

Asset Price Formation and Behavioral Biases

Todd Feldman^a and Gabriele Lepori^b

(ACCEPTED VERSION of a paper published in the Review of Behavioral Finance)

Abstract

We examine the debate on whether psychology affects asset prices using agent-based modelling. We set up three simulation regimes where the first regime contains fundamental investors who invest based on the mean variance framework. The second regime includes purely irrational investors who invest based on behavioral biases. The third regime combines the two types of investors. We test whether the return properties from regime three converge to that of regime one or two. Results suggest that the type of irrationality affects return properties in different ways. Irrational investors who are introspective in their irrationality, only examining their performance and deficiencies, do not have much of a systematic effect on stock returns when combined with rational investors. However, irrational investors that aggregate information in an irrational manner have a systematic effect when combined with rational investors.

Keywords: Financial markets, Behavioral finance, Agent-based finance

JEL Codes: C22, G10, G12

^a Finance Department, San Francisco State University, San Francisco, CA, 94132, USA. tfeldman@sfsu.edu

^b Keele Management School, Keele University, Staffordshire, ST5 5BG, United Kingdom. g.lepori@keele.ac.uk

1. Introduction

The purpose of this paper is to test whether psychology plays a role in the formation of stock prices by exploiting new techniques developed from the agent-based finance literature. More specifically, we ask if different types of irrationality affect market price dynamics in different ways. To do so we create a virtual stock market where rational and different types of irrational agents (investors) coexist. All investors use a variant of the mean-variance framework where irrational investors price risk based on behavioral biases. The net interaction between all investors influences the stock price dynamics. The motivation for the paper is based on the ongoing debate between Efficient Market Hypothesis (EMH) and behavioral theorists. EMH theorists argue that a market which is competitive and reasonably liquid aggregates the information potentially known by millions of diverse investors. This aggregation process occurs such that prices reflect all available information no matter whether irrational investors exist. Irrational investors do not impact the aggregation process for various reasons but two reasons stand out. First, poor trading strategies leave the irrational investor with no wealth leaving only rational investors. Second, the trading of irrational investors may be self-canceling, which would leave the determination of prices in the hands of the rational investors.

Behavioralists argue that the existence of stock market anomalies, such as excessive volatility and stock market bubbles and crashes, provide evidence against rationality. They argue that irrational investors do matter because these anomalies would never occur otherwise. Several empirical papers have provided evidence that investors do exhibit behavioral biases, (See Barber and Odean (2000)). However, behavioralists have still not provided sufficient evidence that these behavioral biases found in the literature actually impact asset prices.

We provide insight into this debate by simulating a virtual framework using three distinct regimes. The first regime includes purely rational investors and the second regime includes purely irrational investors. We explore the return properties of both simulated regimes to understand how differently they impact asset prices. Thereafter, we create a third regime where both rational and irrational investors are combined into one market. Irrational investors should become insolvent and become bankrupt if the EMH group is correct. Rational investors would then dominate the market and the return properties would reflect regime one's properties overtime. Irrational investors should not become insolvent or their trades self-canceling if the behavioralists are correct.

We create the simulated regimes using the pricing process developed by Friedman & Abraham (2009) and modeling of investors with various biases by Feldman (2011). The price process begins when each investor makes an allocation decision between a risky asset such as a S&P 500 portfolio and a safe asset such as cash. All investors make this decision by examining the trade-off between risk and return. The primary difference between a rational and irrational investor is how risk is priced. Each period all investors choose an allocation between a stock portfolio and cash. At that point, the computer program averages each investor's allocation choice across all investors weighting by their portfolio size. Therefore, investors with greater wealth have a greater impact on the mean allocation. As the mean allocation changes, the stock price changes. For example, more demand for the stock portfolio across investors increases the mean allocation resulting in the stock price increasing. As the stock price changes, the new information feeds back into every investor's objective function, resulting in a new allocation choice. The mean allocation is updated thereby changing prices again. This feedback process between allocation choice and price continues until the simulation ends.

This paper extends Friedman & Abraham as well as Feldman's work by combining rational and different types of irrational investors in one market where investors can become insolvent and

get kicked out the market. The primary extension is defining irrational investors into two subgroups, introspective and aggregating. The first type is irrationally introspective in that the source of the irrationality can be tracked back to their own performance. They exhibit well known biases such as overconfidence, loss aversion, and regret. For example, an irrationally introspective investor behavior may be greatly affected if he or she experiences losses. However, they are not affected if they witness others losing money. The second type is irrationally aggregating in that they are affected by the decisions and outcomes of other investors. For example, an irrationally aggregating investor may be affected if they witness other investors losing money even if they are not. That experience will create an incentive to sell their risky asset to prevent them from losing money as well.

A secondary extension includes modeling a new irrational investor, overconfident investors. Overconfidence is one of the most written about biases in the behavioral finance literature. Hirshleifer & Luo (2001) find that overconfident investors may outperform rational investors. We model an overconfident investor via the investor's alpha where the investor believes they can routinely outperform the broad stock market. Another secondary extension includes the introduction of insolvency. Insolvency leaves the possibility for one group to dominate the market. In past agent-based papers, such as Friedman & Abraham (2009) and Brock & Hommes (1998), each investor could not go bankrupt so that no one investor could dominate the market. By allowing bankruptcy we can thereby test whether rational investors dominate in the end as EMH would tell us.

Results suggest EMH theory holds depending on the type of irrationality that dominates in the marketplace. The irrationally aggregating investors are not driven out of the market and still impact asset prices in the long-run when combined with rational investors. A market where irrationally aggregating and rational investors invest together possesses negative skewness, positive excess kurtosis, and greater volatility than the rational regime. On the other hand, asset prices are

not significantly affected when irrationally introspective investors are combined with rational investors. Over the long-run the return properties exhibit low volatility, a slightly skewed normal distribution, and thin tails, just as in the rational regime case. Even so, irrationally introspective together with rational investors tend to keep the stock price lower than its fundamental value, unlike in the purely rational regime where the stock price trades around fundamental value. The intuition is that the poor performance and noise created the introspective investors drives prices down. This distinction between what we call irrationally aggregating and irrationally introspective investors is important to understanding how irrational behavior drives asset prices. For example, how does the view that housing prices always increase prevail in the marketplace? Based on our result, this outcome occurs when people view the losses and gains of other people and slowly become convinced that housing prices will never fall. They get caught into a way of thinking slowly overtime by viewing other investors in an irrational manner. Introspective irrationality is where investors suffer from regret of their own decisions or they are overconfident about their own abilities. In this sense this type of irrationality does not affect other people and therefore it is not as pervasive.

The results have implications for the asset pricing literature as well as the agent-based finance literature. First, we find that irrational behavior does affect asset prices when investors aggregate irrationally but not when they are introspective. Irrationally introspective investors are more akin to noise traders. Therefore, the implication is that it may be reasonable to ignore the impact of noise traders when modeling financial markets. Second, our result provides evidence to support the modeling of irrationally aggregating investors in other agent-based models as opposed to noise trader investors such as irrationally introspective investors.

2. Literature Review

Researchers in finance have explored various ways to test the limits of market efficiency. Fama (1965, 1970) finds support to the conclusion that the stock market is efficient in the sense that stock prices adjust very rapidly to new information based on stock split announcements. However, the behavioral finance literature has grown documenting various behavioral biases that investors exhibit.¹ Some of these biases may have implications for the broad stock market. Even so, there exists little research to date that provides evidence that behavioral biases have a systematic impact on stock returns.² Therefore, we explore this debate further in a new direction by using an agent-based approach.

Recently agent-based modeling literature has grown. It is not our goal to review the entire literature, however, some important contributions include Chiarella (1992), Chiarella & He (2003), Chiarella, He, & Hommes (2006), Farmer (2002), Farmer & Patelli, Kirman (1993), Lux (1995), Lux & Marchesi (2000), Westerhoff (2008), and Hommes, Huang, & Wang (2005). Overall, these financial models model investors that pursue different strategies. Investors switch among these strategies based on past performance adding a positive feedback to the dynamics and are likely to destabilize the market.

Lux (1995) finds that above average returns are reflected in a more optimistic attitude that feedback into the beliefs of other investors. This type of behavior can lead to bubbles. We continue the use of modeling feedback effects.

Brock & Hommes (1998) explicitly model a switching mechanism where investors choose between a fundamental (rational) and trend follower (irrational) approach based on past profits. Their model differs from ours in that the irrational investor is modeled as an integral part of the

¹ See <http://www.behaviouralfinance.net/> for links to the various biases.

² There exist several papers such as Maymin (2009). Maymin shows how simple loss aversion can result in extreme distributions.

pricing process since investors are always switching between the two strategy types. There exists no switching mechanism in our model as we compare the different strategies that survive. Chiarella & He (2003) seek to determine how the results of Brock & Hommes are affected by an alternative market clearing mechanism. The focus is on a number of aspects of the modeling framework. Chiarella & He also examine how agents are allowed to have diverging attitudes to risk. We use different ways of calculating risk as a method to distinguish between various types of investors.

Chiarella, He, & Hommes (2006) seek to introduce traders that use moving average strategies to study the financial market dynamics that result from both fundamental and moving average strategies. We use both a moving average and exponential moving average strategies to distinguish investors in this paper. They find that price dynamics can become less stable as investors increase the window length.

Farmer & Joshi (2002) use a price impact function. Such a function describes the relation between the quantity of an asset bought or sold in a given time interval and the price change caused by these orders. A similar approach is used by Friedman & Abraham (2009) where the demand of the investors has an impact on the resulting price change. We continue the use of this price impact function. Westerhoff (2008) uses an agent-based model to evaluate certain regulatory policies such as transaction taxes, central bank interventions and trading halts. He finds these tools generally have the potential to stabilize financial markets.

We use a similar idea as Westernhoff, using an already established agent-based models to better understand financial market dynamics and alternative policies to stabilize markets. Our question relates to better understanding how irrational strategies impact financial market dynamics. We specifically use the pricing process from Friedman & Abraham and the modeling of irrationally introspective investors from Feldman (2011). Friedman & Abraham develop an agent-based model

where portfolio managers make choices based on the trade-off between return and the aggregate loss of all market participants. These investors price risk based on the concept of loss aversion as opposed to risk aversion. They find bubbles and crashes can occur when investors price risk based on the losses of all market participants. Feldman (2011) creates an agent-based model where both rational and irrational investors invest in a stock market to determine which behavioral bias has the most negative effect on performance relative to rational investors. The irrational investors modeled are of the irrationally introspective type who invest based on the biases of recency, loss aversion, and the disposition effect. In this paper we model three types of irrationally introspective investors and one type of irrationally aggregating investor. Two of the irrationally introspective investors are from Feldman (2011), loss averse and disposition effected investors. We model a new third type, overconfident investors. Loss aversion is where investors are more impacted by their losses than by their gains. Disposition effect is where investors hold on to stocks losing money and sell winning stocks that make them money. And overconfidence is where investors believe their ability to invest in the stock market is better than their actual investing skill. Therefore, they tend to invest in riskier investments.

Why do we model these biases as opposed to other biases? First, Kahneman and Tversky (1979) won the Nobel Prize in economics for formalizing loss aversion, where investors are more impacted by losses than by gains. The concept of prospect theory is now widely used in the literature to explain financial and economic phenomena. In addition, the behavioral finance literature finds evidence that individual investors are lousy investors earning returns well below that of the standard index. Odean (1998) and Barber, Lee, Liu, & Odean. (2007) find evidence to that the lousy returns are due to the disposition effect. In addition, Barber, Lee, Liu, and Odean (2009) find individual investors in Taiwan experience a performance penalty for their trades due to overconfidence and, as a result, trade too aggressively and to their detriment. The evidence suggests

that loss aversion, overconfidence, and disposition effect are some of the stronger biases. Lastly, we use a variant of the Friedman & Abraham model to model irrationally aggregating investors. These investors are impacted not only by their own losses but by the losses of all other participants in the market. Therefore, they aggregate information irrationally.

We use an agent-based framework in order to integrate the various types of biases investors into one market to analyze the price dynamics. Prior to agent-based models there have been authors such as De Long et al. (1990, 1991), that have argued to why imperfectly rational traders may survive in the long run. De Long et al. (1991) examine traders who are overconfident and therefore hold more of the risky asset thereby earning a higher return. As a result noise traders can earn higher expected profits than rational traders. Hirshleifer & Luo (2001) model overconfidence as overestimation of the precision of private information signals and they model prices endogenously. They find that overconfident traders do better than purely rational traders at exploiting misvaluation caused by liquidity or noise trading. We use these results on overconfident investors to create a biased overconfident investor in our simulation.

We use a lot of the same agent-based tools such as modeling price impact and feedback effects. However, the model differs in that there exists no switching in order to evaluate whether irrational investors can survive in a marketplace along rational investors.

3. Model Framework

We use the pricing framework from Friedman & Abraham to create a marketplace that includes both rational and irrational investors. Each type of investor buys and sells a single safe asset with constant return R_0 and a single risky asset with variable return R_1 . The safe asset can be considered cash and the risky asset as the S&P 500. Each investor chooses a single ordered variable

$u \in [0; \infty)$ that represents the allocation to the risky asset. For example, if $u=0.75$, the investor has 75% of their funds invested in the risky asset and 25% invested in cash. If $u = 1$ the investor is fully invested in the stock portfolio and if $u > 1$ the investor is borrowing cash to invest in the stock portfolio. The investor's net portfolio value is denoted by the variable z_i . The net portfolio value, z_i , grows or declines based on the manager's net return. The price of the risky asset results in the following process,

$$P_t = V\bar{u}_t^\delta. \quad (1)$$

where V is fundamental value, \bar{u}_t is the z -weighted mean allocation across all investors, and δ is a parameter that captures the sensitivity of price to buying pressure. Therefore, investors who earn a greater portfolio value, z_i , have a greater influence on the price. The fundamental value V is the present value of the per share earnings. It is calculated as $V(0) = \int_0^\infty 1e^{g_s t} e^{-R_s t} dt = (R_s - g_s)^{-1}$ where earnings growth is g_s and future earnings are discounted at some rate $R_s > g_s$.³

Asset supply comes from fundamental-oriented market participants such as issuers of stocks and bonds, and perhaps other individual investors. Investors exert buying pressure whose intensity is parameterized by δ . Even though supply does not change, one can simulate different environments where a low δ is consistent with a large supply and thus weaker buying pressure or a high δ is consistent with small supply.

The return on the risky asset, $R_{1,t}$, is determined from breaking down the price function into the dividend yield and the capital gains rate. Hence the realized yield on the risky asset is

$$R_{1,t} = (R_s - g)\bar{u}_t^{-\delta} + g + \delta\dot{\bar{u}}_t/\bar{u}_t, \quad (2)$$

³ The theoretical model is continuous. However, we create a discrete model when programming as investors rebalance every week.

where the first term is the dividend yield, the second term captures capital gains due to economic growth, and the third term reflects capital gains due to financial market activity. Note that R_I is equal to the discount rate R_s as in the CAPM when $\bar{u}_t = 1$ and $\dot{\bar{u}} = 0$.

The discount rate is $R_s = R_o + d_R$ where the term $d_R \geq 0$ represents all other factors. These factors include g , since economic growth is known and economy-wide.

The base payoff function is,

$$\Phi(u_{i,t}) = u_{i,t}(R_{1,t} - R_0 + \alpha_{i,t}) - \frac{1}{2}A\sigma^2 u_{i,t}^2. \quad (3)$$

The base payoff function is based on the traditional mean variance frame work where each investor chooses an allocation to the risky asset, u_i , based on the trade off between the excess return, $R_{1,t} - R_0$, and risk, $A\sigma^2$. However, there are three important differences. First, every investor differs by receiving an idiosyncratic shock, α_i . The α_i is a luck component. Some investors outperform the market one week via a positive shock and others underperform the market via a negative shock. For non-overconfident investors it is calculated using the most recent known value, $\alpha_i(t-1)$ and a random variable,

$$\alpha_i(t) = e^{-\tau}\alpha_i(t-1) + \sqrt{\frac{1-e^{-2\tau}}{2\tau}}\sigma v \quad (4)$$

for some given volatility parameter $\sigma > 0$, decay parameter $\tau > 0$, and an independent realization v from the unit normal distribution. (Feller, 1971, p 336).

The idiosyncratic shock follows a mean reverting Ornstein-Uhlenbeck process. If α_i is positive (negative) for investor i , she outperforms (under performs) the market, R_I . The intuition is that every investor holds a diversified portfolio that mimics the S&P 500 but may differ in the

collection of stocks one investor holds. At any point in time one investor may hold a stock that returns a high positive (negative) return that creates a positive (negative) alpha compared to the benchmark S&P 500. This under and out performance is quantified by their α which nets out to zero in the long-run. Second, the return, R_t is endogenously determined by the market participants as expressed in equation 2. Third, A is a constant risk aversion parameter as in the mean variance framework when investors are rational. However, A is endogenously determined by behavioral biases when investors act irrationally. We create irrational investors by turning A on. The manner in how A is modeled is dependent on the kind of behavioral bias the investor exhibits.

Investors adjust their allocation overtime by moving toward the optimal allocation to minimize trading frictions. Large investors and portfolio managers typically sell off or buy in increments over time. Friedman & Abraham call this following a gradient.

$$\Phi_{u_{i,t}} = R_{1,t} - R_0 + \alpha_{i,t} - A\sigma^2 u_{i,t}. \quad (5)$$

3.1. Rational Investors

We use the term rational in the sense of finance and not economics. The rational investor sticks to the traditional approach of the mean variance framework,

$$\Phi(u_{i,t}) = u_{i,t}(R_{1,t} - R_0 + \alpha_{i,t}) - \frac{1}{2}A\sigma^2 u_{i,t}^2, \quad (6)$$

where R_0 and A are constant parameters and $\alpha_{i,t}$ is an idiosyncratic shock. The rational investor uses a static long-term variance based on the historical variance of the S&P 500, 4%. They re-balance

every week as prices change to maintain their desired exposure given the long-run variance. The rational investors move toward their optimal amount of risky asset allocation as in equation five.

3.2. Irrationally Aggregating Investors

We model two sets of irrational investors. The first set of irrational investors are investors who aggregate irrationally.⁴ They are irrational in that they are very sensitive to their own losses as well as the losses of others around them. When they see others losing money their perception of risk increases even if there is no fundamental reason to change their risk perception. In addition, they put too much weight on current losses relative to past losses in assessing risk, a behavioral bias called recency. The base payoff function developed for the aggregating irrational investor is the following,

$$\Phi(u_{i,t}) = u_{i,t}(R_{1,t} - R_0 + \alpha_{i,t}) - \frac{1}{2}A_t\sigma^2u_{i,t}^2. \quad (7)$$

The payoff function is similar to the rational investor except that A has become A_t . The risk in this model, $A_t\sigma^2$, is defined as the perceived loss index, A_t , times the historical variance, σ^2 . The perceived loss index can be thought of as a risk aversion parameter, however, it is endogenously determined and time varying. It is determined by using a loss function and exponential averaging. The loss function is defined as negative gross return, $L_{i,t} = \max\{0, R_{Gi,t}\}$, where $R_{Gi,t} = (R_{1,t} - R_0 + \alpha_{i,t})u_{i,t} + R_0$ is the gross return that investor i currently earns on her portfolio at time t .

We then an exponential average of each investor's loss function. Investors seem to judge investments by the overall historical track record, with greater emphasis on more recent results. The

⁴ Irrationally aggregating investors are modeled similar to the portfolio managers in Friedman & Abraham.

natural formalization is an exponential average. The exponential average is updated from the previous exponential average loss $\hat{L}_i(t - 1)$ as follows:

$$\hat{L}_i(t) = e^{-\eta}\hat{L}_i(t - 1) + (1 - e^{-\eta})L_i(t). \quad (8)$$

The perceived loss index, A_t , is proportional to market-wide losses,

$$A_t = \beta\hat{L}_T(t), \quad (9)$$

the parameter $\beta > 0$ reflects investors' sensitivity to perceived loss, and $\hat{L}_T(t)$ is the perceived loss $\hat{L}_i(t)$ averaged across investors i weighted by portfolio size z_i at time t .⁵ Aggregating irrational investors aggregate the losses of all investors weighting the most current losses more than past losses. In this respect large losses from one or two investors can increase perceived risk among all other investors.

3.3. Introspective Irrational Investors

The second set of irrational investors are emotional about their own performance in three different ways: they are overconfident, loss averse, and disposition effected. Modeling several different types of biases allows for the possibility of self canceling trades as explained in the introduction.

⁵ Parameters such as β are explained in the Appendix.

3.3.1. Overconfident Investors

We formalize overconfident investors via $\alpha_{i,t}$. We assume that overconfident investors overestimate their ability to outperform the overall market. More specifically, overconfident investors tend to remember positive performance shocks more than the negative performance shocks whereas all other investors treat the positive performance and negative performance shocks equally. We formalize this by using a smaller τ (greater memory) parameter for positive shocks and a larger τ for negative shocks, (less memory). We define $\hat{\alpha}_{i,t}$ as the alpha for an overconfident investor. Overconfident investors continuously believe they are outperforming the market and therefore add more leverage to their portfolio. During down markets they may get margin calls because of the overuse of leverage. Their payoff function then becomes

$$\Phi(u_{i,t}) = u_{i,t}(R_{1,t} - R_0 + \hat{\alpha}_{i,t}) - \frac{1}{2}A\sigma^2u_{i,t}^2. \quad (10)$$

and their payoff gradient becomes

$$\Phi_{u_{i,t}} = R_{1,t} - R_0 + \hat{\alpha}_{i,t} - A\sigma^2u_{i,t}. \quad (11)$$

3.3.2. Loss Averse Investors

Loss averse investors are very skittish when it comes to losing money. If they experience heavy personal losses their perception of risk increases dramatically. The loss averse investors price risk using a moving average of past losses called loss aversion risk. The moving average of losses or loss aversion risk, $\hat{L}_{i,t}$ is determined in two steps. First, it is determined by picking up the

investor's current losses, $L_{i,t} = \max \{0, R_{G_{i,t}}\}$, where $R_{G_{i,t}}$ is the gross return at time t . Next, the losses over time are averaged using a moving average where n equals the number of periods used to calculate the moving average. In the base case n is equal to two years. We use two years because research suggests investors weight the most recent year and then the decay ends after another year.⁶ This step is distinct from an exponential averaging in that current and past losses are weighted equally. When calculating successive values, a new value comes into the sum and an old value drops out,

$$\hat{L}_{i,t} = \beta \left(\hat{L}_{t-1} - \frac{\hat{L}_{i,t-n}}{n} + \frac{\hat{L}_{i,t}}{n} \right). \quad (12)$$

The payoff function for the loss averse investor becomes

$$\Phi(u_{i,t}) = u_{i,t}(R_{1,t} - R_0 + \alpha_{i,t}) - \frac{1}{2} \hat{L}_{i,t} \sigma^2 u_{i,t}^2. \quad (13)$$

and the payoff gradient becomes,

$$\Phi_{u_{i,t}} = R_{1,t} - R_0 + \alpha_{i,t} - \hat{L}_{i,t} \sigma^2 u_{i,t}. \quad (14)$$

3.3.3. Disposition effected investors

The disposition effected investors tend to buy or hold when incurring losses and sell when the market increases. They hold on to losses for too long because they feel regret of choosing a losing stock and want to recover their losses. On the flip side they tend to sell too early. To formalize this process we calculate disposition risk. To calculate disposition risk we first pick up the

⁶ See Feldman (2010).

gains, $D_{i,t} = \max \{0, R_{G_{i,t}}\}$. Second a moving average is calculated where the last value moves out and the most current value moves into the calculation of the moving average,

$$\widehat{D}_{i,t} = \beta \left(\widehat{D}_{t-1} - \frac{D_{i,t-n}}{n} + \frac{D_{i,t}}{n} \right). \quad (15)$$

When disposition effected investors experience gains, $D_{i,t}$ is high, and therefore they tend to sell too fast. If they experience losses, disposition risk is low and they tend to hold or buy more shares because of regret. The formalization of this process means that these investors cost average down. For example, as the price of the risky asset falls, investors buy more of the risky asset reducing their cost basis. Similar to the loss averse investors, disposition effected investors average their $D_{i,t}$ for several years weighting observations equally.

The payoff function for the disposition effected investor becomes

$$\Phi(u_{i,t}) = u_{i,t}(R_{1,t} - R_0 + \alpha_{i,t}) - \frac{1}{2} \widehat{D}_{i,t} \sigma^2 u_{i,t}^2. \quad (16)$$

and the payoff gradient is,

$$\Phi_{u_{i,t}} = R_{1,t} - R_0 + \alpha_{i,t} - \widehat{D}_{i,t} \sigma^2 u_{i,t}. \quad (17)$$

3.4. Solvency

The portfolio wealth of the typical investor is $z_{i,t}$. The portfolio size increases or decreases based on the gross return of the investor. The portfolio size can become negative when investors borrow at the risk free rate to invest in the risky asset and the asset price plunges. If the gross return

decreases such that the portfolio wealth becomes negative the simulation program kicks that investor out of the market and the total investor population declines by one.

4. Simulations

We run the simulation using an agent-based program called NetLogo. The interface can be seen in Figure 1. The simulation consists of investors whose allocation u_i (horizontal coordinate) and portfolio size z_i (vertical coordinate). The simulation is at weekly frequency (Freq = 52). The sliders in the interface, center, width, altitude, and height determine the initial u and z for each investor. The "altitude" sets the middle of the initial z distribution, and "width" and "height" control the bounds on the rectangle. At the start of the simulation each investor randomly starts off with a different u and z based on an uniform distribution. The initial population of managers is uniformly distributed in the (u, z) rectangle $[0.2, 1.4] \times [0.4, 1.6]$, set via the sliders labeled "population". The initial conditions change at the beginning of every new simulation run. Each simulation lasts one hundred years or 5,200 weeks. We delete the first twenty years to insure that initial conditions do not impact the results. We run ten simulations for each regime using the baseline parameter configuration explained in the Appendix.

We run various regimes where the total population in each regime remains at thirty, however, the composition of the population per regime changes. Friedman & Abraham find thirty investors is the suitable number of investors to achieve realistic stock price dynamics.

1. Regime 1: Rational: Thirty rational investors.
2. Regime 2A: Irrationally Aggregating: Thirty irrational investors that aggregate losses irrationally.

3. Regime 2B: Irrationally Introspective: Thirty irrationally introspective investors. Ten loss averse investors, ten overconfident investors, and ten disposition effected investors.
4. Regime 3A: Combined 3A: Fifteen rational investors and fifteen irrationally aggregating investors.
5. Regime 3B: Combined 3B: Fifteen rational investors, five loss averse, five overconfident, and five disposition effected investors.

The parameters are configured based on long-run averages of the S&P 500. For example, the long-run standard deviation of the S&P 500 is 20% based on total returns of the S&P 500 for the period from 1926 through 2014. This number is used in the calculation of α_i .⁷ The long-term real US GDP growth rate is approximately 2%. We also define a recession or financial crisis as a 30% fall from the peak in the last six months.

Lastly, we use a baseline parameter configuration used in Friedman & Abraham. The baseline configuration is presented in Table 1 in the Appendix. However, we create eight different parameter configurations for robustness testing. The eight different configurations are created by changing one of the parameters to an admissibly high level and then to an admissibly low level while keeping all other parameters constant. For example, one configuration is where we change δ from the baseline of two to four. The other configuration is where we change δ from two to one.

Table 2 summarizes the number and different types of simulations. For each regime we run ten simulations using a baseline configuration. Each simulation last 5,200 weeks or one hundred years. However, we delete the first 4,160 weeks in order to test the long-run implications of irrational behavior. Therefore, there exists 1,040*10, or 10,400 observations per regime. We

⁷ Source from <http://www.stern.nyu.edu/~adamodar/pc/datasets/histretSP.xls>.

calculate the relevant statistics for each simulation within each regime first before aggregating by regime as opposed to pooling the simulation data and then running statistics.

5. Results

Table 3 displays the return properties between the purely rational and purely irrational regimes. Results suggest, as expected, the rational regime exhibits low volatility, standard deviation equals six percent, and returns are characterized by no fat tails. However, the returns do exhibit slight negative skewness. Both the irrational regimes exhibit high volatility, over 25% standard deviation, large negative skewness, and fat tails. The following sections draw a pairwise comparison of the results across the various regimes.

5.1. Rational and Irrationally Aggregating

We first a look at the null hypothesis that the return properties from the combination of the irrationally aggregating and rational investors simulations, Combined 3A, converge to that of the rational regime in the long-run. Using the Mann-Whitney test, in Table 4, we can reject the null hypothesis that the mean return is the same across the two regimes at the one percent significance level. In addition, the Combined 3A regime's mean price is \$23.33 which is below the mean price of the rational regime of \$28.32. Moreover, the maximum and minimum Combined 3A regime return is similar to that of a large financial crisis experienced in developed markets. The minimum and maximum annual return for the Combined 3A regime are -65% and 50%, respectively, where as the range for the rational regime is between -34% and 15%. Results from Table 5 indicate other distributional features are different as well. The rational regime return distribution is slightly

skewed with no fat tails, whereas regime 3A's return distribution is not normal, being characterized by negative skewness and positive excess kurtosis.

Results from Table 6 displays the results of a Kolmogorov-Smirnov test. The test indicates there exists a significant difference (p-value = 0.0001) in return distributions between the rational regime and Combined 3A.

Lastly, the crash dynamics shown in Table 7 also confirm a significant difference between the two regimes. The average number of 30% crises in a twenty year period for the Combined 3A regime is 1.08 which equates to one crisis every eighteen years. The maximum number of crises recorded in one simulation is four. This is compared to the rational regime which averages 0.20 crises per twenty years or one crisis every century. Formally, a t-test rejects the null hypothesis that the average number of crises is the same between the two regimes at the 1 percent significance level.

5.2. Rational and Irrationally Introspective

We now turn to the irrationally introspective investors. We find there is no significant difference between the rational and combined rational and irrationally introspective regime Combined 3B. The mean, maximum, and minimum return found in Table 4 are very close between the Combined 3B and rational regime. The only difference between the two regimes is that the Combined 3B regime's mean price is significantly lower than the rational regime's mean price even though the return properties are similar. The irrational investors have the effect of lowering the price below the fundamental value. Barber and Odean (2006) find that individual investors consistently under perform market benchmarks. Our results suggest the macro impact of individual investors could be that they lower the market price below its fundamental value.

Results from Table 5 indicate that both the rational and Combined 3B regime's return properties follow a close approximation to a normal distribution. Formal Kolmogorov-Smirnov tests in Table 6, however, reveal that both regimes exhibit non-normal distributions because of the negative skewness. Results suggest that these differences are relatively small in economic terms, though formal tests reveal that the return distributions are statistically different between the two regimes.

Moreover, the crisis statistics in Table 7 suggest there is no significant difference between the two regimes. On average there occurs around one crisis per century in the Combined 3B and rational regime.

5.3. Differences Among Irrational Investors

We also ask whether different types of irrational investors have different effects on asset prices. Results from Table 8 indicate that irrationally aggregating investors alongside rational investors have a greater impact on volatility and fat tails than the irrationally introspective investors alongside rational investors. The irrationally introspective investors impact is that they lower the mean price significantly below fundamental value. The fundamental value in the baseline parameter configuration is \$25 per share. The mean price for the rational regime is \$28.32. The mean price in the combined rational-irrationally introspective, Combined 3B, regime is \$19.84, twenty percent below fundamental value.

Figure 1 displays a histogram of the market's yearly returns for the rational regime, Combined 3A, and Combined 3B. It can be seen that the rational regime and Combined 3B are comparable, annual returns approximate a normal distribution. However, Combined 3A's return distribution is significantly skewed exhibiting large fat tails.

5.4. Performance

Table 9 displays return statistics for the three types of investors rational, irrationally aggregating, and irrationally introspective.⁸ The performance results highlight some of the conclusions we have written about so far. Results indicate the aggregating investor outperforms all other investors on a absolute basis. However, the rational investor outperforms all other investors on a risk adjusted basis.

5.5. Robustness Testing

We explore the long-term properties in the previous section by deleting the first 80% of observations per simulation. The 80% number is an arbitrary deletion point. Therefore, we explore annual return properties by deleting 40%, 50%, 60%, 70%, and 80% of the first observations per simulation. Table 10 indicates that the return properties are relatively stable across the various arbitrary deletion points. We find that the irrationally introspective investors do not have a strong impact on the macro price throughout the simulation.

However, we know from regime 2B that without rational investors the irrationally introspective investors create excessive volatility, negative skewness, and fat tails. Therefore, as EMH theorists argue, irrationally introspective investors do not impact stock returns when combined with rational investors.

Table 11 displays the volatility results for simulations using eight other parameter configurations. In these eight parameter configurations we only change one parameter at a time to a

⁸ We find that every type of introspective investor returns are below the rational and aggregating investor.

very high level and then a very low level. The purpose of this type of parameter configuration is to test whether the results hold under more extreme parameter configurations. If so, the robustness test would be passed.

The volatility in the Combined 3A regime is significantly higher than the other two regimes under all parameter configurations except for when τ is equal to three. The volatility in the Combined 3B regime is closer to the rational regime but modestly higher suggesting the introspective irrational investors play a role at times in asset price formation when investors are sensitive.

6. Conclusion

At one extreme, even if some investors are irrational, aggregate market forces can result in substantial market rationality when irrational investors do not aggregate but act as noise traders. At the other extreme, irrational investors may get richer, not poorer during stock market bubbles when they aggregate information about others irrationally. We find that the type of irrationality is an important factor in understanding how psychology impacts asset prices in the long-run. If the type of irrationality is introspective we find evidence to prove the EMH argument correct. Irrational traders tend to go bankrupt and trades cancel each other out, so that they only have a minor impact on market outcomes as EMH suggests. However, we find that when investors aggregate information in an irrational manner they can have a strong impact on the broad market.

Our results suggest the following for future modeling of financial markets. First, noise traders or what we call irrationally introspective investors have little impact on asset prices in the medium and long-run. Therefore, they could be safely ignored for future work to understand long-run price dynamics. Secondly, EMH theory may for long periods of time, however, there can occur

periods where investors aggregate information irrationally. This process of aggregating information irrationally is a slow drawn out process leading to periods where EMH theory does not hold. Based on these results we argue that behavioral finance and EMH are compatible and not competitive or one is right and the other is wrong. In the simulations the irrationally aggregating investor can appear to look rational. We know which investor is which, however, if one did not know it would be hard to distinguish between the two. However, at other times it can be easy to distinguish. We believe the appropriate modeling is to model market efficiency where biases can interfere in the marketplace from time to time. In other words where biases are turned off and then turned on again. And those biases must be of the irrationally aggregating type.

7. Bibliography

Barber B, Odean T (2000) Trading is hazardous to your wealth: The common stock investment performance of individual investors. *Journal of Finance*, Vol. LV, No. 2: 773-806.

Barber B, Lee Y, Liu Y, Odean, T (2007) Is the aggregate investor reluctant to realize losses? *European Financial Management*, 13: 423-447.

Barber B, Lee Y, Liu Y, Odean, T (2009) How much do investor lose from trade? *Review of Financial Studies*, 22, 2: 609-632.

Brock W, Hommes C (1998) Heterogeneous beliefs and routes to chaos in a simple asset pricing model. *Journal of Economic Dynamics and Control*, Vol. 22(8-9): 1235-1274.

Chiarella C (1992) The dynamics of speculative behaviour. *Annals of Operations Research*, 37: 101-123.

Chiarella C, He X (2003) Heterogeneous beliefs, risk, and learning in a simple asset-pricing model with a market maker. *Macroeconomic Dynamics* 7(4): 503-536

Chiarella C, He X, Hommes C (2006) A dynamic analysis of moving average rules. *Journal of Economic Dynamics and Control* 30(9): 1729-1753.

- De Long B, Shleifer A, Summers L.H., Waldmann R (1990) Noise trader risk in financial markets. *Journal of Political Economy* 99 (4): 703-738.
- De Long B, Shleifer A, Summers L.H., Waldmann R (1991) The survival of noise traders in financial markets. *Journal of Business* 64 (1): 1-19.
- Fama E. F. (1965) The behavior of stock market prices, *Journal of Business* 38: 34105.
- Fama E. F. (1970) Efficient capital markets: A review of theory and empirical work, *Journal of Finance* 25: 34105.
- Farmer D (2002) Market force, ecology and evolution. *Industrial and Corporate Change*, 11(5): 895-953.
- Farmer D, Joshi, S (2002) The price dynamics of common trading strategies. *Journal of Economic Behavior and Organization*, 49: 149-171.
- Farmer D, Patelli P, Zovko I (2005) The predictive power of zero intelligence in financial markets. *PNAS* 102(6).
- Feldman T (2010) A more predictive index of market sentiment. *Journal of Behavioral Finance*, 11: 211-223.
- Feldman, T (2011) Behavioral biases and investor performance. *Algorithmic Finance*, 1:1: 45-55.
- Friedman D, Arbaham, R (2009) Bubbles and crashes: Gradient dynamics in financial markets. *Journal of Economic Dynamics and Control*, 33: 922-937.
- Hirshleifer D, Luo G. Y (2001) On the survival of overconfident traders in a competitive security market, *Journal of Financial Markets* 4: 7384.
- Hommes C, Huang H, Wang D (2005) A robust rational route to randomness in a simple asset pricing model. *Journal of Economic Dynamics and Control* 29(6): 1043-1072.
- Kahneman D, Tversky A (1979) Prospect theory: An analysis of decision under risk. *Econometrica*, 47: 263-291.
- Kirman A (1993) Ants, rationality and recruitment. *Quarterly Journal of Economics* 108: 137-156.
- Lux T (1995) Herd behaviour, bubbles and crashes. *Economic Journal* 105(431): 881-896.
- Lux T, Marchesi M (2000) Volatility clustering in financial markets: A microsimulation of interacting agents. *International Journal of Theoretical Applied Finance* 3(4): 675-702
- Maymin P (2009) Prospect theory and fat tails. *Risk and Decision Analysis*, 1: 187-195.

Odean T (1998) Are investors reluctant to realize their losses? *Journal of Finance*, 53: 1775-1198.

Shefrin H, Statman M (1985) The disposition to sell winners too early and ride losers too long: Theory and evidence. *Journal of Finance*, 40: 791-792.

Westerhoff FH (2008) The use of agent-based financial market models to test the effectiveness of regulatory policies. *Journal of Economics and Statistics (Jahrbuecher fuer Nationaloekonomie und Statistik)*, 228(2):195.

8. Appendix

Table 1: Baseline Parameters

Parameter	Baseline	Definition	Meaning	Theory
R_o	3%	risk-free rate		$\uparrow \Rightarrow$ more to risk-free asset
dR	6%	discount-rate		$\uparrow \Rightarrow$ more to risk-free asset
g	2%	growth-rate	economic growth	$\uparrow \Rightarrow$ more to risky asset asset
σ	20%	standard deviation	variability of idiosyncratic shock	$\uparrow \Rightarrow$ more crises
τ	0.70	persistence of idiosyncratic shock	shock persistence	$\uparrow \Rightarrow$ less crises
η	0.70	memory rate	current losses versus past losses	$\uparrow \Rightarrow$ more crises
β	2	sensitivity to risk	how sensitive investors are to risk	$\uparrow \Rightarrow$ more crises
δ	2	elasticity	sensitivity of demand change to price	$\uparrow \Rightarrow$ more crises

Note. The parameters are chosen based on long-term calculations from 1926 to 2014. The 3% average for Tbill was calculated based on returns for the period from 1926 through 2014 using T-bills represented by the Barclays 3-Month Treasury Bellwether index. Based on the World Bank data the average real GDP growth rate from 1969-2014 was 2.7%. If we lengthen the time frame to 1926 we estimate a growth rate of 2%. The discount rate can be thought of as the US stock market risk premium. The average risk premium from 1926 to 2014 was around 6%. See <http://www.stern.nyu.edu/~adamodar/pc/datasets/histretSP.xls>

9. Tables and Figures

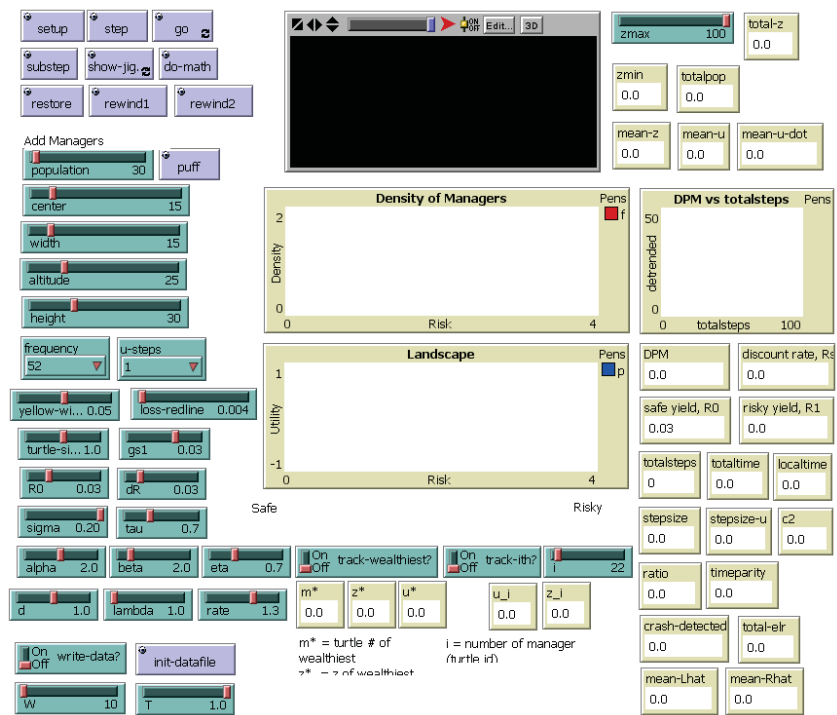


Figure 1: Netlogo Interface

Table 2: Summary Stats

Regime	# of Sims	# of Rational	# of Irrational	Nobs per Sim
1. Rational	20	30	0	1,040
2A. Aggregating Irrational	20	0	30	1,040
2B. Introspective Irrational	20	0	30	1,040
3A. Rational & Aggregating	20	15	15	1,040
3B. Rational & Introspective	20	15	15	1,040

Note. Each simulation lasts one century or 5,200 weeks. We delete the first 80 years or 4,160 observations.

Table 3: Return Properties: Irrationally Introspective vs. Irrationally Aggregating

Regime	Mean Return	Max Return	Min Return	Std. Dev.	Skewness	Kurtosis
Rational	7.76%	15.53%	-34.66	6.56%	-1.10	2.56
Aggregating	12.11%	58.79%	-93.87%	28.42%	-4.82	33.80
Introspective	10.52%	40.95%	-98.34%	25.03%	-3.80	21.76

Note. Returns are annualized. Nobs per simulation = 1,040. Nobs per regime = 10*1,040=10,400. Combined_3A is the regime where irrationally aggregating and rational investors market up the market. Combined_3B is the regime where irrationally introspective and rational investors market up the market.

Table 4: Price Properties

Regime	Mean Price	Max Price	Min Price	Mean Return	Max Return	Min Return
Rational	28.32	41.46	20.42	7.76%	15.53%	-34.66%
Combined_3A	23.33**	42.68	5.90	7.96%**	50.11%	-65.81%
Combined_3B	19.84**	34.08	12.11	7.40%	20.45%	-30.85%

Note. Returns are annualized. Combined_3A is the regime where irrationally aggregating and rational investors market up the market. Combined_3B is the regime where irrationally introspective and rational investors market up the market. Nobs per simulation = 1,040. Nobs per regime = 10*1,040=10,400. Mann Whitney test is used to test difference in means between the two regimes. Z=65.35 for difference in mean price and Z=-17.65 for difference in mean return between Rational and Combined_3A regime. Z=109 for difference in mean price and Z=-0.298 for difference in mean return between Rational and Combined_3B regime. Null hypothesis is that combined_3A and combined_3B regime is not statistically different from the Rational regime. *significant at 5%; **significant at 1%.

Table 5: Distribution Stats

Regime	Std. Dev.	Skewness	Kurtosis
Rational	6.58%	-1.10	2.57
Combined_3A	18.45%	-2.55	12.73
Combined_3B	7.91%	-0.52	0.29

Note. Nobs per simulation = 1,040. Nobs per regime = $10 \times 1,040 = 10,400$. Combined_3A is the regime where irrationally aggregating and rational investors make up the market. Combined_3B is the regime where irrationally introspective and rational investors market up the market.

Table 6: Difference in Distribution

Test	D	P value
Combined_3A K-S	0.17	0.0001
Combined_3B K-S	0.07	0.0001

Note. K-S refers to Kolmogorov-Smirnov test. D refers to the magnitude of the difference in distribution. Nobs per simulation = 1,040. Nobs per regime = $10 \times 1,040 = 10,400$.

Table 7: Crisis Stats

Regime	Avg. # of Crises per 20 Years	Avg. Crisis Time	Max # of Crises
Rational	0.19	9	1
Combined_3A	1.08**	41	4
Combined_3B	0.10	11	1

Note. A crisis is defined as a 30% drop from the most recent high in the last six months. Crisis time is measured in weeks. Combined_3A is the regime where irrationally aggregating and rational investors market up the market. Combined_3B is the regime where irrationally introspective and rational investors market up the market. Nobs per simulation = 1,040. Nobs per regime = $10 \times 1,040 = 10,400$. T test is used to test statistical differences. Null hypothesis is that combined_3A regime and combined_3B regime is not statistically different from regime 1. *significant at 5%; **significant at 1%.

Table 8: Return Properties: Aggregating Irrational vs. Introspective Irrational

Regime	Mean Return	Max Return	Min Return	Std. Dev.	Skewness	Kurtosis
Combined_3A	7.94%	51.11%	-65.81%	18.45%	-2.60	12.76
Combined_3B	7.35%	21.05%	-31.48%	7.85%	-0.52	0.30

Note. The return is an annual return. Nobs per simulation = 1,040. Nobs per regime = $10 \times 1,040 = 10,400$. T stats are used to test statistical differences. Null hypothesis is that regime 2 and 3 are not statistically different from regime 1. *significant at 5%; **significant at 1%.




Figure 2: Differences in Distribution

Table 9: Performance Comparison

Group	Average	St.Dev	Sharpe
1. Rational	5.15%	4.90%	0.083
2. Aggregating	6.10%	6.80%	-0.233
3. Introspective	1.45%	6.02%	-1.52

Note- Return statistics per group. All statistics are annualized over a one hundred year simulation period. The mean, standard deviation, and Sharpe are computed by averaging over the simulation runs per regime and then across the regimes. We assume the risk free rate is 3% in computing the Sharpe ratio.

Table 10: Robustness Test: Deletion Points

Regime	40%	50%	60%	70%	80%
Standard Deviation					
Rational	0.069	0.068	0.067	0.064	0.065
Combined_3A	0.179	0.182	0.165	0.175	0.185
Combined_3B	0.079	0.080	0.081	0.081	0.079
Skewness					
Rational	-0.77	-0.76	-0.91	-0.94	-0.77
Combined_3A	-3.21	-3.31	-2.36	-2.45	-2.59
Combined_3B	-0.50	-0.47	-0.44	-0.44	-0.51
Kurtosis					
Rational	1.772	1.776	2.412	2.338	2.672
Combined_3A	21.031	22.383	12.524	12.276	12.73
Combined_3B	0.503	0.236	0.193	0.112	0.29

Note- The percentages are the percent of observations deleted in each simulation. Each simulation lasts 100 years. Therefore, 40% means that the first 40 years of observations were deleted. Fifty percent means the first 50 years of data were deleted.

Table 11: Robustness Test: Difference in Standard Deviation

Parameter	Number	Rational	Combined_3A	Combined_3B
tau	0.1	18.18%	28.45%	26.31%
	3	2.86%	4.79%	4.96%
eta	0.1	6.70%	14.87%	8.34%
	3	6.22%	15.78%	6.23%
beta	1	6.83%	11.90%	9.24%
	5	7.14%	42.15%	8.92%
delta	1	3.01%	4.84%	3.93%
	3	12.40%	48.58%	23.65%

Note- Tau signifies how long the stochastic shock lasts for in the alpha calculation. Larger tau means less persistence. Eta signifies memory of losses. Larger eta means less memory. Beta signifies sensitivity to losses. The larger the beta the larger the sensitivity. Delta signifies the sensitivity to changes in investor demand on the stock price.