

Human and Artificial Agents in a Crash-Prone Financial Market

Todd Feldman · Daniel Friedman

Accepted: 25 May 2010 / Published online: 17 June 2010
© The Author(s) 2010. This article is published with open access at Springerlink.com

Abstract We introduce human traders into an agent based financial market simulation prone to bubbles and crashes. We find that human traders earn lower profits overall than do the simulated agents (“robots”) but earn higher profits in the most crash-intensive periods. Inexperienced human traders tend to destabilize the smaller (10 trader) markets, but have little impact on bubbles and crashes in larger (30 trader) markets and when they are more experienced. Humans’ buying and selling choices respond to the payoff gradient in a manner similar to the robot algorithm. Similarly, following losses, humans’ choices shift towards faster selling.

Keywords Financial markets · Agent-based models · Experimental economics

1 Introduction

We insert human traders into an agent-based simulation model of a financial market that is susceptible to bubbles and crashes. The results illuminate both the simulation model and the market behavior of humans.

Bubbles and crashes have received a lot of attention recently, but as yet have no widely accepted theoretical explanation. Several agent-based models have been proposed. A leading example is [Brock and Hommes \(1997, 1998\)](#), who find bubble and crash dynamics in simulations when the majority of agents switch from a fundamentalist strategy to a trend-following strategy. A more recent example is the

T. Feldman (✉)
San Francisco State University, San Francisco, CA, USA
e-mail: todd18us@gmail.com

D. Friedman
University of California Santa Cruz, Santa Cruz, USA

Friedman and Abraham (2009) model of agents who adjust their leverage in response to a payoff gradient. Bubbles and crashes arise from an endogenous market risk premium that responds to investors' recent losses.

In his *Handbook of Computational Economics* survey, Duffy (2006) notes that agent-based simulation models are commonly calibrated to *aggregate* field data,¹ e.g., to the mean and variance of weekly returns in the S&P 500 index 1970–2007, but that it is much more difficult to assess the external validity of the model's assumed *individual agent* behavior using field data. For example, how do individual agents respond when others deleverage rapidly as in 2008–2009? Duffy points out that laboratory experiments with human subjects often provide the best available evidence on such adjustment parameters, and thus can illuminate simulation models. Indeed, one of the very first economic experiments sought to calibrate oligopoly simulations. Key parameters for individual reaction functions were difficult to infer from existing field data, so Hoggatt (1959) ran oligopoly markets using networked teletype terminals for humans together with simulated agents that he dubbed “robots.” Garman (1976), in an influential early study of financial market microstructure, had similar motivations and technique.

Duffy's survey notes that, conversely, simulation models can illuminate the behavior observed in laboratory and field markets. In particular, Gode and Sunder (1993) used “zero-intelligence” robots in simulations that help explain the rapid convergence to competitive equilibrium seen in laboratory markets with human traders. Carl Plat (1995) inserted robot liquidity traders into an laboratory financial market. In order to better understand the markets' stability and convergence properties, Cason and Friedman (1997), among others, used robots following equilibrium strategies to train human subjects.

The present paper combines both traditions. Section 2 summarizes the Friedman and Abraham (2009) simulation model, Sect. 3 describes its implementation as a laboratory experiment with human subjects, and Sect. 4 presents the results. The simulation model is robust in the sense that human traders have little impact on bubbles and crashes in larger (30 trader) markets and when they are more experienced. Human trading prowess is demonstrated by the fact that, although human traders earn lower profits overall than do the robots, they earn higher profits in the most crash-intensive periods. Validating the trading algorithm of the simulation model, we find that humans' buying and selling choices respond strongly to the payoff gradient. Also as assumed by the algorithm, humans respond to their own losses (and, to a lesser extent, to market wide losses) by selling.

Concluding remarks are collected in Sect. 5. Appendix A gathers details about implementations and presents supplementary data analysis. Appendix B is a copy of the instructions to laboratory subjects.

¹ Several agent-based financial market models recently have been successfully estimated on aggregate financial market data, e.g., Alfarano et al. (2005).

2 A Model of Bubbles and Crashes

Friedman and Abraham (2009) construct a model of portfolio managers who buy and sell a single riskless (“safe”) asset with constant return R_o and a single risky asset with variable return R_1 . Each portfolio manager chooses a single ordered variable $u \in [0, \infty)$ that represents the allocation to the risky asset. Choices $u > 1$ represent leverage, in which the safe asset is borrowed to purchase more of the risky asset. The manager’s net portfolio value is denoted by the variable z .

The price of the risky asset turns out to be

$$P = V\bar{u}^d, \quad (1)$$

where V is fundamental value, i.e., the present value of future dividends, while \bar{u} is the z -weighted mean allocation across all portfolio managers, and d is a parameter that captures the sensitivity of price to buying pressure. Friedman and Abraham (2009) show that realized yield on the risky asset then is

$$R_1 = (R_s - g_s)\bar{u}^{-d} + g_s + d\dot{\bar{u}}/\bar{u}. \quad (2)$$

where $R_s \geq R_o$ is the discount rate and $g_s < R_s$ is the growth rate. The first term represents the dividend yield, the second term represents capital gains due to underlying growth, and the third term represents short term capital gains (or losses) due to buying (or selling) pressure.

The payoff function of manager i is

$$\phi = R_i(u_i) = (R_1 + \alpha_i - R_o)u_i - \frac{1}{2}c_2u_i^2, \quad (3)$$

where the idiosyncratic excess return α_i is an Ornstein-Uhlenbeck stochastic process, i.e., mean reverting to zero in continuous time. The market-wide price of risk c_2 is determined endogenously, as follows. Denote manager i ’s current loss by $L_i = \max\{0, -R_{Gi}\}$, where the gross return R_{Gi} ignores the risk premium (or sets $c_2 = 0$) in Eq. (3). Let $\hat{L}_i(t)$ denote the exponential average of L_i (see Appendix A for details), and let $\hat{L}_T(t)$ denote the z_i -weighted average of the \hat{L}_i ’s. Then c_2 is proportional to these market-wide perceived losses,

$$c_2 = \beta \hat{L}_T(t), \quad (4)$$

where the parameter $\beta > 0$ reflects investors’ sensitivity to perceived loss.

The key behavioral assumption is that the managers adjust their exposure to risk by following the payoff gradient,

$$\phi_u = R_1 + \alpha_i - R_o - c_2u. \quad (5)$$

In continuous time, the adjustment equation for manager i is $\dot{u}_i = \phi_{u_i}$. That is, each manager continuously adjusts her risk position in proportion to her payoff slope. If

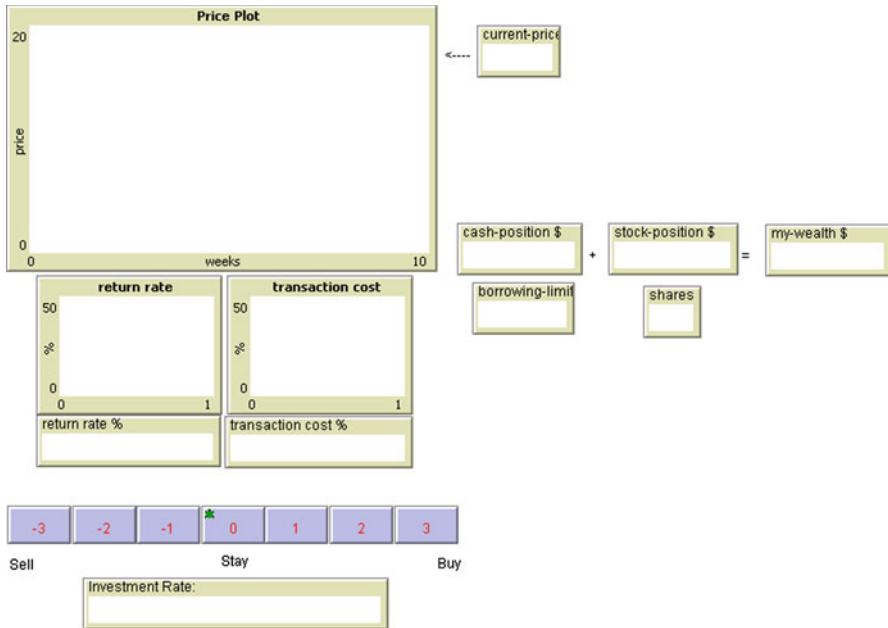


Fig. 1 Basic user interface

ϕ_u is positive (or negative) for a manager, she adds to (or sells off part of) her risky asset position, and does so more rapidly the steeper the payoff function at her current position. As explained in Friedman and Abraham (2009), quadratic transaction costs make it optimal to follow the gradient, and not to jump immediately to the current maximum of the payoff function. Note that the gradient differs across managers due to the α_i term in (5). Note also that the gradient depends on the current strategy choices and adjustments of all managers via the \bar{u} and \hat{u} terms in (2).

For a wide range of parameter configurations, Friedman and Abraham (2009) obtain occasional bubbles and crashes. When losses have been small for a long time, the risk premium shrinks and asset prices rise as fund managers increase leverage. Then slight losses can trigger a crash, as a wider risk premium causes deleveraging, asset price declines, further losses and a still wider risk premium.

3 Experimental Design

The experiment was conducted at University of California, Santa Cruz's Learning and Experiment Economic Projects (LEEPS) lab, using the Hubnet feature of NetLogo. NetLogo is a cross-platform multi-agent programmable modeling environment, and the HubNet feature enables human participation.² In our experiment, each human subject controlled a portfolio manager (called simply a "trader" in the instructions and below) using a screen display similar to one shown in Fig. 1.

² To learn more about NetLogo see <http://ccl.northwestern.edu/netlogo/>.

A typical session lasted 90 min and involved 5 inexperienced human subjects recruited by email from a campus-wide pool of undergraduate volunteers. Subjects silently read the instructions attached in Appendix B and then listened to an oral summary by the experiment conductor. After a couple of practice periods, subjects participated in about 12 periods. Then they were paid in US dollars according to the wealth achieved at the end of each trading period, typically receiving between \$15 and \$25.

In each trading period, each human subject can buy and sell the risky asset. His or her incentive is to maximize wealth by buying units (“shares”) when cheap and selling them when the price is high. Subjects are not told the pricing Eq. 1 nor the values of V , \bar{u} , and d . However, as shown in Fig. 1, they see the current price P and a graph of the price since the period started. The instructions say that price is determined by the growth rate and the interest rate, and by buying and selling pressure, and we write on the board that the growth rate is zero and the interest rate is three percent. We tell subjects that no single trader has much influence on P , but collectively their net buying pressure increases the price and net selling pressure decreases the price.

Each trading period consists of 20 “years,” and the subjects’ screens update on a “weekly” basis, as shown in Fig. 1. Before the trading period begins, each human subject is endowed with 500 “lab dollars” and 70 shares of stock. Current wealth is equal to the agent’s riskless asset holdings (“cash”), plus the number of shares owned times the current stock price. Wealth updates weekly due to price changes as well from interest earned on cash holdings or interest paid on net borrowing; leveraged buying is allowed up to a limit proportional to current wealth. Of course, buying shares decreases the cash position by the amount purchased times the stock price plus a transaction cost. Selling shares increases cash holdings by the proceeds minus a transaction cost. If a human goes bankrupt, i.e., his wealth falls to zero, then he is immediately barred from further trading that period and has a final wealth of -500 Lab Dollars for the period. All human subjects are equally re-endowed and are allowed to resume trading in the next period.

Figure 1 shows the basic interface for human subjects. Besides asset price, the display includes the components of wealth just mentioned. Near the bottom of the screen are seven buttons for buying and selling. Clicking the button labeled -3 sells shares at a rapid rate each week (until another button is clicked or share holdings reach zero), -2 sells shares at a moderate rate, and -1 at rather slow rate. Similarly, by clicking the button 3 (or 2 or 1) the trader purchases shares each week at a rapid (or moderate or slow) rate, and she ceases to buy or sell by clicking button 0.

In more detail, a human trader i ’s current wealth is $z_i = x_i + P y_i$, where P is the current asset price and $x_i \in (-\infty, \infty)$ and $y_i \geq 0$ denote current cash position and share holdings. His risk position then is $u_i = 1 - (x_i / z_i)$. Choosing the slow, moderate or rapid rate ($\pm 1, 2$ or 3) for buying or selling corresponds to a weekly change in u of $0.125, 0.25$ or 0.5 respectively. The corresponding transaction costs are quadratic, and so are equal to square of the weekly change times a constant c . Standard adjustment costs use $c = 1$, where a fast rate incurs a transaction cost of $0.5^2 = 25\%$ of the market value of the transaction, a moderate rate incurs a transaction cost of $0.25^2 = 6.25\%$, and a slow rate incurs a cost of only $0.125^2 = 1.6\%$.

3.1 Simulation Specifications

Robots are distinct from humans in that robots are programmed to react to the payoff gradient which includes the risk cost term based on market-wide losses. Humans may or may not behave similarly.

The integration of human and robot traders into a single market proceeds as follows. Each week the the NetLogo algorithm first updates each robot's allocation u_i based on i 's payoff gradient as defined in Eq. 5. At the same time, each human trader adjusts his or her allocation u_j by changing or staying with one of the seven adjustment rates. Second, the algorithm computes the new values of \bar{u} and \bar{u} by averaging across the new values of u_i and u_j weighted by portfolio size. Third, applying Eqs. 1–2 to these new values, the algorithm updates asset price and return (and portfolio size). Next the algorithm computes losses and updates the parameter c_2 . After a random increment to the the idiosyncratic shock, the algorithm then updates each robot's payoff gradient for the next week. More details can be found in Appendix A. Full documentation, source code and executable code can be found at <http://www.vismath.org/research/landscapedyn/models/markets/>.

Simulations are governed by a parameter vector specifying the discount rate, idiosyncratic shock volatility and so forth, as detailed in Friedman and Abraham (2009). The experiment uses the baseline parameter vector of that paper, except that volatility is increased from $\sigma = 0.2$ to 0.3 to make humans less able to predict price movements. The parameter vector thus is $R_o = 0.03$, $R_s = 0.06$, $g = 0.0$, $\sigma = 0.3$, $\tau = 0.7$, $\eta = 0.7$, $\beta = 2$, $d = 2$, $\lambda = 1$, $d = 1$, and $rate = 1.3$; see Appendix A for all parameter definitions.

3.2 Treatments

We employ three treatment variables, as follows.

- Number of human and robot traders, NT. We examine five levels of NT: (0 humans, 30 robots), (1 human, 29 robots), (5 humans, 25 robots), and (5 humans, 5 robots) (0 humans, 10 robots). The first and last levels are run separately as controls. The other three levels are run as blocks in each session. The sessions all have four blocks of three (1040 week) periods, and the treatments are rotated across blocks with the first and last block the same. The level of NT is publicly announced at the beginning of each block.
- Graphical Interface, IF. In some sessions, human subjects have only the basic user interface described earlier. In other sessions they have the enhanced interface shown in Fig. 2, which also shows the distribution (“Density Chart”) of other traders’ risk positions, and a graph of the payoff function (the “Landscape”). The enhanced interface also includes a NetLogo graphics window that displays the positions (u_i , on the horizontal axis) and portfolio sizes (z_i , on the vertical axis) of each trader, with robots plotted as as small triangles and human traders as round dots. The subject’s own current position and and wealth is indicated by a dot of a different color.
- Transaction Costs, TC. In some sessions, humans incur the standard transaction costs described earlier ($c = 1.0$), while in other sessions they incur low transaction costs ($c = 0.5$), half as large as standard.




Fig. 2 Enhanced user interface with graphics window, landscape, and density chart

3.3 Sessions

We have collected data in 9 sessions, configured as follows.

- Sessions 1, 2 and 3 use only the basic interface (IF = B) and standard transaction costs ($c = 1.0$). The NT rotation in session 1 is (1 human and 29 robots) in the first block, (5 human, 5 robots) in the second block, (5 human and 25 robots) in the third block, and finally (1 human, 29 robots) again in the last block. In session 2, the rotation is (5 human, 5 robots), (5 human and 25 robots), (1 human and 29 robots), and (5 human, 5 robots). In session 3 it is (5 human and 25 robots), (1 human and 29 robots), (5 human, 5 robots), and (5 human and 25 robots).
- Sessions 4, 5, and 6 are the same as 1, 2 and 3 except that they use the enhanced interface (IF = E).
- Sessions 7, 8, and 9 are the same as 1, 2 and 3 except that they use the enhanced interface (IF = E) and low transaction costs ($c = 0.5$).

4 Results

Figure 3 shows a prime example of a bubble and crash. The fundamental value V , given the baseline parameters, is $\frac{1}{R_s - g_s} = \frac{1}{.06 - 0} \approx 17$, but in this period the price rises well above that by week 300 and stays high (with substantial fluctuations) until about week 800 when it crashes down to about 8 where it remains (again with fluctuations) for the rest of the 1,040 week period. One sees many such episodes in other periods, albeit often less dramatic.

There is no universally accepted definition of a crash so, somewhat arbitrarily, the analysis to follow defines it as a decline in price P of at least 50% from its highest point within the previous 26 weeks.




Fig. 3 Asset price in period 7 of session 7

The rest of this section presents answers to five empirical questions: whether humans change the distribution of crashes, outperform robots, herd, follow a payoff gradient and react to exponential average of their own and market wide losses.

4.1 Do Humans Provoke Crashes?

We compare the frequency of crashes in periods with human traders to those with robots only. Figure 4 shows that periods with humans seem to have the same crash distribution as those with no humans in the treatments with 30 traders. However, in treatments with 10 traders, crashes seem more frequent with human traders. Table 1 confirms this impression using the Kolmogorov–Smirnov test.

Is there any time trend? As detailed in the Appendix, we find that the crash frequency declines consistently and significantly over time in the (5 humans, 5 robots)




Fig. 4 Distribution of crashes by population

Table 1 Kolmogorov–Smirnov test


Treatments	Nobs	D	<i>p</i> value
1 Human, 29 robots vs. 30 robot	180	0.06	0.50
5 Human, 25 robots vs. 30 robots	36	0.25	0.15
5 Humans, 5 robots vs. 10 robots	36	0.50	0.00

Nobs report the number of observations in each empirical distribution function, D is the maximum difference between the two cumulative distribution functions, and *p* value reports the significance of the standard Kolmogorov–Smirnov test

treatment with standard transaction costs, but we find no time trend in the other treatments.

4.2 Can Humans Beat Robots?

Figure 5 displays the average end-of-period wealth for human subjects and for the robots used in the experiment. As detailed in Appendix A, the figure adjusts robots' wealth to make it comparable to humans', by removing the random component from initial endowment and by imposing transaction costs. The figure also shows final wealth for the passive no-trade (i.e., buy-and-hold) strategy, given the closing asset prices seen in the experiment. Evidently, gradient adjustment robots outperform both humans and the passive strategy in all three environments, and standard t-tests confirm

**Fig. 5** Wealth comparison

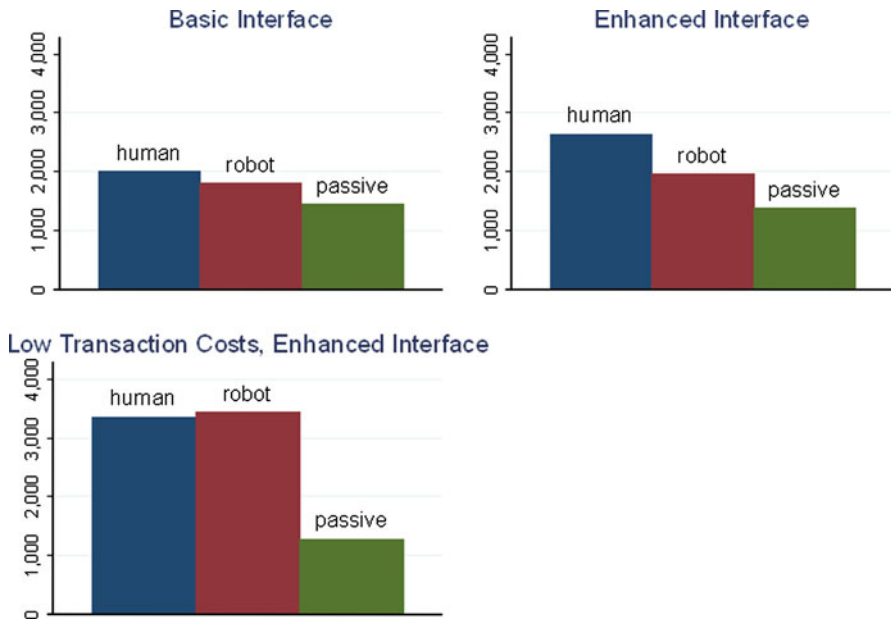


Fig. 6 Wealth comparison during high crash periods

that the differences are significant. The t tests also confirm that humans do significantly better than the passive strategy in the environment with low transaction costs and an enhanced user interface, but show that the differences are insignificant in the other two environments.

Figure 6 shows that human traders have a comparative advantage in crash periods, that is, in the (5 humans, 5 robots) treatment. The advantage is statistically significant in the standard transaction cost treatments with both basic and enhanced interface.

Why do humans do better in crises? Perhaps they can look further down the road, and pursue contrarian strategies. To investigate, we now take a closer look at how humans choose their actions.

4.3 Choosing Adjustment Rates

Figure 7 shows the overall distributions of adjustment rates.

About half the time, human traders remain at adjustment rate 0 and so are inactive. The figure shows clearly that human traders are sensitive to market frictions. They use the slower adjustment rates more often than the moderate rates, and rarely use the fastest rates except when transaction costs are low.

The last panel of the figure sorts robots' continuous gradient adjustment choices according to the closest of the seven choices given to humans. (For example, if a robot chose a gradient of 0.1 then that choice would be coded as a slow buy, 0.125, since it is between 0.0625 and 0.1875.) The result is also a unimodal distribution around 0, but closer to the uniform distribution.

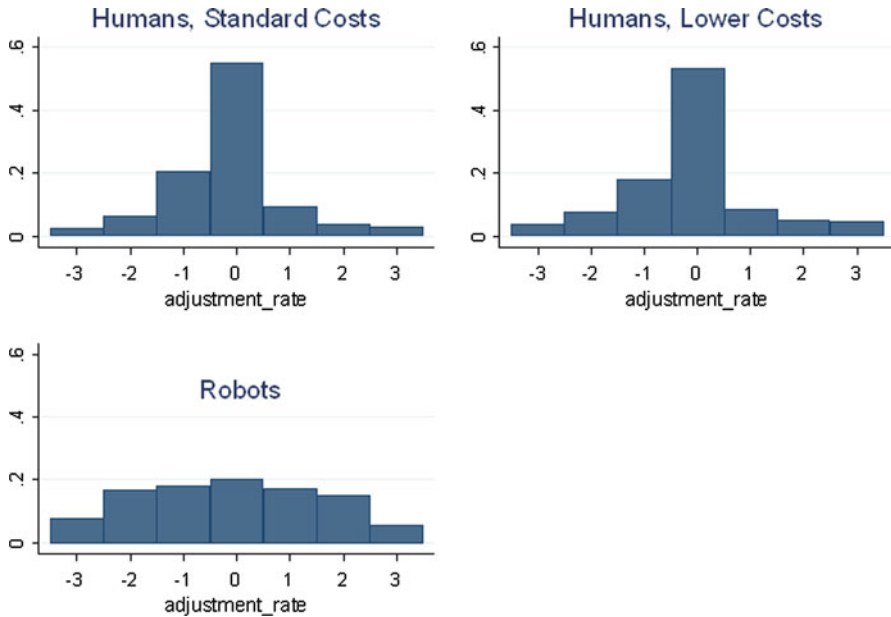


Fig. 7 Frequency of adjustment-rates

Figures 8 and 9 address the question of why humans outperform robots in high crash periods.

We take a closer look at the 5 human, 5 robot data to examine the humans' adjustment rates in the 52 "weeks" prior to the onset of a crash ("Anticipation"), and the weeks for which price is less than half its maximum over the preceding 26 weeks ("Crash"). Figure 8 shows that humans tend to sell in anticipation of a crash and tend to buy at the end of a crash while robots tend toward extremes at the end of a crash. Behavior in other periods resembles that shown in Fig. 7. We also compare adjustment rates of the humans who earn above the average final wealth to humans who earn below average final wealth. Figure 9 shows that the higher earners sell more before the crash and buy more after the crash.

4.4 Herding?

Do human traders influence each others' allocation decisions? The graphics window in the enhanced user interface allows them to see the current leverage position of all traders (robot and human), and Figs. 10 and 11 suggests that this interface may encourage humans to herd together. These figures show typical snapshots taken half way through a trading period, and the human traders (indicated by the colored circles) are much more tightly bunched (at very low values of u) in the Enhanced Interface panel. The dispersion of robots (indicated as red triangles) seems about the same in both panels, as one might expect.




Fig. 8 Humans vs. robot during anticipation and end of crash period




Fig. 9 High vs. low human earner groups




Fig. 10 Basic interface treatment (BT, humans cannot see others). Snapshot taken at week 520 from session 5 period 6



Fig. 11 Enhanced interface treatment (ET, humans cannot see others). Snapshot taken at week 520 session 9 period 3

Figure 12 examines the question quantitatively. It shows that the average correlation coefficient between human adjustment rates are very similar in the relevant population 30 treatments, even in periods with crashes. However, in the population 10 treatments, crashes are more frequent, and here correlation coefficients are much higher, especially in crash periods. Standard one-sided t tests confirm that the differences are insignificant for the entire period ($t = 0.70$) and for the Crash weeks ($t = 1.11$) in the population 30 treatments, but are quite significant in the population 10 treatments ($t = 2.89$). One possible interpretation is that in the more stressful environments, humans learn from each other to be contrarians.

4.5 Do Humans Follow the Gradient?

We are now prepared to investigate one of the primary questions that motivated the experiment: are gradient-following robots at all like human traders? Perhaps the most direct test is to regress human traders' adjustment choices a_h on the gradient evaluated at the trader's current allocation u_h . We specify

$$a_h = \beta_0 + \beta_1 * \widetilde{\phi_{uh}} + e, \quad (6)$$

where $\widetilde{\phi_{uh}}$ for human traders is given by (5) truncated to lie between ± 0.50 . The truncation reflects humans' bounded choice set. Of course, for human traders, α in (5) is set to zero.

A β_1 estimate near 1.0 indicates gradient adjustment, but of course, the discrete nature of humans' adjustment choice precludes attaining that value. To find a more appropriate benchmark value, Table 2 reports regression results for robot data, where




Fig. 12 Correlation of adjustment-rates

Table 2 Gradient adjustment estimates by population treatment

Sample	1 Human, 29 Robots	5 Humans, 25 Robots	5 Humans, 5 Robots
Human	0.82** (0.006)	0.67** (0.004)	0.33** (0.002)
Human nobs	185,132	185,667	181,285
Robots	0.99** (0.000)	0.93** (0.000)	0.64** (0.000)
Nobs	5,418,360	934,200	186,840

Random effects slope coefficient estimates (and standard errors) from Eq. 6. ** Significant at 1%
Variables are as in Table 3

the dependent variable is the sorted gradient choice reported in the last panel of Fig. 7. The Table 2 shows that humans are not exactly following a gradient but are surprisingly close; in each case the estimated coefficient for humans is more than half as large as robots'.

4.6 Do Humans React to Recent Losses?

A second motivating question concerns the endogenous risk premium: do humans respond to an exponential average of their own losses \hat{L} , and/or to an exponential average of market wide losses, c_2 ?

Table 3 Random effects estimation results, dependent variable: human adjustment-rate

Sample	1 Human, 29 Robots	5 Humans, 25 Robots	5 Humans, 5 Robots
Intercept	−0.90** (0.02)	−0.40** (0.02)	0.10** (0.02)
x_h	0.10** (0.01)	0.40** (0.01)	0.30** (0.01)
y_h	0.20** (0.01)	0.13** (0.01)	0.40** (0.01)
z_h	−0.10** (0.01)	−0.70** (0.01)	−0.50** (0.01)
$R1$	4.95** (0.11)	0.20** (0.03)	0.19** (0.01)
$\widetilde{\phi_{uh}}$	0.80** (0.00)	0.65** (0.00)	0.24** (0.00)
\hat{L}	−0.68** (0.03)	−1.50** (0.03)	−0.42** (0.01)
$\hat{L} * IFE$	0.49** (0.04)	0.77** (0.04)	0.34** (0.01)
$\hat{L} * LTC$	−0.62** (0.04)	−0.37** (0.03)	−0.23** (0.01)
c_2	0.89** (0.01)	0.40** (0.01)	0.19** (0.00)
$c_2 * IFE$	−0.23** (0.01)	−0.24** (0.02)	−0.07** (0.01)
$c_2 * LTC$	0.20** (0.01)	0.30** (0.05)	0.06 (0.01)
Crash	0.07** (0.03)	−0.04** (0.00)	0.01** (0.00)
nobs	185,132	185,667	181,285
R^2	0.26	0.20	0.19
$\chi^2_{(8)}$		65,464	57,422

Estimates (and standard errors) from random effects regressions of human traders' adjustment rates a_h on the following explanatory variables: cash holdings x_h , share holdings y_h , wealth z_h , weekly return $R1$, truncated gradient ϕ_{uh} , crash dummy Crash, enhanced interface dummy IFE, and low transactions cost dummy LTC. The number of observations *nobs* differs slightly across treatments due to bankruptcies.

** Significant at 1%

Table 3 presents regression results with controls that include the human traders level of cash, shares, wealth, return, and an indicator variable for crash period. The dependent variable is the adjustment rate chosen by each human each week.

The coefficients reported for \hat{L} and c_2 refer the baseline treatment—standard transaction costs and no ability to view other traders. Consistent with the robot algorithm, these reported coefficients indicate that humans do respond to their own losses by

selling. The c_2 coefficient is positive, suggesting that humans are net buyers as market wide losses increase. (However, more detailed regression results, seen in the appendix, indicate that, when personal losses are removed from the regression equation, humans respond to market wide losses by selling.) The interaction terms indicate that human traders respond more strongly to losses during low transaction cost treatments than in the baseline, but that they respond less strongly to losses when allowed to view other human traders. The Appendix also shows that alternative specifications of exponential losses sharpen these results.

5 Discussion

Agent-human interactions in financial markets can have surprising effects. A cautionary example is the US stock market collapse of October 1987, which the official Brady commission (e.g., [Brady et al. 1988](#)) attributed in no small part to automated agents executing sell orders for “portfolio insurance. “Automated agents today play a larger role than ever, yet little is known about their impact on market performance. For example, a week after the event, investigators of the startling and mysterious downspike in US stock prices on May 6, 2010 began to focus on automated agents that interacted badly with “circuit-breakers” designed for humans ([Wyatt 2010](#)).

This paper takes some first, preliminary steps towards understanding the implications of agent-human interactions. We introduced human subjects as traders (or portfolio managers) into the bubble-and-crash prone simulations of [Friedman and Abraham \(2009\)](#), featuring agents (“robots”) that follow the payoff gradient in a financial market. Notable findings include:

1. Humans adapt reasonably well to the simulated financial market. Overall, human traders earn profits somewhat lower than do the robots but the humans actually outperform the robots in the most challenging, crash-intensive, periods.
2. Human traders tend to destabilize the smaller (10 trader) markets, although less so as they become more experienced. They do not seem to stabilize or to destabilize the larger (30 trader) markets.
3. To a significant extent, human traders follow buying and selling strategies similar to the robots’. Even though human subjects in our experiment are not told the robot strategies, their purchases and sales respond systematically to the payoff gradient, with the same sign as (and more than half of the magnitude of) the robots.
4. The simulation model seems to capture some aspects of how human traders react to losses. Other things equal, humans in our experiment tend to purchase more slowly and sell faster when they personally have recently experienced greater losses. Their reaction to an index of perceived market-wide risk is less responsive than to personal losses. However, it is still consistent with the simulation model.

Are there more general lessons that can be learned from our exercise? Each period of our simple market spans two simulated decades, so, even initially inexperienced undergraduate traders can gain considerable expertise. By design, our markets are far simpler than major modern financial markets. This simplicity allows us to make

clear inferences, and tell us better what to look for in field markets. Our human traders showed that there is some degree of realism in the gradient trading algorithm, and the response to losses, assumed in the simulation. (As noted in the introduction, it is hard to use field data to validate or falsify such individual-level assumptions.)

We believe that our results also shed some light on human behavior. It did not take mob psychology—manias or panics—to create bubbles and crashes. It is sufficient to have a positive feedback loop, as in the simulation, between recent trends in asset returns, realized losses, and the risk premium. The current study also showed that human traders, although not as quick or as profitable overall, were better than the current generation of robots in adapting to the most crash-prone financial markets. The nature of those adaptations suggests that bubbles and crashes in field markets are less likely after human traders have experienced severe turbulence.

Our study opens new territory for exploration. For example, one can compare the abilities of various simulation models (including the [Brock and Hommes \(1997, 1998\)](#) switching models) to account for data generated by other models and by humans.

In a different direction, one might examine existing field data (or even laboratory data) for evidence of gradient-following behavior, or for evidence that sentiment or risk premia respond to recent losses, or for a variety of other human trading strategies. More broadly, we hope our study encourages a more active interplay between analysis of field data, laboratory experiments with human subjects, and agent based simulation models.

Acknowledgements We are grateful to the National Science Foundation for support under grant SES-0436509, Pacific Rim Research Program for a mini grant and its seminar audiences at UCSC, the 2008 ESA conference at Caltech, 2008 Georgia State department seminar, and the 2009 CeNDEF workshop in Amsterdam. We thank Ralph Abraham and Paul Viotti for their invaluable assistance, and two anonymous referees for guidance on final revisions. We retain sole responsibility for remaining idiosyncrasies and errors.

Open Access This article is distributed under the terms of the Creative Commons Attribution Noncommercial License which permits any noncommercial use, distribution, and reproduction in any medium, provided the original author(s) and source are credited.

6 Appendix A

The definitions of simulation parameters are collected in Table 4.

6.1 Wealth Comparison

In order to compare wealth between the two types, we calculate what robot wealth would be under the exact conditions humans face. Therefore, we adjust the robot variable in three ways. First, we set their initial endowment equal to the human endowment. Second, we translate their gradient into one of the seven adjustment rates as explained in the main text. Third, we impute transaction cost based on the translated gradient.

Table 4 Parameter definitions

Parameter	Definition
$R0$	Risk-free rate
dR	Discount rate
g_s	Growth-rate of economy
σ	Variance of idiosyncratic shock
τ	Persistence of idiosyncratic shock
η	Memory rate
β	Sensitivity to risk
d	Price pressure sensitivity
λ	Exponential rate of recruiting from z-pool by a successful manager
d	Portion of holdings to withdraw
rate	Controlling the overall recruiting rate

After making those adjustments we then calculate step by step how wealth for each robot changes overtime until the period ends.

6.2 Integrating Robots & Humans

The decision variables of humans and robots have different scaling properties. The robot allocation choice, u_i , is a number between zero and four while the human's allocation variables, cash and shares, are in the hundreds. The scale is different because it is easier for humans to understand holding dollars and shares in the hundreds rather than in decimals. However, since humans' decisions, alongside robots' decisions, affect the environment, price, their allocation choice and portfolio size variables, u_h and z_h , need to be consistent with robot decision variables. Therefore, we translate the humans' shares into an appropriate u_h , which equals one minus the ratio of their cash to wealth,

$$u_h = 1 - \left(\frac{\text{cash}_h}{\text{wealth}_h} \right), \quad (7)$$

where h refers to a specific human. Under this specification, if cash holdings equals zero then u_h equals 1, meaning they are fully invested in the risky asset, while if cash holdings is negative then $u_h > 1$ meaning they are borrowing the safe asset and if cash holdings is positive then $u_h < 1$ meaning they are investing a portion in the safe and risky asset.

The human's portfolio size, z_h , does not need to be scaled; it changes similar to how the robot's portfolio size, z_i , changes, based on the return, inflow rate, and outflow rate.³ However, there are two subtleties that differentiate the portfolio size z of the

³ See [Friedman and Abraham \(2009\)](#) for an explanation of the inflow and outflow rate.

human and robot trader. First, all humans receive an initial z_h equal to 1. The robots receive their initial risk allocation, u_i , and portfolio size, z_i , randomly via a uniform distribution.⁴ All humans receive the same endowment in order to avoid giving any one an advantage or disadvantage. Second, robots receive an idiosyncratic shock and humans do not. Also, humans pay transaction costs unlike robots, who implicitly face quadratic adjustment costs. These implicit quadratic adjustment costs are the underlying reason why robots follow a gradient. Humans incur quadratic transaction costs because we want to analyze their behavior under the same assumed conditions. Human transaction costs are,

$$t_h = c(\text{adjustment-rate}_h)^2, \quad c = \text{constant}, \quad (8)$$

where given that the constant, c , is set to 1, trading at a fast rate incurs a transaction cost of 25%, a medium rate incurs a cost of 6.25%, and a slow rate incurs a cost of 1.6%. In the low transaction cost treatment, the constant c is reduced 50% to 0.5.

Another integration issue involves buying and selling. The buttons -3 , -2 , -1 , 0 , 1 , 2 , 3 shown on the interface were chosen for ease of viewing. The actual rates are 0.125 for a slow rate, 0.25 for a medium rate, and 0.5 for a fast rate. These rates were chosen based on the standard deviation of the robot's chosen gradient, 0.125, in an all robot simulation using a baseline configuration.

6.3 Frequency of Adjustment-Rates During Crashes

Figure 13 shows that humans sell more at the beginning of a crash and buy towards the end of crash. Humans in the lower transaction cost treatment also use the faster rate more frequently than humans in the standard transaction cost treatment.

We also run each regression with an intercept term. The intercept estimate is statistically different from zero for the humans (Table 5).

6.4 Human and Robot Loss Regressions

Table 5 and 6 rerun the loss regression using human and robot data. In each regression, we use the same explanatory variables.

6.5 Losses Regression Using Different Half Lives

In order to determine how much weight humans put on their current and past losses we run the same regression as in Table 3 using different values of η to calculate \hat{L} . Investors seem to judge managers by the overall historical track record, with greater emphasis on more recent results. The natural formalization is an exponential average,

⁴ The distribution is supported on a $[u; z]$ rectangle set via sliders; the default is $[0.2; 1.4] \times [0.4; 1.6]$.




Fig. 13 Frequency of adjustment rates before crash and toward the end of crash over different treatments

Table 5 Random effects estimation results, dependent variable: human adjustment-rate

Variable	1 Human, 29 Robots	5 Humans, 25 Robots	5 Humans, 5 Robots
Intercept	−0.03** (0.02)	0.10** (0.01)	0.40** (0.03)
u_h	0.09** (0.02)	0.40** (0.01) (0.02)	0.97** (0.03)
$c2 * u_h$	−0.28** (0.02)	−0.20** (0.02)	0.10 (0.01)
$R1$	2.90** (0.02)	0.88** (0.03)	1.56** (0.02)
\hat{L}	−0.89** (0.03)	−1.23** (0.03)	−0.25** (0.01)**
$\hat{L} * IFE$	0.22** (0.04)	0.72** (0.04)	0.42** (0.01)
$\hat{L} * LTC$	−0.76** (0.05)	0.06** (0.04)	−0.15** (0.01)**
$c2$	0.10** (0.01)	−0.08** (0.02)	0.21 (0.01)
$c2 * IFE$	−0.18** (0.02)	0.15** (0.02)	−0.08** (0.01)
$c2 * LTC$	0.01 (0.02)	0.05** (0.02)	0.01 (0.01)
Crash	0.61** (0.03)	0.06** (0.01)	−0.07** (0.00)
Nobs	185,132	185,667	181,285
R^2	0.07	0.01	0.08
$\chi^2_{(8)}$	22,331	7,253	21,933

Estimates are shown with their standard errors in the parenthesis from a regression where the dependent variable is the adjustment-rate, a_h , of human h at time t . u_h refers to the level of risk holding, $R1$ is the weekly return. $c2 * u_h$ is an interaction term of $c2$ times u_h and Crash is an indicator variable that assigns a 1 to the time period of a crash and 0 otherwise. IFE is an indicator variable that equals one for sessions where human can see other traders and 0 otherwise and LTC is an indicator variable that equals one for sessions with low transaction costs and 0 otherwise. Observations between population treatments are slightly different depending on the number of bankruptcies. ** Significant at 1%. Variables are as in Table 3

defined recursively from its previous value $\hat{L}_i(t - h)$ as follows

$$\hat{L}_i(t) = e^{-\eta h} \hat{L}_i(t - h) + (1 - e^{-\eta h}) L_i(t). \quad (9)$$

We run several regressions using different η values to determine how far back human traders remember their losses. We run the following regressions by population

Table 6 Random effects estimation results, dependent variable: robot adjustment-rate

Variable	1 Human, 29 Robots	5 Humans, 25 Robots	5 Humans, 5 Robots
Intercept	0.03** (0.02)	-0.03** (0.02)	-0.08** (0.02)
u_i	0.90** (0.01)	0.96** (0.01)	1.30** (0.01)
$c2 * u_i$	-0.72** (0.01)	-0.60** (0.01)	-0.73** (0.01)
$R1$	0.49** (0.01)	0.40** (0.02)	0.45** (0.04)
\hat{L}	-1.68** (0.06)	-1.67** (0.06)	-0.12** (0.01)
$\hat{L} * IFE$	0.29** (0.04)	0.24** (0.01)	0.04** (0.01)
$\hat{L} * LTC$	-0.12** (0.04)	-0.17** (0.01)	-0.04** (0.01)
$c2$	0.39** (0.01)	0.42** (0.01)	0.01 (0.05)
$c2 * IFE$	-0.01** (0.01)	-0.01 (0.01)	0.05 (0.01)
$c2 * LTC$	0.01** (0.01)	-0.02 (0.01)	0.08 (0.01)
Crash	-0.07** (0.00)	-0.04** (0.00)	0.51** (0.00)
Nobs	5,418,360	934,200	186,840
R^2	0.26	0.25	0.16
$\chi^2_{(8)}$	301,284	286,344	35,371

Estimates are shown with their standard errors in the parenthesis from a regression where the dependent variable is the adjustment-rate, a_i , of robot i at time t . u_i refers to the level of risk holding, $R1$ is the weekly return. $c2 * u_i$ is an interaction term of $c2$ times u_i and Crash is an indicator variable that assigns a 1 to the time period of a crash and 0 otherwise. IFE is an indicator variable that equals one for sessions where human can see other traders and 0 otherwise and LTC is an indicator variable that equals one for sessions with low transaction costs and 0 otherwise. Observations between population treatments are slightly different depending on the number of bankruptcies. ** Significant at 1%. Variables are as in Table 3

treatment because we assume that different population treatments may affect how humans remember their realized losses.

Tables 7, 8, and 9 report regression results using standardized coefficients in order to compare estimates across different values of η . Results from Tables 7 and 8 indicate a value of η between 0.30 and 1.30 characterizes human memory the best for the (1 Human, 29 Robots) and (5 Humans, 25 Robots) treatments. These regressions report the greatest magnitude of coefficients and the highest values of $\chi^2_{(8)}$. This provides

Table 7 1 Human, 29 robots: random effects using different η s

Variable	$\eta = 3$	$\eta = 1.3$	$\eta = 0.7$	$\eta = .3$	$\eta = .1$
$\hat{L}_{h,t}$	-0.09** (0.016)	-0.12** (0.02)	-0.11** (0.028)	-0.16** (0.036)	-0.11** (0.08)
\hat{L}_{h,t_h5_r25}	0.8** (0.021)	5.9** (0.03)	13.5** (0.037)	18.4** 1 (0.09)	18.0** (0.10)
\hat{L}_{h,t_h1_r29}	0.04** (0.004)	0.067** (0.03)	0.08** (0.04)	0.17* (0.03)	0.16* (0.12)
$\chi^2_{(8)}$	78450.58	78935.52	79106.36	80237.31	78406.74

** Significant at 1%; standardized coefficients

Table 8 5 Humans, 25 robots: random effects using different η s

Variable	$\eta = 3$	$\eta = 1.3$	$\eta = 0.7$	$\eta = .3$	$\eta = .1$
$\hat{L}_{h,t}$	-0.25** (0.015)	-0.26** (0.020)	-0.26** (0.027)	-0.23** (0.034)	-0.21** (0.079)
\hat{L}_{h,t_h5_r25}	0.16** (0.019)	0.14** (0.026)	0.13** (0.037)	0.05** (0.074)	0.07** (0.094)
\hat{L}_{h,t_h1_r29}	-0.01** (0.015)	-0.03** (0.022)	-0.04** (0.035)	-0.02** (0.026)	-0.06** (0.089)
$\chi^2_{(8)}$	62980.82	64520.72	65464.79	64898.95	62687.09

** Significant at 1%; standardized coefficients

Table 9 5 Humans, 5 robots: random effects using different η s

Variable	$\eta = 3$	$\eta = 1.3$	$\eta = 0.7$	$\eta = 0.3$	$\eta = 0.1$
$\hat{L}_{h,t}$	-0.16** (0.036)	-0.26** (0.007)	-0.32** (0.008)	-0.20** (0.009)	-0.27** (0.013)
\hat{L}_{h,t_h5_r25}	0.12** (0.004)	0.20** (0.009)	0.20** (0.01)	-0.19** (0.018)	-0.07** (0.02)
\hat{L}_{h,t_h1_r29}	-0.05** (0.004)	-0.07** (0.008)	-0.08** (0.01)	0.02** (0.009)	-0.08* (0.027)
$\chi^2_{(8)}$	54159.50	55693.53	57422.40	63791.14	67095.58

** Significant at 1%; standardized coefficients

some support for using a value of η equal to 0.70 (or about a 1 year half life) in the simulation's baseline configuration. However, Table 9 indicates the lowest η , 0.1, is the correct memory parameter to characterize human memory in the hard hit, high




Fig. 14 Individual human exponential average of losses vs. returns. Each number corresponds to a human trader

crash (5 Humans, 5 Robots) treatment. Humans remember the harder hit crashes better. For example, people remember the Great Depression better than the recession that occurred in 1948.

For a robustness check we run the loss regression by individual and look for estimates that are outliers. Figure 14 shows individual estimates of \hat{L}_h and return by population treatment. We then took out the humans that were noted as outliers and ran all regressions again. Results did not change and thus passed our robustness check.

7 Appendix B: Instructions for Financial Market Experiments

Introduction

You will be a trader in a stock market where you can buy and sell shares of a single stock. The market runs for 12 trading periods. These instructions explain how you trade shares and how you make profits.

All quotes and profits during the experiment will be in lab dollars. At the end of the last trading period, your total profit will be converted into US dollars at a conversion rate written on the board, and paid to you in cash.

Time

Each trading period consists of 20 “years,” and there are 52 “weeks” in each year. A week only takes a second or two of real time in this experiment, so a year goes by

pretty quickly! Each week, your computer screen updates your trades and wealth, and the time remaining in that trading period.

At the end of each trading period, your computer screen reports your profit that period, and the cumulative profit you have earned so far. After the last trading period, it shows your total profit in lab dollars.

Profits and Wealth

Your profit each period is your final wealth less a fixed fee. The fee is written on the board.

Before a trading period begins you will receive \$500 in cash and 75 shares of stock. Your wealth at any point during the trading period equals your cash plus the amount of shares owned times the current share price. Before trading begins:

$$\text{Wealth}_{\text{initial}} = \$500 + 75 \text{ shares} * \text{price}_{\text{initial}}$$

At the end of the trading period:

$$\text{Profit} = \text{Wealth}_{\text{final}} - \text{Fee} = \text{Cash}_{\text{final}} + \text{Shares}_{\text{final}} * \text{price}_{\text{final}} - \text{Fee}$$

Your computer screen helps you keep track of your wealth. The display is similar to this: Your stock position is the number of shares times the share price.

How Wealth Changes

Even when you don't trade, your wealth changes due to:

1. **Interest:** If you hold a positive amount of cash your cash savings will increase with each week. During the trading period your cash grows at 3% per year. However, if you borrow cash, which is the same as having negative cash, your cash position will decrease with every week or worsens at 3% per year.
2. **Stock Dividends:** If you own shares, you may receive additional shares each week at a rate announced at the beginning of each trading period. (Sometimes the rate is zero, and occasionally it is negative.) For example, if the dividend rate is 2%, then the number of shares you hold increases 2% each year.
3. **Share Price:** If you own shares, then your wealth will change with the current share price. As explained below, the price of a share of stock changes every week.

EXAMPLE. Suppose you have \$250 in cash, 50 shares of stock, the interest rate is 3%, the dividend rate is 2%, and the current share price is \$10. Your wealth is $\$250 + 50 * \$10 = \$750$. Suppose that a “year” later the share price is \$15 and you did not buy or sell any stock in the meantime. Then your wealth becomes:

$$\$250 * (1 + 0.03) + 50 * (1 + 0.02) * \$15 = \$1022.$$

Share Price

The share price is determined by the:

1. **Dividend rate:** The higher the growth-rate the higher the share price.
2. **Interest rate:** The higher the interest-rate the lower the share price.

3. **Buying and Selling Pressure:** No one trader has much impact on the share-price, but if many traders are selling at the same time they will tend to temporarily drive down the share price. Likewise, the share price will tend to rise when lots of traders are buying at the same time.

Although the dividend rate and the interest rate remain constant during a trading period, the share price will change every week. Below is an example of a price plot which you will see on your computer screen.

Buying and Selling

Below is an image of how you will buy and sell. When you click a button you are buying or selling shares at a specific rate. Once you click a button (just once will do) your shares will change at that rate each week until you click another button.

A negative number indicates a sell and a positive number indicates a buy. Zero indicates you are neither buying nor selling. The larger the number the faster the rate you are buying or selling. Therefore, if you click “3” then you will buy at the fastest rate possible and if you click “1” you will buy at the slow rate. To stop buying or selling you must click “0”. In order for you to remember whether you are buying or selling and at which rate you are buying or selling, a message box appears underneath the trading buttons.

Trading Limits

Short selling is not allowed in this game. If your stock position is zero, a message will appear on your screen saying that you have zero stock, and the computer will not allow you to sell stock. Margin buying is allowed up to a limit that depends on your current wealth. That is, up to the limit shown on your screen, you can buy stock even when your cash position is zero or negative. If you borrow and the value of your stock position falls too much, then your wealth will hit zero. When that happens, you are bankrupt and can’t trade any more for the rest of that trading period.

Transaction Costs

When you click a button to buy or sell stock you incur transaction costs, which come out of your cash position. These costs are much higher at the highest buying and selling rates.

You can keep track using two charts which are located underneath the buy and sell buttons. The return chart shows how fast your wealth is growing without transaction cost. The transaction costs chart shows how fast the transactions costs are depleting your wealth. The difference between the return and transaction costs is the rate at which your wealth actually changes.



Fig. 15 Graphics window




Fig. 16 Density of managers window

Automated Agents

Before a trading period begins, the experimenter will announce whether you are trading solely with other human participants or whether you are trading with other humans as well as automated agents. Automated agents trade according to a fixed formula. Automated agents can also influence buying and selling pressure that changes the share price.

(For Sessions Where Human Subjects Use the Graphic Windows).

Graphics Window

The graphics window displays automated agents as small triangles, other human traders as round dots, and you as a dot with a specified color. The horizontal axis represents the ratio of stock position to wealth. For example, if you buy more shares then your dot will move to the right. The vertical axis represents the trader's wealth. If your wealth is increasing then your dot will move upward. Zero is indicated by the horizontal magenta strip in the center of the graphics window. The large white circle represents the average share holdings and wealth of all traders (Fig. 15).

Density of Managers

The Density of Managers chart is a histogram of the horizontal position of all traders, automated and human. A large clump of wealthy traders in the graphics window corresponds to a high point in this chart (Fig. 16).




Fig. 17 Landscape window

Landscape Plot

The Landscape chart shows the return rate (profit before transactions costs) each week for traders at every horizontal position (stock position relative to wealth). The graphics window shows where return is positive (green), negative (red), or very close to zero (yellow), but provides less detail than the Landscape chart (Fig. 17).

Summary

1. To earn as much money as possible during this experiment, try to accumulate as much wealth as you can by the end of each trading period.
2. At any time, your wealth is your cash position plus the value of your stock position. That value is the number of shares times the stock price.
3. You begin each trading period with an initial cash and stock position.
4. To purchase or sell stock at a specific rate, you click a button once. To change the rate, or to stop trading, you must click another button.
5. Every time you buy or sell you incur a transaction cost which depletes your wealth. You can have a negative cash position (borrowing at a specified interest rate) but not a negative share position (no short selling allowed).
6. Each trading period lasts a certain amount of real time (e.g., 10 min) and time advances on a weekly basis. Therefore, each trading period begins at week 0 and ends on week 1040. There will be several trading periods in each experiment.

FAQs

Q1. What does the return mean?

A1. Return-rate is the rate in which you are making or losing money. So if your wealth is \$1000 and your return rate is 10% then over one year or 52 weeks your wealth would have increased by $\$1000 * 10\%$ or \$100.

Q2. What does the transaction cost mean?

A2. It means you are paying to make a sell or trade. When you select a rate to buy or sell you will see a return chart and transaction cost chart. If the transaction cost

chart is above the return chart that means you are paying more in transaction costs than you are making in the stock market and thus your wealth is actually decreasing.

Q3. Will I really be paid cash at the end of the experiment?

A3. Yes. The cash payments are funded by the National Science Foundation and other research organizations. These instructions are all on the level. If we do anything deceptive, or don't pay you cash as described, then you can complain to the campus Institutional Review Board and we will be in serious trouble. Our interest is in seeing how people make decisions in market situations.

References

- Alfarano, S., Lux, T., & Wagner, F. (2005). Estimation of agent-based models: The case of an asymmetric herding model. *Computational Economics*, 26(1), 19–49.
- Brady, N. et al., (1988). The Brady Commission Report.
- Brock, W. A., & Hommes, C. H. (1997). A rational route to randomness. *Econometrica*, 65, 1059–1095.
- Brock, W. A., & Hommes, C. H. (1998). Heterogeneous beliefs and routes to chaos in a simple asset pricing model. *Journal of Economic Dynamics and Control*, 22(8–9), 1235–1274.
- Cason, T., & Friedman, D., (1997) Price formation in single call markets. *Econometrica* 65(2).
- Duffy, J. (2006). Agent-based models and human subject experiments. *Handbook of Computational Economics*, 2, 950–1101.
- Friedman, D., & Abraham, R. (2009). Bubbles and crashes: Escape dynamics in financial markets. *Journal of Economic Dynamics and Control*, 33(4), 922–937.
- Garman, M. B. (1976). Market microstructure. *Journal of Financial Economics*, 3(3), 257–275.
- Gode, D. K., & Sunder, S. (1993). Allocative efficiency of markets with zero-intelligence traders: Market as a partial substitute for individual rationality. *Journal of Political Economy*, 101, 119–137.
- Hoggatt, A. C. (1959). An experimental business game. *Behavioral Science*, 4, 192–203.
- Plat, C., (1995). *A mixing model with zero intelligence traders*. Dissertation UC Santa Cruz.
- Wyatt, E., (2010). *Regulators vow to find way to stop rapid dives*. *NY Times*. <http://www.nytimes.com/2010/05/11/business/11agencies.html?ref=business>.