



**NAVAL
POSTGRADUATE
SCHOOL**

MONTEREY, CALIFORNIA

THESIS

**APPLICATION OF A GENETIC ALGORITHM AND
MULTI AGENT SYSTEM TO EXPLORE EMERGENT
PATTERNS OF SOCIAL RATIONALITY AND A
DISTRESS-BASED MODEL FOR DECEIT IN THE
WORKPLACE**

by

Jacob Foster Davis

June 2008

Thesis Advisor:

John Hiles

Co-Advisor:

Steven Iatrou

Approved for public release; distribution is unlimited

THIS PAGE INTENTIONALLY LEFT BLANK

REPORT DOCUMENTATION PAGE			<i>Form Approved OMB No. 0704-0188</i>	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instruction, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188) Washington DC 20503.				
1. AGENCY USE ONLY (Leave blank)		2. REPORT DATE June 2008	3. REPORT TYPE AND DATES COVERED Master's Thesis	
4. TITLE AND SUBTITLE Application of a Genetic Algorithm and Mutli Agent System to Explore Emergent Patterns of Social Rationality and a Distress-Based Model for Deceit in the Workplace			5. FUNDING NUMBERS	
6. AUTHOR(S) Ensign Jacob Foster Davis, United States Navy			8. PERFORMING ORGANIZATION REPORT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Naval Postgraduate School Monterey, CA 93943-5000			10. SPONSORING/MONITORING AGENCY REPORT NUMBER	
9. SPONSORING /MONITORING AGENCY NAME(S) AND ADDRESS(ES) N/A			11. SUPPLEMENTARY NOTES The views expressed in this thesis are those of the author and do not reflect the official policy or position of the Department of Defense or the U.S. Government.	
12a. DISTRIBUTION / AVAILABILITY STATEMENT Approved for public release; distribution is unlimited			12b. DISTRIBUTION CODE	
13. ABSTRACT (maximum 200 words) Organizations rely on the honest operation of its members, but in an environment where individual members cannot be observed the opportunity for individuals to lie can lead to dishonest choices (Grover, 1993). This thesis created and applied a computer-based Genetic Algorithm and Multi Agent System in order to test the predictions of Dr. Steven Grover's distress-based model of the antecedents of lying in organizations. Grover's model blends self-interest theories and uses role theory to identify potential antecedents to lying. The created system provided agents that encountered situations of distress such as those described by Grover's model. The agents' actions were then observed and compared to Grover's hypothesis that an individual's skill will be inversely proportional to his frequency of lying. Social rationality has been shown to emerge in simple self-interested agents. A hypothesis that in an environment where an organization and its members are independently self-interested, the frequency of organization members lying will be inversely proportional to the magnitude of feedback provided to the organization was tested. The results support both Grover's hypothesis and the hypothesis on social rationality. Self-interest individuals with higher skills lied less than individuals with lower skills. Also, self-interested individuals lied less in the presence of a higher magnitude of negative organizational feedback.				
14. SUBJECT TERMS Deceit, Lying, Workplace, Genetic Algorithm, Multi Agent System, Simulation, Subterfuge, Evolution, Genes, Memes, Agents, Economics, Morality, Psychology, Social Rationality			15. NUMBER OF PAGES 125	
			16. PRICE CODE	
17. SECURITY CLASSIFICATION OF REPORT Unclassified	18. SECURITY CLASSIFICATION OF THIS PAGE Unclassified	19. SECURITY CLASSIFICATION OF ABSTRACT Unclassified	20. LIMITATION OF ABSTRACT UU	

THIS PAGE INTENTIONALLY LEFT BLANK

Approved for public release; distribution is unlimited

**APPLICATION OF A GENETIC ALGORITHM AND MULTI AGENT SYSTEM
TO EXPLORE EMERGENT PATTERNS OF SOCIAL RATIONALITY AND A
DISTRESS-BASED MODEL FOR DECEIT IN THE WORKPLACE**

Jacob Foster Davis
Ensign, United States Navy
B.S., United States Naval Academy, 2007

Submitted in partial fulfillment of the
requirements for the degree of

MASTER OF SCIENCE IN SYSTEMS TECHNOLOGY

from the

**NAVAL POSTGRADUATE SCHOOL
June 2008**

Author: Jacob Foster Davis

Approved by: Professor John Hiles
Thesis Advisor

Mr. Steven Iatrou
Thesis Advisor

Professor Dan Boger
Chairman, Department of Information Sciences

THIS PAGE INTENTIONALLY LEFT BLANK

ABSTRACT

Organizations rely on the honest operation of its members, but in an environment where individual members cannot be observed, the opportunity for individuals to lie can lead to dishonest choices (Grover, 1993).

This thesis created and applied a computer-based Genetic Algorithm and Multi Agent System in order to test the predictions of Dr. Steven Grover's distress-based model of the antecedents of lying in organizations. Grover's model blends self-interest theories and uses role theory to identify potential antecedents to lying. The created system provided agents that encountered situations of distress such as those described by Grover's model. The agents' actions were then observed and compared to Grover's hypothesis that an individual's skill will be inversely proportional to his frequency of lying.

Social rationality has been shown to emerge in simple self-interested agents. The hypothesis tested is that— in an environment where an organization and its members are independently self-interested — the frequency of organization members lying will be inversely proportional to the magnitude of feedback provided to the organization was tested.

The results support both Grover's hypothesis and the hypothesis on social rationality. Self-interest individuals with higher skills lied less than individuals with lower skills. Also, self-interested individuals lied less in the presence of a higher magnitude of negative organizational feedback.

THIS PAGE INTENTIONALLY LEFT BLANK

TABLE OF CONTENTS

I.	INTRODUCTION.....	1
A.	BACKGROUND	1
	1. Grover’s Model	1
	2. Social Rationality and the Golden Rule as Emergent Properties of Self-Interested Individuals.....	4
B.	PROBLEM STATEMENT	5
C.	METHODOLOGY	6
	1. Computational Modeling as a Valid Research Mechanism	7
	2. Validation of the Computational Techniques Used in this Thesis.....	8
	<i>a. Multi Agent Systems</i>	<i>8</i>
	<i>b. The Genetic Algorithm.....</i>	<i>9</i>
	3. Method of Computational Technique Integration.....	11
	4. Brief Description of the Experiments.....	12
D.	SCOPE	13
E.	ORGANIZATION OF THESIS	15
II.	THEORY BACKGROUND.....	17
A.	INTRODUCTION.....	17
B.	PSYCHOLOGY	18
	1. Lying as Commonplace	18
	2. Lying in a Western Business Culture.....	19
	3. Role Theory	19
C.	SELF-INTEREST THEORIES.....	20
	1. Agency Theory	20
	2. Ethical Ambivalence Theory.....	21
	3. Grover’s Proposed Limitations	22
D.	SOCIAL RATIONALITY.....	22
	1. Cooperation Emerging from Self-Interest.....	23
	2. Social Rationality Emerging from Self-Interest.....	23
	3. Social Rationality in Economics	24
E.	CONCLUSION	26
III.	SYSTEM DESIGN.....	29
A.	INTRODUCTION.....	29
B.	TOPVIEW	31
	1. The Year	31
	2. The Day	32
	3. The Round	32
	4. The Step	33
C.	THE MULTI AGENT SYSTEM.....	33
	1. The Widget	33
	<i>a. Widget Attributes.....</i>	<i>33</i>

b.	<i>Important Widget Methods</i>	35
2.	The Worker	38
3.	The Organization	39
4.	The Market Environment	44
D.	THE GENETIC ALGORITHM.....	46
1.	Organization Gene Pools (Organizational Internal Environment).....	46
a.	<i>Organization Gene Pool Parameters</i>	47
b.	<i>Additional Organization Design Decisions</i>	48
2.	Worker Gene Pools (Co-Worker Internal Environments)	48
E.	CONCLUSION	49
IV.	VERIFICATION, VALIDATION, AND EXPERIMENTATION.....	51
A.	INTRODUCTION.....	51
B.	VERIFICATION AND VALIDATION.....	51
1.	Characteristics of MASes.....	51
2.	Demonstration of MAS Behavior in the Computer Model of this Thesis	54
3.	Demonstration of Genetic Algorithm Behavior in the Computer Model of this Thesis.....	60
C.	EXPERIMENTATION	62
1.	Overview	62
a.	<i>Experiment One: Manipulation of the Organizations' Attribution Resolution</i>	62
b.	<i>Experiment Two: Manipulation of the Market Environment's Consumer p(detection)</i>	64
c.	<i>Experiment 3: Unaltered Run</i>	65
2.	Data Analysis.....	66
a.	<i>Results of Experiment One: Manipulation of the Organizations' Attribution Resolution</i>	66
b.	<i>Results of Experiment Two: Manipulation of the Market Environment's Consumer p(detection)</i>	71
c.	<i>Results of Experiment Three: Unaltered Run</i>	75
V.	CONCLUSIONS	81
A.	SUMMARY	81
1.	Experiment One: Manipulation of the Organizations' Attribution Resolution.....	81
2.	Experiment Two: Manipulation of the Market Environment's Consumer p(detection)	81
3.	Experiment Three: Unaltered Run	81
B.	RECOMMENDATIONS.....	82
	APPENDIX A: OTHER VERIFICATION CHARTS.....	85
	APPENDIX B: DATA GATHERED FROM EXPERIMENT TWO THAT IS APPLICABLE TO EXPERIMENT ONE	91

APPENDIX C: DATA GATHERED FROM EXPERIMENT TWO, PERMUTATION ONE	95
LIST OF REFERENCES	103
INITIAL DISTRIBUTION LIST	105

THIS PAGE INTENTIONALLY LEFT BLANK

LIST OF FIGURES

Figure 1:	Model of Individual Lying in Organizations (Grover, 1993)	2
Figure 2:	Influence of Role Strength on Lying and Victim Choice (Grover, 1993)	3
Figure 3:	A chart of data from Edmonds' and Hales' M3 model show that individuals acting out of self-interest usually behaved socially rationally (Hales & Edmonds, 2003)	24
Figure 4:	Individual decisions in an El Farol experiment reach a social equilibrium	25
Figure 5:	Development of individual strategies in an El Farol experiment (Edmonds, 1999b)	25
Figure 6:	A top-level view of the computer model	30
Figure 7:	A view of a single Organization	37
Figure 8:	Penalty and reward equations based on the attributes of the Organization policy	42
Figure 9:	Expected Worker payoffs	42
Figure 10:	Cost of inspections	45
Figure 11:	Edmonds' environment contained two states represented by the easy curve and the hard curve (Edmonds, 1999a)	52
Figure 12:	Agents adapt from the easy state to the hard state (Edmonds, 1999a)	53
Figure 13:	Agents adapt from the hard state to the easy state (Edmonds, 1999a)	53
Figure 14:	Verification permutation one (low number of Widgets) average organization profits over time	56
Figure 15:	Verification permutation two (medium number of Widgets) average organization profits over time	57
Figure 16:	Verification permutation three (high number of Widgets) average organization profits over time	57
Figure 17:	Verification permutation one (low number of Widgets) average prices of resold Widgets over time	58
Figure 18:	Verification permutation two (medium number of Widgets) average prices of resold Widgets over time	59
Figure 19:	Verification permutation three (high number of Widgets) average prices of resold Widgets over time	60
Figure 20:	Verification permutation one (low number of Widgets) average Organizational policy changes over time	61
Figure 21:	Verification permutation two (medium number of Widgets) average Organizational policy changes over time	61
Figure 22:	Verification permutation three (high number of Widgets) average Organizational policy changes over time	62
Figure 23:	Results of experiment one, number of lies told vs. Worker skill level	68
Figure 24:	Results from experiment one, number of truths told vs. Worker skill level	68
Figure 25:	Results from experiment one, proportion of truths over lies told vs. Worker skill level	69
Figure 26:	Number of times Workers lied over time, by skill level in experiment one	70

Figure 27:	Average number of times Workers lied over time (averaged by skill level) in experiment one.....	71
Figure 28:	Results from experiment two, number of lies told vs. Market Environment feedback magnitude	73
Figure 29:	Results from experiment two, number of truths told vs. Market Environment feedback magnitude	74
Figure 30:	Results from experiment two, proportion of truths told over lies told vs. Market Environment feedback magnitude.....	75
Figure 31:	Results of experiment three, number of lies told vs. Worker skill level.....	77
Figure 32:	Results of experiment three, number of truths told vs. Worker skill level	77
Figure 33:	Results from experiment three, proportion of truths over lies told vs. Worker skill level.....	78
Figure 34:	Number of times Workers lied over time, by skill level in experiment three..	79
Figure 35:	Average number of times Workers lied over time (averaged by skill level) in experiment three	80
Figure 36:	Verification permutation one average Organization profits by cluster over time	85
Figure 37:	Verification permutation two average organization profits by cluster over time	86
Figure 38:	Verification permutation three average organization profits by cluster over time	86
Figure 39:	Verification permutation one average Worker fitness over time.....	87
Figure 40:	Verification permutation two average Worker fitness over time.....	88
Figure 41:	Verification permutation three average Worker fitness over time.....	88
Figure 42:	Verification permutation one average consumer demand of resold Widgets over time	89
Figure 43:	Verification permutation two average consumer demand of resold Widgets over time	90
Figure 44:	Verification permutation three average consumer demand of resold Widgets over time.....	90
Figure 45:	Number of lies told vs. Worker skill level, experiment two, permutation three.....	92
Figure 46:	Number of truths told vs. Worker skill level, experiment two, permutation three.....	92
Figure 47:	Proportion of truths over lies told vs. Worker skill level, experiment two, permutation three	93

LIST OF TABLES

Table 1.	Payoff matrix for Axelrod's prisoner's dilemma	23
Table 2.	States of the Widget object	34
Table 3.	Other Widget attribute values and purpose.....	35
Table 4.	Important Widget methods	36
Table 5.	Worker Attributes	39
Table 6.	Instances of Organization monetary flow	40
Table 7.	Organization attributes.....	40
Table 8.	Organization business and policy options.....	43
Table 9.	Market Environment attributes	46
Table 10.	Verification Experiment Market Environment attribute values.....	55
Table 11.	Verification Experiment independent Market Environment attribute	55
Table 12.	Experiment one Market Environment attribute values	63
Table 13.	Experiment one manipulated Organization variable.....	64
Table 14.	Experiment two Market Environment variables	64
Table 15.	Experiment two independent Market Environment variables	65
Table 16.	Experiment three Market Environment attribute values.....	65
Table 17.	Contingency table and Chi-Squared test for experiment one	67
Table 18.	Contingency table and Chi-Squared test for experiment two	72
Table 19.	Contingency table and Chi-Squared test for experiment three	76
Table 20.	Contingency table and Chi-Squared test for experiment two, permutation three.....	91

THIS PAGE INTENTIONALLY LEFT BLANK

EXECUTIVE SUMMARY

This thesis created and applied a computer-based Genetic Algorithm and Multi Agent System in order to test the predictions of Dr. Steven Grover's distress-based model of the antecedents of lying in organizations. Grover's model blends self-interest theories and uses role theory to identify potential antecedents to lying. The created system provided agents that encountered situations of distress such as those described by Grover's model. The agents' actions were then observed and compared to Grover's hypothesis that an individual's skill will be inversely proportional to his frequency of lying.

Social rationality has been shown to emerge in simple self-interested agents, and so the hypothesis that in an environment 1) where an organization and its members are independently self-interested and 2) where the organization receives merited negative feedback from its environment and 3) where workers cannot be observed or punished for lying, the frequency of organization members lying will be inversely proportional to the magnitude of that feedback.

Three experiments were run. The first experiment tested Grover's hypothesis and collected data on the decisions of individuals of different skill levels. The second experiment varied the magnitude of negative feedback that was provided to organizations and the frequency of individuals' decisions to lie was recorded. The third experiment involved running the computer model without variations or constraints and observing individuals' skill levels and their decisions to lie.

The results support both Grover's hypothesis and the hypothesis on social rationality. Self-interest individuals with higher skills lied less than individuals with lower skills. Also, self-interested individuals lied less in the presence of a higher magnitude of negative organizational feedback.

THIS PAGE INTENTIONALLY LEFT BLANK

ACKNOWLEDGMENTS

No amount of prose can express how thankful I am for my beautiful wife Andrea. She has been a source of love, encouragement, and hope for me during our first year of marriage. She also selflessly postponed the fulfillment of her own Masters Degree by at least 18 months so that I would have the opportunity to earn mine. She is smart, lovely, and caring and I am lucky to share my life with her.

I would also like to thank my flexible and helpful professors and advisors, especially Mr. Steve Iatrou and Mr. John Hiles. Their guidance helped me shape and refine my ideas from a featureless mess of thought.

THIS PAGE INTENTIONALLY LEFT BLANK

I. INTRODUCTION

A. BACKGROUND

1. Grover's Model

Professor Steven Grover derived a model of the antecedents of deception in organizations by combining two self-interest theories and role theory. This model includes eight hypotheses to predict when individuals are likely to deceive, where deception is defined as a willful misrepresentation of the truth (Grover, 1993).

The first self-interest theory used to derive the model is agency theory. Agency theory is a theory of self-interest and is centered on the relationship of a principal and an agent. In agency theory, the principal is defined as the party that delegates work to the agent party. The agent party is defined as the party receiving and performing the work delegated by the principal (Eisenhardt, 1989). Agents acting in self-interest behave in accordance with personal goals and to external rewards and punishments are received from a principal. Agency theory also introduces information asymmetry, a situation where an agent has more information than his principal. Situations where the agent's behavior cannot be observed might allow the agent to deceive the principal without penalty (Grover, 1993).

The second self-interest theory, role theory, introduces the idea of role conflict. Role conflict is a situation where an agent holds multiple roles that may conflict, which might make satisfying these roles difficult or impossible. A role is a set of actions or behaviors that a person uses to guide their behavior in a situation.¹ Potential difficulty of satisfying conflicting roles can lead to psychological distress, causing the agent to seek ways to reduce the resultant dissonance. Dissonance reducers include: 1) choice, where an agent behaves according to a dominating role and ignores the others, 2) avoidance,

¹ Examples of roles and their respective environments include: the businessman when at work, the husband when at home or with family, and the host during a dinner party. Multiple roles can be imposed simultaneously, and this can lead to conflict.

where the agent quits or no longer participates, 3) compromise, where the agent is able to satisfy both roles via a creative solution, and 4) deception, where the agent satisfies one role and misreports to the other role sender. In the cases of choice and deception, the victim who is ignored or lied to will be the role sender² who is the weakest. (Grover, 1993).

Grover's integration of these two theories yields his proposed models of conflict and lying illustrated by Figure 1. This figure shows the development of dissonance and its subsequent reduction during role conflict. Upon role conflict, and individual's characteristics influence the psychological distress of the individual. Distress or dissonance must be reduced by the individual and can be accomplished in one of the four ways discussed previously.³

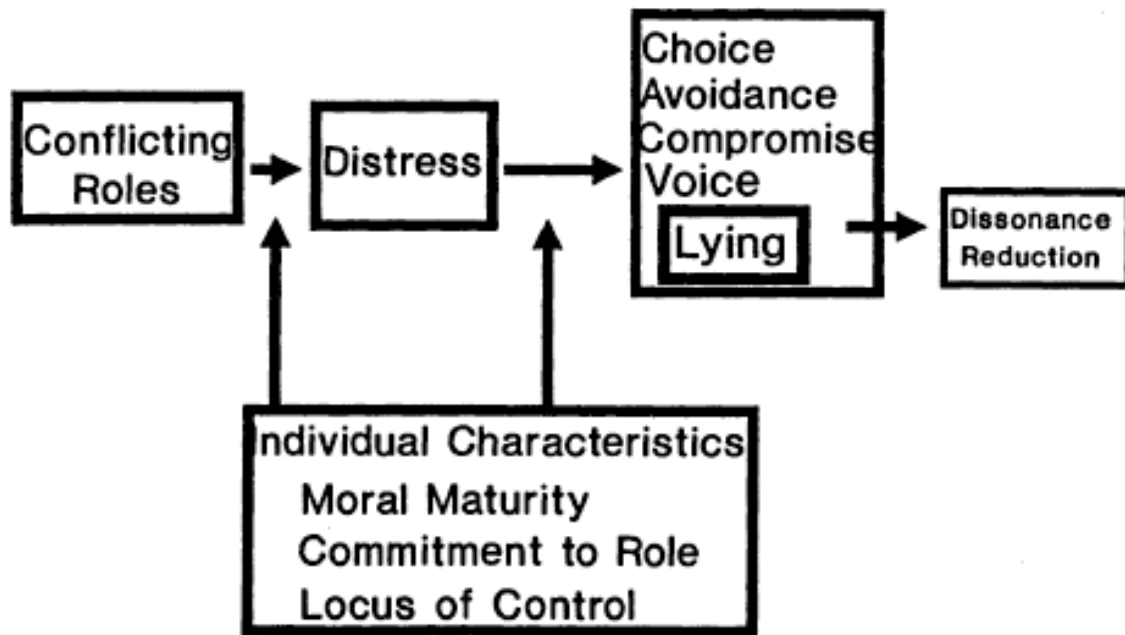


Figure 1: Model of Individual Lying in Organizations (Grover, 1993)

² A role sender is a person or other source of a role. The role sender of the businessman could be a boss or superordinate. The role sender of the husband might be culture or custom. The strength of a role sender is the magnitude to which the individual feels compelled to satisfy roles from that role sender.

³ The Voice option, which was not introduced and will not be explored in this thesis, is the ability of the individual to reduce dissonance by actually changing the requirements of roles. For example, if the businessman and the husband are conflicting because the businessman must work late, but the husband must return home to spend time with his family, the Voice option might suggest that the businessman ask his boss to change his working hours so that he can satisfy his husband role and businessman role simultaneously.

Grover further develops this model into expected outcomes based on the dissonance reducers in Figure 2. This figure shows that when one role sender is stronger than the other, “the individual is likely to choose to fulfill the stronger role at the expense of the weaker” (Grover, 1993). This figure also illustrates Grover’s hypothesis that lying will occur only when the strengths of the role senders is near equivalence. In cases where one role sender is significantly stronger than another, lying will not occur; instead, the individual will simply neglect the weaker role sender (the choice option).

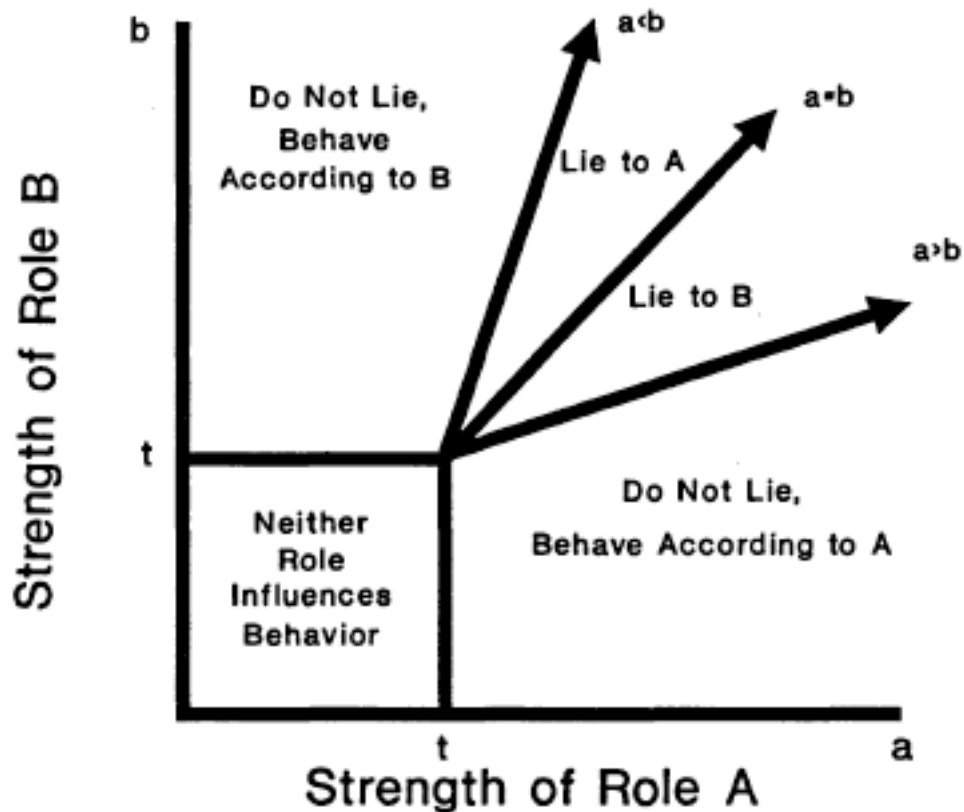


Figure 2: Influence of Role Strength on Lying and Victim Choice (Grover, 1993)

Grover proposes the following hypotheses based on his integration of self-interest and role theory:

Hypothesis 1. Greater organization-individual value incongruity will be associated with higher frequency of lying to agents of the organization...

Hypothesis 2a. Stronger organization values will be associated with less lying to agents of the organization.

Hypothesis 2b. Stronger individual values will be associated with more lying to agents of the organization...

Hypothesis 3a. Individual job skill will be negatively associated with lying to agents of the organization.

Hypothesis 3b. Time available to complete the task will be negatively associated with lying to agents of the organization.

Hypothesis 4. Higher performance expectations will be associated with higher incidence of lying to the sender of the performance expectation...

Hypothesis 5. Individuals holding more than one formal role within the organization are more likely to deceive the organization than individuals holding only one role...

Hypothesis 6. People who have role demands outside the organization are more likely to deceive the organization than those who do not hold such roles...

Hypothesis 7. The level of identification with a role will be inversely related to the likelihood of lying to the sender of that role...

Hypothesis 8. Workers reporting to more than one superior are more likely than those with a single superior to deceive a superior.

(Grover, 1993)

2. Social Rationality and the Golden Rule as Emergent Properties of Self-Interested Individuals

Social rationality is a behavior or action of an individual that results in a joint social gain of a group of individuals. Similar in concept of is the Golden Rule, which has been found to serve as a guide to an individual's behavior in many circumstances. The Golden Rule says: "Do unto others as you would have them do unto you." The Golden Rule can be found in nearly all human cultures and even in non-human species (Pfaff,

2007). Neuroscientist Donald Pfaff theorizes that the Golden Rule is hard-wired in human brains as “reciprocal altruism.” (Pfaff, 2007).

Even in an environment motivated only by the concern for one’s self, the Golden Rule can be found. Reciprocity has been shown to emerge in computer simulations that contain self-interested agents. Others have shown that social rationality can emerge from a collection of self-interested individuals. A more thorough review of this phenomenon is provided in Chapter II.

In an organization, it is not socially rational for individuals to lie. Except in a few instances,⁴ deceit within organizations is harmful to the organization itself. However, it might be individually rational for individuals to lie, depending on the policies of the organization. Since individuals acting individually rational (as opposed to socially rational) would harm the organization and cause harm to the members of the organization, the presence of individual rationality might lead to social rationality under an environment where the most fit organizations and members survive.

One of Grover’s hypotheses and the emergence of social rationality from self-interested individuals will be explored in this thesis.

B. PROBLEM STATEMENT

A computer model was built to test two hypotheses and simulated individuals who were not limited in their decision to lie or not lie. The individuals were contained within organizations that competed in a free market industry.

Hypothesis 3a of Grover’s model states that an individual’s job skill level is inversely related to the frequency with which that individual decides to lie. This was tested in the computer model of this thesis by creating workers who had constant skill levels. As the workers made decisions, their decisions to lie or not lie and skill level were recorded.

⁴ This thesis operates under the assumption that deception in the workplace is neither desirable nor helpful.

This thesis also explored the possibility of an emergent social rationality from self-interested agents, referred to as the Davis Hypothesis and defined as follows:

Davis Hypothesis: *In an environment where 1) an organization and its members are independently self-interested, 2) the organization receives merited negative feedback from its environment, and 3) workers cannot be observed or punished for lying, the frequency of organization members lying will be inversely proportional to the magnitude of that feedback.*

Simply put, the practice of deception in an organization might be an end to itself, even if workers cannot be detected or directly punished. The environment required by Grover's model allows for this possibility to be tested. The amount of negative feedback that an organizational entity would receive as a result of a lie told by a worker was varied over three permutations and the frequency of lies told by workers was recorded.

C. METHODOLOGY

Models that predict behavior and represent real world phenomenon, such as Grover's model of deception in the workplace, are well suited to be expressed in a computer model. Computer models can rapidly explore a problem space and provide the means to observe the aggregate behavior and emergences of individuals. Computer models can also produce large amounts of data quickly and inexpensively to increase statistical confidence.

Since Grover's model is one of deceit in the workplace, a model of a common workplace was created. This workplace was a model of several organizations in a competitive free market. Each organization was the home to several workers who were able to improve the value of raw materials. Testing the Davis Hypothesis required an identical model.

Three experiments were conducted. The first experiment ran the computer model with parameters that prevented workers from being monitored⁵ by the organizations to test Grover's hypothesis. The second experiment was identical to the first except that the

⁵ The fact that a worker told a lie could actually be seen by his organization, but the parameters of the computer model prevented this observation from having any effect on the organization or on the untruthful worker. Thus, as Grover's hypotheses require it, the workers were unmonitored.

organization feedback mechanism was varied to test the Davis hypothesis. The third experiment ran the model without these two constraints, which served as a comparison to a situation where behavior could be observed.

1. Computational Modeling as a Valid Research Mechanism

A model is any representation of a theory about a real-world phenomenon. Traditionally, models constructed out of a natural language have been used by social scientists, but for the past four decades, researchers in the social sciences have also been using computational modeling. Computational models are comprised of a mixture of mathematics and algorithms. While computational modeling itself does not require the use of computers, computers have proven to be effective computational modeling tools. Computer models⁶ are essentially theories rendered as computer programs (Taber & Timpone, 1996).

Formal models⁷ have advantages over natural-language models because they can be precise, have clear assumptions, are easily internally validated, can take advantage of formal deduction, and can be unambiguously communicated among scientists. The major advantage of computational models over formal mathematical models (such as a set of predictive equations, or Game Theory) is that while mathematical models easily suffer from oversimplification, computational models “greatly increase(s) the level of realism one may incorporate in a formal model without sacrificing analytic focus” (Taber & Timpone, 1996). In past social science work, computational modeling has demonstrated its merits: the ability to allow for theoretical uncertainty when necessary, the ability to easily express complex theoretical concepts, and for the ability to combine the implications of several theoretical assumptions (Taber & Timpone, 1996).

Computational modeling is especially well suited for modeling complex or “bottom up” processes even beyond the need for convergence and optimization, as is

⁶ Taber and Timpone use the term *computational model* to describe both a model comprising of mathematics and algorithms and a model written in computer instructions. This thesis uses the term *computational model* to define a model comprising of mathematics and algorithms and the term *computer model* to describe a *computational model* written as computer instructions (a computer program).

⁷ Taber and Timpone imply that computational models are formal models (Taber & Timpone, 1996).

required for empirical modeling. Additionally, “one can induce theory from empirical work, represent this theory as a computational model, and generate new predictions from the model” (Taber & Timpone, 1996).

2. Validation of the Computational Techniques Used in this Thesis

This thesis employed an integration of two approaches to computational modeling, the Multi Agent System and the Genetic Algorithm.

a. Multi Agent Systems

A Multi Agent System (MAS) is a type of complex system with interacting agents. The MAS approach to simulation is founded on the principle that groups of simple programs, or agents, that have simple rule sets can create and model complex behavior. MASes can be observed throughout the real world: the streets are navigated by individual drivers; a free market economy contains individual consumers and suppliers; organizations are made of individual people. Rush-hour traffic, the rise and fall of prices (Sowell, 2007), and lines in the company cafeteria are all complex aggregate phenomenon that occur as a result of individual decisions.

In a MAS, an agent has three types of abilities: 1) he can receive information from his environment, 2) he can process that information and translate it into an action, and 3) he can perform an action that somehow affects his environment. The agent may have multiple instances of each of these abilities.

When multiple agents are placed in an environment, their actions affect their environment and, in turn, might affect other agents. The aggregate behavior of the agents can become complex. If a single agent is in the environment, his actions are somewhat predictable because of his simple rule set. This is analogous to a single vehicle driving through a city’s streets. In this situation, the vehicle’s environment (distance to a turn, maximum speed possible without collision, etc.) is constant or can be easily calculated or represented by a mathematical model. If, however, the streets are full of other vehicles, a single vehicle’s environment is constantly changing, because it is affected by the actions of other vehicles. Even in an environment as simple as a highway

becomes so complex that approaches as sophisticated as fluid and gas dynamics, servomechanisms, and network analysis have been used in attempts to describe its aggregate behavior (Chandler, Herman, & Montroll, 1958). Such a complex environment cannot easily be modeled by mathematics alone (Taber & Timpone, 1996).

A MAS was used in this thesis because, while a single individual's actions based on self-interest might be easily determined with a mathematical model such as Game Theory, the real-world contains more than one individual. Also, Game Theory requires perfect knowledge of not only one's own payoffs, but the payoffs of one's opposition (Krippendorff, 1986). In a real-world scenario, multiple individuals are interacting simultaneously; therefore the environment is non-deterministic. At the same time, the environment is not probabilistic because individuals do not make decisions purely randomly. An environment that lies between being deterministic and probabilistic is considered to be complex, and MASes are able to model complexity.

The utilization of a MAS also allows for aggregate patterns to emerge from simple agents. It is assumed that in real-life situations agents will not be interacting with a principal in a vacuum, i.e. there will be other agents interacting with the same principal, and there will be other principals that interact with other agents. Grover's model is used to make predictions about human social behavior. A model with multiple agents and multiple principals best reflect real-world conditions. The interactions of simple agents who act in self-interest within a social system produce the emergent behaviors that might arise from such an environment. These emergent behaviors are required to test the Davis Hypothesis.

b. The Genetic Algorithm

With computer models that involve many independent variables, exhaustive simulation techniques quickly become too large with current computing resources. It is not possible to take a many-dimensional problem and test the entire

problem space in a reasonable amount of time⁸. The Genetic Algorithm (GA), a technique developed by computer scientist John Holland (and used and tested by many others), has been shown to provide a good solution to this problem.

The GA involves two steps: selection and mating. Selection is a test of an individual's fitness. If the individual fails to pass this test, it dies. The most fit survivors mate (Holland, 1992).

Mating occurs between two individuals and consists of encoding, crossover, and mutation. Each individual's attributes are encoded into binary strings of equal length. These strings are then laid in parallel, and at one or more crossover points the strings swap material. The resulting strings are then mutated. Mutation is a bit-wise function. Each bit in the string has a chance to be mutated. If a bit is mutated, it is simply flipped; a zero would become a one and vice-versa⁹. The two resultant binary strings of this process are the "children" of the original strings.

The children of all matings in a generation replace those individuals who died. The individuals who did mate survive. Individuals who survived but did not mate also survive. All of these survivors make up the population of the next generation.

This whole process of selection and mating is repeated for a number of generations.

The GA has been found in many cases¹⁰ to be a good tool for exploring problems that have many variables, and thus which's problem spaces are enormous (Holland, 1992). While the GA does not test each permutation of a problem space explicitly, each permutation has the potential to be tested. Many random instantiations give good initial coverage of the entire space, and each instantiation will converge on a

⁸ Axelrod's prisoner's dilemma had individuals whose decision table was only 70 binary characters long, but an exhaustive search for strategies in this simple individual would still take more than 100 times the current age of the universe at 100 strategies per second (Axelrod, 1997). This thesis used agents with binary strings of length 86 and 320.

⁹For example, if each bit has a 0.01 chance of being mutated, then about one in every 100 bits would be flipped.

¹⁰ According to John Holland, these cases include jet engine design and gas pipeline system control (Holland, 1992).

local maximum present in the space (Axelrod, 1997). The problem space of the model created to test Grover's hypothesis and the Davis Hypothesis was large and complex.

3. Method of Computational Technique Integration

The computer model contained one environment, 12 organizations, and 96 workers. The environment, organizations, and workers were arranged as a MAS, in which organizations would compete with each other to buy raw resources and then resell them to the environment. Each organization employed eight of the 96 workers.

The GA was used within each organization as a means of adaptation for the workers, and it was also applied to the organizations. To combat the development of superindividuals, the organizations were divided into three sections within which they competed¹¹. This reduced the influence of superindividuals.

Because a MAS and GA were both used simultaneously, there are two domains that should not be confused.

The first domain was that of the MAS. This domain will be referred to as the External Environment. There was only one External Environment. The External Environment contained all of the organizations and their respective workers. When organizations bought raw widgets and assigned them to the workers, and when the workers attempted to complete or lie about a task, the External Environment was in play. This environment provided no adaptation, but it allowed for the model to exhibit complex behaviors.

The second domain was the domain of the GA, referred to as the Internal Environment. There were 13 Internal Environments. One Internal Environment held the organizational policies of each of the 12 organizations and will be referred to as the Organizational Internal Environment. Another Internal Environment was for each organization's workers, which will be referred to as the Co-Worker Internal

¹¹ Under certain conditions, the GA will evolve superindividuals who have very high fitness, but who end up homogenizing the gene pool (Holland, 1992). While some simulations desire this result, the goal of this experiment is to evolve agents who work well together. Separating the organizations into three sections did lead to superindividuals within each section, but the organizations were then heterogeneous at the section-level.

Environment. Each of the Co-Worker Internal Environments held a working strategy for each worker. The purpose of the Organizational Internal Environment was to enable the organizations to develop organizational policies that were effective. The purpose of the Co-Worker Internal Environment was to enable co-workers to develop effective working strategies.

The External Environment and the Internal Environments did not operate in the same time domain. The time units of the External Environment were the Round, Day, and Year, and will be discussed in Chapter III. The time unit of the Internal Environment was the generation. The difference between these time domains is analogous to the way an individual might develop his strategy in the game of Chess. On each occasion of playing Chess, the individual plays with the current strategy (External Environment). After winning or losing several times¹², the individual might modify his strategy (Internal Environment), and on the occasion of the next game he would use the new strategy.

The External Environment provided input for the Internal Environment indirectly. During the events of the External Environment, each agent in the system earned a fitness score. During the events of the Internal Environment, this fitness score was used.

4. Brief Description of the Experiments

The environment provided raw “widgets” for the organizations to purchase according to an elastic demand curve. The organizations then could pass these raw widgets to their respective workers, who then had the choice to either 1) honestly attempt to complete the widget, 2) make the widget appear to be complete (deception), or 3) do nothing. Organizations could then resell completed widgets at a price of their choice, and the environment bought the processed widgets based on price and quantity according to an elastic demand curve.

¹² This is the detail of the analogy that results in two different time domains. The strategy (Internal Environment) is developed in a different domain from the actual implementation (or acting out) of the strategy (External Environment).

The fitness of each organization was determined by the organization's net monetary gain. The fitness of each worker was determined by the number of widgets that were (or appeared to have been) completed by the worker.

The first experiment tested Grover's hypothesis and involved the running of the model with parameters that prevented individuals from being punished or caught for lying. The second experiment tested the Davis Hypothesis and ran the model in identical conditions, but it also varied the amount of feedback that was provided to the organizations. The third experiment ran the model without any parameters that restricted the ability of an organization to detect and punish deception.

The attributes and actions of each agent in the simulation were recorded. To test Grover's hypotheses, the frequency of lying and the related task difficulty and individual skill levels were recorded and analyzed. The frequency of lying was also recorded for the Davis Hypothesis, and the same data was collected in the third experiment.

A Chi-Squared test was used to determine if there was any relationship between skill level and frequency of lying, and between skill level and task difficulty. The data was then plotted to show the direction and strength of the relationship. This same analysis was performed in the third experiment. In experiment two, A Chi-Squared test was used to determine if a relationship between magnitude of organization feedback and frequency of lying. This data was also plotted to show the direction and strength of the relationship.

The results of all three experiments supported Grover's hypothesis and the Davis Hypothesis. Workers with higher skill levels lied less frequently, and a higher magnitude of feedback to the organizations reduced the frequency with which workers lied.

D. SCOPE

Grover's model equates and interchangeably uses the terms lying, deception, and subterfuge (Grover, 1993). He defines lying and further extends its domain to:

“A false statement made with the intent to deceive” is the definition of *lying* given by the *Oxford English Dictionary*, and it matches the definition adopted by social psychologists (DePaulo & DePaulo, 1989;

DePaulo, Stone, & Lassiter, 1985; Ekman, 1985). The essential aspects of this definition are that the perpetrator *knows* the information is false, *wants* to mislead another person, and engages in the behavior proactively (Bok, 1978). Other than by lying, people can deceive by omitting facts or information, by presenting information in a certain order, or by framing statements to guide the listener away from the truth.

(Grover, 1996)

This thesis follows the same scope, definitions, and interchanges.

In the computer model, lying was represented by a worker's decision to untruthfully indicate that tasks assigned to him were completed. In addition to lying, the worker had the choice to either 1) complete the task¹³, or 2) not attempt to complete the task and indicate to the organization that the task was not completed. An attempt to lie about a task completion was done with 100% probability of success, whereas an attempt to truthfully complete a task might have had a lower probability of success. Decisions to do nothing were also performed with 100% success.

Grover's hypothesis 3a requires that a model of individual skill levels be present. The computer model represented individual skill level by giving each individual in the model a random set of skills to complete different types of tasks. In the model, there are three types of tasks and each individual has a skill level that corresponds to each of these types of tasks. Skill level is modeled by defining with what probability of success an individual can complete a task of the corresponding type. Higher-skilled individuals have a higher probability of success, and lower-skilled individuals have a lower probability of success. The range of probabilities of success was approximately normally distributed between zero and one, inclusive.

The scope of the virtual environment included several organizations and their workers. Throughout this thesis, the terms "task", "widget" and "product" are used interchangeably to mean a unit of work that can 1) be attempted and improved by a worker, and 2) bought, sold, and inspected by an organization.

¹³ Or make the task appear to be complete (deception).

Organizations were not allowed to interact with one another. An Organization could only interact with its own workers, and with the single environment. Interaction with the environment was limited to 1) the buying of raw materials, 2) the selling of processed products, and 3) the refunding of money earned from products that were determined to be defective by the environment. Organizations were allowed to assign tasks to workers and receive complete or incomplete tasks. Organizations were also allowed to buy the tasks from the environment, inspect them¹⁴, and sell them.

Workers could not interact with one another. Workers could only interact with their organization. Workers were only allowed to receive tasks and return complete or incomplete tasks. With regard to the tasks, workers were only allowed to make one of three choices: 1) attempt to complete the task honestly, 2) make the task appear to be completed (be dishonest), or 3) do nothing.

Tasks (or widgets) could only be handled by a single organization and a single worker. Widgets could only be assigned once, regardless of a worker's action or inaction. There was no opportunity for rework. Each widget was of one of the three types, and this type was constant.

E. ORGANIZATION OF THESIS

This thesis will provide a brief literature review of psychology and self-interest theories as they relate to Grover's model. The review closely follows Grover's literature review, but it also introduces a few topics not covered by Grover during his development of his model. The review also provides information on social rationality and the workplace.

A Chapter on system design will provide detailed specifications of the computer model that was built. This Chapter will also provide justification for the design parameters and other design decisions that were included in the computer model, as well as their relevance to Grover's model and the Davis Hypotheses.

¹⁴ Experiments one and two did not allow tasks/widgets to be inspected. Grover's and Davis' hypotheses require that workers not be monitored. Experiment three allowed organizations to discover and punish lying workers for a price.

The fourth Chapter will provide verification results for the computer model that was designed in this thesis. Experiment details and analysis of their data will also be provided.

The thesis will conclude with a summary and explanation of the results, as well as recommendations for future study.

II. THEORY BACKGROUND

A. INTRODUCTION

Grover has indicated the importance of studying deceit in the workplace:

People have ample opportunity to either lie or tell the truth in the course of their work. Workers continually report their behavior and give information to peers, superordinates, and others, in written form, orally, and nonverbally. The truck driver records the number of hours on the road, the nurse charts vital signs, the certified public accountant states what has been audited, and the forester reports a tree census. Organizations generally rely on these reports to be honest. However, each of these individuals may have reasons to lie. The truck driver returns home sooner if he says he drove the speed limit when in fact he exceeded it; the nurse may not have time to actually observe vital signs that must be recorded; CPA's may gain partner status by exaggerating their work quantity; and the forester may misreport the tree census to prevent deforestation. Dishonesty in the workplace has implications for how organizations function and therefore demands theoretical attention.

(Grover, 1993)

This Chapter will establish some of the previous research done in the field of lying as it relates to its motivations and origins. Most of the background is congruent with the background theory cited by Grover in the construction of his model. This Chapter will also provide some background information on social rationality and its relevance to the workplace and the Davis hypothesis.

B. PSYCHOLOGY

Grover begins his literature review with some background on lying and deceit as related to psychology:

Psychologists who have studied lying include Hartshorne and May (1928), who conducted a series of studies investigating lying, stealing, and cheating. "The results of these studies show that neither deceit nor its opposite, 'honesty,' are unified character traits, but rather specific functions of life situations" (Hartshorne and May 1928, p. 411). Other psychologists have examined the nonverbal cues of deceit, finding that few untrained people can discern liars from tellers of the truth (Bond, Kahler, and Paolicelli 1985; DePaulo and Pfeifer 1986; Ekman 1985; Riggio and Friedman 1983; Zuckerman, Koestner and Alton 1984). The psychology literature therefore suggests that the difficulty with which deception can be identified and the lack of a trait to explain lying provides ample opportunity for lying in organizations.

(Grover, 1993)

1. Lying as Commonplace

At least one study has shown that people lie as a part of everyday life. DePaulo, et al.'s study of 1996 examined different groups of people that recorded when they told lies. Most often people reported telling lies about "their feelings; their actions, plans, and whereabouts; and their achievements and knowledge" (DePaulo, Kirkendol, Kashy, Epstein, & Wyer, 1996).

This study also gathered data about how people felt before, during, and after lie-telling in terms of distress. The study observed that before a lie, the participant was only under a small amount of distress and that after a lie the distress was not significantly reduced. More distress was reported during the actual telling of the lie, and some of this distress lingered. However, "for more than 70% of the lies in both samples, participants said that if they could relive the situation, they would tell the lie again" and they "...claimed that both they and the targets of their lies would have felt a bit worse if the truth had been told instead of the lie"(DePaulo et al., 1996).

Lies! Lies!! Lies!!!: The Psychology of Deceit by Charles Ford addresses the paradoxical nature of lie-telling: There are many situations where lies are a social or cultural expectation, and others where lies are frowned upon. “The ability to lie, and to tell when a lie is needed, are fundamental human skills. Whereas those who lie too much are considered unreliable, those who lie too little are labeled tactless” (Price, 1996).

2. Lying in a Western Business Culture

Grover established that lie-telling in a business environment is very different than lying in everyday life. Traditionally, the modern capitalist business culture is concerned with the effects of lying and fraud on business. In a casual observation, a common database was searched for the terms “ethics” and then for “business [and] ethics.” Nearly 40% of the results in the entire “ethics” result set also contained the word “business”¹⁵. Grover also alludes to ethics as a significant concern to businesses (Grover, 1993), (Grover & Chun Hui, 1994).

The *Journal of Business Ethics* cites examples of businesses such as Enron and WorldCom, whose deceitful practices “helped to fuel a massive loss of confidence in the integrity of American business and have contributed to a very sharp decline in the U.S. stock market” (Carson, 2003).

3. Role Theory

Role theory is based on the concept that an individual belongs to several types or images which guide his actions. A role is an image to which the individual strives to achieve congruence. Role enactment is the degree to which a person acts out a role appropriately (Solomon, Surprenant, Czepiel, & Gutman, 1985). For example, a 40-year-old man who works for a large company might hold the roles of businessman, friend,

¹⁵ The author made this observation by using the common search engine Google as well as Google’s Scholar (Beta) database, which searches scholarly publications. In Google Scholar, the term “ethics” returned about 2,690,000 results, and the term “business [and] ethics” returned 1,040,000 (“everyday [and] ethics” returned 190,000, and “common [and] ethics” returned 1,360,000 results). In Google’s normal Web search, “ethics” returned 93,600,000 and “business [and] ethics” returned 16,200,000 (“everyday [and] ethics” returned 439,000, and “common [and] ethics” returned 2,252,000). These queries were made in June of 2008.

husband, golfer, and father. Depending on his situation, the man identifies himself with the appropriate role and attempts to appropriately fulfill that role. He would act differently with his buddies on the golf course than he would with his daughter at bed time.

As Grover reports, role conflict is a situation where multiple roles are appropriate, but the actions of the roles do not allow for mutual satisfaction. For example, the businessman is always prepared for next-day meetings, but the husband and father always spends time with his family during the evenings. These conflicting role demands can lead to psychological distress. Alleviation of this distress can be achieved by both choosing one role and ignoring the others, avoiding role conflicts, or by compromising between the demanding roles (Grover, 1993).

C. SELF-INTEREST THEORIES

Grover also addresses self-interest:

The theory and research on deception and ethics in business has predominantly used the self-interest notion to explain deception. The organizational politics literature suggests that people may deceive in order to control resources, conceptualizing resource control as a self-interest commodity. Schein (1979) theorized that people use deception to attain power, especially in high slack organizations, where personal power may be readily pursued under the guise of organizational benefit.

(Grover, 1993)

1. Agency Theory

Agency theory is a type of self-interest theory, and it is based on an agent's relationship to his principal. The principal is defined as the party that delegates work to the agent party. The agent party is defined as the party receiving and performing the work delegated by the principal (Eisenhardt, 1989).

Agency theory is concerned with resolving two problems that can occur in agency relationships. The first is the agency problem that arises when (a) the desires or goals of the principal and agent conflict and (b) it is difficult

or expensive for the principal to verify what the agent is actually doing. The problem here is that the principal cannot verify that the agent has behaved appropriately.

(Eisenhardt, 1989)

2. Ethical Ambivalence Theory

In attempts to reduce the problems associated with the agency-principal relationship, organizations might introduce feedback to workers. This can communicate the values of the organization to the workers. When done perfectly, the worker would be provided feedback that guided his actions towards the goals of the organization. However, in an environment that is very complicated the practice of providing feedback might not be perfect, and ethical ambivalence can emerge.

Ethical ambivalence is a form of sociological ambivalence in which (a) the behaviors, attitudes, and norms that are shaped and maintained by the organizational reward system conflict with (b) the behaviors, attitudes, and norms congruent with the ethical values and judgments of organizational ... classes of individuals who have a vested interest in the outcomes of actions and the solutions of problems.

(Jansen & Mary Ann Von Glinow, 1985)

Organizational reward systems can be a source of ethical ambivalence to its members. In the academic community, openness and access to methods or research and experimentation are essential to the advancement of knowledge. However, business sponsorship of academic researchers can lead to “norms of secrecy that support knowledge as private property” (Jansen & Mary Ann Von Glinow, 1985).

In this example, the academic community underwent a change in reward system. Researchers were now rewarded for secrecy, or rather, secrecy served as a safeguard from “later involvements in priority disputes” (Jansen & Mary Ann Von Glinow, 1985). It might be safely assumed that, if increased knowledge is beneficial to business sponsors, any hindrance to its progress would not be in the best interest of the business. Thus, the effect of loss of information as a result of secrecy is probably not intended.

This misalignment of reward structures and desired results is a key point of Grover's model development:

Goal conflict and self-interest also operate in Jansen and Von Glinow's (1985) theory of ethical ambivalence, which states that behaviors rewarded by the organization may be at odds with those expected by some organizational stakeholders. Unethical behavior therefore may unintentionally result from inappropriate reward systems. That is, employees may engage in unethical behavior because that unethical behavior, rather than the expected behavior, is rewarded.

(Grover, 1993)

The computer model built in this thesis allowed for ethical ambivalence to arise by allowing Organizations to adapt their own reward and penalty structure. These structures are explained in Chapter III.

3. Grover's Proposed Limitations

Grover has identified several reasons why self-interest theories are not sufficient to explain deception in the workplace. His first reason is that self-interest theories "presuppose that individuals have the ability to predict and weigh the differential values of a variety of outcomes," and that organizational reinforcements might be unclear or too complex (Grover, 1993). Grover reports that another problem is that some individuals may not be acting out of self-interest, but out of selflessness or altruism. Also, certain lies may be prosocially motivated, such as the engineer who wants to complete more safety tests and so lies, saying that required testing (that had actually already been completed) had not yet been completed. By lying, he may have acted prosocially to make a safer product (Grover, 1993).

D. SOCIAL RATIONALITY

Grover agrees with researchers who theorize that self-interest cannot explain instances of deception because some people act altruistically or selflessly. In contrast, other research in neuroscience and cooperation suggest that altruism and selflessness are actually products of self-interest.

The Golden Rule says: “do unto others as you would have them do unto you.” This rule is almost universally accepted as a moral compass and is considered natural moral law (Cunningham, 1998). The Golden Rule can be found in the philosophy of Confucius, Islam, Christianity, Plato, Buddhism, Hinduism, and Immanuel Kant. It is even found in non-human species ranging from primates to insects (Pfaff, 2007). The Golden Rule is the foundation of moralistic trust (Cunningham, 1998), and neuroscientist Donald Pfaff theorizes that this trust is hard-wired in human brains as “reciprocal altruism.” Whether instinctive or not, patterns of reciprocity and mutual trust are observed in the real world (Pfaff, 2007).

1. Cooperation Emerging from Self-Interest

Even in an environment motivated only by the concern for one’s self, the Golden Rule can be found. Reciprocity has been shown to emerge in an iterated game of the prisoner’s dilemma, a game played by two people, each of whom individually makes a decision to cooperate or defect for a predetermined amount of points (Axelrod, 1997). Although an individual can earn more points in the short run by defecting, as both players defect more frequently the individual gains diminish. Eventually, both players learn that cooperating, while only moderately beneficial in the short run, is a more effective strategy in the long run.

Table 1. Payoff matrix for Axelrod's prisoner's dilemma

	Cooperate	Defect
Cooperate	3/3	0/5
Defect	5/0	1/1

2. Social Rationality Emerging from Self-Interest

Dr. Bruce Edmonds has performed many experiments that explore social intelligence emerging from evolutionary processes and social simulations. In several of

these experiments, he and Dr. David Hales have shown that social rationality (and some might argue altruism) can emerge from a collection of self-interested individuals who are subject to an evolutionary process or algorithm. In their 2002 experiments, Edmonds and Hales predicted that “If a socially rational agent can perform an action whose joint benefit is greater than its joint loss, then it may select that action” (Hales & Edmonds, 2003). In their M3 model, they found that the majority of individuals (under most circumstances) will perform actions to help others, even if those actions immediately result in an individual penalty (Hales & Edmonds, 2003). A chart of this trend is illustrated in Figure 3.

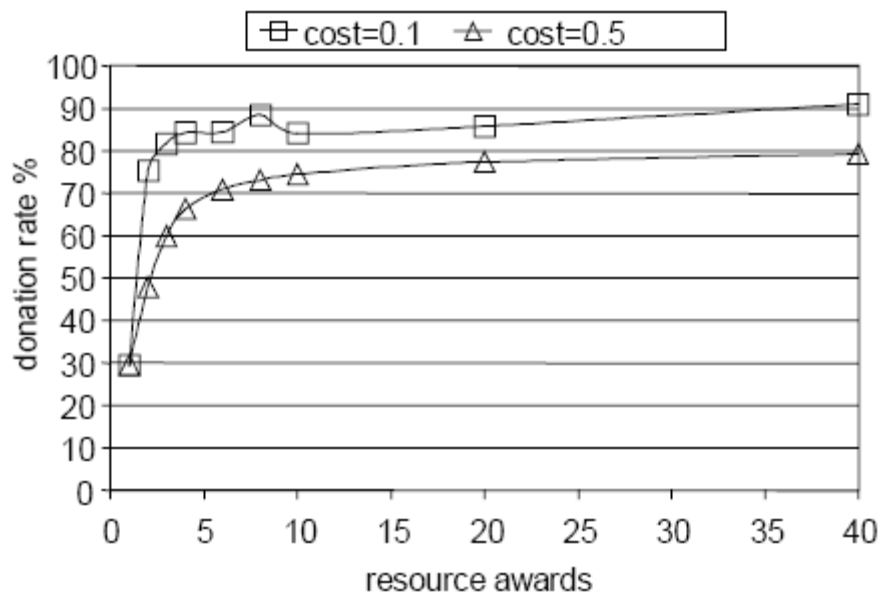


Figure 3: A chart of data from Edmonds' and Hales' M3 model show that individuals acting out of self-interest usually behaved socially rationally (Hales & Edmonds, 2003)

3. Social Rationality in Economics

Economies are examples of complicated environments in which individuals act out of self-interest. The classic El Farol Problem¹⁶ utilizes very simple individuals that

¹⁶ This is a classic economics problem. There is a bar that plays great music every Thursday night, but the bar is small and easily overcrowded. A good night is when the bar is not overcrowded (60% of the total population of potential attendees is considered to be a comfortable capacity), and a bad night is when the bar is overcrowded. The goal of each potential attendee is to attend El Farol on a good night. Each time an individual attends on a good night, they take a bonus to happiness. A penalty to happiness is awarded if the individual attends on a bad night.

act out of self-interest, and many have shown that over time individuals reach equilibrium behaviors (Arthur, 1994) (Edmonds, 1999b). Figure 4 illustrates how quickly individuals in an El Farol experiment can reach a social equilibrium that approximates the comfortable capacity of the bar (60 attendees). Figure 5 shows individual trends of attendance.




Figure 4: Individual decisions in an El Farol experiment reach a social equilibrium (Arthur, 1994)

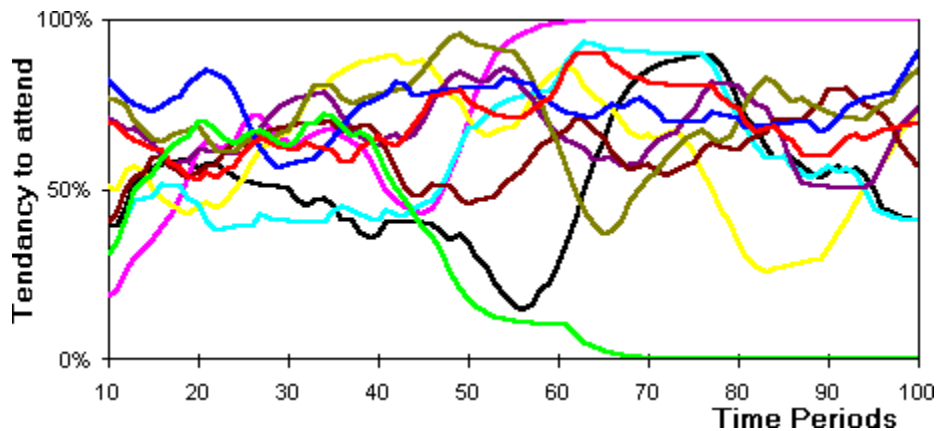


Figure 5: Development of individual strategies in an El Farol experiment (Edmonds, 1999b)

Thomas Sowell describes an economic situation in the real-world. He describes the evolution of the United Auto Workers (UAW) in the United States during the industrialization of the automobile industry. One of the purposes of the labor union was to ensure that auto workers' working conditions were safe and more comfortable, as well as to fight for the desires of the auto workers, which was usually the standard of payment, benefits, and available labor hours. As the UAW achieved more for its members, the effect was a natural rise in the operating costs of automobile makers because they had to pay higher wages and spend more on improving and maintaining working conditions. When foreign automobile makers, who did not allow the formation of labor unions, began to enter the US market, domestic automobile makers began to experience loss of sales and eventually significant financial trouble because the foreign auto makers were able to sell cars cheaper. The financial stress of US automobile makers forced them to find ways of cutting costs, and that led to the firing of hundreds of thousands of members (and non-members) of the UAW(Sowell, 2007).

It was the self-interest of the members of the UAW, then, that led to fewer jobs and was therefore detrimental to UAW members. The goal of the UAW was to improve the condition of auto workers, but in reality it harmed them (Sowell, 2007). While the self-interest of the UAW (better working conditions and higher pay) was different than the self-interest of the auto makers (better profits), it has become apparent that their survival and prosperity requires cooperation.

Modern labor unions have realized the need for cooperation with employers. Today, instead of being insensitive to the competitive market environment, labor unions understand the need to keep their employers profitable and surviving, even if it means lower wages, fewer benefits, and less luxurious working conditions (Sowell, 2007). This adaptation of the labor unions and automobile makers in the United States demonstrates how individual self-interest can lead to social rationality.

E. CONCLUSION

By integrating self-interest and role theory, Grover's model utilizes the advantages of both theories with regard to their explanations of why people might lie.

Grover's position is that neither self-interest nor psychological distress can individually explain a person's decision to lie. The purely self-interested agent (according to agency theory) only acts with the information that is perfectly knowable. Conflicting or overly demanding roles can lead to role conflict and distress. This distress can be alleviated by a compromise or by neglecting all but one role. These methods can be implemented by a lie, but to determine which role would be neglected or supported cannot be determined; it requires self-interest.

The model built in this thesis was able to incorporate the key theories of self-interest and role theory. Individual agents of the model acted in self-interest, and their actions and interactions provided role demands.

Just as the UAW's self-interest led to a significant decline in the welfare of its members, the Davis hypothesis suggests that the practice of deceit in an organization actually leads to a reduction in the practice of deceit, even if the deceivers are not directly punished. The use of multiple agents in an economic setting allows the complexity of reality to be modeled and for the emergent patterns suggested by the Davis hypothesis to be explored.

THIS PAGE INTENTIONALLY LEFT BLANK

III. SYSTEM DESIGN

A. INTRODUCTION

The computer model built in this thesis was a compound MAS (External Environment) that used the GA to give agents the ability to adapt (Internal Environment). There were two types of agents: Organization agents and Worker agents. The Organization agents were placed in a static¹⁷ Market Environment object. The Organization agents interacted with the Market Environment, and their interactions affected one another indirectly by modifying the Market Environment. Eight Worker agents were placed in each Organization agent, and each Worker interacted with his Organization.

The GA was applied to populations of Organization agents and Worker agents. Organization agents contained a gene string that dictated certain behaviors and policies of the organization. Worker agents contained a gene string that dictated actions to be taken for each possible assignment situation¹⁸.

Figure 6 illustrates a high-level view of the computer simulation. Figure 7 is a depiction of a single Organization.

¹⁷ Here, the term static means that the Market Environment was not adaptive and its attributes did not change.

¹⁸ Zero, one, two, or three Widgets could be assigned at any one time. There were four Widget types, and after removing duplicate permutations and further compression, this resulted in 20 possible assignment situations.

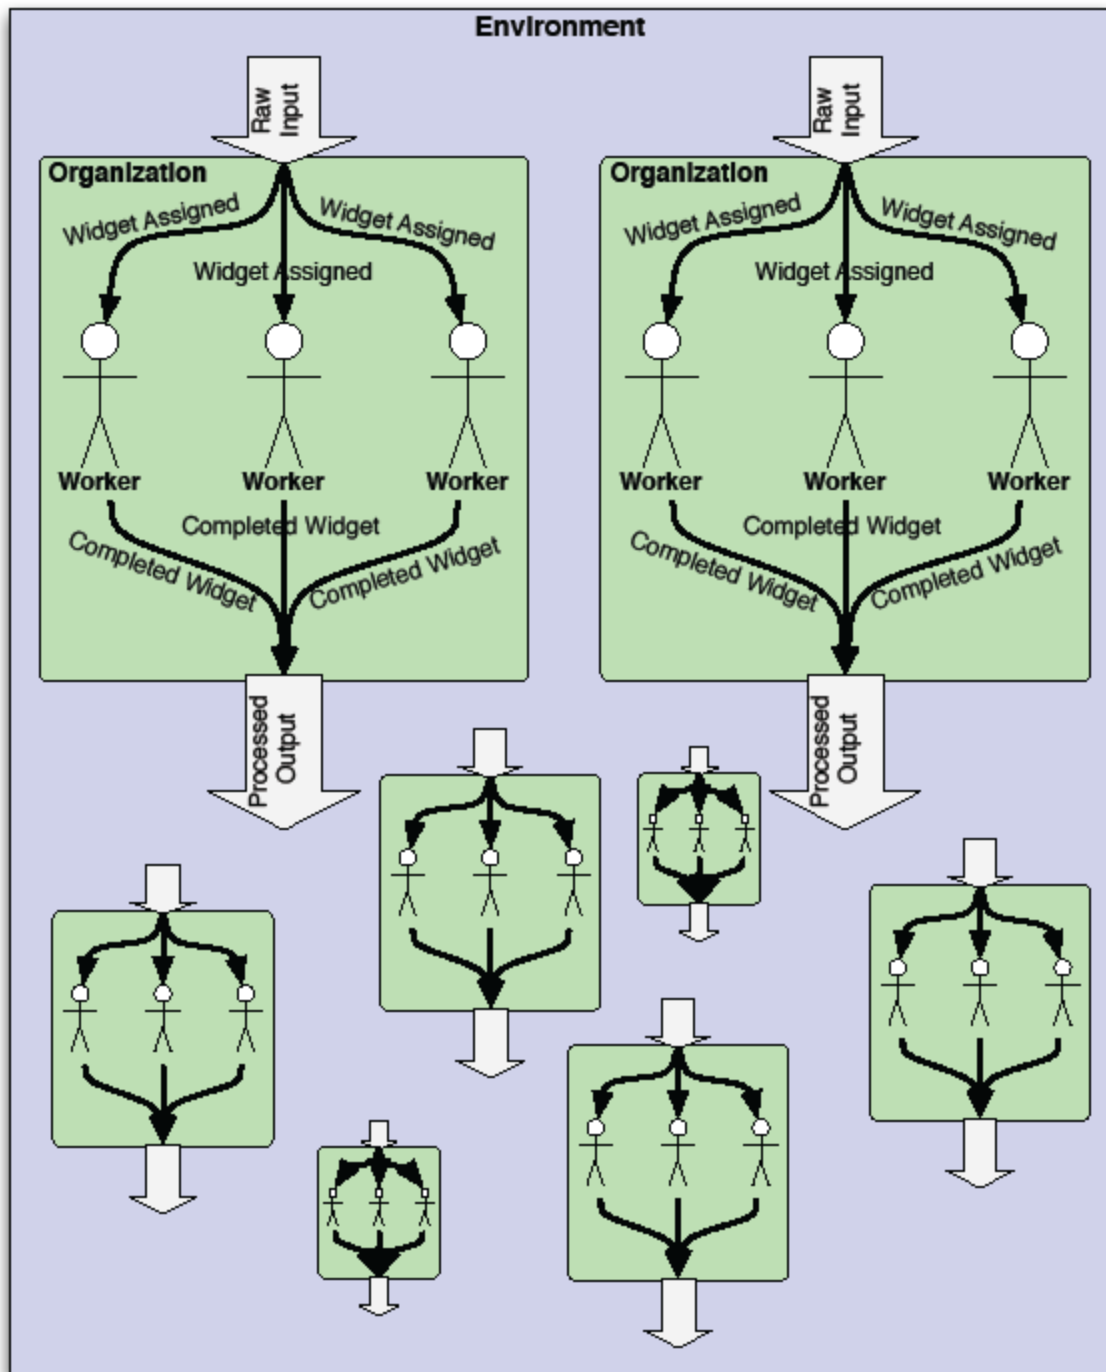


Figure 6: A top-level view of the computer model

B. TOPVIEW

The External Environment had three temporal components: the Year, the Day, and the Round. There were six Rounds in a Day, two Days in a Year, and 50 Years in a trial¹⁹. Each of these temporal components will be defined in this sub-section.

1. The Year

A Year contained two Days and the Organizational Internal Environment events. During the Organizational Internal Environment events, the GA was applied to the Organization gene pools, which allowed them to adapt. Each Organization had a chance to mate that was based on its net monetary flow. Organizations that had a higher net gain had more chances to mate each Year.

Both the Organizations and the Workers adapted (during their respective Internal Environment events) to a “moving target²⁰.” This approach was taken by Axelrod’s iterated prisoner’s dilemma, and his results showed that after an initial period of turbulence, individuals would still be able to adapt²¹. Axelrod discussed the tendency for individuals under a GA to become less random and mobile after a few generations, and that this phenomenon may also help to stabilize such a complex environment (Axelrod,

¹⁹ These were the numbers for the experiments run, and they are simple parameters that can be passed with different values to the computer model. These particular inputs were chosen because: six Rounds amplifies the difference between Workers’ decisions (because those decisions are made six times) and makes it possible for each Raw Widget to be completed with ten Workers; two Days allows successful Workers the opportunity to try previously successful decisions again; 50 Years allows for 50 generations of the Organizational Internal Environment and 100 generations of the Co-Worker Internal Environment. Axelrod ran his prisoner’s dilemma for 50 generations, although he saw peak performance after about 35 and observed a decline in average fitness until about the 9th generation (Axelrod, 1997). Also, early observations of the model also provided sufficient performance and adaptation around 35 generations. Since all data in this thesis, regardless of generation, was used in analysis, more generations helped to reduce noise associated with the initial turbulence and randomness associated with the first 10 or so generations.

²⁰ This is opposed to a condition where the Organization agents would not adapt or change at all, and the Worker agents would be the only agents adapting. This “non-moving target” approach is significantly less complex, but neither allows for the exploration of an Organization’s behavior nor represents reality very accurately.

²¹ As opposed to the supposed alternative, in which the agents would be required to adapt so fast that they could not adapt and the entire simulation would be random and chaotic.

1997). It was found in early testing and observation that both the Organizations and the Workers could adapt simultaneously and still show progressive adaptation.

2. The Day

A Day contained six External Environment events (steps 1-6), and the Co-Worker Internal Environment events (step 7):

1. New Raw Widgets were instantiated
2. Raw Widgets were sold to bidding Organizations
3. The Round begins and ends
4. Completed Widgets were sold by Organizations to the Market Environment
5. An Organization's unsold Widgets were destroyed
6. The Market Environment had a chance to inspect recently purchased Widgets and return Falsified Widgets for a refund
7. The GA was applied to the Worker gene pools (Co-Worker Internal Environment events begin and end)

A Worker's fitness was determined by his choices during the Round (part of the External Environment) and by the Organization's rewards and penalties policies²². In general, Workers who Completed more Widgets usually had a higher fitness. This fitness was then used by the Co-Worker Internal Environment events to adapt the Workers.

3. The Round

A Day contained six rounds. Each Round operated only in the External Environment. The Round was an opportunity for the Organizations to take recently purchased Raw Widgets and assign them to Workers. Up to three Widgets of any combination of types could be assigned to a Worker each round. After assignment,

²² Rewards and penalties policies were only effective during experiment three.

Workers had the opportunity to interact with the Widgets and could attempt to 1) Complete, 2) Falsify or 3) Do Nothing to the Widgets. The Widgets were then returned to the Organization, which might have conducted spot checking for Falsified Widgets²³.

4. The Step

Each time a Worker was given an assignment of Widgets, he was given exactly four steps²⁴ to complete his tasks. Each step, the Worker made a decision. If no tasks/Widgets remained before the conclusion of the four steps, the stepping process was ended.

C. THE MULTI AGENT SYSTEM

The objects in the MAS included Widgets, Worker agents, Organization agents, and the Market Environment.

1. The Widget

Widget objects represented goods that could be 1) purchased raw, 2) improved, and 3) resold. Widgets provided Organization agents a way to earn revenue. Worker agents Completed, Falsified, or Did Nothing to Widgets to earn performance feedback from Organization agents.

a. Widget Attributes

Each Widget object had two key Boolean attributes. The first was the “complete” attribute, and the second was the “honestly_done” attribute. The combinations of these two Boolean attributes defined the possible states of a Widget. Table 2 shows the possible states of the Widget object.

²³ Spot checking was only effective in experiment three.

²⁴ This aspect of the model was one of the few decisions in which definition could not be changed or tweaked after the beginning of software development. Other such decisions were: there were three types of Widgets, type A, type B, and type C; an Organization may assign up to three Widgets of any combination of types.

Table 2. States of the Widget object

Actual State	State as perceived by Organization objects	“complete” flag value	“honestly_done” flag value
Raw ²⁵	Raw	False	True
Complete	Complete	True	True
Falsified	Complete	False	False

Widget objects also had four other attributes, listed in Table 3. The “type” attribute defined the Widget object as one of three types of Widget. Different Workers had different skill abilities to Complete a Widget, and this ability was specific to the Widget’s type. The “p_success_lie” was a value between zero and one (inclusive) and represented the probability that a Worker attempting to Falsify this particular Widget object would succeed. “p_success_lie” was set to one for all experiments.

The “time_spoiling” and “time_consumed” attributes held integer values that told an observer how long a particular Widget had either 1) been waiting in the Market Environment for an Organization to purchase it (spoiling) or 2) how long, after being sold back to the Market Environment, it has been consumed. These values were used by the Market Environment to determine when a Widget would be removed from the simulation. For the experiments in this thesis, Raw Widgets spoiled after one day of not being purchased by an Organization and were consumed after one day after being resold to the Market Environment.

²⁵ A Raw Widget is equivalent to an incomplete Widget. Widgets not Completed or Falsified remain Raw.

Table 3. Other Widget attribute values and purpose

Attribute	Possible Values	Purpose
type	“A”, “B”, or “C”	Defines the type of the Widget object, which is used to provide different levels of difficulty.
p_success_lie	$\mathbb{R} \cap \{0 \leq value \leq 1\}$	Dictates the p(success) of an attempt to Falsify.
time_spoiling	$\mathbb{Z} \cap \{0 \leq value\}$	Represents the amount of time a Widget has been waiting to be purchased by an Organization
time_consumed	$\mathbb{Z} \cap \{0 \leq value\}$	Represents the amount of time a Widget has been in the Market Environment after having been resold to the Market Environment by an Organization

b. Important Widget Methods

Each Widget had three important methods, listed in Table 4. These methods represented an interaction between a Worker and a Widget object. To make an honest attempt at Completing a Widget, a Worker would call the “honest_attempt” method. The Worker provided his skill, represented by a p(success). If successful, the Widget’s state would change to Complete. To attempt to Falsify a Widget, a Worker would call the “attempt_to_decieve” method. If successful, the Widget’s state would be changed to Falisified. Workers were required to interact with Widgets, so the “do_nothing” method represented a Worker’s choice to take no action²⁶, or to keep the Widget in the Raw state.

²⁶ Because the GA is used to adapt the Worker’s decisions (and because a technique was used to reduce the size of the working decision table to make the size of a Worker’s binary string manageable), it was possible (and quite frequent) for Workers to “misfire”, i.e. to indicate that an action for a certain type of Widget be taken, but to then find that no Widget of the indicated type was available. Misfires simply fizzle, and were effectively decisions to “do_nothing” which were analyzed as such.

Table 4. Important Widget methods

Method	Input	Potential Widget State Affect
honest_attempt	$\mathbb{R} \cap \{0 \leq value \leq 1\}$	If successful, can change from Raw to Complete
attempt_to_deceive	None	If successful, can change from Raw to Falsified
do_nothing	None	None

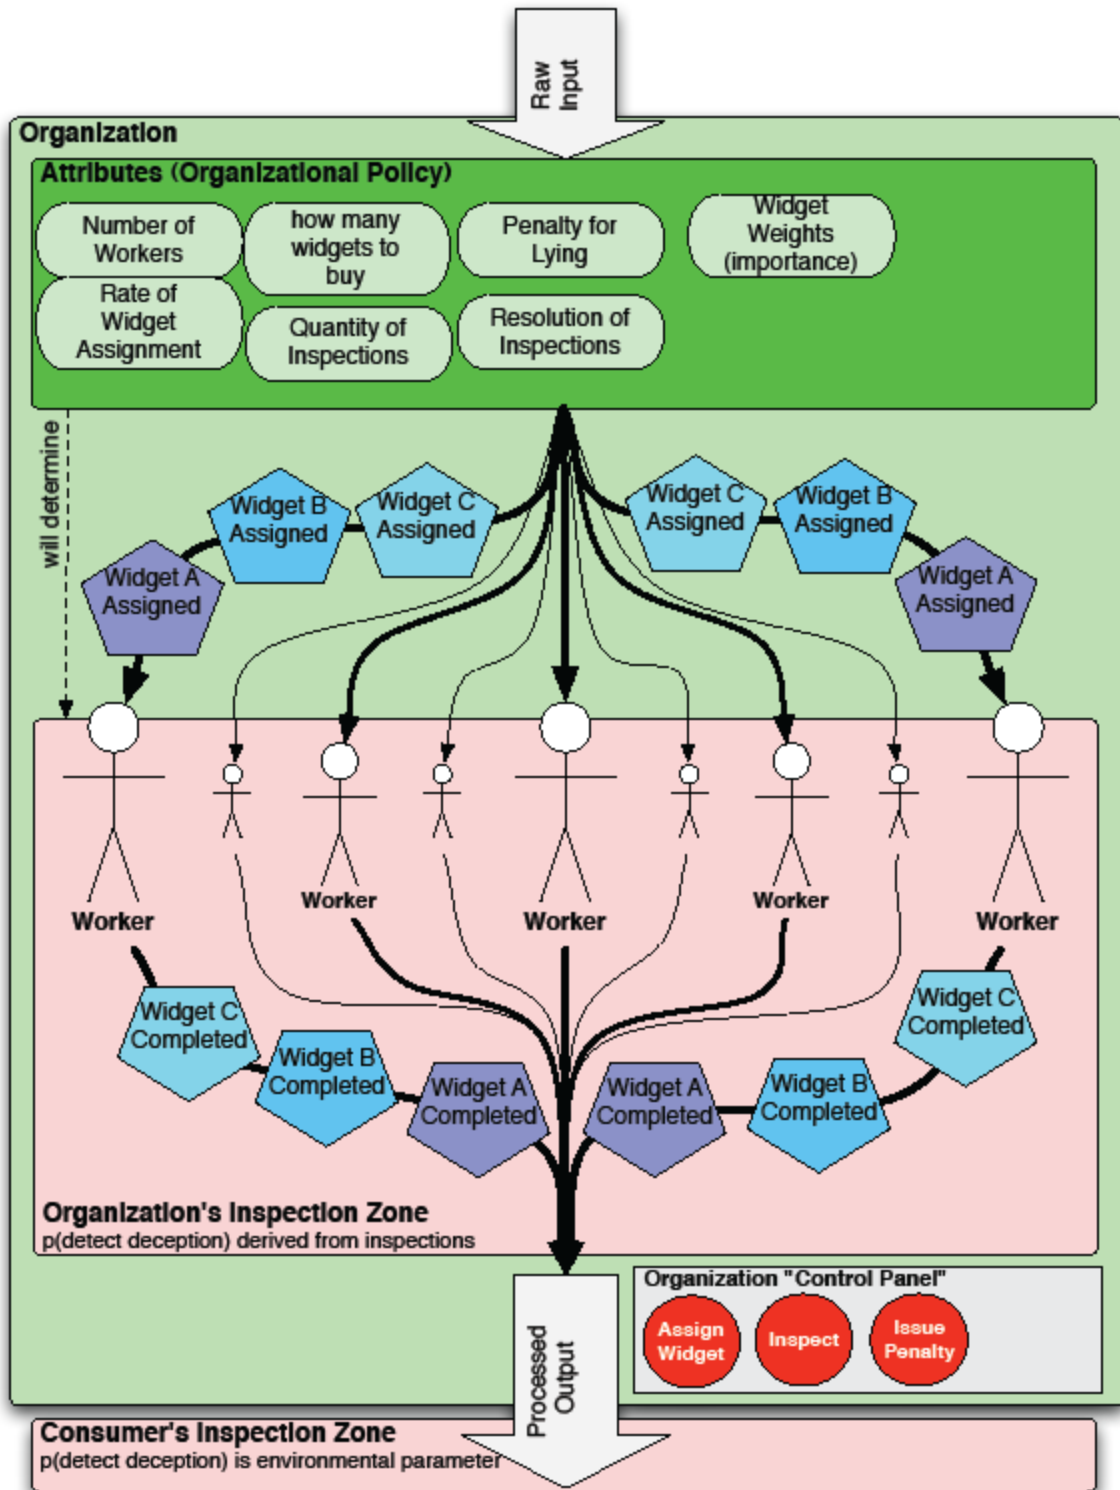


Figure 7: A view of a single Organization

2. The Worker

The Worker object was an agent that represented an individual person in an Organization. The goal of each Worker was to survive, and fitness for survival was determined by the Worker's actions with regard to interacting with Widgets. In general, a Worker that transformed more Widgets into the Complete state than another would have a higher fitness²⁷.

Each Worker object had six attributes, listed in Table 5. Three of the attributes represented a Worker's skill or ability to transform Raw Widgets into Complete Widgets. The fourth attribute was a lookup table that was used by the Worker to make a decision about what action to take when presented with a particular situation. The Rewards and Penalties attributes held a current tally of a Worker's feedback from his organization. These two values were used to calculate a Worker's fitness. Workers were rewarded for Completing Widgets and penalized for not Completing Widgets.

The Skill values of the Worker were randomly assigned at instantiation, and remained static throughout a trial. Even if a Worker was "killed" by the GA, the new Worker's Skill attributes would be identical. The possible values were randomly determined, but they were equally spaced to approximate a normal distribution of skill levels.

The GA was applied to the Decision Table attribute to provide an adaptation mechanism. The Rewards and Penalties attributes were reset to zero after each round of adaption by the GA. The decision table contained every possible situation (20 in all) of Widget assignment by an Organization²⁸ for each step.

²⁷ Or that made the Organization *think* he had done so by Falsifying a Widget and not getting caught.

²⁸ Organizations were allowed to assign zero, one, two, or three Widgets, in any combination of types. For example, two type A Widgets and one type C Widget could be assigned, or zero Widgets, or three type B Widgets.

Table 5. Worker Attributes

Attribute	Possible Values	Purpose
Skill_A	{0, 0.11, 0.22, 0.33, 0.44, 0.56, 0.67, 0.78, 0.89, 1}	Represent a Worker's ability to honestly Complete a Widget object of type "A"
Skill_B	{0, 0.11, 0.22, 0.33, 0.44, 0.56, 0.67, 0.78, 0.89, 1}	Represent a Worker's ability to honestly Complete a Widget object of type "B"
Skill_C	{0, 0.11, 0.22, 0.33, 0.44, 0.56, 0.67, 0.78, 0.89, 1}	Represent a Worker's ability to honestly Complete a Widget object of type "C"
Decision Table	Gene object ²⁹	A lookup table for the individual Worker that provides the decision to be made by the Worker given a particular situation
Rewards	$\mathbb{R} \cap \{0 \leq value\}$	Holds the current value of the rewards given to a Worker by his Organization
Penalties	$\mathbb{R} \cap \{0 \leq value\}$	Holds the current value of the penalties given to a Worker by his Organization

3. The Organization

The Organization object served a dual role as an agent and as a working environment. The Organization was an agent whose actions affect the External Environment, but the Organization also had policies that affected the Worker objects and therefore these policies were effectively the environment in which Workers operated. This working environment will be referred to as the Work Environment. The Organization provided opportunities for Workers to manipulate Widgets, and also provided feedback to the Workers. Organization objects could buy Raw Widgets from the Market Environment and resell Widgets it perceived to be Complete.

The fitness of an Organization was determined by the net monetary gain of that Organization. Table 6 outlines the instances of gain or loss of money for an

²⁹ A Gene object is a set of attributes that can be transformed by the GA. The whole of the object can be represented as a binary string.

organization. The three opportunities for a loss (or investment) of revenue were 1) when purchasing Raw Widgets, 2) when inspecting Complete Widgets, and 3) in the event that the Market Environment found a Falsified Widget. The sole source of income for Organizations was the ability to resell Complete Widgets.

Table 6. Instances of Organization monetary flow

Instance	Flow	Magnitude factors
Buying a Raw Widget	Loss	Current market price is determined by the type of widget, and by the actions of other purchasing Organizations.
Inspection of a Widget claimed to be Complete	Loss	Determined by a global constant for each Widget type multiplied by a global factor determined by the Organization's inspection resolution policy
Market Environment detects a Falsified Widget	Loss	Cost of an inspection as above, plus the price that the Widget sold for multiplied by 1.3
Resale of a Complete ³⁰ Widget	Gain	Price as set by the Organization's policy

Each Organization had three attributes, as listed in Table 7. The monetary gains and losses were stored in these attributes, as was the Organization's business decisions and policy.

Table 7. Organization attributes

Attribute	Possible Values	Purpose
Credits	$\mathbb{R} \cap \{0 \leq value\}$	Represents the sum of monetary gain
Debt	$\mathbb{R} \cap \{0 \leq value\}$	Represents the sum of monetary loss
Policy	Gene object	Provide business decisions/policy

³⁰ As perceived by the Organization

The Policy attribute held many decisions that could be made by the Organization. The details of these decisions are contained in Table 8. The freedom given to the Organizations with regard to policy, especially in the rewards and punishment structure, allowed the computer model to exhibit ethical ambivalence. Therefore, the potential for agents to experience distress due to a misaligned rewards system, as required by Grover's model, was achieved.

The policy decisions that were not related to the rewards structures helped the Organizations behave like real organizations. The Buy Decision and Buy Amount policies allowed the Organization to move between the three Widget markets. If too many Organizations were buying Widgets from the Widget A market (raising the average loss to participating Organizations), the Organization had the freedom to exclusively or also purchase from the Widget B market, where there might have been less competition to buy and thus lower prices.

The Resale Price attribute allowed the Organization to name a price for the resale of processed Widgets. The Market Environment took bids of the Organizations and determined the quantity demanded in accordance with the basic economic principle of elastic demand³¹. As with any free market, the price set by the Organization would not just be an arbitrary value³². Organizations that set their prices very high would usually receive only a few purchases from the Market Environment³³ and Organizations with the lowest prices had the greatest chance of selling each Widget. This free market behavior

³¹ Generally, people will buy more of something (increase demand) when the price is lower. People will buy less of something (decrease demand) as the price of a product rises (Sowell, 2007). The sale of Raw Widgets to Organizations simulates a demand curve shift, and the resale of processed Widgets to the Environment simulates price elasticity of demand $n > 1$.

³² The value of each Organization's resale price is random during the first Year.

³³ An Organization that was a monopoly in a particular market could set a high price and still have enough demand for his product to make great revenue. However, just as in real free markets with elastic demand curves (Sowell, 2007), the entrance of competitors into the market will motivate Organizations to lower their prices in order to sell their inventory of Widgets. Organizations that didn't lower their prices usually died from lack of revenue.

allowed the computer model to achieve a higher level of realism and to explore the higher-order effects of individual Worker and Organization actions³⁴.

The $p(\text{inspection})$ attribute defined the probability that a Completed Widget would be inspected by the Organization. The cost of an inspection is shown in Figure 10. The Attribution Resolution, Lying Penalty, and Weight all contributed to the magnitude of the rewards and punishments given to Workers by Organizations according to Figure 8. The Attribution Resolution would determine how many Workers were penalized when a Falsified Widget was detected³⁵. While the Workers did not calculate expected payoffs, each Round they were subject to the expected payoffs shown in Figure 9.

$$\text{Penalty}_{\text{Falsification Detected}} = \left[\frac{1}{\text{Attribution Resolution}} \times \text{Weight} \right] + \left[\frac{4}{\text{Attribution Resolution}} \times \text{Lying Penalty} \right]$$

$$\text{Penalty}_{\text{Failure to Complete a Widget}} = \text{Weight}$$

$$\text{Reward}_{\text{Widget perceived Complete}} = \text{Weight} \times 2$$

Figure 8: Penalty and reward equations based on the attributes of the Organization policy

$$\text{Expected Gain}_{\text{attempt_to_deceive}(1)} = \text{Reward}_{\text{Widget perceived Complete}} - p(\text{inspection}) \times \text{Penalty}_{\text{Falsification Detected}}$$

$$\text{Expected Gain}_{\text{honest_attempt}} = \text{Reward}_{\text{Widget perceived Complete}} \times \text{Skill} - (1-\text{Skill}) \times \text{Penalty}_{\text{Failure to Complete a Widget}}$$

$$\text{Expected Gain}_{\text{do_nothing}} = \text{Penalty}_{\text{Failure to Complete a Widget}}$$

Figure 9: Expected Worker payoffs

³⁴ In an early iteration in the development of the computer model, a centrally defined/coordinated price system was implemented. This system was extremely hard to tune, because improperly set values would either cause Organizations to fail without question, or to flourish without regard for customer satisfaction or Worker behavior. These problems, and the resultant success of the system when changed to an approximate free market, are congruent with the basic principles of economics according to Sowell (Sowell, 2007).

³⁵ The Attribution Resolution is directly proportional to the number of other Workers that will be penalized. This represents an Organization's ability or relative inability to determine the source of Falsification. As more Workers are penalized, the penalty given to each worker is proportionally lower in magnitude.

Table 8. Organization business and policy options

Attribute	Possible Values	Purpose
Buy Decision*	Boolean	Determine if a type of Raw Widget will be bid for
Buy Amount*	{0.06, 0.13, 0.19, 0.25, 0.31, 0.38, 0.44, 0.50, 0.56, 0.63, 0.69, 0.75, 0.81, 0.88, 0.94, 1} ‡	Used if Buy Decision is True. When offered to buy a Raw Widget, this is the p(buy) to simulate a scalable quantity of Raw Widgets desired
Resale Price*	$\mathbb{Z} \cap \{25 \leq value \leq 51225\}$ †	The asking price for a Completed Widget
p(inspection)*	{0, 0.14, 0.29, .43, 0.57, 0.71, 0.86, 1} ‡	Probability that a Completed Widget will be inspected before resale
Attribution Resolution*	$\mathbb{Z} \cap \{0 \leq value \leq 7\}$	How finely a Falsified Widget can be traced back to an individual Worker, and helps determine how much a Worker who is caught Falsifying a Widget will be penalized
Lying Penalty*	$\mathbb{Z} \cap \{0 \leq value \leq 7\}$	A factor that helps determine how much a Worker who is caught Falsifying a Widget will be penalized
Rate of Assignment	Single value from the set of all 20 possible combinations of assigning from 0 to 3 Widgets	How many Widgets, and of what type composition, will try to be assigned to each Worker
Weight*	$\mathbb{Z} \cap \{0 \leq value \leq 7\}$	A factor that determines how much a Worker will be rewarded or penalized for a Widget

* Each of these actually contains 3 individual instances, one for each of the Widget types: “A”, “B”, and “C”

† Values in this range are in increment of 25

‡ Values in this cell have been rounded from four to two decimal places for readability

4. The Market Environment

The Market Environment object was the top-level object that contained all of the other objects and served as the operating environment of the Organization objects. The Market Environment instantiated new Widgets and provided them to the Organizations for purchase. The Market Environment also accepted the bids of Organizations that wished to resell Completed Widgets, and it could act as the “consumer” who might discover Falsified Widgets and return them to the Organization of origin for a refund.

The Organization objects experienced the Market Environment through its attributes, listed in Table 9. All of these values were held constant in experiments one and three.

One of the Market Environment’s tasks was to instantiate new Widget objects. This was done each Day. Just as milk will spoil if it is not processed quickly enough, Widgets were designed to spoil if they were not purchased and preserved by the Organizations. Widgets were created at the rates defined in the Widget Renew Rates attribute. Widgets that were not purchased spoiled after they exceeded the values stored in the Widget Max Spoil attribute. The spoiling prevented Raw Widgets from piling up past a certain point. In all experiments, this value was actually set to 0, which meant that Raw Widgets had only one chance to be purchased by an Organization.

Each Day, the Organizations had the opportunity to inspect a portion of inventory to look for Falsified Widgets. The Market Environment determined the expense of these inspections with the Inspection Costs and the Attribution Resolution Factors attributes. The Inspection Cost was a cost for all Organizations and was determined according to Figure 10. Each Organization could determine the resolution they would attribute a Falsified Widget, but a greater resolution would cost the Organization more money each inspection. The Attribution Resolution Factors attribute provided this factor. The cost of an inspection was the Inspection Costs multiplied by the Attribution Resolution Factors. In the experiments, the Inspection Costs was 250, and the Attribution Resolution Factors ranged from one for the lowest resolution, to 4.5 for the highest.

$$\text{Cost}_{\text{inspection}} = \text{Attribution Resolution Factor} \times \text{Inspection Cost}$$

Figure 10: Cost of inspections

Widgets were resold to the Market Environment as Completed also had a shelf life. This is analogous to a gallon of milk that must be consumed within a certain amount of time. The Widget Max Consumption defined the number of days after which a Widget would be removed from the Market Environment. If Completed Widgets were not consumed, the demand for Completed Widgets would plummet, and Organizations would have to lower prices to a level below which they could not make a profit. In the experiments, this value was set to zero, which meant that Widgets were consumed within one Day.

After being resold to the Market Environment, Widgets might have been inspected. This was to simulate the real-world “money back guarantee” that some businesses provide and the assumption that if an organization’s products are always broken, fewer people will purchase them. Each Widget might be inspected with a $p(\text{inspection})$ equal to the Consumer $p(\text{detection})$ attribute. In the experiments, this value was arbitrarily set at $\text{Consumer } p(\text{detection}) = 0.5$. Manipulation of this variable was the focus of experiment two.

Table 9. Market Environment attributes

Attribute	Possible Values	Purpose
Widget Renew Rates	$\mathbb{Z} \cap \{0 \leq value\}$	Each time unit, this determines how many new Widgets are instantiated
Widget Max Consumption	$\mathbb{Z} \cap \{0 \leq value\}$	Each time unit, this determines the maximum Widget time_consumed before a Widget is removed
Widget Max Spoil	$\mathbb{Z} \cap \{0 \leq value\}$	Each time unit, this determines the maximum Widget time_spoiling before a Widget is removed
Inspection Costs	$\mathbb{Z} \cap \{0 \leq value\}$	A factor in the cost to an Organization conducting an inspection of Widgets of that type
Attribution Resolution Factors	8 numbers, each in the set $\mathbb{R} \cap \{0 \leq value\}$	A factor in the cost to an Organization conducting an inspection of Widgets of that type
Consumer p(detection)	$\mathbb{R} \cap \{0 \leq value \leq 1\}$	The probability that a Falsified Widget will be detected after resale and returned for a refund

D. THE GENETIC ALGORITHM

The GA was used to provide Worker and Organization objects with the ability to adapt some of their attributes and behaviors. The events of GA were contained within the Internal Environment. The GA was performed on a grouping of subjects, called gene pools. Each trial included three Organization gene pools and twelve Worker gene pools.

1. Organization Gene Pools (Organizational Internal Environment)

The Organizational Internal Environment was divided into three gene pools. Under certain conditions, the GA will evolve superindividuals who have very high

fitness, but who have a tendency to homogenize the gene pool (Holland, 1992). While some simulations desire this result³⁶, the goal of this experiment is to evolve agents who work well together.

To combat this homogenizing, Holland suggests an alteration of the original GA where fitness is not determined by the individual, but by the success of a group of rules (Holland, 1992). This approach requires that success be Boolean, but the nature of this computer model not provide such a clean indicator of success.

Holland's alteration introduces segregation between agents, and so a simple segregation was imposed. Separating the organizations into three sections did lead to superindividuals within each section, but the organizations were heterogeneous at the section-level. Each section could be thought of as an independent entity that was made of four Organizations.

a. Organization Gene Pool Parameters

The length of an individual Organization binary string was 86 bits. A single crossover per mating was used. A constant population of Organizations was maintained so that the same Workers could be studied for the entire trial.

The use of small population sizes with a high mutation rate has been studied by Randy Haupt, who found that populations of four, with a mutation rate of 0.15 per bit, were most effective at finding global maximums in the fewest function calls (Haupt, 2000). This was the reasoning behind the gene pool population of four Organizations with a mutation rate of 0.15. A high mutation rate also helped to fight homogeneity.

The Organization's policies are very sensitive to changes in the binary string, so organizations that were selected for mating were allowed to survive to the next generation. Each mating produced two offspring. In most cases, this resulted one mating and the death of two organizations per generation. Special cases where this did not happen will be discussed in the next sub-section.

³⁶ Such as Axelrod's iterated prisoner's dilemma, where he sought the best strategy to the game.

b. Additional Organization Design Decisions

In early testing and development of the computer model, all of the Organizations performed so poorly that they resorted to buying nothing, assigning nothing, and selling nothing. Generations with every Organization's fitness equal to zero emerged. This may have been because from a random start, the policies of an Organization might not be suitable for a positive net monetary flow. Some trials would develop Organizations that "figured it out" and broke out of debt, but in others, the Organizations seemed to just give up. Since there was no penalty if everybody gave up, Organizations and Workers did not have much incentive to perform.

The design was modified to incentivize performance. To reach maturity (have a chance at mating), an Organization was required to have a fitness of greater than zero. Any Organization with a fitness level of less than zero was killed. If all Organizations were less than zero, the most fit would mate as usual. If all fitnesses were equal to zero, matings were chosen at random. If all were below zero, then selection was made normally.

Under normal circumstances, all mates were required to be unique (Organizations could not exhibit asexual behavior). However, if only one Organization survived to maturity, he was allowed to mate asexually and produce three offspring.

This harsh additional selection mechanism more rapidly produced reasonable Organizational policies. The high mutation rate also helped the Organizations find useful policy combinations.

2. Worker Gene Pools (Co-Worker Internal Environments)

The Co-Worker Internal Environments each contained one Worker gene pool, which was comprised of the eight co-Workers from the same Organization. The length of an individual Worker binary string was 320 bits. Due to the length of this binary string, three crossovers per mating were used. The population size of eight was a compromise between the need for Organizations to have a good distribution of skill within its workforce, and the need to make the population size small and highly mobile.

Since the Workers' genes are in more need of exploration than the Organizations' genes, a much higher mutation rate, 0.35, was used. This is the highest mutation rate recommended by Haupt's study (Haupt, 2000).

The same harsh selection mechanism that was used in the Organization gene pools was also used in the Worker gene pools. In this gene pool, this mechanism also proved to help quickly develop the Workers into rational decision makers. The criterion for asexual mating and random selection used in the Organization gene pools were used in the Worker gene.

Each Worker had three *static* skill values, but his decisions were *adaptive*. Just as with the Organization gene pool, the tendency for the Worker gene pool to develop superindividuals was an issue. This was a problem because the study aimed to evolve an individual decision table based on individual skill level. It was found, however, that individually varying skill sets acted as a heterogeneous-seeking force in the gene pools.

For example, it was common for a Worker with a low Skill_A to die and receive the offspring from a Worker with a high Skill_A. This caused turbulence with the lesser skilled Worker, but after a few generations it was not uncommon for the lesser skilled Worker to receive a more compatible or advantageous set of genes. Also, highly skilled Workers were not very disturbed when they became the recipient of the offspring of a lesser skilled Worker. Since the aim of this study was not to develop probabilities of behavior, only to see if a difference in behavior between skill levels arose, this turbulence was accepted.

E. CONCLUSION

The Worker was the focus of this thesis, and the design of the system was based around giving Worker agents opportunities to experience distress while acting out of self-interest. The self-interest requirement was met by allowing the Worker to not only honestly attempt to Complete Widgets, but by allowing the Worker to take no action, or to take a shortcut (Falsifying) that increases the chances of survival of the Worker but that might not be in line with the goal of the Organization.

Information about skill levels and decisions of each Worker allowed for an analysis of Grover's hypothesis.

The inclusion of the Organization agents was an essential part of creating a complex environment to approximate reality. Organizations facilitated a complex and changing environment in which the Worker agents could adapt. They also allowed for higher-order effects to propagate. For example, if a Worker decided to Falsify a Widget, the Organization might sell that Widget. If the Environment detected this Widget, it would provide negative feedback to the Organization. The Organization might have to adapt or risk death. This adaptation might have come in the form of increased inspections by the Organization, which in turn would have changed the Worker's environment. The Worker might then need to adapt or risk survival.

IV. VERIFICATION, VALIDATION, AND EXPERIMENTATION

A. INTRODUCTION

This section will provide evidence to show that the computer model built in this thesis is both verified and valid. Also, details of the experiments that were run to test the hypotheses of this thesis will be provided, and the data from each of the three experiments will be analyzed.

B. VERIFICATION AND VALIDATION

This sub-section will provide supporting evidence that the MAS built in this thesis is verified and valid. Specifically, this sub-section will show that the computer model demonstrates the complex and adaptive behavior that is characteristic of MASes. By virtue of this demonstration, the process of the GA will also be validated, since the GA provided each agent the ability to adapt. Additional trends in the data will be characteristic of an applied GA.

1. Characteristics of MASes

A Mutli Agent System produces complex and adaptive aggregate trends. This sub-section will provide data from an experiment done by other researchers using MASes.

The agents within a MAS adapt to the conditions of their environment. To illustrate this adaptation, Edmonds' 1999 experiment (where he modeled economic agents with bounded rationality) will be used.

Figure 11 is a graph of an environmental aspect of this experiment. It shows that there were two states of the environment, the easy state and the hard state. The easy state is represented by the curve that allows for very high utility often and at lower quantities of purchased products. The easy state also has only one local maximum. The hard state is represented by the curve that allows for very high utility only in certain areas of the

problem space, and also has two local maxima. At any one time, only one of these environmental states may be active.




Figure 11: Edmonds' environment contained two states represented by the easy curve and the hard curve (Edmonds, 1999a)

One part of Edmonds' experiment was to begin the model with one environmental state and to then change the state of the environment. Figure 12 shows the effect on individual utility when the environment began in the easy state and was then changed to the hard state at date 50. Figure 13 shows the effect when the environment began in the hard state and was then changed to the easy state.

It is important to notice how the system adapts and achieves a new equilibrium state when the environment changes. In both figures, the agents found the global maximum of the easy state within the first 10 or so dates. When the state of the environment was changed to hard, the agents began searching for a new equilibrium.




Figure 12: Agents adapt from the easy state to the hard state (Edmonds, 1999a)




Figure 13: Agents adapt from the hard state to the easy state (Edmonds, 1999a)

2. Demonstration of MAS Behavior in the Computer Model of this Thesis

This sub-section establishes that the computer model built in this thesis demonstrates the characteristics of a MAS. To achieve this, an environmental change was made between three permutations of a run.

This change was to the Market Environment's Widget Renew Rates attribute. This attribute was varied and run in three permutations: The first ran the model with 100 Widgets of each type spawning per Day; the second only had 600 Widgets spawn each day; the third permutation had 1,100 Widgets spawn each Day. This attribute affects the number of Widgets that can be purchased by any Organization, and in turn, the number of widgets that can be manipulated by Workers and then resold to the Environment for a profit.

Table 10. Verification Experiment Market Environment attribute values

Market Environment Attribute	Values
Widget Max Consumption	A: 0 B: 0 C: 0
Widget Max Spoil	A: 0 B: 0 C: 0
Inspection Costs	A: 250 B: 250 C: 250
Consumer p(detection)	A: 0.5 B: 0.5 C: 0.5
Attribution Resolution Factors	0:1; 1:4.5; 2:4; 3:3.5; 4:3; 5:2.5; 6:2; 7:1.5

Table 11. Verification Experiment independent Market Environment attribute

Market Environment Attribute	Permutation 1	Permutation 2	Permutation 3
Widget Renew Rates	A: 100 B: 100 C: 100	A: 600 B: 600 C: 600	A: 1100 B: 1100 C: 1100

Several parts of the model were observed: the Organization and Worker average fitnesses, the number of widgets sold, the price of Completed Widgets, the demand for Completed Widgets, and the average number of changes made each Year to the policy tables of the Organizations.

The Market Environment was not changed during the run, but between permutations, key factors show significant change. Those factors that did not change are also an indication of success. A complete set of charts for this run is included in Appendix A, but the average Organization profits and the average price of resold Widgets will be examined in this section.

Figure 14, Figure 15, and Figure 16 chart the average Organization profits over time. Profit is the net positive flow of money to an Organization. In Figure 14, which

shows the situation where there are very few Widgets (100 per Day) in the Market Environment, shows how the average profit oscillates near 100,000 credits. When the environment is changed to include more Widgets (600 per Day), as depicted in Figure 15, a significant increase in the net monetary gain of the Organizations was observed. This demonstrates the computer models ability to adapt to and take advantage of the higher availability of Widgets. Figure 16 shows another slight increase when the number of Widgets is made very high (1,100 per Day).




Figure 14: Verification permutation one (low number of Widgets) average organization profits over time




Figure 15: Verification permutation two (medium number of Widgets) average organization profits over time




Figure 16: Verification permutation three (high number of Widgets) average organization profits over time

Figure 17, Figure 18, and Figure 19 chart the average price of Widgets that were resold to the Market Environment. These charts show two significant happenings. First, it can be observed that generally, the price of resold Widgets in the Market Environment decreases over time. This supports the verification of the free market economic functions of the computer model. In a free market, the result of different firms competing in the same industry leads to lower prices for that industry (Sowell, 2007).

The second observation is not very easy to determine without further analysis, but it appears that the rate of price decrease is diminishing as more Widgets are introduced into the Market Environment. This might be due to the more fierce competition in an environment where resources are scarcer (Figure 17). In the third permutation (Figure 19), where 1,000 more Widgets per Day are introduced, the Organizations were able to purchase Raw Widgets cheaper and were therefore less desperate to recover losses.




Figure 17: Verification permutation one (low number of Widgets) average prices of resold Widgets over time

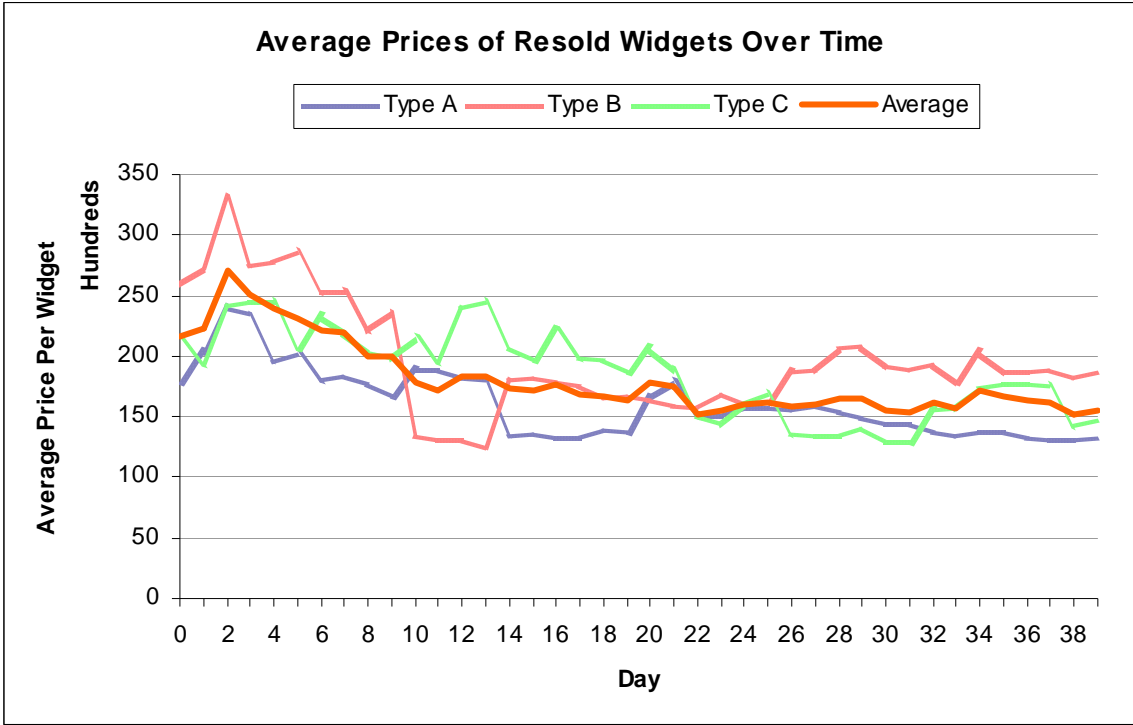


Figure 18: Verification permutation two (medium number of Widgets) average prices of resold Widgets over time




Figure 19: Verification permutation three (high number of Widgets) average prices of resold Widgets over time

3. Demonstration of Genetic Algorithm Behavior in the Computer Model of this Thesis

Indication of a working GA can be observed in Figure 20, Figure 21, and Figure 22. In GAs, the first generations result in very high numbers of changes to the binary strings of the individual. This is due to the initial randomization of individual binary strings. Individuals must quickly find ways to survive, and therefore there is a large amount of changing across the average individual. Over time, individuals in a gene pool will begin to converge and become more homogenous, therefore reducing the number of differences between the individuals each generation.

Since a change was made to the MAS, significant differences between Figure 20, Figure 21, and Figure 22 were not expected. This supports a functional and valid GA.




Figure 20: Verification permutation one (low number of Widgets) average Organizational policy changes over time




Figure 21: Verification permutation two (medium number of Widgets) average Organizational policy changes over time




Figure 22: Verification permutation three (high number of Widgets) average Organizational policy changes over time

C. EXPERIMENTATION

1. Overview

Three experiments were conducted. The first experiment ran the computer model with parameters that prevented workers from being monitored by the organizations to test Grover’s hypothesis. The second experiment was identical to the first except that the organization feedback mechanism is varied to test the Davis Hypothesis. The third experiment ran the model without these two constraints, which served as a comparison to a situation where behavior could be observed. Each experiment was run for 10 trials.

a. *Experiment One: Manipulation of the Organizations’ Attribution Resolution*

In an unaltered run of the computer model, each Organization would be able to set its own Attribution Resolutions, which would determine how precisely a

penalty was awarded upon the finding of a Falsified Widget. A higher resolution factor would result in lower precision. For example, if Worker #5 Falsifies a Widget and the resolution is two, then not only is Worker #5 penalized, but so are Worker #4 and Worker #6. The penalty is also reduced by a factor of three in this instance. If the resolution is one, then only Worker #5 is penalized, and without reduction. A resolution of zero meant that no Worker would be penalized for lying.

In this experiment, however, the Organizations' Attribution Resolution was forced to be zero. An Attribution Resolution equal to zero meant that a Worker could never be punished for lying.

Table 12. Experiment one Market Environment attribute values

Market Environment Attribute	Values
Widget Renew Rates	A: 600 B: 600 C: 600
Widget Max Consumption	A: 0 B: 0 C: 0
Widget Max Spoil	A: 0 B: 0 C: 0
Inspection Costs	A: 250 B: 250 C: 250
Consumer p(detection)	A: 0.5 B: 0.5 C: 0.5
Attribution Resolution Factors	0:1; 1:4.5; 2:4; 3:3.5; 4:3; 5:2.5; 6:2; 7:1.5

Table 13. Experiment one manipulated Organization variable

Organization Attribute	Value
Attribution Resolution	A: 0 B: 0 C: 0

b. Experiment Two: Manipulation of the Market Environment's Consumer p(detection)

The feedback provided to Organizations that sold Falsified Widgets was determined in the Market Environment's consumer detection phase (see Figure 7). Without this feedback, Organizations might not have been motivated to reduce the production and resale of Falsified Widgets.

This experiment manipulated the Market Environment's Consumer p(detection) attribute between zero, 50%, and 100%. This effectively provided different magnitudes of feedback to the organizations in order to test the Davis Hypothesis.

Table 14. Experiment two Market Environment variables

Market Environment Attribute	Values
Widget Renew Rates	A: 600 B: 600 C: 600
Widget Max Consumption	A: 0 B: 0 C: 0
Widget Max Spoil	A: 0 B: 0 C: 0
Inspection Costs	A: 250 B: 250 C: 250
Attribution Resolution Factors	0:1; 1:4.5; 2:4; 3:3.5; 4:3; 5:2.5; 6:2; 7:1.5

Table 15. Experiment two independent Market Environment variables

Market Environment Attribute	Perm. 1	Perm. 2	Perm. 3
Consumer p(detection)	A: 0 B: 0 C: 0	A: 0.5 B: 0.5 C: 0.5	A: 1 B: 1 C: 1

c. Experiment 3: Unaltered Run

The third experiment did not manipulate or vary any attributes. The same Market Environment attributes were used as were used in experiment one. This experiment allowed for the observation of an environment opposite to that of the Grover model and where Workers could not be caught and punished for lying. In this experiment, Workers could be caught and they could be punished for lying, although the detection of a liar cost the Organizations money according to the inspection/attribution cost described in Chapter III.

Table 16. Experiment three Market Environment attribute values

Market Environment Attribute	Values
Widget Renew Rates	A: 600 B: 600 C: 600
Widget Max Consumption	A: 0 B: 0 C: 0
Widget Max Spoil	A: 0 B: 0 C: 0
Inspection Costs	A: 250 B: 250 C: 250
Consumer p(detection)	A: 0.5 B: 0.5 C: 0.5
Attribution Resolution Factors	0:1; 1:4.5; 2:4; 3:3.5; 4:3; 5:2.5; 6:2; 7:1.5

2. Data Analysis

Each of the experiments produced data of when Workers lied and the skill levels of the Workers. The decisions were pooled into two categories: 1) Lied or 2) Did not lie. Lies were counted from the successful attempts at Falsifying Widgets³⁷. Counts for category two (Did not lie) were drawn from the remainder of the Worker's actions.

Data from each experiment was made into a contingency table and a Chi-Squared test was run to check for independence with an alpha of 0.01 ($\alpha=0.01$). The data was plotted against Worker skill level and against time to show strength and direction of relationships.

a. Results of Experiment One: Manipulation of the Organizations' Attribution Resolution

This experiment prevented Organizations from detecting or punishing Workers for deception. The Chi-Squared null hypothesis for this experiment was that Workers' skill levels and their decisions to 1) lie or 2) not lie were independent. The alternative hypothesis was that Workers' skill levels and their decisions to lie or not lie were related.

The Chi-Squared test of the experiment one contingency table (Table 19) provides strong evidence that a Worker's choice to lie or not lie is related to the Worker's skill level when Organizations could not catch or punish lying Workers. The Chi-Squared critical value was 21.666 and the Chi-Squared statistic was 3737.377. The resultant p-value was less than 0.001 (p-value < 0.001). The null hypothesis was rejected.

³⁷ A successful attempt is defined as an instance where the Worker's intention to deceive resulted in a Widget being changed to the Falsified state.

Table 17. Contingency table and Chi-Squared test for experiment one

Contingency Table

Skill Level	Lied	Did not Lie	TOTAL
0	24043	123165	147208
0.11	30445	122109	152554
0.22	28110	120284	148394
0.33	23937	126010	149947
0.44	26013	119343	145356
0.56	23734	114810	138544
0.67	25809	115015	140824
0.78	22437	119399	141836
0.89	18796	124231	143027
1	23355	129437	152792
TOTAL	246679	1213803	1460482

chi-squared Stat	3737.377
df	9
p-value	0
chi-squared Critical	21.666

The weak negative relationship between skill level and Worker’s decisions to lie observed in Table 19 can more easily be seen in Figure 23:




Figure 23: Results of experiment one, number of lies told vs. Worker skill level




Figure 24: Results from experiment one, number of truths told vs. Worker skill level




Figure 25: Results from experiment one, proportion of truths over lies told vs. Worker skill level

Over time, Workers lied less frequently. This was preceded by a brief initial trend of lying, as depicted in Figure 26. An average frequency (skill-independent) is provided in Figure 27.




Figure 26: Number of times Workers lied over time, by skill level in experiment one

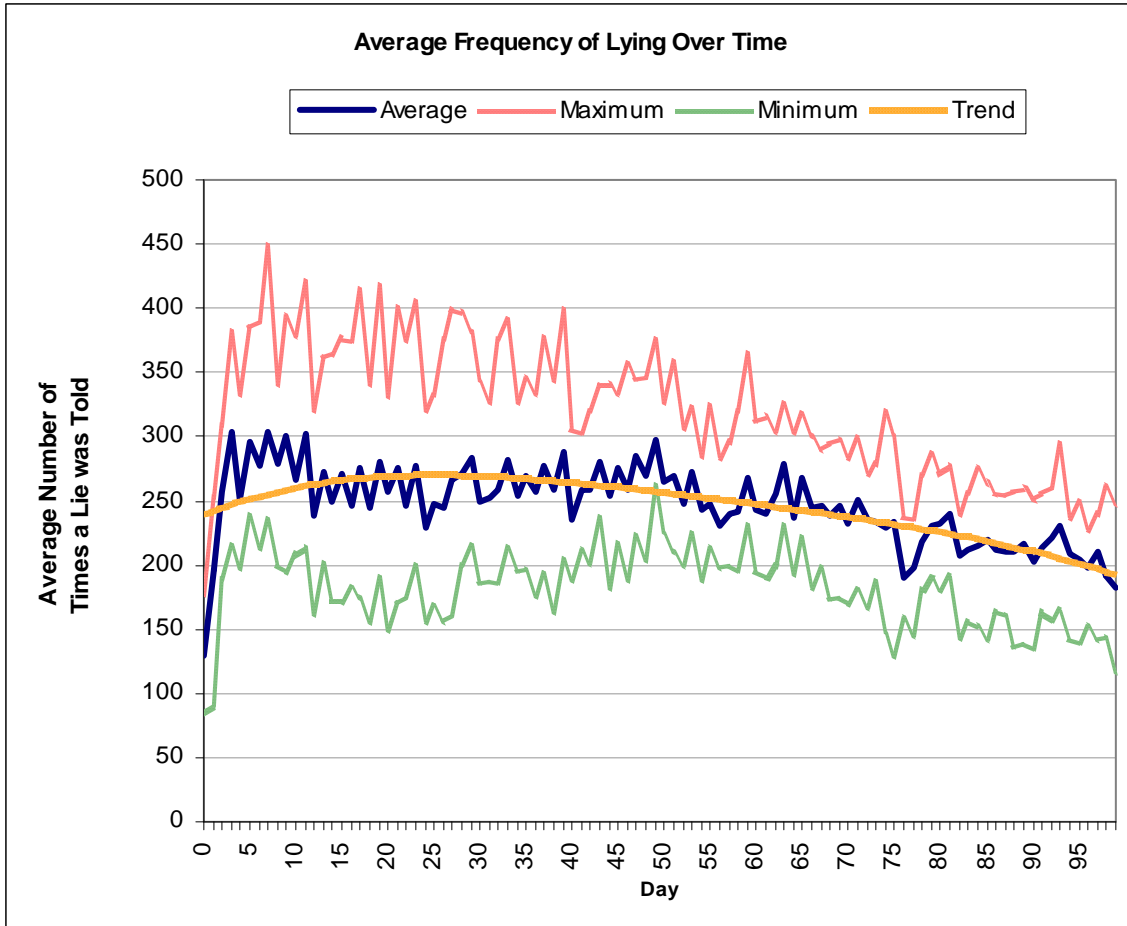


Figure 27: Average number of times Workers lied over time (averaged by skill level) in experiment one

b. Results of Experiment Two: Manipulation of the Market Environment's Consumer $p(\text{detection})$

This experiment manipulated the amount of feedback that was given to Organizations. This experiment also prevented Organizations from detecting or punishing Workers for deception. Because this experiment was very similar to experiment one and because similar data was collected, data applicable to experiment one will also be presented in Appendix B.

The first permutation of this experiment set the ability of the Market Environment to give Organizations negative feedback to zero. After data collection, it was realized that the Davis hypothesis required a condition where non-zero negative

feedback was administered to the Organizations. Therefore, the data collected from this permutation was not used to test the Davis hypothesis. The results of this permutation were interesting; however, as there appeared to be a trend towards honesty as time progressed. These results are provided in Appendix C.

In this experiment, the Chi-Squared null hypothesis was that Workers' decisions to lie or not lie were independent of the Market Environment's Consumer p(detection) attribute magnitude. The alternative hypothesis was that Workers' decisions to lie or not lie were related of the Market Environment's Consumer p(detection) attribute magnitude.

The Chi-Squared test of the experiment two contingency table (Table 20) provides strong evidence that a Worker's choice to lie or not lie was related to the Worker's skill level when Organizations were allowed to catch and penalize lying Workers. The Chi-Squared critical value was 6.634897 and the Chi-Squared statistic was 35915.7. The resultant p-value was less than 0.001 (p-value < 0.001). The null hypothesis was rejected.

Table 18. Contingency table and Chi-Squared test for experiment two

Contingency Table

p(detection)	Lied	Did not Lie	TOTAL
0.5	246679	1213803	1460482
1	122564	1209705	1332269
TOTAL	369243	2423508	2792751
chi-squared Stat			35915.7
df			1
p-value			0
chi-squared Critical			6.634897

A strong negative relationship between the number of lies told and the magnitude of feedback is visible in Figure 28.




Figure 28: Results from experiment two, number of lies told vs. Market Environment feedback magnitude

There did not appear to be a relationship between the number of truths told and the Market Environment Consumer $p(\text{detection})$ attribute value in Figure 29.




Figure 29: Results from experiment two, number of truths told vs. Market Environment feedback magnitude




Figure 30: Results from experiment two, proportion of truths told over lies told vs. Market Environment feedback magnitude

c. Results of Experiment Three: Unaltered Run

This experiment made no alterations to the computer model. Organizations were allowed to attempt to detect deception and punish Workers accordingly.

The Chi-Squared null and alternative hypotheses for this experiment were identical to the null and alternative hypotheses in experiment one.

The Chi-Squared test of the experiment 3 contingency table (Table 21) provides strong evidence that a Worker’s choice to lie or not lie was related to the Worker’s skill level when Organizations were allowed to catch and penalize lying Workers. The Chi-Squared critical value was 21.666 and the Chi-Squared statistic was 3226.733. The resultant p-value was less than 0.001 (p-value < 0.001). The null hypothesis was rejected.

Table 19. Contingency table and Chi-Squared test for experiment three

Contingency Table

Skill Level	Lied	Did not Lie	TOTAL
0	21600	154550	176150
0.11	18819	155462	174281
0.22	17031	126315	143346
0.33	15301	138032	153333
0.44	15232	135209	150441
0.56	14505	151525	166030
0.67	14795	130030	144825
0.78	15195	157390	172585
0.89	13305	149960	163265
1	11205	120613	131818
TOTAL	156988	1419086	1576074
chi-squared Stat			3226.733
df			9
p-value			0
chi-squared Critical			21.666

The weak negative relationship between skill level and Worker's decisions to lie observed in Table 21 can more easily be seen in Figure 31:




Figure 31: Results of experiment three, number of lies told vs. Worker skill level




Figure 32: Results of experiment three, number of truths told vs. Worker skill level




Figure 33: Results from experiment three, proportion of truths over lies told vs. Worker skill level

Over time, Workers lied less frequently. This was preceded by a brief initial trend of lying, as depicted in Figure 34. An average frequency (skill-independent) is provided in Figure 35.




Figure 34: Number of times Workers lied over time, by skill level in experiment three




Figure 35: Average number of times Workers lied over time (averaged by skill level) in experiment three

V. CONCLUSIONS

A. SUMMARY

1. Experiment One: Manipulation of the Organizations' Attribution Resolution

The results of this experiment support Grover's hypothesis 3a: *Individual job skill will be negatively associated with lying to agents of the organization.* Workers in the computer model lied less frequently as their skill increased, and the null hypothesis of independence was rejected.

2. Experiment Two: Manipulation of the Market Environment's Consumer p(detection)

The results of this experiment support the Davis Hypothesis: *In an environment where 1) an organization and its members are independently self-interested, 2) the organization receives merited negative feedback from its environment, and 3) workers cannot be observed or punished for lying, the frequency of organization members lying will be inversely proportional to the magnitude of that feedback.* Workers in the computer model lied less frequently as their Organizations were subjected to higher levels of negative feedback for producing Falsified Widgets, and the null hypothesis of independence was rejected.

The results of this experiment also support Grover's hypothesis 3a. Workers in the computer model lied less frequently as their skill increased.

3. Experiment Three: Unaltered Run

The results of this experiment support Grover's hypothesis 3a in an environment where organizations can detect and punish liars. This suggests that the predictions of Grover's 3a hypothesis might be independent of an organization's negative actions towards liars. The null hypothesis of independence was rejected in this experiment.

B. RECOMMENDATIONS

The Davis hypothesis needs further testing. Time constraints prevented more permutations to be examined, and data should be collected for longer time periods to check for oscillations. While the statistical analysis supported the Davis hypothesis, more data points need to be tested to check for a non-linear relationship.

The computer model in this thesis only tested one of Grover's ten hypotheses. The biggest limitation of this computer model was that peer agents were not allowed to interact with each other, i.e. Worker agents could not interact with other Worker agents. To explore many of Grover's hypotheses, a computer model that allows for peer interaction would be needed.

Other models that might attempt to study this subject would probably not need such a complicated model. The use of the GA on the behavior adaptation of the agents in this model added a great deal of complexity, but a good approach might also be agents whose behaviors were adapted by a simple scoring mechanism. Each agent could have three possible behaviors: lie, be honest, or do nothing, and the behavior with the highest score would be the behavior that the agent actually performed. A difficulty with this approach, however, is that in order to score behaviors that were not actually performed, the agent would be required to make some hypothetical determination of what the score would have been, had he been acting on the other possible behavior. That is why the GA is so useful: it does not require a determination of what could have been, only on what worked or did not work.

In order to capture the emergent behaviors that were observed, such as the average reduction of deceivers, even when no direct feedback occurs, the computer model must maintain a certain level of complexity. This observation might not have been made if, in this computer model, only the Worker agents were allowed to adapt. On the other hand, an interesting experiment would be to test if a co-evolving hierarchical architecture is required to achieve this behavior. Much of the other research with social rationality has been done with simple agents in simple environments.

The luxury of working with a computer model is that it never forgets. Many times during development and testing, unexpected data would be returned. The first instinct is to blame the computer. When development was approached carefully and incrementally, however, most of the time these unexpected returns were found to be correct and logical. The bug in the loop was human, not computer. On the other hand, the pains of working with computers are also present. Any software developer knows the agony of wasting hours to find a single typographical error.

THIS PAGE INTENTIONALLY LEFT BLANK

APPENDIX A: OTHER VERIFICATION CHARTS

Figure 36, Figure 37, and Figure 38 plot the same data as was presented in the Verification section, but here the data is not averaged.




Figure 36: Verification permutation one average Organization profits by cluster over time




Figure 37: Verification permutation two average organization profits by cluster over time




Figure 38: Verification permutation three average organization profits by cluster over time

The average Worker fitness over time (Figure 39, Figure 40, and Figure 41) do not show much change because Workers were somewhat insulated from the Market Environment.




Figure 39: Verification permutation one average Worker fitness over time




Figure 40: Verification permutation two average Worker fitness over time




Figure 41: Verification permutation three average Worker fitness over time

Widget demand was proportional to the number of Widgets that were available for resale.




Figure 42: Verification permutation one average consumer demand of resold Widgets over time




Figure 43: Verification permutation two average consumer demand of resold Widgets over time




Figure 44: Verification permutation three average consumer demand of resold Widgets over time

**APPENDIX B: DATA GATHERED FROM EXPERIMENT TWO
THAT IS APPLICABLE TO EXPERIMENT ONE**

The following data is applicable to experiment one and was collected for the permutation where the Market Environment’s Consumer p(detection) was equal to one. The Chi-Squared test of the experiment two, permutation three contingency table (Table 22) provides strong evidence that a Worker’s choice to lie or not lie is related to the Worker’s skill level when Organizations cannot catch or punish lying Workers. The Chi-Squared critical value was 21.666 and the Chi-Squared statistic was 689.6626. The resultant p-value was less than 0.001 (p-value < 0.001).

Table 20. Contingency table and Chi-Squared test for experiment two, permutation three

Contingency Table

Skill Level	Lied	Did not Lie	TOTAL
0	13751	130323	144074
0.11	12110	104315	116425
0.22	11817	114926	126743
0.33	11540	126120	137660
0.44	11069	117615	128684
0.56	12786	115810	128596
0.67	13826	130523	144349
0.78	12830	129315	142145
0.89	11818	116270	128088
1	11017	124488	135505
TOTAL	122564	1209705	1332269
chi-squared Stat			689.6626
df			9
p-value			0
chi-squared Critical			21.666



Figure 45: Number of lies told vs. Worker skill level, experiment two, permutation three




Figure 46: Number of truths told vs. Worker skill level, experiment two, permutation three







Figure 47: Proportion of truths over lies told vs. Worker skill level, experiment two, permutation three


THIS PAGE INTENTIONALLY LEFT BLANK


APPENDIX C: DATA GATHERED FROM EXPERIMENT TWO, PERMUTATION ONE


Since this data tested a situation where negative feedback was not provided to the Organizations, it did not apply to the Davis hypothesis and was therefore not included in the testing of the Davis hypothesis. This data still showed trends of decreasing frequency of deceit, however.

















This chart shows that truths became more prevalent, but data might need to be collected for a longer time period to determine if this trend is oscillating around a constant value.

THIS PAGE INTENTIONALLY LEFT BLANK

LIST OF REFERENCES

- Arthur, W. B. (1994). Inductive reasoning and bounded rationality (the el Farol Problem). *American Economic Review (Papers and Proceedings)*, 406.
- Axelrod, R. (1997). *The Complexity of Cooperation*. Princeton, New Jersey: Princeton University Press.
- Carson, T. L. (2003). **Some Lessons of the Recent Corporate Scandals** [Abstract]. *Journal of Business Ethics*, 43(4) 389.
- Chandler, R. E., Herman, R., & Montroll, E. W. (1958). Traffic Dynamics: Studies in Car Following. *Operations Research*, 6(2), 165-184.
- Cunningham, W. P. (1998). The Golden Rule as Universal Ethical Norm. *Journal of Business Ethics*, 17(1), 105-109.
- DePaulo, B. M., Kirkendol, S. E., Kashy, D. A., Epstein, J. A., & Wyer, M. M. (1996). Lying in Everyday Life. *Journal of Personality and Social Psychology*, 70(5), 979.
- Edmonds, B. (1999a). Modelling bounded rationality in agent-based simulations using the evolution of mental models. *ADVANCES IN COMPUTATIONAL ECONOMICS*, 11, 305.
- Edmonds, B. (1999b). **Gossip, Sexual Recombination and the el Farol Bar: Modelling the Emergence of Heterogeneity**. *Journal of Artificial Societies and Social Simulation*, 2(3)
- Eisenhardt, K. M. (1989). Agency Theory: An Assessment and Review. *The Academy of Management Review*, 14(1), 57-74.
- Grover, S. L. (1993). Lying, Deceit, and Subterfuge: A Model of Dishonesty in the Workplace. *Organization Science*, 4(3, Focused Issue: The Legalistic Organization), 478-495.
- Grover, S. L. (1996). Lying in Organizations. *Antisocial Behavior in Organizations* () Sage Publications.
- Grover, S. L., & Chun Hui. (1994). The Influence of Role Conflict and Self-interest on Lying in Organizations. *Journal of Business Ethics*, 13(4), 295-303.
- Hales, D., & Edmonds, B. (2003). Evolving Social Rationality for MAS using "tags." *AAMAS '03: Proceedings of the second international joint conference on autonomous agents and multiagent systems* (497). New York: ACM.

- Haupt, R. L. (2000). *Optimum population size and mutation rate for a simple real genetic algorithm that optimizes array factors*
- Holland, J. H. (1992). Computer programs that "evolve" in ways that resemble natural selection can solve complex problems even their creators do not fully understand. *Scientific American*,
- Jansen, E., & Mary Ann Von Glinow. (1985). Ethical ambivalence and organizational reward systems. *The Academy of Management Review*, 10(4), 814-822.
- Krippendorff, K. (1986). **GAME THEORY**. Retrieved June 16, 2008, from http://134.184.131.111/asc/GAME_THEOR.html
- Pfaff, D. W. (2007). *The Neuroscience of Fair Play: Why we (usually) follow the golden rule*. New York: Dana Press.
- Price, J. S. (1996). Lies! Lies!! Lies!!!: The psychology of deceit. *Nature Medicine*, 2(7), 821.
- Solomon, M. R., Surprenant, C., Czepiel, J. A., & Gutman, E. G. (1985). A role theory perspective on dyadic interactions: The service encounter. *Journal of Marketing*, 49(1), 99-111.
- Sowell, T. (2007). *Basic Economics: A common sense guide to the economy* (Third ed.). New York: Basic Books.
- Taber, C. S., & Timpone, R. J. (1996). *Computational Modeling*. Thousand Oaks, Calif.: Sage Publications.

INITIAL DISTRIBUTION LIST

1. Defense Technical Information Center
Ft. Belvoir, Virginia
2. Dudley Knox Library
Naval Postgraduate School
Monterey, California
3. Lt. Col. Karl Pfeiffer
Naval Postgraduate School
Monterey, California
4. Mr. Steven Iatrou
Naval Postgraduate School
Monterey, California
5. Mr. John Hiles
Naval Postgraduate School
Monterey, California
6. Mr. Steven Grover
New Zealand
7. Mr. Dan Boger
Naval Postgraduate School
Monterey, California
8. Mr. Don Nelson
United States Naval Academy
Annapolis, Maryland