Trust and Behavior in a Model of Financial Markets: An Agent-Based Approach

Adam Rhodes-Rogan*

April 18, 2013

Abstract

This thesis explores the question of whether or not trust in a friend's opinion affects general market behavior over time. We propose an agent-based model, programmed with the NetLogo modeling software, to test this question, and allow agents to be influenced by exogenous information and their friends opinions when considering to buy or sell an asset. Through agents' trading behavior, we can observe changes in the price of the asset over time, and observe trends that can be explained by herding behavior. We predict that higher average levels of trust will cause the market to crash more frequently, and will increase the probability of a crash occurring. We find support for these hypotheses, showing that higher average trust causes a significantly greater number of crashes to occur and significantly increases the probability of a crash occurring.

1 Introduction

How does interaction among individuals influence the emergent, aggregate properties of a system? With regards to financial markets, this is a persistent question. To attempt to answer it involves a seemingly endless number of assumptions about individuals' psychological characteristics and motivations, along with more conjecture about the way interact in a market setting. Some theoretical explanations of market behavior have been elevated to the level of a standard in academics. These are far from complete, however, and more nuanced approaches to addressing both the validity of assumptions made, and their implications, are needed.

This thesis explores the question of how a particular type of market behavior, a crash, is to modeled. We do this using a simple, agent-based model, and collect data on the relationship between agent parameters and market behavior. The paper is laid out as follows: we outline a brief history of the theory behind modeling financial markets; the modeling software we use, NetLogo, is explained; we describe the current model, its parameters, and mechanisms; we present our findings and discuss their implications.

2 Literature Review

2.1 The Efficient Markets Hypothesis

The dominant explanation of market behavior for the past half-century has been the efficient markets hypothesis (EMH), first popularized by Fama (Shleifer, 2000; Fama, 1965). The EMH explains how markets function to represent information in the form of a price for any given security. In essence, there are three main arguments made by the EMH, from which the rest of the explanation follows.

*University of Vermont. Version 3. April, 2013. I would like to thank Professor Bill Gibson for his unconditional guidance, patience, and passion. Without his support, and exceptional knowledge of agent-based modeling and economic theory, I would not have been able to this complete this project. I would also like to thank Professor John F. Summa for his exuberance with, and insistence on, adopting a critical perspective of social-scientific research. Additionally, I would like to thank Professor Michael J. Tomas III for teaching me about the finely-detailed nature of the financial industry, and for his efforts to help build my professional network. Contact: arhodesr@uvm.edu.

Those arguments are: 1) that investors are, on average, rational and value securities rationally; 2) if some investors are not rational, their market actions are randomly distributed, thus their actions cancel each other in the aggregate and do not affect a security's price; 3) to the extent that investors are irrational in a similar fashion, their actions are countered by smart investors seeking arbitrage who serve to eliminate the irrational investors' effect on a security's price (Shleifer, 2000). From these arguments, one can see how, in an efficiently functioning market, prices reflect all relevant information about a given security (Shleifer, 2000; Samuelson, 1965). Moreover, the EMH holds that any errors in expectations about a security's price are uncorrelated. This suggests that no individual market participant would consistently make errors with regards to his trading behavior.

The theoretical framework of the EMH provides an ample starting point for empirical evaluation of market behavior. Many have observed market movements that would appear to be sufficient for the EMH to be true. A random walk pattern, or a series of price changes that appear to be random, has commonly been fit to certain periods of financial market activity (Kahneman, 2003; Shiller, 1981). Similarly, there has been a large body of literature focused on identifying instances where market behavior does not follow that predicted by the EMH. Some have addressed questions about the relationship proposed to exist between price and information. Generally, any information about a security can motivate a trader to act, but material information, such as announced changes in dividend offerings, may carry more importance than other kinds (Shiller, 1981). Shiller found that the frequency of exchange on the S&P 500 was too high, relative to his EMH-predicted frequencies, for information about dividends to be the sole influence on trading (Sewell, 2012; Shiller, 1981). In this paper, we do not attempt to address the EMH or its assumptions, but instead propose and evaluate a method for modeling financial crashes, or rapid declines in asset prices over a relatively short period of time. We examine crashes specifically because they are not easily explained by the EMH, and might provide greater insight as to how the market is a function of individual interactions.

2.2 Crashes in Financial Markets

Historically, crashes in financial markets tend to be statistically rare events (Gonçalves, 2003). A crash might be defined as a significantly sharp decline in securities' prices across an entire market. While many can likely recall famous crashes anecdotally, such as the recent 2009 crash, the crash in the late 1990s, or the crash preceding the Great Depression, over the entire history of financial market activity, crashes are often considered outliers (Sewell, 2012; Johansen et al., 2008). Crashes are difficult to explain under the context of the EMH and many questions about the nature of this unique market behavior continue to be unanswered (Stein, 2009; Sharpe, 1991). Some wonder if markets are naturally prone to crashes, while others attempt to identify exogenous factors that caused each crash (Barberis and Shleifer, 2003). Recently, some have modeled crashes as prolonged organized market states (Johansen et al., 2008).

2.3 Crashes as Organized States

The theory behind modeling crashes in this way stems from the idea of herding behavior, which has been used to explain the formation of asset bubbles and crashes by many (Sewell, 2012; Sornette, 2003; Shiller, 1981). This behavioral approach to explaining the aggregate effects of individual traders' actions on markets is not new; consider Keynes' beauty contest explanation and Mackay's statement about the madness of crowds (Shleifer, 2000). A key aspect of herding behavior is that its eventual outcomes can be irrational from the perspective of those involved, even when each individual's behavior is rational to some extent (Shiller, 1981). For example, consider the following allegory. An individual wants to eat dinner at one of two new restaurants that have opened right next to each other in his hometown. He wants to eat at the restaurant with the highest quality food, but he has not been presented with any reliable information about either establishment's quality. He determines instead, that he will eat at the restaurant that seems to be more popular when he arrives. Upon arrival, he sees that restaurant A has many patrons, while restaurant B only has a few. So, this individual rationally chooses to dine at restaurant A. However, by doing so, he has unknowingly influenced the actions of likeminded individuals, who will dine at the restaurant that seems to be more popular. As his initial decision, and others' subsequent decisions about which restaurant to attend cascade down the series of choices made, the result would be that a significantly larger number of people would be eating at restaurant A. Now, should restaurant A actually have objectively lower quality food than restaurant B, this outcome would be considered irrational. One can easily see how traders on financial markets can be prone to a similar form of positive reinforcement. Indeed, others have contextualized the basics of the aforementioned process to financial markets (Ghashghaie, 1996).

To examine this herding process, and its prevalence in financial markets, some have attempted to model it using an agent-based approach (Johansen et al., 2008; Gonçalves, 2003; Barberis and Shleifer, 2003). Depending on the model, there may be different distinctions between classes of investors, such as noise traders and arbitrageurs, or rational and irrational investors (Barberis and Shleifer, 2003). Sometimes, a distinction is made between a natural, disorganized state, and an organized market state (Johansen et al., 2008). We are interested in this distinction in our current study.

In a disorganized state, the price of the asset on the market might display a random walk pattern, as would be predicted by the EMH. However, in an organized state, the asset's price trends upwards or downwards for extended periods of time (Johansen et al., 2008). Using this theoretical framework allows for examining what might cause the market to enter an organized state, how often such states occur, and how quickly they resolve. Crucially, phenomena such as bubbles and crashes can be modeled and examined with this approach, and some have done so effectively (Johansen et al., 2008; Gonçalves, 2003; Sornette, 2003). Traditionally, crashes in these types of models have been defined as a rapid switch from one organized state to another. Should the asset's price experience a prolonged upward trend, a crash could occur if the trend changes direction without any intermittent disorganized states (Johansen et al., 2008). An organized state that resolves with a disorganized state would be considered a correctional movement in this case (Johansen et al., 2008).

The models that frame market behavior as either organized or disorganized often utilize behavioral parameters when specifying how agents choose to behave. When individuals make choices under uncertainty, as is the case with many investment decisions, they tend to use cognitive heuristics and rely on factors other than information directly relevant to the decision at hand (Kahneman, 2003). These can include whether or not a similar investment has performed well in the past, or a trusted friend's opinion on the investment (Baumeister et al., 2001; Slovic, 1993). The question is whether or not utilizing this information can lead to herding behavior and irrational market outcomes. Some might argue that professional traders would not seek to use irrelevant information, and would be able to make decisions about securities based solely on the relevant information at hand (Sharpe, 1991). However, in some models that incorporate contagion risk under an organized/disorganized framework, and further segment the market into professionals and non-professionals, the correlated behavior of the non-professionals ultimately overpowers that of the others, sending the market into organized states (Sornette, 2003).

In this paper, we seek to answer the question of how an individual's trust in the opinions of others affects market behavior, insofar that it does. Specifically, we are interested in whether or not organized states are more likely to occur when traders on a market have more trust in each other's opinions. That is to say, are organized states more likely to occur when the average individual places high importance on the market outlook held by his closest colleagues? Since we are not entirely interested in bubbles that do not result in crashes, we seek to examine how different levels of the average individual's trust affect the probability of a crash occurring. To do this we use an agent-based model, programmed using modeling software called NetLogo. The rest of the paper is laid out as follows: the capabilities of NetLogo are described; the original model adapted for use in this paper is explained; the parameters and functions of the current model are explained in detail; and finally our methods and results are discussed.

3 Modeling a Financial Market

3.1 NetLogo

NetLogo is the modeling software used for this study. Like other modeling software, NetLogo allows the user to design a world in which psychological and behavioral aspects of different agents are shaped by defined

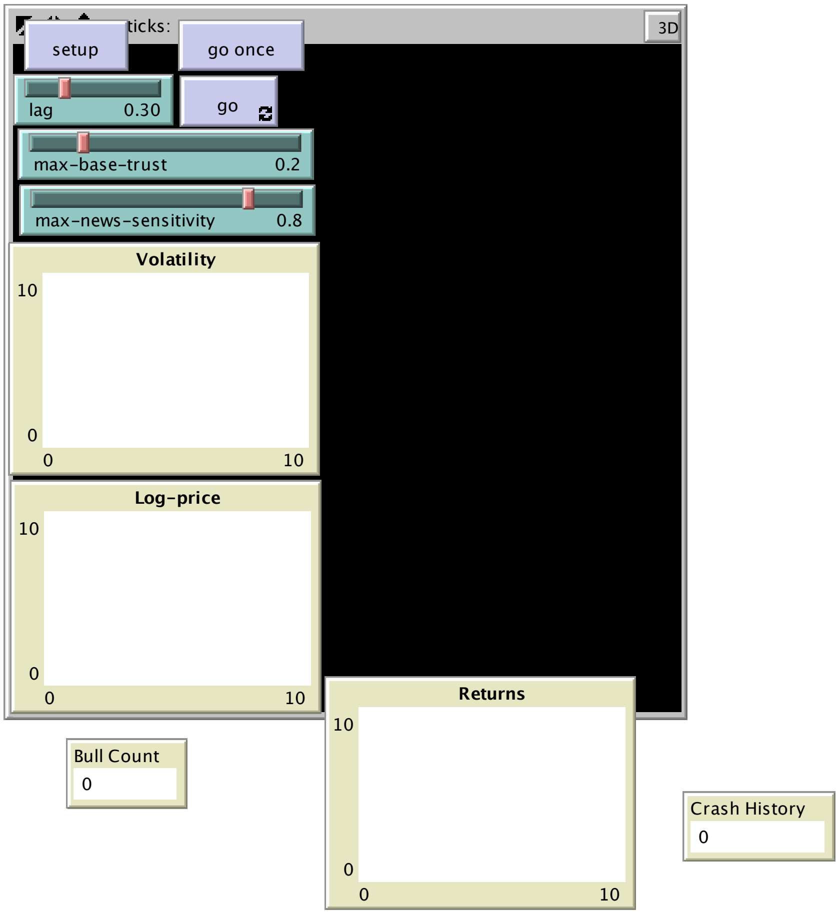

Figure 1: The dashboard of NetLogo. Silders and plots are created by the user.

parameters, and then adjusted to test different research questions (Gonçalves, 2003). It has a multitude of applications, one of which involves examining the behavior of complex systems.

In the NetLogo world, "turtles" exist on a two-dimensional plane; a turtle is the general name for an agent in the NetLogo code. Once the underlying characteristics of the agents are defined, one can create a framework for how any given set of agents will interact with each other and the NetLogo environment. In the current model, we attempt to parameterize these interactions to represent a market, with a function written in for buying and selling an asset, a function for the asset's price determination, and a mechanism for exogenous news to enter the market. The agents' individual specifications are written in such a way as to model how we believe investors behave, and the general structure of the agent does not change without our input. During each run of the model, agents buy and sell an asset, and we set up displays that allow the user to observe the price of the asset over time, the average returns and their volatility, which are all endogenously determined by the agents' behavior [1]. Through the interaction of the agents in the current model, one can observe these emergent properties that approximate the behavior of financial markets (Gonçalves, 2003).

Additionally, one can set other environmental parameters for the agents in the NetLogo world. We allow agents to make decisions within a short time period, and allow the model to run for a set number of time periods. We also specify how and when the model should cease to run; when a crash occurs, the model stops.

[1] Figure 1 shows the basic NetLogo display window. The black area is the NetLogo world, where agents exist and interact. The spaces for plots are identified as well.

Additionally, we specify how many agents exist and are actively making decisions in the NetLogo world. As with the agents' parameters, and the functional process of the model, these environmental specifications are subjectively determined. The shape of the NetLogo world, which determines how many agents can exist, and thus how many interactions there would be, is assumed to be symmetric. However we believe that, through simple observation of a series of runs, it becomes clear that the emergent properties of the system represent the behavior of financial markets, and that our environmental specifications are sufficient.

3.2 Goncalves' Model

The current model is adapted from Goncalves' model of financial market behavior. He originally constructed the model in an effort to examine the efficient markets hypothesis, herd behavior, and some explanation of why bubbles and crashes occur, as they did in his model (Gonçalves, 2003). The changes we have made are only meant to address some deficiencies in the identification of bubbles and crashes in the original specification. The theoretical underpinnings of our work are largely the same as that of Goncalves, however we do not attempt to examine the same questions as he does.

The main influences behind Goncalves' model and, in effect, ours, is the hazard rate model of crashes put forth by Johansen, Ledoit, and Sornette (JLS). Further clarified in a more recent paper, the JLS model examines how herding behavior can pose risks for the stability of financial markets (Johansen et al., 2008). Johansen et al. define a crash as a switch from one highly organized market state, the build, to another, the crash. They define a hazard rate as the risk of a crash occurring at the next step at any given point during a run of the model (Johansen et al., 2008). They test to see what factors influence hazard rate, and find that imitative behavior of noise traders, not smart investors, causes an increase in the hazard rate (Johansen et al., 2008).

Goncalves' model, and the model we use, do not make the same distinctions between non-professional and professional investors. It moves away from traditional dichotomies of noise vs. smart or informed vs. uninformed traders (Gonçalves, 2003). Instead, it makes the assumptions that: 1) traders are informed; 2) traders are rational in a limited manner; 3) traders have heterogeneous underlying preferences (Gonçalves, 2003). The assumption of limited rationality is included because it is assumed that agents both cognitively interpret information presented in a subjective manner, which varies across individuals, and form opinions about the nature of the world based on these interpretations (Gonçalves, 2003). By assuming this, the agent in this model is not like the standard neoclassical agent. We also assume all agents in the model have the same access to news presented every time period (Gonçalves, 2003). This differentiates the agents from traditional noise traders, or non-professional investors, which are traders that are not informed and tend to buy or sell in a manner not consistent with how a rational investor would trade (Shleifer, 2000).

Goncalves' model further differentiates itself from the JLS model by assuming that coupling, or trust between individuals, varies from agent to agent and is endogenously determined through agents' interactions (Gonçalves, 2003). The social communication among agents that are near each other is an important component of the information gathered by each agent when forming a sentiment about the asset. This form of communication may be desirable to the investor who is rational in a limited manner, as other investors may have valuable experience or insight as to whether or not the agent should buy or sell. Agents also have differences in the sensitivity to other agents' forecasts, i.e. carrying levels of trust, which, when varied for all agents, may cause changes in the observed market behavior.

The way trust is determined depends on an assumption made about the social structure of the market. That is, which agents influence other agents' decisions depends on how they are connected. The network of traders exists on a two-dimensional lattice, and each agent, as represented by a point in the lattice, is connected to his four bordering neighbors (Gonçalves, 2003). The restrictions placed on the size of the lattice determine how many agents there are, and the boundaries where agents necessarily cannot be connected to four bordering neighbors. For instance, an agent occupying a corner space will only be connected to two other agents, which both occupy spaces along perpendicular edges of the NetLogo world.

Next, we describe the current model in detail. It is important to note that agents are not affixed with any gender in our model, or many other models used to explore questions about financial markets. We include a heterogenous term that accounts for specific individual differences in how agents make decisions, and we assume that this controls for any gender based differences. However, this may not be the case, as the general decision making process may not be constant for all investors in the real world. From here on, we refer to agents as being male. We start with a description of the agents' parameters. Then, we lay out the how the model can produce emergent properties that are representative of a financial market, focusing on an agent's decision rule, and the resulting effects. Finally, we define a crash and explain how the model searches for patterns that mimic would fit the crash criteria.

4 The Current Model: Agent's Parameters

4.1 My-sentiment ($Si(t)$)

Each agent in the NetLogo world is considered an informed participant in the market, and has an interest in his own financial performance (Gonçalves, 2003). In order to make decisions about how to interact with his neighbors, an agent forms a belief about the direction of the asset's price. This variable is representative of that belief.

The process through which this value is determined is given by the following equation:

$$Si(t) = sgn(lag \times Si(t) + Ki \times NSi(t) + (nsi \times Q(t) + ei(t)))$$

If the outcome of this decision is greater than 0, My-sentiment is set to 1, and the agent will be bullish, expecting the asset price to rise. As such, he will buy one share of asset and his color will be set to green on the display. If the outcome of an agent's decision is less than or equal to 0, My-sentiment will be set to -1, and the agent will be bearish, expecting the asset price to fall. In this case, he will sell the one share of asset and his color will be set to red on the display.

It is important to point out that, in the agent's decision equation, there is a variable that represents his previous sentiment towards the asset. This is meant to take into account the agent's former market behavior. So, if in one time period, an agent determines his sentiment is bullish, in the next time period, he will remember that he was previously bullish. This will influence the formation of his new outlook.

4.2 Opinion-vol ($ei(t)$)

This variable is included to aid in the representation of the heterogeneous preferences held by agents. It is allowed to float along a uniform distribution, with the maximum allowable value set by the user, for each agent in the NetLogo world. It is set at the beginning of each run, based on the aforementioned distribution, and is fixed for the duration of the run.

Instead of the initial value of opinion-vol being set for each agent, it acts as the standard deviation for a normal distribution with a mean of zero. A value is then selected at random from this distribution for each agent. The result is that agents preferences vary along a normal distribution, not a uniform one, and the standard deviation of each agent's distribution function is the same. The reason for this roundabout specification is due to the nature of the NetLogo programming language.

We add a scalar to this variable in an effort to increase its importance in the agents decision. The justification for doing so is that our change better represents the bounded-rationality condition that Goncalves sought to include in the model. Another way to interpret this variable is the underlying outlook an agent brings to the market, whether or not he is inherently bullish or bearish, optimistic or pessimistic, regardless of market behavior.

There is substantial literature regarding the robust nature of these outlooks (White et al., 2003). Individuals experiences and psychological framing of those experiences cause people to adopt certain world-views and outlooks about the future. Once adopted, these are not as susceptible to change compared to their views about particular events, and they persist for long periods of time (White et al., 2003). Kahneman and Tversky, along with others, put forth the idea of several cognitive heuristics, or short cuts, individuals employ when faced with decisions under uncertainty (Kahneman, 2003). This variable could represent the anchor, or initial piece of information used, in the anchoring heuristic. An anchor is an implicitly determined or suggested reference point, from which one can make decisions about the probability of an event (Kahneman, 2003).

It is important to note that this variable is not meant to test any hypotheses about the nature of cognitive heuristics. Instead, it is included because we believe the theoretical explanations for its inclusion are sound, and Goncalves original specification did not adequately account for the importance of this factor. Since it is a key variable in determining the actions of agents in this model, further research about the ideal specification of this variable is warranted.

4.3 Trust (Ki)

This variable is representative of the susceptibility of an agents trading decision to those of his nearest neighbors. This is the primary behavioral variable of the model explored in this study. Its inclusion is what makes the model similar to the JLS model in that it allows for contagion throughout the NetLogo world to take hold.

The value of this variable is initially set to base-trust (defined below), and then changes at successive time period. The reinforcement mechanism involved in this process is described in the following section. Once an agents trust value is determined, it acts as a scalar to the sum of the sentiments formed by his nearest neighbors. In this way, each agents sentiment influences, and is influenced by other agents outlooks on the asset price. This interconnectedness allows us to test whether or not this behavioral variable affects market behavior, and to what degree.

4.4 Base-trust

This parameter is used as a reference point to set the value of the trust variable for each agent. Like the heterogeneous preference variable, base-trust is allowed to float randomly along a uniform distribution, with the maximum allowable value set by the user on each run (Gonçalves, 2003). The user can adjust the base-trust parameter via a slider, and its value can range continuously from 0 to 1. When the user sets up each run, agents are assigned a base-trust value that persists for that run's duration. For the first time period of each simulation, agents trust values are set to their initial base-trust values, and then vary based on a spin-glass type mechanism (Johansen et al., 2008; Gonçalves, 2003).

We use the base-trust parameter to test for the effects of trust on the behavior of the market. Since its value is randomly determined on a uniform distribution, a higher value of base-trust would allow for a greater number of agents to have high base-trust. Conversely, with a lower value of base-trust, a smaller number of agents with high trust would be allowed. It is this key assumption that lets us to draw inferences about the market behavior by adjusting the base-trust parameter.[2]

4.5 News-sensitivity ($nsi(t)$)

This is another behavioral variable, originally included by Goncalves. Set by the user at the beginning of each run, it represents the importance an agent places on the tone of incoming news when making a decision. Like the other behavioral variables, it is allowed to float randomly on a uniform distribution, with the maximum allowable value set by the user on each run.

5 The Current Model: Global Parameters

5.1 Log-price, Returns, and Volatility

Log-price measures the price of the asset, expressed in percentage terms. We can interpret a log-price of 1.15 as being 15% greater than a log-price of 1.00. It is initially set to 0, and varies based on the actions of

[2]Figure 2 shows the display of a run in progress. The agents are represented by red and green squares in the NetLogo space. We can see that the market is currently in a disorganized state, and has been for some time. Note the Bull Count is 48, indicating that 48% of agents are bullish.

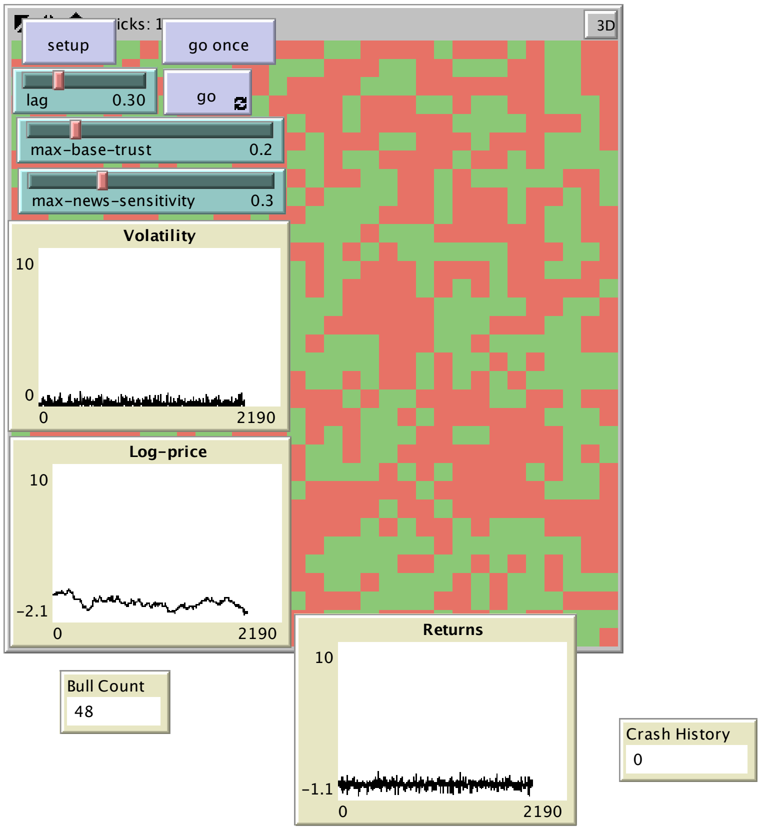

Figure 2: The NetLogo display. Max-base-trust is set to 0.20.

the agents during each run. By observing changes in the plot of log-price over time, we can approximate the average market sentiment, and witness trends in agent behavior. Importantly, the log-price is used in our crash definition.

Returns are the values of gains or losses realized, on average, by agents in the market. Aggregating each agents sentiments, which can take a value of either positive or negative 1, and then dividing by the total number of agents in the NetLogo world, determines the value of returns. This variable acts to determine the log-price at each time period through a market clearing mechanism described below.

The volatility term is meant to be an indicator of the level of returns. It is not meant to measure the standard deviation in returns, nor is it meant to act as a proxy for risk. Instead, the inclusion of this variable is largely for observational purposes. Volatility is set to be the absolute value of returns, thus is acts to show the observer when relatively large swings in the asset price occur.

By observing the changes in the volatility variable on the NetLogo display, the user can see how the model produces dynamics that are similar actual financial market behavior (Gonçalves, 2003). These include volatility clusters, jumps, correctional movements, and excess volatility Shleifer (2000). Some argue that the presence of these types of market behaviors conflicts with the behavior implied by the EMH, but we do not examine volatility in such a manner here.

5.2 News-meaning ($Q(t)$)

This is the valence, either positive or negative, of the news about the asset entering the market at each time period. It acts as a signal about the asset to the agents in the NetLogo world. It is assumed that news about the asset can either be good or bad. Furthermore, it is assumed that the probability distribution of the news is normal, with a mean of 0 and a standard deviation of 1. With this specification, any news value presented to the market that is greater than 0 is attributed a value of 1, otherwise it is set to -1.

When combined with an agents news-sensitivity, the news meaning will always have the same relative importance in the agents decision. As such, agents cannot discount the value of some news and place greater value on other news. This is likely a poor representation of how traders actually behave on markets, and there is some evidence that suggests interpretation of information can change with its context (Kahneman, 2003). However, to simplify this model, we do not allow for any changes in the relative importance of the news in an agents decision.

5.3 Crashes

Perhaps the most important global variable in the model, for the purposes of our inquiry, is a crash. We define a crash as a rapid onset of an organized state with a downward trend, and describe in detail how the model is programmed to search for this pattern of market behavior. After the model has been allowed to run for a specified number of time periods, it begins to search for crashes. If one occurs, that particular run of the model stops, and a crash is recorded.

The method used for finding crashes comes from another, more complex model that integrates financial market activity into the economy as a whole. The financial market in that model bears some resemblance to our's in its structure, a two-dimensional lattice, and how its agents forecasting parameters are specified (Gibson and Setterfield, 2013). Agents form opinions about the direction of an assets price based on relevant information and subjective, agent-specific factors. These similarities are important to be aware of, as they provide ample justification for using a similar crash specification. In both our model and Gibson & Setterfields model, crashes can be determined endogenously through agents interactions.

By examining historical patterns of the S&P 500, an index of 500 large-cap stocks commonly used to gauge financial market performance in the US, typical build and crash patterns were determined (Gibson and Setterfield, 2013). The build phase, a period generally consisting of steadily rising stock prices, typically lasted 200 weeks. Crashes occurred over shorter time periods after the build phases, lasting about 25 weeks. By plotting the two phases on a chart of price and time, one could see how the behavior fits a triangle pattern [3] [4].

To search for crashes, we create a matrix for prices starting from the current price, p0, and extending to the price 224 ticks prior to the current price, p224[5].This matrix updates at every week during each run, once the run has reached 1000 weeks, or about 19 years and 3 months of market activity. We then identify the build period, existing from p224 to p24[6]; this is denoted as Δp . The value of Δp is determined by the following:

$$\Delta p = p24 - p224$$

This parameter, Δp measures the change in price over the 200-week period, which we aim to represent as the build. It is possible that p24, meant to represent the height of the build, is not the highest price over

[3] Figure 3 graphically depicts what we expect the typical crash to look like

[4] For further discussion of this crash identification method see Gibson and Setterfield (2013)

[5] "Tick" is the formal NetLogo term for a time period in the model. One tick is recorded once all of the model's mechanisms (described below) have occurred. We consider each tick in this model to be equivalent to one week of financial market activity.

[6] p24 is the log-price 24 weeks prior to the current period.

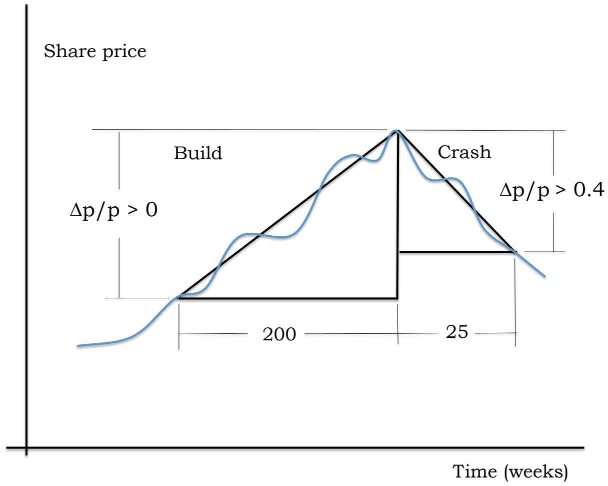

Figure 3: The triangle approach to identifying crashes.

the entire period. As such, this method is not well suited for identifying local maxima or minima in the build period. However, the following rule still applies:

- If $\Delta p \leq 0$, there has not been a build.
- If $\Delta p \geq 0$, a build of some degree has occurred.

This specification of Δp is designed to ignore a downward price trend that begins to fall more rapidly. In our examination, we are interested in whether or not a sizable build has occurred, and so we adjust the parameters for Δp to approximate the size of the build in the S&P 500 over the 200 weeks prior to the crash period.

We define the size of the drawdown in price from p24 to p0 as a crash, if one was to occur. We set the drawdown threshold as a 40% decrease in the assets price from the height of the build. By observing annualized return data for the S&P 500 one can see that the 40% decline is a rough estimate of the size of the recent crash in the 2008-2009 period. The result is a two-part definition for a crash in the current model.

- If 1) $\Delta p > 7.5\%$ [7] and 2) the drawdown $\geq 40\%$, then a crash has occurred.

The model will continually search for crashes defined this way until the end of each run, stopping the run if a crash should occur.

In the next sections, we summarize and demonstrate the processes involved in a typical run of the model. We then describe the methods used in our evaluation, present our findings, and discuss their implications.

6 The Current Model: Mechanisms

To begin, the user assigns values for max-base-trust and max-news-sensitivity by adjusting the sliders on the models display. Then the user clicks setup. This causes the model to assign opinion-vol, base-trust,

[7] This value was determined by looking at S&P 500 price data from November 12, 2004 - September 12, 2008. This would be considered the 200-week build period, and it is important to note that the value of the index on 12-Sep, 2008 (1251.69) is significantly lower than its high, reached on 12-Oct, 2007 (1561.80).

and news-sensitivity values to all the agents on the market in accordance with the specifications previously described.

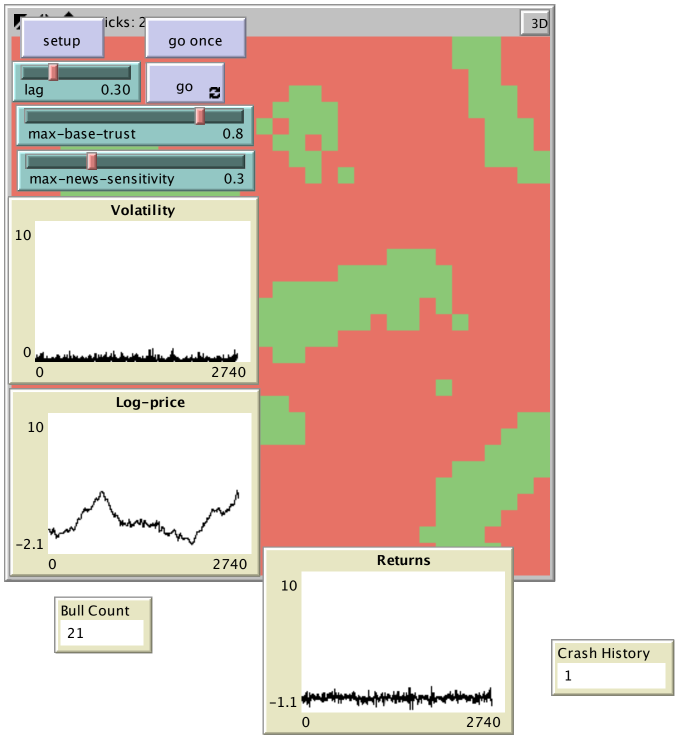

Figure 4: The NetLogo display immediately following a crash.

To start a run, the user clicks go. Immediately, one can observe the log-price, returns, and volatility charts begin to change. More importantly, one can see the dynamics of individual agents, as their colors change from green to red on the NetLogo display based on their sentiment. One can look at this display to determine when prolonged organized states occur, as they are obvious when they do; the system will display a somewhat constant pattern of green and red agents. It is interesting to note the shapes that organized states tend to take on during the run[8] [9]

At every time period, each of the following processes occurs in its corresponding order.

6.1 News Arrival

News enters the market. Its value is randomly selected from a normal distribution, with a mean of 0 and standard deviation of 1. If its initial value is greater than 0, the qualitative meaning of the news to all agents

[8]In Figure 4 we can see some small clusters of bullish agents surrounded by bearish agents. There are a few agents that seem to have strong opinions that outweigh the importance of their neighbors' opinions.

[9]Also note that the Bull Count is 21, indicating that 79% of agents were bearish just prior to the crash.

will be set to 1. Otherwise, it is set to -1.

6.2 Agent's Decision Rule

The agent interprets the news, discounts its importance, consults his neighbors, and forms an opinion about the nature of the asset that is either bullish or bearish.

$$Si(t) = sgn(lag \times Si(t) + Ki \times NSi(t) + (nsi \times Q(t) + ei(t)))$$

The agent weights the qualitative value of the news, previously determined to be either 1 or -1 by his news-sensitivity. This number is added to his opinion-vol, the term that represents his inherent outlook on the market, which is heterogeneous across agents. Then, the agent consults the sentiments of his 4 nearest neighbors, or those that share a border on the lattice with him. These sentiments are summed to obtain a value called sentiment of neighbors. This value is then discounted by the agents trust value, which is initially set to his base-trust. Finally, the value of his previous sentiment, weighted by a lag factor, is added. The result of this decision process is a number, often a decimal value. If it is greater than 0, the agents sentiment is set to 1, he buys one share of the asset that time period, and his color is set as green. Otherwise, his sentiment is set to -1, he sells one share of the asset, and his color is set as red.

6.3 Market Clearing

This calculates the average returns earned by agents in the current time period, and sets the price of the asset. First, all the sentiment values, either 1 or -1, of the agents in the market are summed then divided by the number of traders on the market. The quotient, noted as returns is then multiplied by a logarithmic factor, and added to the previous price. The result is a price for the asset, which is interpreted over time as percentage changes from the origin.

6.4 Updating the Market Sentiment

This process is key for the model to be consistent with the theory behind markets acting as organized states. We have previously described how agents trust parameters are formed, by assigning a value at random from a uniform distribution with the minimum value being 0 and the maximum value being controlled by the user. We also described how agents consider their neighbors sentiments about the asset when forming their own. However, we did not explain if agents trust parameters change over time, as they do in this model.

We assume that agents can learn from their behavior in some way, and will modify the importance they place on others opinions over time. We also assume that agents are aware of the returns they earn in each time period. So, we define a cognitive rule for agents trust dynamics, or the propensity to be influenced by others sentiments. We assume that, if good news about the asset reaches the market, and returns were positive, that is to say he was bullish when the price was increasing, then his trust parameter will be set equal to his base-trust plus his returns. Similarly, if bad news enters the market, and the agent correctly predicted the market movement, having a bearish outlook, his trust parameter will be set equal to his base-trust minus returns (which will be negative in this case). In both of these cases, when a market movement confirms an agents sentiment, the trust in his neighbors opinions increases, and he will value their input more. Conversely, if the agent is bullish when returns are negative, or bearish when returns are positive, trust in his neighbors sentiments diminishes.

- If returns > 0 and news-meaning > 0, then trust = base-trust + returns
- If returns > 0 and news-meaning < 0, then trust = base-trust returns
- If returns < 0 and news-meaning < 0, then trust = base-trust returns
- If returns < 0 and news-meaning > 0, then trust = base-trust + returns

With this mechanism in place, organized states can arise. A spin-glass type model can explain these states (Gonçalves, 2003). In a spin-glass model, some exogenous force is applied to a system at a particular time, which causes it to enter a highly organized state. The system is then slow to recover from the organized state, even when the exogenous force is removed. Many proposed behavioral aspects of financial markets are well explained by a spin-glass type model, including herding behavior. When agents sentiments are coordinated, that is, when trust is high, the system is more prone to entering organized states. Agents sentiments can become coordinated to some degree if appropriate market movements confirm enough of their forecasts.

Importantly, an agents trust parameter is bounded by his base-trust, and at every time period he refers back to his base-trust and either adds or subtracts returns to form a new trust value. If we allowed the trust parameter to gather momentum in another way, we would likely see trends that build quickly and do not ever resolve. Moreover, we think the current specification better approximates how investors behave in the real world. Like the heterogeneous term, people may have a finite range for the level of importance they place on outside opinions, and allow themselves to trust others more or less based on experience. It is important to note, however, that the dynamics observed in the model are entirely dependent on specifications such as these, subjective as they are.

By observing price dynamics, a user can see large builds or declines, that sometimes do not resolve (Gonçalves, 2003). When they do, the system will either become disorganized, or switch to an organized state moving in the opposite direction. We are interested in the periods where the system switches from one organized state to another, as this is how we have defined crashes. We anticipate these switches will occur when agents trust in each others sentiments is high, and a series of bad news about the asset enters the market. If our anticipations are confirmed, then we may be able to say something about the rationality of the outcome, as determined by the choices of agents that have limited rationality.

7 Methods: Testing for Behavioral Effects on Financial Markets

In this study, we are primarily concerned with how agents trust parameters affect market behavior, insofar that it does. We hypothesize that high trust levels cause crashes. This is based on others findings about how correlated agent behavior can cause crashes, as well as the structure and specifications of the current model (Gonçalves, 2003).

To test this hypothesis, we examine financial market activity over an extremely long period of time, at various base-trust levels. We create three separate conditions of financial markets, one with low trust levels (max-base-trust = 0.2), one with moderate trust levels (max-base-trust = 0.5), and one with high trust levels (max-base-trust = 0.8). The low trust condition serves as our control, and we seek to examine the effects of moderate and high trust conditions on crash frequency.

We believe the low trust condition serves as an adequate control for several reasons. First, Goncalves finds that crashes are statistically rare events, and the price, returns, and volatility dynamics are most representative of an actual financial market when max-base-trust is set around 0.19 (Gonçalves, 2003). Second, we believe that coupling does indeed occur in markets to some extent. This is supported by findings of clustered return volatility and price movements that cannot be adequately explained by directly relevant information (Sewell, 2012; Shiller, 1981). Furthermore, intuition about the behavioral aspects of markets would suggest that traders share information, placing some level of importance on each others sentiments (Kahneman, 2003; Baumeister et al., 2001).

Next, we set the max-news-sensitivity to 0.3. This level is suggested to be representative of actual financial market functioning (Gonçalves, 2003). We do not vary this parameter in any of the runs, as we are not examining this behavioral aspect of market activity in this study.

Then, we set the duration of agents' interactions for each run. We choose 3750 weeks, or about 72 years. Thus, if the market reaches 3750 weeks without crashes, the run is considered over, no crashes are recorded, and a new run begins.

7.1 How Frequent Are Crashes?

We collected data over 150 runs for each trust condition for a total of 450 runs across all conditions. In total, if there were no crashes in any of the runs, we would expect 1,237,500 weeks (approximately 23,798 years) of financial market activity, after truncating the data set. We omit the first 1000 weeks from each run to control for differences in initial conditions, a common practice when examining the behavior of complex systems over long periods of time (Gibson and Setterfield, 2013). Thus, if a crash had not occurred, a typical run of our final dataset would consist of 2750 weeks.

The average duration of a run in our model is reported in Table1, along with the shortest run (1001 weeks), and the longest run (3750 weeks, the maximum allowable duration). N is the total number of weeks across all runs, and we can interpret this to mean that our data displays fewer weeks than anticipated. Since some runs did not extend to the maximum duration, crashes must have occurred.

Table 1: Average Duration of Runs and Total Weeks

Variable	**Mean**	**Std. Dev.**	**Min.**	**Max.**
weeks	2257.747	789.36	1001	3750
N		986664		

In Table 2, we show the total number of crashes that occurred across all runs in our data set.

Table 2: Crash Summary

Variable	**Mean**	**Std. Dev.**
crash_count	1	0
N	173	

Table 3: Crashes Per Condition

Condition	Low Trust	Moderate Trust	High Trust
Crash Count	10	67	96

We observed a total of 173 crashes, a large, and unusual finding. Traditionally, crashes are considered to be statistically rare events in normally functioning markets (Gonçalves, 2003). We parse the crashes into their different base-trust conditions in Table 3. Here we can see that in the low-trust condition, there were only 10 crashes, and crashes become considerably more frequent in the moderate trust condition (67 crashes) and the high-trust condition (96 crashes). This gives us some confidence that our control group, the low-trust condition, is operating properly.

7.2 Does Trust Affect Financial Market Behavior?

Using this data set, we extend a panel using "runs" as our ID variable and "weeks" as our time variable. This allows us to examine each run as a single data point, as well as examine heterogeneous differences among runs. Moreover, since we are looking longitudinally, any important omitted variables that do not vary over time, such as news-sensitivity, are accounted for.

To see if trust affected market behavior, we run a regression of base-trust on crash-count, our indicator for crashes. The results are displayed in Table 4. We can see that there is a positive coefficient on base-trust, and it is highly significant (β= 0.00112, $p < 0.001$). This would seem to suggest that higher levels of trust did indeed cause more crashes to occur, confirming our hypothesis.

We create dummy variables for the different trust conditions. We are careful to avoid the "dummy variable trap", or a case when dummies for all the conditions being tested are specified. Failure to do this would result in the dummies having a perfect correlation with the dependent variable (Stock and Watson,

Table 4: Trust's Effect on Crashes

	(1)
	crash_count
base_trust	0.00112***
	(10.44)
Constant	-0.000170***
	(-4.54)
Observations	986664

t statistics in parentheses
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

2010). As such, the maximum number of dummies we can include is one less than the total number of trust conditions. So we create a dummy for moderate trust, which is given a value of 1 if base-trust = 0.5, and is 0 otherwise. We also create a dummy for high trust, which is given a value of 1 if base-trust = 0.8, and is 0 otherwise.

Table 5: Effects of Trust on Crashes

	(1)	(2)
	crash_count	crash_count
base_trust	0.00112***	
	(0.000107)	
moderate_trust_dummy		0.000409***
		(0.0000582)
high_trust_dummy		0.000667***
		(0.0000649)
Constant	-0.000170***	0.0000301**
	(0.0000374)	(0.00000929)
Observations	986664	986664

Robust standard errors in parentheses
Moderate-trust levels were 0.5, high-trust levels were 0.8. Otherwise, max-base-trust was 0.2.
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

We regress these dummies on our crash indicator, and the results are shown in Table 5. Here, we can see that both the coefficient on the moderate trust dummy ($\beta 1 = 0.000409$, p < 0.001) and the high trust dummy ($\beta 2 = 0.000667$, p < 0.001) are positive and highly significant. This reinforces our hypothesis, that higher levels of trust cause more crashes, with the highest levels of trust causing the greatest number of crashes.

7.3 Does Trust Affect the Probability of Crashes?

In our examination of market behavior, we have focused on a particular market phenomenon, a crash. Importantly, our two-part definition for a crash incorporates a build period that will always precede the crash, should one occur. However, we observe a binary outcome; either there was a crash, as it is defined, or there was not.

Since our dependent variable is binary, and not continuous, we must interpret our regressions differently. These types of regressions generally model the probability that the dependent variable equals 1 (Stock and

Table 6: Linear and Non-Linear Probability Models using Maximum Likelihood Estimators

	LPM	Probit	Logit
main			
moderate_trust_dummy	0.000409***	0.533***	2.146***
	(0.0000582)	(0.0806)	(0.339)
high_trust_dummy	0.000667***	0.684***	2.708***
	(0.0000649)	(0.0790)	(0.332)
Constant	0.0000301**	-4.060***	-10.61***
	(0.00000929)	(0.0738)	(0.316)
Observations	986664	986664	986664

Robust standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Watson, 2010). Instead of viewing the trust coefficients as the influence on crash frequency, we can interpret them as the change in the probability that a crash occurred associated with a change in the maximum base-trust level. The regression estimate is considered to be a linear probability model (LPM) in this case (Stock and Watson, 2010).

Furthermore, we use a maximum likelihood estimation to answer this question. This process models a likelihood function, which acts as the joint probability distribution of the data, treated as a function of the unknown coefficients (Stock and Watson, 2010). The maximum likelihood estimator (MLE) approximates the coefficients on our independent variables that would maximize this likelihood function. As such, we would interpret the MLE coefficients on trust to be the values that maximize the probability of drawing the distribution of crashes that are actually observed (Stock and Watson, 2010).

We use two maximum likelihood methods, probit and logit, to test the robustness of our findings in the LPM. These tests are nonlinear regressions, specifically designed for binary variables, like our crash indicator (Stock and Watson, 2010). Since the predicted value of a crash can only lie between 0 and 1, a nonlinear is warranted. The probit model uses a cumulative standard normal distribution function, while the logit model uses a cumulative standard logistic distribution function. We report our findings in Table 6.

Our findings are robust in both non-linear probability models. We find that the moderate and high trust dummies are both positive and significant. We interpret the coefficients to mean that, when trust is higher, the market is significantly more likely to reach criticality and crash than at lower trust levels [10]. This makes intuitive sense, as we know that crashes are endogenously determined through agent behavior. Since we have allowed for the possibility of herding behavior, and then manipulated the potential for herding behavior to occur, by varying the max base-trust level, we were able to observe how changes in agents' propensity to value their neighbors' opinions affected the behavior of the market system as a whole.

8 Discussion

Given our results, we find that we can reject the null hypothesis that trust had no effect on crash frequency. Clearly, higher trust levels were associated with a greater number of crashes. This finding persisted when we parsed trust levels into their three categories, low, moderate, and high. By doing this, we also showed that the largest number of crashes occurred when max base-trust was the highest. We went on to model the joint-probability distribution functions of trust and crashes. Through maximum likelihood estimation, we found that there was a significant increase in the probability of a crash when base-trust was increased.

[10] The size of the coefficients on the moderate and high-trust dummies are difficult to interpret in the profit and logic specifications. This is because they examine the joint-probability distribution of trust and crashes. However, we are mainly concerned with the sign, not necessarily the size of these coefficients.

Taken together, our findings show that, in our model, trust has a significant impact on agent behavior, and increases the risk that, through their interaction, agents create market outcomes that model bubbles and crashes. This provides further empirical support of the idea that correlated behavior in the market can be detrimental to its performance. By increasing the degree to which investors act in a similar manner, based on their intrinsic inability to evaluate an asset without consulting others, the market becomes more prone to enter organized states, and switch between them rapidly.

We now focus our attention on the limits of this model, and thus the robustness of our findings. First, as is often the case with agent-based modeling, the stylization of agent's characteristics is entirely controlled by the user. Goncalves' specification, and the JLS model from which it draws its structure, both make critical assumptions about the nature of the agents and how they interact. In our model, we retain the assumptions made by Goncalves, and add a few of our own. One major assumption is that individuals are rational in a limited manner (Gonçalves, 2003). This applies to all the agents in our model, and so does not represent the classic dichotomy of rational and irrational investors (Gonçalves, 2003). But the limited nature of rationality is not well defined, and we take it to mean that individuals cannot interpret information in a vacuum, and necessarily bring their own idiosyncratic biases to bear on their decisions. While there is psychological literature to support this (Kahneman, 2003; Baumeister et al., 2001), we cannot be sure about the actual tendencies of the financially active population we are trying to model. Moreover, we are unsure as to how these tendencies vary across the population of actual investors operating in the market.

The way we model limited rationality is through the news interpretation and contagion parameters. Certainly, there may be more factors involved in an actual investor's sentiment formation process. When modeling how base-trust and news-sensitivity are distributed across agents in the market, we code for a random selection process along a uniform distribution, with the minimum being 0 and the maximum being controlled by the user. We assume that these parameters are uniformly distributed. In fact, many human traits tend to be normally distributed. Thinking anecdotally, it may be the case that there is a tendency for the importance investors place on their friends' opinions, and the importance they place on the news, to vary along a normal distribution. If we were to change how these parameters were determined in our model, with the mean and standard deviation being controlled by the user, we might observe some critical differences in our observations. However, we think that trust would still be positively related to crashes, and that a uniform distribution may actually underestimate the effects of trust on crashes. This is because, on a normal distribution, agents with base-trust and news-sensitivity values that are very high or low, relative to the mean, would be less likely to be present on the market. Still, further testing with these parameters could be done to determine if a normal distribution of trust and news sensitivity would better represent individuals trading on a real financial market.

In this model, we did not vary the news-sensitivity parameter. We did this for two reasons: one was to restrict our examination to the effects of trust on crashes, and the other was because of how the news interpretation mechanism is specified. Looking at the agent's decision rule, we can see that, for some agents, the qualitative value placed on the news through interpretation can always exceed the trust and heterogeneous parameters. Since agents place a weight, as determined by their trust value, on their neighbors decisions, and the maximum number of neighbors an agent can have is 4, there can arise a situation where the trust-weighted sentiment takes a value of 0. Consider an agent who has two neighbors that are bullish, and two that are bearish. In this case his cumulative, trust-weighted sentiment would be 0; only the news and his own market outlook would determine his decision. When our model enters organized states, we can observe many agents like this as different patterns emerge; they tend to lie on the border that exists between bullish and bearish groups of agents. However, by not varying news-sensitivity along with base-trust in some way, we may have overestimated the effect that trust, on its own, has crash frequency and the probability of crashes. Further research can be done to determine how news-sensitivity and trust might be varied together a priori. A $3x3$ model of low, moderate, and high base-trust and news-sensitivity immediately comes to mind, and, in future studies, it may be worthwhile to see if we did indeed overestimate the effects of trust on crash frequency and probability by failing to vary news-sensitivity with base-trust in some way.

Another issue with the news interpretation mechanism is that agents weight gains and losses in a uniform manner. Literature on prospect theory and other psychological aspects of decision makers shows that, on average, people tend to be loss averse, but more so than they are comfortable with seeking gains (Kahneman, 2003). We do have a process that is meant to take this phenomenon into account, somewhat; it is the way agents update their market sentiment. By looking at their returns, agents adjust their trust levels based on whether or not they correctly predicted the market movement. However, when agents make an incorrect prediction, such as when an agent is bullish but average returns are negative, they may not adjust their trust levels anymore than they would have if they had correctly predicted the market. We might observe changes in our findings if we add a scalar to negative changes in agents' trust levels.

Importantly, there are no costs for agents to participate in their trading activities. Furthermore, the only real costs faced by agents in this model are when he incorrectly predicts a market movement. Regardless of their previous performance, agents trade at every time period. This may not be representative of an actual market for several reasons. For one, investors face multiple investment opportunities in the real world, and in the face of consistently poor performance on the stock market, might seek to invest else where. For another, it is clear that trading costs affect the real returns earned by individuals (Sharpe, 1991). Instead of actively trading each week, investors might buy and hold an asset, or a security that represents an index of assets, for many weeks, even years. While we think our current specification of how agents trade, buying or selling one asset every week, is a reasonable assumption, it could be argued that it does not adequately differentiate between high and low degrees of bullish or bearish sentiments. Future examination of crashes with this model may benefit from such differentiation if it modifies how much agents trade over a given period of time.

Finally, our model may not be well suited to evaluate the behavior of markets where more than one asset is traded. In our model, agents trade a single asset, receiving and interpreting news and opinions only about that asset. However, we model crashes with historical parameters that were obtained about the S&P 500 index. This would not be a problem if we thought that investors only traded securities that tracked the entire index in a uniform manner, but in reality this does not occur, and there is large variation in the volume and volatility of the stocks that make up the index (Shiller, 1981). We assume then, that the asset being traded in our model is either perfectly correlated with the S&P500, or is in fact a security that tracks its performance. In this way, the generalizability of our findings is limited further. We think that the model may be better suited to examine the market behavior of a highly stylized sector. If we were to think about the model in this way, some of the assumptions begin to seem more realistic. For instance, it is reasonable to expect that investors can all be equally aware of news about a single company they are all interested in. In this case, it would also be reasonable to assume that investors cannot know everything about the company, and thus would be interested in seeking the opinions of their neighbors when making their own decision. This might lead us to change our crash parameter values to better suit the historical performance of such markets.

When we try to model a financial market, we always face the trade off between better specification and generalizability. There will always be issues with the subjectivity of parameter definitions in modeling social phenomena. However, our findings still offer insight into how the cognitive limitations of individuals influence market behavior.

References

Barberis, N. and A. Shleifer (2003). Style investing. *Journal of Financial Economics 68*, 161–199.

Baumeister, R. F., B. E., C. Finkenaur, and K. D. Vohs (2001). Bad is stronger than good. *Review of General Psychology 5*(4), 323–370.

Fama, E. (1965). The behavior of stock market prices. *Journal of Business* (34-105).

Ghashghaie, S. (1996). Turbulent cascades in foreign exchange markets. *Nature 381*(6585), 767–770.

Gibson, B. and M. Setterfield (2013, January). Real and financial crises.

Gonçalves, C. P. (2003, December). Artificial financial market model.

Johansen, A., O. Ledoit, and D. Sornette (2008). Crashes as critical points.

Kahneman, D. (2003). Maps of bounded rationality:. *The American Economic Review*.

Samuelson, P. (1965). Proof that properly anticipated prices fluctuate. *Industrial Management Review*, 41–49.

Sewell, M. (2012). Behavioral finance. Technical report, University of Cambridge.

Sharpe, W. F. (1991, January). The arithmetic of active management. *The Financial Analysts' Journal 47*(1), 7–9.

Shiller, R. (1981). Do stock prices move too much to be justified by subsequent changes in dividends? *The American Economic Review 71*(3), 421–436.

Shleifer, A. (2000, April). *Inefficient Markets: An Introduction to Behavioral Finance.* Oxford University Press.

Slovic, P. (1993). Perceived risk, trust, and democracy. *Risk Analysis 13*(6), 675–682.

Sornette, D. (2003). *Why Stock Markets Crash.* Princeton University Press.

Stein, J. C. (2009, August). Presidential address: Sophisticated investors and market efficiency. *The Journal of FInance 64*(4), 1517–1584.

Stock, J. H. and M. W. Watson (2010). *Introduction to Econometrics (3rd Edition)*. Addison-Wesley.

White, M. P., S. Pahl, M. J. Buehner, and A. Haye (2003). Trust in risky messages: The role of prior beliefs. *Risk Analysis 23*(4), 717–726.