation to have dropped so that a trader who buys at H^t acts *against* the movement of prices.

In our experiment each trader receives a private signal, which is one of three possible signals (S_1, S_2, S_3). It is important to note that herding according to our definition does not imply that, after some history, all traders take the same action irrespective of their private signal. Such a situation would imply an informational cascade and would be inconsistent with moving prices and an informationally efficient financial market. To see why, consider the role of prices as reflecting the information contained in the traders' actions. If all informed types act alike then their actions would be uninformative, and as result, prices would not move. Therefore, such uniformity of behaviour cannot explain prices movements, which is a key feature of financial markets. Moreover, if the uniform action involves trading, then a large imbalance of trades would accumulate without affecting prices — contrary to common empirical findings.[5]

Our definition is the same as that commonly employed in the literature on informational herding, which allows comparability of our results. There are, however, other plausible definitions of herding and contrarianism.[6] Instead of defining herding and con-

[5]See, for instance, Chordia, Roll, and Subrahmanyam (2002).

[6]See also Brunnermeier (2001), Chamley (2004), or Vives (2008) for related definitions of herding,

trarianism as switches of behaviour relative to price movements since the beginning of trading as in (*iii-h*), one could define herding and contrarianism as switches relative to the most recent price movements or actions. For instance, someone with negative information would engage in buy herding if she buys after the price rose by x units or after observing a sequence of y buys. The difficulty of such a classification is to find the indisputable criterion for the "right" number of recent actions or the right size of recent price movements. One could also define herding and contrarianism relative to the majority action. Herding would then be defined as a switch to adopting the majority action; contrarianism would be defined as a switch to taking the opposite of the majority action. Herding and contrarian behaviour according to such a definition would, however, occur under similar circumstances as under our definition because, at least loosely, if the majority buys, prices rise, if the majority sells, prices fall so that, for instance, buy herding would arise when prices rise.[7]

3 The Guiding Theory and Testable Predictions

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

Trade Direction. The idea behind informational herding is best explained by example. "FI", a financial institution, has a key competitor that has just declared bankruptcy. This may be good for FI because they may be able to attract the failed competitor's customers. If this situation materializes, a share of FI is worth V_h. The competitor's failure may also be bad since FI may have made the same mistakes as the failed bank; a share of FI is then worth $V_l < V_h$. We are interested in the behaviour of a privately informed investor, who received a (noisy) signal S and who observes several sales. Sales, loosely, convey that the sellers had negative information. We ask if it is possible that a trader sells when he observes many sales, even though his private (noisy) signal alone tells him that FI is worth V_h, $\mathsf{Pr}(V_h|S) > \mathsf{Pr}(V_l|S)$. Suppose that the price $p \in (V_l, V_h)$ is fixed. Then for any signal S with $\mathsf{Pr}(V_h|S) < 1$ there exists a number of sales x large

including for non-financial market environments. See also Park and Sabourian (2011) for an extensive discussion of the definition.

[7] One can imagine further, broader definitions that, for instance, dispense with the initial benchmark and define herding or contrarianism only relative to the actions of recent predecessor(s). For example, one could classify a trader as engaging in herding if the trader takes the same action compared to the most recent predecessor. Yet dropping the benchmark is problematic. Suppose, for instance, that two traders both have favourable information and buy one after the other. Arguably, they buy because of their information, not because everyone else or their immediate predecessor takes the same action.

enough such that, upon observing x sales, $\mathsf{E}[V|S, x \text{ sales}] < p$. In other words, even if the private signal favours state V_h, $\mathsf{Pr}(S|V_h) > \mathsf{Pr}(S|V_l)$, the trader sells because his private information is swamped by the negative information derived from observing early sales.

There are two shortcomings to this argument. First, the price in financial markets is not fixed. Second, traders cannot choose when to trade. The second point will play a significant role in our experimental analysis. On the first point, if, as is a common in financial market models, the price would be such that $p = \mathsf{E}[V|\text{information contained in all past trades}]$, then we would have that for all past trading activity $p \lesseqgtr \mathsf{E}[V|S]$ if and only if $\mathsf{Pr}(S|V_l) \lesseqgtr \mathsf{Pr}(S|V_h)$. In other words, if the price responds to information derived from trading, someone with favourable information would never sell, ruling out herding. Experimental evidence has confirmed this (see Drehmann, Oechssler, and Roider (2005) and Cipriani and Guarino (2005))). Park and Sabourian (2011), however, have found that when there are more than two possible outcomes, herding in the sense of traders trading against their information with the majority is possible. Our experiment is guided by the qualitative ideas of their model.

The idea in Park and Sabourian (2011) can be best explained, once again, by example. Consider the above banking example with a third outcome, one in which FI is unaffected by its competitor's failure, associated with value V_m with $V_l < V_m < V_h$.[8] Assume all outcomes are equally likely. We are interested in the behaviour of an investor, who has a private signal S, after different public announcements. Specifically, consider a good public announcement G that rules out the worst state, $\mathsf{Pr}(V_l|G) = 0$, and a bad public announcement B that rules out the best state, $\mathsf{Pr}(V_h|B) = 0$. Assume that the price of the stock is equal to the expected value of the asset conditional on the public information and that the investor buys (sells) if his expectation exceeds (is less than) the price. Note that the price will be higher after G and lower after B, compared to the ex-ante situation when all outcomes are equally likely.

Both G and B eliminate one state, so that, after each such announcement there are only two states left. In two state models, an investor has a higher (lower) expectation than the market if and only if his private information is more (less) favourable towards the better state than towards the worse state. Thus, in the cases of G and B, $\mathsf{E}[V|G] \lesseqgtr \mathsf{E}[V|S, G]$ is equivalent to $\mathsf{Pr}(S|V_m) \lesseqgtr \mathsf{Pr}(S|V_h)$ and $\mathsf{E}[V|S, B] \lesseqgtr \mathsf{E}[V|B]$ is equivalent to $\mathsf{Pr}(S|V_m) \lesseqgtr \mathsf{Pr}(S|V_l)$. Hence, for example, after good news G, an investor buys (sells) if he thinks, relative to the market, that it is more (less) likely that FI will thrive than be unaffected. It follows from the above that the investor buys after G and sells after B if and only if $\mathsf{Pr}(S|V_h) > \mathsf{Pr}(S|V_m)$ and $\mathsf{Pr}(S|V_l) > \mathsf{Pr}(S|V_m)$. Such an investor, loosely, herds in the sense that he acts like a momentum trader, buying with rising and selling with falling

[8] It is immaterial that the value is between V_l and V_h — any third value does the trick.

prices. The private information (conditional probabilities) that is both necessary and sufficient for such behaviour has thus a U-shape. Conversely, the investor sells after G and buys after B if and only if $\mathsf{Pr}(S|V_h) < \mathsf{Pr}(S|V_m)$ and $\mathsf{Pr}(S|V_l) < \mathsf{Pr}(S|V_m)$. Such an investor, loosely, trades contrary to the general movement of prices. The private information that is both necessary and sufficient to generate such behaviour has thus a hill-shape. Formally, we distinguish four possible shapes of signal likelihood functions (LF):

increasing: $\mathsf{Pr}(S|V_l) < \mathsf{Pr}(S|V_m) < \mathsf{Pr}(S|V_h)$; *decreasing:* $\mathsf{Pr}(S|V_l) < \mathsf{Pr}(S|V_m) < \mathsf{Pr}(S|V_h)$;
U-shaped: $\mathsf{Pr}(S|V_i) > \mathsf{Pr}(S|V_m)$ for $i = l, h$; *Hill-shaped:* $\mathsf{Pr}(S|V_i) < \mathsf{Pr}(S|V_m)$ for $i = l, h$.

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_l or V_h relative to the other extreme state. We thus define the *bias* of a signal S as $\mathsf{Pr}(S|V_h) - \mathsf{Pr}(S|V_l)$. A U-shaped LF with a negative bias, $\mathsf{Pr}(S|V_h) - \mathsf{Pr}(S|V_l) < 0$, will be labeled as an nU-shaped LF and a U-shaped LF with a positive bias, $\mathsf{Pr}(S|V_h) - \mathsf{Pr}(S|V_l) > 0$, will be labeled as a pU-shaped LF. Similarly, we use nHill (pHill) to describe a Hill-shaped LF with a negative (positive) bias. 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 what follows, we will refer to signals with a particular shape of likelihood function or a trader who receives such a signal by the shape only (e.g. a signal S with an increasing LF is referred to as an increasing signal).

Note that in the example, G and B are exogenous public announcements. In general, however, public announcements or, more rather, public information is created endogenously by the history of publicly observable trading with, loosely, G signifying a prevalence of buying, B a prevalence of selling.

Experimental Parameters. There are three states, $V \in \{V_l, V_m, V_h\} = \{75, 100, 125\}$, all are equally likely, $\mathsf{Pr}(V_l) = \mathsf{Pr}(V_m) = \mathsf{Pr}(V_h)$. We have two types of traders: informed traders (our laboratory subjects, who make up 75% of the trading population and who can buy, sell or hold as they wish); and noise traders (controlled in the lab by the computer, accounting for 25% of the trading population who buy or sell with equal probability). While not necessary for the result, we use noise traders since we worried that subjects' ability to count the number of trades and compare them to the number of subjects in the room might distort the results.[9] The experimental implementation of noise traders was as follows: for a given number of possible trading decisions, noise traders were added so that the ratio of noise to informed was roughly 1/3 (e.g. with 15 subjects, we added 5 noise traders). A coin toss for each noise trader determined whether this trader would buy or sell. Their trade time (see below) was then a uniform draw for the possible trade

[9]Noise traders also add an element of realism by simulating the inclusion of traders with exogenous reasons to buy or sell, who might be of particular interest in an endogenous-time setting.

times (in seconds) from $[0, 180]$.

Subjects can observe all previous prices H^t. In addition, each receives one of three signals, S_1, S_2, S_3, which are private and conditionally i.i.d. informative. Subjects further receive information about the signal likelihood function (hereafter: LF). Each treatment had an increasing signal (S_3), a decreasing signal (S_1) and a non-monotonic signal (S_2).

The rational choice for informed traders (assuming indifferent agents buy) is to buy if their expectation of the value exceeds the price and to sell otherwise. To simplify the experimental setting, buys and sales happen at a single price, which is set by a computer. Subjects know that the price will adjust upwards after a buy and downwards after a sell and they can thus infer actions from past prices (an up-tick indicates that there was a buy, a down-tick indicates a sale).

Trade Timing. The herding theory that we allude to above is based the assumption that traders act in an exogenous, predetermined sequence. In reality and in our experiment they can trade whenever they want and we thus add to the theory by studying this important aspect. Generally, there is no tractable model to provide us with firm predictions about traders' timing decisions. We can, nevertheless obtain some theoretical guidance. Smith (2000) provides a model in which a single trader who can make a single trade at one of two points in time (early or late). His results intuitively extend to the case of multiple traders.[10] Smith shows that a trader with a "good" or "bad" news signal will trade early. In the Park and Sabourian (2011) setup, such signals have monotonic LFs. Moreover, Smith also presents an example with a U-shaped signal and shows that within his framework the recipient of such a signal would delay. In addition, we also ran simulations (available upon request) of the trade-timing decision for the parameters used in our experiments that indicate that for recipients of hill- and U-shaped signals there exists a set of parameters for which delay is optimal.

The basic intuition of the timing decision is that traders generally expect prices to move in the direction of the state which they consider most likely to occur. Traders who think that the highest or lowest state is most likely (they have increasing/decreasing LFs) should then act earliest because early in the round prices are (in their opinion) furthest from their favoured state and thus profit opportunities are largest.[11] Hill-shaped signals

[10]In Smith (2000) the trader obtains a piece of information about a public signal that will be released. After the release of this public signal, prices will adjust instantaneously to the fundamental value implied by the signal. To see the equivalence to our setting assume that people choose a trading action (buy, sell or pass) in accordance with the optimal actions prescribed by the "static" exogenous-time theory discussed earlier in this section. Then their actions will affect the price and (noisily) reveal traders' information. Thus the price at the end of the trading day is, loosely, a sufficient statistic for all traders' private information. Moreover, the price is public information. Thus the price at the end of the trading day is the same as the price in Smith (2000) after the release of the public signal. A trader's information in our model can thus be understood as a signal about the information that will be revealed through trading.

[11]Formally, traders with increasing (decreasing) signals expectations of the public expectations are sub

are slightly different: when trading starts, prices will first be close to the hill-shaped type's favoured value. As prices move away from their favoured value, trading against the movement of prices becomes most profitable. So even though their signal is quite informative they may delay trading. At the beginning of a treatment, U-shaped types are least sure about the direction that prices might take and they may thus delay to learn first from the behaviour of others.

Qualitative hypotheses implied by the theory. Combing the theoretical insights, we can develop a number of hypotheses.

Hypothesis 1 *Decreasing types sell, and increasing types buy. If we observe herding then this is most likely to be caused by U-shaped types. If we observe contrarianism, then this is most likely caused by hill-shaped types.*

In analyzing their timing behaviour we will look at the *distribution* of trading times and we are interested in the relative ordering of the distributions of trading times for the different signal types. This yields the following hypothesis.

Hypothesis 2 *U-shaped types will act later than monotonic types. Hill-shaped types will act before U-shaped types. Hill-shaped types act after the first few trades have occurred.*

Multiple Trades. In much of what is to follow we are concerned with a setting where traders can act only once. Since, in reality, people can trade multiple times, and since there is no information-based theory that is able to describe behaviour and the influence of information, we have also explored an experimental setting in which people can trade twice. With two trades possible, we conjecture that types with bad and good news should still sell and and buy respectively, and they should do so rather sooner than later. Moreover, as prices move, traders' information rents are reduced in expectation and thus the greater the number of trades, the more intense is the competition for information rents. A very straightforward assertion is thus that trades should occur earlier when people can trade more often. We will explore the two-trade case in detail in the penultimate section.

Implementation of the market price in the experiment. One conceptual difficulty that arises in the experimental implementation is the manner in which the price is set and updated by the central computer. For lack of a theory we used the reasonable updating rule by which prices adjust assuming that the most recent trade was taken in accordance with the optimal decision under exogenous-timing as implied by Park and Sabourian (2011). For instance, in a setting with a negative U-shaped signal and absent herding, a buy would have been assumed to come from either a noise trader or an informed

(super) martingales and they thus believe that future prices will move against them.

trader with an increasing signal.[12] See section 4.1 for more details. Note, however, that to detect herding and contrarian behaviour, *it is not necessary that the price is set according to any theory*— the definitions speak only of increasing or decreasing prices. We are, in fact, not testing a particular theory in all its details, but, instead, are interested in the qualitative implications with regards to the impact of information.

We could have used other price-setting mechanisms. For instance, we could have followed Cipriani and Guarino (2005) and used subjects as market makers; they did this in one treatment and had two participants set prices. Our particular framework, however, is rather complex and with human traders on both sides (market making and active trading) we would need to worry about how traders form expectations about market makers and vice versa. Notwithstanding this point, Cipriani and Guarino (2005) found behaviour to be robust with respect to variations in the price-setting rules.[13] Another possible variation of the price setting mechanism is to include a transaction cost, such as a bid-ask-spread.[14]

4 Experimental Design

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

4.1 Overview

Each group of traders was made up of 13-25 experimental subjects. At the beginning of the session, subjects were given information about the trading system, the kind of signals that they may receive, and the functioning of the market. Subjects were told that they would not interact directly with each other but rather that the trades were with a central computer. Subjects were explicitly told that the price would be increased by the central computer following a buy decision, and would be decreased following a sell decision and that they would have access to the full price history generated by the central computer (provided by the experimental software in the form of a price-chart).

[12]This rule does, however, become problematic with multiple trades because two trades by the same person will be counted as two separate pieces of information.

[13]We are grateful to an anonymous referee for making this point.

[14]An earlier working paper version of Drehmann, Oechssler, and Roider (2005), dated April 2003, evaluates the impact of transaction costs. They find that transaction costs make "not trading" optimal for some prices (when at such prices the gain from trade is lower than the cost) and so trading breaks down. Since this effect is well-understood we ignored transactions costs and instead focused on the information structure as the key differentiating factor between treatments.

Prior to running sessions, for each of the six rounds, we had a random draw for the true value (and thus in all sessions, subjects received signals according to the same LF). For each round, each subject received a private signal, either S_1, S_2, or S_3, where signals were drawn using the LFs for the underlying treatment given the true value draw (e.g., the practice round was based on a treatment with an nU-shaped LF and true value $V = 100$).[15] Prior to the start of each treatment, subjects were provided with an information sheet detailing the signal likelihood function (i.e. information about all possible signals) and the posteriors that each signal would imply. The information on the sheet was common knowledge to all subjects. The subjects were *not* told anything about the implications of U-shaped, hill-shaped or monotonic LFs.

All of the experimental subjects took part in all of the rounds in each session (the practice round plus all 6 incentivized rounds) and they were all made aware of this fact. The practice round was special in two senses. Firstly, it was not incentivized and this was stressed, so subjects could use it to get used to the software, and as such it is not included in the results. Secondly, they were allowed to ask questions pertaining to understanding, after the practice round. Information for the round was presented in the same format as for the incentivized rounds.

The subjects received financial incentives to ensure that they took the tasks seriously as is standard practice within experimental economics (excluding the unpaid practice round). First, subjects were provided with a show-up fee of £5 (in UK currency, or equivalent in Canadian currency) and this was known to the subjects. Second, it was explained to the subjects that their bonus payment would depend precisely upon the actual value of each share and the price at which they sold or bought the share (both denominated in virtual currency units or VCU) and examples were given to make this clear, for example:

"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."

It was further explained to the subjects that the central computer would maintain a running total of their winnings in VCU after each round and this would be converted to real currency at the end of the session and that this could reach a possible maximum of £25 (in UK currency, or equivalent in Canadian currency). Subjects made an average of £7.70 in total bonus payments (across all 6 rounds) or £12.70 in total including their show-up fee. The subjects were informed that the rounds would last 3 minutes and that

[15]Figures 2-4 in the supplementary appendix describe the signal LFs.

they would receive announcements about the remaining time after 2:30 minutes, and 2:50 minutes. Each round was followed by a cool-off period of about 1-2 minutes and then subjects were given information for the next round. They had about 2-3 minutes to digest this new information. They were told in advance (as early as at the recruitment stage) that each session would last approximately 1 hour.

The existence and proportion of noise traders (roughly 25% of trades, see Section 3) was made known to the subjects in advance. Subjects were also aware that noise traders randomized 50:50 between buying and selling and that they traded at random times.

We considered two classes of treatments: in the first, subjects were allowed to trade once, in the second they could trade twice. The software allowed subjects to trade at most this specific number of times. The sequence of transactions produced a history of actions and prices, H_t with $t \in [0, 180]$, that recorded the timing (in seconds), price, and direction of each transaction. Subjects were shown the history in the form of a continuously updating price chart during each treatment, and they were also given the current price, P_t, where $P_0 = 100$.

Subjects were told that they had three possible actions $a = \{\text{sell, pass, buy}\}$ one (or two) of which they could undertake during the 3 minutes of trading time.[16] For the treatments in which two trades were allowed, subjects were additionally informed that they could trade twice, so they could "buy and buy", "sell and sell", "sell and buy", etc. as and when they wished during the three minute period. They were instructed that pressing the "pass"-button would count as one of the actions that they were allowed. It was stressed to the subjects that their virtual profits per treatment were generated based on the difference between the price at which they traded, P_t, and the true value of the share, V. It was emphasized that the price at the end of the trading round would not be relevant for their payoffs.

The subjects themselves were recruited from the Universities of Toronto, Cambridge and Warwick. No one was allowed to take part twice. We ran 13 sessions in all: 3 at the University of Cambridge (13 subjects each), 6 at the University of Warwick (18, 19, 22, 22, 22, and 25 subjects) and 4 at the University of Toronto (17, 18, 13, and 13 subjects). We collected demographic data only for the Warwick sessions: of the subjects there, around 49% were female, around 73% were studying (or had already taken) degrees in Economics, Finance, Business, Statistics, Management or Mathematics. 53% claimed to have some prior experience with financial markets, and 23% claimed to own or have owned shares at some point.[17]

[16] The "passing" decision is a relict of a setting with exogenous timing. Instead of pressing a hold button, traders could also just not trade. We have not analyzed this decision separately.

[17] Appendix E details the questions asked in the questionnaire. When asked what motivated their decisions (across different rounds) 44% of subjects mentioned a combination of prices and signals, 31%

4.2 Treatments

In our setup, each round of play corresponds to a unique treatment.[18] In Treatment 1,2, and 3, subjects were allowed to trade at most once, in treatments 4,5, and 6 subjects were allowed to trade at most twice. Each treatment included signals with an increasing, a decreasing and a non-monotonic LF, where Treatment 1 and 6 had a signal with an nU-shaped LF, Treatment 2 and 4 a signal with a nHill-shaped LF, Treatment 3 and 5 had a signal with a pU-shaped LF. The underlying parameters and signal LFs are listed in the supplementary appendix together with the instructions given to subjects.

The specific signal structures by *rounds* were as follows:[19]

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

The same order of treatments was maintained across all sessions.[20]

4.3 Behavioral, non-rational predictions for the static decision

To complete the analysis, we considered the possible impact of risk aversion and loss aversion on decision making, and various behavioural alternatives to Bayesian updating. First, we considered a model in which subjects do not update their beliefs as prices change but act solely on the basis of their prior expectation. Second, we considered one setting in which subjects update their beliefs on the basis of changing prices at a slower rate than they should and one setting in which people overweigh their own private information.

only price, 18% only signal and the remaining 7% had other motivations. 38% thought that in general the current price was more important than the signal, 36% thought the signal was more important than the current price and the remaining 26% felt they were of similar value. Roughly 24% claimed to have carried out numerical calculations.

[18]So treatment 1 occurred in round 1, treatment 2 in round 2, etc.

[19]The practice round was a one-trade treatment with a negative U-shaped signal structure.

[20]Since each round was only played once, and since prior to any rounds a practice round was undertaken we believe that learning between rounds would be minimal. There was no econometric evidence of learning between rounds. For instance, there was no discernible trend in behavior. In untabulated regressions we also clustered standard errors by rounds and double-clustered standard errors by rounds and subjects and found that the results were unchanged.

Finally, we developed error correction models in which subjects account for errors made by their peers and react rationally to these errors; these models are in the spirit of level-k beliefs (see Costa-Gomes, Crawford, and Broseta (2001)) and the Quantal Response Equilibrium (see McKelvey and Palfrey (1995) and McKelvey and Palfrey (1998)).

The full details of the alternative specifications and tests can be found in Section A of the supplementary appendix.

5 Signals and Herding or Contrarianism

5.1 Observations from Summary Statistics

The total number of trades made by experimental participants was 1991 spread over all 6 treatments; broken up by trader type we have 623 (S_1), 786 (S_2) and 584 (S_3). For treatments 1 to 3 we had 683 trades (197 S_1, 276 S_2 and 210 S_3), for treatments 4-6 there were 1308 (425 S_1, 510 S_2 and 373 S_3) trades. Overall, we observe that the monotonic types, S_1 and S_3, always sell/buy in 85% of all their trades.

One of the main questions that we want to answer is whether U-shaped signal types switch from selling to buying if prices rise and whether they switch from buying to selling if prices fall, i.e. that they herd in the sense of following the "majority" action. The "static" theory, applicable under exogenous sequencing, of Park and Sabourian (2011) suggests that only U-shaped types would herd. Yet the definition of herding and contrarianism does not rule out that other types herd or act as contrarians. Table 1 gives the raw numbers on trading behaviour, split by signal types.

Finding 1 (Hypothesis 1, summary statistics) *U-shaped types herd in 26% of possible cases, hill, increasing and decreasing types herd in 4%, 10% and 21% of all cases. Contrarianism arises more frequently, most stemming from hill-shaped types (67%). Decreasing, increasing, and U-shaped types act as contrarians in 31%, 19% and 37% of all possible cases.*

The fraction of herding actions is larger than that observed in Drehmann, Oechssler, and Roider (2005) or Cipriani and Guarino (2005); these papers also found persistent evidence of contrarianism, albeit to a smaller extent.

5.2 Regression Analysis of the Trading Direction Decision

We ask the following questions:

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

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

The random assignment of signals to traders allows us to interpret mean differences in signal-specific effects as the average causal effect of the signal. Formally, we estimate the following equations to test whether a type of signal is a significant cause for herding or contrarian behaviour respectively:

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

where the dependent variables herd_i and contra_i are dummies that apply the Herding/Contrarianism Definition in the sense that they are set equal to 1 if the current trade by individual i is a herding and contrarian trade respectively and 0 otherwise, α is a constant, and $\mathsf{u\text{-}shape}_i$ and $\mathsf{hill\text{-}shape}_i$ are signal dummies that are set equal to 1 if the individual who performs the current trade received a U-shaped (for the herding estimation) or hill-shaped (for the contrarian estimation) signal. Parameter fixed_i is an individual fixed effect that controls for specific traders who persistently err.[21] Given the random assignment of signals, we can assume that $\mathsf{E}[\mathsf{u\text{-}shape}_i \cdot \epsilon_i] = 0$ and $\mathsf{E}[\mathsf{hill\text{-}shape}_i \cdot \epsilon_i] = 0$, which are the main identifying assumptions. In our estimations we cluster standard errors at the trader (individual) level to correct for unobserved components at the trader level. We run the regressions restricted to the cases for which herding and contrarianism respectively are possible.

In each scenario we estimated the model by logit without fixed effects and by a linear model with fixed effects. For logit estimations, we report the marginal effects at the mean. Further, we report the results across all six treatments; in the penultimate section, we expand the analysis to check if the number of trades affects the estimates (it does not).

For the herding specification, β represents the impact of the signal on the average individual's choice of whether or not to herd. If it is positive and significant, then a U-shaped signal increases the probability of herding relative to all other signals. The first two columns in Table 2 summarize the result from our regression. Overall, obtaining a U-shaped signal increases the probability of herding by about 12% relative to any other signal and it is significant at all conventional levels. Parameter estimates are large unaffected by fixed effects. Overall the regression confirms the hypothesis that recipients of U-shaped herding-type signals are generally more likely to herd.

For the contrarianism specification, coefficient β represents the impact of the signal on the average individual's choice of whether or not to act as a contrarian. If it is positive and significant then the hill-shaped signal increases the probability of contrarianism relative

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

to all other signals. The first two columns in Table 3 summarize the results from our regression. Receiving the hill-shaped signal increases the chance of acting as a contrarian by 36-40% relative to any other kind of signal. All coefficients are significant at all conventional levels. Overall we confirm that the hill-shaped signal is the significant source of contrarianism relative to all other signals.

Finding 2 (Hypothesis 1, regression analysis) *The U-shaped and hill-shaped signals are the significant sources for herding and contrarianism respectively.*

Our findings here are noteworthy for two reasons. First, (irrational) contrarianism has been observed in other experiments before (e.g. Drehmann, Oechssler, and Roider (2005), Cipriani and Guarino (2005), Alevy, Haigh, and List (2007)). Thus arguably, people exhibit a general tendency to act against the crowd. Here we show that despite this tendency, contrarianism is still most likely caused by recipients of signals that admit contrarianism theoretically: *so observed contrarianism is not necessarily irrational, but may be the result of sensible, information-based considerations.* Second, the marginal effect of a U-shaped signal on the probability of herding is stronger than that in our exogenous timing companion paper, Park and Sgroi (2009). Combined with the fact that U-shaped types do not herd as much as they should theoretically, this implies that *due to the timing of actions* non-U-shaped types herd proportionately less and U-shaped types herd proportionally more than with exogenous timing: *so the importance of signals that are conducive to herding increases with the potential to delay.*

6 Signals and the Timing of Actions

6.1 Absolute Timing

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

We computed the cdfs for a large variety of subsamples, such as treatments 1-3, 4-5, 4-5 (first trades), and so on. The timing pattern for increasing and decreasing signals showed no differences, neither did positive and negative U-shaped signals. In presenting the results, we thus combine increasing and decreasing signals as "monotonic" signals, and we combine positive and negative U-shaped signals as "U-shaped signals". Moreover, we also aggregate trading times for the respective types across all treatments. Figure 1 provides plots of the relevant differences of cdfs.[22] We find the following.

Finding 3 (Hypothesis 2) *U-shaped types trade later than monotonic types.*

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

Finding 4 (Hypothesis 2) *On almost the entire domain (apart from the first few seconds) the hill-shaped type trades systematically earlier than the U-shaped types.*

To consider Hypothesis 2, one can compare the top right and bottom right panels. As can be seen, for the first few seconds, more trades stem from non-hill-shaped types. Yet after these first few trades, the hill-shaped types trade strongly (this follows as their cdf rises strongly relative to the other cdfs).

Finding 5 (Hypothesis 2) *The trades by hill-shaped types are concentrated after the first few transactions have occurred.*

In the supplementary appendix we further examine if pure information theory can explain the timing of decisions. As is common in the literature on information theory we use the entropy of posteriors to measure the informativeness of a signal. We observe, however, that information theory does not seem sufficient to explain the timing of decisions.

6.2 Relative Timing: Clustering

The word "herding" semantically suggests not only that people take the same action, but also that people act at almost the same time. Definitions of herding (such as ours) and models of herding do not capture the timing decision and, since the models typically force actions to be taken in a strict exogenous sequence, thus have no built-in simultaneity. During the experiments, however, we did observe that traders often acted at almost the same time. This behaviour, which has not been identified before in laboratory experiments, is in the spirit of the mass behaviour that one may associate with herding.

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

We categorize this trading at almost the same time as "stimulus-response" driven trading in the sense that one trade triggers others in short succession. We thus define a trade to be *triggering* if it happens 5 or 10 seconds after its predecessor (this time separation avoids spurious proximities of trades) and at least 5 seconds after the first trade in the round. We then define a *cluster* as a situation where at least 2 more trades occur within 3 seconds of the triggering trade.[23] Table 4 provides summary statistics and indicates that there are a sizeable number of clusters. For instance, those with 5-second delays occur, on average twice in each treatment, those with 10 seconds about once. Furthermore, a large fraction of trades is involved in a cluster (between 21% for 5-second delays and 10% for 10-second delays), which is remarkable because more than 25% of trades occur so early that are excluded by design.

There are several questions to ask. First, why do people trade in a cluster? Second, are clustered trades herding or contrarian trades in the sense of trade-direction? Third, do signals plays a role in the decision to cluster?

As for the first question, the simplest explanation for why clusters arise is that some traders play a (delay-) strategy which includes a conditioning of the form "wait until the next trade and then act". If traders play such a strategy then, naturally, one trade may trigger another or several others,[24] and it is important to understand to what extent information affects this type of behaviour.

A more complex explanation as to why it may pay to trade in a cluster is as follows. One feature of our experimental setup is that prices are set assuming that each trade is performed by a noise trader with constant probability. Although it would be difficult to argue that a trade that triggers a cluster is more or less likely to be informed, a trade that follows another in close succession may well be more likely by an informed trader. In this case, the price adjustment following this trade is too small because the price adjustment accounts for the possibility of noise. Consequently, one may argue that it may be profitable to be the *third* person in a cluster and to trade in the same direction as the second.[25] We try to capture this idea in a regression where we control for trades

[23] The timing numbers used in our definition of a cluster were based on giving traders enough time to observe a change in price on their screens, infer the direction of trade required to produce this price change, make a trading decisions in response and then initiate it: 5-10 second seems enough time to do this for the first trade to reply to the triggering trade, while a further 3 seconds seems enough to capture trades in the same cluster, but not so long as to admit trades generated in reply to a new triggering trade. We performed the analysis with several variations of these numbers and found behaviour to be similar.

[24] Argenziano and Schmidt-Dengler (2010) show that in an N-player pre-emption game clusters can arise naturally as part of an equilibrium strategy. Their model is, however, a fixed-investment framework with an exogenous delay benefit (which we do not have) and without moving prices.

[25] On theoretical explanation for clustering is given by Admati and Pfleiderer (1988) who show that clustering can be rational as informed traders trade more aggressively when they believe that there are more uninformed traders around. Their idea is, however, not directly applicable to our model, because we do not have the simultaneous order submission that is the basis of Admati and Pfleiderer (1988).

that are in the same direction as their predecessor.

To understand all the above questions, we ran the following regression:

$$\begin{aligned}\mathsf{cluster}_i &= \beta_0 + \beta_1\mathsf{herd}_i + \beta_2\mathsf{contra}_i + \beta_3\mathsf{u\text{-}shape}_i + \beta_4\mathsf{increasing}_i \\ &\quad +\beta_5\mathsf{decreasing}_i + \beta_5\mathsf{round\ trip}_i + \beta_6\mathsf{same\ as\ before}_i + \mathsf{fixed}_i + \epsilon_i, \qquad (2)\end{aligned}$$

where $\mathsf{herd}_i, \mathsf{contra}_i, \mathsf{u\text{-}shape}_i, \mathsf{increasing}_i$, and $\mathsf{decreasing}_i$ are the usual herding, contrarianism and signal dummies, $\mathsf{cluster}_i$ is a dummy that is 1 when the current trade is in a cluster and 0 otherwise, $\mathsf{round\ trip}_i$ is a dummy that is 1 if the trader who made this trade makes his other trade in the opposite direction (buy-sell or sell-buy), and $\mathsf{same\ as\ before}_i$ is a dummy that is 1 if the current trade is in the same direction as the trade just before it and zero otherwise. These covariates are, in essence, all effects that could play a role in our analysis. We ran these regressions, as before, for a linear probability model with and as a logit model without fixed effects. Moreover, we classified trade clusters in two ways: the first included the triggering trade as part of the cluster, the second omits the triggering trade. Overall, we find the following:

Finding 6 *With the exception of the U-shaped signal for 5-second quiet periods before trades, none of the covariates in (2) is persistently statistically significant.*

This finding is, of course, a negative result: although one can argue that clustering may be caused by a delay strategy ("act after the next trade"), one may have suspected that, for instance, signals should have played a role, i.e. that some types of signals are more likely to delay and thus act in clusters than others. Yet we found essentially no persistence or explanation for traders' behaviour, except that the occurrence of clusters themselves and we are thus left with the stylized fact that subjects tend to trade in unison, a finding which represents an important area for future research.

7 The Second Trade

7.1 Herding and Contrarian Estimates with One vs. Two Trades

In half of our treatments, subjects have the option to trade twice. One natural question is whether this option affects the impact that signal have on the chance of engaging in herd behaviour. To answer this question, we ran the following regression:

$$\begin{aligned}\mathsf{herd}_i &= \alpha + \beta_1\mathsf{u\text{-}shape}_i \times \mathsf{1\text{-}trade}_i + \beta_2\mathsf{u\text{-}shape}_i \times \mathsf{2\text{-}trade}_i + \beta_3\mathsf{1\text{-}trade}_i + \epsilon_i, \qquad (3)\\ \mathsf{contra}_i &= \alpha + \beta_1\mathsf{hill\text{-}shape}_i \times \mathsf{1\text{-}trade}_i + \beta_2\mathsf{hill\text{-}shape}_i \times \mathsf{2\text{-}trade}_i + \beta_3\mathsf{1\text{-}trade}_i + \epsilon_i, \ (4)\end{aligned}$$

where $\mathsf{herd}_i, \mathsf{contra}_i, \mathsf{u\text{-}shape}_i$, and hill_i are the herding, contrarianism, U- shape and hill-shape indicators from (1), $\mathsf{1\text{-}trade}_i$ and $\mathsf{2\text{-}trade}_i$ are 1 if the trade was made in a one- and two-trade treatments respectively. Parameters β_1 and β_2 then reveal the marginal effect of a U- and hill-shaped signal respectively in the one- and two-trade treatments. The third and fourth columns in Tables 2 and 3 display the estimates. At the bottom of the table we present the results of an F-Test for equality of the coefficient estimates $\hat{\beta}_1$ and $\hat{\beta}_2$. Columns five and six perform the same analysis, where we further differentiate between the first and second trade in the two trade treatments.

Finding 7 *The coefficient estimates for the impact of signals on the probability of herding and contrarianism are robust to the number of trades in that we cannot reject the hypothesis that they coincide. Signals do not, however, have the same impact on the second trade being herding or contrarian.*

7.2 The Impact of Round-Trip Trades

The fact that signals have a reduced effect on the second trade is noteworthy. One possibility is that subjects followed an altogether different strategy when making their second trade. Namely, with two trades, traders have the opportunity to make so-called "round-trip" or "return" trades by selling first and then buying later or vice versa. This way, they can realize a trading profit in the process.

Table 6 provides summary statistics for the second trade in general, and shows that about 23% of second trades are part of a round trip transaction. About 76% of the return trades yielded a trading profit which suggests that return-trades were performed on the basis of "buy low, sell high" (or "sell high, buy low"). Furthermore, most of round-trip trades are performed by the hill- and U-shaped types.

All this indicates, that traders may well have a particular, possibly non-information-based strategy, in their trading and that this may affect our estimate of herding and contrarianism. A trader who merely aims to buy low and sell high may thus act for reasons that have little to do with his information. Yet in our analysis thus far, this trader's actions may be classified as herding or contrarian and we would thus obtain spurious estimates.

We thus analyze to what extent our estimates in Tables 2 and 3 change when we take account of this possible misclassification. We ask the following question: what is the probability that a first/second trade is a herding trade conditional on the trade being a return trade (when herding is possible) relative to the case where it is not a return trade?

To answer this question, we ran the following regressions

$$\begin{aligned} \mathsf{herd}_i &= \alpha + \beta_1 \mathsf{u\text{-}shape}_i + \beta_2 \mathsf{return\ trade}_i + \beta_3 \mathsf{u\text{-}shape}_i \times \mathsf{return\ trade}_i + \epsilon_i, && (5)\\ \mathsf{contra}_i &= \alpha + \beta_1 \mathsf{hill\text{-}shape}_i + \beta_2 \mathsf{return\ trade}_i + \beta_3 \mathsf{hill\text{-}shape}_i \times \mathsf{return\ trade}_i + \epsilon_i. && (6) \end{aligned}$$

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

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

Table 7 summarizes our findings and indicates that our herding estimates from Section 5 are biased *downwards* by round-trip trades (the coefficients on the product term are negative and significant) and that our contrarian estimates are unaffected. This is good news for our analysis as it indicates that, if anything, the effect of a herding signal as a source for herd behaviour is underestimated by the possibility of round trip transactions.

Finding 8 (Impact of Return Trades on Estimates for Hypothesis 1) *The estimates underlying Finding 2 for herding become stronger and those for contrarianism remain unaffected when we correct for round trip transactions.*

7.3 The Timing of Actions with One vs. Two Trades

Our final question concerns the timing of trades of one-trade relative to two-trade treatments. Since informed traders compete to exploit their private information, more trades imply higher competition for information rents which, under our price-setting regime, should speed up trading. The panels in Figure 2 plot the differences of cdfs of timing, where we aggregated all trades in treatments 1-3 and 4-6 as well as first and second trades in treatments 4-6.

Finding 9 *Allowing people to trade twice accelerates their trade-times: (1) The first trade in treatments 4-6 occurs earlier than the single transaction in treatments 1-3. (2) The single trade in treatments 1-3 occurs earlier than the second trade in treatments 4-6. (3) All trades together in treatment 4-6 occur earlier than in treatments 1-3.*

To assess this finding, suppose subjects' timing strategies for their trade time T in the single trade treatment could be described by some density f on $[0, 180]$ and consider the following two timing strategies as benchmarks.[26] In the first, traders choose the times for their two trades τ_i, τ_j according to some joint density $f(\tau_i) \cdot f(\tau_j)$ over $[0, 180]$. In the second, traders choose the time t_1 of their first trade according to $f(t_1)$ and then choose the time t_2 for their second trade on $[t_1, 180]$ according to $f(t_2|t_2 \geq t_1)$. Applied to our trading setup, the first specification loosely implies that the subjects apply their single-trade timing strategy as independent draws to the two trades; the second specification implies that traders apply the same strategy for the single and first trade and then apply the same strategy of their first trade to their second trade, conditional on the execution of the first trade. Intuitively, the first specification would then imply that the distribution of trade times is such that the first trade for the two-trade specification occurs before the single trade, but that the distributional order for all trades is unclear.[27] The second specification would imply that the distribution of trade times is such that trades for the two-trade specification occur before the single trade, but that there should be no order when comparing the first trade for the two-trade specification with the single trade.

Neither of these benchmarks implies that the first trade of the two-trade specification and all trades taken together from the two-trade specification occur earlier than the single trade from the one-trade specification. Our finding thus indicates that there is an accelerating effect when traders can trade more often that is distinct from the pattern that would emerge from the two benchmarks that we discuss above.

8 Conclusion

Herding has long been suspected to play a role in financial market booms and busts. Recent theoretical work shows that informational herding is possible if the signal likelihood function for traders has a specific shape. Other work shows that when timing is endogenous to the decision, traders with good or bad news should trade earlier than those with less informative signals. Giving traders a choice of when to act is not only natural, but

[26]The idea is that players play a symmetric mixed strategy with full support on the available time interval; implementing this strategy, the probability that a trader has played up to time t can be described by a distribution, and, for simplicity, we assume here that it has density f.

[27]We are grateful to an Associate Editor for making this point. In support of this argument we ran the following simulation which assumes that traders play a uniform timing strategy. We first generated 1 million uniformly distributed trades on the [0,180] interval. These observations are used as the single trade times. We then generated another 2 million observations which are interpreted as trade-times for the two-trade settings. We randomly form 1 million pairs with the smaller element being the first trade, and the larger the second trade. For these trades, we carried out the same distribution computations as for our sample and observed the described pattern. The simulations invoke the Mersenne-Twister method (designed to generate a high level of pseudo-randomness and to avoid serial correlation).

there are also important insights that can be gleaned from such an analysis.

It is not clear ex ante, how the decision to time one's trades should affect herding and contrarianism. One possibility is that when herding-prone types delay their actions systematically, herd behaviour can become more pronounced and significant compared to exogenous timing settings. On the other hand, research by Drehmann, Oechssler, and Roider (2005) and Cipriani and Guarino (2005) has revealed that people have a general tendency to act as contrarians. Another possibility thus is that by removing the artificial friction of exogenous timing, herding disappears. Our work directly addresses this open question.

Having collected almost 2000 trades, we found that subjects' decisions were generally in line with the qualitative predictions of the information theory learning theory when that theory admits rational herding and contrarianism. For example, types theoretically prone to herd or be contrarian are the significant and important source of this kind of behaviour when it does arise. Furthermore, types with extreme information about an asset (both good or bad) trade systematically earliest, and those with signals conducive to contrarianism trade earlier than those with information conducive to herding. We thus find strong evidence for the impact of the type of information both with respect to the direction and the timing of trades.

We can break our findings down further into four key messages. First, we find additional and qualitatively novel support for information-based motives for herding theory in the laboratory. Second, adding endogenous-timing leaves the key predictions of sequential herding theory unchanged as far as the direction of trade is concerned. Therefore, our results suggest that earlier work which forces subjects to act in a strict sequence remains valid even though the timing assumptions impose an artificial friction. Third, we combine two literatures by linking information-based trade directions and timing and show that signals that push subjects towards herd or contrarian behaviour also push them towards delay, relative to the signals that guide subjects towards clear buy or sell decisions. This point is a potentially important avenue for future research as the combination of herding/contrarianism in decision-making and clustering in time can work together to potentially exacerbate/counter prices movements which drift away from fundamentals. Finally, we also identify a new experimental stylized fact in that traders tend to cluster their actions in time. This final key finding represents a potentially important avenue for future research.

References

Admati, A., and P. Pfleiderer (1988): "A Theory of Intraday Patterns: Volume and Price Variability," *Review of Financial Studies*, 1, 3–40.

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.

Argenziano, R., and P. Schmidt-Dengler (2010): "Clustering in N-Player Preemption Games," Working paper, University of Essex.

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

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

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

Bloomfield, R., M. O'Hara, and G. Saar (2005): "The 'make or take' decision in an electronic market: Evidence on the evolution of liquidity," *Journal of Financial Economics*, 75(1), 165–199.

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

Chamley, C. (2004): *Rational Herds*. Cambridge University Press, Cambridge, United Kingdom.

Chamley, C., and D. Gale (1994): "Information Revelation and Strategic Delay in a Model of Investment," *Econometrica*, 62, 1065–1085.

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.

Costa-Gomes, M., V. P. 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.

Gale, D. (1996): "What have we learned from social learning?," *European Economic Review*, 40, 617–628.

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.

Levin, D., and J. Peck (2008): "Investment dynamics with common and private values," *Journal of Economic Theory*, 143(1), 114 – 139.

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

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

Park, A., and H. Sabourian (2011): "Herding and Contrarian Behavior in Financial Markets," *Econometrica*, 79(4), 973–1026.

Park, A., and D. Sgroi (2009): "Herding and Contrarian Behaviour in Financial Markets: An Experimental Analysis," Working paper, Universities of Toronto and Warwick.

Sgroi, D. (2003): "The Right Choice at the Right Time: A Herding Experiment in Endogenous Time," *Experimental Economics*, 6(2), 159–180.

Smith, L. (2000): "Private Information and Trade Timing," *American Economic Review*, 90, 1012–1018.

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

Welch, I. (1992): "Sequential Sales, Learning, and Cascades," *The Journal of Finance*, 47(2), 695–732.

Ziegelmeyer, A., K. B. My, J.-C. Vergnaud, and M. Willinger (2005): "Strategic Delay and Rational Imitation in the Laboratory," Working paper, Max Planck Insitute of Economics.

Table 1
Herding and contrarian trades for all traders by treatment.

The first row in each treatment grouping lists how many herding trades were observed, the second row entries list the number of possible herding trades. An S_1 type cannot herd-sell and can herd-buy only if the price has risen. An S_3 type cannot buy-herd and can sell-herd only if the price has fallen. Similarly, an S_1 type cannot be a sell-contrarian and acts as a buy-contrarian only when buying after prices have fallen; conversely for the S_3 types. The description for the herding and contrarian actions for the S_2 types are more involved, but they are described in detail in Section 3.

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

Table 2
The Effect of U-Shaped Signals on the Probability of Herding.

The table represents regressions of the occurrence of a herding trade on the trader receiving a U-shaped signal as expressed in equation (1). Logit regressions report the marginal effects. Linear probability fixed effects regressions control for trader-fixed effects. The data is restricted to include only trades that could potentially be considered as herding trades. At the bottom of the table we include results for F-tests for equality of coefficients when testing whether the impact of a U-shaped signal is different for one vs. two trade treatments. For most cases, we cannot reject the hypothesis that they coincide. For all tables that follow, standard errors are in parentheses and are clustered at the trader (individual) level, * indicates significance at the 10% level, ** at the 5% level, and *** at the 1% level.

	Dependent Variable: Herding trade indicator					
Treatments	All		Single Trade vs Two Trades		Single Trade vs. First Trade Vs. Second trade	
Fixed Effects	No	Yes	No	Yes	No	Yes
U-shaped signal	0.12039*** (0.024)	0.12754*** (0.033)				
U-shaped × single trade			0.12600*** (0.034)	0.11958*** (0.044)	0.12577*** (0.034)	0.11960*** (0.044)
U-shape × second trade			0.11756*** (0.029)	0.13250*** (0.042)		
U-shaped × first trade					0.14584*** (0.032)	0.18329*** (0.051)
U-shaped × second trade					0.08463** (0.036)	0.0805 (0.050)
Control			-0.02346 (0.026)	-0.00519 (0.024)	-0.02341 (0.026)	-0.00516 (0.024)
Constant	-0.25017*** (0.008)	0.11899*** (0.008)	-0.24230*** (0.012)	0.12052*** (0.013)	-0.24186*** (0.012)	0.12052*** (0.013)
Observations	1142	1142	1142	1142	1142	1142
R-squared		0.40854		0.40869		0.41215
			F-Tests for equality of coefficients			
Single=Two trades			Yes	Yes		
Single=First					Yes	Yes
Single=Second					Yes	Yes
First=Second					No*	No*

Table 3
The Effect of Hill-Shaped Signals on the Probability of Acting as a Contrarian.

The table represents regressions of the occurrence of a contrarian trade on the trader receiving a hill-shaped signal, as expressed in equation (1). Logit regressions report the marginal effects. Fixed effects regressions control for trader-fixed effects. The data is restricted to include only trades that could potentially be considered as contrarian trades. Standard errors and significance levels are denoted as in Table 2.

	Dependent Variable: Contrarian trade indicator					
Treatments	All		Single Trade vs Two Trades		Single Trade vs. First Trade Vs. Second trade	
Fixed Effects	No	Yes	No	Yes	No	Yes
Hill-shaped signal	0.35781*** (0.059)	0.40184*** (0.095)				
Hill-shaped × single trade			0.26423*** (0.079)	0.29987** (0.120)	0.26430*** (0.079)	0.30245** (0.120)
Hill-shaped × second trade			0.44029*** (0.083)	0.50308*** (0.144)		
Hill-shaped × first trade					0.49889*** (0.135)	0.56684*** (0.183)
Hill-shaped × second trade					0.39504*** (0.112)	0.44264** (0.196)
Control			0.05286 (0.035)	0.06607 (0.047)	0.05288 (0.035)	0.06565 (0.047)
Constant	-0.21680*** (0.013)	0.25692*** (0.008)	-0.23468*** (0.017)	0.23436*** (0.018)	-0.23474*** (0.017)	0.23450*** (0.018)
Observations	743	743	743	743	743	743
R-squared		0.43545		0.43983		0.44039
			F-Tests for equality of coefficients			
Single=Two trades			No*	Yes		
Single=First					Yes	Yes
Single=Second					Yes	Yes
First=Second					Yes	Yes

Table 4
Summary Statistics on Clusters.

The table lists the number and proportion of times that there is one trade preceded by 5,10, and 15 seconds of no trades and that is then followed by two additional trades within the next 2 seconds

no trade before trigger	5 seconds	10 seconds	15 seconds
Cluster triggered by informed	49	17	9
Cluster triggered by noise	74	36	8
Informed trades involved	1043	673	520
in cluster (including at time 0)	(52%)	(34%)	(26%)

Table 5
The Probability of Being in a Cluster

The table shows the results from a logit regression of the probability of being in a cluster on a variety of explanatory variables. Standard errors are given in parentheses below the coefficients; these are clustered at the group level.

No trading before the cluster	5 seconds	10 seconds	5 seconds	10 seconds	5 seconds	10 seconds	5 seconds	10 seconds
Triggering trade included?	Yes	Yes	No	No	Yes	Yes	No	No
Fixed effects included?	No	No	No	No	Yes	Yes	Yes	Yes
Herd	-0.02559	-0.03687	-0.01303	-0.00341	0.00542	-0.0475	-0.00495	-0.01868
	(0.053)	(0.041)	(0.045)	(0.033)	(0.075)	(0.051)	(0.062)	(0.041)
Contrarian	-0.02089	-0.01409	0.00656	-0.00484	0.045	0.01533	0.04853	0.00636
	(0.051)	(0.041)	(0.043)	(0.036)	(0.064)	(0.049)	(0.052)	(0.036)
U-shaped signal	0.17797***	0.02459	0.11308*	0.01415	0.15349**	0.00567	0.08459	-0.01217
	(0.066)	(0.043)	(0.060)	(0.037)	(0.075)	(0.055)	(0.067)	(0.047)
decreasing signal	0.15648**	0.02938	0.08132	0.00855	0.12771	-0.02346	0.04732	-0.0369
	(0.065)	(0.044)	(0.060)	(0.037)	(0.078)	(0.055)	(0.070)	(0.046)
Increasing signal	0.08914	0.00974	0.03323	-0.00887	0.10883	-0.00678	0.07149	-0.01268
	(0.073)	(0.049)	(0.066)	(0.042)	(0.084)	(0.057)	(0.071)	(0.042)
Same trade as preceding trade	-0.01329	-0.00784	0.00535	0.0028	-0.02781	-0.00499	-0.0016	0.00983
	(0.030)	(0.022)	(0.024)	(0.017)	(0.036)	(0.027)	(0.027)	(0.019)
Round-trip trade	0.02146	-0.01007	-0.01806	-0.03028	0.03941	-0.03717	-0.01708	-0.06048*
	(0.038)	(0.030)	(0.035)	(0.026)	(0.065)	(0.044)	(0.054)	(0.036)
Constant	-0.32155***	-0.21958***	-0.28725***	-0.17001***	0.14636**	0.14729***	0.11436*	0.11236***
	(0.060)	(0.039)	(0.054)	(0.035)	(0.069)	(0.048)	(0.059)	(0.039)
Observations	996	996	996	996	996	996	996	996

Table 6
Return Trades.

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

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

Table 7
Impact of Return Trades on Herding and Contrarianism.

The table condenses six regressions of the equations in lines (5) and (6) (by signal type and then with respect to herding and contrarian behaviour separately). When cells are empty, there was insufficient data or the variable was dropped. Constants were omitted from the report. Standard errors and significance levels are denoted as in Table 2.

	Herding	Contrarian
U-shaped signal	0.14067*** (0.03979)	
return × U-shaped signal	-0.10634** (0.04417)	
Hill-shaped signal		0.39800*** (0.11415)
return × Hill-shaped signal		-0.1513 (0.15627)
Return	0.22477*** 0.02003	0.37721*** 0.02619
Constant	-0.28673*** (0.01383)	-0.33828*** (0.01315)
Observations	764	489

Figure 1

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

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

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

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

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

Andreas Park
University of Toronto

Daniel Sgroi
University of Warwick

December 21, 2011

Abstract

Part A of this document examines the performance of alternative behavioral models discussed in Section 4.3 of the main paper. Part B discusses whether information theory may help in understanding the timing decisions. Parts C through G detail the experimental time-line, instructions, information for participants about signals, the software and the questionnaire. The information sheets on signals outline the signal likelihood functions. References to sections are with regards to those in the main paper. This appendix is not intended to be published with the paper, but rather provides additional information for the benefit of an interested audience.

A Alternative Explanations for Trading Behavior

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

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

A.1 Risk and Loss Aversion

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

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

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

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

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

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

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

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

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

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

A.2 Decision Rule: Prior Actions or No Updating

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

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

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

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

[1] Arguably, we are only using one part of the tools developed in prospect theory, S-shaped valuations, and ignore that other component, decision weights. However, the latter relate to re-scaled probabilities which we analyze separately so as to be able to distinguish the effects of the two components.

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

A.3 Probability Scaling and Shifting

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

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

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

Consequently, when modeling an overconfident trader who puts more weight on his prior signal we would apply an $\alpha > 1$ re-scaling on the initial probabilities. Alternatively, one can also model slow updating directly by applying an $\alpha < 1$ re-scaling to the posterior probabilities. Of course the effect will be similar: in both cases the histories or updated probabilities would be less important to traders than under the theory from the main text. We considered both specifications.

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

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

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

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

A.4 Error Correction Provisions

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

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

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

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

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

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

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

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

A.5 Summary of Alternative Behavioral Explanations

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

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

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

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

(subjective) expectation for which they would herd. However, with slow updating, they will also be slower to reduce prices and thus it is conceivable that they herd when prices move "against" the herd.

B Information Theory and Timing

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

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

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

Signal Distribution

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

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

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

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

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

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

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

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

C Time-line

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

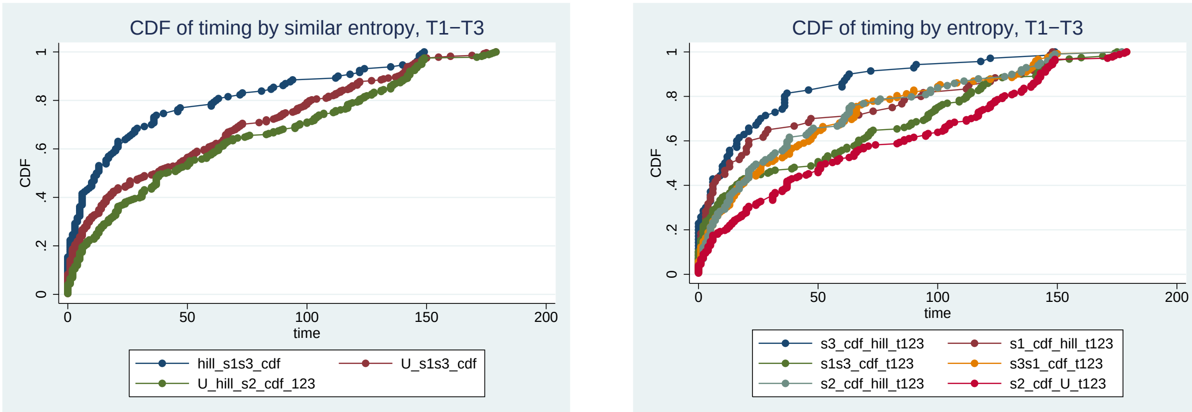

Figure 1
Cumulative distributions ordered by signal entropy.

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

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

10. Pause to answer final questions.

11. Run treatments 2-7.

12. Read instructions 8, which ends the experiment.

13. Calculate and distribute payments while participants complete receipts and questionnaires.

D Instructions

Note that the parts of the instructions in bold indicate that a name, number or currency be included in the instructions which vary by session. Words in italics are emphasized. The instructions are long, and took an average of around 25 minutes to deliver including typical questions. Payment calculations typically took around 5 minutes during which subjects were asked to shut down open software and complete a questionnaire. Note that in the instructions the example round is called "round 1" with the true experiment encompassing rounds 2-7. In the main text of the paper we instead call the rounds "treatments", and ignoring the example round, renumber them to be treatments 1-6.

D.1 Instructions 1 (Welcome)

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

D.2 Instructions 2 (After Seated)

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

D.3 Instructions 3 (The Experimental Setting)

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

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

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

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

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

D.4 Instructions 4 (The Software)

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

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

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

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

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

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

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

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

D.5 Instructions 5 (Compensation)

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

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

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

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

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

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

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

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

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

D.6 Instructions 6 (The Information Setting)

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

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

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

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

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

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

D.7 Instructions 7 (Summary)

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

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

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

D.8 Instructions 8 (Experiment End)

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

E Information Sheets

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

F Questionnaire

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

About you

1. Your age:

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

About your decisions today

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

About the experiment

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

G The Software

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

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

The client software provided a simple to use graphical interface which enabled subjects to observe private information (their signal), and public information (the movement of prices and the current price), as well as indicating to them when they could trade (flashing red and enabling trading buttons) and providing the means of trade (buy, sell and pass buttons). Figure 6 below shows a screen shot of the software in action.

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

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

References

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

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

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

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

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

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

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

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

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

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

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

Table I.
Risk-Aversion Analysis.

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

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

Table II.
Loss-Aversion Analysis.

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

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

Table III.
No Updating and Prior Actions.

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

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

Table IV.
Overweighting of the Prior.

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

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

Table V.
Variations in the Perception of Noise Trading.

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

Figure 2

Information Sheet for a Positively U-Shaped Signal Likelihood Function

Round

Signals: **Case 2**

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

If the true effect will be POSITIVE then you receive

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

If the true effect will be NEGATIVE then you receive

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

If indeed the effect will be NO EFFECT then you receive

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

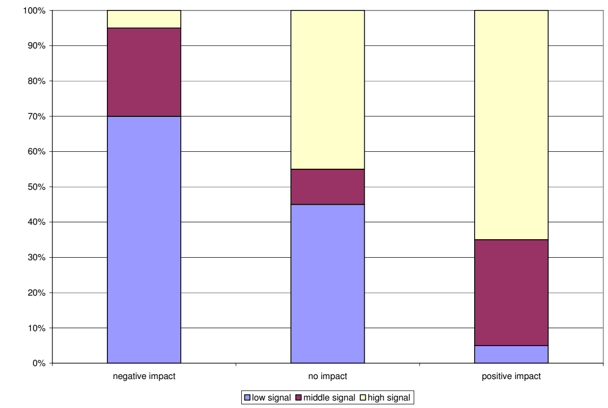

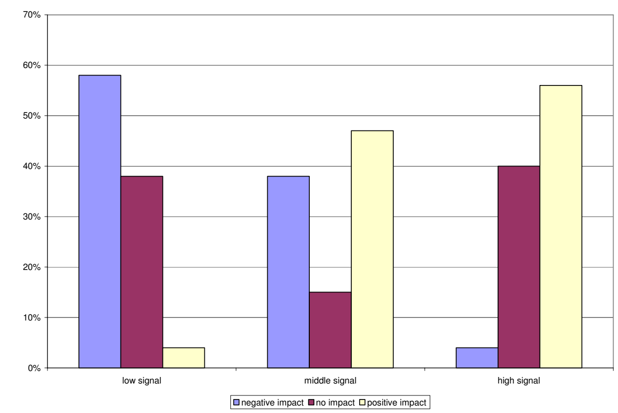

Figure 3
Information Sheet for Negatively U-Shaped Signal Likelihood Function

Round

Signals: **Case 2**

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

If the true effect will be POSITIVE then you receive

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

If the true effect will be NEGATIVE then you receive

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

If indeed the effect will be NO EFFECT then you receive

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

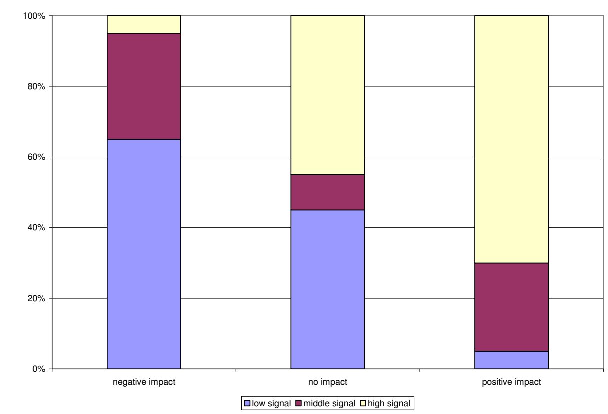

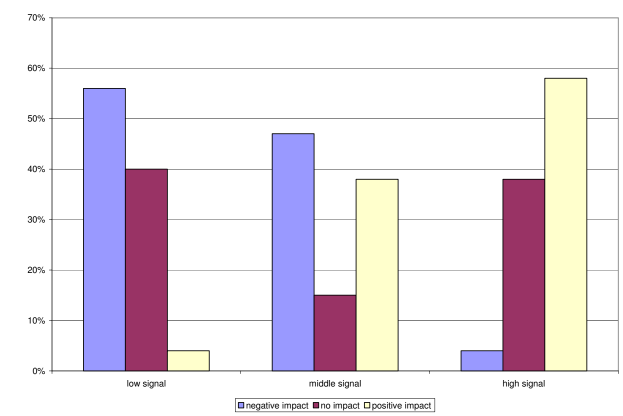

Figure 4
Information Sheet for Negatively Hill-shaped Signal Likelihood Function

Round

Signals: **Case 1**

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

If the true effect will be POSITIVE then you receive

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

If the true effect will be NEGATIVE then you receive

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

If indeed the effect will be NO EFFECT then you receive

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

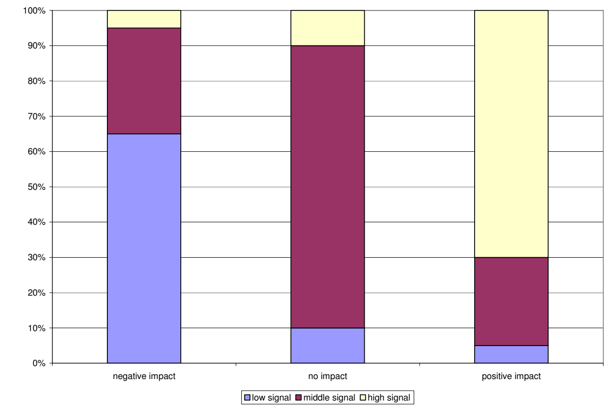

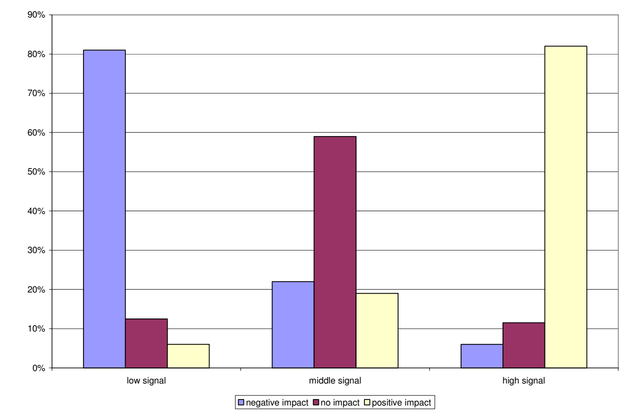

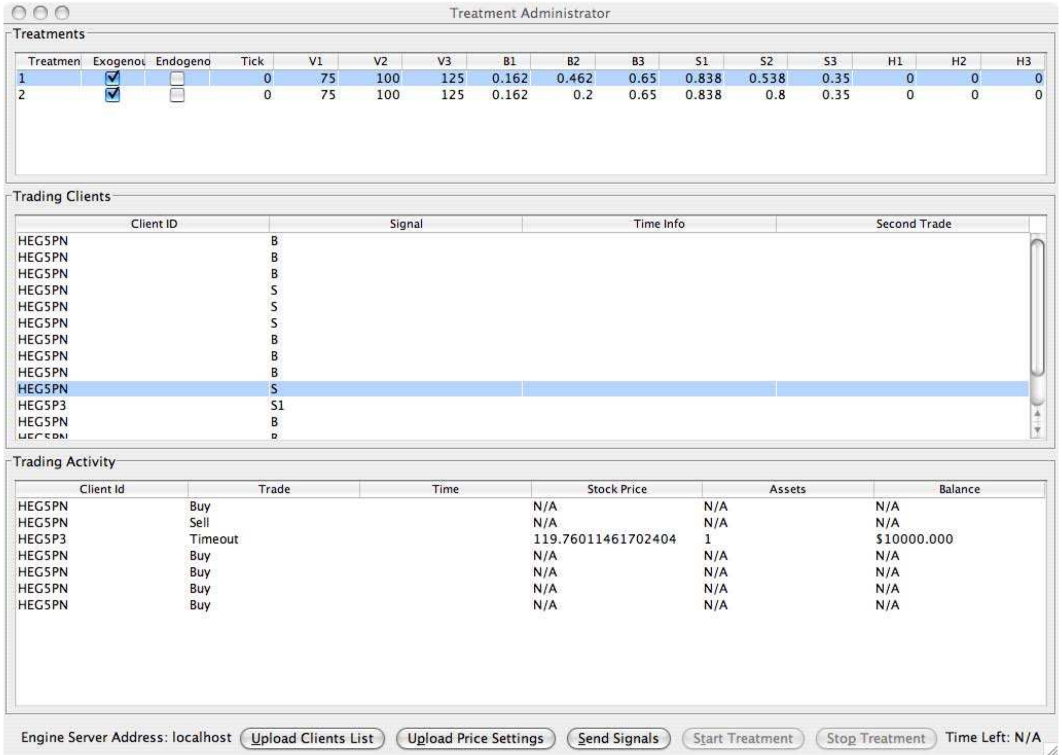

Figure 5
The Administrative Interface

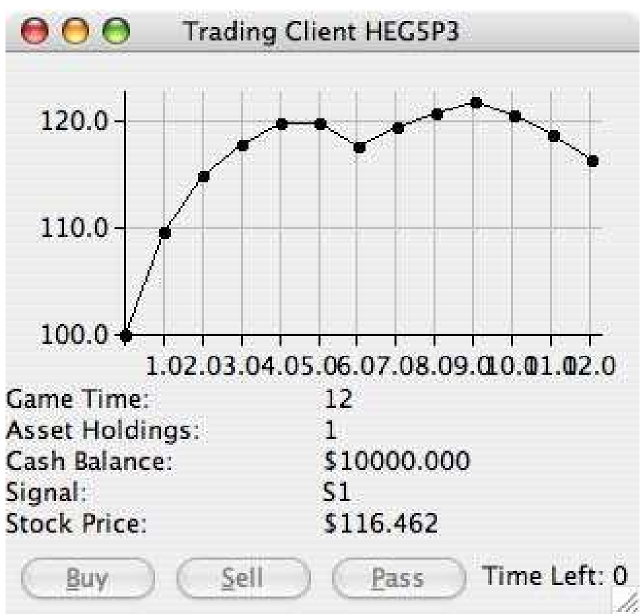

Figure 6
The Trading Client