Bubbles & Crashes: An Experimental Approach *

Todd Feldman† Daniel Friedman‡ Ralph Abraham§

May 1, 2008

Abstract

We test a bubbles and crash model (Friedman & Abraham, 2008) in a controlled laboratory environment. The experiment uses agent-based modeling to create a virtual financial market where human subjects act as stock market traders alongside automated robots. We use the experimental data to first test whether humans adjust their exposure to risk in response to a payoff gradient and to test second whether humans perceive risk by responding to an exponential average of their losses. We find that humans do not exactly follow a gradient but are very close. We also find that humans strongly respond to losses putting more weight on the most current losses. However, how they respond to losses depends on the frequency and predictability of crashes.

Keywords: Financial markets, escape dynamics, time varying risk premium, constant gain learning, agent based models, experimental economics.

*We thank Paul Viotti, Andy Sun, Matt Draper and Don Carlisle 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*

§Mathematics Department, University of California, Santa Cruz, CA 95064. *abraham@vismath.org*

1 Introduction

Despite their intrinsic interest, financial bubbles and crashes as yet have no widely accepted theoretical explanation. In response, Friedman & Abraham develop an out of equilibrium agent-based model focusing on portfolio managers who adjust their exposure to risk in response to a payoff gradient. Bubbles and crashes occur for a wide range of parameter configurations in extended models incorporating an endogenous market risk premium based on investors' historical losses and exponential averaging. Even though the simulations confirm bubbles and crashes, simulation models are more valuable when they work in tandem with empirical studies and/or laboratory experiments with human subjects. Therefore, we devise an experiment where human subjects interact with automated robots to test the assumptions driving Friedman & Abraham's model.

2 Experimental Design

We conducted an experiment 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. HubNet is a technology that lets you use NetLogo to run participatory simulations. In a participatory simulation, a group can take part in enacting the behavior of a system as each human controls a part of the system by using an individual device. The LEEPs' laboratory has 14 computers each linked to Hubnet via a server where subjects interact in a virtual market as seen in Figure 1.

A typical experiment lasted 90 minutes and involved 5 inexperienced human subjects recruited by email from a campus-wide pool of undergraduate volunteers. Humans silently read the instructions and then listened to an oral summary by the conductor. After a couple of practice rounds, they played about 12 periods. Humans subjects are paid based on the average of their wealth achieved at the end of each trading period which is redeemed at a couple of cents of real money, typically between $15 and $25.

During the trading period each human acts as a trader in a stock market alongside other humans and automated robots. Their objective is to maximize their wealth by buying and selling shares of a single stock at price P,

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

where V is the fundamental value, $\bar{u}$ is the mean distribution of allocation choices among robots and humans, and α is a positive parameter that measures sensitivity to buying pressure. Humans do not know the price equation (1) nor the values of V, $\bar{u}$, or α. However, as shown in Figure 1, they can see the current price and price plot. We do tell them that price is determined by the growth rate, interest rate, and buying and selling pressure. More specifically, we tell them the growth rate is zero, the interest rate is three percent, and that no one individual can move the stock price, but collectively, net buying pressure increases the price and net selling pressure decreases the price.

Each trading period consists of 20 "years," where the computer screen updates the trades and wealth on a weekly basis as shown in Figure 1. Before each trading period begins humans are endowed with five hundred dollars and seventy shares of stock. Their wealth at any point in time is equal to their cash plus their number of shares owned times the stock price. Human cash and wealth change based on several factors. First, humans earn interest on cash savings as well as pay interest if they borrow. Margin buying is allowed up to a limit that depends on their current wealth. Second, a buy reduces their cash position by the amount purchased times the stock price plus a transaction cost. A sell increases their cash position by the amount sold times the current stock price minus a transaction cost. Third, human shares grow based on the growth rate.[1] In addition, humans can go bankrupt. If a human goes bankrupt, they are banned from trading and incur a loss of $500 for the period. However, they are allowed to resume trading in the next period.

Humans view events on the monitor screen and respond by clicking one of seven buttons called adjustment-rates. Of the seven adjustment-rates, 3 accumulates shares at a very fast rate, 2 accumulates at a medium rate, 1 accumulates shares at a slow rate, 0 refrains from trading, -1 sells at a slow rate, -2 sells at a medium rate, and -3 sells at a fast rate. A message box reminds them which button is active. In addition, there exists monitors to view their

[1]Shares do not grow in the base case where the growth rate is zero.

holdings of cash, shares, wealth, transaction cost, and rate of return.

2.1 Treatments

We use two types of treatments. The first type varies the number of traders who participate in a market. The second type involves information about the other participants:

- Number of humans and robots. As the population size declines price volatility and the frequency of crashes increase. The three different population treatments include (1 human and 29 robots), (5 humans and 25 robots), and (5 humans and 5 robots). Each experiment runs four blocks. We run three blocks of each treatment where the fourth block repeats the treatment run in the first block. For every experiment we rotate block order so the final data set contains the same number of observations per treatment.[2] In addition, all blocks are known to everyone.

- Information. Depending on the experiment, humans are able to see (or not see) a graphics window, density of traders plot, and a landscape plot as shown in Figure 2 (or not, as in Figure 1). The graphics window displays automated robots as small triangles and human traders as round dots where humans can identify themselves by a specified color. The graphics window allows humans to see the ratio of stock position to wealth of every human and robot. The density of managers chart is a histogram of the horizontal position of all traders, automated and human. 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 baseline configuration values in the simulated model are $R0 = dR = 0.03$, $g = 0.0$, $\sigma = 0.2$, $\tau = 0.7$, $\eta = 0.7$, $\beta = 2$, $\alpha = 2$, $\lambda = 1$, $d = 1$, $rate = 1.3$, and $c = 1$. We use the same parameter values for the experiments except we increase σ to 0.3 to induce sufficient variability so that humans cannot predict future price movements. Humans are not specifically told parameter values. The meanings of these parameters are described on our websites.

[2]For example, the first experiment ran (1 human and 29 robots), (5 human, 5 robots), (5 human and 25 robots), and (1 human, 29 robots). The second experiment ran (5 human, 5 robots), (5 human and 25 robots), (1 human and 29 robots), and (5 human, 5 robots). And the third experiment ran (5 human and 25 robots), (1 human and 29 robots), (5 human, 5 robots), and (5 human and 25 robots).

2.2 Integrating Robots & Humans

The humans' shares and wealth are translated into an appropriate u and z. Humans' risk allocation, u_j, equals one minus the ratio of their cash to wealth and the portfolio size, z_j, changes based on their gross return, inflow rate, and outflow rate[3],

$$u_j = 1 - (\frac{cash_j}{wealth_j}), \tag{2}$$

$$\dot{z}_j = [R_o + (R_1 - R_o - t_j - \delta \hat{L}_j + \rho z_o e^{\lambda \hat{R}_j}]z_j. \tag{3}$$

where j refer to humans and t refers to the transaction cost. The human's initial z_j is equal to 1 and subsequently changes based on equation (3). The robots receive their initial risk allocation, u_i, and portfolio size, z_i, randomly via a uniform distribution in the (u, z) rectangle $[0.2, 1.4] \times [0.4, 1.6]$, set via the sliders. The u and z possible range is between 0 and 4. The distinction between how portfolio size changes for a robot and a human is that robots receive an idiosyncratic shock and do not pay transaction costs where as humans pay transaction costs and do not receive an idiosyncratic shock. The transaction cost is determined as,

$$t_j = c(adjustment\text{–}rate_j)^2, \quad c = constant. \tag{4}$$

As in real markets humans face transaction costs where larger orders incur larger trading costs. The constant, c, is set to 1 such that 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%. Humans are able to see a monitor that tracks their transaction costs. For every trade transaction costs reduce humans' cash savings. We use transaction costs for humans in order to analyze whether humans are sensitive to market frictions or whether they thrash between buying and selling at fast rates. 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. We then scale the adjustment rates up in order to

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

accurately affect human cash and share holdings.

Figure 1: Human Interface

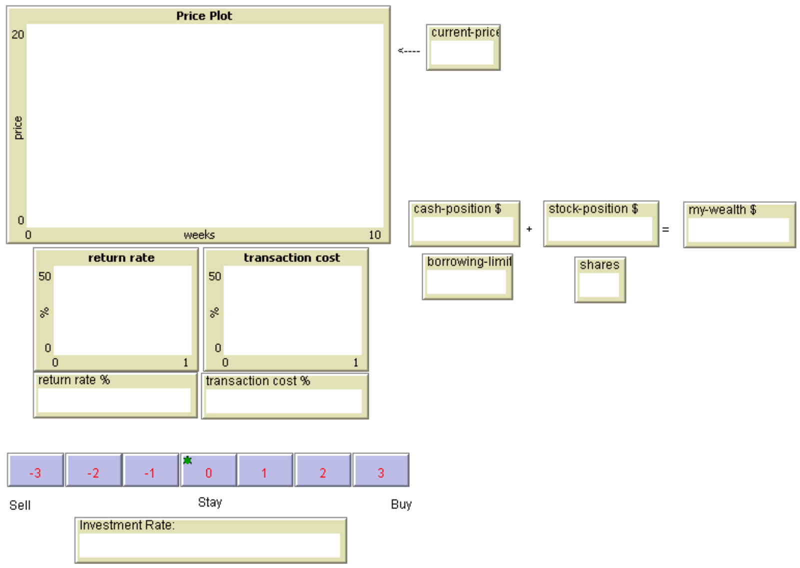

Figure 2: Human Interface With Graphics Window, Landscape, and Density Chart

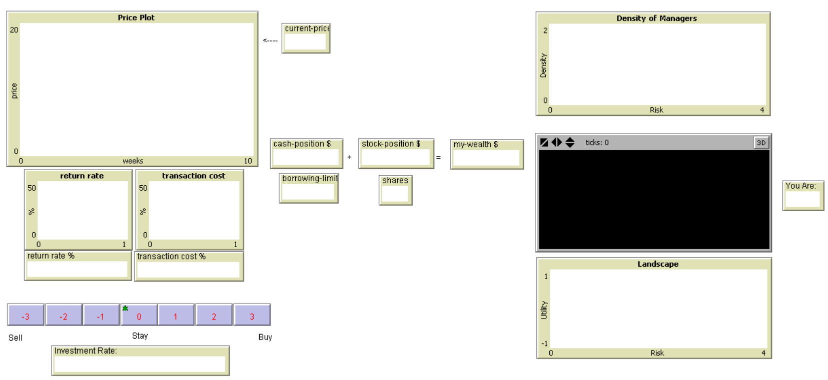

3 Results

To investigate these assumptions from Friedman & Abraham (2008) , and to check their robustness, we analyze data from all five experiments. We define a crash as as a decline in price P of at least 50% from its highest point within the last half year.

3.1 Do Humans React to an Exponential Average of their Losses?

In order to investigate whether humans react to an exponential average of their losses we run the following regression,

$$adjustment\text{–}rate_t^h = \beta_0 + \beta_1 * cash_t^h + \beta_2 * shares_t^h + \beta_3 * wealth_t^h + \beta_4 * return_t + \beta_5 * h5\text{–}r25$$
$$+\beta_6 * h1\text{–}r29 + \beta_7 * \hat{L_t^h} + beta_8 * \hat{L_t^h}\text{–}h5\text{–}r25 + \beta_9 * \hat{L_t^h} - .15cm\ \text{–}h1\text{–}r29 + \beta_{10} * crash\text{–}period + \epsilon. \quad (5)$$

where h refers to human data and the dependent and explanatory variables have the following meanings,

- The dependent variable, $adjustment\text{–}rate_t^h$, is the average rate at which humans choose to trade at time t.
- $Cash_t^h$ represents the level of average cash holdings at time t.
- $Shares_t^h$ represent the average number of shares at time t.
- $Wealth_t^h$ represents the level of average wealth at time t.
- $Return_t$ is the log first difference in price.
- $\hat{L_t^h}$ is the average of humans' exponential average of losses at time t.
- The intercept represents the base treatment, (5 human, 5 robot).
- *h5-r25* is an indicator variable that assigns a 1 to the (5 human, 25 robot) treatment and 0 otherwise.
- *h1-r29* is an indicator variable that assigns a 1 to the (1 human, 29 robot) treatment and 0 otherwise.
- *Crash-period* is an indicator variable that assigns a 1 to the time period of a crash and 0 otherwise.[4]

[4] The crash period is determined as the point near the top of a bubble to where crash detection ends, where crash detection is a decline in price P of at least 50% from its highest point within the last half year.

The dependent and explanatory variables are determined using an equally weighted average at each week. [5] The $\hat{L_t^h}$ is determined by setting η equal to 1.3 because regression results show that the model specification where $\hat{L_t^h}$ is calculated using an η between 1.3 and 3 produces the highest R-square. Lastly, the interaction variables, $\hat{L_t^h}$*-h5-r25* and $\hat{L_t^h}$*-h1-r29*, tell us how humans respond to an exponential average of losses relative to the $\hat{L_t^h}$, (5 human, 5 robot) baseline treatment.

Table 1: Human OLS Regression: All Experiments

Parameter	Estimate	Standard Error	Pr > $\lvert t \rvert$
Intercept	-0.085	0.0035	<0.0001**
$Cash_t^h$	0.0001	0.0000012	<0.0001**
$Shares_t^h$	0.0022	0.000022	<0.0001**
$Wealth_t^h$	-0.0001	0.00001	<0.0001**
$Return_t$	1.014	0.027	<0.0001**
h5-r25	0.056	0.003	<0.0001**
h1-r29	0.057	0.002	<0.0001**
$\hat{L_t^h}$	0.188	0.028	<0.0001**
$\hat{L_t^h}$*-h5-r25*	-1.368	.128	<0.0001*
$\hat{L_t^h}$*-h1-r29*	-1.002	.034	<0.0001*
Crash-period	-0.026	0.008	0.0026**

∗ significant at 5%; ∗∗ significant at 1%

The results in Table 1 indicate humans do respond to losses. However, how they respond to losses depends on the treatment. Humans responded to losses by buying in the (5 human, 5 robot) treatment and selling in the (5 human, 25 robot) and (1 human, 29 robot) treatments. The theory says that as losses accumulate humans should sell. Results from the (5 human, 25 robot) and (1 human, 29 robot) treatments confirm the theory but the (5 human, 5 robot) treatment does not. One reason the (5 human, 5 robot) treatment does not is due to the number and predictability of crashes. After the first period in the (5 human, 5 robot) block, humans realized that a crash was inevitable and therefore waited for a crash in order to accumulate shares at low prices. Lastly, the crash-period estimate reveals humans sold

[5] The (1 human, 29 robot) treatment, obviously does not average the human data since each human participates in its own virtual market, therefore those observations stand alone.

modestly during crashes.

3.1.1 Do Humans Follow a Gradient?

Figure 3: Frequency of Adjustment-Rates

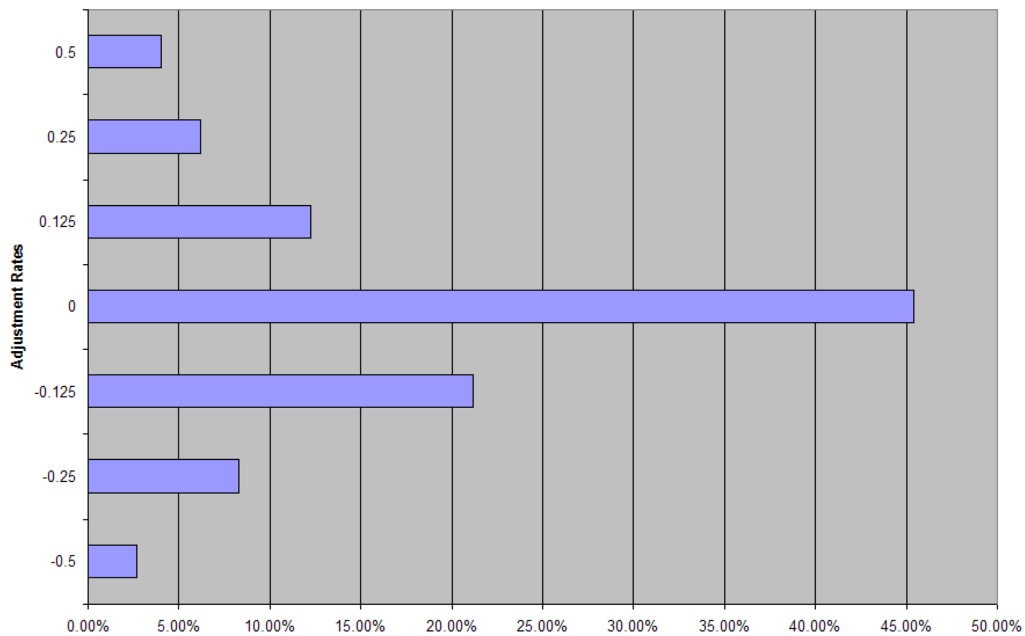

Figure 3 shows how frequently humans choose one of the seven adjustment rates. The distribution of choices is relatively symmetric with humans choosing to hold 45% of the time. This provides evidence that humans are sensitive and aware to market frictions. Humans do not jump back and forth between buying at a fast rate and then selling at a fast rate which confirms gradient dynamic behavior versus adaptive dynamic behavior.

In order to test whether humans follow a gradient similar to robots we assume humans see the same gradient as do robots, and regress their choices, $adjustment\text{–}rate^h$, on the gradient evaluated at the humans' current u_j, called $gradient^h$. Theoretically, if humans are exactly following a gradient then the $gradient^h$ estimate should equal 1.00. However, comparing the $gradient^h$ estimate to 1.00 is not appropriate in our study. By design the $gradient^h$ estimate will be less than 1.00 because humans can only choose seven different adjustment-rates and not a continuous set of adjustment rates. In order to find a more appropriate comparison estimate, we run the same regression using robot data where the explanatory variable is the robots' actual chosen gradient, $gradient^r$, and the dependent variable is the robots' actual gradient translated into one of the seven adjustment-rates humans face. We translate the robot's gradient by using midpoints to construct seven ranges that correspond to the seven adjustment-rates. We then assign an adjustment-rate to each of the robots' gradient depending on which of the seven ranges their gradient falls into.[6] Results from Table 2 report the appropriate comparison estimate ,$gradient^h$, is .69 and not 1.00. Therefore, the closer the estimated coefficient using human data is to the coefficient using robot data, 0.69, the more evidence that humans follow a gradient.

Table 2: Human's Adjustment Rate vs. Human's Gradient

Parameter	Estimate	Standard Error	Pr > $\|t\|$	R^{sq}
$gradient^h$	0.32	0.0035	<0.0001∗∗	.10
$gradient^r$	0.69	0.0016	<0.0001∗∗	.82

∗ significant at 5%; ∗∗ significant at 1%

According to Table 2 humans are not exactly following a gradient but are very close. In addition, the ability to view the graphics window plays a role. We achieve a $gradient^h$ estimate of 0.47 using only data from the graphics window experiments.

We also ran an additional regression to the one in equation (5) using robot data in order to compare estimates between the human and robot regressions. The regression results indicate the signs and level of significance for all estimates are the same. Robots and humans only differ in the magnitude of the estimates. For example, humans actually respond stronger to losses than do robots and robots sell ten times more aggressively that humans during

[6]For example, if a robot chose a gradient of 0.1 then that choice would be defined as a slow buy, 0.125, since it is between 0.0625 and 0.1875.

crashes.

4 Conclusions

We conduct one of the first studies to integrate agent-based modeling and experimental economics. The experiment consists of a virtual financial market that includes automated robots who follow a gradient (but are distinct in that each receives an idiosyncratic shock) and humans subjects. We use the experimental data to test two important questions: do humans react to market frictions by following a gradient and do humans perceive risk by reacting to an exponential average of their losses? From our analysis we can conclude several results. First, humans do respond to losses by generally selling. However, when crashes are more frequent and predictable, humans respond to losses by buying. Overall, since bubbles and crashes are not predictable the analysis of experimental data provides evidence that an exponential average of losses can be used as a way to measure the perception of risk. In addition, humans tend to weight the most current losses heavily. The model specifications using a high η to calculate $\hat{L}$ performed better than the ones using a low η. Second, humans do not exactly follow a gradient as compared to robots but are very close to following one. In addition, humans follow a gradient more closely when allowed to view all market participants in the graphics window. Lastly, experiments where humans are able to view other market participants in the graphics window tend to herd around each other and hardly ever disperse very far from the group. It is interesting how humans do not follow the robots, who drive the majority of the price dynamics in two of three population treatments.

There is still more work to be done. We would like to use the experimental data to test the assumptions of other agent-based financial models in order to determine which agent-based model fits the experimental data better. In addition, we would like to ask whether humans are a stabilizing or destabilizing force. Moreover, we would like to run more experiments with different constants on the transaction costs to see how human behavior changes as we reduce transaction costs.

5 Bibliography

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.

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.

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

Feldman, Todd, Dan Friedman, and Ralph Abraham. Bubbles & Crashes: An Ex- perimental Approach. Future volume.

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.