– Not for Publication –

Supplementary Appendix for

Herding, Contrarianism and Delay in Financial Market Trading

Andreas Park
University of Toronto

Daniel Sgroi
University of Warwick

June 18, 2009

Abstract

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

A Alternative Explanations for Trading Behavior

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

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

A.1 Risk and Loss Aversion

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

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

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

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

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

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

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

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

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

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

A.2 Decision Rule: Prior Actions or No Updating

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

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

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

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

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

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

A.3 Probability Scaling and Shifting

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

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

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

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

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

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

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

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

A.4 Error Correction Provisions

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

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

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

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

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

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

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

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

A.5 Summary of Alternative Behavioral Explanations

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

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

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

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

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

B Information Theory and Timing

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

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

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

Signal Distribution

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

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

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

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

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

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

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

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

C Time-line

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

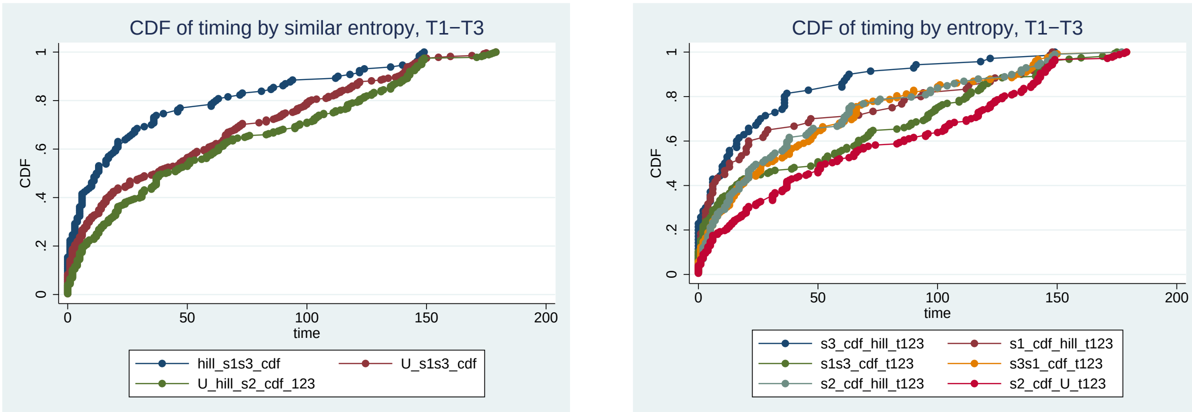

Figure 1
Cumulative distributions ordered by signal entropy.

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

1. The room is prepared and software pre-loaded into the machines to be used, which are allocated each to one ID number.

2. Read instructions 1 including random distribution of ID cards and seat subjects on the basis of the allocated ID cards.

3. Read instructions 2 including the completion and collection of permission forms.

4. Read instructions 3 which explains the experimental setting.

5. Read instructions 4 which explains the software.

6. Read instructions 5 which explains the compensation.

7. Read instructions 6 which explains the information setting.

8. Read instructions 7 which summarizes the instructions and pause to answer any questions.

9. Run treatment 1 (the example round).

10. Pause to answer final questions.

11. Run treatments 2-7.

12. Read instructions 8, which ends the experiment.

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

D Instructions

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

D.1 Instructions 1 (Welcome)

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

D.2 Instructions 2 (After Seated)

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

D.3 Instructions 3 (The Experimental Setting)

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

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

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

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

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

D.4 Instructions 4 (The Software)

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

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

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

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

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

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

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

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

D.5 Instructions 5 (Compensation)

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

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

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

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

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

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

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

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

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

D.6 Instructions 6 (The Information Setting)

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

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

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

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

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

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

D.7 Instructions 7 (Summary)

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

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

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

D.8 Instructions 8 (Experiment End)

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

E Information Sheets

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

F Questionnaire

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

About you

1. Your age:

2. Your gender:

3. Your degree subject:

4. Have you ever owned shares?

5. Do you have any experience of financial markets? (if so, what are your experiences)

About your decisions today

6. What made you decide to buy, sell or pass?

7. How important was the current price?

8. How important was the past price data (the graph)?

9. How important was your "broker's tip"?

10. What else mattered?

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

About the experiment

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

G The Software

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

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

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

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

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

References

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

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

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

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

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

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

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

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

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

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

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

Table I.
Risk-Aversion Analysis.

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

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

Table II.
Loss-Aversion Analysis.

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

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

Table III.
No Updating and Prior Actions.

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

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

Table IV.
Overweighting of the Prior.

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

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

Table V.
Variations in the Perception of Noise Trading.

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

Figure 2
Information Sheet for positive U Shape

Round

Signals: **Case 2**

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

If the true effect will be POSITIVE then you receive

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

If the true effect will be NEGATIVE then you receive

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

If indeed the effect will be NO EFFECT then you receive

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

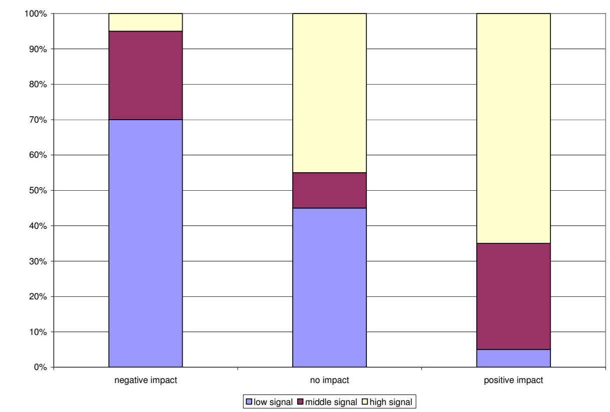

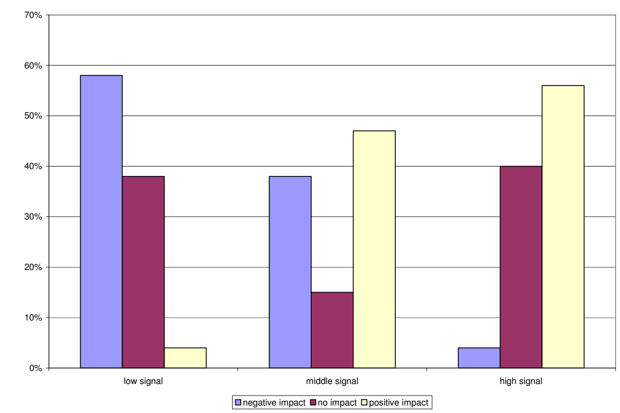

Figure 3
Information Sheet for negative U Shape

Round

Signals: **Case 2**

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

If the true effect will be POSITIVE then you receive

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

If the true effect will be NEGATIVE then you receive

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

If indeed the effect will be NO EFFECT then you receive

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

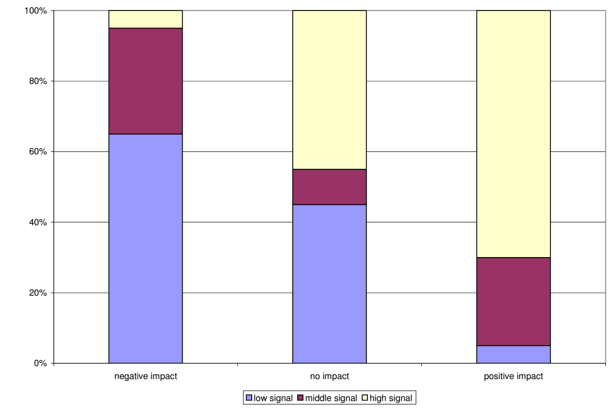

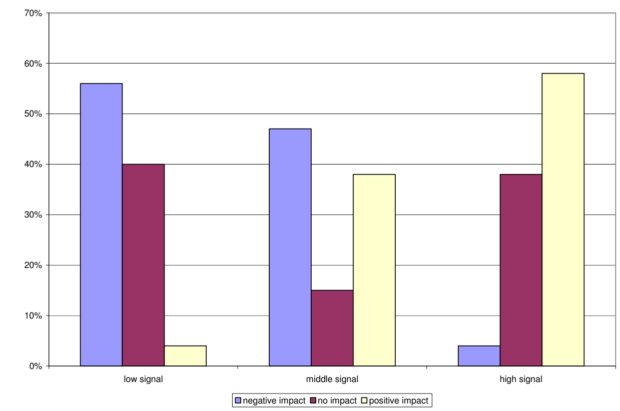

Round

Signals: **Case 1**

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

If the true effect will be POSITIVE then you receive

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

If the true effect will be NEGATIVE then you receive

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

If indeed the effect will be NO EFFECT then you receive

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

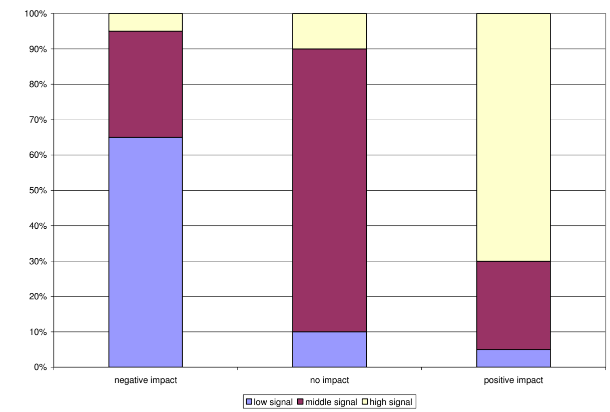

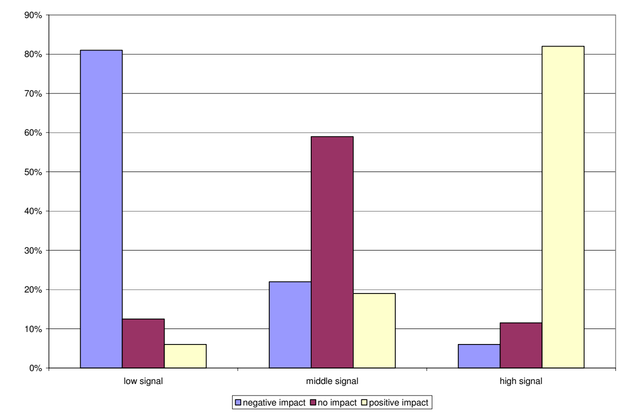

Figure 4
Information Sheet for negative Hill shape

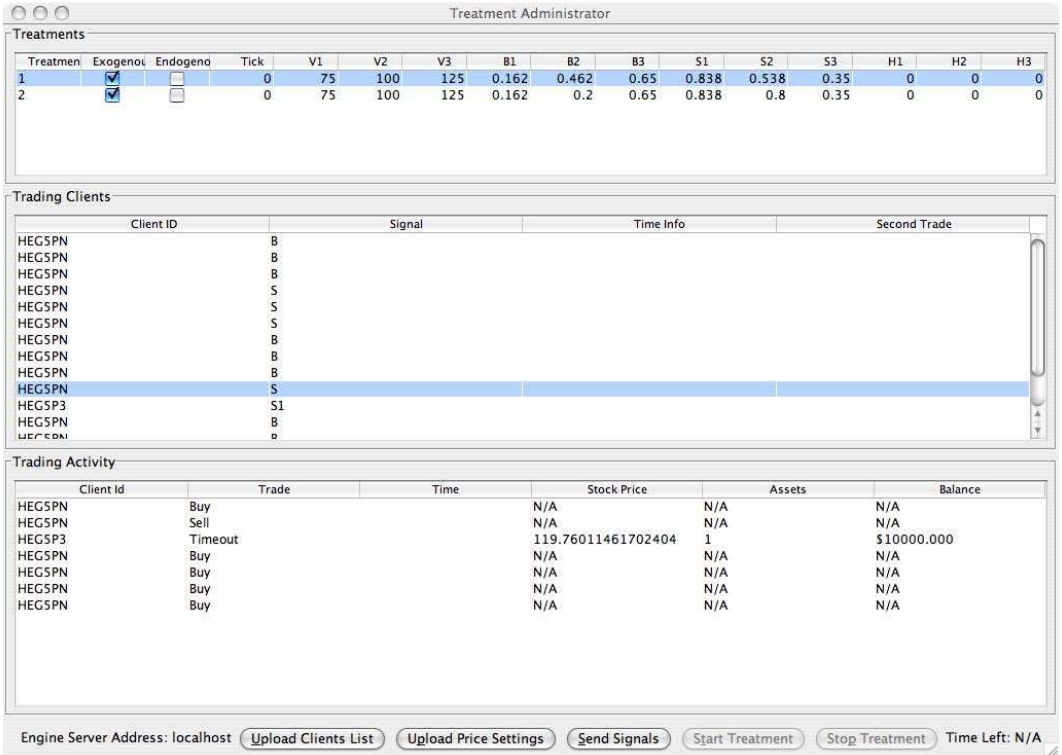

Figure 5
The Administrative Interface

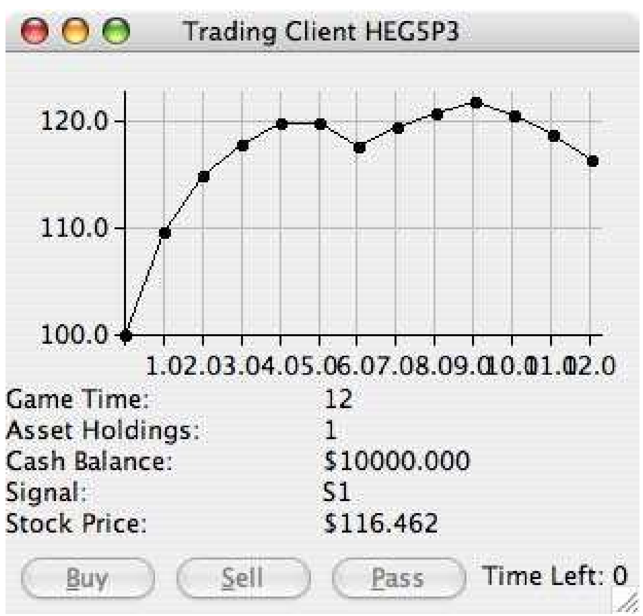

Figure 6
The Trading Client