4-24-2013

Decision-Making Amplification Under Uncertainty: An Exploratory Study of Behavioral Similarity and Intelligent Decision Support Systems

Merle Campbell
Georgia State University

PERMISSION TO BORROW

In presenting this dissertation as a partial fulfillment of the requirements for an advanced degree from Georgia State University, I agree that the Library of the University shall make it available for inspection and circulation in accordance with its regulations governing materials of this type. I agree that permission to quote from, to copy from, or publish this dissertation may be granted by the author or, in his/her absence, the professor under whose direction it was written or, in his absence, by the Dean of the Robinson College of Business. Such quoting, copying, or publishing must be solely for the scholarly purposes and does not involve potential financial gain. It is understood that any copying from or publication of this dissertation which involves potential gain will not be allowed without written permission of the author.

Merle W. Campbell

NOTICE TO BORROWERS

All dissertations deposited in the Georgia State University Library must be used only in accordance with the stipulations prescribed by the author in the preceding statement.

The author of this dissertation is:

Merle W. Campbell
800 Autry Landing Cir
Johns Creek, GA 30022

The director of this dissertation is:

Mark Keil
Department of Computer Information Systems
Robinson College of Business
Georgia State University
P.O. Box 4015
Atlanta, GA
30302-4015

Decision-Making Amplification Under Uncertainty: An Exploratory Study of Behavioral Similarity and Intelligent Decision Support Systems

BY

Merle Wayne Campbell

A Dissertation Submitted in Partial Fulfillment of the Requirements for the Degree

Of

Executive Doctorate in Business

In the Robinson College of Business

Of

Georgia State University

GEORGIA STATE UNIVERSITY
ROBINSON COLLEGE OF BUSINESS
2013

ACCEPTANCE

This dissertation was prepared under the direction of the Merle Wayne Campbell Dissertation Committee. It has been approved and accepted by all members of that committee, and it has been accepted in partial fulfillment of the requirements for the degree of Executive Doctorate in Business in the J. Mack Robinson College of Business of Georgia State University.

H. Fenwick Huss, Dean

DISSERTATION COMMITTEE

Mark Keil (Chair)
Peter Zhang
Yusen Xia

ABSTRACT

Decision-Making Amplification Under Uncertainty: An Exploratory Study of Behavioral Similarity and Intelligent Decision Support Systems

BY

Merle Wayne Campbell

April 25, 2013

Committee Chair: Mark Keil

Major Academic Unit: Robinson College of Business

Intelligent decision systems have the potential to support and greatly amplify human decision-making across a number of industries and domains. However, despite the rapid improvement in the underlying capabilities of these "intelligent" systems, increasing their acceptance as decision aids in industry has remained a formidable challenge. If intelligent systems are to be successful, and their full impact on decision-making performance realized, a greater understanding of the factors that influence recommendation acceptance from intelligent machines is needed.

Through an empirical experiment in the financial services industry, this study investigated the effects of perceived behavioral similarity (similarity state) on the dependent variables of recommendation acceptance, decision performance and decision efficiency under varying conditions of uncertainty (volatility state). It is hypothesized in this study that behavioral similarity as a design element will positively influence the acceptance rate of machine recommendations by human users. The level of uncertainty in the decision context is expected to moderate this relationship. In addition, an increase in recommendation acceptance should positively influence both decision performance and decision efficiency.

The quantitative exploration of behavioral similarity as a design element revealed a number of key findings. Most importantly, behavioral similarity was found to positively influence the acceptance rate of machine recommendations. However, uncertainty did not moderate the level of recommendation acceptance as expected. The experiment also revealed that behavioral similarity positively influenced decision performance during periods of elevated uncertainty. This relationship was moderated based on the level of uncertainty in the decision context. The investigation of decision efficiency also revealed a statistically significant result. However, the results for decision efficiency were in the opposite direction of the hypothesized relationship. Interestingly, decisions made with the behaviorally similar decision aid were less efficient, based on length of time to make a decision, compared to decisions made with the low-similarity decision aid. The results of decision efficiency were stable across both levels of uncertainty in the decision context.

TABLE OF CONTENTS

ABBREVIATIONS AND ACRONYMS

AF-ANN	Adaptive Feed-forward Artificial Neural Network
AI	Artificial Intelligence
ANN	Artificial Neural Network
DA	Decision Aid
DM(s)	Decision Maker(s)
DSS	Decision Support Systems
ES	Expert System
HCI	Human Computer Interaction
ITM	In-the-Money
IDSS	Intelligent Decision Support Systems
MLP	Multiple Layer Perceptron
MSE	Mean Squared Error
OTM	Out-of-the-Money
RA	Recommendation Agent
VIX	Chicago Board of Options Exchange Volatility Index

LIST OF TABLES

LIST OF FIGURES

1. INTRODUCTION AND THEORETICAL BACKGROUND

1.1 Introduction

Intelligent systems are an increasingly important strategic component of an organization's information systems portfolio (Hayes-Roth, 1997; Hayes-Roth & Jacobstein, 1994). The importance of these systems has grown exponentially in the last decade due to advances in the fields of computer science and artificial intelligence (AI), where powerful new “intelligent” technologies have evolved. Intelligent systems have the potential to greatly expand support to decision-makers (DMs) in complex problem-solving domains (Roth, Bennett, & Woods, 1987). The question that we continue to face in practice, however, is how to deploy the power available through these intelligent systems to improve human decision-making. One of the key challenges in this pursuit is the fact that DMs are often reluctant to accept advice from these intelligent decision aids in practice. This phenomenon is particular problematic in the field of investment management, where it is unlikely that the true benefits of intelligent decision support technologies will ever be fully realized if DMs fail to listen to them.

Improving support for DMs under uncertainty is of particular interest in the investment management profession, where the allocation of capital across a seemingly unlimited number of investment alternatives is one of the most critical decisions in the investment management process. In this context, the long-term viability of an investment portfolio depends on the trader’s ability to minimize portfolio volatility (risk) while simultaneously working to capitalize on viable opportunities. To accomplish this successfully, portfolio-trading decisions must be made both accurately and without delay. The basic mechanics of this process currently falls on the shoulders of human decision makers, where emotion, biases, intuition, and heuristics often play a large role in the portfolio decision-making process. This is in direct contradiction to the neoclassical postulate in economics that individuals possess rational expectations and strive to maximize expected utility in financial decision-making (Lo, 2005; von Neumann & Morgenstern, 1947). Furthermore, the high-velocity and uncertain nature of market data continuously challenges the cognitive capacities of those involved in the investment management domain. In this context, Hayes (1962) highlighted the fact that when making decisions by evaluating evidence, and the number of decision parameters is greater than four, human decision-making rapidly deteriorates.

The fact that individuals often make sub-optimal investment decisions due to cognitive limitations is also well established based on empirical and laboratory research (Bhandari, Hassanein, & Deaves, 2008; Kroll, Levy & Rapoport, 1988). This is due in part to the known issues that individuals have with representing probability and risk. Individuals often have the tendency to base their decisions on subjective probabilities, rather than the more concrete and fact based objective probabilities that result from careful and deliberate analysis of available data. As a result, the asymmetry that often exists with respect to an individual's perceptions of probabilities and actual probabilities can adversely impact decision making (Harrison & Rutstrom, 2008; Savikhin, Lam, Fisher, & Ebert, 2011).

As highlighted by the recent financial crisis, financial market phenomena can evolve rapidly, and sometimes without detection. Traders are constantly responding to market shocks of unknown origin, which can further impact their performance as decision makers. In addition, the flow of data and information continues to grow in scope and complexity given the increased use of algorithmic and high-frequency trading, making it increasingly difficult for DMs to maintain rational behavior under periods of extreme stress. Camerer et al. (2004) point out that these types of perturbations in the decision context make it difficult for DMs to adhere to the normative axioms of inference and choice. Similarly, Kahneman and Tversky (1982) referred to this particular phenomenon as an application error, in that the DM possesses the requisite cognitive skills to make an appropriate decision, but exogenous factors in the decision context inhibit the effective application of these skills. The literature on human error also provides insight into this phenomenon, and refers to this type of error as a slip (Zhao & Olivera, 2006). Slips are situations where an individual has the requisite knowledge on how to execute a specified task, but does not carry it out appropriately due to internal or external distractions (Rizzo et al., 1987; Stewart & Chase, 1999).

Sub-optimal decision making, resulting from application errors and slips is a major detractor of decision performance in the investment management profession, where even modest amounts of downside variance can compromise overall portfolio performance. Over time, poor decision making in portfolio management and trading can place a financial services firm at a competitive disadvantage. This is due in part to the proliferation of financial services firms offering trading advice in the last decade, where the competition for investor assets is fierce. In this context firms place a great deal of marketing emphasis on their trading performance relative to that of their competitors. In addition, published market proxies and indices are used by investors as benchmarks to evaluate the relative performance of an investment strategy. In order for a firm to attract and maintain investor assets it has to perform well relative to its peers and its assigned benchmark. As a result, improving decision performance in investment management is a major initiative for the industry at large.

In an effort to improve decision making performance, many firms in the financial services industry have experimented with various forms of Decision Support Systems (DSS). The adoption of such systems has been an effort to improve decision performance by reducing the cognitive demands placed on DMs, and to assist them in volatile markets. Of more recent interest in the financial services industry is the utilization of artificial intelligence (AI) based-technologies like Expert Systems (ES), knowledge bases, fuzzy logic, multi-agent systems, natural language, genetic algorithms, and artificial neural networks (Sousa et al., 2007). Often referred to as Intelligent Decision Support Systems (IDSS), these systems are intended to mimic and capture certain salient and beneficial characteristics of human decision-makers, such as approximate reasoning, intuition, and common-sense (Jackson, 1999). The distinguishing feature of IDSS is that they are often designed to provide a computerized representation of both tacit and explicit human knowledge (Gregor & Benbasat, 1999). However, in practice, these intelligent systems have failed to provide any meaningful performance results given their low acceptance rates as decision aids in securities trading. This is due in no small part to the fact that the designers of these systems have focused on how to build better performing systems, as opposed to focusing on techniques to increase their utilization as part of a joint human-machine cognitive system.

The impact of intelligent decision aids on decision-making process and outcome have been studied extensively by researchers (Gupta et al., 2006; Linger & Burstein, 1997; Moreau, 2006; Phillips-Wren & Jain, 2005; Roth et al., 1987). However, the research done so far is by no means exhaustive, and the influence of IDSS as part of a joint human-machine cognitive system on decision performance and efficiency under conditions of uncertainty remains a topic of interest for practitioners and researchers. Moreover, since the value of information systems and technology tends to be influenced by their actual use in decision-making (Devraj & Kohli, 2003), a greater understanding of the utilization of IDSS in a real-time decision context is also needed.

Given the heterogeneous nature of the design and architectural elements used in many modern IDSS, a greater understanding of the use of specific types of design features in IDSS will also contribute to both theory and practice. Intelligent systems in general, and IDSS in specific, will never be fully accepted in industry if we do not have a more robust tacit and theoretical understanding of their impact on decision-making performance.

In an attempt to explore the impact of IDSS as part of a joint human-machine cognitive system, a theoretically grounded prototype system was developed for use in an experiment in the financial services industry. The IDSS prototype for this experiment builds on the theoretical foundation established by the Computers are Social Actors (CASA) paradigm (Nass , Steuer , Tauber & Reeder, 1993 ; Nass , Steuer & Tauber , 1994 ; Reeves & Nass , 1996). The CASA research demonstrates that many of the social rules and dynamics that apply in human – human interaction can apply to human –computer interaction. The CASA research also provides support for the notion that technological artifacts can be perceived as social actors by their human usersand as a result, users often project behavioral attributions towards them (Reeves & Nass, 1996).

Building on the CASA research, the IDSS prototype was specifically designed using a theoretically grounded design element intended to evoke the perception of behavioral similarity between the human user and the IDSS. The theoretical basis for this design element is derived from the "similarity-attraction hypothesis" which predicts that people prefer to interact with others who are perceived to be similar to themselves.

Similarity attribution is well founded in the literature with respect to on-line e-commerce websites, and on-line product recommendation agents (Al-Natour et al., 2005; Aksoy 2006; Komiak & Benbasat, 2006). Similarity attribution has also enjoyed a great deal of interest in practice in domains like marketing and e-commerce where companies like Amazon, Apple iTunes, and Groupon are working to personalize product recommendations and marketing offers to consumers, based on the purchase behavior and perceived preferences of similar consumers.

Despite the growing support for similarity in research and practice, the concept of similarity as a design element has not been meaningfully extended beyond the e-commerce domain. In addition, much of the research on similarity has been conducted in relatively structured settings, without considering the decision context. In an attempt to further the use of behavioral similarity as a design element this dissertation aims to contribute to the literature by exploring the hypotheses that the use of a behaviorally similar IDSS should positively influence (1) the acceptance rate of artifact recommendations, (2) decision performance and (3) decision efficiency. Market volatility, as a surrogate measure for uncertainty in the decision context, is expected to moderate the hypothesized relationships.

The dissertation is organized as follows. The first chapter provides the background, rationale, and significance of the study. Chapter 1 also provides the theoretical and conceptual framework for the study, as well as the research questions. Chapter 2 provides a review of the literature supporting the study. The literature review encompasses four primary domains: Human-Computer Interaction, Trust in Information Technology, Decision Support Systems, and Artificial Intelligence for Knowledge Acquisition. Chapter 3 explicates the research model and hypotheses used for the study. Chapter 4 then provides an overview of the research methods used to test the respective hypotheses. This chapter provides an overview of the design of the experiment, as well as an introduction to the prototype IDSS that will be used in the execution of the experiment. The findings and conclusions are presented in Chapter 5, and limitations are presented in Chapter 6. Chapter 7 highlights the expected contributions to theory and practice, while providing some perspective on opportunities for future research.

1.2 Background

1.2.1 Motivation for the Study

Over the last two decades, advances in the fields of computer science and artificial intelligence (AI) have provided powerful new computational tools to DMs. These “intelligent” tools have the potential to greatly enhance cognitive capability and decision-making in complex problem-solving domains. Furthermore, these intelligent decision aids are increasingly being considered as "partners" and "teammates" that support or assist the human DM in performing complex functions and tasks. However, the question that we continue to face both in theory and practice is how to design these new tools in ways that increase their acceptance by DMs. These intelligent decision aids will be incapable of positively influencing investment decision-making if they are not designed in a way that encourages utilization. Yet despite this open question, only a limited amount of research has been done to explore the actual impacts of these intelligent systems on user decisions (Wong and Monaco, 1995). As a result, the overriding motivation for this study is to explore the use of a specialized IDSS to improve the acceptance of advice from an intelligent system, as well as to improve both decision performance and efficiency under conditions of uncertainty.

1.2.2 Significance of the Study

The development of an intelligent decision aid that is capable of improving decision performance and efficiency holds a great promise for the financial services industry, where the environment is punctuated by uncertainty, extreme complexity and growing competition. To better contend with this environment portfolio traders need decision aids that are adaptive, can cope with variability and that are capable of providing support in times of extreme entropy. Most importantly, these systems need to be designed in a way that traders will actually use them. As a result, this study develops and tests a prototype in a real-time, semi-structured decision context. A particularly significant element of this prototype relates to the design of its knowledge base (KB). The KB is constructed through the replication of a trader's individual decision strategy in order to foster the perception of behavioral similarity. The basis for this approach to KB construction is derived from the work of Malakooti and Zhou (1994), who highlight the fact that a Multiple Attribute Utility Function (MAUF) exists for all decision makers, and that an individual's MAUF can be captured and replicated using Adaptive Feed-forward Artificial Neural Networks (AF-ANNs). If found effective, this type of intelligent system can provide new business opportunities for firms in high-velocity markets where the decision domain is often uncertain and semi-structured, and the acceptance rate of decision technologies is low. From a theoretical perspective, HCI and DSS researchers could benefit from a greater understanding of design features that can impact recommendation acceptance and decision-performance of intelligent systems in semi-structured domains.

1.2.3 Theoretical and Conceptual Framework

With respect to theoretical background, this thesis describes an exploratory approach to the design, implementation, and evaluation of an IDSS based on the theoretical premise of behavioral similarity. In developing the prototype IDSS as part of the research study, there are four distinct streams of literature that serve as a foundation. HCI is the basis for the theoretical foundation because understanding how humans perceive and interact with intelligent agents in the exchange of knowledge and reason is critical to the design of the prototype. The second theoretical element relates to the construct of trust in information technology. Trust is an important multi-dimensional construct that has a tremendous influence as an antecedent in IT adoption behavior (Mayer et al., 1995; McKnight et al., 2002; Muir, 1987; Xiao & Benbasat, 2002). The literature on DSS represents the third theoretical element, supporting the notion that DSS and related technologies can be used in support of all stages of the decision-making process. Lastly, theoretical elements regarding the use of AI to enhance DSS were used. This element is important to highlight the use of AI technology to create intelligent and adaptive systems for handling complex semi-structured problems. This body of literature will also support the design of the behaviorally similar knowledge base (KB) for the prototype IDSS.

1.3 Theme and Research Question

Presently available decision support technologies make it possible to significantly amplify the intellect of a human decision-maker. In addition, it has been established that humans and computers possess complementary information processing capabilities, and therefore, significant advantages may be achieved by fostering a symbiotic relationship between human and machine (Felsen, 1975). As a result, computers should be used to complement rather than substitute human judgment when solving complex non-linear decision problems. While the literature is replete with studies of human-computer interaction (HCI) and IT adoption, what appears to be less studied is the use of specific design features to influence the acceptance of advice from an IDSS for purposes of improving decision-making in an uncertain decision context. As a result, this study explored the following research questions:

1. Can perceived behavioral similarity positively influence the frequency by which a human DM relies on advice from an IDSS under conditions of uncertainty?
2. Can perceived behavioral similarity positively influence the decision-making performance of a joint human-IDSS cognitive system under conditions of uncertainty?
3. Can perceived behavioral similarity positively influence the decision-making efficiency of a joint human-IDSS cognitive system under conditions of uncertainty?

2. REVIEW OF EXISTING LITERATURE

2.1 Human-Computer Interaction

A variety of theoretical perspectives are used in the design and implementation of this research study. In particular, the field of human-computer interaction (HCI) represents the nucleus of the theoretical structure. Many of the theoretical elements of HCI are derived from a multitude of research domains, and owe their theoretical origin to studies of human-human interaction (HHI). These theories provide valuable insight on how humans trust, perceive others' behavior, exchange knowledge, share opinions, and coordinate activities. Elements of HCI, coupled with the field of distributed artificial intelligence, provide insight as to how humans perceive and interact with intelligent agents to exchange knowledge and reason about goals and actions. Research on HCI, particularly elements from cognitive engineering, computer-supported cooperative work, and anthropological perspectives highlight features of computer systems that have the ability to engender effective joint problem solving (Jones & Mitchell, 1995). Theory related to the human user's perception of a technological artifact is also an important building block in the theoretical structure of the current research study.

2.1.1 Technological Artifacts as Social Actors

Given the rapid increase in the level of sophistication and intelligence of modern computer systems, as well as the integral role they play in our daily lives, these systems are increasingly ascribed attributes which are often analogous to those of humans. Some researchers have argued that such human-computer teams function similarly to human-human teams (Bowers, Oser, Salas, & Cannon-Bowers, 1996). Furthermore, researchers have provided evidence suggesting that people do enter into relationships with computers, robots, and interactive machines in a manner similar to other humans (Nass et al., 1996; Reeves & Nass, 1996).

Al-Natour, Benbasat, and Cenfetelli (2005) found that humans can perceive technological artifacts as social actors, and that human users can make personality and behavioral attributions towards them. An example of utilizing elements of human personality to measure personalities of inanimate objects is found in Nass et al. (1995). Nass et al. conducted a number of experiments endowing technology artifacts with human-like personalities. In an experiment with 48 subjects, dominant and submissive subjects were randomly matched with a computer with either a dominant and submissive trait. When asked to work with a computer on a problem-solving task, subjects were attracted to the computer that demonstrated a personality characteristic similar to their own. Furthermore, subjects found the interaction with the computer more satisfying, when they were utilizing a machine that had a similar personality trait. The results of these experiments reveal that personality attributions can be based on certain system attributes like voice, text, or physical representation, and even the most superficial manipulations are sufficient to produce personality. Reeves and Nass (1996) found that even technologically sophisticated users treat technological artifacts as if they were human beings, as opposed to being simple tools.

Qui and Benbasat (2009) investigated the effects of integrating anthropomorphic interfaces, like humanoid embodiment and voice output, on users' perceived social relationship with a technological artifact. In the design of their experiment, Qui and Benbasat (2009) utilized an animated avatar and voice output in an e-commerce website for selecting a digital camera. The findings from this laboratory experiment indicated that using humanoid embodiment and human voice–based communication significantly influenced users' perceptions of social presence in the artifact. This increased users' intentions to use the anthropomorphized artifact as a decision aid through enhancing users' trusting beliefs and perceptions of enjoyment.

2.1.2 Behavioral Similarity Theory

One of the key findings from the Nass et al. (1995) experiment was the fact that a user's perception of a technological artifact could be manipulated in way that created a feeling of similarity between the user and the artifact. The basis for this relationship is the "similarity-attraction hypothesis" which predicts that people will prefer to interact with others who are similar in personality to themselves. Byrne et al. (1967) substantiated the claim that attraction between humans is a positive linear function of the proportion of similar characteristics. Similarity is attractive to humans because a shared belief structure can provide validation of personal views, and can result in fewer disagreements and conflicts among parties (Byrne et al., 1967).

Furthermore, research on this theory suggests that similarity plays an important role in persuasion, cooperation, commerce, and the formation of opinion (Aksoy et al., 2006). Furthermore, endorsers who are perceived as similar to their audience have been shown to have more influence in changing attitudes and opinions (Haas 1981; Simons, Berkowitz, & Moyer 1970). In this context, Mathews, Wilson, and Monoky (1972) conducted an experimental study of cooperative behavior in a buyer-seller dyad, focusing on the effect of perceived similarity of characteristics upon cooperative behavior. It was hypothesized that buyer-seller dyads in which the individuals perceive themselves as being similar would achieve more cooperation, in contrast to dyads in which the individuals perceive themselves as being dissimilar. Their study revealed that perceived similarity between negotiators can increase the number of cooperative responses. Mathews, Wilson, and Monoky (1972) also highlighted the fact that the illusion of similarity is an important consideration in buyer-seller interaction. Evans (1963) found that similarity between buyers and sellers increases the probability of a successful sale. And Busch and Wilson

(1976) found that perceived similarity has a positive effect on salesperson trust and influence.

Much of the research on perceived similarity has been extended from HHI to HCI, in the evaluation of the dynamics between human user and advice giving technological artifact. The majority of this research is in the form of e-commerce recommendation agents and websites, as well as on-line customer decision aids (Aksoy et al., 2006; Comic & Benbasat, 2006; Al-Natour et al., 2005). These studies provide a great deal of theoretical support to the notion that perceived similarity can transcend the HHI context to influence human-machine interaction dynamics. In particular, computers that seemingly behave in ways that are similar to humans may promote more cooperative behavior from consumers (Moon 2000).

E-commerce and recommendation agent (RA) researchers have recently posited that websites should be designed with the goal of building relationships and improving the end-user's experience (Al-Natour & Benbasat, 2009). Aksoy et al. (2006) proposed that if an online recommendation agent (RA) is perceived to behave in ways that are similar to a human consumer, based on a seemingly similar decision-making process, consumers should be more likely to accept the RA's product recommendations. This topic was explored via two laboratory experiments in which participants searched and chose cellular phones from an online website. In the experiment, the similarity of attribute weights and perceived decision strategy similarity were manipulated in a database to determine their influence on participants.

The results of the Aksoy et al. experiment indicated that the perceived benefits of the RA were higher when a decision strategy similar to that of the consumer was used, resulting in higher choice quality and reduced search. Another important finding from the study was the impact of perceived dissimilarity between the consumer and an RA. Specifically, dissimilarity in both attribute weights and decision strategies were found to have negative effects on consumer choices and website loyalty. While these results provide useful insight as to the benefits of perceived similarity in consumer interaction with a website RA, additional insight is needed in terms of evaluating perceived decision strategy in more complex artifacts, like intelligent systems for decision support. And while the results were found tractable in an relatively structured e-commerce domain, very little literature exists with respect to exploring perceived decision similarity in semi-structured decision domains punctuated by uncertainty and risk.

Fostering the human DM's perception of similarity with respect to the machine's decision process represents a non-trivial element of this thesis. To accomplish this, the theoretical foundation established by Al-Natour et al. (2008) was utilized. Based on a review of the similarity attraction literature, the study by Al-Natour et al. (2008) represents one of the most comprehensive investigations of perceived decision process and outcome similarity in decision aids (DAs). Specifically, the authors investigated the impact that the constructs of perceived decision process and outcome similarity had on a human DM's evaluation of an e-commerce DA. To conduct their investigation, the authors used a laboratory setting in which subjects performed an online shopping task for a laptop computer. A DA was provided to offer product-specific information and recommendations to a user, and the DA was manipulated to investigate the effect of users' perceptions of the similarity between their own decision process and that followed by the DA to arrive at a product recommendation. The outcome of this study showed that perceived process similarity resulted in positive and significant effects on users' perceptions of usefulness and trustworthiness in a DA. The Al-Natour et al. (2008) study advanced the earlier efforts of Al-Natour et al. (2006), who investigated the role of design characteristics in forming social perceptions about an on-line shopping assistant. In an experiment using an on-line shopping assistant in a structured decision-task, it was found that both perceived personality similarity and perceived behavioral similarity between the human user and the technology artifact, positively affected users' evaluations of the DA.

While much of the aforementioned research provides support for an important theoretical element of the current research study, several voids in knowledge remain. Principally, the use of perceived similarity was found to be an antecedent to users' perceptions of a DA in a relatively structured, low-velocity domain (e-commerce shopping interaction). In addition, the aforementioned studies investigated and measured the perceptual and cognitive interaction of users and e-commerce decision aids, leaving the actual effects on decision performance and recommendation acceptance largely unexplored.

2.1.3 Uncertainty and the Decision Context

The concept of uncertainty and its impact on economic behavior has intrigued both economists and scholars for more than a century. And despite the number of theoretical and technological advancements over this corresponding period, quantifying the impact of uncertainty on economics still remains a formidable challenge (Pellissier & Fusari, 2007). Stewart (2000) highlights the fact that uncertainty in prediction simply means that, given current knowledge and information, there can be multiple possible future states. Uncertainty plays a major role in financial markets where human DMs are often charged with making some form of prediction regarding a future state, based on current knowledge and available information. These predictions can often take the form of economic data and trend forecasting, market levels, bond yields and even securities prices. A majority of the time, humans perform reasonably well regarding these predictions. Stewart (2000) highlights the potential that humans have for performing impressive mental feats.

So, if human cognitive competence can be robust in many endeavors, why then is decision making performance often suboptimal in practice? There are fundamentally two streams of research that can assist in answering this question. The first relates to the literature on biases and heuristics, where it is well known and documented that humans are fallible, are subject to making errors and don't always perform up to their full potential (Kahneman & Tversky, 1974;1982; Camerer et al., 2004). The second literature stream relates to the environment in which the individual is operating. According to Stewart (2000), humans are often forced to function in environments that do not foster optimal performance. The situation or context in which human judgment is exercised and predictions are made can play an integral role in the quality of the outcome. This is particularly true in the presence of uncertainty, where human error can routinely be found at the core of many accidents and disasters. This phenomenon is also a common fixture in the financial markets, where human error can adversely impact financial and economic outcomes. As a result, the problem of suboptimal performance is not completely a problem of biases, heuristics and limited human ability, rather it is the product of a combination of these innate characteristics and the state of the decision context (Rizzo et al., 1987; Stewart & Chase, 1999).

Kahneman and Tversky (1982) are credited with providing one of the seminal works in the field of uncertainty and decision making. They highlight the fact that uncertainty is pervasive, and can extend to represent uncertainty about signs or stimuli in the external environment as well as the potential consequences of a course of action. The influence of uncertainty in the external environment can be particularly salient in decision-making and can have adverse consequences for DMs (Stewart, 2000; Tversky, 1974). In this context, external factors in the decision environment can interfere with the effective application of a DM's skills and knowledge. Referred to as an application error, this situation arises when a DM has the requisite skills to make an appropriate decision, but is unable due to the influence of uncertainty in the external decision context (Kahneman & Tversky, 1982).

As previously established by Kahneman and Tversky (1982), uncertainty in the decision context can increase the number of errors that are made by a DM. The literature on human error can be integrated to provide additional insight into this phenomenon. Zhao and Olivera (2006) refer to errors of this type as a slip. Slips are classified as situations where an individual DM has the requisite knowledge on how to execute a prescribed task, but is unable to carry it out effectively due to either internal or external distractions (Rizzo et al., 1987; Stewart & Chase, 1999). Camerer et al. (2004) in their study on risky decision making, highlight the fact that instability in the decision environment can make it difficult for DMs to adhere to normative decision making measures. Camerer et al. (2004) also point out that a key dimension of risky choice is ambiguity, where uncertainty is based on a lack of information regarding probabilities. This is extremely common in the field of investment management, where the decision context is often punctuated by ambiguity and uncertainty.

Until recently, many financial and economic models have largely ignored the influence of uncertainty as an external factor. This is due in part to a long standing assumption in the economics and finance literature that human decision-makers operate primarily as rational utility maximizing individuals (Markowitz, 1954). However, these assumptions of rationality, and their underlying consequences for financial market efficiency, have been called into question over the last decade. As previously mentioned, psychologists and economists have documented numerous departures from market rationality in the form of specific behavioral biases and heuristics that are innate to the process of human decision-making under conditions of uncertainty and risk. The presence of these behavioral biases can lead to less predictable and undesirable economic outcomes for market participants. This perspective further supports the notion that individual preferences may not be entirely stable over time, but rather are likely to be influenced by a number of factors, both internal and external to the individual DM. One of the key external factors is related to specific environmental conditions in which the individual is situated when making a decision. As a result, when these environmental conditions shift, it should be expected that individuals' behavior deviate in response (Lo, 2005).

The adjustment of individual behavior, in response to shifting environmental conditions, is of particular interest in this research study. More specifically, the influence of an uncertain external environment on decision maker performance is an important moderating variable in this study, based on the foundation established by Lo (2005). Specifically, Lo (2005) points out many cited examples regarding violations of rationality that occur based on a changing environment. For example, loss aversion, overconfidence, overreaction and other behavioral biases are consistent with individuals using heuristics to adapt to an uncertain environment. Despite the research done in this domain, an empirical investigation of the decision context and decision-making outcomes in financial markets has been largely unexplored.

2.2 Decision Support Systems

Simon (1997) highlighted the fact that decision making is one of the most critical activities conducted within an organization. Since the late 1960's a variety of independent and standalone IT artifacts have been developed and deployed to support the complex activity of decision making. Referred to as Decision Support Systems (DSS), these systems are classically defined as computer-based tools used to support users in complex decision-making and problem solving tasks (Shim et al., 2002). DSS first started to populate the corporate landscape in the early 1970s (see Scott-Morton, 1978). Soon after their arrival, it was realized that DSS could be beneficial in solving poorly defined and non-structured decision problems (Holtzman,1989). The motivating principle underlying DSS is that resource-intensive, but standardizable, information and data processing tasks can be performed effectively by a computer-based system, thus increasing the availability of some of the human decision maker's mental processing capacity (Haubl & Trifts, 2000).

The literature on DSS also supports the notion that human DMs can be good at selecting the relevant variables for use in the decision process, but they are often ineffective at integrating and retaining large quantities of information (Haubl & Trifts, 2000). As a result, effective DSS should be designed to capitalize on the inherent strengths and compensate for the inherent weaknesses of their users (Hoch & Schkade, 1996).

More advanced variants of DSS, referred to as expert systems (ES), entered the domain of decision support in the mid 1980's. An ES attempts to capture and model the knowledge of human experts, thus making that knowledge accessible in problem solving tasks. Therefore, obtaining and coding the necessary knowledge from an expert is a prerequisite for constructing a tractable system. Although beneficial in many respects, ES do possess certain limitations in terms of supporting decisions (Yoon, 1994). The difficulty of programming and maintaining the knowledge and rule-base of the system, and the enormous time and effort required to extract the knowledge base from human experts, are but a few examples. The inability of an ES to use inductive learning and inference to adapt to dynamic situations is also a limitation (Hawley et al., 1990). In addition, an ES only knows what it has been programmed, and since it is not possible to program everything into the rule-base, the ES may be rendered useless in extraordinarily dynamic and semi-structured information domains, like those encountered in portfolio trading.

2.2.1 Foundations of IT in Decision Support

The decision-making process is theoretically concerned with generating and evaluating multiple alternatives and choosing the decision, or alternative, that satisfies expected utility. However, in the majority of decision-making problems, conflicting criteria for judging the possible alternatives often exist. The primary concern of the DM, therefore, is to maximize utility while operating within the constraints of the problem domain. For complex problems in which many tractable alternatives exist, the task of selecting the optimal alternative becomes difficult for the DM without some form of assistance (Malakooti, 1993). However, receiving decision-making assistance, in the form of advice or recommendations, presents a challenge in the effective use of DSS. Yaniv and Kleinberger (2000) provide insight into this issue, by highlighting the fact that a DM's perspective has a substantial influence on the weighting of their opinion, as well as the weighting of external advice. DMs normally have privileged access to the rationale that lead them to make their own decisions, but only limited access to the rationale that lead others to make their decisions. This fundamental asymmetry between the access to the logic for one's own decision, and the access to the logic used for another's decision strategy, sets the stage for a biased weighting of the DM's own decision versus the advice received. Therefore a DM may not consider the two respective opinions, or decisions, to be equivalent (Yaniv & Kleinberger, 2000).

Making decisions under uncertainty and risk is an increasingly difficult task when alternatives are numerous and when the complexity of the decision environment is high (Payne et al., 1993). Hawley et al. (1990) highlight the fact that a majority of decisions encountered by top-level financial managers lack complete structure. This factor complicates the decision-making process for DMs, challenging conventional methods of computer-aided decision support. Further complicating the use of DAs in investment management is the fact that many of the decisions in portfolio trading are largely unique in character, requiring an element of judgment and discernment to arrive at an appropriate decision in timely manner. The aforementioned factors present a void in the existing literature with respect to DAs in highly unstructured and semi-structured domains, like those encountered in portfolio trading.

2.2.2 DSS Utilization and Decision Performance

Critical to the foundation of this research study is the idea that a behaviorally similar IDSS will be relied on more frequently by decision-makers, and that the increased use of this system will positively influence decision performance. Hoch and Schkade (1996) support this notion in showing that decision makers who are provided with a DSS will utilize it to analyze problems in greater detail, and as a result, make better decisions. This is supported by the concept of bounded rationality (Simons, 1955). Based on this concept, it is commonly believed that decision makers would like to conduct a more comprehensive analysis when making decisions, but are unable to do so due to their innate cognitive limitations (Taylor, 1975).

In the domain of investment management, Felsen (1975) highlighted the fact that investment performance can be improved by at least partial automation of the investment decision-making process. However, despite the literature supporting the connection between DSS utilization and decision performance, the empirical evidence supporting this relationship are by no means conclusive. For example, some researchers have provided empirical evidence that the use of a DSS does not necessarily improve decision-making performance (Benbasat & Nault, 1990). Furthermore, Todd and Benbasat (1992) provide evidence that decision performance may even be reduced as a result of using DSS. The lack of empirical support for the relationship between DSS use in general, and IDSS use in particular, and decision performance represents an important gap in the literature to be explored in this study.

2.2.2.1 Perceived Usefulness

Perceived usefulness (PU) is seen as a fundamental, and often necessary, determinant of user acceptance of a DSS. This notion of performance expectancy, which owes its origin to the early Technology Acceptance Model (TAM), specifies "the degree to which a person believes that using a particular system would enhance his or her job performance" (Davis, 1989). Davis et al. (1989) theorized PU to be an important determinant of intention behavior, compared to other cognitive factors. Their theory was supported on the basis that in an organizational context, emphasis is often placed on productivity as a motivating factor. As a result, an individual's assessment of the performance benefits associated with technology use (i.e., PU) will be the single most important determinant of usage intentions and behavior (Davis, 1989). Of the many determinants of IT and decision aid (DA) adoption that have been explored in the literature, this particular construct has received a great deal of theoretical support (Cooke et al., 2002; Wang & Benbasat, 2005).

Empirical studies representing a range of IT systems and platforms have found PU to be a strong determinant of intention and usage patterns (Davis & Venkatesh, 2004). Specifically, PU has been shown to be an influential antecedent to the adoption and use of on-line DAs, as well as other types of DSS in which the decision domain lacks complete structure and is subjective (Dhaliwal & Benbasat, 1996; Arnold et al., 2004).

In terms of specific technologies, PU has been used extensively in the literature on e-commerce interactions. In this context, researchers have embraced a perspective in which the extrinsic cognitive beliefs of the users are critical in determining the adoption of IT artifacts like websites and recommendation agents (Al-Natour et al., 2011). PU has also been researched with respect to Knowledge-Based Systems (KBS), where Gregor and Benbasat (1999) found that the use of an explanation facility can lead to favorable perceptions, including the perceived usefulness of the artifact. Jones and Mitchell (1995) conducted an experiment to test PU in an intelligent associate system in a real-time decision context. In this study the DSS was perceived to be useful by its operators, and was able to provide performance benefits for certain portions of the experimental control task. PU and similarity has also been evaluated in the literature. Al-Natour and Benbasat (2005) and Al-Natour et al. (2008) found that perceived process similarity had a significant positive effect on perceived usefulness and trust in an e-commerce DA.

PU serves as an integral theoretical element in this thesis, given the theory that the DM's cognitive beliefs of the IDSS artifact will be critical to its utilization in a decision-making context. And while numerous examples exist in the literature regarding the use of PU in decision aids and support systems, what appears to be less studied is the influence that cognitive beliefs like PU have in situations involving decision-making under uncertainty, in semi-structured decision domains.

2.3 Trust in Information Systems

No matter how robust or "intelligent" a DSS may be, the system’s advice and guidance may be rejected by a DM who does not trust it, disrupting the potential benefits of the system in terms of decision performance or efficiency. Furthermore, if asked to use a DSS in which they do not trust, DMs may use any means available, even at the expense of efficiency and effort, to direct the output of the system toward their own decision (Muir, 1987). As a result, in order to realize the performance benefits of IDSS, system designers and researchers must first design DAs that decision makers will trust enough to use. As previously mentioned, there is ample theoretical support for the treatment of technological artifacts as recipients of social and relational aspects of trust (Wang & Benbasat, 2005). Furthermore, numerous studies have extended the attribute of trustworthiness to technical systems, as well as intelligent computer agents (Komiak & Benbasat, 2004; Muir & Moray, 1996).

Based on the multi-dimensional nature of trust, a universally accepted definition of what constitutes "trust" in the HCI literature has remained elusive, with many of the definitions originating from the domain of HHI. Mayer, Davis, and Schoorman (1995) define trust as, "*the willingness to be vulnerable to the actions of another party based on the expectation that the other will perform a particular action important to the trustor*". Madsen and Gregor (2000) defined human-computer trust as "*the extent to which a user is confident in, and willing to act on the basis of, the recommendations, actions, and decisions of an artificially intelligent decision aid*." This definition has been chosen from the literature to operationalize the working definition of trust in this study. This is based on the fact that it encompasses both the user's confidence in a system, as well as their willingness to act on the advice and guidance of the system.

Chopra and Wallace (2003) highlighted the fact that trust is important in situations where there is a state of dependence between two parties, and when this dependence entails a certain element of risk. This is a particularly salient point in the context of portfolio trading, where the decision domain is punctuated by uncertainty and financial risk, and a DM (trader) may be asked to rely on a DSS for trading support. An important consideration here is the fact that the literature shows that people generally decide to trust others when facing situations that have a high degree of uncertainty (Dasgupta, 1988; Kollock, 1994; Sniezek & Van Swol, 2001). As a result, uncertainty is an antecedent to the decision to trust another (Mayer et al., 1995).

With respect to the way that trust develops, research provides support for the notion that users may initially be predisposed towards distrust in decision support artifacts. In supervisory control environments, Sheridan and Hennessy (1984) found that system operators may also be initially biased toward distrust. Following this notion, Muir (1987) proposed that trust evolves over time, and is dependent on a human's ability to estimate the predictability of a machine's behavior. He further states that a human's ability to estimate a machine's predictability will depend on his/her own limitations as a decision-maker, and on elements and characteristics of the machine and the environment in which it is operating. The study also emphasized the fact that in order for user trust in an IT artifact to develop, the behavior of the machine must be observable.

A large amount of research on trust between humans and machines has emerged in the field of e-commerce (Gefen et al., 2003). In this context, Wang and Benbasat (2005) built upon the definitions of trust from Xiao and Benbasat (2002) and McKnight et al. (2002), to define trust in an RA based on users' beliefs in an agent's competence, benevolence, and integrity. Trust is an important multi-dimensional construct that can help consumers overcome perceptions of uncertainty and risk and engage in "trust-related behaviors" with a web-based agent (McKnight et al., 2002).

2.3.1 The Influence of Similarity on Trust

The DM's perception of trust in the prototype system is an integral part of the theoretical foundation of this study. Zuckers (1986) highlighted the effects of personality similarity with respect to influencing feelings of trust. In a study of online shopping assistants, Al-Natour et al. (2005) showed that perceived behavioral similarity had significant effects on trust and perceived usefulness of the artifact by its user. Komiak and Benbasat (2006) conducted a study extending the research on trust and IT adoption by investigating how RA personalization and familiarity affected the adoption of the RA in an e-commerce transaction. They investigated the adoption of an RA through enhancing cognitive and emotional trust in the artifact by personalizing the RA to the customer. The perception of similarity was fostered through the RA asking questions to better identify the customer’s personal needs for a particular product. The study revealed that customer trust is particularly important in e-commerce transactions. Specifically, the Komiak and Benbasat (2006) study revealed how perceived personalization could be used to increase a customers’ intention to adopt an RA by positively influencing cognitive trust and emotional trust.

Cognitive Trust

The construct of cognitive trust in competence will be an important element in the study because trust plays an important theoretical role in the acceptance and utilization of the IDSS for purposes of decision-making under uncertainty. The concept of cognitive trust is consistent with the concept of trusting beliefs (McKnight et al., 2002), and can be defined as a trustor’s rational expectations that a trustee will have the necessary attributes to be relied upon (Komiak & Benbasat, 2004). Specifically, cognitive trust is developed when the trustor believes that a valid foundation to trust is fundamentally present (Lewis & Weigert, 1985). Komiak and Benbasat (2006) showed that utilizing emotional trust and cognitive trust as part of IT adoption models in e-commerce contexts is beneficial with respect to influencing adoption behavior.

Utilizing an adaptation of the definitions of trust from Xiao and Benbasat (2002) and McKnight et al. (2002), this thesis defines trust in a technological artifact as a DM’s beliefs in the system's competence to make accurate recommendations. Referred to as competence-belief, this concept has been well accepted in many recent studies. Competence-belief means that an individual believes that the trustee (technological artifact) has the ability, skills, and expertise to perform effectively in a specific domain (McKnight et al., 2002).

2.3.2 Trust and Task Delegation

Another important aspect of the study is the application of existing theory related to task delegation. Relying on the guidance of an IDSS artifact requires a willingness to delegate certain cognitive elements of the decision-making process. Since the delegation of a task involves ceding a certain degree of responsibility, but retaining accountability for the ultimate decision, trust is critical to any delegation-oriented interaction (Milewski, 1997). Mulken et al. (1999) highlighted the fact that delegation depends on the trustworthiness of an agent. In the context of HHI, trust determines how a person decides whether to delegate, what to delegate, and to whom to delegate (Axley, 1992). Trust in technological artifacts exhibits many of the same dynamics with respect to delegation. The concept of user trust has been an omnipresent issue in the design of decision support and control systems (Sheridan, 1980), and will be a necessary antecedent in order for DMs to accept the advice from the IDSS prototype in this study.

2.4 Artificial Intelligence for Decision Support

Advances in the fields of computer science and artificial intelligence (AI) have provided many theoretical and practical improvements to the design of modern DSS. The potential contributions of these intelligent elements to DSS have been described as enormous (Whinston, 1997). The origins of this technology go back many years, to include research in the fields of expert systems, robotics and supervisory control systems (e.g. Negroponte, 1970; Roth, Bennett, & Woods, 1987; Turban & Watkins, 1986; Woods, Johannesen & Potter, 1991). While a universally agreed upon definition of AI remains elusive in the literature, most experts agree that AI is associated with two basic premises. The first premise relates to studying the thought process of humans. The second premise relates to representing these human thought processes via machines (Turban, Aronson, & Liang 2004). The notion of "intelligent behavior" is a key theoretical element from the AI field, and was used as a supporting concept in this study. Specifically, the ability of a machine to learn and develop an adaptive KB based on experience will be key a component of the research artifact.

Advances in AI techniques and methods have resulted in many improvements in the DSS field (Dahr & Stein, 1997; Turban et al., 2004; Jackson, 1999). As an example, advancements in knowledge base design and structure, fuzzy logic, multi-agent systems, natural language processing, genetic algorithms, and neural networks are but a few such examples found in the literature (Sousa et al., 2007). The utilization of AI technologies to create IDSS is an effort to develop systems that have the capability to imitate certain human characteristics, such as intuition, approximate reasoning, and common-sense (Jackson, 1999). This is an important element in the design of the prototype IDSS for this study given the fact that the KB is intended to imitate the investment selection process of a human decision-maker.

2.4.1 Artificial Neural Networks

A promising development in the field of AI research is what is referred to as the artificial neural system, also commonly referred to as artificial neural networks (ANN). An ANN is a computer algorithm that simulates the neural process by which human learning takes place. ANN technology was developed in an attempt to replicate the knowledge acquisition and organization processes of the human brain. ANN can provide significant support in terms of organizing, classifying, and summarizing data. ANN can be effectively used to discern patterns in data with a high degree of prediction accuracy, using a limited number of a priori assumptions (Wong & Selvi, 1998). Haykin (1999) highlights the fact that a neural network structure has a natural propensity for storing experiential knowledge and making it available for use. Knowledge is distributed over the ANN with a structure of processing units called neural nodes, which are connected by weighted connections, or weights. ANN knowledge is acquired through a learning process referred to as training.

Unlike the rule based ES mentioned previously, the ANN approach to knowledge-base design is not programmed with any preexisting rules or structure, rather it actually learns through experience as well as trial and error (Hawley et al., 1990). This adaptive capability enables ANN to be applied to problem domains that are lacking in structure, require some form of pattern recognition and may involve incomplete or noisy data (Desai & Bharati, 1998). As a result, an increasing amount of application and development efforts have concentrated on using ANN in the finance and capital markets sector (Wong & Selvi, 1998). Financial services organizations are second only to the Department of Defense with respect to sponsoring research efforts in ANN (Trippi & Lee, 1996). While the literature is replete with examples of using ANN for pattern recognition and for solving problems of a non-linear nature in a business context (Wong & Selvi, 1998; Ainscough, et al., 1997; Trippi & Turban, 1996; Haykin, 1998), very little research has been conducted to-date with respect to using ANN technology to replicate the decision process of a human DM for purposes of creating an IDSS.

An important element of ANN design is the learning mechanism. ANNs can be classified into one of two categories: supervised and unsupervised. Supervised learning is based on an external teacher or DM who provides feedback in terms of evaluating a given set of alternatives. Unsupervised learning does not require the input from an external DM, and is typically performed without direct evaluations by DMs (Malakooti & Zhou, 1994). The most widely used ANN based on supervised learning are multiple layer perceptrons (MLP). Nodes in the MLP are structured in three hierarchical layers: input layer, hidden layer(s), and output layer. Information travels from the input layer to hidden layer(s), and then to the output layer. Hecht-Nielson (1989) provided support for the fact that any continuous function could be implemented with a three-layer perceptron. An ANN with no recursive loops is known as a feed-forward neural network, and MLPs are classified as feed-forward ANNs (Chen & Lin, 2010). The supervised learning MLP will be used in this study, consistent with Quaha and Srinivasan (1999), who point out that algorithms designed for supervised learning are ideal. Among the available training algorithms, the Levenberg-Marquardt (LM) back-propagation algorithm designed by Rumelhart, Hinton and Williams (1986) was selected for this study given its prevalence in Finance.

2.4.1.2 Capturing Decision-Maker Preference Structures

The ability to effectively capture and mimic the decision-making process of a human DM using ANN is a critical design feature of the proposed prototype system. This element is important because most of the existing research and state-of-the-art in IDSS rely upon single predefined KB, to support DMs across a specified domain. However, since the introduction of utility functions by the economists von Neumann and Morgenstern (1944), it has been accepted that "rational" decision makers, confronting the same decision, may make two different decisions based on their subjective probabilities (Pomerol, 1995). Therefore, capturing a DM's individual utility function and preference structure represents an important, yet often overlooked element, to the decision support process. Chen and Lin (2003) successfully accomplished this by utilizing an ANN approach for solving multiple criteria decision-making (MCDM) problems. In their study, a modification of ANN called a decision neural network (DNN) was utilized. The DNN was used to capture and represent the DM’s preference structure using the multi-attribute utility function (MAUF) method. The findings of the Chen and Lin (2003) study illustrate the advantages of ANN as a promising tool in terms of approximating the MAUF and representing the preference of a DM.

Malakooti and Zhou (1994) presented an Adaptive Feed-forward Artificial Neural Network (AF-ANN) approach to solve discrete MCDM problems. In their study, Malakooti and Zhou (1994) utilized an AF-ANN to successfully capture and represent the preferences of a DM, in order to select the preferred alternative. An essential benefit of the AF-ANN is that it can adjust and improve its representation of the decision space, as more information from the DM is captured. The aforementioned studies provide a foundation for the tractability of using AF-ANN to capture a DM's MAUF. While this work is by no means exhaustive, the theoretical and experimental evidence support the use of a feed-forward MLP AF-ANN as a KB for the prototype IDSS.

2.4.2 Intelligent Decision Support Systems

Many of the limitations found in DSS and ES could be overcome with advances in the field of artificial intelligence (AI). Complementing the suite of existing decision-making systems are what are referred to as intelligent decision support systems (IDSS). Just as in DSS and ES, there is no universally agreed upon definition for IDSS. Essentially, IDSS are constructed by combining a DSS with elements of AI, like evolutionary and adaptive algorithms. The rationale behind this basic design is to combine the knowledge reasoning capabilities of AI and the basic capabilities of DSS. Turban and Watkins (1986) further defined IDSS as decision support systems with inbuilt ES technology.

IDSS are intended to provide a system that is capable of supporting a decision maker in all phases of the decision-making process through a set of recommendations that have the capability of replicating domain expertise (Wang, 1997). The IDSS literature provides numerous examples illustrating that the use of IDSS can improve the decision-making process and outcomes (Gupta et al., 2006; Phillips-Wren & Jain, 2005). The literature also highlights the fact that IDSS can support cognitive tasks by playing an active role in aiding task and data processing performance, supporting the premise that the use of IDSS can result in better decisions in terms of the outcome of the decision itself (Linger & Burstein, 1997). Despite the recent advances in computational capability of IDSS, there is limited knowledge as to how to effectively deploy the power available through these new capabilities to improve human decision-making performance (Roth et al., 1987). In addition, given the nascent state of the research on IDSS and AI based knowledge systems, few studies have been conducted to investigate the impact that IDSS knowledge-base design may have on decision-making performance (Moreau, 2006). Furthermore, few studies in the existing literature have focused on the decision itself as the unit of analysis.

3. Model and Hypotheses

In this section, the theoretical model, variables and related hypotheses are introduced. The research model for this study consists of a high-level conceptualization of the relationships between behavioral similarity, recommendation acceptance and decision performance of a human-machine decision-making system.

While behavioral similarity theory provides the theoretical foundation, this study is not focused on explicating the psychologically oriented belief of perceived similarity as an antecedent to IT adoption, as this is well established in the literature. Rather this study is focused on investigating the use of perceived similarity to construct a behaviorally similar IDSS to influence recommendation acceptance and human decision-making in a semi-structured and uncertain information domain. In so doing, this research study contributes towards bridging the theoretical gap between psychologically oriented cognitive beliefs and IDSS design characteristics for purposes of improving decision performance and efficiency in an uncertain decision context.

3.1 Research Model

Drawing on the HCI, trust, DSS, and AI literature, this study developed and empirically tested a prototype IDSS that utilized an AF-ANN knowledge base to elicit the perception of behavioral similarity between a decision aid and a DM. The perception of decision process similarity should elicit the psychological constructs of perceived usefulness and trustworthiness in the IDSS, consistent with Al-Natour et al. (2008). In the study, perceived usefulness and trustworthiness were not measured; rather they were used to provide theoretical support for the design of the IDSS. Building on this theoretical foundation, this study made the assumption that an increase in the perceived usefulness and trustworthiness of the IDSS artifact, resulting from perceived behavioral similarity, should positively influence the acceptance rate of IDSS recommendations, as well as the decision performance and efficiency of the joint human-IDSS cognitive system. In terms of the research model (Figure 1), a derivation of the behavioral similarity model outlined in Al-Natour, et al. (2005) is utilized. In the proposed research model the influence of the decision context is an integral element, and is expected to moderate the influence of perceived behavioral similarity.

Figure 1: Conceptual Framework

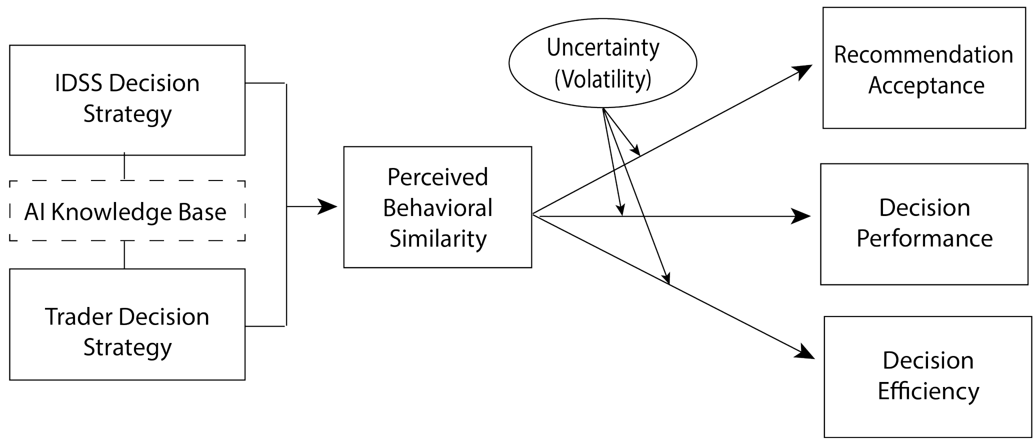

In the proposed research model, perceived behavioral similarity will be fostered through decision outcome similarity. In this context, decision outcome similarity refers to the degree to which the trading recommendation of the IDSS and human trader match.

3.2 Measures and Operationalization

3.2.1 Independent Variables

Consistent with the exploration of the aforementioned research questions, there are two primary dimensions of interest explored in this study: *Similarity State* and *Volatility State*. These two dimensions were operationalized as respective treatment conditions (independent variables) for purposes of conducting the experiment. *Similarity State* was manipulated in the experiment based on the design of the KB in the IDSS artifact, and *Volatility State* was tested as a moderator.

Similarity State

The first treatment, titled *Similarity State*, is based on similarity-attraction theory (Byrne & Stefaniak, 1967) and is operationalized in this experiment as the level of perceived behavioral similarity between a human trader and a specialized IDSS artifact. The similarity state treatment is fostered by the human trader's perception of similarity with respect to the machine's underlying decision process. To facilitate the use of this variable in the experiment the theoretical foundation and approach established by Al-Natour et al. (2008) was utilized. As a result, perceived decision process similarity was manipulated based on outcome similarity in equity covered-call option trading decisions.

The variable *Similarity State* represents (2) treatment levels: high-similarity and low-similarity. The two respective treatment conditions are created based on the type of knowledge-base used by the IDSS. The high-similarity KB is constructed using an Adaptive Feed-forward Artificial Neural Network (AF-ANN) based on the IDSS user to create a behaviorally similar artifact (KB_H). The second treatment condition, low-similarity, is established through the utilization of an IDSS with a KB trained using the AF-ANN knowledge acquisition approach for a different user (KB_L). The exact mechanics of how the respective treatment levels were implemented is detailed in Chapter 4.

Uncertainty in the Decision Context

The second experimental condition, titled *Volatility State*, is a moderator used to operationalize the level of uncertainty in the decision context for subjects in the experiment. Uncertainty is a pervasive fixture in the equity trading market environment. It is well documented in the literature that uncertainty in the decision context can subject humans to cognitive biases, slips, and application errors (Camerer et al., 2004; Rizzo et al., 1987; Stewart & Chase, 1999; Kahneman & Tversky, 1982; Zhao & Olivera, 2006).

The level of market volatility, as measured by the Chicago Board of Options Exchange Volatility Index (CBOE VIX), is intended to operationalize the level of uncertainty in the decision context, and is a moderating variable in the study. Market volatility was selected as an operational measure of uncertainty consistent with Bloom (2009), where financial market implied share-return volatility was used as a canonical measure for uncertainty.

The VIX measures the expected volatility of the Standard and Poor's 500 stock index over the next thirty days. The VIX was selected as an observable measure of aggregate uncertainty due to its prevalence in financial markets consistent with Basu and Bundick (2011) and Bekaert, Hoerova and Lo Duca (2010). The VIX is a forward-looking indicator of the expected volatility of the Standard and Poor's 500 stock index, and is a broadly utilized metric used by options and equity traders in practice to evaluate the level of expected volatility in the equity markets. Bekaert, Hoerova and Lo Duca (2010) highlight the fact the VIX can be bifurcated into two components, a proxy for risk aversion as well as expected stock market volatility ("uncertainty"). The VIX is well suited as a surrogate for uncertainty in this study given the fact that it is a forward-looking metric, measuring volatility that investors expect to see, as opposed to measuring volatility that has been previously realized (Whaley, 2008).

The moderating variable *Volatility State* represents (2) levels: high-volatility and low-volatility. The classification schema for the two volatility states is based on a median split of the CBOE Volatility Index (VIX) for the respective trading cycle. Trades are classified as either high-volatility (above median VIX) or low-volatility (below median VIX).

3.2.2 Dependent Variables

DSS research has focused on a myriad of metrics and variables to determine the influence of a DSS on decision-making outcomes. For example, DeLone and McLean (1992) evaluated individual impact through effectiveness, efficiency, estimated value of the information and the system, and changes of behavior based on system use. Keen and Scott-Morton (1978) also identified the variables of effectiveness and efficiency as useful in assessing the impact of DSS use.

The dependent variables selected for the evaluation of the IDSS prototype are derived from the literature on information systems success, and have been adapted to this particular study. The primary variable of interest is the frequency by which the traders accept or override the trading advice of the IDSS. The two primary variables used to evaluate decision-making amplification are based on the categories of effectiveness and efficiency outlined by Keen and Scott-Morton (1978). Effectiveness refers to the quality or performance of the decision, and efficiency is typically measured as the speed or reliability of the decision (Sharda et al., 1988).

Based on these evaluation criteria, it is expected that the use of the prototype IDSS will have a direct effect on three general aspects of trader decision-making in a semi-structured and uncertain trading environment: (1) trader acceptance of the machine's recommendation, (2) decision performance (effectiveness) and (3) decision efficiency.

Table 1: Summary of Measures

Type	Dependent Variable	Definitions	Measure
Objective	Recommendation Acceptance	Primary measure of agreement between IDSS and DM	The frequency by which the trader agrees with the recommendation of the IDSS
Objective	Decision Performance	Primary measure of decision performance for trader, artifact, and combined system	The mean option premium generated per trade, as well as portfolio standard deviation
Objective	Decision Efficiency	Secondary measure of decision performance based on the efficiency of trading decisions	The cumulative time to evaluate and execute a trading decision (measured in minutes)

Recommendation Acceptance

Trader acceptance of the machine's recommendation is an important element with respect to evaluating the influence of the IDSS. Langlotz & Shortliffe (1983) highlight the fact that decision-making performance and user acceptance of system recommendations can be independent issues, and therefore should be evaluated separately. In this study, recommendation acceptance relates to user acceptance of the IDSS trading recommendations. Lack of user acceptance (where acceptance means the trader adopts the machine's advice) is seen as a major problem in the design and deployment of modern DSS. Relying on the advice of another, in many respects, involves a willingness to delegate certain elements of the decision-making process. Since delegation in a decision-task involves ceding partial responsibility, but retaining full accountability for the ultimate and final decision, a DM must trust the entity providing the advice (Milewski, 1997). It is therefore proposed that using behavioral similarity as a design element should increase the amount of trust a DM has in the IDSS. Based on this, traders should be more willing to accept the recommendations of the prototype IDSS.

The unwillingness to take another's advice is a common phenomenon experienced in human-human interaction, due to the fact that accepting advice from another party often exposes a DM to a potential conflict between their initial decision and advice from another party. As a result, DMs often encounter cognitive friction in reconciling these two diverse views in order to make a decision (Yaniv & Kleinberger, 2000). Behavioral similarity should mitigate the effects of this phenomenon by reducing the level of uncertainty and opacity regarding the way the IDSS processes investment data and arrives at trading recommendations.

In terms of measuring recommendation acceptance, the frequency by which the DM concurs with the advice of the system, as well as the frequency by which the trader maintains their original decision will be recorded. Based on the design of the experiment, the perception of behavioral similarity in the IDSS should positively influence the acceptance of machine recommendations. As a result of integrating behavioral similarity into the IDSS, traders will receive support for a decision from an intelligent system that they perceive as both trustworthy and useful (Zuckers, 1986; Nass et al., 1995; Al-Natour, et al., 2005; Komiak & Benbasat, 2006).

Decision Performance

Decision performance is based on the decision outcome of the selected equity derivative trades for the combined human-machine system. Several researchers have claimed that outcome is one dimension of DSS performance measurement (Sainfort et al., 1990; Kanungo & Sharma, 2001). The outcomes for this experiment will be measured in terms of (1) individual trade performance, which is the gross option premium generated for each option trade; and (2) the realized standard deviation of the gross option premium generated for each option trade. A behaviorally similar IDSS should positively influence decision performance by reducing the influence of application errors and slips in the trading process during periods of elevated volatility.

Portfolio standard deviation is an important metric in the evaluation of investment and portfolio performance (Markowitz, 1952; Sharpe, 1987). With respect to Modern Portfolio Theory (MPT), lower standard deviation portfolios are preferable to investors in cases where the expected returns are equivalent. Portfolio standard deviation is often the result of sub-optimal trade selection and inconsistency in security selection methodology, as well as abrupt responses to environmental factors. Reducing portfolio standard deviation is a much sought after goal in securities trading. Based on the aforementioned benefits of the prototype IDSS, trades executed with the behaviorally similar IDSS should have lower standard deviation than those executed with the low-similarity system. In this experiment, portfolio standard deviation is evaluated based on the gross option premium generated for each trade.

Decision Efficiency

Decision-making efficiency has been previously evaluated as a characteristic of MIS success (Raymond, 1985). In defining system success, Seddon (1997) defined efficiency as more work done in the same time, or less time for more work of equivalent quality. In this study, decision-making efficiency is defined as the total time required to make a trading decision, measured in minutes. The time required to evaluate trading factors and make a trading decision is important given the fact that decision makers can often be described as “cognitive misers” who strive to reduce the amount of cognitive effort associated with decision-making (Shugan 1980). This phenomenon is particularly salient in instances when alternatives are numerous and/or difficult to compare, and the complexity of the decision environment is high (Payne et al. 1993). Furthermore, given the velocity of news, information, and data flow in today's markets, traders must be able to make decisions quickly and with minimal reservation if they want to capture opportunities in the market. Reducing the time required to make a trading decision will enable traders to focus on additional profitable opportunities and threats to their portfolio.

3.3 Hypotheses

Critical to the design of this experiment is the notion that outcome similarity in trading recommendations between a human trader and the prototype IDSS will result in the perception of behavioral similarity. As a result of the perception of behavioral similarity, traders should view the IDSS as more trustworthy and useful, positively influencing recommendation acceptance, decision performance and efficiency. This influence will be most salient in times of uncertainty or perturbation in the market environment, when the human DM can become distracted. In addition, since humans seek to confirm their own decisions (Al-Natour et al., 2008), an IDSS recommendation that appears similar to the trader's will be viewed as more credible by the DM. Based on these factors an increase in the amount of decision-making delegated to the IDSS should result, as measured by the level of recommendation acceptance by the human trader. This interaction should become particularly evident in situations where the decision context (equity markets) experience heightened levels of uncertainty (volatility). This interaction is grounded in the fact that humans generally decide to trust, and rely on others when encountering situations involving uncertainty (Dasgupta, 1988; Kollock, 1994; Sniezek & Van Swol, 2001). As a result, traders will be more likely to rely on their behaviorally similar "teammate" in periods of higher uncertainty.

When uncertainty and distraction are more pervasive, i.e. when market volatility increases, DMs are more prone to cognitive biases, slips, and application errors (Rizzo et al., 1987; Stewart & Chase, 1999; Kahneman & Tversky, 1982; Zhao & Olivera, 2006). However, the IDSS trained on the DM will remain rational under elevated levels of uncertainty, and will not be influenced and distracted by the conditions in the environment. Based on this premise the IDSS should provide a trading recommendation that would be consistent with that of a rational trader, irrespective of the decision context.

3.3.1 Recommendation Acceptance (H1)

Based on the premise that behavioral similarity matters in the design of IDSS, the following hypotheses are presented regarding perceived behavioral similarity and the acceptance of recommendations:

> **H1a:** The utilization of a behaviorally similar IDSS will increase the acceptance rate of machine recommendations
>
> **H1b:** Uncertainty (volatility) in the decision context should strengthen the influence of a behaviorally similar IDSS on the acceptance rate of machine recommendations

3.3.2 Decision Performance (H2)

Based on the premise that behavioral similarity matters in the design of IDSS, the following hypotheses are presented regarding perceived human-machine similarity and decision performance:

> **H2a:** The utilization of a behaviorally similar IDSS will positively influence trading performance
>
> **H2b:** Uncertainty (volatility) in the decision context should strengthen the influence of a behaviorally similar IDSS on trading performance

Based on the premise that behavioral similarity matters in the design of IDSS, the following hypotheses are presented regarding perceived human-machine similarity and portfolio volatility:

> **H2c:** The utilization of a behaviorally similar IDSS will decrease the standard deviation of trading performance
>
> **H2d:** Uncertainty (volatility) in the decision context should strengthen the influence of a behaviorally similar IDSS on the standard deviation of trading performance

3.3.3 Decision Efficiency (H3)

Based on the premise that behavioral similarity matters in the design of IDSS, the following hypotheses are presented regarding perceived human-machine similarity and decision-making efficiency:

> **H3a:** The utilization of a behaviorally similar IDSS will decrease the time it takes a trader to make a trading decision
>
> **H3b:** Uncertainty (volatility) in the decision context should strengthen the influence of a behaviorally similar IDSS on the time it takes a trader to make a trading decision

4 RESEARCH METHODOLOGY

The research method serves as the foundation for the advancement of knowledge in any given domain. For this reason, careful consideration was given in this dissertation not only to the theoretical constructs, but also to a rigorous and methodological research approach. To explore this particular research study a controlled experiment was conducted to evaluate the above hypotheses regarding the effects of perceived behavioral similarity as a decision aid design element on the three dependent variables of interest. More specifically, a quantitative approach was used to investigate the effects of utilizing a prescribed IDSS (treatment) on recommendation acceptance of machine solution, decision performance, and decision efficiency under two volatility states.

4.1 Experimental Design

An experimental design was selected in order to provide answers to the aforementioned research questions, and explore the hypotheses of this thesis. This research approach was selected due to the fact that developers of DSS and AI based systems often lack the empirical data needed to support the proposed merits of their systems. This is particularly true in the financial services industry where advances in AI technology can provide substantial benefits. As a result, it was decided to utilize an experiment to capture the necessary quantitative performance metrics for the prototype system. It was important to test for statistical differences between the decisions made with the use of a behaviorally similar IDSS, compared to decisions made with a low-similarity artifact. Data for each of the dependent variables (DVs) was captured in a database that was specifically designed for this experiment. The individual trades evaluated and executed with the prescribed IDSS under the two respective treatments is the unit of analysis for the experiment.

To test the underlying hypotheses, two primary statistical techniques were utilized. For the first hypothesis, *Recommendation Acceptance*, the Cochran's Q test was utilized. This particular test was selected due to the fact that recommendation acceptance with the IDSS is a binary response measured under two treatments (high-similarity and low-similarity) with two conditions (high-volatility and low-volatility). Cochran's Q tests that the marginal probability of a positive response, in this case agreement with the IDSS, is unchanged across the repeated-measures treatments and conditions.

With respect to hypotheses two and three, a two-way within-subjects counterbalanced repeated-measures ANOVA was utilized. This particular type of ANOVA was utilized in an effort to discern statistically significant differences in the means between the trades of the two treatment groups under the two volatility states. This approach allows for the interaction effects of volatility to be analyzed, while controlling for potential confounds based on any differences in the individual traders. The repeated measures ANOVA was conducted in the statistical software package SPSS.

A counterbalanced approach to the experimental design was implemented to mitigate the potential confounds of market factors, carryover effects or learning bias in the results. Counterbalancing can be useful in distributing any outside effects over the two respective treatment conditions. In addition, pilot tests with the IDSS were employed to test the software, but also to orient the users to the IDSS. Since the ANN KB was the same for all participants during the pilot tests, this training period should help minimize the impact of learning bias in the actual experiment. A total of 3 pilot tests on approximately 35 trades were conducted prior to the beginning of the experiment.

In terms of calculating an a priori sample size for the analysis, a medium effect size (Cohen .25), Alpha=.05, and Power = .95, was used. Based on these inputs a target sample size of 54 trades per trader was calculated. Although the unit of analysis is the individual trade and not the trader, it was felt this was the most tractable approach to generate a statistically powerful sample size. As a result of this methodology, 56 trades per trader were actually evaluated, for a total sample size of 112 trades across the respective treatment conditions. Table 2 highlights the experimental design and sample size for each of the respective treatment conditions.

Table 2: Experimental Design

		Similarity State	
		Low-Similarity	High-Similarity
Uncertainty in the Decision Context (***Volatility State***)*	Low-Volatility	112	112
	High-Volatility	112	112

*Moderator based on median VIX split.

4.1.1 Research Setting

It is important to understand the phenomenon in question within a real-life context to fully elucidate an understanding of the perception of AI technology and its impact on decision-making process and performance in varying degrees of uncertainty (volatility). In this case, the context was an equity call option trading operation at a major investment management firm. The experiment was conducted during normal market hours. The participants were evaluated in their natural work environment, and the IDSS prototype was deployed on their individual workstations.

4.1.2 Study Participants

Because of the limited availability of subjects trained in equity call-option trading strategies in a real working environment, in addition to the time and effort involved in training an ANN to accurately depict a user's decision-making process, the experiment was limited to four subjects. The selection of the investment professionals for the study was based on the subjects' familiarity in option trading, and their level of technological proficiency. These factors were deemed important for consistency and sampling.

Participants were financial professionals from a large financial services firm. The experiment was conducted in the subjects' natural work environment during business hours. No incentives were provided to the subjects for participating in the experiment. The average age of the subjects was 32, and subjects were all male with an average tenure in their current role of 4.2 years. The subjects participated in the experiment after completing one training and orientation session, and three pilot trading sessions with the IDSS artifact. Each participant was proficient in the use of normative trading software and applications, and had exposure to the same workstation and market information during the duration of the experiment.

4.1.3 Manipulations

Similarity State was the key manipulation in the experiment. The manipulation was based on the perception of decision process similarity and was fostered in the experiment by the outcome similarity of trading decisions between a human DM and the prototype IDSS. While many of the previously mentioned studies have explored behavioral similarity based on perceptual measures, the influence of this construct on the actual decision outcome itself remains unexplored.

Volatility State, a moderator, was based on using a median split of the CBOE Volatility Index (VIX) to create a "high" and "low" volatility classification scheme for option trades. Trades with comparable technical characteristics and option tenor were classified into one of the two groups based on the market state in which they were executed. The classification approach was used to create two trade groups: high-volatility (above median VIX) and low-volatility (below median VIX). This classification schema helped control for sample size and individual differences of the underlying option positions presented to the traders. This approach also ensured homogeneity of option maturity and technical characteristics for the respective trades (Unit of Analysis).

4.2 Procedure

4.2.1 Strategic Trading Artificial Neural Network (STANN)

In terms of manipulating the independent variables in the experiment, an IDSS prototype system was developed. The IDSS was designed to support traders in making equity covered-call option trading decisions. The prototype IDSS artifact is referred to as the Strategic Trading Artificial Neural Network (STANN). The interface of STANN is based on a simple 2-dimensional avatar, consistent with Qiu and Benbasat (2004). The knowledge base for the prototype IDSS is derived from an Adaptive Feed-forward Artificial Neural Network (AF-ANN). The AF-ANN was used to capture and represent each trader's preferences and utility function based on a historical dataset of trading decisions derived from each DM.

4.2.1.1 Knowledge Base

In this study, the research questions relate to evaluating the effects of integrating two different types of KB design (independent variable) into an IDSS for purposes of evaluating the acceptance rate of recommendations from the machine, as well as decision performance and efficiency under two volatility states. The high-similarity KB design is based on similarity-attraction theory (Byrne & Stefaniak, 1967). Under this treatment condition the KB is constructed using an AF-ANN trained on the primary IDSS user to create a behaviorally similar artifact (KB_H). The second treatment condition involved the utilization of an IDSS with a KB trained using the AF-ANN for a different user (KB_L), to create a behaviorally dissimilar artifact.

Two IDSS artifacts were utilized to support traders in the two treatment groups:

high-similarity and low-similarity. Both systems were identical, with the exception of the AF-ANN KB. Under the USER + KB_H condition, representing the high-similarity state, subjects made trading decisions with an IDSS equipped with a KB trained on their individual trading decisions. In the USER + KB_L condition, representing the low-similarity state, subjects made trading decisions with an IDSS equipped with a KB trained on another trader's trading decisions.

In terms of manipulating the KB for the treatment of low-similarity, a trade matching algorithm was utilized. The trades from the trader KB with the most extreme option from that generated by the KB of the actual trader were used to populate the IDSS recommendation queue. In situations when all AF-ANN KBs generated the same recommendation, the algorithm randomly selected an option recommendation 2 strike prices away from the recommendation of the AF-ANN KB. During the experiment, the number of trades where all AF-ANN KBs agreed on the exact same trade was less the 2% of the total. The algorithm was designed to ensure that the recommendations presented to the traders in the low-similarity state were different than those that were generated by the AF-ANN KB, yet were realistic and not extreme.

In terms of AF-ANN architecture, a two-layer feed-forward neural network with sigmoid hidden and output neurons was developed and tested with the Matlab software application. Selecting the optimal number of hidden neurons and hidden layers is highly problem dependent, and is often the product of experimentation (Azoff, 1994). In this study the hidden number of neurons was selected based on the approach described by Tan (2001). A small number of hidden neurons was first used and then gradually increased. The procedure started with 1 hidden layer, containing 10 hidden neurons. Training of the network was conducted until a maximum of 100 epochs were completed without achieving a new low mean-squared error (MSE). An epoch represents each cycle in the training of the ANN, or more specifically, each instance in which the network is presented with a new input pattern. A new neural network was then developed with the number of hidden neurons increased by 1. The training and in-sample validation and performance measurement process was then repeated. After each successive trial, the performance of the network was assessed to determine if the new network structure was superior to its predecessor. This iterative process continued until the subsequent network structure reached an asymptote in performance based on mean-squared error (MSE), or produced inferior in-sample results. A total of 18 neurons were ultimately selected for use in the hidden layer. This architecture was considered to be robust to generalize with out-of-sample datasets without concern of over-fitting.

The network was trained with a Levenberg-Marquardt algorithm. The network consisted of 6 input units, and one output unit, which represents the dichotomous trade opinion (0 = no trade) or (1 = trade). The fundamental objective of option trading is pattern recognition (or nonlinear discriminant analysis). The objective of the network therefore was to classify an option trade as either a "trade" or "no-trade" based on the level of certain key technical metrics in the data input-vector.

In terms of the data input-vector, a total of 6 technical metrics were selected for purposes of training the ANN (Table 3). Each of the technical metrics are common to what a trained and experienced option trader would use in order to evaluate an option candidate and subsequently make a trading decision. The technical analysis metrics can be broadly classified into three major areas: (1) Trend Analysis, (2) Momentum Analysis, and (3) Option Metrics. Trend Analysis consists of the relationship of the underlying stock price to its 60-Day Moving Average (DMAVG). Traders often use a common stock's relationship to its 60-DMAVG as a technical indicator to discern trend and direction. If a stock is climbing above its 60-DMAVG, option traders may be more aggressive in increasing the number of call option contracts against the position. The converse of this relationship is also true with respect to option coverage.

The second technical metric is Momentum Analysis, which consists of both the upper and lower Bollinger Bands. Bollinger Bands are used by traders to measure the technical “highs” or "lows” of a stock’s price relative to a previous trading range. Specifically, the 2-Standard Deviation Bollinger Bands are often used to evaluate a stock’s trading range relative to its 20-Day Moving Average. By evaluating the width of the Bollinger Bands, option traders can gauge the tractability of increasing or decreasing their position size. Option traders are often more active when the width of the Bollinger Bands is wide, and are often less active when these bands move closer together. The relationship (ratio) of the stock’s underlying price to its respective upper and lower Bollinger Bands was used in the input data-vector, since this a normative trading metric used by option traders.

The third technical metric used in training ANN relates to three option specific metrics. The first consists of the ratio of the stock’s current price to the underlying strike price of the selected option. This relationship is often used to assess the risk of a particular option candidate. The second option specific metric is the yield of the underlying option premium relative to the strike price of the stock. This particular metric is used in conjunction with the stock price to strike price ratio in order to evaluate the risk-to-reward relationship for an option trade. The higher the level of this relative yield, the more likely the trader is to select the trade. The third option specific metric is the log-normal probability of the likelihood that the stock will expire in-the-money (ITM) at expiration. Often referred to as ITM probability, this metric is used extensively by option traders to evaluate the statistical risk of making a profitable trading decision. If this probability is low, say below a threshold like 20%, then the trader interprets this as 80% chance the stock option will expire out-of-the-money (OTM), allowing them to capture the entirety of their option premium at expiration with a relatively high degree of confidence. An option is classified as ITM or OTM based on the relationship of the underlying stock price to the option strike price. ITM options occur when the underlying stock trades higher than the strike price. The inverse is true for OTM options. Covered call option traders prefer positions to be OTM at expiration.

Table 3: ANN Input Data Vector

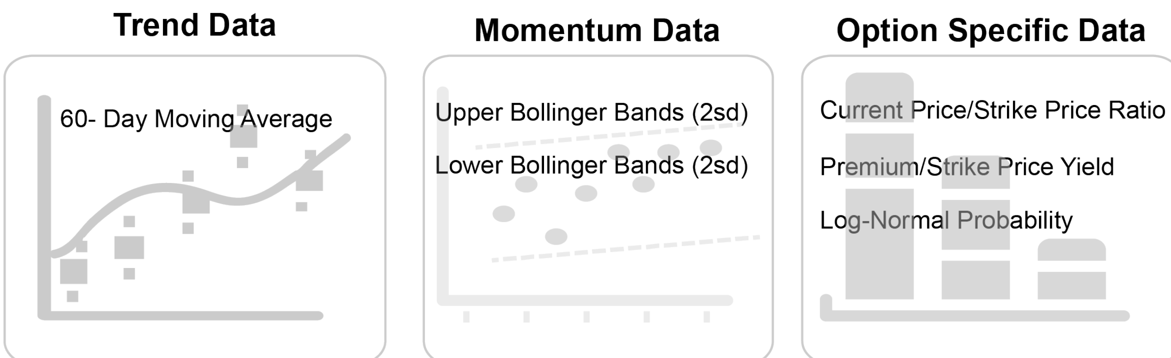

In order to construct the knowledge base for the prototype IDSS, a historical data set with approximately 354 trades was used to construct the input data-vector for use with the AF-ANN. The AF-ANN was used to capture and represent each trader's preferences and utility function based on a dataset of trading decisions derived from each DM over the preceding 6-month period. Specifically, the 6 data points from the aforementioned metrics of Trend Analysis, Momentum Analysis, and Option Metrics were used to train the AF-ANN, with the dichotomous trader response for trading decision representing the output vector. Figure 2 below provides an overview of the ANN architecture used to create the KB of STANN.

With respect to training the AF-ANN, the historical trading dataset was segmented into a training sample (70%), a validation sample (15%) and a test sample (15%) consistent with Kaufman (1998). The input vectors and the corresponding target vectors (trade decision) were used to train the network until it could reasonably associate the input vectors with the specific output vector based on minimizing the MSE of the ANN.

Figure 2: AF-ANN Architecture

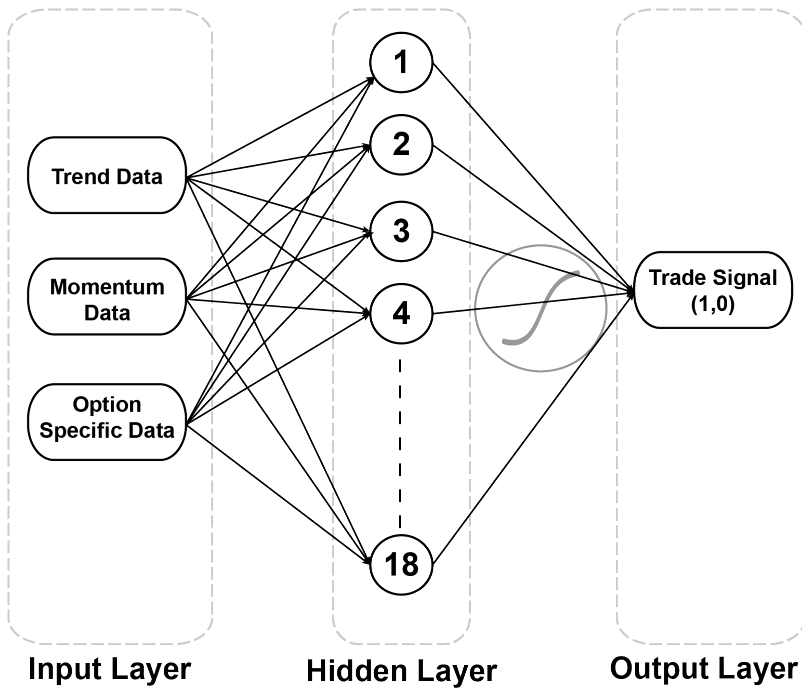

4.2.1.2 User Interface

The interface and graphics for both IDSS artifacts was identical for each treatment group, with the exception of the name label on the display avatar. STANN(1) represented the high-similarity KB_H state, and STANN(2) represented low-similarity KB_L state. Only the researcher knew the meaning of the IDSS titles. In addition, the datasets containing actual trading data in real-time market conditions were consistent across subjects and IDSS artifacts, with traders evaluating trades at the same time and interval. Each treatment group was provided underlying equity positions with comparable implied volatility and option tenor for trade consideration. In addition to the recommendation page, the IDSS provided a series of displays the traders could utilize to analyze the presented trades (Figure 3). Careful consideration was given to the data and displays that STANN provided, based on the three technical analysis metrics used in the underlying AF-ANN KB: (1) Trend Analysis, (2) Momentum Analysis, and (3) Option Metrics.

Figure 3: STANN Display

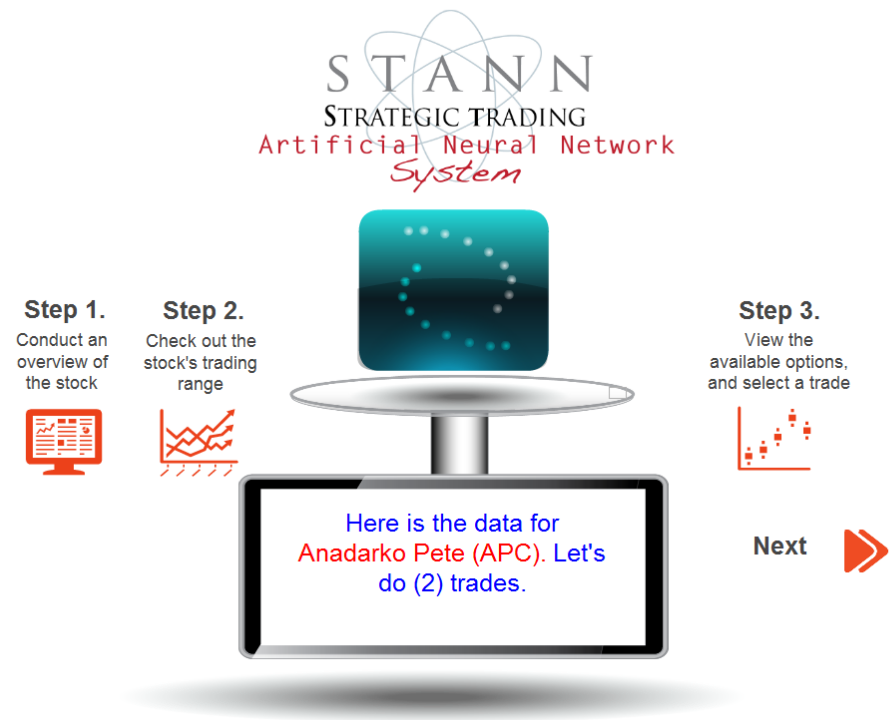

4.2.2 Implementation Procedures

To control for potentially confounding effects in the design of the experiment, and to avoid bias due to the learning effect, careful consideration was given to the design of the experiment. A repeated-measures design was selected in order to control for potential confounds based on individual trader characteristics. With respect to implementation, all subjects were randomly assigned to one of the two treatment groups, USER + KB_H and USER + KB_L (2 participants to each group). All subjects evaluated the same trades at the same time, one group of (2) in the high-similarity condition, and the other group of (2) in the low-similarity condition. Each trading group continued trading throughout the course of the experiment until the target number of trades was achieved (54). Once the prescribed number of trades was reached, traders were switched to the other treatment group. The traders were informed that they were being switched to a different system, but the difference was not revealed. Figure 4 below provides an outline of the implementation of the experiment.

Figure 4: Trade Sequencing

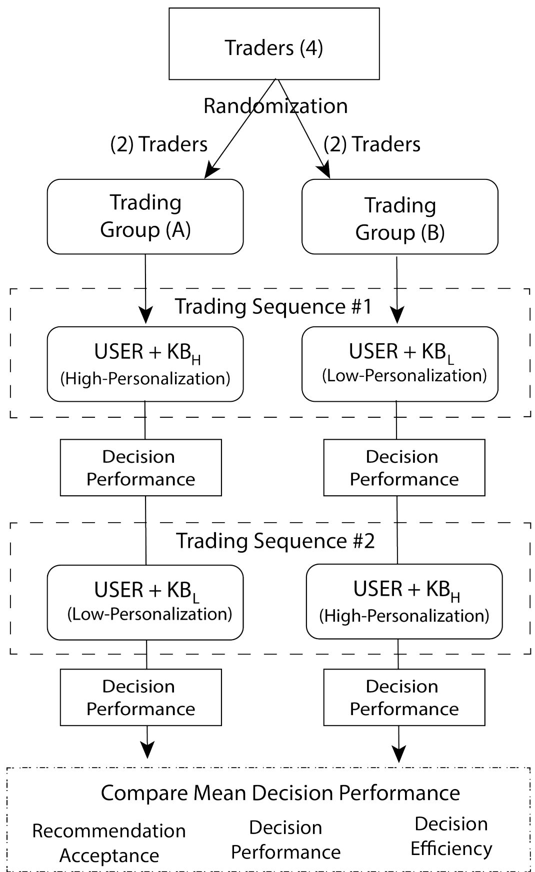

4.2.2.1 Trade Selection Mechanics

Selecting the trades for use in the experiment was an important consideration with respect to design. Given that the goal was to allow the traders to operate in a real-time environment without artificially manipulating the evaluation and selection of underlying securities, certain assumptions were made in order to maintain the use of underlying portfolio positions. The first relates to the assumption of temporal stability, where the value of an observation does not depend on when the treatment is delivered to the subjects (Rubin, 1974; 1978). In an option trading context, both implied and intrinsic volatility should be perceived equivalently by traders, irrespective of the time and/or sequence in which a position is evaluated by a trader.

In order to control for potential differences in position data presented to the traders, each trader was assigned the same security at the same time during the respective trading sessions. In order to ensure a homogeneous trading opportunity set for the experiment, portfolio positions with comparable levels of intrinsic volatility and option tenor were selected from the strategy portfolio.

The timing and selection of trades for evaluation by each trader was based on elements of technical analysis for the underlying equity position, as well as the underlying strategic objectives of the portfolio. As the portfolio positions reached certain technical trading levels, call option candidates were populated into the IDSS trading queue for evaluation and execution by the traders. The option trade candidates were out-of-the money (OTM) contracts from the two nearest terms option expiration months. This selection and queuing methodology was designed to ensure that each trader was processing the same market information, and equity and call option data simultaneously. The trade selection process continued until the minimum required number of trades was achieved for each trading run.

4.2.2.2 Decision-Maker Integration

The experiment was conducted over a 129-day period, with each subject evaluating a minimum of 112 trades (56 trades under each treatment combination). In addition to the actual experiment, each trader participated in a total of 3 pilot tests on approximately 35 trades prior to the official experiment. The pilot studies were designed to orient the subjects to the use of the system, as well as to test the software prototype and evaluate the tracking database for the dependent variables.

Each trading session started with a queuing session, where each subject was asked to load the prescribed IDSS on their workstation. During each trading session subjects were presented with a series of option trades to evaluate and execute. The IDSS artifact first presented a display of the current technical and quantitative metrics that are used under normative conditions to evaluate an option trade. At this point a timer was started and subjects were asked to review the respective trades as presented in the IDSS trading queue. Once the trades were evaluated and selected, subjects were asked to record their decision in the trading queue on the IDSS.

After making their initial trading decision, subjects were directed to access the "recommendation" page of the IDSS. On this page, the technical and quantitative metrics of the position were displayed once again. At this point, the subjects were asked to select an icon labeled "provide guidance". After approximately 15-20 seconds of processing time the IDSS presented its recommendation to the trader. This step in the process served as the decision-process outcome similarity manipulation. The subject was then allowed to conduct additional research using normative measures if desired. In the final step of the process the subject was then asked to record their final trading decision. At this point the trader could either accept the IDSS recommendation, or override the machine’s solution in deference to their initial decision. Figure 5 provides an outline of the sequencing of decisions for the experiment.

Figure 5: Decision Sequencing

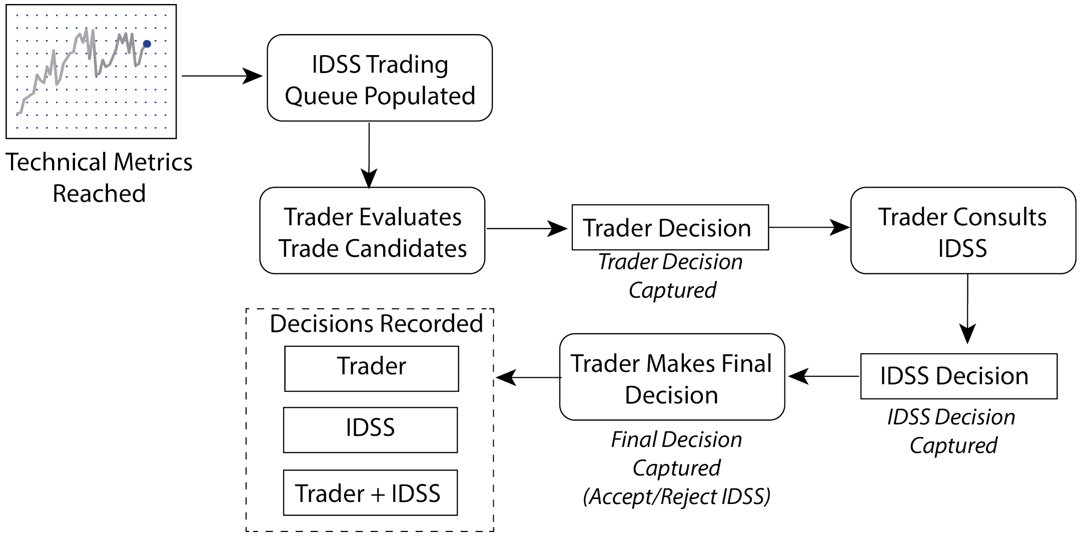

5 FINDINGS

5.1 Evaluation of Hypotheses

The objective of this study was to explore the use of behavioral similarity as a design element to positively influence certain dimensions of a human-machine investment decision system. Specifically, this study was designed to explore the acceptance rate of trading advice from an IDSS for purposes of improving decision performance and efficiency in an uncertain decision context. The following sections outline the results of the experiment with the prototype IDSS across the three primary hypotheses.

5.1.1 Recommendation Acceptance (H1)

To date, the concept of using behavioral similarity in an intelligent decision system to increase recommendation acceptance has been unexplored in the literature. Recommendation acceptance is a dichotomous variable that records when a trader agrees with, or accepts, the recommendation from the IDSS. Table 4 illustrates recommendation acceptance across the respective conditions based on the percentage of trades where the advice of the artifact was accepted.

Table 4: Recommendation Acceptance Results

	Similarity State			
	Low-Similarity		High-Similarity	
Volatility State	*N*	*M*	*N*	*M*
Low-Volatility Market	112	14.30%	112	43.80%
High-Volatility Market	112	16.10%	112	42.90%

A non-parametric Cochran's Q analysis test was first used to test for statistical differences in recommendation acceptance (i.e., number of trades in which the trader accepted the recommendation of the IDSS) across the two similarity and volatility states (k=4). This particular test was selected based on the fact that in the analysis of repeated-measures two-way randomized treatments, where the response variable consists of a dichotomous outcome (0 and 1), Cochran's Q test is a powerful non-parametric statistical test to determine if the 4-treatment levels have identical effects (Conover, 1999). In this analysis *Recommendation Acceptance* was coded as 1= agree, and 0 = disagree.

Based on an overall significance test, a systematic difference in the level of recommendation acceptance with the IDSS was found, [$Q(3, N = 112) = 44.38$, $p < .001$]. However, this omnibus test does not provide specific information about the pattern of differences across the individual conditions. In order to evaluate the pattern with respect to the research hypotheses, it was necessary to conduct two pair-wise comparisons among the respective conditions. Consistent with the research hypotheses, comparisons are made with respect to recommendation acceptance based on similarity state (high and low) and the level of underlying volatility (high and low). In order to extend the level of analysis based on a pair-wise comparison, a form of the chi-square test for within-subjects designs called McNemar's chi-square was selected. McNemar's test is a non-parametric test that is often used to compare two population proportions that are often related to each other, where the response variable consists of a dichotomous outcome. This test was selected to further evaluate the results of the omnibus Cochran's Q test for k=4 levels, to now evaluate the dichotomous response variable based on a k=2 comparison of similarity state (high and low) based on the level of underlying volatility (high and low).

Based on the results of a McNemar Test in the low-volatility market state, trades executed with the behaviorally similar IDSS (high-similarity state) agreed with the recommendations provided by the artifact 43.8% of the time, compared with only 14.3% for trades evaluated with the behaviorally dissimilar IDSS (low-similarity state), $p<.001$. With respect to the high-volatility market state, trades executed with the behaviorally similar KB (high-similarity state) agreed with the recommendation provided by the artifact 42.9% of the time, compared with only 16.10% for trades evaluated with the behaviorally dissimilar KB (low-similarity state), $p<.001$. An evaluation based on volatility state failed to provide any meaningful results based on a McNemar Test for differences in volatility in the high-similarity state ($p = .88$), and in the low-similarity state ($p=.85$).

Based on the results of the analysis, it is clear that the utilization of the behaviorally similar IDSS improved the acceptance rate of machine recommendations, compared to the low-similarity IDSS (Figure 6). As a result, H1a was fully supported by the results of the analysis. However, the hypothesis (H1b) that volatility in the decision context should strengthen the influence of a behaviorally similar IDSS on the acceptance rate of machine recommendations was not supported.

Figure 6: Recommendation Acceptance

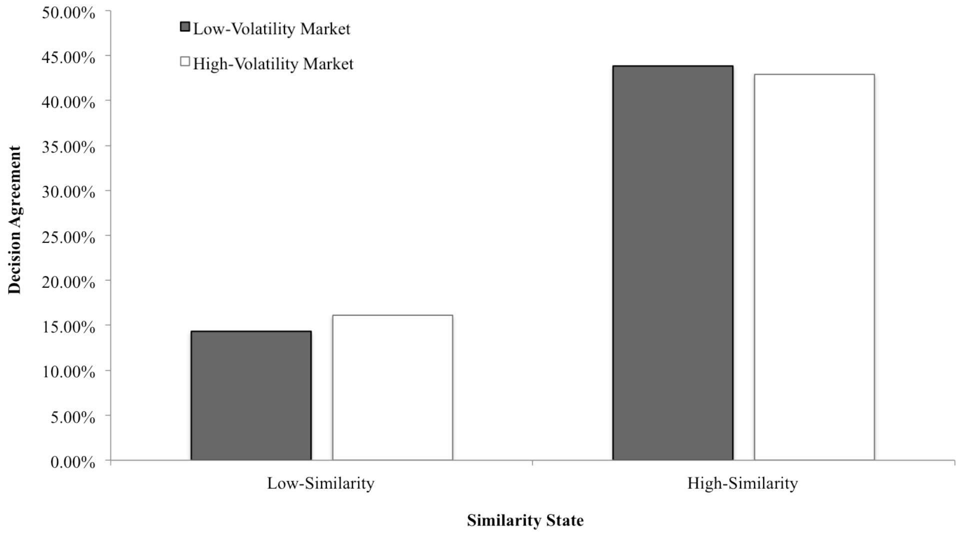

5.1.2 Decision Performance (H2)

As previously mentioned, despite the fact that research has shown a causal relationship between involvement with technology and user attitudes and acceptance of IS, the effects of decision aid recommendation acceptance on human decision-making outcomes has been largely unexplored (Hess et al., 2006). As a result, this dissertation is intended to provide support for the notion of positively influencing decision performance based on increasing the acceptance of recommendations from an IDSS.

The purpose of H2 is to investigate the influence of the two types of IDSS similarity-state on decision performance. In the information systems literature this construct is often referred to analogously as effectiveness, and is evaluated based on the accuracy of a decision compared to a normative solution for an individual DM (Payne et al, 1993). In this study, decision performance is an objective measure defined as the mean option premium generated from the trades executed by the combined man-machine system under the two volatility states. This definition of decision performance is comparable to the normative measure by which individual traders are evaluated in practice.

Data Exploration

One of the important assumptions in the repeated measures analysis of variance procedure is that the variance/covariance matrix of the observed data follows a particular pattern. Referred to as sphericity, this pattern is typically characterized with equal variances in the diagonal, and equal covariance in the off-diagonal elements. However, it is not necessary to test the sphericity assumption in this experiment since a repeated measures factor with only two levels is utilized. As a result, the sphericity assumption is satisfied since there is essentially only one covariance.

Another important assumption in the use of a parametric test like the analysis of variance is the assumption of normality in the distribution of the dependent variable. Normality is an assumption that the data are derived from a normal distribution. To test the assumption of normality an exploratory data analysis was conducted across the respective effects. Fisher Skewness Coefficient (Z=skewness/standard error; Z between+/-1.96) was used to evaluate the normality of variables (Tabachnick and Fidell, 2001). The results of the Z-skewness tests highlighted the fact that most Z-values for decision time fell outside of +1.96 to -1.96, implying that the data failed to meet the normality assumption (Hung et al., 2005). Figure 9 shows the distribution of each level of the dependent variable. A visual inspection of Figure 9 also highlights the fact that the data are significantly non-normal, with noticeable positive skew. Furthermore, the results of a Quartile-Quartile (Q-Q) plot revealed reasonably significant deviations from the observed line across the respective levels of the dependent variable (Figure 10).

A second goal of the exploratory data analysis was to test for the presence of outliers in the dependent variable. The presence of outliers in the data can be a potential threat to the validity of the results by further contributing to skew and non-normality across levels of the dependent variable. Furthermore, outliers may result in biased parameter estimation, misspecification, and misleading results from the data analysis. It is therefore important to identify outliers prior to conducting modeling and analysis of data (Williams et al., 2002; Liu et al., 2004). To check for the presence of outliers in the data box plots were used to identify values >1.5 times the interquartile range away from the median. Figure 11 shows a significant number of outliers in the underlying data, along with significant positive skew. Skewness is an extremely common phenomenon in financial data. Most financial datasets, including asset prices, asset returns and option premia, intrinsically have either positive or negative skew. As a result, outlier removal and data transformation are often employed in practice in order to utilize parametric statistical techniques.

To minimize the adverse impact of extreme values in the data, outliers at or over 2.5 standard deviations from the mean were evaluated (Brase & Pellillo, 2012). Based on this approach a total of 9 individual outliers were identified. Upon closer observation of the outliers, these extreme values appeared to be the result of aberrant trades that were executed at extremely high stock-price to strike-price ratios. In a normative trading context, this type of trade is often considered to be a trading error. As a result, a total of 9 outliers and extreme values were removed from the data. Despite the removal of these extreme values the assumption of normality was still violated based on subsequent Z-skewness tests.

To contend with the violation of the normality assumption for use in the repeated-measures analysis of variance, a log10 transformation was first utilized to contend with the observable positive skew in the data. However, subsequent results of Z-skewness tests showed that the transformed data also failed to satisfy the assumption of normality. As a result, the Aligned Rank Transform (ART) method was utilized. Rank transformation procedures were originally proposed as a technique for dealing with violations of normality and sphericity (Conover and Iman, 1981). For repeated-measures designs, the analysis of variance *F*-test was found to be robust to violations of normality when performed on ranked data (Zimmerman and Zumbo, 1993). The ART procedure was conducted on the decision performance data consistent with Wobbrock, et al. (2011).

Results

The mean scores and standard deviations for decision performance, measured as gross option premium received for each trade, are outlined in Table 5. In the low-volatility state, trades made with the low-similarity IDSS had higher performance (M=1804, SD=944) compared to trades made with the high-similarity IDSS (M=1568, SD=1149). However, as highlighted in Table 5 below, the performance of trades made with the low-similarity IDSS deteriorated significantly in a higher volatility environment (M=1402, SD=1201). In contrast to the performance of the low-similarity IDSS, trades made with the high-similarity IDSS exhibited a higher degree of relative stability and overall performance in the higher volatility environment (M=1512, SD=1062).

A two-way within subjects repeated-measures analysis of variance was used to investigate the statistical significance of differences in performance. The within-subjects factors were *Similarity State*, *Volatility State,* and the interaction of the two respective conditions.

Table 5: Mean Trade Performance

	Similarity State					
	Low-Similarity			High-Similarity		
Volatility State	*N*	*M*	*SD*	*N*	*M*	*SD*
Low-Volatility Market	103	1804	944	103	1568	1149
High-Volatility Market	103	1402	1202	103	1512	1062

The results of the analysis of variance are presented in Table 6. The analysis revealed that the mean performance scores for *Similarity State* were not significantly different, F(1,102) = 1.125, p = .291. The results of the analysis of variance on the main effect of *Similarity State* failed to support the hypothesis that the utilization of a behaviorally similar IDSS will increase trading performance (H2a). *Volatility State* as a main effect was found to be statistically significant in the expected direction F(1,102)= 7.57, p = .007, $\eta_p 2$ =.069. With respect to H2b, the interaction of *Volatility State* and *Similarity State* was found to be statistically significant, F(1,102)= 5.537, p = .021, $\eta_p 2$ = .051.

Table 6: Mean Performance ANOVA Results

Source of Variation	*MS*	*df*	*F*	*p*	η_p^2
Similarity State (KB)	15365	1	1.13	.291	.011
Volatility	83633	1	7.57	.007**	.069
Similarity State (KB) x Volatility	60825	1	5.54	.021*	.051

* p< .05, **p< .01

Based on the results of the ANOVA, it was found that volatility in the decision context strengthened the influence of a behaviorally similar IDSS on trading performance, fully supporting H2b. While the effect size is small based on $\eta_p 2$ = .051 (Cohen, 1988), the findings of the analysis are particularly beneficial to practitioners where stability in performance during elevated periods of volatility is a critical element in the long-term viability of a portfolio. Figure

7 below graphically illustrates the interaction of mean performance based on *Volatility State*.

Figure 7: Mean Performance Interaction

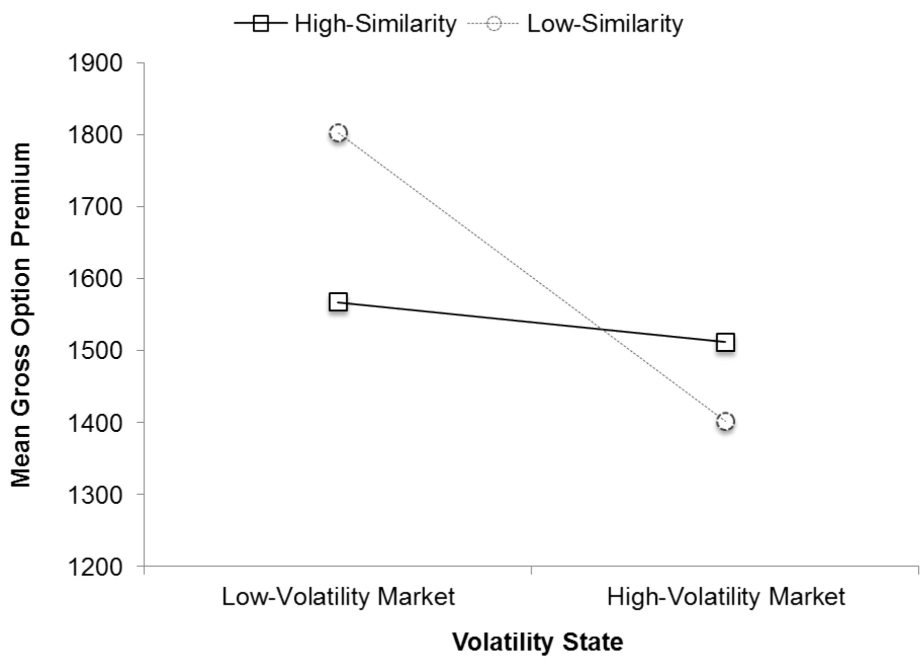

With respect to H2c, the utilization of a behaviorally similar IDSS did not consistently result in a lower standard deviation of trading performance across both volatility states. As a result H2c was not supported. However, as shown in Table 5, the utilization of a behaviorally similar IDSS did result in a lower standard deviation of trading performance in the high-volatility environment (*SD*=1062), compared to the low-volatility environment (*SD*=1201). As a result, H2d was fully supported based on lower standard deviation of trading performance in the high-volatility environment.

5.1.3 Decision Efficiency (H3)

The time required to make a decision (decision efficiency) has been previously evaluated as an important metric of MIS success (Raymond, 1985). In defining system success, Seddon (1997) defined efficiency as more work done in the same time, or less time for more work of equivalent quality. In this study, decision-making efficiency is defined as the total time required to evaluate and execute a covered-call option trade. The time was recorded from the launch of the IDSS by the trader, to the completion of the final trading decision. Based on the design of this study, trades evaluated and executed with a high-similarity IDSS should have a lower mean decision time per trade, compared to trades executed with a low-similarity IDSS. Reducing the time required to make a trading decision can enable traders to focus on additional opportunities and execute more profitable trades over the course of a trading day. Based on the premise that behavioral similarity matters in the design of IDSS, the hypotheses regarding the effects of perceived human-machine similarity and decision-making efficiency are explored.

Data Exploration

To test the assumption of normality in the dependent variable of decision time, an exploratory data analysis was conducted across the respective effects. Based on the Fisher Skewness Coefficient, the data for decision time violated the assumption of normality. Figure 12 shows a histogram of the underlying data. As evidenced by Figure 12, the data exhibits significant positive-skew, with the presence of outliers. The results of a Quartile-Quartile (Q-Q) plot also revealed a significant deviation from normality (Figure 13).

With respect to decision time, there are a number of reasons that outliers in the data may be present. Since the experiment was conducted in a noisy real-time trading environment, traders may have received a phone call, a distraction from a colleague, or momentarily switched to another more pressing task prior to completing the assigned trade. Box plots were used to visually inspect for the presence of outliers in the data, defined as values >1.5 times the interquartile range away from the median. Figure 14 shows a significant number of outliers in the underlying decision time data.

To alleviate the adverse impact of extreme values in the decision time data, outliers at or over 2.5 standard deviations from the mean were removed (Brase & Pellillo, 2012). In total, 13 outliers and extreme values were removed from the data, leaving a total of 99 trades in the sample. Despite the removal of these extreme values, the data still violated the assumption of normality based on subsequent Z-skewness tests. As a result, a log10 transformation was applied to the dependent variable after the extreme data points were removed. A log10 transformation was conducted in accordance with Tabachnick and Fidell (2007) and Howell (2007), based on the presence of positive-skew in the decision time data. The results of a subsequent Z-skewness test on the transformed data revealed values between +1.96 and -1.96, implying that the data satisfactorily met the assumption of normality (Hung et al., 2005).

Results

The mean scores and standard deviations for decision time, measured in decimal form as the total time in minutes required for a trader to fully explore and execute each trade, are outlined in Table 7. In the low-volatility state, trades executed with the low-similarity IDSS had a lower mean decision time (M=0.86, SD=.31) compared to trades executed with the high-similarity IDSS (M=1.02, SD=0.19).

Table 7: Mean Decision Time (minutes)

	Similarity State					
	Low-Similarity (S2)			High-Similarity (S1)		
Volatility State	*N*	*M*	*SD*	*N*	*M*	*SD*
Low-Volatility Market	99	0.86	0.31	99	1.02	0.19
High-Volatility Market	99	0.83	0.37	99	1.06	0.21

Trades executed in the high-volatility state exhibited a similar patter with respect to decision time. A lower mean decision time resulted for the low-similarity IDSS (M=0.83, SD=.37) compared to trades made with the high-similarity IDSS (M=1.06, SD=0.21).

A two-way within subjects repeated-measures analysis of variance was used to investigate the statistical significance of differences in decision time. The within-subjects factors were *Similarity State*, *Volatility State,* and the interaction of the two respective conditions. The results of the analysis of variance are presented in Table 8. The analysis of variance showed that the mean scores for *Similarity State* were significantly different, $F(1,98) = 8.02$, $p = .006$, $\eta_p 2 =.076$, with a medium effect size (Cohen, 1988).

Table 8: Mean Decision Time ANOVA Results

Source of Variation	*MS*	*df*	*F*	*p*	η_p^2
Similarity State (KB)	0.828	1	8.02	.006**	.076
Volatility	0.00007	1	0.002	.969	0
Similarity State (KB) x Volatility	0.024	1	0.58	.447	.006

** p< .01

The results of the analysis of variance on the main effect of *Similarity State* failed to support the hypothesis that the utilization of a behaviorally similar IDSS will increase decision efficiency by reducing the amount of time required to execute a trade (H3a).

While the main effect of Similarity State was found to be statistically significant, trades executed with the high-similarity IDSS were actually slower than the trades executed with the low-similarity IDSS (Figure 8). *Volatility State* as a main effect was not found to be statistically significant $F(1,98)= 0.002$, $p = .969$. With respect to H3b, the interaction of *Volatility State* and *Similarity State* was also found to be statistically insignificant, $F(1,98)= 0.58$, $p = .447$. Thus, neither H3a nor H3b were supported.

Figure 8: Graph of Mean Decision Time

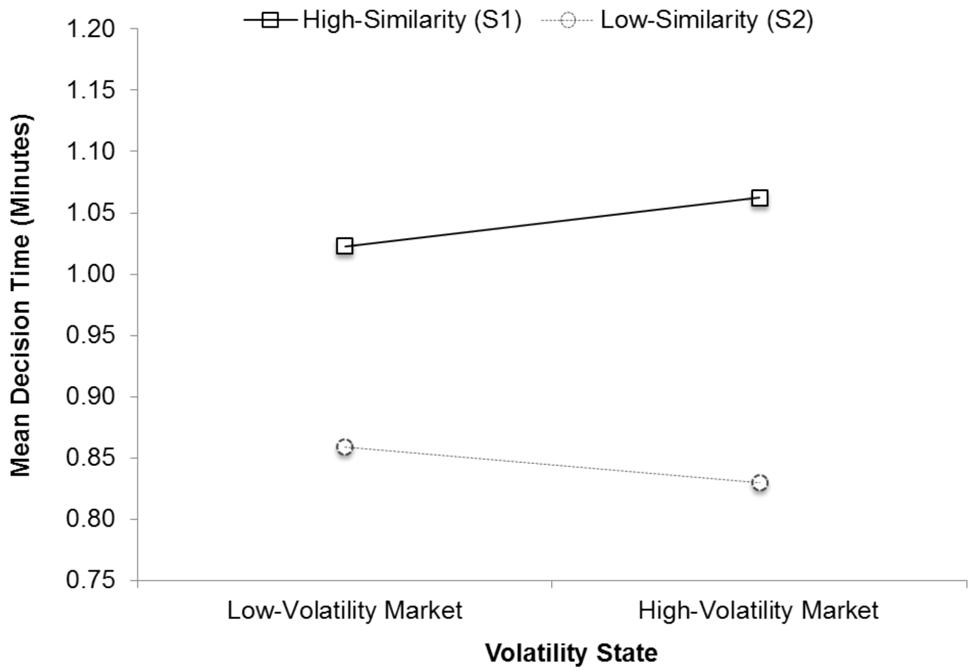

6. DISCUSSION AND IMPLICATIONS

6.1 Discussion

The findings of this study have meaningful implications for both theory and practice. In terms of contributions to theory, the findings of this study show promising results with respect to intelligent decision aid adoption based on behavioral similarity as a design trait. In addition, this study showed that an increase in IDSS adoption for purposes of decision support positively influenced decision-making performance in periods of increased uncertainty.

While the literature is replete with studies of similarity in an e-commerce context, an empirical understanding of behavioral similarity with respect to intelligent systems and their impact on decision performance and efficiency under conditions of uncertainty is needed in the HCI literature. HCI researchers could especially benefit from a greater understanding of the interaction dynamics between human and machine in a real-time semi-structured decision domain, like financial services. This study is an important contribution to the HCI literature because human-computer cooperative problem solving has been an omnipresent issue in the field.

In terms of contributions to practice, the results of this study provide some interesting insights for the design of intelligent decision aids. First, designers of intelligent systems could greatly benefit from design features that increase the acceptance rate of machine recommendations. This would represent an important contribution for the use of decision aids since system adoption is often a necessary antecedent to performance amplification. Any improvement in recommendation acceptance will be beneficial for improving decision performance. In addition, a greater understanding of human-machine integration dynamics during periods of increased market uncertainty will ultimately allow systems architects to design more effective trading tools. Any modest improvement in performance during periods of higher volatility could provide a significant competitive advantage in the market place. Even small improvements in performance during these periods could add up to be a significant monetary value.

The application of AF-ANN technology to capture the decision preferences of human traders also represents a potential advancement for the financial services industry. While not a direct research hypothesis, the successful application of this KB architecture in the experiment could greatly benefit AI practitioners by providing empirical support for ANN powered decision aids in actual trading environments.

6.2 Limitations

While this research study has the potential to contribute to both theory and practice, expectations should be tempered with respect to some of the inherent limitations of the experiment. Conducting the experiment in a real-time trading domain alone contains a confluence of variables and exogenous factors that are hard to control. Secondly, the limited number of subjects participating in the experiment also limits the study. Although the unit of analysis is the individual trade, a larger number of subjects may help increase the power of the experiment and insulate the study from threats to validity.

The design of the IDSS itself is also a potentially limitation of the study. There could be elements in the AF-ANN design topology that may adversely impact the performance of the artifact based on market condition and direction. For example, the exclusion of the VIX as an ANN input could impact the ability of the ANN to generalize across market conditions. While the CBOE VIX Index is used by traders to evaluate volatility, it was excluded from the ANN topology since it is was not deemed important from a trade classification perspective. As a result, the IDSS will be unable to elicit the hypothesized perception and behavior from the traders if trades generated by the IDSS are substantially different from those of the traders themselves. This could also result from issues in the training of the AF-ANN itself, along with the potential for extreme out-of-sample input vectors for the artifact during the experimental period.

The presence of market volatility as a surrogate measure for uncertainty in the decision context may also be problematic. Given the relatively muted level of volatility during the experimental period it is possible that the two volatility states were not strongly defined with respect to the traders' perception of uncertainty. In an ideal situation trades would have been selected for the study on an a priori basis, screening for days when the VIX exceeded a moving average for example (i.e., trades would be assigned when there were large upward or downward moves in the VIX). However, this approach was not tractable based on a reasonable time frame for completion of the experiment. As result of the relatively lower levels of volatility, the hypothesized effects on the dependent variables may not be as pronounced.

The timing of each trading session is also a potential threat to the experiment. Since the market's direction is unpredictable, and as a result difficult to control for in the experiment, the results of the experiment could be questioned if it was determined that the direction of the market had an influence on participants in the study. Controlling for market direction and volatility simultaneously would be difficult given the market's unpredictable nature, and would represent a formidable implementation challenge. However, threats to results of this experiment could be minimized if the results pattern remained consistent during each of the tested market periods.

The selection of the within-subjects design presents a number of limitations, despite its strength as an experimental design. With respect to strengths, this particular design was selected primarily to reduce error variance that could result from individual differences in the participating subjects. The within-subjects design helps to guard against this given the fact that each subject essentially serves as their own control by being exposed to all treatment levels. While the unit of analysis in the experiment is the individual trade, and not the human subjects, there was a concern that individual differences in performance could compromise the design inflating the Type I error rate.

There are however a number of intrinsic disadvantages of the within-subjects design. One of the most common limitations is referred to as carryover effects. In general, this means a subjects' participation in one condition may affect performance in the other condition. For example, there could be carryover effects based on which decision system the traders used first. This would potentially mitigate the ability to detect a difference in the decision similarity of the respective IDSS artifacts. In addition, given the duration of the experiment, subjects could fall victim to both practice and fatigue effects. Subjects may possibly be more fatigued towards the end of the experiment and therefore were less responsive to the treatment effects than at the beginning of the experiment. Despite the limitations of the repeated measures design, it is felt that is the most tractable with respect to the domain of the experiment.

7. CONCLUSIONS

The impact of intelligent decision aids on decision-making process and outcome has been studied extensively by researchers (Gupta et al., 2006; Linger and Burstein, 1997; Moreau, 2006; Phillips-Wren and Jain, 2005; Roth et al., 1987). However, despite the scope and depth of this existing research, the influence of IDSS as part of a joint human-machine cognitive system on decision performance and efficiency under conditions of uncertainty remains unexplored. As a result, this dissertation explored a topic that is of particular interest to both practitioners and researchers.

In pursuit of this research topic a theoretically grounded prototype system was developed and implemented for use in covered-call options trading in a large financial services firm. The IDSS prototype was developed based on the theoretical premise that technological artifacts are often perceived as social actors, and as a result, users often ascribe behavioral characteristics to inanimate machines (Reeves & Nass, 1996). In the IDSS prototype, these behavioral characteristics were manipulated using a specialized knowledge base in order to engender the perception of similarity between system users and the IDSS. The basis for this approach is the "similarity-attraction hypothesis" which predicts that humans prefer to interact with others who are perceived to be similar to themselves.

To evaluate the prototype IDSS an experiment in a real-time decision domain was conducted. The experiment consisted of four human subjects using the prototype IDSS to make equity call-option trades over a 129 day period. The IDSS was used to investigate the effects of perceived behavioral similarity as a design element on the dependent variables of recommendation acceptance, decision performance and efficiency under varying conditions of uncertainty (volatility) in the decision context. Uncertainty in the decision context, operationalized by using the level of the CBOE VIX Index, was a moderating variable in the study.

Recommendation Acceptance

Of the three dependent variables in the study, recommendation acceptance is arguably the most important given the design of the research model in this experiment. This is due to the notion that in order for any of the potential benefits of the IDSS to be realized, the trader would first have to take its advice. Since the value of information systems tends to be influenced by their actual use in decision-making (Devraj & Kohli, 2003), greater knowledge of the utilization of IDSS in a real-time decision context is needed. Furthermore, intelligent systems have failed to generate any meaningful performance results in practice given their low acceptance rates as decision aids in securities trading. Any underlying improvement in the acceptance rate of machine recommendations could potentially represent a significant advancement in the design of intelligent decision aids for use in financial services. As a result, this study investigated the following research question: can perceived behavioral similarity positively influence the frequency by which a human DM relies on advice from an IDSS under conditions of uncertainty?

Recommendation acceptance is also important given the design of the KB that powers the IDSS. While the AF-ANN approach is by no means new, this appears to be one of the first instances in the literature where this approach was used to capture human decision-making preferences for purposes of fostering the perception of similarity in an intelligent machine. A statistically discernible difference in the acceptance rate of machine recommendations further validates using ANN for this purpose, and could be an indication that the human DM perceived some level of behavioral similarity with their machine trading partner.

Based on the results of the analysis presented above, the behaviorally similar IDSS significantly improved the acceptance rate of machine recommendations compared to the low-similarity IDSS. However, while the results were as hypothesized for the main effects of *Similarity State*, volatility in the decision context did not appear to moderate the influence of a behaviorally similar IDSS on the acceptance rate of machine recommendations. While these results were not exactly as hypothesized, the findings of the main effect did provide an answer to the first research question. This is an important finding and should contribute to the literature on human-computer interaction (HCI) and IT adoption by highlighting the use of specific design features to influence the acceptance of advice from an IDSS.

Decision Performance

Research in DSS has focused primarily on measures of decision performance to evaluate system success (Todd & Benbasat, 1992). As previously discussed, the benefits of DSS on decision performance have been fairly mixed in the literature. Todd and Benbasat (1992) point out that some DSS studies have reported an improvement in decision performance, while others studies reported no improvement based on DSS use. Furthermore, some studies actually revealed degradation in decision performance based on DSS use (Benbasat & Nault, 1990; Sharda, et al., 1988).

Decision performance is an important metric for DSS evaluation in general, and for the evaluation of intelligent decision aids (IDSS) in particular. The mixed performance results in the literature helped motivate the second research question: can perceived behavioral similarity positively influence the decision-making performance of a joint human-IDSS cognitive system under conditions of uncertainty? The pursuit of this question was intended to provide support for the notion of positively influencing decision performance based on increased recommendation acceptance from the IDSS. The second research question was also motivated by the desire to improve the current state of the art in intelligent systems for use in industry, where the efficacy of these systems is poor.

Decision performance was objectively measured in the experiment as the mean option premium generated from the trades executed by the combined man-machine system under the respective treatment conditions. The unit of analysis consisted of all of the trades evaluated and executed by the combined man-machine trading system. The results of the analysis of variance on the main effect of the *Similarity State* failed to provide conclusive support for the hypothesis that a behaviorally similar IDSS will increase decision performance. However, a difference in performance based on *Volatility State* as a main effect was found to be statistically significant. Of particular note in the results was the interaction effect between *Similarity State* and *Volatility State,* where differences in decision performance were detected based on the underlying state of volatility (Figure 7 above).

As hypothesized, the analysis revealed that decision performance was better for trades executed with the high-similarity IDSS during the high-volatility environment, compared to trades executed with the low-similarity IDSS. This is an important finding given the role of the decision context in the design of this study. As previously established, the decision context can play a major role in the quality of human decision-making. DMs typically have the potential to make accurate decisions, but often fail to as the result of external distractions or interference.

This phenomenon is referred to as an application error, in that a DM possesses the requisite cognitive skills and ability to make an appropriate decision, but distraction in the decision context inhibit the effective application of these skills (Kahneman and Tversky, 1982). Application errors are a routine and troublesome phenomenon in the investment management domain, especially during periods of elevated uncertainty (volatility).

A well-trained trader in practice typically performs well by following normative axioms of decision making during tranquil market periods. *Ceteris paribus*, in an environment without distraction and perturbation, slips and cognitive decision errors are relatively muted, and traders perform consistently. However, as the level of distraction and uncertainty in the decision context rises, traders often fail to follow these normative axioms. To cope with this issue, a specialized IDSS was used in the experiment to provide support to traders during volatile markets. It was hypothesized that a behaviorally similar IDSS would be a valuable member of the human-machine dyad during periods of elevated uncertainty since it would be immune to cognitive slips and application errors, unlike its human counterpart. Furthermore, traders should be more inclined to follow the recommendations of their behaviorally similar IDSS teammate. As a result, the behaviorally similar IDSS should help improve human-machine trading performance during periods of elevated uncertainty (volatility). The results of the experiment were consistent with the underlying hypothesis, supporting the notion that behavioral similarity matters as a design element in IDSS.

Another interesting result for this dependent variable occurred in the high-volatility state, where the data revealed that there was a significant degradation in decision performance for the low-similarity IDSS, compared to the trades of the high-similarity artifact. With the low-similarity IDSS, the mean gross option premium generated fell by approximately $402, compared to a decline of only $56 for the high-similarity IDSS. Based on these results, it appears that the behaviorally similar IDSS provided a level of decision support that was more consistent than the alternate system. This is an extremely important finding since stability in performance is a sought after goal in the field of investment management.

One of the most interesting results of the experiment was the fact that trades executed with the low-similarity IDSS actually performed better than the trades executed in the high-similarity IDSS in the low-volatility condition. A possible explanation can be derived by further examining what happens in a low-volatility market state, given the fact that the level of volatility (VIX) is operationalized as uncertainty in the decision context. A low-volatility market state represents a higher level of relative certainty in the decision-context. It is possible during this environment that traders feel more certain, and possibly complacent. As a result, it is possible that the low-similarity artifact provided some form of alternative perspective to the trades in question. While the traders did not explicitly agree with the low-similarity artifact as evidenced in the data in H1, perhaps the alternative perspective provided by the IDSS enabled them to evaluate a broader search space in periods of anemic uncertainty.

Another possible explanation for the performance of trades in the low-volatility state relates to the state of the market. In an environment with nominal levels of uncertainty and disruption, traders are able to make more consistent trading decisions. In this environment, traders are also less likely to make slips and application errors, further improving their performance.

As previously discussed, both performance and consistency of performance are necessary factors for a successful trading strategy. Consistency of performance, measured by the standard deviation of the portfolio, is an important metric in the evaluation of investment portfolios in the investment management industry. With respect to consistency in performance, measured in this experiment as the standard deviation of the mean gross option profits, the same results were found with respect to the influence of decision context. The utilization of a high-similarity IDSS did not consistently result in a lower standard deviation of trading performance across both volatility states. However, the utilization of a high-similarity IDSS did result in a lower standard deviation of trading performance in the high-volatility environment, compared to the low-volatility environment.

The findings with respect to decision performance represent an important contribution with respect to decision support in financial services. This is due to the fact that increased uncertainty in the decision context can subject humans to more cognitive biases, slips, and application errors. An IDSS design element that can attenuate the influence of these factors during periods of uncertainty holds great promise for applications in financial services, where consistency in performance across market conditions is in an important factor in portfolio management.

Decision Efficiency

The third dependent variable evaluated in the study related to how efficiently traders evaluated and executed covered-call option trades. The time required to make a decision is an important metric of MIS success (Raymond, 1985). Decision efficiency can be evaluated based on completing more work in the same amount of time, or less time for more work of equivalent quality (Seddon, 1997). In this study, decision efficiency was defined as the total time required, measured in minutes, to evaluate and execute a covered-call option trade. The total time was measured from time the trader launched the IDSS, to the completion of trade selection.

Similar in many respects to the literature on decision performance, the research that has been conducted regarding decision efficiency has been mixed. As a result, the following research question was posed: can perceived behavioral similarity positively influence the decision-making efficiency of a joint human-IDSS cognitive system under conditions of uncertainty? The answer to this research question is also important to practitioners in the financial services sector. Due to the rapidly evolving nature of the market environment, along with the prevalence of high-frequency and algorithmic trading, traders must make decisions quickly and without reservation in order to maximize their trading performance.

The results of this study revealed an interesting finding: trades evaluated and executed with the high-similarity IDSS actually took longer than the trades executed with the low-similarity IDSS. *Volatility State* as both a main and moderating effect was not found to be significant, meaning that trades with the high-similarity system took longer to evaluate and execute across both volatility conditions.

A number of explanations are possible for the observed results. Huse (1980) highlights that it may be expected that users of a DSS will often require more time to reach a decision as they orient and familiarize themselves with the system. However, given the duration of the experiment, it is unlikely that orientation effects resulted in a higher decision time. Another potential explanation could be based on the notion of involvement with the IDSS. Hess et al. (2006) found that personality similarity between a user and the decision aid resulted in increased involvement with the decision aid. As a result, it could be that perception of behavioral similarity positively increased the involvement with the IDSS, encouraging the trader to dedicate more time to information processing and trade evaluation. In many respects this would be a positive contribution of behavioral similarity as a design element, particularly if a causal connection between the amount of time the DM dedicated to information processing and decision performance could be established. In this context the additional improvement in performance may more than compensate for the additional time dedicated to evaluating a trade.

Directions for Future Research

While the results of this study revealed a few findings of interest, the research provided in this dissertation is only a first step. Future research should take into consideration and possibly correct for many of the limitations of this study. While the presented research explored and attempted to answer three research questions, it simultaneously opened up a number of new opportunities for contribution to theory and practice.

Although a link between behavioral similarity and the decision context with the underlying dependent variables has been established, the nature of that relationship across various levels of uncertainty is still not entirely clear. Future insight could be provided by conducting the experiment across more extreme ranges of volatility, where it is expected that decision errors are likely to be the most pronounced. For option trading in particular, degradation in performance is most prevalent when volatility reaches relative extremes. An exploration of a broader range of uncertainty in the decision context would serve to supplement the experimental findings of this study.

One of the most surprising findings of the experiment was the fact that decision amplification was found for traders using the low-similarity IDSS during periods of low volatility. Exploring this finding would be an interesting follow-on study from both a theoretical and practical perspective. While the focus of this study was based primarily on decision amplification during periods of elevated uncertainty, the exploration of the research questions at low levels of uncertainty would also be interesting. Given the fact that the low-similarity IDSS KB was derived from a trader other than the IDSS user, perhaps this system provided some form of alternative perspective to the traders during periods of low volatility. An exploration of this observed phenomenon could also provide additional theoretical support for the field of HCI in general, and for DSS researchers in particular.

The performance of the IDSS in absolute terms also represents an interesting research opportunity. As opposed to evaluating the combined performance of the man-machine synergistic decision system in this study, an interesting extension would be to evaluate the performance of man versus machine individually. Given the emerging trend towards cost savings and automation in financial services post the Great Recession of 2007, along with the amount of technological progress in AI technologies, institutional firms may look to complement their existing trading desks with stand-alone AI based trading applications. This would certainly be true if the efficacy of such technologies could be established.

Investigating the performance differential between man and machine in varying environments of uncertainty would also be a useful research endeavor. Since this experiment evaluated the decision performance and efficiency of the combined man-machine decision system, it would also be interesting to disaggregate the two in a man versus machine trading competition. A better theoretical and practical perspective as to when man and machine is most effective would assist systems architects in developing an optimal model for deployment of intelligent decision systems.

Future experiments with more subjects could also provide additional insight into the findings presented in this study. Although difficult to accomplish given the time and programming required to construct and implement an AF-ANN based on an individual, expanding the number of subjects and increasing the sample size of trades could provide additional support for the findings of this study.

REFERENCES

Abdelkader, A. (2006). A Cooperative Intelligent Decision Support System for Contingency Management. Journal of Computer Science 2 (10): 758-764, 2006

Ainscough, T.L., and J.E. Aronson (1999). A Neural Networks Approach for the Analysis of Scanner Data. *Journal of Retailing and Consumer Services*. Vol. 6.

Aksoy, L. (2006). Should Recommendation Agents Think Like People? *Journal of Service Research*, 8(4), 297–315.

Al-Natour, S. Benbasat, I., and Cenfetelli, R. (2011) The Adoption of Online Shopping Assistants: Perceived Similarity as an Antecedent to Evaluative Beliefs. Journal of the Association for Information Systems: Vol. 12: Iss. 5, Article 2

Al-Natour, S., and Benbasat, I. (2005). The role of similarity in e-commerce interactions: The case of online shopping assistants. … of the Fourth Annual Workshop on ….

Al-Natour, S., Benbasat, I., and Cenfetelli, R. (2006) The Role of Design Characteristics in Shaping Perceptions of Similarity: The Case of Online Shopping Assistants. *Journal of the Association for Information Systems:* Vol. 7: Iss. 12, Article 34.

Al-Natour, S., Benbasat, I., and Cenfetelli, R. (2008). The effects of process and outcome similarity on users' evaluations of decision aids. *Decision Sciences*, 39(2), 175-211.

Atoyan, H., Duquet, J., and Robert, J. (2006). Trust in new decision aid systems, Proceedings of the 18th International Conference of the Association Francophone d'Interaction Homme-Machine, p.115-122, April 18-21, 2006, Montreal, Canada

Azoff, M.E. (1994). *Neural Network Time Series Forecasting of Financial Markets*. Chichester: Wiley.

Basu, S., and Bundick, B. (2011). Uncertainty Shocks in a Model of Effective Demand. *Boston College* WP, 774.

Bates, J. (1994). The Role of Emotion in Believable Agents, Communications of the ACM, 37,122-125.

Bell, D.E. (1982). Regret in decision making under uncertainty, Operations Research 30 (5) 961-981.

Benbasat, I., and Nault, B.R. (1990). An evaluation of empirical research in managerial support systems, Decision Support Systems 6 203–226.

Berry, D. C., and Broadbent, D. E. (1987). Expert Systems and the Man-machine Interface. PartTwo: The User Interface, *ExpertSystems* (4:1), pp. 18-27.

Bhandari, G., Hassanein, K., and Deaves, R. (2008). Debiasing investors with decision support systems: An experimental investigation. *Decision Support Systems*. 46(1), 399-410.

Bloom, N., Floetotto, M., Jaimovich, N., Saporta-Eksten, I., and Terry, S. J. (2012). Really uncertain business cycles (No. w18245). National Bureau of Economic Research.

Bowers, C. A., Oser, R. A., Salas, E., and Cannon-Bowers, J. A. (1996). Team performance in automated systems. In R. Parasuraman, and M. Mouloua, (Eds.), *Automation and Human performance: Theory and Applications*. (pp. 243-263). Mahwah, NJ: Lawrence Erlbaum Associates.

Brase, C.H., and Pellillo, C. (2012). *Understandable Statistics: Concepts and Methods*. Boston, MA: Brooks/Cole, Cengage Learning. Print.

Brynjolfsson, E., and McAfee, A. (2011). *Race against the machine: How the digital revolution is accelerating innovation, driving productivity, and irreversibly transforming employment and the economy*. Amazon Digital Services.

Burns, R., and Burns, R. (2008). *Business Research Methods and Statistics Using SPSS*. Sage Publications LTD.

Busch, P., and Wilson, D. (1976), An Experimental Analysis of a Salesmans Expert and Referent Bases of Social Power in the Buyer-Seller Dyad. *Journal of Marketing Research*, 1 (February), 3-11.

Byrne, D., Griffitt, W., and Stefaniak, D. (1967) Attraction and similarity of personality characteristics. *J. of Personality and Social Psychology* (5), 82-90.

Camerer, F., Loewenstein, G., and Prelec, D. (2004). Neuroeconomics: Why economics needs brains. *Scandinavian Journal of Economics*, Vol. 106, no. 3, 555-79.

Carruth, A., Dickerson, A., and Henley, A. (2002). What do we know about investment under uncertainty? *Journal of Economic Surveys*, *14*(2), 119-154.

Chen, J., and Lin, S. (2010). An interactive neural network-based approach for solving multiple criteria decision-making problems. *Decision Support Systems*. pp. 1–10, Aug.

Cleary, T.A. (1968). Test bias: Prediction of grades of Negro and white students in integrated colleges. *Journal of Educational Measurement*, 5: 115-124

Cohen, J., and Cohen, P. (1983). *Applied multiple regression/correlation analysis for the behavioral sciences, 2nd ed*. Hillsdale, NJ: Erlbaum

Conover, W. (1999). *Practical Nonparametric Statistics*. Third Edition ed. Wiley, New York, NY USA. pp. 388–395.

Conover, W.J., and Iman, R.L. (1976). On Some Alternative Procedures Using Ranks for the Analysis of Experimental Designs. *Communications in Statistics*. A5, 14, 1349- 1368.

Dahr, V., and Stein, R. (1997). *Intelligent Decision Support Methods - The Science of Knowledge Work*. Prentice-Hall.

Dasgupta, P. (1988). Trust as a commodity. In D. Gambetta (Ed.), Trust (pp. 49–72). New York: Basil Blackwell.

Davis, F. (1989). Perceived Usefulness, Perceived Ease of Use, and User Acceptance of Information Technology. *MIS Quarterly* (13:3) pp. 319-340.

Davis, F. D., and Venkatesh, V. (2004). Toward pre-prototype user acceptance testing of new information systems: implications for software project management. *IEEE Transactions on Engineering Management*, 51(1), 31–36.

DeLone, W.H., and McLean, E.R. (1992). Information systems success: the quest for the dependent variable, *Information Systems Research*. 3 (1) (1992) 60–95.

Dhaliwal, J. S., and Benbasat, I. (1996). The use and effects of knowledge-based system explanations: Theoretical foundations and a framework for empirical evaluation. *Information Systems Research, 7,* 342–362.

Einhorn, J., and Hogart, R.M. (1981). Behavioral Decision Theory: Processes of Judgment and Choice. *Journal of Accounting Research* Vol. 19, No. 1 (Spring), pp. 1-31.

Einhorn, J.. and Hogarth, R.M. (1978). Confidence in judgment: persistence of the illusion of validity. *Psych. Rev*. **85** 395–416.

Ford, F.N. (1985). Decision Support Systems and Expert Systems: A Comparison, Information and Management, pp. 21-26.

Frey, D., Schulz-Hardt, S., and Stahlberg, D. (1996). Information seeking among individuals and groups and possible consequences for decision making in business and politics. In E. Witte and J.

Friedrich, R.J. (1982). In defense of multiplicative terms in multiple regression equations. *American Journal of Political Science*, 26: 797-833.

Gefen, D., Karahanna,E., and D.W. Straub (2003). Trust and TAM in Online Shopping: An Integrated Model. *MIS Quarterly*, (27)1, pp. 51-90.

Gelderman, M. (1995). Factors affecting the success of management support systems: analysis and meta-analysis. *Proceedings of the IS/MAS forum of the Annual Conference of the American Accounting Association*. Orlando, FL.

Gelderman, M. (1998). The relation between user satisfaction, usage of information systems and performance. *Information and Management* 34 (1) (1998) 11–18.

Gorry, G.A., and Scott-Morton, M.S. (1971). A framework for management information systems, *Sloan Management Review* 13(1)

Greenhouse. S. W., and Geisser, S. (1959). On methods in the analysis of profile data. *Psychometrika*, 24, 95-112.

Gregor, S., and Benbasat, I. (1999). Explanations from intelligent systems: theoretical foundations and implications for practice, *MIS Quarterly*, v.23 n.4, p.497-530.

Gregor, S., and Benbasat, I. T. (1999). Explanations from Intelligent Systems: Theoretical Foundations and Implications for Practice. *MIS quarterly*, 23(4), 497–530.

Guerlain, S., Brown, D., and Mastrangelo, C. (2000). Intelligent Decision Support Systems. *Proceedings of the IEEE Conference on Systems, Man, and Cybernetics*. Nashville, USA.

Gupta, N., Forgionne, G., and Mora, M. (2006). *Intelligent Decision-Making Support Systems Foundations, Applications and Challenges*. Springer-Verlag, Germany.

Haas, R.G. (1981). The Effects of Source Characteristics on Cognitive Responses and Persuasion. *Cognitive Responses in Persuasion*. Hillsdale, NJ: Lawrence Erlbaum, 141-72.

Hagan, M., Demuth,H., and Beale, M.H. (1986). *Neural Network Design*. PWS Publishing Company, Boston, MA.

Harrison, G., and Rutstrom, E. (2008). Risk aversion in the laboratory. *Research in Experimental Economics*.12:41–196.

Haubl, G., and Trifts, V. (2000). Consumer decision making in online shopping environments: The effects of interactive decision aids. *Marketing Science*, 19(1), 4–21.

Hawley, D., Johnson, J., and Raina, D. (1990). Artificial Neural Systems: A New Tool for Financial Decision-Making. *Financial Analysts Journal*. 46(6), 63–72.

Hayes-Roth, F.(1997). Artificial Intelligence: What Works and What Doesn't. *Al Magazine*, Summer 1997, pp. 99-113.

Hayes-Roth, F., and Jacobstein, N.(1994). The State of Knowledge-based Systems. Communications of the ACM (37:3), 1994, pp. 27-39.

Haykin, S.S. (1999). *Neural Networks: A Comprehensive Foundation, 2nd ed*. Upper Saddle, NJ.: Prentice Hall.

Hecht-Nielsen, R. (1989). Theory of the backpropagation neural networks. *Proc. Int. Joint Conf. Neural Networks*, pp. 593–611.

Hess, T. J., M. A. Fuller, and J. Mathew. (2006). Involvement and Decision-Making Performance with a Decision Aid: The Influence of Social Multimedia, Gender, and Playfulness. *Journal of Management Information Systems* (22) 3, pp. 15-54.

Hoch, S.J., and Schkade, D.A. (1996). A psychological approach to decision support systems. *Management Science*. **42** (1) 51–64.

Holtzman, S. (1989). *Intelligent Decision Systems*. Addison Wesley.

Howell, D. C. (2007). *Statistical methods for psychology (6th ed.)*. Belmont, CA: Thomson Wadsworth.

Hunga, S. Y., Kub, Y. C., Liangb, T. P. and Leea, C. J. (2005). Regret Avoidance as a Measure of DSS Success.

Huse, E. (1980). *Organizational Development and Change*. West Publishing Co., New York.

Jackson, P. (1999). *Introduction to Expert Systems*. 3rd edition, Addison-Wesley

Jones, P., and Mitchell, C. (1995). Human-Computer Cooperative Problem Solving: Theory, Design, and Evaluation of an Intelligent Associate System. *IEEE Transactions on Systems, Man, and Cybernetics* 25(7): 1039–1053.

Kaastra, I., and Boyd, M. (1996). Designing a neural network for forecasting financial and economic time series," *Neurocomputing, vol. 10*, pp. 215–236, 1996.

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

Kahneman, D., Slovic, P., and Tversky, A. (1982). *Judgment Under Uncertainty: Heuristics and Biases*. Cambridge, UK: Cambridge University Press.

Kanungo, S., Sharma, S., and Jain, P.K. (2001). Evaluation of a decision support system for credit management decisions. *Decision Support Systems*. 30. 419–436.

Kaufman, P.J. (1998). *Trading Systems and Methods*. Wiley Trading Advantage. New York: Wiley.

Keen, P. (1981). Value Analysis: Justifying Decision Support Systems, *MIS Quarterly*, 5 March, 1-15.

Keen, P.G.W., and Scott-Morton, M.S. (1978). *Decision Support Systems: An Organizational Perspective*. Addison-Wesley Publishing.

Keren, G., and Bruine de Bruin, W. (2003). On the assessment of decision quality: considerations regarding utility, conflict and accountability. In: Hardman, D.J., Macchi, L. (Eds.). *Thinking: Psychological Perspectives on Reasoning, Judgment and Decision-making*. John Wiley and Sons, New York, pp. 347–363.

Kollock, R. (1994). The emergence of exchange structures: An experimental study of uncertainty, commitment, and trust. *American Journal of Sociology*, 100, 313–345.

Komiak, S. Y. X., and Benbasat, I. (2004). Understanding Customer Trust in Agent-Mediated Electronic Commerce, Web-Mediated Electronic Commerce, and Traditional Commerce. *Information Technology and Management* (5:1-2). pp. 181-207.

Komiak, S.Y. X., and Benbasat, I. (2006). The Effects of Personalization and Familiarity on Trust and Adoption of Recommendation Agents. *MIS Quarterly*. (30: 4).

Kroll, Y., and Levy, H. (1992). Further tests of the separation theorem and the capital asset pricing model. *American Economic Review*. 82:664–670.

Kroll, Y., Levy, H., and Rapoport, A. (1988). Experimental tests of the separation theorem and the capital asset pricing model. *American Economic Review*. 78:500–519.

LeDoux, J.E. (1996). *The Emotional Brain: The Mysterious Underpinnings of Emotional Life*. New York: Simon and Schuster.

Levi, K. R., Shalin,V., and Perschbacher, D. (1990). Learning plans for an intelligent assistant by observing user behavior. *International Journal of Man-Machine Studies*, 33 (5):489-503, 1990.

Lewis, D. J., and Weigert, A. (1985). Trust as a Social Reality. *Social Forces* (63:4), June. pp. 967-985.

Linger, H., and Burstein, F. (1997). Intelligent decision support in the context of the modern organisation. In: *Proceedings of the 4th conference of the international society for decision support systems* – ISDSS'97, Lausanne, Switzerland, July 21–22.

Liu H., Shah S., and Jiang W. (2004). On-line outlier detection and data cleaning. *Computers and Chemical Engineering*, 28, 1635–1647.

Lo, A. (2005). Reconciling Efficient Markets with Behavioral Finance: The Adaptive Markets Hypothesis, *Journal of Investment Consulting* 7, 21–44.

Madhavan, P., and Wiegmann, D. A. (2004). A new look at the dynamics of human-automation trust: is trust in humans comparable to trust in machines? *Proceedings of the Human Factors and Ergonomics Society*. 48th Annual Meeting. Santa Monica, CA: Human Factors and Ergonomics Society.

Madsen, M., and Gregor, S. (2000). Measuring human-computer trust. *In Proceedings of 11th Australasian Conference on Information Systems* (December 6-8), Brisbane.

Malakooti, B. (1993). A decision support system for discrete multiple-criteria problems: certainty, uncertainty, and hierarchical. *Applied Mathematics and Computation*, *54*(2-3), 131-166.

Markowitz, H. (1952). Portfolio selection. *Journal of Finance*: 77-91.

Mathews, H. Lee, D., Wilson, T. and Monoky, J.F. (1972). Bargaining Behavior in a Buyer-Seller Dyad, *Journal of Marketing Research*, 9 (February), 103-5.

Mayer, R. C., Davis, J. H., and Schoorman, F. D. (1995), An intergrative model of organizational trust. *Academy of Management Review*, 20, 709–734.

McKnight, D.H., Choudhury, V. and Kacmar, C. (2002). Developing and Validating Trust Measures for e-Commerce: An Integrative Typology. *Information Systems Research*, (13)3, pp. 334-359.

Miller, C. (1984). The Role of Performance Related Similarity in Social Comparison of Abilities: A Test of the Related Attributes Hypotheses, *Journal of Experimental Social Psychology*, 18 (November), 513-24.

Moon, Y. (2000). Intimate Exchanges: Using Computers to Elicit Self-Disclosure from Consumers. *Journal of Consumer Research*, 26 (March), 323-39.

Moreau, E. M. F. (2006). The impact of intelligent decision support systems on intellectual task success: An empirical investigation. Decision Support Systems, 42(2), 593–607.

Morry, M. (2005). Relationship satisfaction as a predictor of similarity ratings: A test of the attraction–similarity hypothesis, Journal of Social and Personal Relationships, (22) 4, pp. 561–84.

Morton, A.S, (1978). *Decision Support Systems: An Organizational Perspective*. Addison-Wesley, Reading, MA.

Muir, B. (1987). Trust between humans and machines, and the design of decision aids. *International Journal of Man-Machine Studies*.

Muir, B.M., and N. Moray (1996). Trust in Automation: Part II. Experimental Studies of Trust and Human Intervention in a Process Control Simulation. Ergonomics, (39)3, pp. 429-460.

Mulken, S., Andre, E., and Muller, J. (1999). An Empirical Study on the Trustworthiness of Life-Like Interface Agents. Saarbrucken, Germany: DFKI.

Nass ,C.I., Steuer, J.S., and Tauber, E. (1994). Computers are social actors . *Paper presented to CHI'94 conference of the ACM /SIGCHI*. Boston, MA, USA .

Nass ,C.I., Steuer, J.S., Tauber, E., and Reeder, H. (1993). Anthropomorphism, agency, and ethopoeia: computers as social actors. *Proceedings of the International CHI Conference*. Amsterdam, The Netherlands.

Nass, C., Moon, Y., Fogg, B. J., Reeves, B., and Dryer, C. (1995). Can computer personalities be human personalities. *International Journal of Human-Computer Studies*, 43:223–239.

Nass, C.L., Fogg., B. J., and Moon, Y. (1996). Can computers be teammates? International Journal of Human-Computer Studies, 45, 669–678.

Negroponte, N. (1970). *The Architecture Machine: Towards a more Human Environment*. Cambridge, MA: MIT Press.

Palma-dos-Reis, A. and Zahedi, F. M. (1999). Designing Personalized Intelligent Decision Support Systems: The Impact of Risk Attitude on Using Security Selection Models. *Decision Support Systems*. 26 (1), 31-47.

Payne, J., Bettman, J., and Johnson, E. (1993). *The Adaptive Decision Maker*. New York: Cambridge University Press.

Payne, J.W, Bettman, J., and Johnson, E. (1993). *The Adaptive Decision Maker*. Cambridge University Press, Cambridge, UK.

Pedro A C Sousa, J. P. P. R. A. R. (2007). Operational Risk Management: How an IDSS may help. *Journal of Decision Systems*, 16, 197–212.

Pellissier, M., and Fusari, A. (2007). Volatility as an Indicator of Uncertainty and its Impact on the Realizations of Industrial Business Expectations. Third Joint European Commission - OECD Workshop on International Development of Business and Consumer Tendency Surveys.

Phillips-Wren, G., Jain, L. (2005). *Intelligent Decision Support Systems in Agent-Mediated Environments*. IOS Press, Amsterdam.

Phillips-Wren, G., Mora, M., Forgionne, G.A., and Gupta, J.N.D. (2009). An integrative evaluation framework for intelligent decision support systems, *European Journal of Operational Research*, 195, 3, 642-652.

Pomerol, J.-Ch. (1995). Artificial intelligence and human decision making, *Proc. of Euro XIV*, Jerusalem, R. Slowinski Ed., 169-196.

Qiu, L., and Benbasat, I. (2004). The Effects of Text-To-Speech Voice and 3D Avatars on Consumer Trust in the Design of Live Help Interface of Electronic Commerce. *AMCIS 2004 Proceeding*s. Paper 387.

Qiu, L., and Benbasat, I. (2009). Evaluating Anthropomorphic product recommendation agents: a social relationship perspective to designing information systems, *Journal of Management Information Systems* (25:4), pp 145-181.

Reeves, B., and Nass, C. (1996). *The Media Equation : How People Treat Computers, Television, and New Media Like Real People and Places*. Cambridge: Cambridge University Press.

Roth, E M, Bennett, K B., and Woods, D D. (1987). Human interaction with an "intelligent" machine. *INT. J. MAN MACH. STUD*. Vol. 27, no. 5-6, pp. 479-525.

Roth, E. M., Patterson, E.S. and Mumaw, R. J. (2002). Cognitive Engineering: Issues in User-Centered System Design. In J. J. Marciniak (Ed.), *Encyclopedia of Software Engineering,* 2nd Edition. New York: Wiley- Interscience, John Wiley and Sons.

Rubin, D. B. (1974). Estimating causal effects of treatments in randomized and non-randomized studies. *Journal of Educational Psychology*, 66, 688-701.

Rumelhart, D. E., Hinton, G. E., and Williams, R. J. (1986). Learning the internal representations by error propagation. *Parallel distributed processing vol 1 and 2,* Massachusetts: MIT Press.

Sainfort, F.C., Gustafson, D.H., Bosworth, K., and Hawkins, R.P. (1990). Decision support systems effectiveness: conceptual framework and empirical evaluation, *Organizational Behavior and Human Decision Processes* 45 (2) (1990) 232–252.

Savikhin, A., Lam, H. C., Fisher, B., and Ebert, D. S. (2011). An Experimental Study of Financial Portfolio Selection with Visual Analytics for Decision Support. *In System Sciences (HICSS)*. 2011 44th Hawaii International Conference (pp. 1-10). IEEE.

Scott Morton, M.S. (1978). *Decision Support Systems: An Organizational Perspective*, Addison-Wesley, Reading, MA.

Seddon, P.B. (1997). A respecification and extension of the Delone and Mclean model of IS success, *Information Systems Research* 8 (3) 240 – 253.

Sharda, R., Barr, S.H., and McDonnell, J.C. (1988). Decision support system effectiveness: a review and an empirical test, Management Science 34 (2) 139–159.

Sharpe, W.F. (1992). Asset allocation: Management Style and Performance Measurement. *Joumal of Portfolio Management*, Winter, 7-19.

Sheridan, T.B. (1980). Computer control and Human alienation. Technology Review, 83, 60-76.

Sheridan, Z. B. and Hennessy, R. T. (1984). Research and modeling of supervisory control behavior: Report of a workshop. Washington, DC: National Academy Press.

Shim J.P., et al. (2002). Past, present and future of decision support technology. *Decision Support Systems*. Vol 33, pp111-126.

Shortliffe, E. H. (1976). *Computer-based Medical Consultations: MYCIN*. Elsevier Computer Science Library, New York.

Simon, H. A. (1955). A Behavioral Model of Rational Choice. *Quarterly Journal of Economics,* (69). 1955, pp. 99-118.

Simon, H. A. (1956). Rational Choice and the Structure of the Environment. *Psychological Review,* (63). 1956, pp. 129-138.

Simon, H.A. (1997). *Administrative behaviour: a study of Decision-Making process in Administrative Organizations*. New York, Free Press.

Simons, H.W., Berkowitz, N., and Moyer, R.J. (1970). Similarity, Credibility and Attitude Change: A Review and a Theory. *Psychological Bulletin*, 73 (January), 1-16.

Sniezek, J. A., and Van Swol, L. M. (2001). Trust and expertise in a judge advisor system. *Organizational Behavior and Human Decision Processes*, 84, 288–307.

Sousa, P., Pimentão, JP., and Ribeiro, R.A. (2007). Operational Risk Management: How an I-DSS may help. *Journal of Decision Systems,* Vol. 16, No 2 (2007) pp 197-212

Statman, M. (2003). A century of investors. *Financial Analysts Journal* 59 (3) 52–59.

Stewart, T. R. (2000). *Uncertainty, judgment, and error in prediction. Prediction: Science, Decision Making, and the Future of Nature*. Island Press, Washington, DC, 41-57.

Tabachnick, B. G., and Fidell, L. S. (2001). *Using Multivariate Statistics (4th ed.)*. Boston, MA: Allyn and Bacon.

Tabachnick, B. G., and Fidell, L. S. (2007). *Using multivariate statistics* (5th ed.). Boston: Allyn and Bacon.

Tan, C.N.W. (2001). *Artificial Neural Networks: Applications in Financial Distress Prediction and Foreign Exchange Trading*. Gold Coast, QLD: Wilberto Press.

Taylor, R.N. (1975), Psychological Determinants of Bounded for Decision Rationality Implications Making Strategies. *Decision Sciences* (6:3), July, pp. 409-429.

Todd, P. and Benbasat, I. (1992). The use of information in decision making: an experimental investigation of the impact of computer-based decision aids. *MIS Quarterly,* 16 (September) 373– 393.

Tong-Seng Quaha, T.S. and Srinivasan, B. (1999). Improving returns on stock investment through neural network selection. *Expert Systems with Applications*, 17, 295–301.

Trippi, R. R., and Turban, E. (1992). *Neural Networks in Finance and Investing: Using Artificial Intelligence to Improve Real World Performance*. McGraw-Hill, Inc..

Trippi, R., and Turban, E. (1996). Neural Networks in Finance and Investing. Burr Ridge, Il.: Irwin.

Trippi, R.R. and Lee, J.K. (1996). *Artificial Intelligence in Finance and Investing: State-of-the-Art Technologies for Securities Selection and Portfolio Management*. Chicago, IL: Irwin Professional.

Turban, E., Aronson, J. E., and Liang, T-P. (2004). *Decision Support Systems and Intelligent Systems*, 7th ed., Prentice-Hall.

Tversky, A. (1974). Assessing uncertainty. *Journal of the Royal Statistical Society. Series B (Methodological)*, 148-159.

Van Swol, L. M., and Sniezek, J. A. (2005). Factors affecting the acceptance of expert advice. *British Journal of Social Psycholog*y, 44, 443–461.

Vanstone, B., and Finnie, G. (2009). An empirical methodology for developing stockmarket trading systems using artificial neural networks. *Expert systems with applications*. 36(3), 6668-6680.

von Neumann, J., and Morgenstern, O. (1944). *Theory of games and economic behavior*. Princeton University Press, Princeton.

von Neumann, J. and Morgenstern (1947). *Theory of Games and Economic Behavior, Ed. 2*. Princeton University Press, Princeton, N.J.

Wang, H.Q. (1997). Intelligent agent assisted decision support systems: integration of knowledge discovery, knowledge analysis, and group decision support. *Expert Systems with Applications*, 12 (3) 323–335.

Wang, J., and Malakooti, B. (1992). A feedforward neural network for multiple criteria decision making. *Computers and Operations Research,* 19 (2).

Wang, W., and Benbasat, I. (2005). Trust in and adoption of online recommendation agents. *Journal of AIS*, 6(3), 72–101.

Wang, W., and Benbasat, I. (2007). Recommendation agents for electronic commerce: Effects of explanation facilities on trusting beliefs. *Journal of Management Information Systems*, 23, 217–246.

Wang,W., and Benbasat, I. (2008). Attributions of trust in decision support technologies: A study of recommendation agents for e-commerce, *Journal of Management Information Systems*, 24. pp. 249–273.

Whaley, R. (2008). Understanding VIX. Available at SSRN 1296743.

Williams G. J., Baxter R. A., He H. X., and Hawkins S., Gu L. (2002) A Comparative Study of RNN for Outlier Detection in Data Mining. IEEE International Conference on Data mining (ICDM'02), Maebashi City, Japan, CSIRO Technical Report CMIS-02/102.

Wobbrock, J.O., Findlater, L., Gergle, D., and Higgins, J.J. (2011). The Aligned Rank Transform for Nonparametric Factorial Analyses using only ANOVA Procedures. *In Proc*. CHI 2011, ACM Press. 143-146.

Wong, B. and Selvi,Y. (1998). Neural network applications in finance: A review and analysis of literature (1990–1996). *Information & Management*.

Wong, B., and Monaco, J. (1995). Expert systems applications in business: a review and analysis of the literature (1977–1993). *Information and Management,* 29. 141–152.

Wong, B., and Selvi,Y. (1998). Neural network applications in finance: A review and analysis of literature (1990–1996). *Information and Management*.

Woods, D. (1985). Cognitive technologies: The design of joint human-machine cognitive systems. *AI Magazine*.

Woods, D. (1988). Cognitive Engineering: *Human Problem Solving with Tools*.

Woods, D.D., Johannesen, L., and Potter, S.S. (1991). Human Interaction with Intelligent Systems: An Overview and Bibliography, *SIGART Bulletin*, 2(5).

Woodside, A.G., and Davenport, W. Jr. (1974). The Effect of Salesman Similarity and Expertise on Consumer Purchasing Behavior. *Journal of Marketing Research*, 11 (May), 198-202.

Xiao, S. and Benbasat, I. (2002). The Impact of Internalization and Familiarity on Trust and Adoption of Recommendation Agents. Unpublished Working Paper 02-MIS-006, University of British Columbia, Vancouver, Canada.

Yaniv, I., and Kleinberger, E. (2000). Advice taking in decision making: Egocentric discounting and reputation formation. *Organizational Behavior and Human Decision Processes*, 83, 260–281.

Ye, L. R., and Johnson, P. E. (1995). The impact of explanation facilities on user acceptance of expert systems advice. *MIS Quarterly*, 157-172.

Zhao, B., and Olivera, F. (2006). Error reporting in organizations. A*cademy of Management Review*, 31, 1012–1030.

Zimmerman, D.W., and Zumbo, B.D. (1993). Relative Power of the Wilcoxon Test, the Friedman Test and Repeated Measures ANOVA on Ranks. *Journal of Experimental Education*, 62, 75-86.

Zipf, G. K. (1949). *Human Behavior and the Principle of Least Effort: An Introduction to Human Ecology*, Addison-Wesley, Reading, MA.

Zuckers, L. (1986). Similarity Based Trust. *Research in Org. Behavio*r (8), 53-111.

APPENDIX

Figure 9: Distribution of Decision Performance

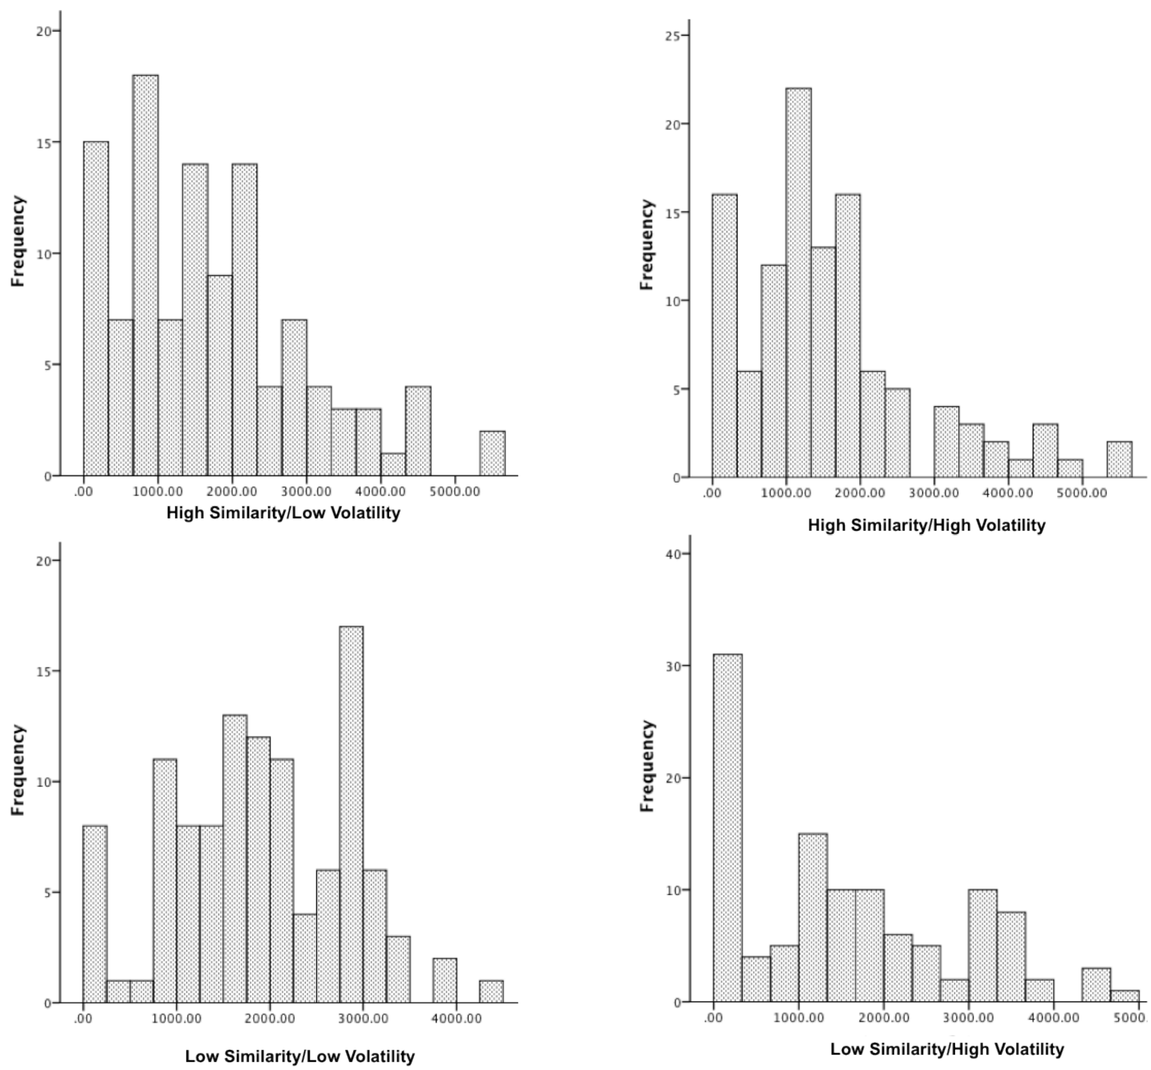

Figure 10: Decision Performance Q-Q Plot

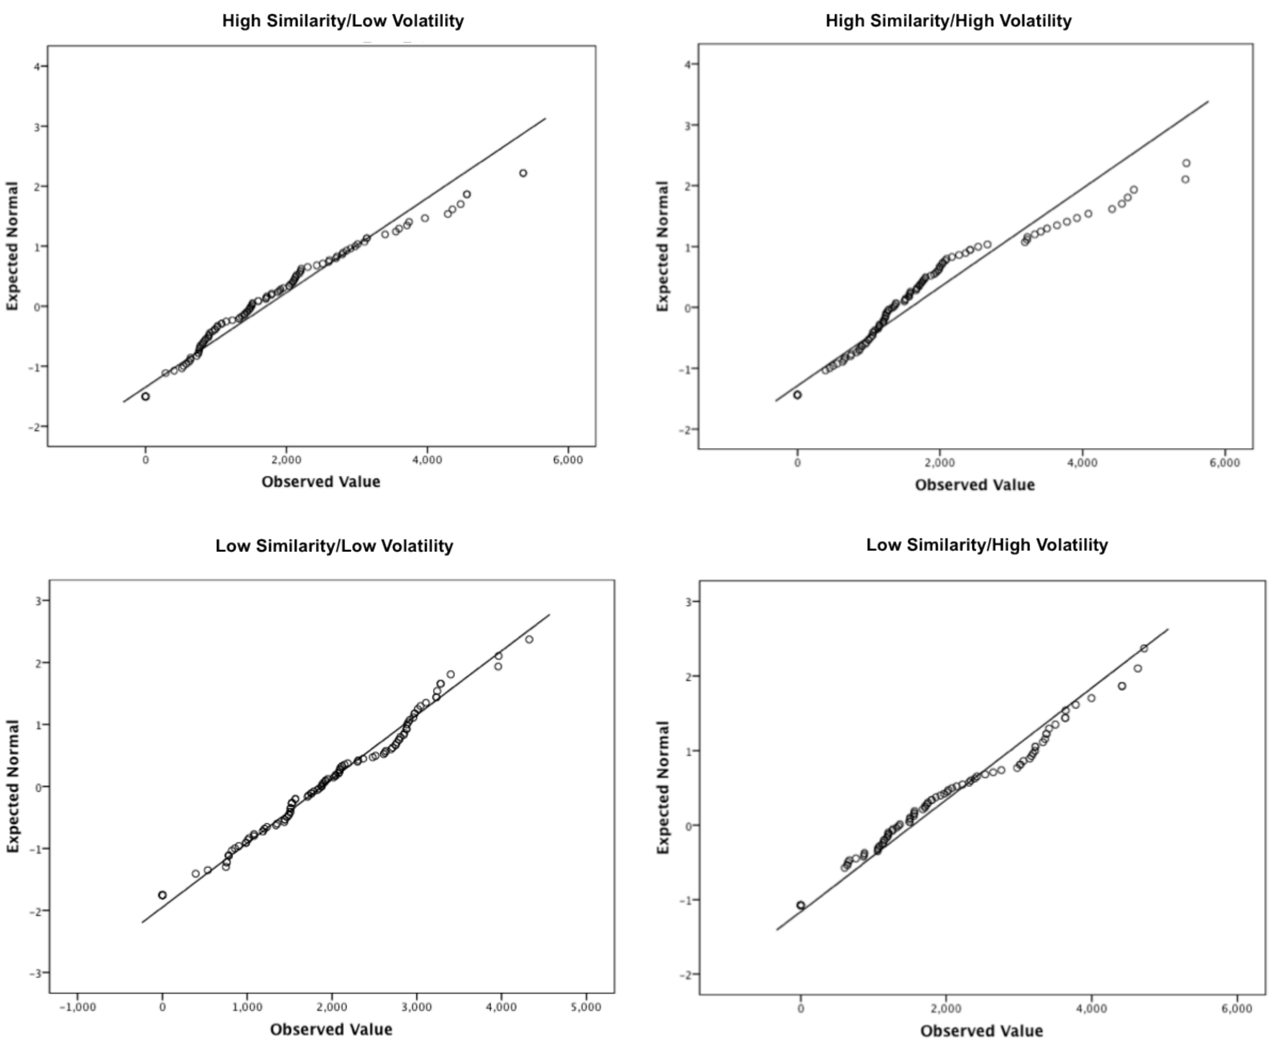

Figure 11: Decision Performance Box Plot

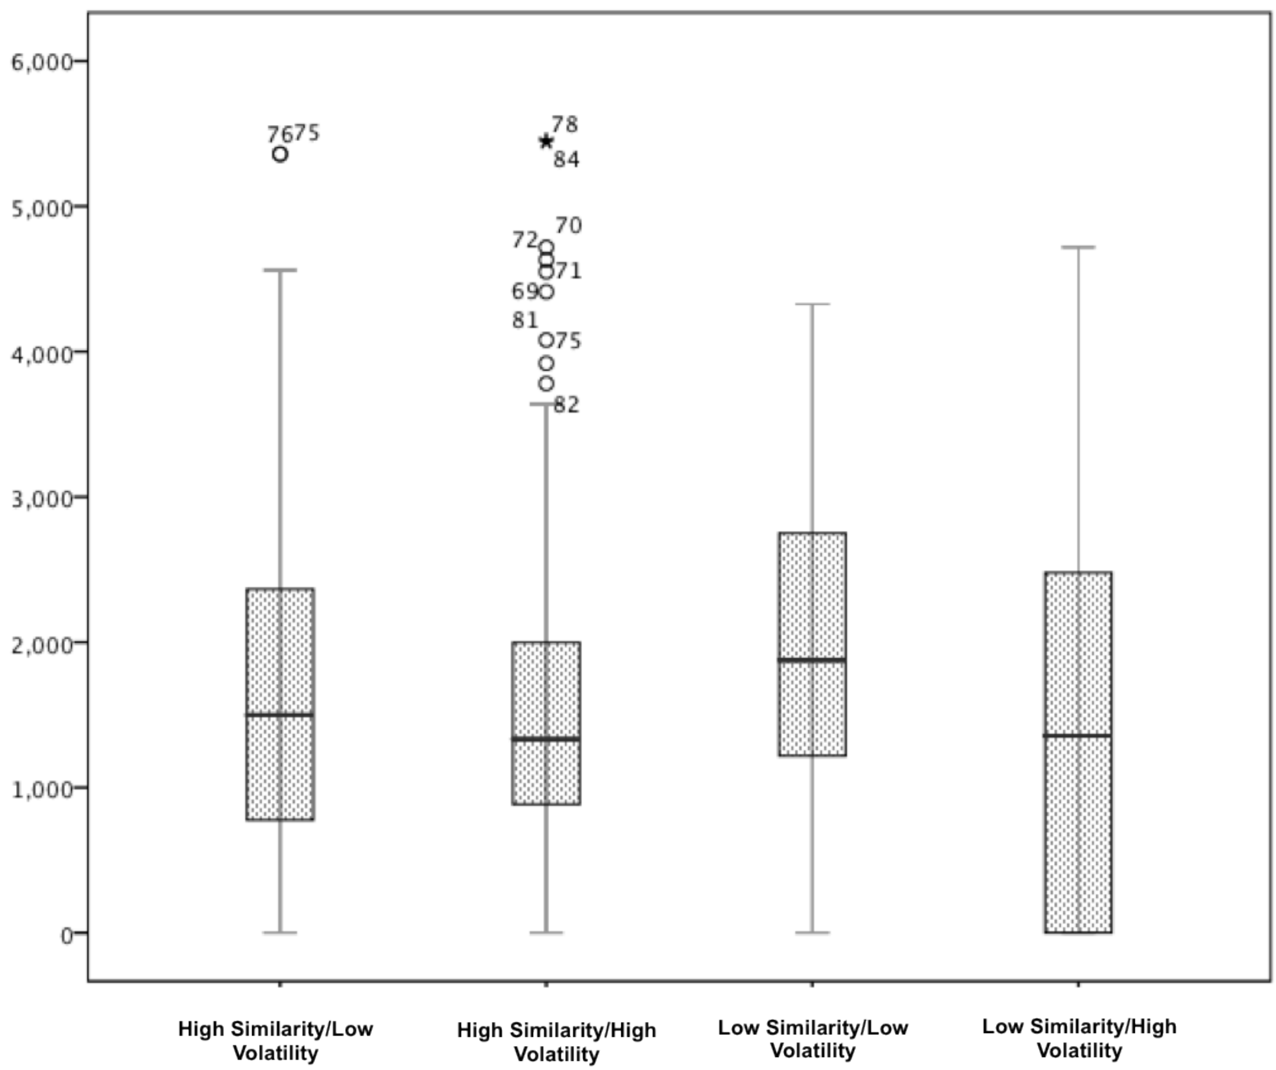

Figure 12: Decision Time Distribution Plot

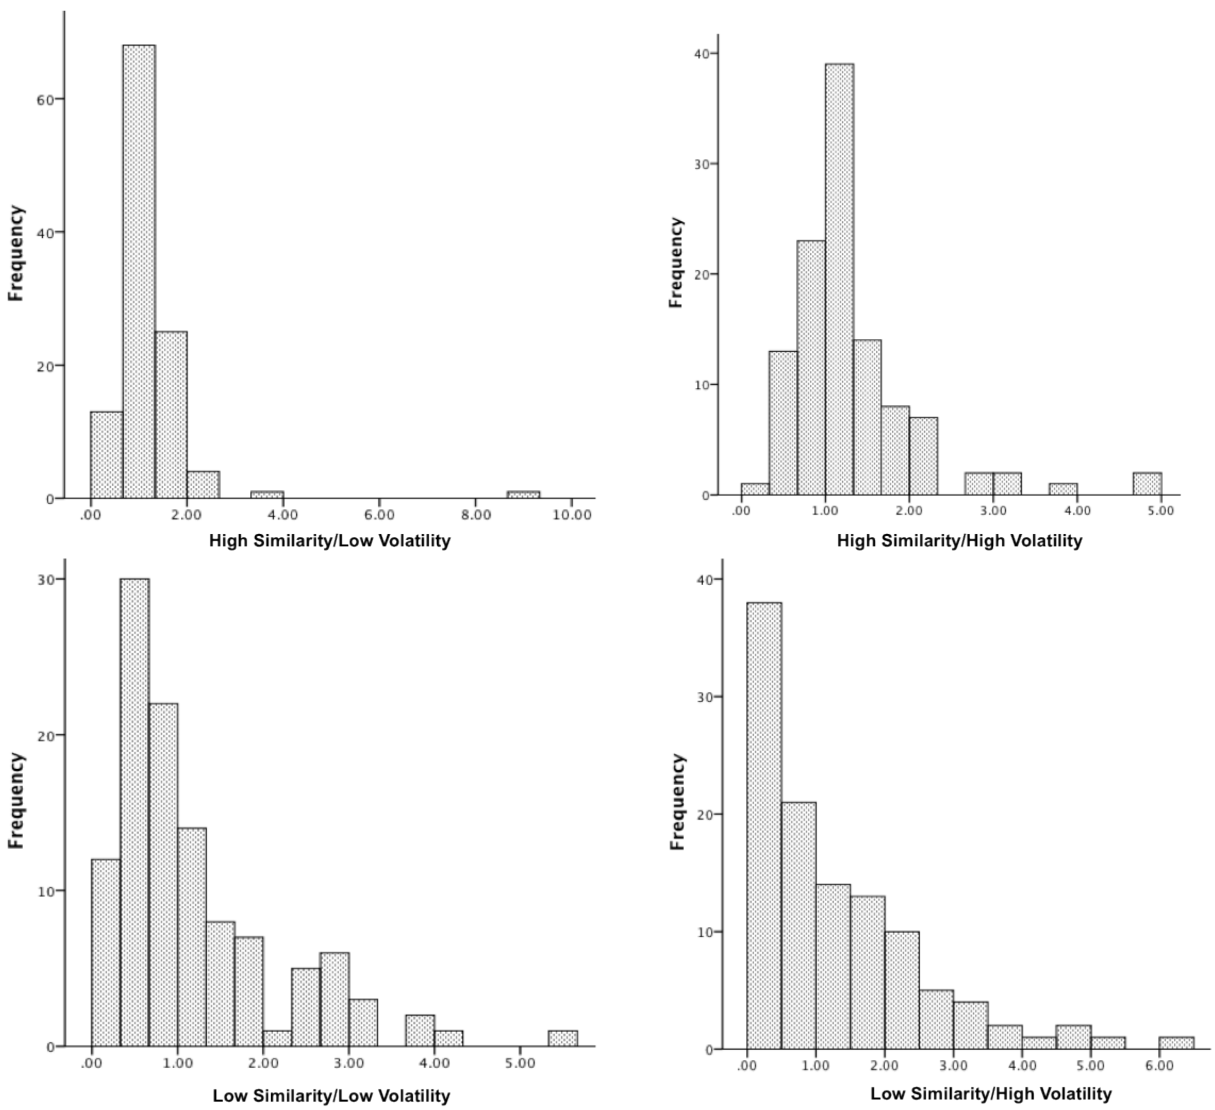

Figure 13: Decision Time Q-Q Plot

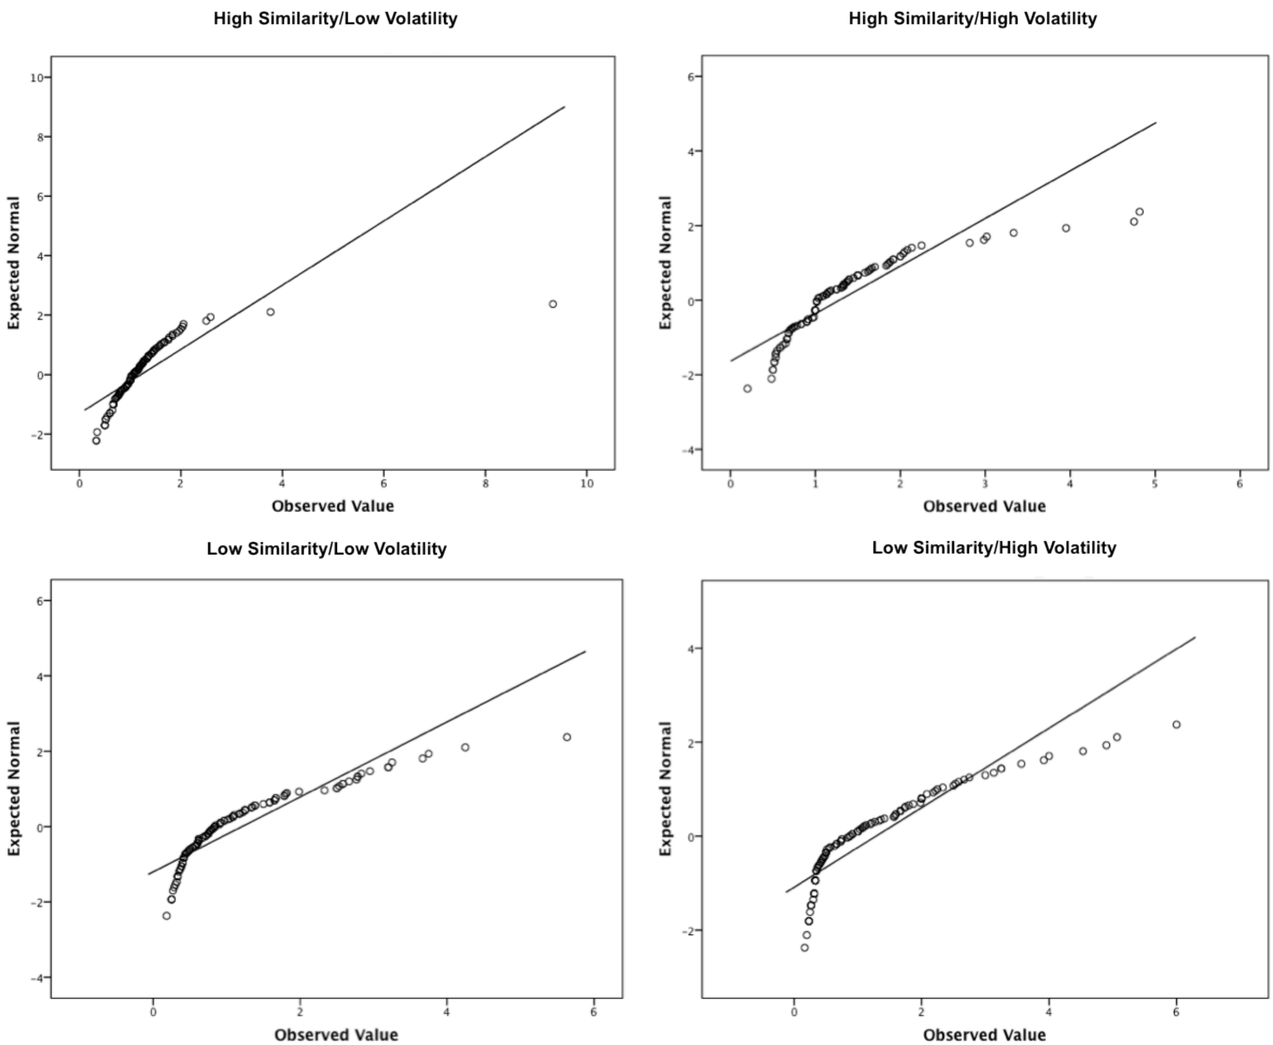

Figure 14: Decision Time Box Plot

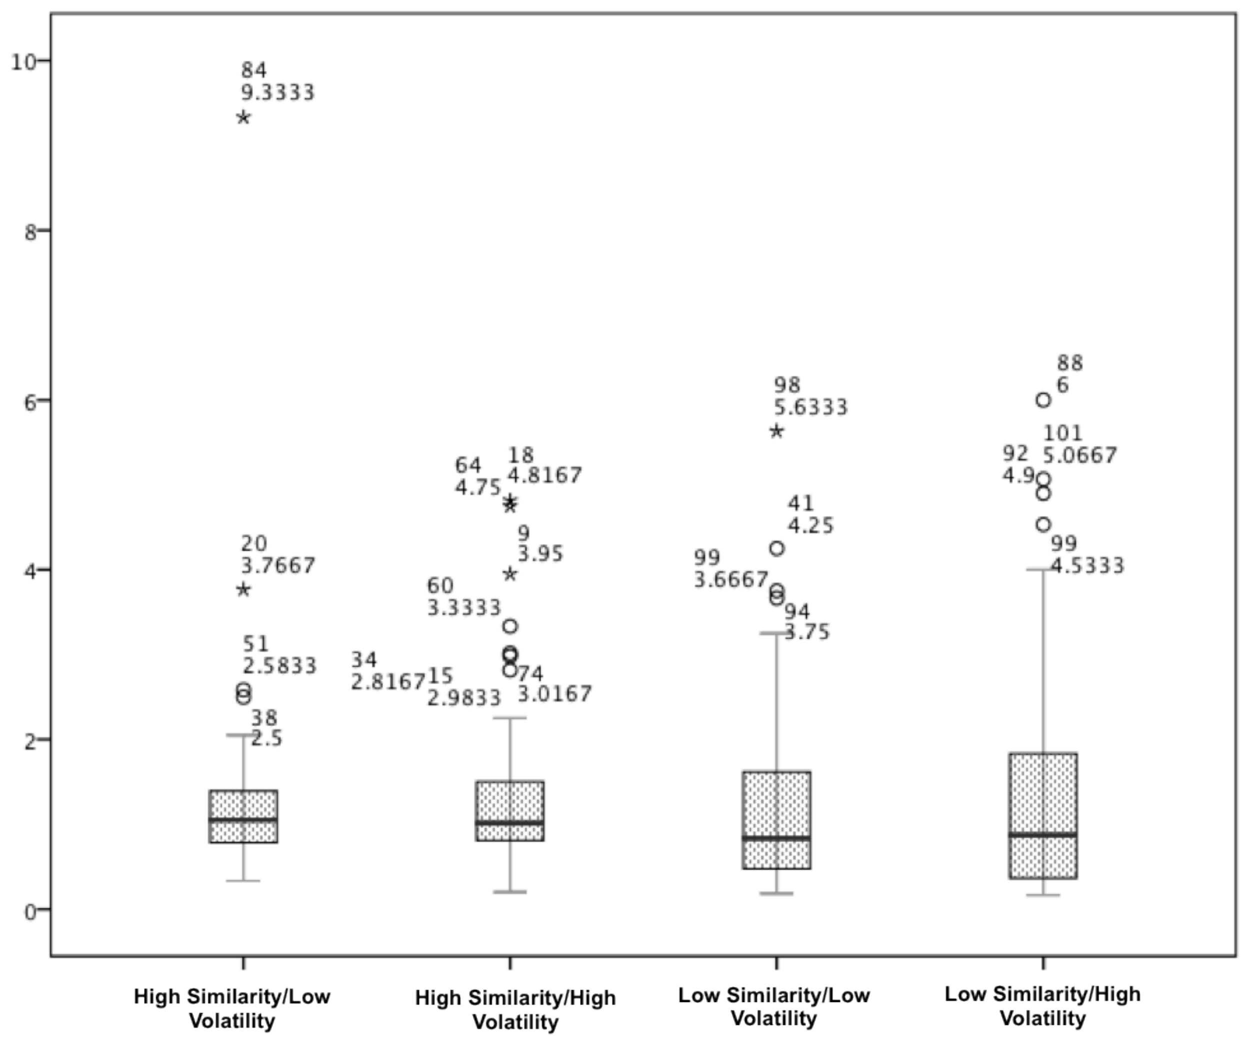

Table 9: Summary of Hypotheses

Dependent Variable	Hypothesis	Hypothesized Effect	Significance	Support
Recommendation Acceptance	H1a	The utilization of a behaviorally similar IDSS will increase the acceptance rate of machine recommendations	p<.001	Yes
Recommendation Acceptance	H1b	Uncertainty (volatility) in the decision context should strengthen the influence of a behaviorally similar IDSS on the acceptance rate of machine recommendations	n.s.	No
Decision Performance	H2a	The utilization of a behaviorally similar IDSS will positively influence trading performance	n.s.	No
Decision Performance	H2b	Uncertainty (volatility) in the decision context should strengthen the influence of a behaviorally similar IDSS on trading performance	p <.05	Yes
Decision Performance	H2c	The utilization of a behaviorally similar IDSS will decrease the standard deviation of trading performance	n.s	No
Decision Performance	H2d	Uncertainty (volatility) in the decision context should strengthen the influence of a behaviorally similar IDSS on the standard deviation of trading performance	p <.05	Yes
Decision Efficiency	H3a	The utilization of a behaviorally similar IDSS will decrease the time it takes a trader to make a trading decision	p<.01	No
Decision Efficiency	H3b	Uncertainty (volatility) in the decision context should strengthen the influence of a behaviorally similar IDSS on the time it takes a trader to make a trading decision	n.s.	No