Humans, Robots and Market Crashes: A Laboratory Study *

Todd Feldman† Daniel Friedman‡

October 7, 2008

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 otherwise they 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. Likewise, following losses, humans' choices shift towards faster selling. There are problems in properly identifying fundamentalist and trend-following strategies in our data.

Keywords: Financial markets, agent-based models, experimental economics. **JEL codes:** C63, C91, D53, G10

* **Preliminary Draft.** We are grateful to the National Science Foundation for support under grant SES-0436509, and to seminar audiences at UCSC and at the 2008 ESA conference at Caltech. We thank Ralph Abraham and Paul Viotti for their invaluable assistance. We retain sole responsibility for remaining idiosyncrasies and errors.

† Economics Department, University of California, Santa Cruz, CA, 95064. *tf61194@ucsc.edu*

‡ Economics Department, University of California, Santa Cruz, CA, 95064. *dan@ucsc.edu*

1 Introduction

We insert human traders into an agent-based simulation model of a financial market prone 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 interest, but as yet have no widely accepted theoretical explanation. In response, several agent-based models have been proposed. Brock & Hommes (1997, 1998) find bubble and crash dynamics in simulations when the majority of agents switch from a fundamentalist strategy to a trend-following strategy. Friedman & Abraham (2008) model agents who adjust their exposure to risk in response to a payoff gradient, and obtain bubbles and crashes arising from an endogenous market risk premium based on an exponential average of investors' recent losses.

Such simulation models use specifications and parameters that are difficult to verify from field data, but are amenable to testing in the laboratory. To illuminate the simulation models, we create a laboratory environment based on the Friedman & Abraham (2008) model, and examine how human traders respond to the payoff gradient and to losses. We also examine whether their actions are consistent with fundamentalist or trend-following strategies, or with switching.

Section 2 discusses some of the relevant literature. We mention earlier papers, some of them quite influential, that combined simulation with laboratory experiment. We also say a little more about agent-based models of financial markets. Section 3 summarizes the simulation model of Friedman & Abraham (2008), and section 4 describes its implementation as a laboratory experiment.

Section 5 presents the results. We find that human traders have little impact on bubbles and crashes in larger (30 trader) markets and when they are more experienced. In general, human traders earn lower profits overall than do the robots but earn higher profits in the most crash-intensive periods. Humans' buying and selling choices respond to the payoff gradient in a manner similar to the robot algorithm. In addition, humans respond to their own losses (and, to a lesser extent, to market wide losses) by selling. It turns out that there are problems in properly identifying fundamentalist and trend-following strategies in our data and perhaps also in field data.

Concluding remarks are collected in Section 6. Appendix A gathers details about implementations and presents supplementary data analysis. Appendix B is a copy of the instructions to laboratory subjects.

2 Earlier Literature

Economic studies combining human subjects and simulations have been rare but distinctive. One of the very first economic experiments sought to calibrate oligopoly simulations. Key parameters for 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 market microstructure, had similar motivations and technique. Conversely, Gode and Sunder (1993) used "zero-intelligence" robots in simulations intended to explain the rapid convergence to competitive equilibrium in laboratory markets with human traders. Plat (1995, 2005) inserted robot liquidity traders into an laboratory financial market. Cason and Friedman (1997), among others at around the same time, used robots following equilibrium strategies to train human subjects. Duffy (2006) points up general complementarities between simulations and laboratory experiments.

Our own study begins with the agent based model of financial markets reported in Friedman and Abraham (2008), as explained in the next section. The study also sheds some light on the quite different models of Brock and Hommes (1997, 1998), which consider the interplay of two basic strategies. In the fundamentalist strategy, traders take positions that will be profitable if stock prices tend to revert towards the fundamental value. In the trend-following strategy, traders take positions that will be profitable if deviations from fundamental value persist and grow. The agents in these models tend to switch their strategy when the alternative strategy has been more more profitable recently. Boswijk et al. (2007) report that the S&P 500 data support the existence of both strategies and the model of switching; some details are provided in Appendix A.

3 A Model of Bubbles and Crashes

Friedman and Abraham (2008) 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

$$P = V\bar{u}^{\alpha}, \tag{1}$$

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 α is a parameter that captures the sensitivity of price to buying pressure. Friedman and Abraham (2008) show that realized yield on the risky asset then is

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

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

$$\phi = R_i(u_i) = (R_1 - R_o + \epsilon_i)u_i - \frac{1}{2}c_2 u_i^2, \tag{3}$$

where ϵ_i is an idiosyncratic shock and c_2 is the current market-wide price of risk. The shock ϵ_i is an Ornstein-Uhlenbeck stochastic process, i.e., mean reverting to zero in continuous time. The 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 equation (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,

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

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,

$$\phi_u = R_1 - R_o + \epsilon_i - c_2 u. \tag{5}$$

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 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 (2008), 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 ϵ_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 $\dot{\bar{u}}$ terms in (2).

4 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 (see http://ccl.northwestern.edu/netlogo/ and the HubNet feature enables human participation. 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 Figure 1.

A typical session lasted 90 minutes 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 the price is high. Subjects are not told the price equation (1) nor the values of V, $\bar{u}$, or α. However, as shown in Figure 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 Figure 1. Before the trading period begins, each human subject is endowed with five hundred lab "dollars" and seventy 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 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 (or portfolio size) is $z_i = x_i + Py_i$, 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 $c = 1$, so trading at a fast rate incurs a transaction cost of $0.5^2 = 25\%$ of the market value of the transaction. A moderate rate then incurs a transaction cost of $0.25^2 = 6.25\%$, and a slow rate incurs a cost of only $0.125^2 = 1.6\%$.

Simulations are governed by a parameter vector specifying the discount rate, idiosyncratic shock volatility and so forth, as detailed in Friedman and Abraham (2008). 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 $R_o = 0.03$, $R_s = 0.06$, $g = 0.0$, $\sigma = 0.3$, $\tau = 0.7$, $\eta = 0.7$, $\beta = 2$, $\alpha = 2$, $\lambda = 1$, $d = 1$, and $rate = 1.3$; see Appendix A for all parameter definitions.

Figure 1: Basic User Interface

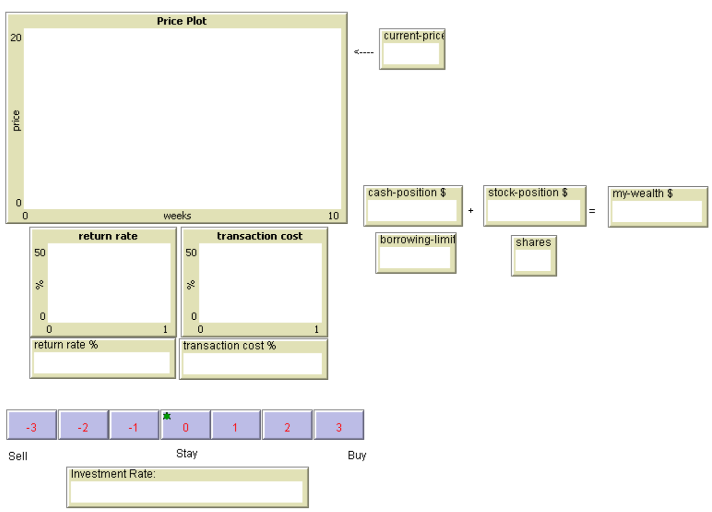

4.1 Treatments

We employ three treatment variables, as follows.

Figure 2: Enhanced User Interface with Graphics Window, Landscape, and Density Chart

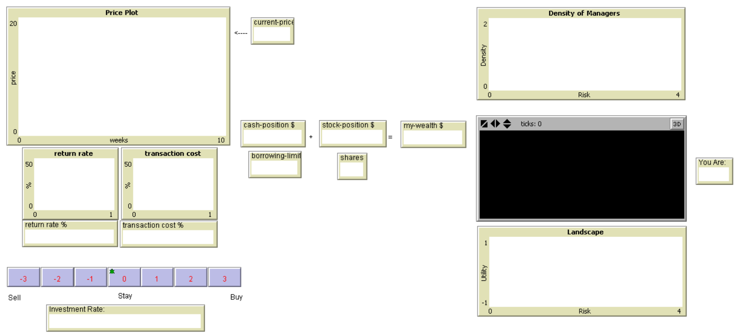

- Number of human and robot traders, NT. We examine five levels of NT: (0 humans, 30 robots), (1 human, 29 robots), (5 humans, 25 robots), and (5 humans, 5 robots) (0 humans, 10 robots). The first and last levels are run separately as controls. The other three levels are run as blocks in each session. 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 Figure 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.

4.2 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 ($c = 0.5$).

5 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 1040 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.

The rest of this section presents answers six empirical questions: whether humans change the distribution of crashes, outperform robots, herd, follow a payoff gradient, react to exponential average of their own and market wide losses, and switch between strategies based on past profits.

Figure 3: Asset Price in Period 7 of Session 7.

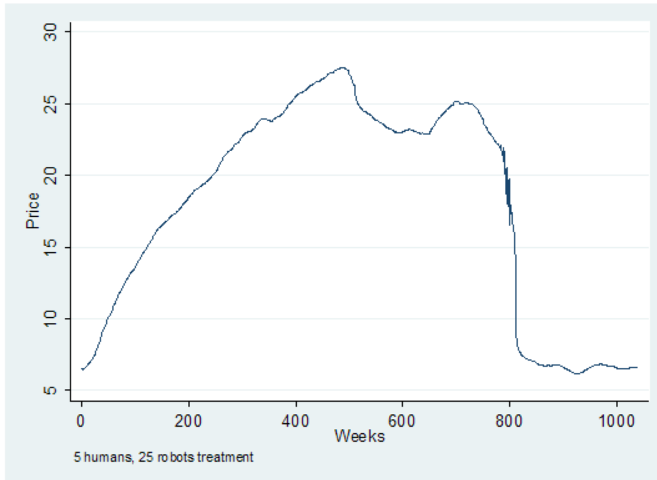

5.1 Do Humans Provoke Crashes?

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 Kolmogrov-Smirnov test.

Figure 4: Distribution of Crashes by Population

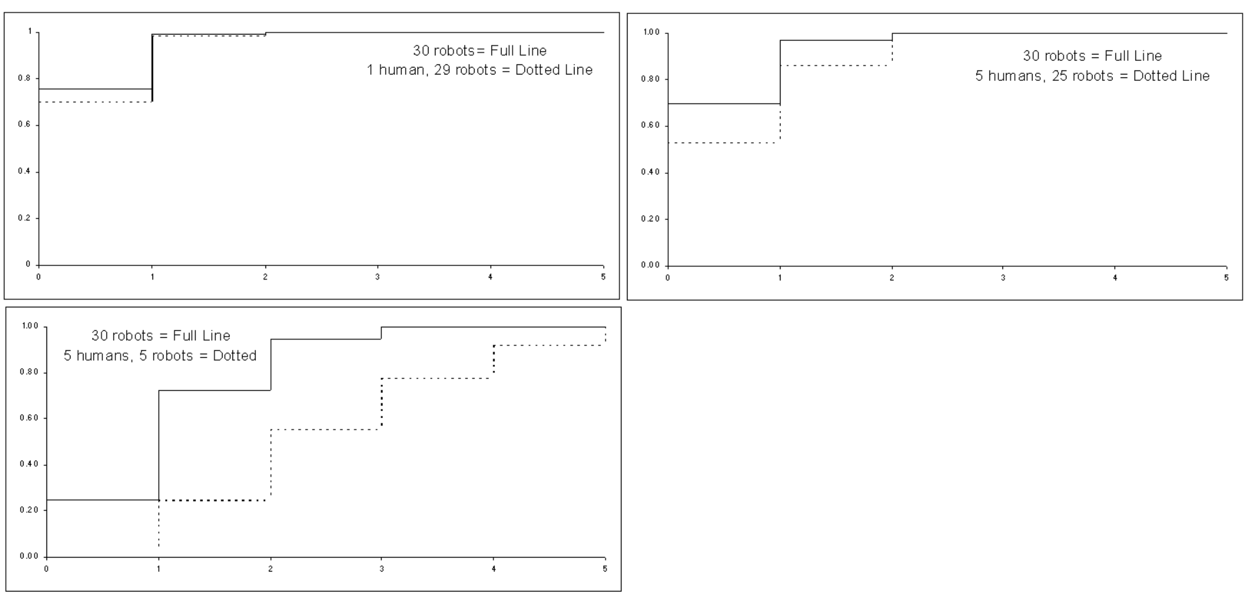

Table 1: Kolmogrov-Smirnov Test

Treatments	NOB	D	P-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

Note.NOBs reports the number of observations in each empirical distribution function, D is the maximum difference between the two distributions, and P-Value reports the significance of the standard 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) treatment with standard transaction costs, but find no time trend in the other treatments.

5.2 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 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, the (5 humans, 5 robots) treatment. The comparative advantage is statistically significant in the basic and enhanced interface standard transaction cost treatments.

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.

Figure 5: Wealth Comparison

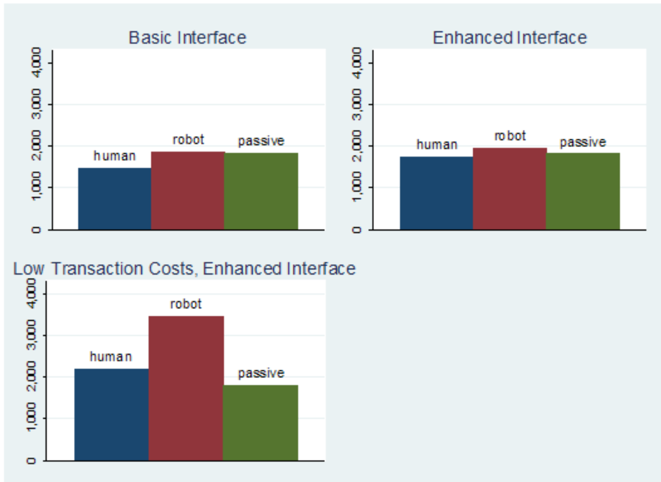

Figure 6: Wealth Comparison During High Crash Periods

5.3 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 that 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.

Figure 7: Frequency of Adjustment-Rates

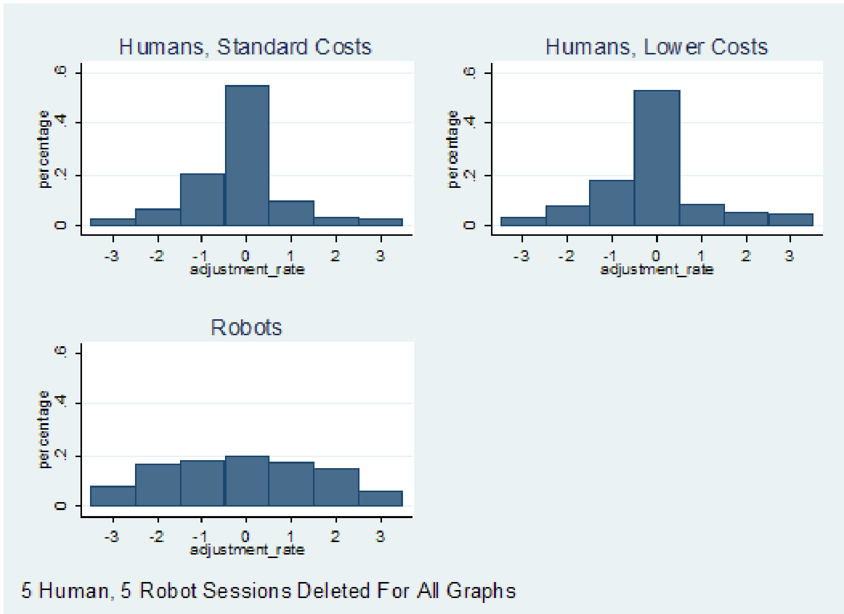

5.4 Herding?

Do human traders influence each others' allocation decisions? The graphics window in the enhanced user interface allows them to see what other human traders (and robots) are currently doing, and Figure 8 suggests that it may encourage humans to herd together. The figure shows 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.

Figure 9 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 t-tests confirm that the differences are insignificant (T=0.50) in the population 30 treatments and highly significant (T=2.30) in the population 10 treatments. One possible

Figure 8: Graphics Window Photos (Taken at Week 520)

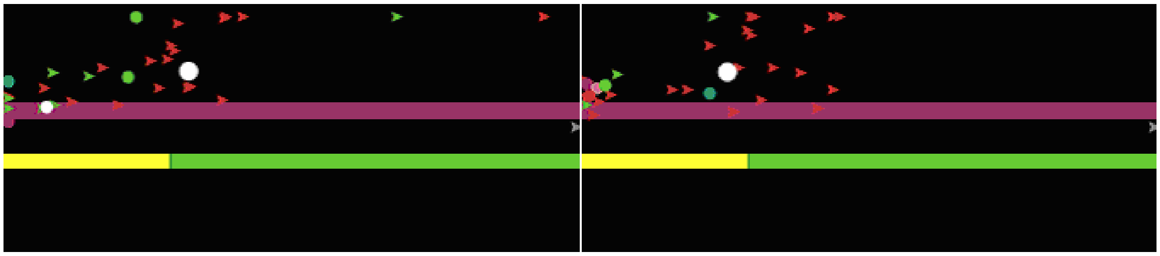

Basic Interface Treatment Enhanced Interface Treatment

interpretation is that in the more stressful environments, humans learn from each other to be contrarians.

Figure 9: Correlation of Adjustment-Rates

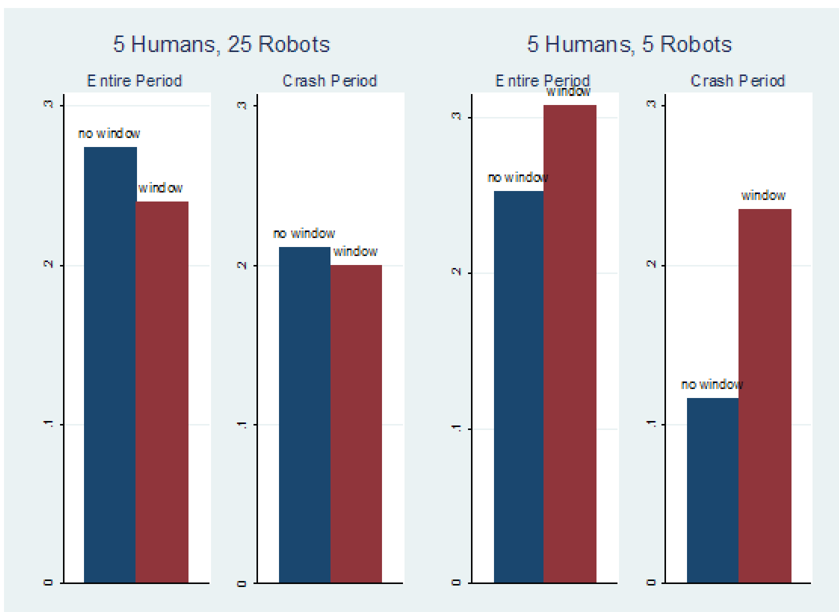

5.5 Do Humans Follow the Gradient?

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 traders' adjustment choice a_h on the gradient evaluated at the trader's current allocation u_h. We specify

$$a_h = \beta_0 + \beta_1 * \tilde{\phi_{uh}} + e, \tag{6}$$

where $\tilde{\phi_{uh}}$ for human traders is given by (5) truncated to lie between $\pm$ 0.50. The truncation reflects humans' bounded choice set. Of course, for human traders, ϵ 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 the dependent variable is the sorted gradient choice reported in the last panel of Figure frequencyallhumansrobots. The Table 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'.

Table 2: Gradient Adjustment Estimates by Population Treatment

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

Note. Random effects slope coefficient estimates (and standard errors) from equation (6). **significant at 1%.

According to Table 2 humans are not exactly following a gradient but are surprisingly close. Humans' coefficients are in every case are more than half as large as robots'.

5.6 Do Humans React to Recent Losses?

A second motivating question concerned the endogenous risk premium: do humans respond to an exponential average of their own losses, $\hat{L_i}$, and/or to an exponential average of market wide losses, c_2? The next regression investigates whether adjustment rates a_h chosen by humans can be explained by these variables. The regression includes as control variables the human trader's level of cash, shares, wealth, return, and an indicator variable for crash period.

The estimates on $\hat{L}$ and c_2 refer to how humans respond to losses in the baseline treatment, standard transaction costs and no ability to view other traders. Results from Table

8 indicate, as theory predicts, humans respond to an exponential average of their personal losses by selling. Unlike theory predicts, humans slightly buy as market wide losses increase. However, this result is misleading because regression results, seen in the appendix, indicate that humans at times respond to market wide losses by selling when personal losses are removed from the regression equation. The interaction terms explain how humans respond to personal and market wide losses across the different treatments. The magnitude of the personal loss estimates are partially offset in the enhanced interface treatment. For example, in the (5 humans, 25 robots) treatment, the estimate on $\hat{L} * IFE$ is -0.73 instead of -1.50 as in the base case. However, the magnitude of the personal losses are reinforced in the low transaction treatment. For example, in the (5 humans, 25 robots) treatment the estimate on $\hat{L} * LTC$ is -1.87 compared to the the base case estimate of -1.50. In addition, we find that humans respond more strongly to personal losses relative to market wide losses.

The Appendix also shows that alternative specifications of exponential losses sharpen these results.

Table 3: Random Effects Estimation Results, Dependent Variable: Human Adjustment-Rate

Variable	**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)**	0.42 ± (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)**
N	185,132	185,667	181,285
R^2	0.26	0.20	0.19
$\chi^2_{(8)}$		65,464	57,422

Note. 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 $\widetilde{\phi_{uh}}$, crash dummy $Crash$, enhanced interface dummy IFE, and low transactions cost dummy LTC. The number of observations N differs slightly across treatments due to bankruptcies. **significant at 1%

5.7 Trend Following and Fundamentalist Strategies

Recall that Brock and Hommes (1997, 1998) model switching between two basic strategies, trend-following and fundamentalist. Appendix A.2 below summarizes the empirical techniques of Boswijk et al. (2007, henceforth denoted BHM07) use to detect such behavior in S&P price data.

We now apply the same techniques to our laboratory and simulation data. Table 4 reports the results. For comparative purposes, the first column simply transcribes the S&P results from BHM07. The second column runs the same model on annual prices in our robot simulations using (30 robots) and (10 robots) treatments. Here one does not expect to see significant coefficients, because the robots are programmed as gradient followers, not as fundamentalists or trend followers or switchers. Surprisingly, the evidence for the two strategies, and perhaps for switching, seems as strong or stronger as in the S&P data. The last column reports model estimates for price data from our experiments combining humans and robots, using the (1 humans, 29 robots), (5 humans, 25 robots), and (5 humans, 5 robots) treatments. The strategy estimates are not much different from the robot data, but the switching coefficient is significantly larger.

Table 4: Estimates of the BHM07 Model

Variable	**BHM's Results**	**Robot Data**	**Human-Robot Data**
Trend Follower	$1.13 \pm 0.04^{**}$	$1.96 \pm 0.12^{**}$	$1.86 \pm 0.11^{**}$
Fundamentalist	$0.76 \pm 0.06\ ^{**}$	$0.75 \pm 0.03^{**}$	$0.88 \pm 0.01^{**}$
Switching	10.29 ± 6.94	$8.44 \pm 1.52^{**}$	$12.98 \pm 2.23^{**}$

Note- This tables reports estimates and standard errors in parenthesis. An estimate above one refers to a trend follower and an estimate below one refers to a fundamentalist. The number of observations in each sample is 132 in BHM's paper, 1440 using robot data, and 2160 using human-robot data. **significant at 1%

Results from Table 4 indicate that the estimates are similar across all three data sets. This could have two implications.

The results suggest that the techniques of BHM07 might not adequately identify the behavior they investigate. There is one encouraging note: the estimates for switching and for the fundamentalist strategy are somewhat larger in the experimental data than in the robot only data.

5.7.1 Correlation Coefficient & Counting

We look further into whether humans are using a fundamentalist strategy by calculating the correlation coefficient between the lag of deviation of price from fundamental value and adjustment-rate for humans and robots. We find the human's correlation coefficient is -0.31

and the robot's correlation coefficient is -0.11, indicating humans and robots on net are fundamentalists. In addition, the human's and robot's correlation coefficient are comparable indicating similar behavior between the two groups.

We also count the number of times humans and robots used a trend following strategy versus a fundamentalist strategy. Results from Table 6 show that humans act as fundamentalists

Table 5: Percentage Used Fundamentalist & Trend Following Strategy

Strategy	**Human**	**Robot**
Trend Follower	37%	47%
Fundamentalist	63%	53%

the majority of the time. In addition, the percentages between humans and robots are not very different. Therefore, from the above analysis we conclude humans are not using more than one strategy, and if anything they are using a fundamentalist strategy.

6 Conclusions

Ours may be the first study to use humans to better understand agent-based simulations and also to use the agents to better understand human behavior. We introduced human subjects as traders (or portfolio managers) into the bubble-and-crash prone simulations of Friedman and Abraham (2008), 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 a third 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 with personal losses. However, it is still consistent with the simulation model.

5. A fundamentalist vs trend-following switching model seems to have an identification problem in our data. It misidentifies some gradient-following robots as fundamentalists, others as trend-followers and switchers. On the other hand, it plausibly detects more switching and more fundamentalism when humans are present.

Our study is far from definitive, but it opens new territory for exploration. We hope it inspires other investigators to look for gradient following behavior and response to sentiment indexes based on recent losses in field and laboratory data. It may inspire the development of better techniques for identifying various sorts of trader strategies. More broadly, we hope it encourages a more active interplay between laboratory experiment with human subjects and agent based simulation models.

7 Bibliography

Arthur, W.B., Holland, J.H., LeBaron, B., Palmer, R., Taylor, P. (1997). Asset pricing under endogenous expectations in an artificial stock market. In: Arthur,W.B., Durlauf, S.N., Lane, D.A. (Eds.), The Economy as an Evolving Complex System II, Proceedings Volume XXVII, Santa Fe Institute Studies in the Science of Complexity. Addison-Wesley, Reading, MA, pp. 1544.

Abraham, Ralph, Dan Friedman, and Paul Viotti. Complex Dynamical Systems and the Social Sciences. http://www.ralph-abraham.org/articles/full-list.html, MS#123.

Axelrod, R. (1997). The Complexity of Cooperation: Agent-BasedModels of Competition and Collaboration. Princeton University Press, Princeton, NJ.

Axelrod, R., Tesfatsion, L. (2006). A guide for newcomers to agent-based modeling in the social sciences, this handbook.

Boswijk, H.P.,Hommes, C.H.,Manzan, S., (2007), Behavioral Heterogeneity in Stock Prices, Journal of Economic Dynamics and Control, special issue WEHIA workshop, Essex, 2005,

Cherkashin, D., J. D. Farmer, and S. Lloyd. A Simple Evolutionary Game with Feedback between Perception and Reality. Santa Fe Institute Working Paper 07-08-032, Santa Fe, NM, 2007.

Duffy, J., Unver, M.U. (2006). Asset price bubbles and crashes with near zero-intelligence traders. Economic Theory 27, 537563.

Duffy, J. (2006). Agent-Based Models and Human Subject Experiments. Handbook of Computational Economics, Volume 2, 950-1101.

Farmer, J. D. Market Force, Ecology and Evolution. Ind. & Corp. Change 11(5) (2002): 895-953.

Friedman, Daniel. (2005). Conspicuous Consumption Dynamics. UCSC working paper: http://www.vismath.org/research/landscapedyn/models/articles/

Friedman, Daniel. (2006). Bubbles and Crashes: Escape Dynamics in Financial Markets. UCSC working paper: http://www.vismath.org/research/landscapedyn/models/articles/

Friedman, Daniel and Yin-Wong Cheung. Speculative Attacks: A Laboratory Study in Continuous Time. UCSC working paper.

Friedman, Daniel and Joel Yellin, (1997). Evolving Landscapes for Population Games, University of California, Santa Cruz draft manuscript.

Garman, Mark 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, 119137.

Gode, D.K., Sunder, S. (1997a). What makes markets allocationally efficient?. Quarterly Journal of Economics 112, 603630.

Hoggatt, Austin C. 1959. An Experimental Business Game. Behavioral Science 4, 192-203.

Smith, V.L., Suchanek, G.L., Williams, A.W. (1988). Bubbles, crashes, and endogenous expectations in experimental spot asset markets. Econometrica 56, 11191151.

8 Appendix

8.1 Implementation Details

Table 6: Paramter Definitions

Paramter	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
α	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

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 calculate what their transaction cost would based on the translated gradient. After making those adjustments we then calculate step by step how wealth for each robot changes overtime until the period ends.

Integrating Robots & Humans The decision variables of both 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,

$$u_h = 1 - (\frac{cash_h}{wealth_h}), \tag{7}$$

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.[1] However, there are two subtleties that differentiate the portfolio size z of the 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.[2] 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,

$$t_h = c(adjustment\text{–}rate_h)^2, \quad c = constant, \tag{8}$$

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 fifty percent 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.

[1] See Bubbles & Crashes (2008) for an explanation of the inflow and outflow rate.

[2] The distribution is limited to the (u, z) rectangle $[0.2, 1.4] \times [0.4, 1.6]$, set via the sliders in NetLogo.

8.2 BHM's Empirical Model

A feature of the Brock & Hommes model is that it has been formulated in deviation from a benchmark fundamental. The dynamic asset pricing model they estimate is as follows,

$$Rd_t = n_t\phi_1 d_t - 1 + (1 - n_t)\phi_2 d_{t-} + \epsilon_t \tag{9}$$

where $R = \frac{1+r}{1+g}$, dt denotes the deviation of price from fundamental value, ϕ_1 and ϕ_2 denote the coefficients of two types of beliefs, n_t represents the fraction of investors that belong to the first type of traders and ϵ_t represents a disturbance term. If the value of ϕ is positive and smaller than one it suggests that investors expect the stock price to mean revert towards the fundamental value. If the belief parameter is greater than one, it implies the deviation of the stock price to grow over time at a constant speed. The fraction of investors is determined as,

$$n_t = \frac{1}{1 + exp(-\beta((\phi_1 - \phi_2)d_{t-3}(d_{t-1} - Rd_{t-2})))} \tag{10}$$

where β is the speed of switching strategies. The empirical model estimates the above two equations. The estimation results in Table 4 show that there are two significantly different regimes: one characterized by a coefficient less than one and the other by a coefficient greater than one. The qualitative interpretation of the regimes is: one group of fundamentalists believing that the stock price will mean revert towards the fundamental value and another group of trend followers believing that prices will persistently deviate from the fundamental valuation.

8.3 Robustness Checks

Frequency of Crashes Over Time

Figure 10: Mean Crashes per 20 Yrs by Trading Period

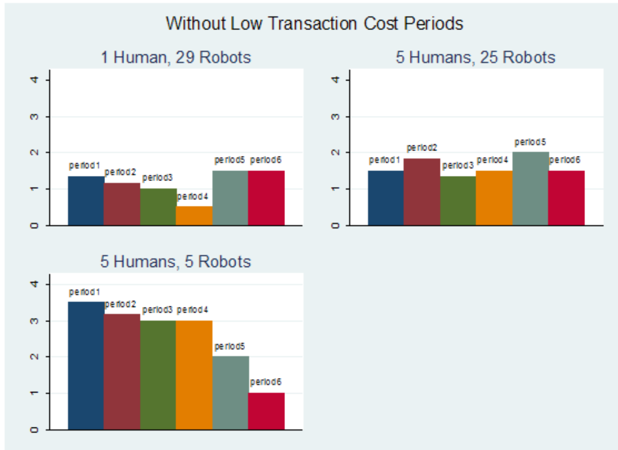

Frequency of Adjustment-Rates During Crashes

We find 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.

Figure 11: Frequency of Adjustment-Rates During Crashes

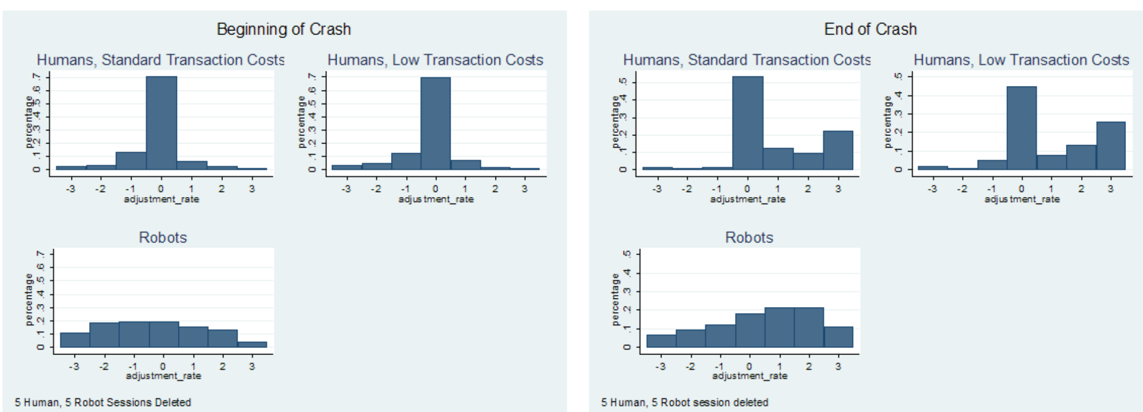

We also run each regression with an intercept term. The intercept estimate is statistically different from zero for the humans.

Human and Robot Loss Regressions

Table 7: 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)**	0.97 ± (0.02)**
$c2 * u_h$	-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 ± (.04)**	0.42 ± (0.01)**
$\hat{L} * LTC$	-0.76± (0.05)**	0.06 ± (0.04)**	-0.15 ± (0.01)**
c_2	0.10 ± (0.01)**	-0.08 ± (0.02)**	0.21 ± (0.01)
$c_2 * IFE$	-0.18± (0.02)**	0.15 ± (.02)**	-0.08 ± (0.01)**
$c_2 * LTC$	0.01± (0.02)	0.05 ± (.02)**	0.01 ± (0.01)
$Crash$	0.61 ± (0.03)**	0.06 ± (0.01)**	-0.07 ± (0.00)**
Human Nobs	185,132	185,667	181,285
R^2	0.07	0.01	0.08
$\chi^2_{(8)}$	22,331	7,253	21,933

Note-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%

These two tables runs the above loss regression using human and robot data. In each regression, we use the same explanatory variables.

Table 8: 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)**
c_2	0.39 ± (0.01)**	0.42 ± (0.01)**	0.01 ± (0.05)
$c_2 * IFE$	-0.01 ± (0.01)	-0.01± (0.01)	0.05 ± (0.01)
$c_2 * LTC$	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)**
N	5,418,360	934,200	186,840
R^2	0.26	0.25	0.16
$\chi^2_{(8)}$	301,284	286,344	35,371

Note-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%

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 above using different values of η to calculate $\hat{L}$. We run the following regressions by population treatment because we assume that different population treatments may affect how humans remember their realized losses.

Tables 9, 10, and 11 report regression results using standardized coefficients in order to compare estimates across different values of η. Results from Tables 9 and 10 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 some support for using a value of η equal to 0.70 in the simulation's baseline configuration. However, Table 11 indicates the lowest η, 0.1, is the correct memory parameter to characterize human memory in the hard hit, high 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.

Table 9: 1 Human, 29 Robots: Random Effects Using Different Etas

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 ± (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

onepc significant at 1%; standardized coefficients

Table 10: 5 Humans, 25 Robots: Random Effects Using Different Etas

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 ± (.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 ± (.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 ± (.089)*
$\chi^2_{(8)}$	62980.82	64520.72	65464.79	64898.95	62687.09

**significant at 1%; standardized coeffcients

For a robustness check we run the loss regression by individual and look for estimates that are

Table 11: 5 Humans, 5 Robots: Random Effects Using Different Etas

Variable	$\eta = 3$	$\eta = 1.3$	$\eta = 0.7$	$\eta = .3$	$\eta = .1$
$\hat{L}_{h,t}$	-0.16 $\pm$ (0.036)**	-0.26 $\pm$ (0.007)**	-0.32 $\pm$ (0.008)**	-0.20 $\pm$ (0.009)**	-0.27 $\pm$ (0.013)**
$\hat{L}_{h,t}\_h5\_r25$	0.12 $\pm$ (0.004)**	0.20 $\pm$ (0.009)**	0.20 $\pm$ (0.01)**	-0.19 $\pm$ (0.018)**	-0.07 $\pm$ (0.02)**
$\hat{L}_{h,t}\_h1\_r29$	-0.05 $\pm$ (0.004)**	-0.07 $\pm$ (0.008)**	-0.08 $\pm$ (0.01)**	0.02 $\pm$ (0.009)**	-0.08 $\pm$ (0.027)*
$\chi^2_{(8)}$	54159.50	55693.53	57422.40	63791.14	67095.58

**significant at 1%; standardized coefficients

outliers. Figure 12 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.

Figure 12: Lhat Regression by Individual

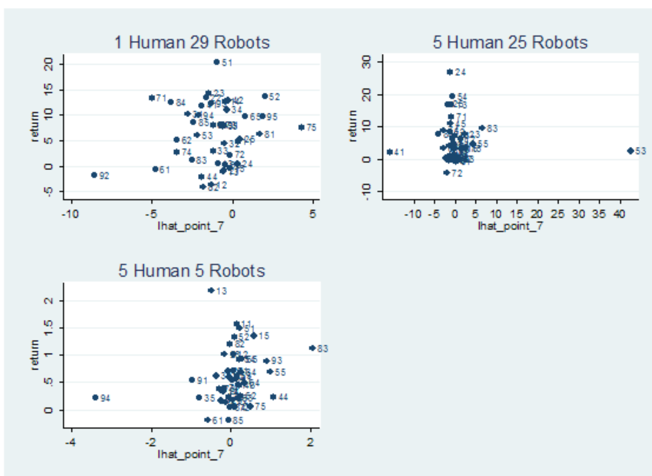

We also test how the enhanced interface and low transaction treatments affects experiment groups react to an exponential average of their losses. Again, there are nine experiment group, three for each treatment. Figure 13 in the Appendix shows that experiment groups do differ overall in how they react to an exponential average of their losses when able to view a graphics window. However, they progress in time simarily, and only change differently overtime in the later trading period where there are less observations. Figure 14 show that groups do not differ overall in how they react to an exponential average of their losses when facing lower transaction costs. In addition, they do not differ how they progress in time or change overtime. These results, confirm that the enhanced interface treatment does affect how humans respond to losses but the lower transaction treatment plays a smaller role if at all.

Figure 13: Random Effects by Experiment, Enhanced Interface Treatments

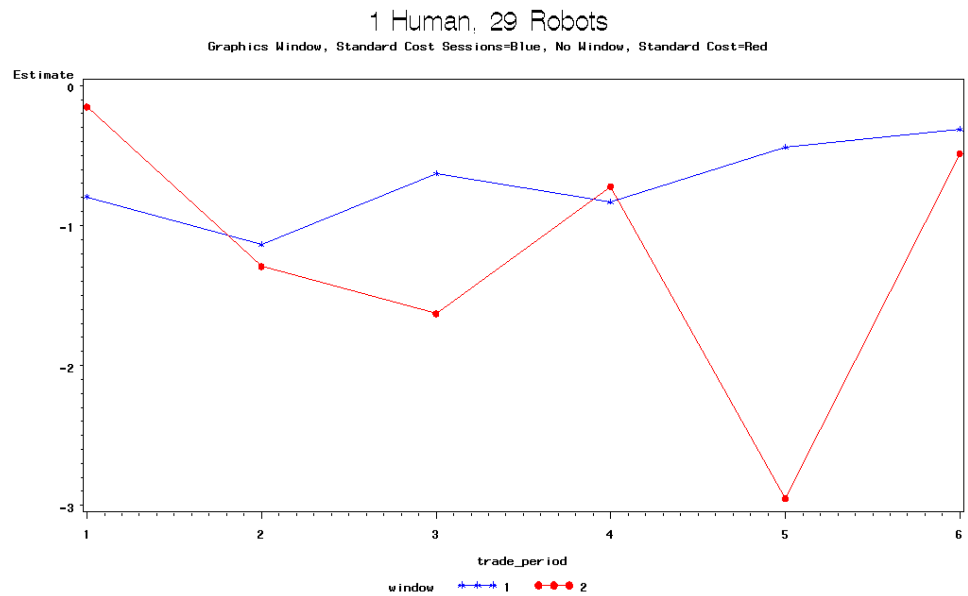

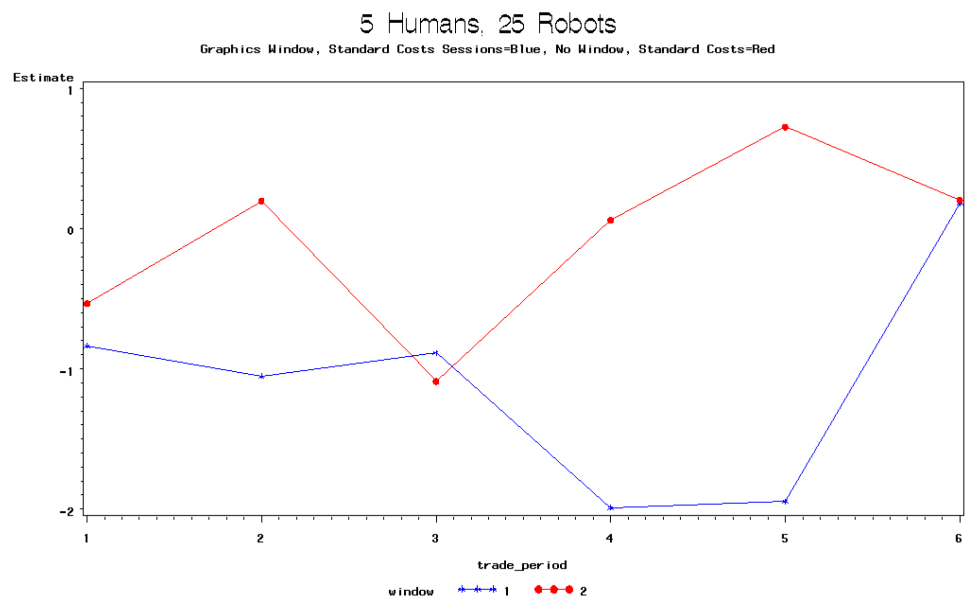

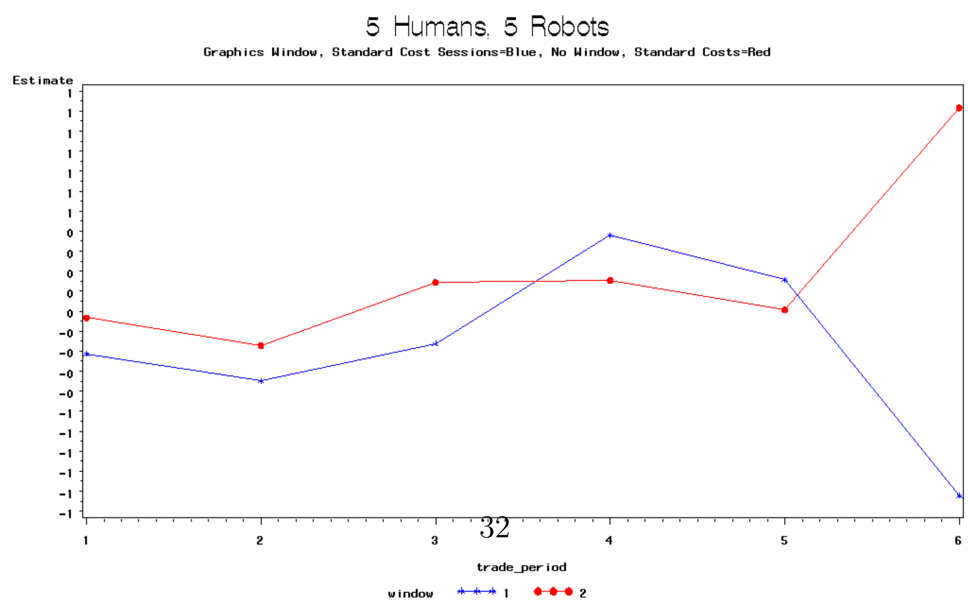

Figure 14: Random Effects by Experiment, Transaction Cost Treatments