Bard

Senior Projects Spring 2011   Bard Undergraduate Senior Projects

Spring 2011

Reflexivity in Financial Markets: A Neuroeconomic Examination of Uncertainty and Cognition in Financial Markets

Steven Pikelny
Bard College, Spike716@gmail.com

Follow this and additional works at: https://digitalcommons.bard.edu/senproj_s2011

Part of the Banking and Finance Law Commons, Behavioral Economics Commons, Behavioral Neurobiology Commons, Cognitive Psychology Commons, Computational Neuroscience Commons, Economic Theory Commons, Finance Commons, Finance and Financial Management Commons, Management Information Systems Commons, Social Psychology Commons, and the Theory, Knowledge and Science Commons

Recommended Citation

Bard

Introduction

For financial markets to function optimally they must, at a fundamental level, interpret reality with some degree of accuracy. Although money, bonds, equities, derivatives, and other financial instruments are abstract products with no value in and of themselves, they do not exist in a vacuum; their value is ultimately determined – in some form or another – by sections of the real economy. In the absence of a mechanistic cause-and-effect relationship, this determination is indirect. Financial markets are distinct and separate from – though highly influential upon – markets involved with the production and sale of goods and services, and can only operate by processing all relevant information and translating it into financial transactions.

In his examination of the market process, Hayek (1948) considers how a central planner – a single cognizant agent or a group of cognizant agents – would organize an economy. Hayek notes, "*If* we possess all the relevant information, *if* we can start out from a given system of preferences, and *if* we command complete knowledge of available means, the problem which remains is purely one of logic" (77, emphasis in original). Hayek emphatically concludes that this would be an extremely difficult task, as the centralized authority does not possess nearly enough knowledge to decide upon the most efficient flow of capital. The central planner(s) can be benevolent and extraordinarily intelligent, but they lack critical information to make the most optimal decision.

If not through rational, centralized planning, how does an economy allocate its resources? Hayek points to the decentralized competitive process of markets. Markets are comprised of many cognizant, interacting agents, each with a mental model of the world and a hypothesis of how it will evolve. This system generates more accurate knowledge because agents directly experience the production process. They presumably have "unique information" of "the

1

particular circumstances of time and place," adding a higher degree of accuracy to their knowledge. Additionally the evolutionary process of competition ensures that, in the case of contradictory knowledge, the more accurate tend to survive on average.

Hayek implicitly makes the simplifying assumption that relevant information and knowledge is quantitative, as the problem becomes one of logic. Markets ultimately translate the qualitative observations and theories of agents into a price, but there is no systematic way to do this. This task provides the central planner with another difficulty, that of interpreting and processing information. Each agent faces this problem as well, but the market as a whole presumably has more computational power than does a committee of central planners. A market's processing capacity is much larger, as agents' spheres of knowledge interact with each other, influencing market prices. Markets allocate resources by interpreting reality through a bottom-up process, in which individual agents interpret the world and subsequently interact with other agents to create an emergent order. None of the agents, however, can comprehend the system in its entirety. Viewed in the context of information processing systems, both the planner(s) and markets undertake computational work in translating information into economic transactions; but the former does this in a centralized manner, and the latter in a decentralized manner.

Financial markets, however, add another layer of complexity to Hayek's story; investors are not necessarily concerned with the most efficient allocation of capital, but rather what other investors think the most efficient allocation of capital is. The Keynesian beauty contest thought experiment elucidates this problem: a newspaper runs a beauty contest where its judges are told, not to pick the prettiest girl in a line-up (as in a traditional contest), but which contestant they think the other judges will choose as the prettiest. It is possible – and in many cases likely – that

the winner of this contest would differ from the winner of a traditional beauty contest. But even in this thought experiment the outcome is in some sense tied to reality. Consider the contest from the viewpoint of a judge, J. J's answer will depend on the answers of the other judges, which are informed by the other judges opinions of the other judges' answers. To make a decision, J must possess a *theory of mind* – that is, a theory about the workings of other judges' thought processes. If J suspects that its own cognitive makeup is similar to – but different from – the other judges, then J can assume that the other judges go through a similar evaluative process. J might (implicitly) assume that all the judges possess similar mechanisms for processing visual stimuli. This assumption alone means that the complex game-theoretic reasoning of the judges is not completely isolated from the actual contestants. The bets are directly determined by the expectation of other bets, rather than reality itself; but these bets are, at least partially, influenced by information coming from reality.

Put simply, it is unlikely that a llama in a dress would win either contest. Given the contest's rules, it is certainly possible that a llama could win, but if judges are all human (and possess human-centric theories of mind) J might hold the slightly more controversial assumption that they all possess the same basic standards of beauty (i.e., facial symmetry, the appearance of vitality, etc.). This second assumption, known in the Game Theory literature as a "focal point," expresses a belief held by all agents in a game about all other agents' beliefs without the need for communication. Although the judges' individual theories of mind might differ from each other, the collectively perceived overlap would act as a focal point. J can therefore assume that its interpretation of reality is at least partially shared by the other judges.

The inclusion of reality into this model ensures that the system is an open one, rather than a closed system of feedback loops. If the judges were instead told to pick a number between one

and ten, they would focus entirely on their expectations of the other judges. While convoluted, chaotic, and perhaps unpredictable, this system would be deterministic. Given an initial set of conditions, the judges should theoretically come to the same conclusion every time. Adding an exogenous factor ensures that the system is, at least in part, determined by forces external to that system. The important implication for financial markets is that, while its movements are largely determined by internal forces, reality acts as a focal point. Although the same piece of information can be interpreted differently by different investors, a focal point exists as to the information's meaning.

It is interesting that Keynes chose a beauty contest to make his point because there are, after all, no objective standards of beauty. Notions of beauty differ between individuals and across cultures; there is no systematic way to determine what is beautiful and what is ugly. The biggest difference between the Keynesian beauty contest and financial markets is that the former is based on opinion, while the latter is based on the potential occurrence of an event. Regardless, the interpretation and prediction of highly complex systems is similar to determining beauty in that it is far from concrete. It is difficult to say that one investor's mental model of financial markets is more accurate than another's.

In lieu of a systematic way to discern future events through given information, mental models can only be formed through induction, rather than deduction. Inductive principles are arrived through trial and error, rather than step-by-step, unambiguous logic. Deductive logic can play a large role in determining an investment strategy, but this method only works when given complete information. The present value of a financial security, for example, is derived by taking the sum of the expected future discounted values of the security's cash flow, plus the discounted expected sale price. This tool can price any security with complete accuracy; the trouble comes

when one must pick the values of all future cash flows, the future sale price, and the discount rates. There is no unambiguous way to find the quantitative values of these three variables, so investors can only make educated guesses based on the information available to them.

Since cognition plays a central role in the context of interpreting information, it is important to examine the cognitive in the context of financial markets. Economic agents must, in the following order, observe reality, formulate an understanding of it, develop a prediction or set of predictions, devise a reaction, and implement its reaction. Investors, more specifically, must acquire information about the current state of the economy, society, and financial markets; formulate a mental model; predict the direction of sections of the financial market; devise an investing strategy based on this model; and implement the strategy. While the central planner(s) in Hayek's model are limited by the amount of knowledge they possess, the financial market as a whole is limited by investors' perception of reality. Human perception acts as the medium through which information must pass to become a price. If investors exhibit systematic biases in processing information, this will subsequently affect the efficiency of the emergent order.

Because of the ambiguity and lack of deductive reasoning in the investment process, it is difficult to say that an agent's perceptions or actions are "biased". This is a contentious issue in the financial academic literature, spurring debate between two main schools of thought. The efficient market hypothesis (EMH) states that prices fully reflect the risk-adjusted value of securities, given the available information about them. Conversely, proponents of behavioral finance (BF) argue that humans are subject to psychological biases in their decision making process, which subsequently make markets less efficient. EMH proponents counter that the biases cancel each other out on average, or that the mere presence of rational investors – who would notice market inefficiencies – would ensure that price discrepancies are arbitraged out of

existence; BF proponents respond that transaction costs are often too high to take advantage of arbitrage opportunities.

Both schools of thought are somewhat vague, however, in their notions of rationality and efficiency. The most specific definition of efficient prices given by the EMH is the price formulated by the "agreed upon" interpretation of information. This, however, is a tautology. Any conceivable interpretation of a set of information would lead to efficient prices by this definition, no matter how accurate the risk associated with them proved to be. The vagueness of the theory protects it from empirical falsification, but, in doing so, greatly limits its descriptive and predictive power. Conversely, BF assumes that there exists an efficient equilibrium price for all securities, but investors systematically exhibit irrational psychological biases and obscure real market prices from these equilibrium values. However, this viewpoint assumes that "rational" financial decisions based on the given information exist – from which investors are deviating.

Hayek (1968) notes, "If anyone actually knew everything that economic theory designated as 'data,' competition would indeed be a highly wasteful method of securing adjustment to these facts" (9). However, no single cognitive entity possesses all of the relevant "data," nor would it possess a method for processing this data. Without a "rational" decision-making process for comparison, BF's claim that investors do not exhibit rational behavior does not hold much weight. Moreover, BF overlooks the processing of financial information at a market-wide level. Hayek argues that competition through markets acts as a "discovery procedure," where the system as a whole works towards a best guess of what prices should be with respect to available information. In an uncertain market environment, it is more useful to ask, not whether investors are rational or irrational, but whether investors demonstrate bounded

rationality.[1] The question with respect to market behavior becomes one of whether the market acts *as if* investors are boundedly rational, and whether market prices are as accurate as possible subject to available information *and* with respect to the time and energy used to process that information.

This question may not have a clear empirical answer. While it is easy to call the financial crisis of 2008 – with the advantage of perfect hindsight – an episode of market inefficiency, it is more difficult to say that a computationally efficient financial market could have reacted more appropriately, given the amount of information and processing power available in 2006. Since there is no perfect, computationally efficient mechanism to compare financial markets with, there is no way to reliably discern how efficient financial markets act at any given time. Testing for computational efficiency runs into the same problems as testing for regular efficiency, with the added difficulty of measuring computation of a social system.

Nonetheless, a careful analysis of this question can lead to useful conclusions for both investors and academic economists. I take a bottom-up approach towards formulating a hypothesis of market computational efficiency under conditions of uncertainty, first by examining the epistemological problems of the investor (with special attention paid to a case study describing how a group of investors copes with these problems), then by reviewing the neuroscience literature relating to risk and uncertainty, and finally by using philosophy of mind as a framework to interpret these observations.

I propose a model, the reflexive market hypothesis, in which investors attempt to fill their perceived gaps in information by using their perception of other investors as informational inputs. This strategy can lead towards two possible outcomes: in the first scenario, this reflexive

---

[1] Bounded rationality, a term coined by Herbert Simon, examines rationality with respect to the computational capacity of, and information held by, the agent in question. This will be discussed more extensively in chapter 5.

mechanism helps financial markets fully integrate and process information, and prices move towards their efficient value; in the other scenario, when a key piece of information is systematically overlooked, the mechanism inspires overconfidence in investors, and market "inefficiencies" ensue. While the EMH and BF describe markets in polarized states, one advantage of the reflexive market hypothesis is that it describes markets in all states between these extremes. More importantly, the reflexive market hypothesis addresses many of the critiques of the EMH, while also addressing certain methodological problems of BF.

My project is split into three sections and a conclusion. Section 1 is devoted towards laying the groundwork of financial analysis and describing the significant overlap between academic and investor analysis. This will involve outlining and clarifying the epistemological issues associated with investor decision making. I discuss the benefits and drawbacks of deductive and inductive reasoning, risk and uncertainty, and how analysis of financial phenomena not only descries, but alters financial markets. Sections 2 provides an overview of the two leading schools of thought related to market efficiency and investor rationality, the EMH and BF. I discuss the strengths and weaknesses of their respective methodologies, and how these methodologies influence their conclusions. Section 3 discusses the rising sub-discipline of neuroeconomics, with particular attention paid to how decisions are made in financial contexts and under uncertainty. I then discuss the implications and philosophical issues of this approach, and expand on it with the vantage point of cognitive science. Finally, a concluding section outlines an alternative to the EMH and BF, the Reflexive Market Hypothesis.

1. Investor Epistemology

Academia does not have a privileged perspective of financial markets. Investors and academic financial economists face the same epistemological problem; both must simplify an extraordinarily complex and uncertain system down to a model with predictive capabilities. The former intends to interact with financial markets while the latter intends to observe and categorize financial markets; as these goals are quite different, each party has adopted similar, but different, methodologies. The two parties should take interest in each other because they examine the same problem from opposite ends; the investor takes an inductive, implicit, bottom-up approach, while the academic takes a more deductive, explicit, top-down approach. As such, the following chapter views the questions of investor rationality and financial market efficiency from the internal perspective of investors, and highlights its overlap with the perspective of academic economists.

1.1 Mental and Mathematical Models

Investors have two primary tools for analyzing financial markets: deduction and induction. The former involves explicit step-by-step logic and is *a priori*; given an initial state (or set of assumptions) and desired outcome, deduction is the process that leads towards an objective conclusion. The latter, induction, is a process of trial and error; it involves observation, inference of conclusions based on the observation, construction of a hypothesis, noting whether further observations match the initial hypothesis, altering the initial hypothesis to match recent observations, and so forth.

Deduction provides a logical structure for the investment process, describing the objective relationship between variables. When looking to describe general relationships, the explication of a line of deductive reasoning leads to the formation of a model, defined by MacKenzie (2006) as "verbal and mathematical representations of markets or of economic processes" (6). Explication has the benefit of helping to maintain consistency, and allowing for communication. A basic model, for example, is the formula for the present value of a security, defined as the sum of all discounted payments plus the discounted sale price. Given each payment, each discount rate, and the sale price, the present value of a security can be found by merely manipulating the variables. Investors often use private models to clarify their own hypotheses, but academically-derived models, often altered, are the most widely propagated. For example, the Capital Asset Pricing Model (CAPM), Modern Portfolio Theory, and the Black-Scholes option-pricing model are used to find the necessary return for a security, the most efficient allocation in a portfolio, and the present value of an option contract, respectively.

These models can be interpreted in two ways. In one way, they are descriptive, positive theories – they can be used to predict the behavior of investors, and subsequently the behavior of

markets. In another way, they act as normative financial engineering; they describe the optimal portfolio allocation, or a benchmark for buying or selling securities. MacKenzie (2006), borrowing a term from economic sociology, describes these models as *performative*. Although these models were originally devised as descriptions of markets, they were performative in the sense that their formation altered the way markets function. MacKenzie continues to distinguish between two types of performativity. In *Barnesian performativity*[2], "practical use of an aspect of economics makes economic processes more like their depiction by economics." Conversely, in *counterperformativity*, the "practical use of an aspect of an aspect of economics makes economic processes less like their depiction by economics" (17).

Although difficult to observe empirically, MacKenzie notes the Black-Scholes model as one of the clearer examples of Barnesian performativity. When originally proposed the model intended to describe the behavior of option prices, but was only an "approximate fit." During the 1970s, however, option prices increasingly began to resemble the price predicted by Black-Scholes. MacKenzie attributes this to two factors: first, as option markets gained popularity, the market environment began to more closely resemble the assumptions made by the model; second, the model became a ubiquitous component of investors' trading strategies.

Evidently, the effect of the latter was significant, as the aftermath of the 1987 stock market crash saw the subsequent fit of the model diminish. A key component of the model is the probability of the option expiring in-the-money; although past volatility and a normal distribution were used as estimates for this value, they proved to be inadequate. The functional form of the model was not specifically violated, but the most commonly used device for determining a key input was. Nevertheless, MacKenzie notes, "Option theory has left its

---

[2] MacKenzie refers here to sociologist Barry Barnes, who he notes, "has emphasized ... the central role in social life of self-validating feedback loops" (19).

permanent imprint on the options markets: the theory is embedded in how participants talk and in technical devices that are essential to their markets. But what is performed in patterns of prices in those markets is no longer classic option-pricing theory" (33).

The utilization of the Black-Scholes model highlights a key flaw in relying on models too heavily. These formulas are useless on their own; they describe relationships between unknown, future variables. In lieu of a method to determine the future with certainty, expected values are used as substitutes. Expected values are a function of the security's risk, or the probabilities associated with certain events multiplied by the value of the events' respective outcome. High-risk securities, with lower probabilities of success, require higher yields as compensation. US treasury bills have low interest rates because they are perceived as riskless. Conversely, junk bonds (a.k.a. high-yield bonds) have relatively higher chances of defaulting – investors therefore require a higher risk premium to represent the probability of default.

However, risk is an ordinal concept, not a cardinal one. It is possible to say US Treasury Bills are less risky than junk bonds, but not by how much. It is easy to account for higher and lower risk (either increase or decrease a coefficient in expected value of a variable), but difficult to quantify amounts of risk. While conceptually very useful, defining variables in terms of risk is not as useful in practice; it is often nearly impossible to determine the explicit probabilities associated with different factors.

Knight (1921) famously made the important distinction between risk and uncertainty: *risk* refers to the concept used to quantify an unknown future in terms of probability; *uncertainty* refers to the concept used to describe an unknown future where probability cannot be determined. For example, betting on a 3 when rolling a six-sided die involves risk; there is an equal chance that the die will land on any of the six sides, so the probability of winning is about

0.1667. If one bets $10 on a dice roll with a $60 payoff for winning and a $0 payoff for losing, the expected value of the outcome ($10) is self-evident. Consider this bet proposed by a street hustler, as opposed to a casino employee. The former scenario is obviously the riskier one due to the increased probability of cheating, but it is difficult to determine the change in expected value of the bet. One can look at the cheating history of the street hustler[3] to determine the probability of a weighted die in the current wager, but this is an imperfect measure. Perhaps the hustler owes more money to a loan-shark than usual, and his preference to make money at present is much higher than in the past. Even if this piece of information is available, it is difficult to determine the internal state of the street hustler, let alone his current risk preference, because this scenario involves a grossly incomplete set of information; it is plagued with uncertainty.

Uncertainty is essentially a problem of incomplete information. In the example of the street hustler, uncertainty is only reduced as more information becomes available; moreover, it is difficult to imagine explicit probabilities even in this hypothetical scenario. Uncertainty is never completely reduced, as the problem becomes more complicated when it is examined more closely. An investor can reason that a stock's price is partially a function of the company's future earnings, which are determined by multiple qualitative factors, such as management quality, future political and legal environments, future consumer sentiment, potential competitors, etc. This investor can, hypothetically, create an intricate model of all these factors, but would run into two problems. The first problem being that it is significantly more difficult to determine the relationships between variables without mathematical values; the second problem being that an extraordinary amount of information would be required to give the model accurate predictive power – information would be needed down to the quantum level of detail before an explicit

---

[3] Assuming this information is somehow available.

probability of future events could be found.[4] If this were possible, however, the information problem introduced by Hayek (1948) would be solved, and markets would be unnecessary to determine prices.

Outside of a casino, it is difficult to find examples where probabilities are explicitly known, if they exist at all. Nevertheless, investors make decisions regarding unquantifiable risk with quantifiable levels of capital, and financial markets set prices in a highly uncertain environment. Various qualitative variables are somehow translated into quantifiable security prices, and while these prices may not perfectly correspond to risk, they are undoubtedly correct (or close to correct) more often than pure chance would predict.

But before aggregate financial markets can determine prices, investors must process the information and somehow draw conclusions from it. This requires, at the very least, a general idea as to what the objective probability of an event is. Arriving at the 'best guess' of unknown objective probabilities involves the formation of informal mental models of variables. These models cannot be arrived at deductively, but rather inductively. Although the scientific method is essentially an inductive process, forming hypotheses is far from a scientific process; this involves inferring imprecise, general patterns from past observation, for the sake of applying the same pattern for future observations. The imperfection of this method has been well-noted by epistemologists for centuries:

> That there is nothing in any object, considered in itself, which can afford us a
>
> reason for drawing a conclusion beyond it; and, that even after the observation of
>
> the frequent or constant conjunction of objects, we have no reason to draw any

---

[4] Moreover, as the amount of information in the model increases, the computational time necessary to process the information would increase dramatically.

inference concerning any object beyond those of which we have had experience.

(Hume, 1748/2000)

In other words, the confirmation of a pattern in the past does not necessarily mean that it will be confirmed in the future. Although deduction can be used to determine mathematical relationships between quantitative variables, ensuring that they are correct in any circumstances, determining the variables themselves requires some knowledge of the relationship between qualitative variables.

On a practical level, induction acts as a reasonable substitution for deduction in many cases. If one observes the sun rising every day of his life, he can infer with a high degree of certainty that it will rise the next morning. In terms of constructing accurate models (implicitly or explicitly), induction works best if there is swift, reliable, and clear feedback as to how well theory does; this allows for a quick update to the initial theory. Unfortunately, the problem of induction noted by Hume is amplified in the context of financial markets. Investors do not receive swift, reliable, or clear feedback.

Investors have no systematic way to determine whether their models were successful (unsuccessful) because they were accurate (inaccurate) or because they were lucky (unlucky). For example, consider the hypothesis that the S&P 500 index will increase two days after the New York Yankees win a baseball game, and will decrease two days after a loss. If an investor bases his investment strategy on this theory, and is successful, should he conclude that the theory is correct and continue to base her trading strategies on it? It is certainly possible that this relationship exists, but it is also possible that the strategy was implemented during a time frame in which the Yankees had an exceptional team and the S&P 500 was in the middle of a bull

market; the relationship could merely be coincidental. In any case, it is safe to conclude that this hypothetical relationship, if observed, would be ambiguous.

Another epistemological problem faced by investors is the accuracy of information. Investors can observe balance sheets, view financial metrics, and read news stories, but this is all secondary information. Companies occasionally misreport or misestimate the state of their finances, governments can only estimate macroeconomic figures, and news outlets misestimate social indicators. Like Hayek's central planners, the accuracy of individual investors' models are limited by the accuracy of the information they possess.

In short, investors must create and consistently uphold a model of the world with limited information and no completely reliable way to determine when they are correct. Moreover, this model must be complex enough to be at least somewhat reliable, yet simple enough for the investor to understand. This model must be fluid and able to cope with an ever-changing global economy. To make matters more complicated, investors must account for the mental models of all other investors – as described by the Keynesian beauty contest – and take into account how other investors take into account other investors' models.

For the sake of clarity, it helps to distinguish between two sets of facts that investors must analyze. Herrmann-Pillath (2009) does this by borrowing terms from Searle (1995): *observer-relative* facts describe facts that exist independently from the observer. These include fundamental metrics, current legal and political environments, and so forth. Each investor will have a different interpretation of the facts, but the facts will stay the same no matter whom or what views them. In the Keynesian beauty contest, the contestants themselves would constitute observer-relative facts. Conversely, *observer-dependent* facts describe facts that exist based exclusively on the state of the observer. The opinions of the judges in the Keynesian beauty

contest constitute observer-dependent facts. Although an investor's opinion constitutes an observer-dependent fact, if it leads him to buy a security, the resulting price increase would constitute an observer-relative fact. Determining the observer-dependent facts of investor opinion in the current period, however, has significant influence over price movements in future periods. Hermann-Pillath (2009) summarizes:

> A stock market is a highly subjective phenomenon in the sense that, for example, traders experience a wide variety of emotions and cognitive states related to it, and in a fundamental sense its driving forces are subjective […] However, the stock market clearly also is accessible to objective knowledge, such as statistics and econometric research built on it. (8)

Although correctly priced securities in financial markets should be based on the objective information of observer-relative facts, incomplete information in an uncertain environment render these facts difficult to interpret. In this sense, trying to discern observer-dependent facts may act as useful informational inputs into the individual investor's model. This is especially useful when attempting to construct explicit probabilities of uncertain events. While this can potentially lead to herding and financial bubbles, as investors using other investors as basis points can lead to a positive feedback look in prices, this is not necessarily true in every case. In an ethnographic study of derivative traders, Buenza and Stark (2010) describe how investors at a major investment bank were able to discern explicit probabilities of merger events from the positions of other traders. While this proved to be useful for many of their positions, however, Buenza and Stark (2010) observe one significant loss: since other major market players used the same strategy, they all suffered from "cognitive interdependence" and overlooked key information relating to their positions.

In examining the trading strategies of investors specializing in merger arbitrage, Buenza and Stark (2010) observe that traders were acutely aware of their cognitive limitations in observing and processing information, and devised ways to "reduce their cognitive overload." This was accomplished by making extensive use of an "assembly of electronic scaffolding to supplement [their] mental processes: a PowerPoint presentation, followed by a Word memorandum, followed by an Excel spreadsheet, all of it condensed into a single live cell on a Trading Summary" (20). Traders also collected and meticulously organized information on mergers. This allowed them to draw on previous, similar mergers in order to gain insight into the probabilities of future mergers.

Despite this organization, the traders were aware of the fallibility of their models, and considered the possibility of suffering from a myopic viewpoint. They observed the price difference –the spread – between the two merging companies as a way of measuring the market's uncertainty of the merger, thus providing a check on their own model. Observing the actions (indirectly) of other investors is useful in that it provides another viewpoint of the same set of information, and potentially points to overlooked information. Although the arbitrage traders meticulously constructed their models, they did not rule out the possibility that they were mistaken in their analysis. While viewing a spread that contradicted a model did not immediately lead to its abandonment, it caused the traders to reconsider their position.

Buenza and Stark (2010) observe two deals. The first deal involved a merger between Career Education Corporation and Whitman Education Group. The traders noticed that the spread between the two companies remained high. Although this initially caused the traders to look for more obstacles faced by the merger, they found none. They subsequently increased their position, but slowly so as to not affect prices significantly and catch the attention of competing

traders using similar techniques. When asked why the spread implied a different probability than the traders' models, they replied that it was due to a different interpretation of the information.

Buenza and Stark (2010) describe how this "reflexive modeling" led to a $2.8 billion loss amongst arbitrage traders after the collapse of the planned merger between General Electric and Honeywell International. Although the media exhibited concerns over a European opposition to the merger[5], the merger was approved in the United States. Since a merger approved in the United States had never been contradicted by a European regulatory entity, the traders did not account for it in their models; this led to the overestimation of the probability of the merger. Other traders overestimated the probability as well, so the spread between the two companies did not move significantly when the concerns were expressed. The lack of an increase in the spread subsequently confirmed the traders' initial hypothesis, and likely confirmed the hypothesis of competing traders. This ultimately caused the spread to narrow further, as the probability of the merger falling apart was considered miniscule, and investors increased their positions.

This ethnography shows that taking other investors' opinions into account does not necessarily lead to incorrect prices on the aggregate level. When the strategy is successful, the price converges to a single implied probability. However, this strategy, when implicated among many market participants, can lead to false confidence and ultimately prices that imply an incorrect risk. Buenza and Stark (2010) note, "a socio-technical conception of markets reveals that arbitrage disasters are not a fault of character, but an unintended consequence of a mostly functional system" (46).

Relating model reflexivity to merger arbitrage provides a clear description of how using observation-dependent facts can be used to check initial hypotheses; since mergers are binary

---

[5] Buenza and Stark (2010) observe a spike in articles in *The Wall Street Journal*, *The Financial Times*, and *The Economist*.

events – either they happen or they do not – it is easier to apply an explicit probability. However, this concept of reflexivity in financial markets can be applied to general investor behavior as a way of dealing with uncertainty. Specifically, using perceived observer-dependent facts as information is a way for investors to protect themselves against unknown uncertainty. While it may be difficult to devise a quantitative probability from qualitative variables, some uncertainty can be accounted for implicitly. In analyzing currency swaps, for instance, an investor can determine that the probability of one country's central bank raising interest rates is somewhere between zero and one. It is uncertain because it is impossible to explicitly quantify in a systematic way, but the investor recognizes the uncertainty's existence. If China declares nuclear war on one of the countries, however, the investor's model might not account for the event. This type of event can be classified as a Black Swan (Taleb, 2007), where the event seemingly came out of nowhere but is somewhat clearer in retrospect; but from the investor's perspective, this constitutes a different type of uncertainty.

2. Academic Thought

In this section, it is necessary to review the past academic literature of investor rationality and market efficiency for two reasons. First, from the academic standpoint, it is important to document what has already been said on the subject; this will help identify the questions that have been answered and the questions that remain unanswered. I pay particular attention to the methodologies used by both the EMH and BF in their attempt to overcome the problem of learning about a complex and uncertain system, not unlike the problem faced by investors. Second, as noted by MacKenzie (2006), many of the models inspired by the EMH are performative, and have significant influence over how markets operate. Moreover, as the ubiquity of BF literature increases, this may have a stronger performative effect over investors' theories of mind.

2.1 The Efficient Market Hypothesis

The EMH represents the heterodox school of thought relating to financial markets. While the EMH has its roots in the Random Walk model going back to Bachelier (1900), Eugene Fama famously outlines the EMH explicitly, first in Fama (1970) and then in Fama (1991). At its most succinct, the EMH is "the simple statement that security prices fully reflect all available information" (Fama, 1991: 1575). "Security prices" are well defined, but the rest of this simple statement is highly ambiguous on multiple levels, and requires clarification of numerous points. Fama splits "all available information" into three interpretations. The weak form of the EMH considers available information to be a security's set of historical prices; the semi-strong form considers all information that is available to the public (e.g., earning reports, market news); the strong form considers privately held information. Fama (1970) is acutely aware of the ambiguity of this definition, and admits that the definition of the EMH in this form is so general that it has no empirically testable properties without the help of a joint-hypothesis to lend a concrete definition of "full reflection."[6]

"Fully reflected" prices could be interpreted as equilibrium prices on a theoretical level, but Fama (1970) notes, "the equilibrium expected return on a security is a function of its 'risk.' And different theories would differ primarily in how 'risk' is defined" (384). He continues to note three sufficient (but not necessary) conditions for the EMH to hold: (i) transaction costs are zero; (ii) the cost of attaining information is zero; and (iii) all market participants agree on the pricing implications of information. These conditions are obviously not descriptive of financial markets, but market efficiency is still obtainable if they are met in sufficient numbers. For

---

[6] Fama (1970) even goes so far as to put quotations around every mention of "fully reflect" and "efficient" when discussing these concepts.

example, only a "sufficient number" of investors need to obtain relevant information for prices to be "efficient."

Although Fama notes the many ambiguities that make the EMH empirically untestable on its own, the crippling vagueness of the theory makes itself almost unintelligible. In his analysis, Fama makes a convincing argument that the word "fully" in full reflection is so ambiguous as to avoid critical examination, but he does not succeed with his choice of the word "reflect." Fama (1970) makes two important comments on this word, once by associating it with "utilize," and then by introducing sufficient condition (iii). The interesting implication of this is that the information has no value in and of itself; its relevance only comes once it is utilized by financial markets. Moreover, by condition (iii)[7], it does not matter how information is utilized or interpreted; so long as information is merely being utilized by market participants, the market is considered efficient. Even with the help of a joint hypothesis to define "full reflection," the theory is a tautology; markets are efficient by definition. If every market participant underwent a frontal lobotomy and made their trading decisions based on how pretty they thought the words in quarterly earnings reports were, the market would still be considered efficient by a literal interpretation of Fama's criteria.[8]

The ambiguity of the EMH and its assumptions, however, are somewhat irrelevant to its validity as a theory of positive economics. The theory uses a methodology that allows it to sidestep the question of decision-making and cognition under conditions of uncertainty. By only describing the general behavior of aggregate markets, the questions become irrelevant. Friedman (1953) argues in *The Methodology of Positive Economics* that the discipline of economics should

---

[7] It is difficult to interpret condition (iii) in the context of sufficiency, but positive economics as outlined by Friedman (1953) might suggest that this could refer to the *average* interpretation of information.
[8] Once could argue that the aesthetics of the words do not constitute information, but this requires a more formal definition of "information," which will come in chapter 5.

follow the scientific method of logical positivism, and defines the goal of the positive scientist as "the development of a 'theory' or 'hypothesis' that yields valid and meaningful (i.e., not truistic) predictions about phenomena not yet observed" (7). This theory should "abstract central features of complex reality," and comprise of language intended to "promote systematic and organized methods of reasoning." Friedman continues to make the more controversial claim that the theory's value is determined entirely by its predictive power; the assumptions it makes are irrelevant and need not be realistic (or even internally consistent). Theory alone is a tautology; whether expressed in language or a mathematical equation, the initial conditions contain the conclusion. It does not contain "substantive content," but rather acts as an "analytical filing system" to help clarify aspects of reality. A theory's ability to provide this is secondary, however, to its predictive power.

Friedman gives the example of expert pool players. One might formulate a hypothesis stating that the players play as if they are aware of, and continuously calculate, the correct angle and force to make any shot. Being expert pool players, they make a large percentage of their shots. Does this mean that the players are *actually* aware of the complex mathematics of their shots? Not necessarily. This would be a good positive theory because it describes reality with a high degree of predictability; but the theory only states that shots are taken *as if* the players are aware of the complex mathematics. Extrapolated to economic phenomena, a business owner does not necessarily have to be explicitly aware of the fact that marginal operating cost is (approximately) equal to marginal revenue. The business owner might not even know what the two curves look like. So long as businesses generally exhibit this behavior, however, the theory is useful. In the case of the EMH, all three of Fama's (1970) conditions need not be satisfied; as

long as the market acts *as though* information and trading costs are zero, and all market participants agree on the implications of information, the EMH is useful.

Fama (1970) makes an attempt at satisfying the secondary goal by translating the EMH it into three sets of mathematical notations:

(1) $E(\tilde{p}_{j,t+1}|\Phi_t) = [1 + E(\tilde{r}_{j,t+1}|\Phi_t)]p_{j,t}$

where $p_{j,t}$ represents the price of security j at time t, $r_{j,t+1}$ represents the one-period percentage return, tildes indicate random variables, and Φ_t represents "a general symbol for whatever set of information is assumed to be 'fully reflected' in the price at time t." This states that future prices, subject to currently available information, are equal to current prices plus the expected return (conditional upon available information). He continues:

(2) $x_{j,t+1} = p_{j,t+1} - E(p_{j,t+1}|\Phi_t)$

(3) $E(\tilde{x}_{j,t+1}|\Phi_t) = 0$

Where $x_{j,t+1}$ represents the difference between actual future prices and future prices predicted by presently available information. The crux of the EMH rests on equation (3), where the expected value of $x_{j,t+1}$ is zero.

While the conversion to mathematical notation makes the EMH more closely resemble a testable hypothesis, the Φ term is still quite vague. In this form, Φ could be interpreted as "factors which determine security prices." This definition, however, is useless, as it would state that security prices are determined by the factors that determine security prices. Even using the three forms' definitions of "all available information" does not clarify the Φ term, as they are only slightly less vague. Fama (1991) identified this as the "joint-hypothesis problem," where market efficiency can only be tested with a model that explicitly defines "full reflection," and "all available information." Any test of market efficiency is therefore model-dependent.

A more modest thesis for market efficiency, endorsed by Jensen (1978) and Malkiel (2003) is that, while market prices are not always perfect, efficient markets "do not allow investors to earn above-average returns without accepting above average risks" (Malkiel, 2003: 60) with respect to the information they use. Considering the type of efficiency being tested, this means that investors cannot make risk-adjusted economic profit based on past pricing data, publically available information, or privately held information.

Friedman (1953) notes that theories can never be proven by empirical observations, but can only fail to be disproven. But it is unclear whether a test finding contradictory evidence actually disproves the EMH, or whether it is just a poor model. This has been the running theme of market efficiency research for most of the theory's existence, and much of the research in favor of the EMH consists of proponents noting apparent anomalies, and explaining it through use of the joint hypothesis problem. Those arguing in favor of the EMH believe that the empirical evidence is overwhelming; Jensen (1978) goes so far as to state, "I believe there is no other proposition in economics which has more solid empirical evidence supporting it than the Efficient Market Hypothesis" (95).

The empirical implications of the EMH (and subsequently, its testable properties) are generally interpreted with respect to the three forms of market efficiency. Weak form efficiency is measured through tests of how well past data predicts future prices. Semi-strong form efficiency is measured through tests of how quickly prices update to news. As new information is introduced to the market, the EMH would predict that prices adjust rather quickly. These tests are trickier, as it is difficult to determine which news is relevant to price changes. Tests of this nature can only observe market reaction relative to the introduction of individual pieces of information, such as stock splits and earning reports. It is difficult to disprove semi-strong form

efficiency, however, because of the relative ease of citing an undocumented piece of information as contributing to price movement.

Fama (1965) examines serial correlations of the 30 Dow Jones Industrial Average stocks between 1957 and 1962, and concludes that a relationship between past and future stock prices exists, but is miniscule. Much of the Random Walk research that Fama (1970) cites (Kendall, 1953; Samuelson, 1965; Mandelbrot, 1966) come to similar conclusions. Fama (1991) reiterates the research from his 1970 literature review, but also mentions the growing contrarian research that had come about since.

DeBondt and Thaler (1985, 1987) find that stocks that perform well over three and five year periods tend to perform poorly relative to the market in subsequent years; conversely, stocks that underperform the market in three and five year periods tend to outperform the market in subsequent years. They attribute this to investors systematically overreacting to news. This conclusion contradicts both the weak and semi-strong forms of market efficiency. Banz (1981) responds to this conclusion, arguing that this phenomenon is likely due to firm size; small stocks, which generally perform worse than the market, have greater upside potential. Fama (1991) offers his analysis, "we may never be able to say which explanation of the return behavior of extreme winners and losers is correct, but the results of DeBondt and Thaler and their critics are nevertheless interesting" (1582).

Tests of semi-strong form efficiency find similarly ambiguous results as weak form tests; a large portion of the literature focuses on how prices respond to events and new information, but tend to continue drifting in the same direction for a period after the event (Ball and Brown, 1968; Foster, Olsen, and Shevlin, 1984). However, these same authors note that this is not necessarily inconsistent with semi-strong efficiency; Ball (1978) argues that, in the case of post-earnings

drift, the CAPM model is insufficient to model the risk associated with the excess returns of the drift.

Strong form efficiency is largely believed to be false, but useful for the sake of theoretical completeness. Tests tend to focus on the prevalence of private information held by investors, such as documentation of insider trading (Black, Jenson, and Scholes, 1972) leading to excess returns. While this may be the only testable form of private information, it is unclear what, exactly constitutes private information on a theoretical level.

Yet another ambiguity of the EMH is its treatment of information regarding other investors. As the Keynesian beauty contest points out, security prices are directly determined by other investors, and only indirectly by market fundamentals. For an individual investor to have a hypothesis of a security's future price, he must also have a hypothesis of the market's interpretation of the given information. Even if the investor expects all other investors to interpret the information in the same way (on average, at least) it requires a theory of mind, which is presumably informed partially through observations of how other investors react to information. Fama does not explicitly classify this form of information – likely due to the difficulty of quantifying it – but an orthodox interpretation of market efficiency might classify it as either the private information of strong-form efficiency, or being encoded in past prices (weak form efficiency).

Ultimately, this is not necessarily problematic for the EMH; as long as markets acts *as if* they are efficient with respect to past price information or earnings reports, the theory holds regardless of the underlying mechanisms for price formation. If an investor holds a mistaken belief about the market, for example, but makes an investment based on how he thinks other investors behave, his investment might better reflect the fundamental information. But if

investors believe that they are "correctly" interpreting the information, while others are "incorrectly" interpreting it, or vice versa, aggregate prices might exhibit interesting behavior.

2.2 Behavioral Finance

Starting in the early 1980s, the financial literature began to present a continuously increasing amount of empirical evidence that highlighted anomalous financial phenomena. Along with the contradictory evidence cited in Fama (1991) literature review, the literature of Behavioral Finance (BF) has since noted many apparent anomalies and patterns in financial markets that are difficult to view under the scope of weak form market efficiency. Banz (1981) finds evidence that small cap stocks consistently earn better returns than large cap stocks; Shiller (2005) finds evidence that stocks with low price/earnings ratios consistently outperform stocks with high ones; Jegadeesh and Titman (1993) find evidence that stocks' price movements exhibit positive and negative momentum. Moreover, there is considerable evidence that prices for certain securities exhibit higher returns in Januarys (Haugen and Lakonishok, 1988), and on Mondays (French, 1980).[9] But, as noted in the previous chapter, the EMH school systematically reviews these observations and concludes that either the deviation from efficiency is insignificant, or that the model is too ambiguous to discredit market efficiency. Nevertheless, these phenomena have been systematically documented, and the EMH does a poor job of explaining them.

Even if the EMH describes the underlying force of financial markets, BF argues that the EMH ignores the many frictions getting in the market's way. Moreover, BF argues that these frictions (investor irrationality and limits to arbitrage) are the root causes of the documented phenomena. Barberis and Thaler (2005) summarize this view: "The theory of limited arbitrage shows that if irrational traders cause deviations from fundamental value, rational traders will often be powerless to do anything about it" (12). BF attempts to explain market phenomena by

---

[9] It should be noted that French is not necessarily a proponent of BF, but notes that it would be difficult to systematically profit from the "Monday Effect" due to transaction costs.

departing from the positive economics methodology outlined by Friedman (1953)[10], and directly observing the underlying mechanism of financial markets by taking into account psychology and sociology literature. In doing so, they bring to light some useful information regarding of investor decision making.

Much of the BF literature consists of studies documenting "irrational" behavior in humans, especially in the context of risk analysis and uncertainty. Many of these studies present subjects with various situations that have rationally correct solutions, and observe how actual behavior deviates from the predictions of the *homo economicus* model.[11] This behavior can be explained by two viewpoints. One explanation is that natural selection honed specific traits in the human brain to survive in a natural environment, but these traits are ill-equipped to analyze the complexities of financial markets; human investors are thus subject to systematic biases in their analyses. Another explanation is that the human brain possesses a limited amount of computational power, and uses heuristics (i.e., rules of thumb, or generally applicable problem-solving techniques) to offer a best guess of the correct answer while expending smaller amounts of energy. The two explanations are not necessarily mutually exclusive, and ultimately lead to the same implications; while these deviations from rationality may be anywhere between miniscule and substantial on an individual basis, BF argues that they lead to significantly different outcomes in aggregate. The terms 'bias' and 'heuristic' are frequently used interchangeably throughout much of the literature, although the former implies a systematically wrong conclusion, and the latter implies a strategy that can either lead to a correct or incorrect conclusion.

---

[10] That is, by challenging the notion that a theory's internal mechanisms are irrelevant. BF justifies its argument in part by directly observing whether the larger theory's assumptions are realistic.
[11] *Homo economicus* is a theoretical, perfectly rational and self-interested economic agent whose behavior is meant to approximate that of humans.

Kahneman and Tversky (1974) present three oft-cited heuristics: *representativeness*, *availability*, and *anchoring*. Representativeness suggests that humans, when asked to determine the probability that object A belongs to set B, base their projection on how well object A resembles set B. The most famous example of this is the description of a woman named Linda:

> Linda is thirty-one years old, single, outspoken, and very bright. She majored in
>
> philosophy. As a student, she was deeply concerned with issues of discrimination
>
> and social justice, and also participated in anti-nuclear demonstrations.

When research participants were asked whether it is more likely that (A) "Linda is a bank teller," or (B) "Linda is a bank teller and is active in the feminist movement," they typically chose (B), as that description is more representative of a feminist than a bank teller. However, the probability of (A) is larger than (B), as the latter option is contained within the former; Linda is a bank teller in both descriptions, but (B) merely eliminates the scenarios where Linda is not a feminist.

Kahneman and Tversky (1974) note that the representativeness heuristic is also present when evaluating sample sizes; they find that many human subjects weight small sample sizes the same as larger ones. In the context of financial markets, Barberis and Thaler (2005) note that this heuristic could pose a problem in the consideration of fund managers or financial analysts. If an analyst makes four good stock predictions, for example, he might be perceived as talented; this string of predictions is more representative of a good analyst than a bad analyst. Conversely, the gambler's fallacy, where humans expect small samples to always represent the same probabilities as larger samples, describes the same heuristic working in the opposite direction. For example, Kahneman and Tversky (1974) note that if a (fair) coin is flipped seven times and

lands on heads each time, many subjects expect a tails to be "overdue." The probability of tails, however, is still 50%, independently of the previous flips.

The availability bias (Kahneman and Tversky, 1974) suggests that people tend to weight information that is more easily retrievable heavier than other information. For example, someone whose friend has recently suffered a heart attack may disproportionally estimate the probability of having one. While the memory is, statistically, a single data point, it is more easily retrievable than a mental calculation of the actual probability. In a financial setting, this bias may lead people to choose securities that seem more familiar than securities with better fundamental valuations.

Kahneman and Tversky (1974) also argue that humans exhibit anchoring; that is, they make evaluations relative to an arbitrary starting value. When asked to estimate the percentage of United Nation countries that were African, subjects estimated an average of 25 percent when first asked whether the number was higher or lower than 10 percent. In contrast, subjects estimated an average of 45 percent when initially asked if the number was higher or lower than 60 percent. When evaluating security prices, the current market price could easily distort an investor's perception of the security's true value by acting as a reference point.

Since Kahneman and Tversky (1974), researchers have noted many other biases and heuristics. Lord, Ross, and Lepper (1979) find that people are very reluctant to give up previously held beliefs, and are subject to a *confirmation bias* – that is, they weight information that conforms to their beliefs more heavily than information that contradicts it. Alpert and Raiffa (1982) note an overconfidence bias, where people overestimate their predictive capabilities. For example, they find that when people are asked to construct 98 percent confidence intervals, they are correct only 60 percent of the time.

The BF literature also notes some peculiar preferences held by human subjects. Prospect theory, proposed by Kahneman and Tversky (1979), notes that, on average, people are risk averse when dealing with gains, but are risk seeking when dealing with losses. A disproportionally strong preference for certainty over uncertainty, which is known as ambiguity aversion, has also been observed. In the Ellsberg paradox (Ellsberg, 1961) discusses a situation in which participants are given the option of choosing one of two games. In the first game subjects pick a ball from an urn containing 50 black balls and 50 red balls, and can bet on either color. In the second game, an urn contains 100 balls, but the amount of each color is unknown. Although subjective utility theory predicts that subjects should be indifferent between the two games, Camerer and Weber (1992) observe that they systematically choose the first.

While this approach emphasizes the problems of the EMH, and provides useful information in regards to the behavior of individual investors, it ultimately suffers from some theoretical and empirical problems. In terms of the macro empirical data, Malkiel (2003) notes that, while the noted phenomena occasionally occur, they do not occur predictably. He notes a Wall Street joke that the January Effect "is more likely to occur on the previous Thanksgiving." If these effects were, in fact, reliable and dependable, investors who read the academic literature would exploit them, thus distorting the effect or eliminating it all together. The BF literature itself would be counterperformative. The more profitable – and identifiable from an academic standpoint – are these patterns, the more likely investors are to find and exploit them. BF argues that arbitrage costs prevent investors from exploiting patterns in most cases, but this means that the patterns themselves are not large or distinctive enough to exploit. If this is the case, it is difficult to argue that the patterns are significant in an academic context.

Moreover, BF does a poor job of explaining the circumstance in which these biases arise on both a micro and a macro level. On the macro level, it is unclear how the observed behavioral patterns explain aggregated market phenomena. BF compiles a long list of circumstantial evidence, but does not show a definitive link between micro and macro financial behavior. While certain biases may lead some investors to overreact to information, this does not necessarily imply that the entire market will overreact to information; perhaps the confirmation bias would lead other investors to discredit the same information, and react equally in the opposite direction.

From a micro standpoint, these conflicts manifest themselves within the mind of each investor, and it is difficult to determine which biases will exhibit themselves when not isolated from one another. If an investor observes a stock's movements for five days, will he commit the representativeness bias and buy it based on the small sample size of observations, or will his ambiguity aversion prevent his from drawing any conclusions from it? After seeing poor performance in their portfolio, will an investor succumb to prospect theory and exhibit risk-seeking-in-losses behavior, or will loss-aversion prevail? In analyzing a security, will overconfidence and confirmation bias cause an investor to take on excessive risk, or will their risk aversion prevail? Will an investor succumb to availability bias and over-account for the financial disaster that stimulates vivid imagery, or will he write off the small probability of such an event as impossible?

It is also unclear how these observations relate to actual investors making decisions in the context of actual financial markets. Most BF studies, and nearly all behavioral psychology studies, primarily test subjects' ability to process risk, rather than uncertainty, and do so using tests of the general public instead of practicing investors. Barberis and Thaler (2005) note that there is little evidence that experience related to investing diminishes these biases (and even

suggest that experience may lead to more overconfidence in financial analysis), but this does not fully resolve the question. It is difficult to draw conclusions about how investors process risk, as these subjects are not given the same tools that investors frequently utilize. As established in the first chapter, investors make extensive use of formal models when making decisions; it is not unlikely that they would utilize calculators and internet search engines to aid them if faced with the same problems of behavioral psychology experiments. Perhaps these biases exhibit themselves when investors determine implicit probabilities in situations of uncertainty, but BF has yet to establish this.

With respect to uncertainty, BF's reliance on the concept of investor rationality as a comparison against human investors poses a problem as well. The first sentence of Barberis and Thaler (2005) states, "The traditional finance paradigm … seeks to understand financial markets using models in which agents are 'rational'" (1). They continue to define rationality as agents' ability to "update their beliefs [with respect to new information] correctly, in the manner described by Baye's law," while holding these beliefs consistent with each other. Their first footnote further elaborates the consistency of beliefs:

> Consistent beliefs means that agents' beliefs are correct: the subjective distribution they use to forecast future realizations of unknown variables is indeed the distribution that those realizations are drawn from. This requires not only that agents process new information correctly, but that they have *enough* information about the structure of the economy to be able to figure out the correct distribution for the variables of interest (1).

Barberis and Thaler (2005) argue that, in light of the evidence compiled against complete investor rationality, BF offers a better understanding of financial phenomena. However, this

view makes critical assumptions about rationality, irrationality, and correct pricing. It assumes that there exists an unambiguous interpretation of all available information, which implies a "correct" price, that rational investors interpret the information in this way, and that irrational investors interpret the information "incorrectly," subsequently influence the price in a direction away from the "correct" one.

In the context of uncertain financial markets, it is quite difficult to define what "rational" behavior entails. In light of the complex set of information investors use, the methodology of BF might be inadequate to properly analyze whether investors conform to the notion of "rationality." While many tests show that humans are poor calculators of risk[12] where probabilities are known and easily identifiable, it is much more difficult to describe whether a reaction to uncertainty is rational or irrational. Even in the context of the experiments there is, from the participant's perspective, a great deal of ambiguity and uncertainty.

The Ellsberg paradox game, for example, has different outcomes depending on whether the experimenter or a friend compose the ambiguous urn (Kühberger and Perner, 2003). This implies that participants may be taking into account information about the uncertainty (and implicit probability) of the aspects of the experiment itself. In another example, the test for representativeness involving Linda the bank teller, in concluding that participants are irrational, assumes that the two options are clear and convey a definitive amount of information. When comparing option (A) "Linda is a bank teller" with option (B) "Linda is a bank teller and is active in the feminist movement," it is unclear what information is conveyed to the participant. Is it irrational to conclude that the inclusion of the "feminist movement" information in option (B) implies that Linda is *not* active in the feminist movement in option (A)? If the participant interprets the options as (A) "Linda is a bank teller and is not active in the feminist movement"

---

[12] Without access to calculators or formal models.

and (B) "Linda is a bank teller and is active in the feminist movement," then this does not imply irrationality – it implies that the participant interpreted the information in a way deemed incorrect by the experimenter. However, in situations where the implications of information are ambiguous, this type of interpretation may not necessarily be incorrect. In the context of laboratory experiments, the experimenters have the informational upper hand, as they were the ones who constructed the experiment. In this sense, the participants are working with a different set of information than is *homo economicus*.

Taking uncertainty into account, some of the biases might be perfectly rational; the availability bias, for instance, might not be a bias if information has varying degrees of validity. The most easily available information might prove to be the most valid. It is extremely important to consider uncertainty when looking at these experiments because investing involves a great deal of uncertainty. Unlike financial markets, many of these experiments involve games with clearly defined risk and probability. It is therefore questionable as to how well they relate to actual investment patterns.

Moreover, Barberis and Thaler (2005) set up a straw-man argument against the EMH. In discussing the Malkiel (2003) definition of market efficiency, Barberis and Thaler (2005) note that "prices are right" implies "no free lunch," but "no free lunch" does not imply that "prices are right." Their point is that, if prices are wildly inefficient, arbitrage costs could still prevent investors from making a positive risk-adjusted profit. In other words, a security that is priced efficiently does not necessarily imply that it is priced correctly. But it is important to note that Fama (1970, 1991) describes the theory of market efficiency without any mention of rationality. The EMH sidesteps the issue of rationality and "correct" pricing entirely by *defining* efficiency in terms of the market's interpretation of information.

What Barberis and Thaler (2005) point out is precisely the problem noted by Fama in testing the EMH: the joint-hypothesis problem. Malkiel (2003) even states, "What I do not argue is that market pricing is always perfect" (61). Since information is often incomplete or inaccurate, the price of a security can only be as accurate as the information associated with it implies. In his later discussion of the internet bubble, Malkiel (2003) notes that the "irrationality" of markets is only clear in hindsight:

> While it is now clear in retrospect that such professionals were egregiously wrong, there was certainly no obvious arbitrage opportunities available. One could disagree with the projected growth rates of security analysts. But who could be sure, with the use of the Internet for a time doubling every several months, that the extraordinary growth rates that could justify stock valuations were impossible? After all, even Alan Greenspan was singing the praises of the new economy. Nothing is ever as clear in retrospect. The extent of the "bubble" was only clear in retrospect. (75).

Although Barberis and Thaler's (2005) dissection of "no free lunch" efficiency may be correct, it is unclear which theory, exactly, they are contradicting. The difference between Barberis and Thaler's (2005) use of the term "correct" and Fama's (1970) use of the term "efficient" may just be semantics, but I believe that it highlights the fact that BF fundamentally misunderstands the EMH's claims.

2.3 Methodological Comparison

Friedman (1953) notes that, since the effects of air resistance on a falling object are often miniscule, the simple equation to find the speed of an object falling in a vacuum[13] is frequently used as an approximation of the force of gravity. The assumption that air resistance does not exist is not meant to be taken literally, but makes the model more efficient in the sense that it is greatly simplified without significantly compromising the accuracy of its prediction. Certain scenarios, such as the weight and shape of the object, can amplify the model's shortcomings, but Friedman argues that it does not always pay to compromise the simplicity of the model. In any case, the basic model provides a useful framework. Proponents of the EMH argue that it provides an efficient[14] model of market behavior with respect to the methodology of positive economics. Much like the formula for the speed of an object falling in a vacuum, the EMH describes the outcome of a process working in a frictionless environment without describing how the process itself works. The testability of the EMH is model-dependent, however, as it offers no testable hypothesis on its own.

While the EMH might potentially be accurate descriptor of markets on average, it does a poor job of explaining market behavior under all circumstances. The dot-com bubble may have been a result of efficient prices, for example, but the theory does not offer any insight as to why these phenomena occur. The period in which the EMH was formulated may have necessitated the sacrifice of some accuracy due to technological constraints[15], but these obstacles are no

---

[13] $s = \frac{1}{2}gt^2$; where s = speed, g = the gravitational constant, and t = time.

[14] That is, it is efficient in its explanation of reality with respect to the theory's complexity. This is not to be confused with market efficiency.

[15] For much of the discipline's history, it was very difficult to observe economic behavior directly. While hard sciences, and some soft sciences, could turn to controlled experiments – in which isolating variables is comparatively simple – economists hoping to achieve the same degree of rigor turned to mathematical models and statistical analysis.

longer as large. In this context, the EMH provides a useful foundation to build a theory upon, but it now may be worth sacrificing model efficiency for precision.

BF takes the first step towards this goal, first by observing apparent market inefficiencies, and then by attempting to explain them through observed experimental behavior. This approach is highly beneficial, as it catalogues the many ways in which actual human behavior tends to deviate from the behavior predicted by theory. The approach is incomplete, however, as it suffers from three main problems. First, BF's refutation of the EMH misinterprets its claims. This primarily stems from a methodological disagreement, rather than an empirical one, but neither side explicates it; ultimately, the two schools of thought talk past each other.

Second, the behavioral experiments used by BF decontextualize financial decision-making theoretically and empirically. It is unclear whether the observed behavior constitutes "irrationality," given the amount of information the participant has and the tools at their disposal within the confines of the experiment.

Thirdly, it is unclear how these biases would manifest themselves outside of the laboratory, as many situations might prompt the biases to overlap and contradict each other. This confusion is evident in explaining individual investor behavior, and in explaining aggregate market behavior. BF fails to draw a clear link between the behavior of individual investors and the aggregate behavior of markets. Moreover, observing persistent aggregate market inefficiency from an academic standpoint is extremely difficult, as the observations are likely counterperformative.

Until these three problems are addressed, the EMH continues to provide the only coherent working model of financial markets. Though vague, the general prediction that it is impossible to make above-average returns without accepting above-average risk provides

investors and academic economists with a consistent and mostly true assumption. But as the accuracy of models gains an increasing amount of influence (from both private and policy standpoints) over the way financial markets actually function, "mostly true" may not be adequate.

3. The Mind of the Investor

Despite the proliferation of EMH and BF literature, neither offers an entirely comprehensive theory of financial markets. Proponents of the EMH argue that the theory acts as an accurate description of markets most of the time, but do not offer an alternative description for when markets do not act efficiently. Moreover, by viewing the actual market mechanism as a black box, the EMH makes it difficult to determine when markets are efficient and when they are inefficient. In this context, it is surprisingly difficult to evaluate market efficiency with respect to the seemingly inefficient phenomena of bubbles and financial crises. Even if these phenomena can, in fact, be classified as efficient, the EMH does not offer any useful insight into why they occur. Conversely, BF does not offer any cohesive description of financial markets, but instead relies on the mere existence of investor fallibility to argue that market prices are inaccurate. However, this view fails to explain why (and if) market prices are accurate, at least some of the time. Both views can be used to explain the alternating behavior of market prices – efficient and inefficient, accurate and inaccurate – but there lacks a cohesive perspective of why sometimes the market is efficient (accurate), and why it is sometimes inefficient (inaccurate).

A more cohesive theory must address the three problems faced by BF, and the literature of neuroeconomics and philosophy of mind provide a good starting point for this. This section will selectively review the literature of neuroeconomics[16], as it offers a more relevant description of decision making under uncertainty. I then discuss philosophy of mind literature as a possible interpretation and clarification of neuroeconomics' findings, and use this interpretation to respond to the methodological questions of the EMH and BF. Finally, I conclude with a proposal for a theory of market efficiency that adequately accounts for cognition in financial markets

---

[16] That is, economic decisions analyzed through the tools of neuroscience.

3.1 Neurofinance

For reasons already stated, it is fundamentally difficult to draw conclusions from behavioral experiments relating to finance. BF attempts to learn about workings of financial markets by using reductionist methods; by examining the individual parts of financial markets – investors – in isolation, BF infers conclusions about how the entire system functions. Investors, however, make decisions based on the resources and information at their disposal. The decisions they make are highly context-dependent, and altering the context in which they make decisions can lead to drastically different behavior.

Assuming that BF studies adequately isolate cognitive tasks such that participants have the necessary information to make optimal decisions, sub-optimal behavior only sheds light on human deductive faculties. Since the investment process involves aspects of both deduction and induction, testing the latter is just as important, if not more so; while investors have access to a wide selection of mathematical and technological tools to aid them with deduction, the same cannot be said with respect to induction. However, there are no studies that test the inductive faculties of humans with *homo economici* as a control group, thus rendering a test of optimal induction impossible. Even if the behavior observed in a behavioral experiment examining induction under conditions of uncertainty can be generalized towards investor behavior in financial markets, no information is provided as to whether the behavior is rational, boundedly rational, or irrational; the experiment only states that, given the circumstances of the experiment, certain behavior manifests itself.

From the economist's standpoint, however, these studies are partially useful in that they make observations that could potentially increase the predictive power of models. Even though using observed experimental behavior may be unreliable in the context of actual markets, it

provides a clearer set of criteria than does asking investors about their motivation and thought processes. With the advent of technology with the precision to measure neural phenomena, however, it has recently become possible to observe the underlying process of human decision-making, all without having to venture into the murky topic of subjective experience and motivation.[17] Camerer, Loewenstein, and Prelec (2005) note in their literature review of neuroeconomics that the methodological boundaries of economics "have constantly been reshaped by tools such as mathematical, econometric, and simulation methods. Likewise, the current surge of interest in neuroscience by psychologists emerged largely from new methods, and the methods may productively blur the boundaries of economics and psychology" (12).

Camerer et al. (2005) continue to note some of the tools neuroscience has at its disposal. Brain imaging techniques, such as the electro-encephalogram (EEG), positron emission topography (PET), and functional magnetic resonance imaging (fMRI), measure electrical activity and blood flow in the brain, elucidating the time course of processes or brain regions associated with different tasks and functions. Measuring the activities of individual neurons allows for more precise observations, although this is highly intrusive and is confined to non-human testing. Electrical stimulation of specific brain areas, as well as experiments involving patients with lesions on certain areas, allow researchers to view results of the stimulation and inhibition of these areas. Measurements of psychophysical indicators (i.e., heart-rate blood, pressure, skin conductance, etc.) provide less precise observations, as they examine the effects of the central nervous system on the peripheral nervous system, but are also less intrusive and can be used outside of the laboratory setting.

Despite these technological advancements, neuroscience largely does not completely solve the decontextualiztion problem, as the experiments take place in the imperfectly simulated

---

[17] This will come in the following chapter.

environment of the laboratory. But, in the same way that economists find it useful to examine labor and capital as components of GDP, providing more precise measurements of the components of behavior – and of how these components interact – may prove to be useful regarding behavior outside of the laboratory. In order to discuss and appreciate the findings of neuroeconomics, however, a cursory understanding of neuroscience is required.[18]

The human brain is composed of a network of billions of interacting neurons, communicating with each other through the use of synthetic chemicals called neurotransmitters. Neurons receive neurotransmitters through dendrites (branch-like structures), send an electrical pulse through its axon (a thin fiber), and releases neurotransmitters into a synapse (the space between it and another cell), at which point another neuron detects the neurotransmitters and responds. Neurons are highly interconnected, and subsequently comprise a vast network of communication between one another. Cognition results as an emergent property of this network.

Hebbian theory[19] suggests that, as specific pairs of neurons communicate more often, the neural pathways connecting the two become stronger. This theory has relevant implications for learning, memory, and induction. As the neural patterns recur multiple times, the neural circuits associated with the pattern strengthen, making the patterns of neural communication more likely to occur. In terms of behavior, observations associated with positive outcomes might stimulate the associated pathways, while observations associated with negative outcomes might activate another pathway. For example, if an investor observes a pattern between companies reporting higher-than-analyst-expected earnings and an increase in its stock price, the neural pathways between the sections of the brain that correspond to these events would become stronger, and the investor would have an easier time associating the two concepts in the future, potentially leading

---

[18] I use Kalat (2009) for reference in this very general overview of neuroscience.
[19] First introduced by Donald Hebb in Hebb (1949).

to profitable investments. The concepts of earnings reports and stock prices are admittedly abstract at the neural level. This information must first be observed translated from either light or sound, and then processed by multiple parts of the brain before these concepts can even be formed.

Like an economy, the brain does not operate through a chaotic mess of interactions, but rather specializes its functions by region. The hindbrain controls vital body functions. The limbic system, including the amygdala, thalamus, hippocampus, and hypothalamus play important roles in emotion, memory, directing sensory input, and motivation, respectively. The amygdala in particular plays a strong role in fear and anxiety. The cerebral cortex is involved with decision-making and processing information at a higher level; the temporal lobe processes hearing and language; the occipital lobe is responsible for processing visual input; the parietal lobe processes sensory input; and the frontal lobe is associated with planning and voluntary movement. The prefrontal cortex, a subsection of the frontal lobe, is involved with the higher processing of information and decision-making. Furthermore, the prefrontal cortex can be divided into multiple specialized sections.

Much of the neuroscience literature, as well as the BF literature, distinguish between two types of neural processes: automatic and controlled processing. Camerer et al. (2005) characterize "controlled" processing as a serial and effortful process that can be deliberately evoked, and typically allows for good introspection. This activity typically occurs in the prefrontal cortex, and can be characterized by the step-by-step logic of deductive reasoning. Conversely, "automatic" processing works in parallel, and is effortless, reflexive, and may not provide any introspection. Image and language recognition are automatic processes, working through the occipital and temporal lobes. This plays an important role in pattern recognition, and

subsequently induction. Camerer et al. (2005) further distinguish between the "cognitive" functions just described and "affective" functions, which deal with emotions, feeling, and motivation. Most affective functions are automatic and may not even be recognized by the controlled process.[20]

These regions do not act in isolation. For an investor to analyze the balance sheet of a corporation, he must first translate visual stimuli into words and numbers, and process the meaning of these words and numbers. On a more abstract level, investors might automatically recall the balance sheet of a similar company for comparison. In an uncertain environment deduction cannot act alone, and the ability to automatically recall relevant information from the past is extremely important.

For investors, who are primarily concerned with making investments as well as analyzing them, the interaction between cognitive and affective systems is important as well. In terms of motivation, traditional economic theory would invoke a utility-based explanation for investors' desire to make money. Montague and Berns (2002) note that, in situations where different potential outcomes are hard to compare, the brain must use some sort of "internal currency." They suggest that the dopaminoceptive[21] system of the orbitofrontal cortex – a subsection of the prefrontal cortex – may serve the function of weighing the expected outcomes of predicted scenarios when faced with uncertainty. Montague and Berns (2002) note that the computational model they use to describe the resulting predictor-valuation process bares a remarkable resemblance to the Black-Scholes option-pricing model in functional form. They comment, "This odd connection may simply be a coincidence; however, it is possible that the connection between the two approaches is symptomatic of a more fundamental biological connection" (280).

---

[20] Camerer, et al. (2005) note that controlled-affective processes are difficult to distinguish in a pure form, but method acting provides a good characterization.

[21] Dopamine is a neurotransmitter associated with reward.

Hsu et al. (2005) also find evidence that the orbitofrontal cortex, as well as the amygdala, is active when participants face uncertainty. Participants were presented with a similar game as proposed in the Ellsberg paradox, with two decks of cards – one possessing half red cards and half blue cards, and the other composed of an unspecified proportion of each. Although they replicate the findings of Camerer and Weber (1992) – that is, participants systematically choose the deck with known risk – Hsu et al. (2005) provide a more precise snap-shot of decision making by using an fMRI.

While providing more useful information than Camerer and Weber (1992), the Ellsberg paradox game only tests known uncertainty; participants are not aware of the probabilities of drawing either color card, but are aware of the magnitudes of the consequences of drawing each. Bechara et al. (1994) provide a better test of uncertainty in observing participants perform the Iowa Gambling Task: participants are presented with four identical decks of cards and $2,000 of play money. The participants are then told to draw a card from whichever deck they want, and that some cards will lead to a monetary payoff, and some will lead towards a monetary loss. The participants are not told the probabilities of drawing each card, nor the magnitude of the gain or loss, for any of the decks. Participants face an uncertain scenario, and must induce the value of drawing cards from each of the decks. The decks, however, are weighted such that decks A and B have high payoffs, but incur even higher losses; the expected value of the decks is negative. Conversely, decks C and D have lower payoffs than the previous two decks, but even lower penalties, yielding a positive expected value. Participants are not told how many rounds they will play, but the game ends after 100 trials. Along with the control group, a number of ventromedial prefrontal cortex lesion patients play the game. The ventromedial prefrontal cortex, a subsection of the orbitomedial prefrontal cortex, is believed to be associated with integrating emotional

information between the amygdala and the prefrontal cortex. Bechara et al. (1999) found that, while both groups initially chose cards from all four decks, the control group eventually tended towards decks C and D, while the lesion patients tend toward A and B.

Bechara et al. (1997) replicated the experiment, this time measuring skin conductance and asking the subjects to explain their thought process intermittently. They found that the control participants exhibit anticipatory responses before choosing from the bad decks after a number of draws. This occurs before the participants are able to articulate their hypothesis that decks C and D have negative expected values. Conversely, the lesion patients did not exhibit the same anticipatory responses, nor do they stop picking cards from the bad decks.

In both experiments the "general intellect and problem solving skills" of the lesion patients were deemed intact. They more closely resemble the characterization of *homo economicus* than did the control group, as they were reliant purely on their logical capabilities without interference from their emotions; yet, the lesion patients perform worse. If the ventromedial prefrontal cortex does, in fact, play a large role in integrating emotional information from the amygdala, these results imply that emotions play a large component in successful decision-making under conditions of uncertainty.

Damasio (1996) outlines the Somatic Marker Hypothesis as an explaination this phenomenon. The hypothesis states that, in the face of uncertain decisions, the brain utilizes information relating to the body's physical state as a way of predicting the value of the potential outcome. More specifically, the hypothesis states that, after an event occurs – inducing an emotional reaction – the body experiences an affective, or somatic, state. When a similar event occurs in the future, the body enters an anticipatory somatic state; this anticipatory state is then interpreted by the ventromedial prefrontal cortex, which can re-trigger the emotional response

associated with the event's outcome, potentially leading to the recall of information associated with the event. This theory, coupled with the Iowa gambling task experiments, suggests that emotional responses comprise a large role in intuitive decision-making under uncertainty. Given enough time, participants may use conscious observations to supplement their decisions, but it appears that affect provides the necessary information for making quicker, yet still accurate, decisions.

Lo and Repin (2002) provide contextual evidence for the somatic marker hypothesis by measuring the somatic activity of actual traders during live trading sessions. The body temperature, skin conductance, respiration, blood pressure, and muscle movements of ten derivatives traders were measured, and were compared with real-time prices. Lo and Repin (2002) found that affective responses were significantly correlated with market events, though these responses were stronger for less experienced traders. Lo et al. (2005) replicated these findings while using a larger sample size with varying degrees of trading experience. They extend upon the results of the first study, and also found that traders with more extreme emotional reactions to events exhibit worse performance.[22]

Noting the potentially overwhelming effects of emotion on decision-making, Loewenstein et al. (2001) formulate the risk-as-feelings hypothesis. This hypothesis extends the somatic marker hypothesis, postulating that anticipatory emotions have the potential to override the rational decision making of the prefrontal cortex, as well as supplement it. Taken with the literature regarding emotions and uncertainty, this implies that investors faced with highly uncertain and fear-inducing scenarios may make less optimal decisions than they would otherwise. If financial markets take a quick negative turn, for instance, investors may exacerbate the downfall by exhibiting poor decision making. Shiv et al. (2005) observe that, when betting on

---

[22] Normalized with respect to daily standard deviation of profits and losses.

coin flips, patients with orbitofrontal lesions are less likely to stop betting after losing bets, thus leading to more profitable outcomes than control participants. Although this is primarily a test of risk, rather than uncertainty, it shows that emotional reactions to losses have the potential to impair decision-making.

Hsu et al. (2005) provide similar observations; while control patients become risk- and ambiguity-averse while playing the Ellsberg Paradox game, orbitofrontal cortex lesion patients are risk- and ambiguity-neutral, making no distinction between the two decks. "This is behaviorally abnormal," Hsu et al. note, "but is consistent, ironically, with the logic of subjective utility theory" (1682). De Martino, Camerer, and Adolphs (2010) also find evidence that amygdala lesion patients exhibit less risk aversion than do control patients.

Alternatively, Kuhen and Knutson (2005) and Knutson et al. (2005) both find evidence that heightened activity in the nucleus accumbens is associated with gain prediction. Activity in this area also predicted risky choices and risk-seeking mistakes. Interestingly, Breiter et al. (2001) find evidence that the dopamine-based neural pathways in this area are involved with the expectation of monetary gain, as well as the craving for cocaine; this implies that high levels of excitement and anticipation are correlated with activity in this area.

These studies taken together provide strong evidence that information processing, especially with respect to uncertainty, occurs outside of the serial, logical thought process. Emotional reactions towards events – driven primarily by the amygdala – appear to be encoded in the body via affective markers, and play a large role in weighing outcomes as the orbitomedial prefrontal cortex interprets relevant cues. This process provides quick decision-making capabilities before the controlled, rational part of the brain recognizes a general pattern. However, if the emotional reaction is too strong, the affective system can override the prefrontal

cortex, leading to suboptimal decisions. In particular, evidence suggests that activity in the brain areas associated with fear can lead to risk aversion.

In terms of financial markets, the alternating emotions of fear and excitement may play contrasting roles. As conditions of uncertainty stimulate a fear-response in investors, this evidence suggests that they may become risk-averse, and either terminate or refuse to enter positions. Conversely, as investors become more confident – and perhaps overconfident – leading to excitement and anticipation of monetary gain, this evidence suggests that their risk appetite may increase, potentially leading to riskier positions.

3.2 Finance and Philosophy of Mind

Thus far, I have spent a considerable amount of time discussing the methodologies of behavioral experiments and mathematical models as they relate to economics and finance, both from the perspective of the investor and the academic. This discussion would be incomplete, however, without mention of the methodology in which neuroscience relates to economics and finance. In moving psychological analysis from behavior to cognition, the natural question that arises is one of the relationship between the subjective experience of consciousness and neural states. This question is well beyond the scope of this project. Aside from taking for granted that Cartesian dualism[23] is false, I will not venture into this issue. Nevertheless, investors do not think in terms of their neural connections, but rather in terms of their beliefs and predictions about society, the economy, and financial markets. The connection between these two categories is a murky area, but using a similar methodology as outlined by Friedman (1953) can be theoretically useful.

Ross (2005), in relating philosophy of mind to general economic theory, notes the theoretical problem of representing the subjective and qualitative phenomena of *propositional attitudes* – which refer to mental states that are "somehow 'about' states of affairs in the world" (41) such as intentions, beliefs, and desires – in terms of objective neural states. For example, explaining an investor's decision to take a long position on gold by stating, "the investor wants to make money, and he believes gold will rise in value," would prescribe him *intentionality*[24]*,* as the explanation invokes the propositional attitudes of desires and beliefs that are "about" things external to the investor. However, while describing the subjective, internal state of the investor,

---

[23] This is the philosophy proposed by René Descartes, which states that all mental phenomena are completely separate from the body.
[24] The Stanford Encyclopedia of Philosophy defines intentionality as "the power of minds to be about, to represent, or to stand for, things, properties and states of affairs."

this sentence says nothing of the actual neural process underlying the behavior. Philosophers refer to this characterization as "folk psychology," as the description is not rooted in scientific observation, but rather common-sense beliefs about psychology and behavior. Ross (2005) notes, "One obvious possible idea here is that every particular intentional state – that is, every particular propositional attitude state or 'mental state' – might be identical with a particular brain state, directly identifiable and describable in the language of neuroscience" (39). Taken to its logical conclusion, however, Ross (2005) notes that this approach implies "no two creatures with significantly different brain states could, *by definition*, ever share the same belief" (40, emphasis in original). *Eliminative materialists*, in their view that it is impossible to map intentional states onto neural states in any sort of coherent manner, argue that propositional attitudes plainly do not exist.

The *externalist* school of thought argues that thought is a process that transcends isolated neural systems, and suggests that the intentionality of an investor has meaning only in relation to his environment (i.e., financial markets). A neural state can be "about" the price of gold, but this "aboutness" does not have any meaning outside of its relationship with the specific environment of the gold market. Since neural systems represent adaptations to certain environments, decontextualizing them from these environments would produce an incomplete picture of mental phenomena. Herrmann-Pillath (2009) notes, "There is no such thing as an ultimate neuronal cause of behavior, as all neuronal phenomena operate in conjunction with phenomena that reach beyond the physical boundaries of the skull" (26).

Daniel Dennett, an externalist, concedes that intentionality does not exist in any meaningful physical form, yet he takes an approach that bears a striking resemblance to the

methodology outlined in Friedman (1953).[25] Dennett takes the intentional-stance – that is, the stance in which intentionality is used as a predictive tool for cognitive systems – and posits that humans take the intentional-stance towards themselves and other humans. In this sense, the illusion of intentionality has a very real effect on financial markets, as investors' respective theories of mind draw quite heavily on "folk psychology." Because investors act *as if* other investors have intentionality, and because they make decisions based on this assumption, "folk psychology" has a significant influence over market behavior. Even if investors' thoughts can be isolated from their environment, they do not interpret the actions of other investors in terms neural connections, but in terms of beliefs, desires, and intentions; their respective theories of mind draw quite heavily on "folk psychology."

The viewpoint of externalism helps clarify a point of contention between the EMH and BF: the two schools adopt opposing views of investor intentionality. By modeling entire financial markets as a black box, in which the only input towards efficient prices is information, the EMH characterizes financial markets themselves as information processing systems. This approach is closer to an externalist viewpoint because, in viewing financial markets as a black box, the EMH makes it impossible to isolate individual investors from the larger computational process. The EMH uses efficiency in the same way Dennett uses intentionality – as a predictive, rather than descriptive, tool. In this context, improving the rationality of market participants would not make expected market prices subject to available information any more or less efficient; the behavior of prices is primarily due to forces external to the investor. Switching the market participants between human investors and *homo economici* would not alter the expected value of a security, though it might alter the error term.

---

[25] This approach describes the outcome of a process, rather than how the process actually functions. As Ross (2005) puts it, this provides "fictions that facilitate predictions" (162).

In isolating experimental participants from financial markets, and in arguing that these experiments alone are indicative of investor behavior, BF confines itself to a mostly *internalist* viewpoint.[26] The resulting implication of claiming that prices deviate from their efficient values is that this occurs due to irrational forces internal to individual investors. BF acknowledges that the structure of financial markets plays a role in inefficient pricing, but only insofar as transaction costs prevent the elimination of inefficiencies. Switching market participants from human investors to *homo economici* in this framework would have a drastically different impact on market prices – they moving the expected price closer to their efficient values.

Neuroeconomics, like BF experiments, taken by itself does not escape a confinement to internalism. Both behavioral and neuroscientific experiments provide useful observations, but extrapolating these observations directly to financial markets provides only a tenuous connection. Any theory of financial markets that hopes to provide a more accurate framework than the EMH must first place these observations in the proper context.

---

[26] As opposed to externalism, internalism states that thought can be described in terms internal to neural systems.

3.3 Information and Investor Cognition

Before an alternate framework to the EMH can be offered, the distinction between two types of efficiency, *price efficiency* and *computational efficiency*, must be made. The former describes the accuracy of prices with respect to available information, while the latter describes the accuracy of prices with respect to available information *and* computation. Although the EMH formally makes the claim that prices follow the first type of efficiency on average, it implicitly makes the claim that the market as a whole follows the second form of efficiency.

Many economic models view information as single-dimensional; they assume that information exists, and there is only one correct interpretation of it. This might be a helpful simplifying assumption in certain models, but may not be worth holding when examining information in the context of uncertainty. The EMH remains agnostic as to how information is processed – information goes into the system, something happens, and then efficient prices come out – but never gives a satisfying explanation of what constitutes information.

Although Hayek appears to use *information* and *knowledge* interchangeably, he does not take the time to draw a distinction between them. It is difficult to clearly define the two concepts in relation to each other because of their intimate relationship; but for the sake of formulating working definitions, *information* can be thought of as an unprocessed input of data. *Knowledge* can be thought of as higher-order information; that is, information that has been processed and contextualized at one level, and acts as information at a higher level. These definitions are admittedly vague, but are context-dependent relative to the scope of inquiry. For example, at the neurological level, the frequency of a sound wave constitutes information with respect to the cochlea within the ear, as can the subsequent electrical impulse sent from one neuron to another. As this information is ultimately sent to – and processed by – the temporal lobe, it is

conglomerated into knowledge (i.e., language), which can subsequently be used as information in higher thought processes.

Information, in its most basic form, represents communication between neurons and the external world. Sound waves constitute information with respect to the ear. Knowledge, or higher order information, represents propositional attitudes, such as the belief that this pattern of sound represents an aspect of security's value. Taking a broader viewpoint, the interaction between two investors, in which investor I tells investor J, "the CEO of Goldman Sachs is a fraud," can be viewed in the neurological context as a communication of propositional attitudes. On an even broader level, if investor J uses their composite knowledge of Goldman Sachs to make a decision regarding their stock, taking a short position for instance, this acts as information in relation to the stock price.

Although the two investors might understand each other perfectly well in this interaction, it is unlikely that they have precisely the same conception of "the CEO of Goldman Sachs" and fraudulence. They may understand each other when referring to these concepts, but their corresponding neural states are likely quite different. As the investors possess different *information sets*[27]*,* fraudulence and Goldman Sachs are placed in slightly different contexts once processed. Even as investors contain overlapping sets of information, certain pieces of information would play different roles in their analytical processes and trading strategies. Consider two investors – one specializing in international currency and one specializing in domestic equities – who, upon viewing the same news program, acquire information that the President intends to raise corporate taxes. Although both investors now possess the same piece of information, their total information sets are different – one contains detailed information relating

---

[27] That is, the collection of first- and higher-order information the investor is able to retain and make computational use of.

to equities, and the other contains detailed information relating to currencies. Subsequently, the currency investor may take a long position on foreign currencies, and the equity investor may take a short position on domestic equities. However, it is unlikely that the equity investor would take a position on currencies, and vice versa.

First-order information must go through some sort of computational process for it to be placed in a useful context. Computation, however, requires time and energy; as time and energy are scarce resources, an efficient computational mechanism must take these into account, along with information. Bounded Rationality, a term coined by Herbert Simon, provides description for computational efficiency. As opposed to the omniscient, substantive rationality of *homo economicus*, where "the consequences of any tautology are known as soon as the premises are stated" (Simon, 1979: 69), bounded rationality takes into account the computational efficiency of an agent. This distinction highlights the fact that systems with high degrees of computational efficiency can process information in shorter time periods, but not instantaneously. While an optimal solution might be evident from a set of information, it may take even a highly efficient computational device an extraordinary amount of time to arrive at a solution. Bounded rationality acknowledges the tradeoff between solution optimality and the resources expended in the computational process. As there are often diminishing marginal returns to computation, a boundedly rational solution is one that optimizes the expected outcome of a decision with respect to the amount of time and energy spent processing the information. Moreover, Simon argues that cognitive capacity is often limited. An investor might have access to countless numbers of financial metrics for various securities, but might not be able to retain each one simultaneously in memory. A boundedly rational decision in this context would entail, not investing in the most

optimal security[28], but rather investing in the best security given the amount of time, cognitive capacity, and information at the investor's disposal.

In this sense, the *price-efficient* value of a security represent an upper bound on the potential accuracy of the price; it is the price that is reached with respect to available information given infinite amounts of time and computational power. Both the EMH and BF use price efficiency to describe the accuracy of market prices with respect to available information at any given time. EMH proponents, such as Malkiel and Fama, note that, while prices may not always be price-efficient, prices are price-efficient on average. However, when describing historical incidences of obvious deviations "rational" values, it is not clear whether EMH proponents argue that this is caused by lack of first-order information, or lack time and energy to contextualize the available first-order information to create higher-order information. That is, where prices price-efficient, or computationally efficient? In any case, computational efficiency is a more adequate characterization of markets. In observing price-inefficient behavior, BF argues that prices sometimes deviate from their price-efficient values, but this does not say anything about whether financial markets are computationally efficient.

---

[28] That is, the security with the highest yield given the investors risk preference.

Conclusion – The Reflexive Market Hypothesis

Neither BF nor EMH proponents would dispute the fact that security prices are not always in a state of price efficiency. The disagreement between the two schools comes in explaining why prices deviate from price-efficient values; the EMH proposes that deviations occur randomly, while BF proposes that deviations are a function of investor irrationality. The empirical observations invoked by BF suggest that deviations from price efficiency are not random, but I propose that the typical BF explanation of this, investor irrationality, is a red herring. In the context of uncertainty, irrationality is an empty term. A more constructive question is whether investors are boundedly rational, and whether financial markets are computationally efficient. If investors are not boundedly rational, this implies that the BF explanation – that price-inefficiency is caused mainly by forces internal to investors – is correct. However, if investors are boundedly rational, and prices deviate from their price-efficient values in a non-random way, this implies that the deviation is due to forces external to the investor.

I propose the Reflexive Market Hypothesis (RMH), which states that deviations from price efficiency are caused by market uncertainty, but the strategy investors undertake to deal with this uncertainty, while boundedly rational, can potentially lead to non-random deviations from price efficiency. The amount of uncertainty investors face, and how they cope with this uncertainty on an emotional and cognitive level, has tremendous implications for financial markets. I borrow the term 'reflexive'[29] from Buenza and Stark (2010), who observe that derivative traders utilize reflexive modeling; that is, traders use their perceptions of each other as basis points for their investment decisions. While Buenza and Stark (2010) provides the central

---

[29] Investor George Soros, in his 1987 book *The Alchemy of Finance*, uses reflexivity in a different context – as investors affect security prices in potentially biased ways, the underlying fundamentals of the security change; this, in turn, affects the price of the security. While it is true that this relationship exists, I use reflexivity specifically in the context of market prices influencing investor's valuation of securities.

framework for the RMH, I believe model reflexivity is applicable to financial markets in a more general sense, and should not be confined to formal modeling. My incorporation of the findings of the neuroeconomics literature, along with my analysis of computational efficiency and uncertainty, further separates the RMH from the hypothesis put forward in Buenza and Stark (2010).

Despite Simon's progressive views of decision-making, he believed that, "like a modern digital computer's, Man's equipment for thinking is basically serial in organization" (Simon 1979: 72). In light of the neuroscientific research since 1976, however, this is resoundingly not the case. I will temporarily hold this as a simplifying assumption, though, along with the assumptions that all investors are homogenous in their computational capabilities, and they each possess a homogenous, complete set of information.

Consider an investor, *I*, who begins analyzing a financial position, *x*, in time *t*. Given no information in time *t*, *I* faces complete uncertainty; *I* knows nothing about *x*. After compiling a set of (first-order) information in time *t+1*, *I* still faces complete uncertainty regarding *x* because *I* has not yet had time to process the information. Only after *I* processes the information in time *t+2* will *I*'s uncertainty begin to diminish – the first order information is contextualized and transformed into higher order information. *I*'s uncertainty is further reduced in time *t+3*, but at a diminishing rate. Since *I* faces diminishing returns to retrieving and processing information, *I* must decide how many computational periods it wants to undertake.

Lifting the assumption of serial processing, the dual processing model of Camerer et al. (2005), where investors utilize either a controlled or automatic process, provides a good framework. The controlled strategy involves the serial, deductive faculties of the prefrontal cortex. This is characterized by the use of models, both formal and informal, as investors carry

out an internal step-by-step logic, based first on their (first-order) information sets, and then on their knowledge of financial markets. While this strategy is effective in reducing uncertainty, it requires a considerable amount of time and energy.

The automatic processing strategy involves parallel processing, and is characterized mainly by intuition[30] and heuristics. This strategy is much faster than the controlled strategy (Bechara et al., 1997), but does not succeed in reducing uncertainty as much as the controlled strategy during longer time intervals. The somatic marker hypothesis and risk-as-feeling hypothesis provide key components of this strategy; in using affective cues as a predictor, investors have a source of information external to their direct cognitive faculties. However, these cues can also lead towards strong emotions that can overpower their cognitive faculties. Neuroscientific evidence suggests that situations in which investors face large amounts of uncertainty might stimulate fear, and subsequently cause them to make risk-averse decisions, such as eliminating positions or refusing to enter positions. Alternatively, a perceived lack of uncertainty may stimulate a sense of excitement or anticipation of monetary reward, and cause excessively risky behavior.

In lifting the assumptions of complete, homogenous information sets and agents with homogenous computational power, financial market more closely resemble the market outlined by Hayek (1948), where the total amount of information across the market is fragmented, but overlaps amongst all market participants. This leaves investors with two problems. First, even if given infinite computational efficiency, their incomplete sets of information ensure that individual investors will face the induction problem outlined in chapter 1, and will thus always face uncertainty in their decisions. They will never be able to determine whether their profitable

---

[30] Although intuition is a complex process that this project does not completely explore, I here refer mainly to the brain utilizing affective cues.

(unprofitable) investments are due to their advantageous (disadvantageous) information set, or to the effectiveness of their computation. Second, even if investors are able to identify the price-efficient value of a security or position from their information set, the potential for less fortunate investors[31] to skew the price for long periods of time exists.

As a hedge against both of these problems, the investor supplements the object-relative facts contained in his information set with his perception of the object-dependent facts of other investors. Buenza and Stark (2010) provide an example of how using the state of others as informational inputs can be helpful:

> Consider the decision to carry and umbrella to work. Looking from one's own apartment window and seeing a mostly clear sky, one might decide it unnecessary to prepare for rain. But if one glanced below and found pedestrians carrying umbrellas, one might be prompted to check, from another vantage point, for an impending storm (31).

If an investor comes to the conclusion that the price of security x will rise from the current market price of $50 to $100 in one year, but reads multiple analyst reports that predict the price will fall to $10 in one year, this might be a cause for concern. It is possible that the investor stumbled upon a fantastic profit opportunity, but it is also possible that his initial computation was mistaken. The analysts give the investor another perspective, and might cause him to reevaluate his initial valuation. Alternatively, if the investor is not mistaken, viewing the analyst reports still provides useful information – it indicates that other investors might expect the price to fall to $10, and that the price may fall, even if its price-efficient value is $100.

Using current market prices may act as another useful informational input since it represents a weighted average of other investors' perception of the security's value. However, it

---

[31] Investors with either lesser endowed computational power, or less complete information sets.

is important to note that discerning object-dependent facts may not always be as obvious as reading analyst reports or viewing a single metric, as many of the people analyzing financial markets do not make their thoughts well known. Investors must infer how much information other investors have, and invoke a theory of mind to predict how the other investors will react to this information. In this case, an investor's perception of object-dependent facts must suffice.

An investor *I*'s perception of the object-relative facts of a security can differ from its perception of object-dependent facts for three reasons. Other investors can hold different beliefs about a security because they possess different information sets, other investors can possess similar information sets but come to more accurate conclusions about the security, or other investors can possess similar information sets and come to less accurate conclusions about the security. Suspicion that one of the first two scenarios is correct would increase *I*'s uncertainty, causing *I* to spend more time collecting information, spend more time processing information, or to not enter the position. Suspicion of the third scenario, however, might prompt *I* to enter into the position, or increase its position.

As noted by Buenza and Stark (2010), the strategy of using other investors as informational inputs can help protect against incomplete information and errors in processing information, but it also has the potential to amplify the effects of information that is systematically overlooked. In more general terms, this process can correct deviations from price-efficient values when it works as it is intended to, but also has the potential to amplify a deviation from price-efficient values. One benefit of the RMH is its flexibility in explaining these two scenarios. Rather than explain security prices in a singular state, as does the EMH and BF, the RMH provides a framework to view security prices in a variety of states as they move towards and away from price-efficient values in different magnitudes. For the sake of clarity,

however, I will describe this process working at two extremes: first, when the mechanism accomplishes its goal of reducing uncertainty, thus moving prices towards their price-efficient values, and then when the mechanism works in the opposite direction, pushing prices away from their price-efficient values. The first scenario exhibits negative feedback with regard to deviations, while the second scenario exhibits positive feedback.

In the first scenario, investors possess limited information sets, but, as in Hayek's (1948) model, they overlap. If security prices resemble their price-efficient values, then investors using their perceptions of object-dependent facts will have a reliable informational input. If they are missing a piece of information, or make an error in their valuation process, this mechanism will provide them with a check. This causes the investor to either search for more information or to reevaluate their initial conclusions. Alternatively, if market prices begin to deviate from their price-efficient values, the same process takes effect. However, as investors search for more information and spend more time processing their information sets, they confirm their initial conclusions. As they increase their positions, the market price for the security moves closer to its price-efficient value. In the short term, as new information is introduced, automatic processing provides the market with a quick schema of pricing securities subject to the new information. The extensive literature cited by EMH proponents suggests that the reflexive mechanism leads the market to resemble this scenario more often than not. However, the empirical studies cited by BF proponents note that, in at least some instances, prices deviate from their price-efficient values in non-random ways.

In the second scenario, the reflexive market mechanism either inspires a false sense of confidence in investors, or leads to an overreaction to new information. Both of these outcomes are functions of uncertainty, and this scenario will occur in situations where perceived

uncertainty is unusually high, or unusually low. If a piece of information is systematically overlooked due to fraud, an institutional friction, or by chance, use of object-dependent facts can inspire a false sense of confidence in investors. If an investor comes to an incorrect, but similar, conclusion as other market participants, using his perception of other investors as a reference point will fail to alert him of a computational error or lack of information. Conversely, if an investor comes to a correct conclusion, but this conclusion differs significantly from that of other market participants, he might become less confident in this conclusion. Whereas the previous scenario led to a confirmation of his original conclusion, and a subsequent increase in his position, this may not occur if he already faces significant uncertainty; viewing the contradictory conclusion of other investors might increase his uncertainty further, triggering a fear response and resulting in either a lack of entrance into, or termination of, his position. Alternatively, viewing contradictory object-dependent information might cause the investor to search for information that conforms to the incorrect conclusion. In this case, investors that originally arrive at correct conclusions fail to push prices towards their price-efficient values.

Deviations from the price-efficient value are not usually sustainable[32], however. Once conflicting information eventually becomes unambiguously apparent, investors' models are disrupted. As introduction of this information contradicts their model, investors are suddenly faced with a large amount of uncertainty. If market prices react quicker than the investor can adequately process this information, they must rely on their quicker, automatic strategy. As outlined by Loewenstein et al. (2001), strong emotions in this context can overwhelm the prefrontal cortex. If uncertainty becomes very high, investor may become fearful, and ultimately become risk-averse in terminating positions or refusing to enter into positions that would move prices back towards their price-efficient value.

---

[32] See Soros (1987).

As a descriptive theory, the RMH is supported by the neuroscience literature, and can be continually tested as neuroimaging technology advances. Specifically, as neuroimaging technology becomes less cumbersome and invasive, researchers will be able to undertake studies that are similar to, but more thorough than, Lo and Repin (2002) and Lo et al. (2005). The ability to study investors inside of financial environments is imperative for understanding market behavior for reasons expressed by the school of externalism in section 3.1. A research agenda in this vein may help confirm the micro aspect of the RMH, but the observations of neuroscientific studies alone do not necessarily imply that the predicted macro phenomena manifest themselves.

Agent-based modeling[33] has the potentially carify the RMH as a theory; as agents are programmed with computational and decision-making strategies similar to that specified in the RMH, the resulting phenomena could show whether the macro phenomena predicted by the RMH is consistent. But this would only show that the theory is tautologically correct. In order to test the empirical accuracy of the RMH, Lo and Repin (2002) and Lo et al. (2005) again provide a potential methodology, by comparing their observations with real-time market data. However, this method is restricted by its small sample size, and would only provide an accurate comparison if the participant group is representative of the general investment population.

Assuming that an experiment could satisfy the concern of a representative sample, the RMH – under an econometric time series analysis between neuroimaging data and market data – would imply that activity in the amygdala signals a perception of uncertainty, and this activity

---

[33] Agent-based models are computer models where the modeler programs multiple agents with different computational strategies. These strategies can range from simple to complex, and can incorporate various modes of information processing. As the agents interact with each other inside of a defined space, emergent phenomena manifests itself. This form of modeling, however, has two major problems that might potentially obscure it from reality. First, these models are incredible sensitive to initial conditions. Slight changes in agent strategies or environments have the potential to lead to drastically different phenomena. Because of this, very complex models obscure reality more than it explains it. Second, it would be quite difficult to account for human inductive capabilities, especially in the context of a simplified model.

acts as an accurate predictor of excessively risk-averse behavior. Alternatively, activity in the nucleus accumbens should act as an accurate predictor of excessively risk-seeking behavior.

The question that remains is one of whether this process is computationally efficient as a whole – that is, whether the reflexive market mechanism induces a convergence on price-efficient values in a high proportion to it its induction of a divergence from these values when controlling for time and energy. The second scenario highlights how financial markets can enter into and sustain – at least briefly – financial bubbles. But, are prices close to their price-efficient values often enough with respect to time and energy to classify the processing mechanism of financial markets as computationally efficient? Non-random deviations highlight the fact that this reflexive mechanism may not satisfy the efficiency criteria outlined by the EMH, but even though prices may sometimes deviate from their price-efficient values in non-random ways, this only means that financial markets are not *perfectly* computationally efficient. However, this does not mean that financial markets are necessarily computationally *in*efficient.

For academic economists, this potential discrepancy between market prices and their price-efficient values represent a direction for future research. For investors, this discrepancy represents potential opportunities for risk-adjusted profit. Identifying the magnitude of this discrepancy would be performative for both parties, as the belief in the existence of this discrepancy is a major underlying force of both the advancement of academic knowledge, and the motivation to exploit these discrepancies.

Works Cited

Alpert, M., & Raiffa, H. (1982). A Progress Report on the Training of Probability Assessors. In A. Tversky & D. Kahneman (Eds.), *Judgment Under Uncertainty: Heuristics and Biase*. Cambridge: Cambridge University Press.

Bachelier, L. (1900). *Théorie de la spéculation*. Paris: Cauthier-Villars.

Ball, R. (1978). Anomalies in Relationships Between Securities' Yields and Yield-Surrogates. *Journal of Financial Economics*, *6*(2-3), 103-126.

Ball, R., & Brown, P. (1968). An Empirical Evaluation of Accounting Income Numbers. *Journal of Accounting Research*, *6*(Autumn), 159-178.

Banz, R. W. (1981). The Relationship Between Return and Market Value of Common Stocks. *Journal of Financial Economics*, *9*, 3-18.

Bechara, A., Damasio, A. R., Damasio, H., & Anderson, S. W. (1994). Insensitivity to Future Consequences Following Damage to Human Prefrontal Cortex. *Cognition*, *50*, 7-15.

Bechara, A., Damasio, H., Tranel, D., & Damasio, A. R. (1997). Deciding Advantageously Before Knowing the Advantageous Strategy. *Science*, *275*, 1293-1295.

Black, F., Jensen, M., & Scholes, M. (1972). The Capital Asset Pricing Model: Some Empirical Tests. In M. C. Jensen (Ed.), *Study in the Theory of Capital Markets*. New York, NY: Praeger.

Buenza, D., & Stark, D. (2010). Models, Reflexivity and Systemic Risk: A Critique of Behavioral Finance. Retrieved from http://ssrn.com/abstract=1285054

Camerer, C., & Weber, M. (1992). Recent Developments in Modeling Preferences: Uncertainty and Ambiguity. *Journal of Risk and Uncertainty*, *5*, 325-371.

Camerer, C., Loewenstein, G., & Prelec, D. (2005). Neuroeconomics: How Neuroscience Can Inform Economics. *Journal of Economic Literture*, *43*, 9-64.

Damasio, A. R. (1996). The Somatic Marker Hypothesis and the Possible Functions of the Prefrontal Cortex. *Philosophical Transactions: Biological Sciences*, *351*(1346), 1413-1420.

De Martino, B., Camerer, C. F., & Adophs, R. (2010). Amygdala Damage Eliminates Monetary Loss Aversion. *PNAS*, *107*(8), 3788-3792.

DeBondt, W., & Thaler, R. H. (1985). Does the Stock Market Overreact. *Journal of Finance*, *40*(3), 793-805.

DeBondt, W., & Thaler, R. H. (1987). Further Evidence on Investor Overreaction and Stock Market Seasonality. *Journal of Finance*, *42*(3), 105-137.

Ellsberg, D. (1961). Risk, Ambiguity, and the Savage Axioms. *The Quarterly Journal of Economics*, *75*(4), 643-669.

Fama, E. F. (1965). The Behavior of Stock Market Prices. *Journal of Business*, *38*(1), 34-105.

Fama, E. F. (1970). Efficient Capital Markets: A Review of Theory and Empirical Work. *The Journal of Finance*, *25*(2), 383-417.

Fama, E. F. (1991). Efficient Capital Markets: II. *The Journal of Finance*, *46*(5), 1575-1617.

Foster, G., Olsen, C., & Shevlin, T. (1984). Earnings Releases, Anomalies, and the Behavior of Security Returns. *The Accounting Review*, *59*(4), 574-603.

French, K. (1980). Stock Returns and the Weekend Effect. *Journal of Financial Economics*, *8*(1), 55-69.

Friedman, M. (1953). The Methodology of Positive Economics. In *Essays in Positive Economics* (pp. 3-43). Chicago: University of Chicago Press.

Friedman, M. (676). The Methodology of Positive Economics. In *Philosophy and Economic Theory*. (Original work published 1953)

Haugen, R. A., & Lakonishok, J. (1988). *The Incredible January Effect: The Stock Market's Unsolved Mystery*. Homewood, IL: Dow Jones-Irwin.

Hayek, F. A. (1948). The Use of Knowledge in Society. In *Individualism and Economic Order* (pp. 77-91). London: The University of Chicago Press.

Hayek, F. A. (1968). *Competition as a Discovery Procedure*. Lecture presented in University of Kiel, Kiel.

Hebb, D. O. (1949). *The Organization of Behavior: A Neuropsychological Theory.* New York, NY: John Wiley & Sons.

Herrmann-Pillath, C. (2009). A Neurolinguistic Approach to Performativity in Economics. *Frankfort School - Working Paper Series*, 123rd ser.

Hsu, M., Bhatt, M., Adolphs, R., Tranel, D., & Camerer, C. (2005). Neural Systems Responding to Degrees of Uncertainty in Human Decision-Making. *Science*, *310*, 1680-1683.

Hume, D. (2000). *A Treatise on Human Nature*. Oxford: Oxford University Press.

   (Original Work Published 1748)

Jacob, P. (2010, August 31). Intentionality. In *Stanford Encycolpedia of Philosophy*. Retrieved April 23, 2011, from http://plato.stanford.edu/entries/intentionality/

Jegadeesh, N., & Titman, S. (1993). The Returns to Buying Winners and Selling Losers: Implications for Stock Market Efficiency. *Journal of Finance*, *48*(1), 65-91.

Jensen, M. C. (1978). Some Anomalous Evidence Regarding Market Efficiency. *Journal of Financial Economics*, *6*(2), 95-101.

Jensen, M. C. (1978). Some Anomalous Evidence Regarding Market Efficiency. *Journal of Financial Economics*, *6*(2), 95-101.

Kalat, J. W. (2009). *Biological Psychology* (10th ed.). Belmont, CA: Wadsworth, Cengage Learning.

Kendall, M. G. (1953). The Analysis of Economic Time-Series, Part I: Prices. *Journal of the Royal Statistical Society*, *116*(1), 11-25.

Kühberger, A., & Perner, J. (2003). The Role of Competition and Knowledge in the Ellsberg Task. *Journal of Behavioral Decision Making*, *16*(3), 181-191.

Knight, F. H. (1921). *Risk, Uncertainty, and Profit*. Boston: Houghton Mifflin.

Knutson, B., Kaufman, M., Taylor, J., Peterson, R., & Glover, G. (2005). Distributed Neural Representation of Expected Value. *The Journal of Neuroscience*, *25*(19), 4806-4812.

Kuhnen, C. M., & Knutson, B. (2005). The Neural basis of Financial Risk-Taking. *Neuron*, *47*(5), 763-770.

Lo, A. W., & Repin, D. V. (2002). The Psychopysiology of Real-Time Financial Risk Processing. *Journal of Cognitive Neuroscience*, *14*(3), 323-339.

Lo, A. W., Repin, D. V., & Steenbarger, B. N. (2005). Fear and Greed in Financial Markets: A Clinical Study of Day-Traders. *Cognitive Neuroscientific Foundations of Behavior*, *95*(2), 352-359.

Loewenstein, G. F., Weber, E. U., Hsee, C. K., & Welch, N. (2001). Risk as Feeling. *Psychological Bulletin*, *127*(2), 267-286.

Lord, C. G., Ross, L., & Lepper, M. R. (1979). Biased Assimilation and Attitude Polarization. *Journal of Personality & Social Psychology*, *37*(11), 2098-2109.

MacKenzie, D. (2006). *An Engine, Not a Camera: How Financial Models Shape Markets*. London: The MIT Press.

Malkiel, B. G. (2003). The Efficient Market Hypothesis and Its Critics. *Journal of Economic Perspectives*, *17*(11), 59-82.

Mandelbrot, B. (1966). Forcasts of Future Prices, Unbiased Markets, and Martingale Models. *Journal of Business*, *39*(1), 245-255.

Montague, P. R., & Berns, G. S. (2002). *Neuron*, *36*, 265-284.

Ross, D. (2005). *Economic Theory and Cognitive Science: Microexplanation*. Cambridge, MA: MIT Press.

Samuelson, P. A. (1965). Proof That Properly Anticipated Prices Fluctuate Randomly. *Industrial Management Review*, *6*(1), 41-49.

Searle, J. R. (n.d.). What Is an Institution. *Journal of Institutional Economics*, *1*(1), 1-22.

Shiller, R. J. (2005). *Irrational Exuberance*. Princeton, NJ: Princeton University Press.

Shiv, B., Loewenstein, G., Bechara, A., Damasio, H., & Damasio, A. R. (2005). Investment Behavior and the Negative Side of Emotion. *Psychological Science*, *16*(6), 435-439.

Simon, H. (1979). From Substantive to Procedural Rationality. In F. Hahn & M. Hollis (Eds.), *Philosophy and Economic Theory* (pp. 63-86). Oxford: Oxford University Press. (Original work published 1976)

Soros, G. (1987). *The Alchemy of Finance: Reading the Mind of the Market*. New York, NY: Simon and Schuster.

Taleb, N. (2007). *The Black Swan: The Impact of the Highly Improbable*. New York, NY: Random House.

Thaler, R. H., & Barberis, N. (2005). A Survey of Behavioral Finance. In R. H. Thaler (Ed.), *Advances in Behavioral Finance* (Vol. 2, pp. 1-65). New York, NY: Princeton University Press.

Tversky, A., & Kahneman, D. (1974). Judgment under Uncertainty: Heuristics and Biases. *Science*, *185*(4157), 1124-1131.

Tversky, A., & Kahneman, D. (1979). Prospect Theory: An Analysis of Decision Under Risk. *Econometrica*, *47*(2), 263-292.