Bubbles and Crashes: a Cyborg Approach

Ralph Abraham [*] Todd Feldman [†] Daniel Friedman [‡]

June 17, 2008

Dedicated to K. Vela Velupillai on his 60th Birthday

Abstract

In joint work since 2004 we have created a family of agent-based models for financial markets in which bubbles and crashes occur in imitation of real markets. The evolution of behavioral rules in these models has shed light on some possible mechanisms used by human account managers or traders. Our programming environment, NetLogo, has proved ideal for this work, and also offers a feature, HubNet, capable of extending simulations to include human as well as robot traders. Recently we have used this feature to test a bubbles and crash model 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: Bubbles, crashes, agent-based models, NetLogo, financial markets, escape dynamics, experimental economics.

[*]Mathematics Department, University of California, Santa Cruz, CA 95064. *rha@ucsc.edu*
[†]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

Our goal in this contribution is to introduce the basic assumptions and features of our market models, and to give some early results of one of our "cyborg" experiments – that is, experiments involving human as well as robotic agents. But we will begin with a brief history of the project, written by Ralph.

In 1968, after moving from the math department of Princeton University to that of the University of California at Santa Cruz, I met Dan, then a grad student. After Dan's Ph.D. and some early positions, he became professor of economics at UCLA. During the 1980s I visited frequently at UCLA, and we used to meet for lunch at the faculty club. While I had no actual involvement in mathematical economics, I nevertheless kept up on the news through Dan. Then in 1985, Dan moved to UC Santa Cruz, and we continued meeting for the occasional lunch at the faculty club.

Meanwhile, chaos theory was heating up as a new style of applied math, and the economics community was becoming curious. Richard Goodwin, as assistant professor at Harvard, fellow of Peterhouse College (Cambridge) and professor in Siena, had led a long-term project on non-equlibrium economics, nonlinear dynamics, and so on. He became an early adopter and harbinger of chaos theory. His 1988 lectures in Siena appeared as a book in 1990, spreading his enthusiasm.

All this led to my invitation to give the Jacob Marshak Lecture at UCLA in January, 1987. After my talk on "Nonlinear Systems, Complex Dynamics, and the Social Sciences there followed a lively Q&A, during which there were several questions from an Indian gentleman revealing a deep knowledge and understanding of dynamical systems theory. At the end, I went up to him to inquire, "Who are you?" – and thus met Vela Velupillai.

Shortly thereafter, I received an invitation to a Workshop on Mathematical Economics at the Certosa di Pontignano, Siena (May, 1991). This was the occasion of my meeting the wonderful Richard Goodwin, having a ride in his pet car, seeing Vela again, and also meeting Leonello Punzo, both former students of Goodwin. In addition, I met a group of mathematical economists reporting exciting research in bifurcations of iterated mappings of the plane, especially Laura Gardini of Urbino, with whom I did joint work in the 1990s.

After 2000, Dan began telling me of his work in evolutionary game theory, and an opportunity arose early in 2004 to jointly apply for a grant from the National Science

Foundation. The NSF program involved was aimed at new mathematical methods for the social and behavioral sciences. After reading a special issue of *Nonlinear Dynamics, Psychology, and Life Sciences* on agent-based modeling, we saw a way to apply agent-based modeling to extend evolutionary game theory. In our grant application we wrote:

> Our proposed research will extend the class of models called evolutionary games by allowing the set of strategies (actions) of each player (trader, agent) to be a continuous space, rather than just a finite set. This continues a line of study begun in joint work of Dan Friedman and Joel Yellin in 1997. The central concept of this work is the adaptive landscape.

This grant proposal was funded, and since mid 2004 we have used NetLogo – an agent-based modeling software system – to create a sequence of financial market models. Next, we will explain the basic concepts as they have evolved to date. Later, we will describe the cyborg experiment and its results.

2 The basic model: math

Our project website, http://www.vismath.org/research/landscapedyn/, presents several models. Here we describe the simplest one, Market Model 9.0. These concepts are basic to all of our market models.

We envision a number of money market managers (typically 20 to 100 in our simulations) trading in a financial market with two kinds of assets, riskless (safe) and risky. Each manager has a choice from a continuum of strategies characterized by a nonnegative real number, u. This is her *risk parameter*, and defines the division of her portfolio between the two kinds of assets. The minimum value, $u = 0$, indicates no risk (all assets are riskless), the value, $u = 1$, indicates all risk (all assets are risky), and $u > 1$ indicates leveraged investment (borrowing on the safe asset).

Further, each manager has a portfolio of total worth, z, which we normally assume to be between zero and four, with $z = 1$ indicating a typical starting value. With each step, each portfolio's worth is adjusted according to its risk parameter. The safe portion u earns at rate R_0, which we have fixed at $R_0 = 0.03$, while the risky portion $(1 - u)$ earns at rate R_1, with typically $R_1 \geq R_0 \geq 0$. The manager's gross annual return is thus,

$$R_G = (1 - u)R_0 + uR_1 \tag{1}$$

Financial math leads us to pose,

$$R_1 = R_s/\bar{u}^2 + 2\dot{\bar{u}}/\bar{u} \tag{2}$$

where $R_s = R_0 + R_d$ and $R_d = 0.03$, or $R_s = 0.6$; $\bar{u}$ is the mean value of u choices for all managers; and $\dot{\bar{u}}$ is the time rate of change of $\bar{u}$. Full details may be found in (Friedman and Abraham, 2006).

Also, we assume that the gross return is decreased by a risk cost,

$$c(u) = c_2 u^2/2$$

where $c_2 = 0.02$. (In the research version of this basic model, Model 8.0, the constant c_2 may be varied by a slider.) Then the net return is,

$$R(u) = u(R_1 - R_0) - c_2 u^2/2 \tag{3}$$

Combining (1) and (2) we obtain the payoff function,

$$\phi(u, F) = u(R_s/\bar{u}^2 + 2\dot{\bar{u}}/\bar{u} - R_0) - c_2 u^2/2 \tag{4}$$

where $F(u)$ denotes the dependence of net payoff on the distribution of u choices of all managers.

The simulation proceeds in steps of discrete time intervals of size "stepsize", which the operator may choose as days, weeks, and so on. With each step, each manager's worth, z, is adjusted (depending on the stepsize) according to the net annual return $R(u)$ according to her current choice of risk, u. Additionally, her strategy choice, u, is adjusted according to the assumption of landscape dynamics, a gradient rule. That is, we assume that each manager is hill-climbing up the gradient of the payoff function (3) which depends on the current strategy choices of all managers (and their changes) through $\bar{u}$ and $\dot{\bar{u}}$ in (3).

3 The basic model: NetLogo

We now explain the graphical user interface of our simplest NetLogo model, Market 9.0, shown in Figure 1. The "population" slider (default setting, 30) determines the number of managers for the simulation. The "center" slider (default 20%) determines the mean u for the initial distribution of managers. The "setup" button creates the chosen number of managers with chosen mean u, and with random values of (u, z)

within a medium-sized rectangle in the upper half of the black graphics window. The "frequency" drop down menu (default 52, or weekly steps) determines the number of updates per year. The "go" button begins the simulation, which continues until the "go" button is pushed once more. The small triangles in the upper half of the graphics window indicate the managers, each positioned according to its (u, z) coordinates, so they are seen to move smoothly about as the simulation progresses.

The other features of the interface shown in Figure 1 are three plots and three monitors, that collectively show the position of managers, the landscape function, the market price as a ticker-tape and as current value, the total elapsed time in years, and the current value of net risky yield, R_1.

It has been proven (Friedman and Abraham, 2006) that this model always converges to a heap of managers all in one spot, and indeed, that is we what we observe as the simulation progresses. To obtain bubbles and crashes we need a more sophisticated model, such as Market Model 9.1. It implements two innovations: *surprise* (stochastic variations in payoff) and the c_2-*dynamic* (varying the c_2 coefficient in the gradient rule in response to losses, as explained below).

All of of our models are posted on our website with documentation. The NetLogo models posted there function as applets, that is, you may run the model within your web browser. In addition, the NetLogo models may be downloaded and run in the NetLogo programming environment, which may be freely downloaded from the NetLogo website, http://ccl.northwestern.edu/netlogo. We encourage you to try out the applets.

4 The advanced models

In the course of our project, we made a succession of extensions to the basic model, in search of dynamical features underlying the bubble and crash behavior of real financial markets. Our more sophisticated models have provided many insights into market forces contributing to bubbles and crashes, as reported in our articles published on our website. Three successive extensions, called Model 8.1, 8.2, and 8.3, extend the research version of the basic model, Model 8.0. In parallel, we prepared simplified models of two of these, Models 9.0 and 9.1.

The first of these extensions was successful in exhibiting bubble and crash behavior, and most of our research (reported in the papers mentioned in our bibliography below) has been done with this extension, Model 8.1. The chief dynamical feature of this extension, the c_2-dynamic, has the coefficient c_2 in the the risk cost (see equation 2 above) controlled by an algorithm, rather than by a slider. Unlike the basic model, here we have endogenous perturbations affecting each manager's payoff separately, that we call *surprise*. Our model determines surprise by an Ornstein-Uhlenbeck process. Due to the occasional negative surprise, the managers accrue losses, from which we calculate (for each manager independently) a weighted sum, $\hat{L}$, with higher weights for recent losses, and declining weights for older losses. Our algorithm for the c_2-dynamic makes use of all the individual $\hat{L}$ values, combined in a global, z-weighted mean, L_m. Recall that z is a variable (for each manager) measuring the current worth of that manager's portfolio. The rule to update c_2 is $c_2 = \beta L_m$, where β as a constant.

The user interface for Model 9.1 is shown in Figure 2. Note there are several additional sliders, one of which is "beta", which sets the constant β. The others are described in the User Manual for Model 8.1.

5 The cyborg experiment using Hubnet

The arrival in the 1980s of agent-based modeling in general, and NetLogo in particular, has stimulated a new wave of simulation research in economics, and more generally in the social and behavioral sciences. And it is in this context that we have situated the work performed under our recent NSF grant. However, during our work with NetLogo we discovered that it has various unique features that extend beyond the spectrum of other agent-based modeling systems. One of these unique extras is the HubNet system. This provides NetLogo client interfaces, so that a local net of computers may share control of the graphical user interface of a simulation as it runs. Originally developed for classroom use, we have found it useful for experiments involving human subjects interacting with a market of robot managers, and in other experiments as well. We feel that this work advances the programs initiated by Richard Goodwin and his students into new levels, and that many future agent-based simulations and experiments will follow.

Despite their intrinsic interest, financial bubbles and crashes as yet have no widely accepted theoretical explanation. In response, we developed an out-of-equilibrium agent-based model focusing on portfolio managers who adjust their exposure to risk in response to a payoff gradient, as described above. Bubbles and crashes occur for a wide range of parameter configurations in our advanced 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 devised an experiment where human subjects interact with automated robots to test the assumptions driving our NetLogo models.

6 Experimental design

We conducted an experiment at the University of California at Santa Cruz Learning and Experimental Economic Projects (LEEPS) lab using the Hubnet feature of NetLogo. In a participatory simulation, a group of human subjects can take part in enacting the behavior of a system as each human controls a part of the system by using an individual interface, the HubNet client. The LEEPS laboratory has 14 computers each linked to Hubnet via a server where subjects interact in a virtual market as seen in Figure 3.

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,

$$P = V\bar{u}^{\alpha}, \tag{5}$$

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. See equations (1) and (2) above. Humans do not know the price equation (1) nor the values of V, $\bar{u}$, or α. However, as shown in Figure 3, 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 3. 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 holdings of cash, shares, wealth, transaction cost, and rate of return.

6.1 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

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

[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

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 4 (or not, as in Figure 3). 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 sessions 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.

6.2 Integrating robots and 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],

$$u_j = 1 - (\frac{cash_j}{wealth_j}), \tag{6}$$

$$\dot{z}_j = [R_o + (R_1 - R_o - t_j - \delta\hat{L}_j + \rho z_o e^{\lambda\hat{R}_j}]z_j. \tag{7}$$

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

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).

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

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,

$$t_j = c(adjustment\text{–}rate_j)^2, \quad c = constant. \tag{8}$$

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 accurately affect human cash and share holdings.

7 Results

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

7.1 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,

$$adjustment\text{–}rate_{j,t} = \\ \beta_0 + \beta_1 * cash_{j,t} + \beta_2 * shares_{j,t} + \beta_3 * wealth_{j,t} + \beta_4 * return_t +$$

$$\beta_5 * h5\text{–}r25 + \beta_6 * h1\text{–}r29 + \beta_7 * \hat{L_{j,t}} + beta_8 * \hat{L_{j,t}}\text{–}h5\text{–}r25 +$$
$$\beta_9 * \hat{L_{j,t}} - .15cm\ \text{–}h1\text{–}r29 + \beta_{10} * crash\text{–}period + \epsilon$$

where the dependent and explanatory variables have the following meanings,

- The dependent variable, $adjustment-rate_{j,t}$, is the trading rate of human j at time t.
- $Cash_{j,t}$ represents the level of cash holdings of human j at time t.
- $Shares_{j,t}$ represent the number of shares of human j at time t.
- $Wealth_{j,t}$ represents the level of wealth of human j at time t.
- $Return_t$ is the log first difference in price.
- *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.
- $\hat{L_{j,t}}$ is the humans' exponential average of losses of human j at time t.
- *Crash-period* is an indicator variable that assigns a 1 to the time period of a crash and 0 otherwise.
- The intercept represents the base treatment, (5 human, 5 robot).
- The $\hat{L_{j,t}}$ is determined by setting η equal to 0.7. The interaction variables, $\hat{L_{j,t}}$-*h5-r25* and $\hat{L_{j,t}}$-*h1-r29* tells us how humans respond to an exponential average of losses relative to the $\hat{L_{j,t}}$, (5 human, 5 robot) baseline treatment.

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 selling less in the (5 human, 5 robot) treatment and selling more in the (5 human, 25 robot) and (1 human, 29 robot) treatments. The theory says that as losses accumulate humans should sell. Results from all treatments confirm the theory. One reason why humans respond less to losses in the the (5 human, 5 robot) treatment is possibly 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 bought slighlty during crashes.

Table 1: Human OLS Regression: All Sessions

Parameter	Estimate	Standard Error	Pr > $\|t\|$
$Intercept$	-0.008	0.0009	<0.0001$**$
$Cash_{j,t}$	0.00008	0.0000012	<0.0001$**$
$Shares_{j,t}$	0.0015	0.000022	<0.0001$**$
$Wealth_{j,t}$	-0.0001	0.00001	<0.0001$**$
$Return_t$	0.625	0.012	<0.0001$**$
h5-r25	0.033	0.0007	<0.0001$**$
h1-r29	0.054	0.0007	<0.0001$**$
$\hat{L}_{j,t}$	-0.047	0.004	<0.0001$**$
$\hat{L}_{j,t}$-*h5-r25*	-0.633	0.016	<0.0001$*$
$\hat{L}_{j,t}$-*h1-r29*	-1.137	0.021	<0.0001$*$
Crash-period	0.013	0.001	<0.0001$*$

$*$ significant at 5%; $**$ significant at 1%

7.1.1 Do Humans Follow a Gradient?

Figure 5 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, $adjustment\text{–}rate_j$, on the gradient evaluated at the humans' current u_j, called $gradient_j^h$. Theoretically, if humans are exactly following a gradient then the $gradient_j^h$ estimate should equal 1.00. However, comparing the $gradient_j^h$ estimate to 1.00 is not appropriate in our study. By design the $gradient_j^h$ 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, $gradient_i^r$, 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.[4] Results from Table 2 report the appropriate comparison estimate ,$gradient_j^h$, is 0.87 and not 1.00. Therefore, the closer the estimated coefficient using human data is to the coefficient using robot data, 0.87, 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}
$gradient_j^h$	0.32	0.0017	<0.0001**	.11
$gradient_i^r$	0.87	0.0004	<0.0001**	.91

* significant at 5%; ** significant at 1%

According to Table 2 humans are not exactly following a gradient but are very close.

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 crashes.

8 Conclusion

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

[4]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.

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 selling. However, when crashes are more frequent and predictable, humans respond less to losses. 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. 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, sessions 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 sessions with different constants on the transaction costs to see how human behavior changes as we reduce transaction costs.

9 Acknowledgements

We are grateful to the National Science Foundation for support under grant SES-0436509. Pablo Viotti, our first graduate research assistant, has been a great asset in all aspects of the project, including the NetLogo programming. He was instrumental in the development of our human subject software based on the HubNet feature of NetLogo, which has been marvelously adapted by our graduate student, Todd Feldman, to a test of our model assumptions with human managers. Also, we have been ably assisted by Andy Sun, Matt Draper and Don Carlisle.

10 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.

Elliot, Euel and L. Douglas Kiel, Agent-based modeling in the social and behavioral sciences. *Nonlinear Dynamics, Psychology, and Life Sciences* 8:2 (April 2004), pp. 121-130.

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, and Ralph Abraham. (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:
http://www.vismath.org/research/landscapedyn/models/articles/

Goodwin, Richard M. *Chaotic Economic Dynamics.* Oxford: Clarendon Press, 1990.

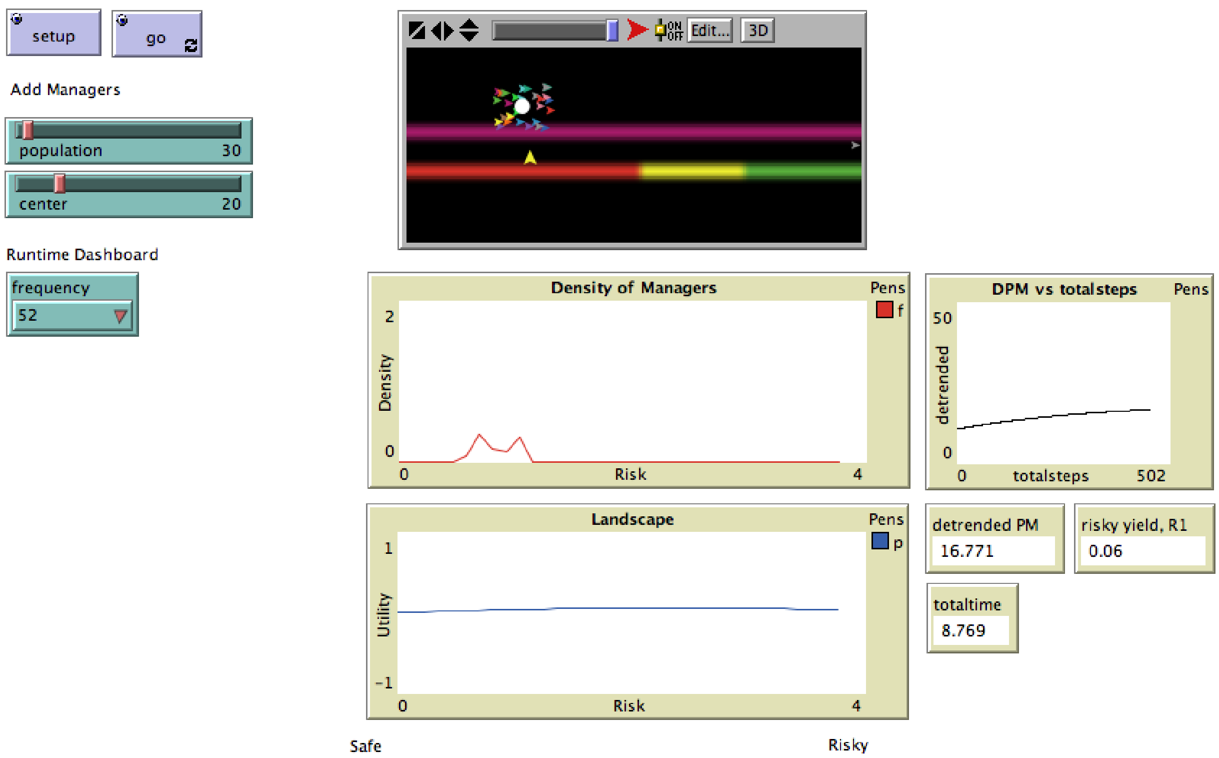

Figure 1: Interface of NetLogo Market Model 9.0.

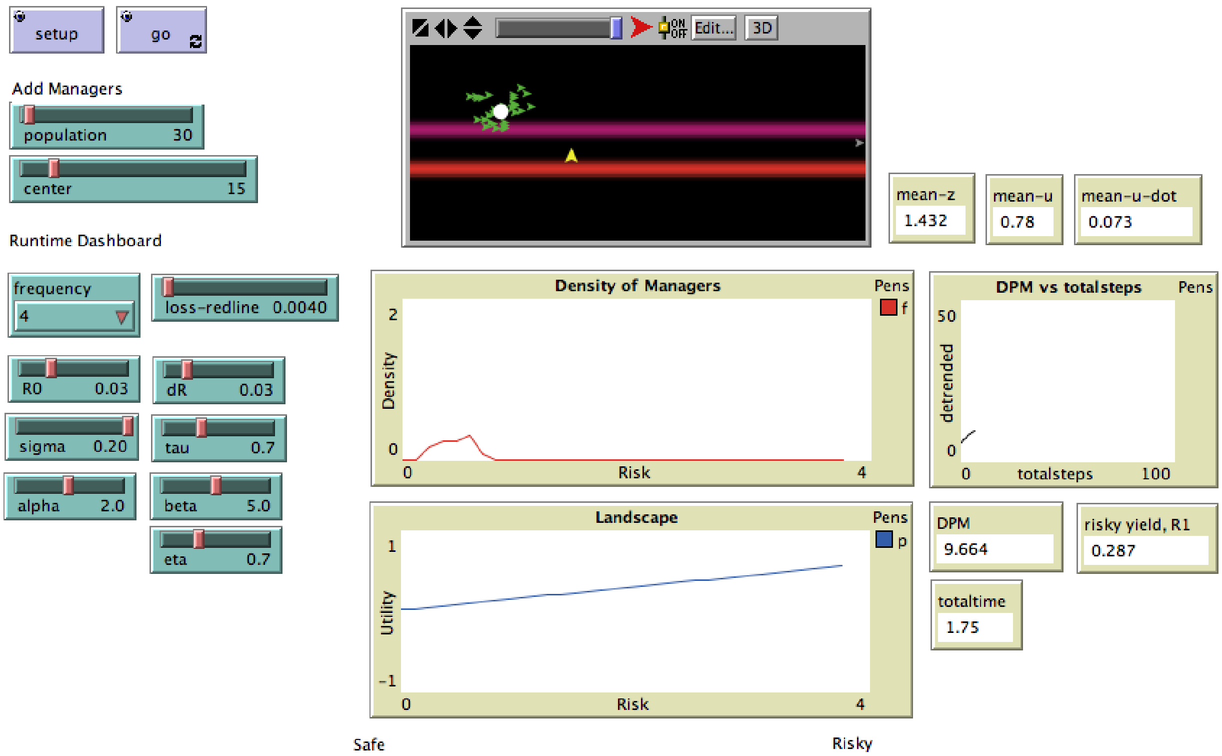

Figure 2: Interface of NetLogo Market Model 9.1.

Figure 3: Human Interface

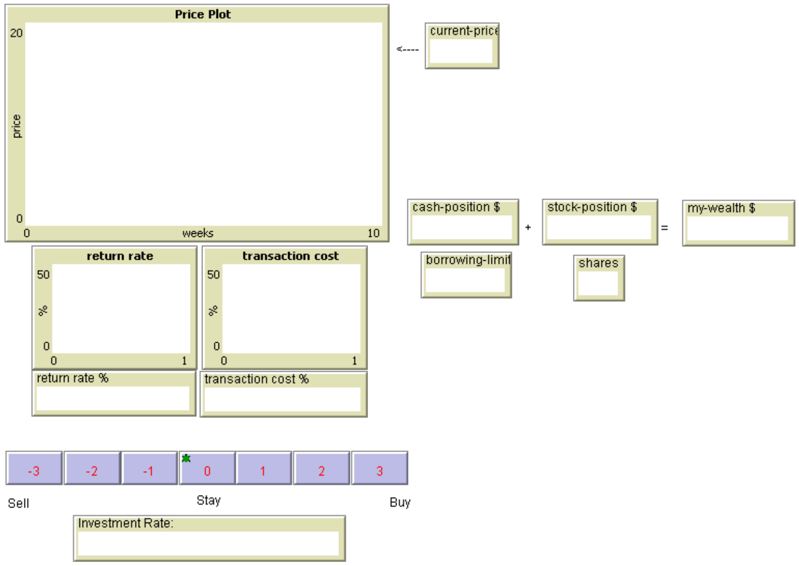

Figure 4: Human Interface With Graphics Window, Landscape, and Density Chart

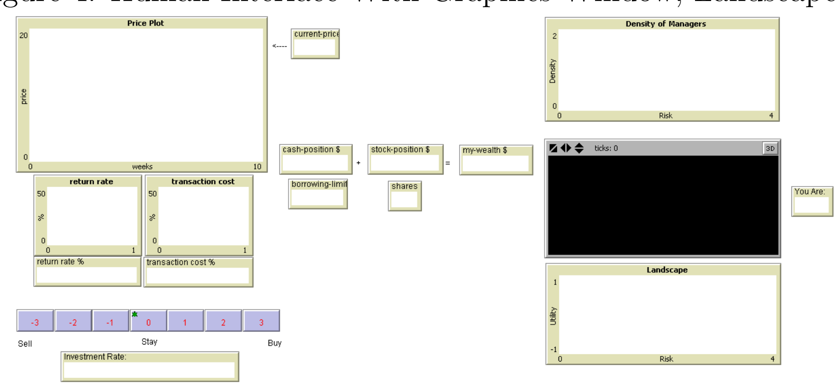

Figure 5: Frequency of Adjustment-Rates