

# **Belief Propagation Analysis in Two-Player Games for Peer-Influence Social Networks**

by

Matthew E Bradwick

B.S. Operations Research  
United States Air Force Academy, 2010

SUBMITTED TO THE SLOAN SCHOOL OF MANAGEMENT IN PARTIAL  
FULFILLMENT OF THE REQUIREMENTS OF THE DEGREE OF

MASTER OF SCIENCE IN OPERATIONS RESEARCH

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2012

©2012 Matthew E Bradwick. All rights reserved

The author hereby grants to MIT and Draper Laboratory permission to reproduce and to distribute publicly paper and electronic copies of this thesis document in whole or in part in any medium now known or hereafter created.

Signature of Author: \_\_\_\_\_

Sloan School of Management  
Interdepartmental Program in Operations Research  
May 18, 2012

Approved by: \_\_\_\_\_

Dr. Natasha Markuzon  
The Charles Stark Draper Laboratory, Inc.  
Technical Supervisor

Certified by: \_\_\_\_\_

Professor Marta C. González  
Gilbert Winslow Career Development Assistant Professor  
Civil & Environmental Engineering; Engineering Systems Division  
Thesis Advisor

Accepted by: \_\_\_\_\_

Professor Dimitris Bertsimas  
Boeing Professor of Operations Research  
Co-Director, Operations Research Center

[This Page Intentionally Left Blank]

# **Belief Propagation Analysis in Two-Player Games for Peer-Influence Social Networks**

by

Matthew E Bradwick

Submitted to the Sloan School of Management on May 18, 2012 in Partial Fulfillment of the Requirements for the Degree of Master of Science in Operations Research

## **ABSTRACT**

This thesis considers approaches to influencing population opinions during counterinsurgency efforts in Afghanistan. A discrete time, agent-based threshold model is developed to analyze the propagation of beliefs in the social network, whereby each agent has a belief and a threshold value, which indicates the willingness to be influenced by the peers. Agents communicate in stochastic pairwise interactions with their neighbors. A dynamic, two player game is formulated whereby each player strategically controls the placement of one stubborn agent over time in order to maximally influence the network according to one of two different payoff functions. The stubborn agents have opposite, immutable beliefs and exert significant influence in the network. We demonstrate the characteristics of strategies chosen by the players to improve their payoffs through simulation. Determining strategies for the players in large, complex networks in which each stubborn agent has multiple connections is difficult due to exponential increases in the strategy space that is searched. We implement two heuristic methods which are shown to significantly reduce the run time needed to find strategies without significantly reducing the quality of the strategies. Lastly, we introduce population-focused actions, such as economic stimulus projects, which when used by the players result in long-lasting changes in the beliefs of the agents in the network.

Thesis Advisor: Professor Marta C. González  
Gilbert Winslow Career Development Assistant Professor  
Civil & Environmental Engineering; Engineering Systems Division

Technical Supervisor: Dr. Natasha Markuzon  
Principal Member of the Technical Staff  
The Charles Stark Draper Laboratory, Inc.

[This Page Intentionally Left Blank]

## Acknowledgments

I would like to thank Dr. Natasha Markuzon of Draper Laboratory for her guiding mentorship over these past two years. Her continued patience, broad expertise, and constant support proved invaluable in helping me develop this thesis.

I would like to thank Professor Marta González of MIT for accepting me as one of her students and offering constructive feedback to improve this work.

I would like to thank Dr. Stephan Kolitz of Draper Laboratory for his time and vast knowledge in the field of operations research which helped me further this research.

I would like to thank Draper Laboratory for the fellowship which funded my graduate degree as well as the United States Air Force Academy for allowing me to pursue a graduate degree these last two years at MIT.

I would like to thank some of my fellow students for all the help on the problem sets we completed during our time at MIT: Brian Crimmel, Jane Evans, Mike Giordano, and Andrew Lee. I learned a lot from each of you and greatly appreciate the time you spent helping me on our class work.

Finally, I would like to thank my family and girlfriend, Andrea, for their continuous support. You were the source of greatly needed encouragement, and your belief in my abilities maintained my motivation and focus to keep pushing even when things looked grim.

The views expressed in this thesis are those of the author and do not reflect the official policy or position of the United States Air Force, Department of Defense, or the U.S. Government.

# TABLE OF CONTENTS

<i>List of Figures</i> .....	8
<i>List of Tables</i> .....	10
<b>1 Introduction</b> .....	<b>13</b>
<b>1.1 Motivation for Research</b> .....	<b>13</b>
<b>1.2 Thesis Organization</b> .....	<b>13</b>
<b>2 The Battle for Popular Support between Insurgencies and Counterinsurgencies</b> .....	<b>17</b>
<b>2.1 The Nature of Insurgencies</b> .....	<b>18</b>
<b>2.2 The Nature of Counterinsurgencies</b> .....	<b>21</b>
<b>2.3 Nonlethal Targeting used to Win Popular Support</b> .....	<b>22</b>
<b>3 Modeling Approach and Formulation</b> .....	<b>27</b>
<b>3.1 Literature Review</b> .....	<b>27</b>
3.1.1 Influential Leaders in Social Networks.....	27
3.1.2 Opinion Dynamics .....	28
3.1.3 The Need for a More Realistic Model .....	31
3.1.4 Cascade and Threshold Models .....	33
3.1.5 Transient Behavior of Belief Propagation in Social Networks .....	34
<b>3.2 Proposed Threshold Network Model Formulation</b> .....	<b>35</b>
3.2.1 Scope of the Model .....	35
3.2.2 Network Characteristics.....	36
3.2.3 Agents Characteristics .....	36
3.2.4 Pairwise Interactions.....	41
3.2.5 Formulation of the Threshold Model.....	44
<b>3.3 Transient, Dynamic Two-Player Game Description</b> .....	<b>51</b>
3.3.1 Assumptions .....	51
3.3.2 Dynamic Game Parameters.....	52
3.3.3 Payoff Function .....	56
3.3.4 Finding Strategies to Improve Payoff .....	61
<b>3.4 The Necessity of Simulations in Calculating Belief Propagation</b> .....	<b>69</b>

3.5	<b>Modeling Formulation Summary .....</b>	<b>69</b>
<b>4</b>	<b><i>Experiments and Analysis .....</i></b>	<b>72</b>
4.1	<b>Results .....</b>	<b>74</b>
4.1.1	Characteristics of Strategies.....	74
4.1.2	How Different Payoff Functions Influence Strategies .....	89
4.1.3	Sensitivity to Model Parameters .....	93
4.2	<b>Analysis of ‘Selective Search’ and ‘Sequential Search’ Heuristics .....</b>	<b>102</b>
4.2.1	Run Time Comparison: Exhaustive Enumeration vs. ‘Selective Search’ vs. ‘Sequential Search’ .....	102
4.2.2	Payoff Maximization: Exhaustive Enumeration vs. ‘Selective Search’ vs. ‘Sequential Search’ .....	104
4.3	<b>Population-Focused Actions Which Affect Belief Propagation .....</b>	<b>110</b>
4.3.1	U.S. Stimulus Projects .....	110
4.3.2	Assassinations.....	123
4.4	<b>Summary of Experiments.....</b>	<b>129</b>
<b>5</b>	<b><i>Future Work Recommendations .....</i></b>	<b>133</b>
5.1	<b>Dynamic Networks .....</b>	<b>133</b>
5.2	<b>Imperfect Information Games .....</b>	<b>134</b>
5.3	<b>Tipping Points.....</b>	<b>134</b>
5.4	<b>Applying Iterated Reasoning to Predict Opponent Strategies.....</b>	<b>137</b>
5.5	<b>Additional Population-Focused Actions Which Affect Belief Propagation .....</b>	<b>138</b>
<b>6</b>	<b><i>Summary and Conclusions .....</i></b>	<b>140</b>
	<b><i>Appendix A: Variable Notation and Definitions .....</i></b>	<b>145</b>
	<b><i>Appendix B: Experimental Data and Results .....</i></b>	<b>147</b>
	<b><i>References .....</i></b>	<b>152</b>

## List of Figures

Figure 2.1 – Spectrum of Popular Support [5].....	23
Figure 3.1 – Agent Belief Spectrum .....	37
Figure 3.2 – Example Network Diagram and Symbol Notation.....	40
Figure 3.3 – Alpha, Beta, and Gamma Default Values .....	43
Figure 3.4 – Step 0 (Initial Strategy) and Step 1 (First Strategy) Example .....	54
Figure 3.5 – Non-penalized and Penalized $\alpha$ , $\beta$ , and $\gamma$ Matrices.....	55
Figure 3.6 – Absolute Payoff vs. Relative Payoff Example .....	59
Figure 3.7 – Line Graph Example with Regular Agent Chosen as Strategy .....	64
Figure 3.8 – Example 73 Agent Network .....	68
Figure 4.1 – Experimental Network Diagrams .....	76
Figure 4.2 – Experiment 5: First Chosen U.S. Strategy.....	80
Figure 4.3 – Experiment 5: First Chosen Taliban Strategy .....	81
Figure 4.4 – Experiment 5: Before Second Chosen Taliban Strategy .....	82
Figure 4.5 – Experiment 5: After Second Chosen Taliban Strategy.....	83
Figure 4.6 – Experiment 5: End Game Network Beliefs .....	84
Figure 4.7 – Ending Network Diagram: Mutable Thresholds = 0 Example.....	85
Figure 4.8 – Ending Network Diagram: Mutable Thresholds = 1 Example .....	86
Figure 4.9 – Payoff Function Comparison: Mean Belief and Number of Agents Won Payoff Plots .....	92
Figure 4.10 – Payoff Function Comparison: Payoff Plots for 20-Step Experiment 2 .....	92
Figure 4.11 – Different Society Types ( $\alpha$ , $\beta$ , and $\gamma$ values): Mean Belief Payoff Plot.....	96
Figure 4.12 – Stimulus Project Default Belief Shocks .....	112
Figure 4.13 – Ending Network Diagram: Stimulus Project Belief Propagation Example.....	116
Figure 4.14 – Ending Network Diagram: No Stimulus Project Example.....	116
Figure 4.15 – Relative Payoff Comparison over Time: Stimulus vs. No Stimulus .....	117
Figure 4.16 – Strategic Planning with Limited Resources: Experiment Set 2 Payoff Plots .....	122
Figure 4.17 – Ending Network Diagram: Assassination Belief Propagation Example .....	126
Figure 4.18 – Ending Network Diagram: No Assassination Project Example .....	127
Figure 4.19 – Relative Payoff Comparison over Time: Assassination vs. No Assassination ....	128

Figure 5.1 – Tipping Point: Network Diagram..... 135  
Figure 5.2 – Tipping Point: Linear Relationship Plot..... 137

## List of Tables

Table 3.1 – Variable Notation Used In Threshold Model Simulation .....	46
Table 3.2 – Number of Agents Won Point System.....	57
Table 3.3 – Example Strategy Profile and Payoff Table .....	60
Table 3.4 – Strategy Space Comparison: Exhaustive Enumeration vs. ‘Selective Search’ .....	65
Table 4.1 – Default Experimental Parameters .....	73
Table 4.2 – Characteristics of Strategies: Experiment Set Description .....	74
Table 4.3 – Characteristics of Strategies: Strategy Profile and Payoff Table Using Howard’s Model .....	77
Table 4.4 – Characteristics of Strategies: Experiments 5 and 6 Strategy Profile and Payoff Tables Using the Proposed Threshold Model.....	78
Table 4.5 – Observation #4: Experiment 1 Strategy Profile and Payoff Table .....	87
Table 4.6 – Observation #4: Experiment 2 Strategy Profile and Payoff Table .....	88
Table 4.7 – Payoff Function Comparison: Experiment Descriptions .....	89
Table 4.8 – Payoff Function Comparison: Mean Belief Strategy Profile and Payoff Tables.....	90
Table 4.9 – Payoff Function Comparison: Number of Agents Won Point System .....	90
Table 4.10 – Payoff Function Comparison: Number of Agents Won Strategy Profile and Payoff Tables .....	91
Table 4.11 – Alpha, Beta, Gamma Values for Different Society Types .....	94
Table 4.12 – Different Society Types ( $\alpha$ , $\beta$ , and $\gamma$ values): Experiment Set Description .....	95
Table 4.13 – Different Society Types ( $\alpha$ , $\beta$ , and $\gamma$ values): Effect of Evolution of Beliefs .....	96
Table 4.14 – $\mu$ Sensitivity: Experiment Set Description .....	97
Table 4.15 – $\mu$ Sensitivity: Experiment Set Payoff Table .....	98
Table 4.16 – Initial (Step 0) Strategy for Stubborn Agent: Experiment Set Description .....	99
Table 4.17 – Initial (Step 0) Strategy for Stubborn Agent: Strategy Profile and Payoff Tables ..	99
Table 4.18 – Penalty for Identical Strategy ( $\lambda$ ): Experiment Set Description .....	100
Table 4.19 – Penalty for Identical Strategy ( $\lambda$ ): Experiments 1 and 2 Strategy Profile and Payoff Tables .....	101
Table 4.20 – Run Time Approximation Comparison Table .....	104
Table 4.21 – Payoff Maximization: Experiment 1 Description .....	105

Table 4.22 – Payoff Maximization: Experiment 1 Exhaustive Enumeration and ‘Selective Search’ Results.....	105
Table 4.23 – Payoff Maximization: Experiment 1 ‘Sequential Search’ Results .....	105
Table 4.24 – Payoff Maximization: Experiment 2 Description.....	106
Table 4.25 – Payoff Maximization: Experiment 2 Exhaustive Enumeration and ‘Selective Search’ Results.....	107
Table 4.26 – Payoff Maximization: Experiment 2 ‘Sequential Search’ Results .....	107
Table 4.27 – Payoff Maximization: Experiment 3 Description.....	108
Table 4.28 – Payoff Maximization: Experiment 3 Exhaustive Enumeration and ‘Selective Search’ Results.....	108
Table 4.29 – Payoff Maximization: Experiment 3 ‘Sequential Search’ Results .....	108
Table 4.30 – Payoff Maximization: Run Times.....	109
Table 4.31 – Phaseshift and Amplitude Sensitivity Experiment Description.....	113
Table 4.32 – Relative U.S. ‘Mean Belief’ Payoffs for Various Phaseshift and Amplitude Values .....	114
Table 4.33 – Stimulus Project Belief Propagation Experiment Description .....	115
Table 4.34 – Strategic Planning with Limited Resources: Experiment Set 1 Description .....	118
Table 4.35 – Strategic Planning with Limited Resources: Experiment Set 1 Case Description	119
Table 4.36 – Strategic Planning with Limited Resources: Experiment Set 1 Payoff Table .....	119
Table 4.37 – Strategic Planning with Limited Resources: Experiment Set 2 Description .....	121
Table 4.38 – Strategic Planning with Limited Resources: Experiment Set 2 Case Description	121
Table 4.39 – Strategic Planning with Limited Resources: Experiment Set 2 Payoff Table .....	122
Table 4.40 – Assassination Belief Propagation Experiment Description.....	126
Table 5.1 – Tipping Point: Results .....	136

[This Page Intentionally Left Blank]

# 1 Introduction

During the 20<sup>th</sup> century, nations on all continents have experienced or intervened in insurgencies. While several counterinsurgency (COIN) operations have proven successful, many have failed miserably. However, the strategic victory over an insurgency does not instantly validate the operational and tactical methods used by the counterinsurgents, nor does it make them applicable for universal use in all COIN operations [1]. Since 2001 the United States has been consumed with ongoing counterinsurgency operations, most notably in Afghanistan and Iraq. Due to the nature of insurgencies, whereby the insurgents limit the effectiveness of conventional warfare by blending in with the civilian population in order to protect themselves, counterinsurgents must adapt their traditional military tactics in order to successfully defeat the insurgents. Thus, substantial emphasis has been placed on winning the support of the civilian population in order to deny the ability for insurgents to have refuge amongst the masses.

## 1.1 Motivation for Research

The motivation for this research stems from the focus that is placed on winning popular support during counterinsurgency operations. The U.S. military is confronted with the difficult task of gaining the popular support of the populace during its COIN operations in Afghanistan and Iraq. Because the United States has limited resources and personnel, crucial decisions must be made by military commanders on which key civilian leaders to engage in an effort to optimally influence the entire civilian population. Therefore, this research analyzes the complex problem of choosing which key leaders to engage by implementing a dynamic two-player “chess” game on a social network, whereby the two players (US military and Taliban insurgents), who have immutable (stubborn) beliefs try to influence the beliefs of the mutable agents in the network toward supporting their cause.

## 1.2 Thesis Organization

The rest of this thesis is divided into five additional chapters. In Chapter 2 we discuss the nature of insurgencies and counterinsurgencies, and specifically, the inherent struggle to win popular support during counterinsurgency operations. In Chapter 3 we first discuss previous

work which has been devoted to the field of belief propagation in social networks. We present enhancements to previous modeling approaches by introducing the formulation of a peer-pressure threshold model and dynamic two-player game. In the network model, mutable agents are assigned a scalar belief between  $-0.5$  (heavily pro-Taliban/anti-US) and  $+0.5$  (heavily pro-US/anti-Taliban). The two stubborn agents (Taliban and US) have immutable beliefs of  $-0.5$  and  $+0.5$ , respectively. All agents in the network are assigned different levels of influence depending upon their position in society. The stubborn agents exert the most influence in the network as they try to influence the mutable agents in the network to adopt their extreme beliefs. Stochastic pairwise interactions occur between neighboring agents in the network, whereby influence is spread through the network as agents change their beliefs according to interaction-type probabilities and the characteristics (i.e. peer pressure effect) of the threshold model. Due to the presence of stubborn agents, which have immutable beliefs of different values, and the stochastic pairwise interactions, the beliefs of the mutable agents end up being random variables over the course of the discrete-time model.

During the two-player game, a US agent and a TB agent are concerned with finding connection strategies to mutable agents in the network in order to maximally influence the network toward support their cause. Both stubborn agents are given alternating opportunities during the game to reevaluate and update their strategy given the current strategy of their opponent and the beliefs of the mutable agents in the network. In Chapter 3 we also introduce two heuristic methods ('Selective Search' and 'Sequential Search') which drastically reduce the computational time needed to find solutions during the two-player game. We show that as the complexity of the problem increases (i.e. each player is allowed multiple strategy connections to agents in the network), the 'Sequential Search' method is necessary as exhaustive enumeration and 'Selective Search' do not scale well due to exponential increases in the potential strategy space.

In Chapter 4 we present key observations supported by empirical evidence about the characteristics of strategies chosen by the players to maximize their payoffs. Experiments also show the sensitivity of the solutions to various model parameters as well as different payoff functions used by the players. Moreover, we show that the use of 'Sequential Search' to find the best solutions for complex problems does not significantly sacrifice the quality of those obtained

solutions when comparing them to the solutions found by the exhaustive enumeration and ‘Selective Search’ methods. The last section of Chapter 4 is devoted to the introduction of population-focused actions available to the US and TB agents in order to influence the propagation of beliefs in the network. The US agent is allowed to conduct stimulus projects, such as building schools, roads, or wells, as well as giving money for economic projects, while the TB agent is allowed to use assassinations. Experimental evidence shows that the use of stimulus projects or assassinations can have long-lasting impacts on the network beliefs. These long-lasting impacts are due to the peer pressure effect which is built into the threshold model and makes agents highly unlikely to change their beliefs away from their neighbors’ beliefs. Finally, Chapter 5 is devoted to introducing interesting areas for future research while Chapter 6 summarizes the work presented in this thesis and offers some conclusions.

[This Page Intentionally Left Blank]

## **2 The Battle for Popular Support between Insurgencies and Counterinsurgencies**

The United States is known as having one of the most highly capable, conventional military forces in the world due to the presence of high-technology weapons and superior air, land, and sea forces. World War II and Operation Desert Storm showcase prime examples where the United States exerted conventional military tactics with much success. However, since the 1960s, the United States has been actively involved in supporting counterinsurgency efforts in countries all over the world, such as Vietnam, Somalia, Colombia, Peru, El Salvador, Guatemala, Nicaragua, and most recently, Afghanistan and Iraq. Due to the landscape of irregular warfare during counterinsurgency operations, whereby the use of military forces in a conventional sense is limited, the United States has been forced to alter and adapt the use of its military force in order to accomplish strategic, operational, and tactical goals against the insurgencies. Political power is the key issue which is sought by insurgents and counterinsurgents as both sides strive for popular support of the people. As such, the U.S. Army states that the “primary objective of any COIN operation is to foster development of effective governance by a legitimate government,” which is obtained by winning the support of the general populace [2].

Lieutenant General William B. Caldwell, IV, who is a general in the U.S. Army tasked with writing doctrine for the Army, recently wrote:

The future is not one of major battles and engagements fought by armies on battlefields devoid of population; instead, the course of conflict will be decided by forces operating among the people of the world. Here, the margin of victory will be measured in far different terms than the wars of our past. The allegiance, trust, and confidence of populations will be the final arbiters of success [3].

As General Caldwell mentions, future wars will not be fought on open battlefields with large conventional military forces, but will rather be fought over gaining the allegiance and support of the population. Therefore, the success of the United States during its current (and future) counterinsurgency efforts will undoubtedly hinge upon the ability for the United States to perfect its strategies for engaging the populace, while simultaneously thwarting the ability for insurgents to gain prominent footholds as legitimate authorities.

## 2.1 The Nature of Insurgencies

The United States military defines an *insurgency* as an “organized movement aimed at the overthrow of a constituted government through the use of subversion and armed conflict” [2]. Insurgents conduct organized politico-military operations in a protracted environment in order to weaken the control and legitimate power of the established government. Insurgents use all methods available to them in order to gain political power and support from the population as legitimate authorities. These tools include, but are not limited to, political, informational (such as appeals to ideological, ethnic, or religious beliefs), militaristic, and economic. The most common approaches to warfare used by insurgents are terrorist and guerrilla tactics. While insurgents would prefer quick and overwhelming victory rather than a long and protracted struggle, the latter often occurs due to the significant disadvantage insurgents face in terms of resources and technology to fight on the open battlefield [2].

Although each insurgency is unique in its specific motives, many similarities exist among them. In every case, insurgents strive to cause political change and the use of military force in order to achieve such change is secondary in nature and only used as a means to an end. Moreover, protracted conflicts only serve to favor insurgents as no other approach seems to make better use of such asymmetry than conducting a protracted popular war. Many historical insurgencies used the protracted warfare approach in order to eventually wear down the enemy and secure victory—examples include the Chinese Communists in conquering China after World War II, as well as the North Vietnamese and Algerians. Additionally, the United States has similarly been exposed to the protracted mindset of the insurgents they face as some Al Qaeda leaders have suggested such approaches in their propaganda writings [2].

### **Mao Tse-tung’s Theory of Protracted War**

Mao Tse-tung is known as being the ‘father of modern insurgency’ due to his 1937 book *On Guerilla Warfare* which outlined how the Chinese people should organize and conduct unlimited guerilla warfare on the Japanese invaders of China during the Second Sino-Japanese War. His theory of protracted war outlines three strategic phases that should be used by

insurgents in order to eventually wear down and eventually defeat the enemy, namely (1) Strategic Defensive, (2) Strategic Stalemate, and (3) Strategic Counteroffensive [2, 4].

### **Phase I – Strategic Defensive**

During Phase I, the insurgency is in a dormant state which is focused on survival and building support for its cause. Insurgents in this phase are much weaker than the government which has superior strength in forces. Thus, insurgent leaders establish base camps and gain support from the population through the use of propaganda, boycotts, sabotage, and demonstrations. Major military combat with the enemy is avoided during this phase as the primary focus is establishing the insurgent movement among the populace by creating logistical support networks, finding recruits, and gaining intelligence on the enemy. Army Field Manual (FM) 3-24 highlights a list of key objectives which are accomplished by the leaders of the insurgency during this phase [2]:

- Recruit, organize, and train cadre members
- Infiltrate essential government organizations and civilian groups
- Establish cellular intelligence, operations, and support networks
- Solicit and obtain funds to support the movement
- Develop sources for outside support

Lastly, during the later stages of Phase I, the insurgency may decide to establish a *counterstate* (or shadow government) which is in opposition to the established government authority and designed to replace the existing government [2].

### **Phase II – Strategic Stalemate**

Phase II is characterized by the expansion of insurgent military forces and overt guerrilla warfare. The objective of the insurgency during this phase is to undermine the people's support of the government while at the same time expanding their areas of control. The insurgents attempt to delegitimize the government by showing the people that the government is incapable of maintaining control of the country and protecting its people. They accomplish this by attacking government forces to show their ineffectiveness while simultaneously conducting

aggressive propaganda campaigns and information operations to influence the populace. Phase II is a combination of military actions by both sides against each other and a simultaneous fight for the support of the people. The military actions of the insurgents have two main purposes— (1) to decrease the fighting capability of the other side, and (2) to convince the people that the insurgents are stronger and more worthy of support. Due to the massive recruiting efforts during Phase I (and continuing throughout the insurgency), the insurgents have enough military force to enable them to concentrate their resources for short periods in order to locally overwhelm government forces. However, the insurgents cannot concentrate their forces indefinitely in any one area and must be able to effectively hide their forces among the people in order to prevent their destruction by the government's superior firepower [2].

In Phase II the insurgency gains popular support in many ways. First they conduct attacks against government targets (security forces, infrastructure, and individuals), which show the impotence and ineffectiveness of the government, while highlighting the strengths of the insurgents. Shadow governments established by the insurgents in areas they control provide services (mainly through conflict resolution, legal, and financial means, not infrastructure) and security for their supporters and work to supplant local government by showing the incompetence of the government. Large information campaigns against the government further erode popular support. Throughout Phase II the insurgent is focused on simultaneously strengthening himself and weakening the government through protracted conflict [2].

### **Phase III – Strategic Counteroffensive**

Finally, Phase III is designated by the insurgency as a counteroffensive period whereby the insurgents sense weakness in the established government authority. The insurgents transition from guerilla warfare to conventional warfare during the counteroffensive movement as they aim to destroy the enemy's military capability and will to fight. Success in Phase III is determined largely by armed conflict, and not propaganda or popular support. However, failure of the insurgent in Phase III does not mean its destruction. Success or failure is often decided quickly in Phase III as it is readily apparent that one side or the other is stronger and military success often rapidly follows [2].

If the counteroffensive movement is successful and the insurgents begin to gain control of portions of the country they become responsible for the population, resources, and territory under their control. Thus, as the balance of power shifts from the government to the insurgents, it is important for the insurgents to maintain the support of the people by:

- Creating an efficient civil administration
- Establishing an effective military organization
- Providing balanced social and economic development
- Mobilizing the population to support the insurgency
- Protecting the population from hostile actions

A triumphant insurgency does not, however, depend on the sequential or successful application of all three stages. The overall objective of any insurgency is to achieve political power. Therefore, if the insurgency is unsuccessful in a later phase it can easily revert back to an earlier stage until a new opportunity presents itself for transition into the next stage. Moreover, insurgents are not required to follow Mao Tse-tung's exact formula for protracted war in order to be successful. Recent insurgencies, such as the Algerians, have been successful without the use of large-scale conventional warfare. Thus, counterinsurgents are faced with a significant challenge in defeating insurgents due to the nature of insurgencies which, although have similar end objectives of gaining political power, can be very different in their specific characteristics, strategies, and tactics used against the government [2].

## **2.2 The Nature of Counterinsurgencies**

The United States Army Field Manual 3-24 defines a *counterinsurgency* as “military, paramilitary, political, economic, psychological, and civic actions taken by a government to defeat insurgency” [2]. Because political power is the central issue fought over by insurgencies and counterinsurgencies, counterinsurgents will use all available instruments of national power in order to maintain the established or fledgling government from losing its legitimacy. Thus, even though the main purpose of America's military is to fight and win wars for the United States, while at the same time securing and defending the nation from potential adversaries, the

unique nature of counterinsurgency efforts requires the U.S. military to adapt to a different landscape of conflict.


Counterinsurgents must not only be ready to fight, but also to build and provide security for the population. COIN operations consist of a mix of offensive, defensive, and stability operations, whereby the exact proportion of effort devoted to each will change over time depending on the current situation. The protection, welfare, and support of the populace are crucial to the success of any counterinsurgency effort due to the intelligence gained from the local people in identifying and eventually rooting out the insurgents. Thus, counterinsurgencies should place significant focus on isolating the insurgents from their cause and support because it is “easier to separate an insurgency from its resources and let it die than to kill every insurgent” [2]. Although killing and capturing insurgents is necessary during COIN operations, attempting to eliminate every insurgent is impossible and can also be counterproductive in terms of producing popular resentment and/or creating martyrs which increase the number of new recruits seeking revenge. Therefore, only after popular support has been won by the counterinsurgents will the insurgency be conquered as (1) the reliance on external support, (2) the need for secrecy in base camp operations, and (3) the capability of hiding amongst the populace are three key vulnerabilities facing insurgencies which can be thwarted by intelligence gained from the population supporting the cause of the counterinsurgency [2].

### **2.3 Nonlethal Targeting used to Win Popular Support**

Any insurgent movement requires the support of the populace in order to survive, let alone to succeed in overthrowing the government. In order to grow and gain strength, an insurgency requires a sufficiently sized population base which will actively support the growth. Thus, one of the most effective ways for counterinsurgents to defeat the insurgents is to efficiently shrink the population base supporting the insurgency by causing the local populace to become hostile or at least apathetic towards the insurgents [5]. As previously discussed, insurgents heavily rely on safe havens within the communities in order to remain hidden from government forces. These safe havens, however, only exist whenever sufficient numbers of the local population support the insurgents [2].

During the struggle between insurgents and counterinsurgents, there will be a minority of people who actively support each side and a majority of people who are either passive in their support or are indifferent toward either side. This concept is best illustrated below in Figure 2.1 from Army Field Manual 3-24.2 [5]. The three main components (active support, passive support, and indifferent) are described in further detail below.

- **Active Support:** Active supporters view the side they support as the legitimate authority, and they personally or publically align themselves with this side. Besides joining the militia, active supporters of an insurgency may participate in a wide array of other activities in support of their side including: spreading propaganda, giving financial support, offering medical assistance or safe havens, providing intelligence or logistical support, and recruiting others to join the cause [2]. Meanwhile, active supporters of the existing government may join the military or police force, provide intelligence to the counterinsurgents, and spread propaganda denouncing the insurgents [6].
- **Passive Support:** Passive supporters are those people who are sympathetic toward one side yet who remain inactive and non-hostile towards those you are not sympathetic toward their cause. Passive supporters of the insurgency are those who allow insurgents to conduct operations in their local proximity and remain silent (or give false information) if asked by counterinsurgents about intelligence concerning insurgent activity [2]. Passive supports of the counterinsurgents obey the rule and laws of the government and may support COIN operations against the insurgents if there is minimal risk [6].
- **Indifferent:** Indifferent individuals are those who represent a large portion of the population. They are unsure of which side to support and have a tendency to either (1) remain neutral until there is clearly a victor, or (2) cater to whoever has a stronger presence in their proximity [2].



*Figure 2.1 – Spectrum of Popular Support [5]*

While both the insurgents and counterinsurgents are vying for support of the populace, the proportion of popular support that is needed for each side to be victorious is very different. Because insurgents can easily create disorder and havoc, the counterinsurgents must usually obtain much greater than 50 percent of the popular support. Meanwhile, a largely passive populace may be all that is necessary for an insurgency to be successful and seize political power [2]. Thus, in order to effectively eliminate an insurgency, counterinsurgents must focus their efforts on achieving a great majority of the popular support through nonlethal targeting methods, such as personal communications, negotiations, and meetings with the population, especially with those individuals who are passive or indifferent.

Although lethal targeting may be necessary, the success of any counterinsurgency ultimately hinges on the ability to win the ‘hearts and minds’ of the populace through nonlethal targeting strategies. ‘Hearts’ refers to the ability to persuade the people that their best interests are served by counterinsurgency success, whereas ‘minds’ means convincing the population that the government force can protect them and that resisting the COIN effort is pointless [2]. However, because the U.S. military has limited resources and personnel, critical decisions must be made by ground forces on whom to selectively target with the end goal of winning the ‘hearts and minds’ of the most people. Although the U.S. Army’s two primary field manuals on counterinsurgency (FM-3-24 [2] and FM 3-24.2 [5]) combine for more than 500 pages of COIN doctrine, of which roughly 50 pages discuss targeting, the advice given for how an Army unit should specifically determine which individuals to target and how to engage them is quite vague. Together, these two manuals state the following:

- Identify leaders who influence the people at the local, regional, and national levels.
- Win over passive or neutral people.
- Nonlethal targets include people like community leaders and insurgents who should be engaged through outreach, negotiation, meetings, and other interactions.
- Meetings conducted by leaders with key communicators, civilian leaders, or others whose perceptions, decisions, and actions will affect mission accomplishment can be critical to mission success.
- Start easy...Don’t try to crack the hardest nut first—don’t go straight for the main insurgent stronghold, try to provoke a decisive showdown, or focus efforts on villages

that support the insurgents. Instead, start from secure areas and work gradually outwards. Do this by extending your influence through the locals' own networks.

The task of determining which people to target at any given time is difficult. Successful nonlethal targeting tactics used in one COIN effort may not produce the same results when used in another COIN operation, which is why it is impossible for a manual to give deterministic formulas for nonlethal targeting strategies. Thus, the Army manuals only provide general, experienced-based guidelines and approaches for how Army commanders should determine which local leaders to target.

This thesis focuses on arguably the most important aspect of any counterinsurgency effort—the extremely difficult task of determining which individuals to engage through nonlethal means in order to gain the most support of the populace. The objective of this thesis is to formulate a realistic, peer-influence social network to model the population of a rural Afghanistan district, and subsequently develop decision tools to help military commanders decide which individuals to engage through nonlethal targeting techniques. The proposed methods seek to determine the impact of different targeting strategies on the population's attitudes over time with the goal of maximizing the beliefs of the populace in support of the U.S. counterinsurgency cause.

[This Page Intentionally Left Blank]

### **3 Modeling Approach and Formulation**

We begin Chapter 3 by with a literature review of belief propagation in social networks. More specifically, we discuss the foundational modeling approach proposed by Acemoglu et al. [7] and Hung [8] as well as shortcomings of the model which have inspired the development of our model. Lastly, we discuss the formulation of a dynamic, two player “chess game” approach whereby U.S. forces and Taliban insurgents attempt to maximally influence the beliefs of the people in the network in a transient environment.

#### **3.1 Literature Review**

Below we discuss the major work done in the field of modeling societal networks and the diffusion of information in such networks. This literature serves as the foundation of our modeling approach.

##### **3.1.1 Influential Leaders in Social Networks**

The key to understanding the problem of belief propagation in social networks is first to understand the characteristics of the network and how information flows from one person to the next. Katz and Lazarsfeld [9] led the way in the field of social science and public opinion formation through the creation of their ‘two-step flow’ model of communication. They theorized that individuals in society may in fact be influenced more by the exposure to one another than by their exposure to the mass media. Thus, they argued that a small minority of influential ‘opinion leaders’ act as liaisons between the media and the rest of society—hence, the ‘two-step’ process. It is these influential leaders who have the greatest access to media as well as a better understanding of the content. In turn, each opinion leader shapes and molds the media’s content and diffuses the information to the people they wield influence over—as the followers of opinion leaders tend to share the same interests, demographics, personalities, and/or socio-economic factors. As such, the idea of focusing political campaigns on influencing the attitudes of influential leaders, and not the attitudes of the critical masses, stemmed from the work done by Katz and Lazarsfeld as it is the influential leaders who are responsible for spreading information and shaping the beliefs of the masses.

Since the foundational work of Katz and Lazarsfeld, many have offered both criticisms and follow-on support for the two-step communication model. Watts and Dodds [10] question how precisely influential leaders exert their influence over their followers in society. They seek to determine “exactly how, or even if, the influentials of the two-step flow are responsible for diffusion processes...or other processes of social change.” Watts and Dodds do not entirely discount Katz and Lazarsfeld’s two-step model, since their research finds examples supporting the idea that opinion leaders are responsible for generating significant ‘cascades’ of influence. However, they believe this is the exception and not the general rule, as they find that influential leaders are “only modestly more important than average individuals.” As such, they claim that the main factor which drives belief propagation in social networks is not the opinion leaders, but rather by easily influenced individuals in turn persuading other easily influenced individuals. This notion directly relates to the ideas behind the cascade and linear threshold models, which will be discussed shortly.

### **3.1.2 Opinion Dynamics**

While several researchers have developed methods for understanding the social aspect of belief propagation and the roles of opinion leaders, others have developed mathematical models to simulate and analyze how information spreads through a social network. In 1964, Abelson [11] created a network model such that pairwise interactions occur between agents in the network which are adjacent to one another (i.e. interactions are only able to occur between agents who know each other, and such adjacency connections resemble those relationships and enable a way for information to be shared or obtained). As two agents interact with each other, the resulting outcome is a function of the two agents’ original beliefs as well as their persuasiveness. While Abelson’s work was an initial step towards a field of much uncertainty, he does mention the limitations of his model since all agents eventually reach a consensus agreement on their opinions which is not necessarily realistic in society.

Later, in 1974, DeGroot [12] used Markov chain theory to model opinion dynamics in society. Each agent predetermines weights for each of his neighbors (adjacency connections) from which he/she obtains information. For instance, if agent  $i$  believes agent  $j$  is a knowledgeable expert on the subject matter, then agent  $i$  will choose a sufficiently large weight

$(p_{i,j})$  for his connection between agent  $j$ . For all agents and their edges, these weights are obtained and comprise what is known as a one-step transition probability matrix of a Markov chain. What makes the transition probability matrix interesting is that the matrix is stochastic since all rows must sum to one. Because of this, DeGroot notes that “if all recurrent states of the Markov chain communicate with each other and are aperiodic, then a consensus is reached.” This consensus opinion is calculated as the sum product of the steady-state probabilities (which can be determined because all recurrent states of the chain communicate and are aperiodic) and the initial opinions. Even though DeGroot mentions how a consensus is not reached should the recurrent and aperiodic condition not be met, his consensus-driven approach to the attitudes of agents in a network is not necessarily appropriate in a realistic setting as people have diverse opinions on all matters and do not always reach a consensus.

Moving forward, Acemoglu et al. [7] developed the spread of (mis)information model which is a stochastic, agent-based network model with two types of agents—regular and forceful. As such, forceful agents in the network are those agents which have a positive probability of forcefully imparting their belief on other agents, while regular agents do not have such a capability. Pairwise interactions between adjacent agents occur according to two probabilistic features—(1) the frequency of their interactions and (2) the type of interaction which occurs, whether it be averaging (both agents average their beliefs), forceful (one agent influences another), or identity (no change in belief). The authors mention the importance of the assumption that “no man is an island” which essentially states that no agent is completely unaffected by the beliefs of those around him. Because of this assumption, in which even forceful agents (with possibly vastly different starting opinions) are capable of being affected by the beliefs of the rest of society, the network beliefs arrive at a consensus among all agents; however, such a consensus is a random variable as the beliefs converge to a convex combination of the initial scalar beliefs.

More recently, Hung [8] added to the existing model developed by Acemoglu et al. [7] by creating a ‘very forceful’ agent class (who are even more influential than the ‘forceful’ agents) as well as a set of ‘stubborn’ agents with beliefs that are static and never change. Thus, the ‘stubborn’ agents are the most forceful (influential) agents in the network and effectively diffuse their beliefs throughout the network. Because of the presence of ‘stubborn’ agents, the beliefs of

the agents in the network never converge to a single value. However, Hung was able to characterize a well-defined first moment for the random variables of each agent's belief. Over time, the expectation converges to an equilibrium value, which means that the belief of each agent converges in expectation to a fixed value in a sufficient amount of time. Lastly, although it is mentioned that the beliefs of the agents do not converge to fixed values in simulation (once again, due to the presence of 'stubborn' agents), Hung does show that the beliefs converge to a type of stochastic equilibrium such that the average of all agent beliefs in the network oscillate around the expected mean belief of the network.

Hung also made two other important contributions—(1) he created a network generator tool and (2) he developed a non-linear, non-convex, mixed integer mathematical formulation to solve the complex problem of finding stubborn agent placements in order to maximally influence the network. The network generator tool was created in a collaborative effort between Hung and the MIT Political Science Department. The tool is designed to create realistic social interaction networks among Pashtun local leaders in a rural Afghanistan district. We will extensively use this network generator tool for the experiments and analyses conducted in Chapter 4. The inputs to the generator include information and intelligence that would hopefully be accurately gathered by counterinsurgents during interactions with the local population. These inputs would include: (1) determining who the local leaders are and their specific roles in society, (2) an estimation of the attitudes of the local population (the specific context to be discussed later), and (3) how many Taliban agents there are and with whom do they interact with and influence. The network generator tool takes these inputs and creates homophily connections between agents who are likely to interact with each other based on socio-demographic features. The resulting network created contains a list of agents and their connections in an undirected graph denoted as  $G = (V, \epsilon)$ . Next, based on the network which is created (including the locations and connections of the Taliban 'stubborn' agents), Hung's mathematical program would determine the optimization-based placement for U.S. 'stubborn' agents in order for the U.S. agents to maximally influence the network based on the long-term expected attitudes of the agents.

Howard [13] subsequently determined that Hung's placement solutions for U.S. stubborn agents were local optimums based on his results analyzing the same problem with a new approach. Howard created a two player game which used a simulated annealing heuristic to look

for Nash equilibrium solutions (whereby neither stubborn agent can improve his payoff by unilaterally deviating from the current strategy). His algorithm employed ‘best response dynamics’ in a back-and-forth updating game until Nash equilibria were either found to exist or not exist. Although Howard mentions the limitations of simulated annealing (no guarantee of finding global optimum solutions and problems dealing with cases that contain multiple equilibria), he does point out that not only were his solutions better than Hung’s, but that his simulated annealing heuristic was much faster, too.

Both Hung and Howard contributed greatly to the problem of modeling rural Pashtun villages in Afghanistan and formulating programs to determine optimization-based placements of stubborn agents in order to maximally influence the network. However, in this thesis we take a new approach to this problem, and we present reasons for improving their models in the following section as we attempt to reach a more realistic setting in terms of the modeling approach and the two player game.

### **3.1.3 The Need for a More Realistic Model**

There are several areas of improvement concerning the previous models presented by Hung and Howard which we address here. First, one of the characteristics of the previous models was that the long-term expected beliefs are independent of the initial beliefs of the mutable agents in the network. Hung and Howard both found that the long-term expected attitudes only depend on the underlying influence parameters and topology of the network (particularly the placement of stubborn agents). From a real world point of view, this finding seems unrealistic, especially if we view two extreme cases of the same network with different initial conditions—(1) every agent’s initial belief is  $-0.5$  (pro-Taliban), except for the U.S. stubborn agent which has an immutable belief of  $+0.5$ , and (2) every agent’s initial belief is  $+0.5$  (pro-United States), except for the Taliban agent which has an immutable belief of  $-0.5$ . In the previous implementations, both extreme cases would yield the same expected long-term beliefs through both the analytical expression and simulation, assuming enough interactions are allowed to occur for the network beliefs to converge to a quasi-equilibrium state. Reasonably, one would anticipate entirely different long-term beliefs for these two extreme cases as other factors, such as peer pressure, will affect the degree to which an agent is willing to change his or her belief.

Thus, if a network is highly concentrated with like-minded agents, we expect the agents in the network to remain close to these initial beliefs despite outside influence from a stubborn agent of an opposing belief. This leads us to propose incorporating peer pressure into the model which will affect belief propagation during pairwise interactions.

An important aspect which was not included in the previous models was the effect peer pressure can play in how influence propagates through the network. For example, one may be more likely to change their current opinion about an issue given 90% of their friends now hold a new opinion versus the same scenario whereby only 10% of their friends hold the new opinion. We believe peer pressure plays an integral part in affecting attitude dynamics and should thusly be included in the social influence model.

Next, Howard's two-player game approach was concerned with finding Nash equilibria. In this situation, opposing stubborn agents would play a back-and-forth game connecting to agents until (hopefully) a pure Nash equilibrium strategy was found whereby both agents would settle on the same targeted agent and exert equal, polar-opposite influence on the particular agent. The opposing stubborn agents (Taliban and United States) would not be inclined to unilaterally deviate from the Nash equilibrium strategy as doing so would be unbeneficial towards their respective payoffs. The problem with this implementation is two-fold. First, it is a very strong and possibly unrealistic assumption that Taliban insurgents and U.S. ground forces would be content on 'talking' to the same local/district leader in order to influence them (and the rest of the Pashtun village(s)) toward supporting their cause. Second, even if the U.S. and Taliban agents were to settle on the same agent to influence, one would not expect equal influence to be exerted by both stubborn agents on the targeted agent as an inherent bias toward supporting one side would most likely occur rather than indifference toward both sides. We therefore will consider enhancements that will put pressure on stubborn agents to select different strategies during the course of the two-player game.

Finally, both Howard and Hung concentrated on modeling and understanding the long-term expected attitudes of the network as the network beliefs reach an asymptotic equilibrium state. In a constantly-changing, dynamic world, analyzing the long-run equilibrium state of the network seems unrealistic. We should instead be concerned with the transient state of the

network to enable quick strategy changes given short-term information on the state of the network. Therefore, we seek to implement a two player game centered on the transient environment in a back and forth ‘chess game’ whereby the opposing stubborn agents consistently reevaluate and update their strategies as necessary.

Based on the proposed areas of enhancement for the model, we seek guidance from literature on different social influence models which may be better suited for creating more realistic models of society. Furthermore, we also seek literature which focuses on the transient behavior of belief propagation in social networks rather than the long-term asymptotic equilibria.

### **3.1.4 Cascade and Threshold Models**

Schelling [14, 15] discusses how discriminatory individual behavior can lead to segregation, separation, or sorting in society. Such discriminatory behavior can be a conscious or an unconscious awareness to age, religion, sex, color, etc. that influences the decisions people make about where to live, where to sit, what job to get, and with whom to talk to and associate with. Schelling creates models in which members of two distinct ‘color’ groups make discriminatory choices about their locations based on the identities of their neighbors. Thus, an individual will move if he is not satisfied with the ‘color’ mixture of his neighborhood. Schelling mentions some individuals may be more tolerant than others, and thus, it takes a higher percentage of their neighbors to be of a different ‘color’ group before they become unsatisfied and change their location. Because of the inherent preference for people to be surrounded by others who are similar to themselves, Schelling shows how people can promote segregation through their individual discriminatory choices of everyday life. Lastly, Schelling touches on a topic called neighborhood ‘tipping’ which occurs when a new minority enters a particular neighborhood in sufficient enough numbers to cause the previous majority group to begin evacuating and relocating as the ‘tipping point’ is reached and they eventually become the minority.

In 1978, Granovetter [16] expanded on the initial segregation models created by Schelling and proposed the threshold model as an approach for modeling societal networks. He argued that an individual’s behavior depends on the number of people who are already engaging in such behavior, which thus gives notion to the way peer pressure influences behavior and

opinions. Therefore, every person has their own ‘behavioral threshold’ which is the proportion of friends which must participate in some action before he or she will be convinced or persuaded to participate in the same action. Thus, an individual with a ‘threshold’ of 0% would be classified as an ‘instigator’ since 0% of that person’s friends would need to previously be partaking in the desired behavior before he or she decides to as well. Meanwhile, people who are more conservative in nature will have higher thresholds—e.g. 80% to 90%. Since people are assumed to act rationally in order to maximize their utility, their individual ‘threshold’ values should be the point at which the perceived benefits of holding some opinion or doing some action outweigh the cost of not doing so.

Recently, Kempe et al. [17] discussed the problem of maximizing the expected spread of an idea or behavior within a societal network due to the presence of ‘word-of-mouth’ referrals. He introduced two stochastic influence models – the *independent cascade (IC) model* and the *linear threshold (LT) model*. In the IC model, Kempe et al. assigned each edge an activation probability and influence is spread in the network by activated nodes independently activating their inactive neighbors based on the activation probabilities. Meanwhile, the LT model assigns weights to each edge as well as threshold values to each node. In this process, a node will become activated if the weighted sum of its active neighbors exceeds its individual threshold value. Kempe et al. proved that the influence maximization problem in both the IC and LT models are NP-hard. They also developed a greedy algorithm for the models which successively selects the node with the maximum marginal belief propagation. Lastly, they showed that their greedy algorithm approximates the optimal solution within a ratio of  $1 - (1/e)$ , or 63%. This approximation ratio stems from the nice properties (*monotonicity* and *submodularity*) of the belief propagation function used by the models.

### **3.1.5 Transient Behavior of Belief Propagation in Social Networks**

Ozdaglar and Yildiz [18] studied the optimal placement of stubborn agents in order to maximize the spread of opinions in the transient setting. They assumed an undirected network  $G$ , with a set of stubborn agents of belief zero,  $V_0$ , with known locations, and well as the initial beliefs of the remaining nodes in the network (either type 0 or type 1). Their goal was to choose  $k$  nodes (not including  $V_0$ ) to form the stubborn agent set  $V_1$  such that the *bias* of the network is

maximized after  $l$  steps, where  $k > 0$  and  $l > 0$ . Ozdaglar and Yildiz produced several conclusions for the influence maximization problem for the transient case—(1) it is always optimal to target nodes with type 0 beliefs rather than type 1 beliefs, (2) optimal target nodes ideally have a large number of neighbors with small neighborhoods, and (3) the neighborhoods are dominated by nodes with belief type 0. They noted that their results are valid when maximizing the expected belief of the network for one step into the future.

## 3.2 Proposed Threshold Network Model Formulation

In the previous section, we mentioned areas of enhancement to be explored from the preceding models developed by Hung and Howard. Thus, the intent of the current modeling formulation is to address those areas with the following measures—(1) the long-term beliefs of the mutable agents in the network should be *dependent* on the initial beliefs of those mutable agents, despite the presence of stubborn agents, (2) peer pressure should affect the outcome of belief exchanges during pairwise interactions, (3) the appeal of identical strategies among opposing stubborn agents (U.S. and Taliban) should be diminished during the two player game, and most importantly (4) we concentrate our focus on studying the transient behavior of the two-player game whereby continual readjustments can be made by both opposing stubborn agents as the game progresses. The proposed formulation adds additional insights into the field of social influence models and how belief propagation is affected due to various modeling parameters.

### 3.2.1 Scope of the Model

Hung's large, 73 node (excluding stubborn agents) network model, which is created through the use of his network generator tool, is designed to accurately represent an Afghanistan district composed of several villages. Districts in Afghanistan are typically comprised of 10-20 villages, each of which contains approximately 100 households or more. We limit ourselves to modeling one Pashtun district in Afghanistan throughout this thesis for three reasons. First, U.S. Army company commanders deal with the problem concerning this thesis—how to maximally influence a rural Pashtun district—and thus, we attempt to mimic the role of an Army company commander by focusing our analysis on only one district. Second, a much larger network spanning many districts in Afghanistan can make the optimization aspect of the problem very

tedious and time consuming. Finally, Afghanistan districts tend to follow historical tribal and social boundaries of which it is believed that the opinions of villagers in one district are greatly independent of the opinions of villagers in another district [13]. Thus, networks containing multiple districts might be deemed unnecessary as sequential analyses of individual districts may be just as accurate.


### **3.2.2 Network Characteristics**

We assume we are given an undirected network,  $G = (V, \mathcal{E})$ , where  $V$  represents the set of agents (i.e. individuals, nodes, or vertices) and  $\mathcal{E}$  is the set of connections (i.e. edges or arcs) between those agents. Throughout our analysis we assume the topology of the network is static such that no agents are added to or removed from the network over the time horizon we are concerned with in the experiments. Because we are focusing our analysis on the transient (short-term) behavior of the network, rather than the long-term behavior (as Hung and Howard had previously analyzed), the assumption of a static network structure is more believable. We do make two exceptions to the notion of creating a static network structure. First, stubborn agents are allowed to change their connections (arcs) to different mutable agents over the course of the two player dynamic game in an attempt to maximally influence the network beliefs toward supporting their cause. Second, special cases involving the Taliban stubborn agent whereby they ‘assassinate’ an agent in the network will temporarily decrease the forcefulness (influence) level of that agent. Both of these exceptions will be discussed later.

### **3.2.3 Agents Characteristics**

#### **3.2.3.1 Belief Spectrum**

We assign each agent in the network a scalar belief on the continuous spectrum from -0.5 to +0.5 (see Figure 3.1 below). This belief indicates how favorable or unfavorable an agent is toward the United States (US) or Taliban (TB). For example, an agent who has a belief of -0.5 strongly favors the Taliban (and is heavily opposed to the US), while an agent with a belief of +0.5 is strongly supportive of the US (and is heavily opposed to the TB). Meanwhile, an agent with a belief of 0 (neutral) is indifferent toward both sides.



*Figure 3.1 – Agent Belief Spectrum*

### 3.2.3.2 Forcefulness (Influence) Level

Moreover, each agent is assigned a particular forcefulness or influence level depending on the type of person they are in society, which directly relates to their job or role in society. We describe the four types of influence levels below:

#### **Regular**

A regular agent is an agent who is a head of household in a rural Pashtun village. Typically, the head of household is the eldest male within a family and is the only voice of the house that matters during important tribal meetings. In Pashtun society, it is not uncommon for more than one generation to live under the same roof, which can amount to up to 40 people living in the same house [19]. While there may be minor disputes between individuals of the same household, we assume that the belief of the eldest male (head of household) accurately represents the belief of everyone living in the house, especially due to the fact that the society is highly patriarchal. As such, regular agents have the lowest level of influence of all agents in the network, but they realistically represent all of the family members who reside under one roof. Regular agents are represented as squares in the network diagrams.

#### **Forceful**

Forceful agents are village leaders which hold more influence than the ‘regular’ head of household agents within a village. Although forceful agents are also most likely to be the head of a household, their role in society gives them further influence than regular agents. Some examples of forceful agents include local (village) religious figures, members of the Shura

(advisory council in rural Pashtun society), or wealthy merchants. Forceful agents are depicted as upside-down triangles in the network diagrams.

### **Forceful+**

Forceful+ agents are even more influential than forceful agents. They represent more central and powerful individuals in Pashtun society, such as district government officials, district religious figures, or the district police chief. Because forceful+ agents are often more centrally located within the network, they wield influence beyond just the village setting. Forceful+ agents are represented as triangles in the network diagrams.

### **Stubborn**

Stubborn agents are the most influential agents in the network. Also, stubborn agents are the only agents in the network with immutable (unchangeable) beliefs. There are two stubborn agents in the network (one US agent representing military ground forces with belief +0.5 and one Taliban agent representing insurgents with belief -0.5). We assume that the US and TB have immutable beliefs of +0.5 and -0.5 because it is unlikely that either opposing stubborn agent can effectively influence and change the belief of the other stubborn agent. The goal of each stubborn agent is to maximally influence the beliefs of the mutable agents (regular, forceful, and forceful+) in the network toward supporting their cause. Stubborn agents are depicted as diamonds in the network diagrams.

Lastly, we represent all agents in the network through the following set notation found below. Agents belong to specific sets based on their forcefulness level.

$V_R = \text{set of all regular agents}$

$V_F = \text{set of all forceful agents}$

$V_{F+} = \text{set of all forceful + agents}$

$V_{US} = \text{set of all US agents}$

$V_{TB} = \text{set of all TB agents}$

$V_M = \text{set of all mutable agents} = V_R \cup V_F \cup V_{F+}$



$V_S = \text{set of all stubborn (immutable) agents} = V_{US} \cup V_{TB}$

$V = \text{set of all agents} = V_R \cup V_F \cup V_{F+} \cup V_{US} \cup V_{TB} = V_M \cup V_S$

$|V| = \text{total number of agents (nodes) in the network} = n$

In Figure 3.2, we represent a general network diagram to illustrate the various concepts discussed thus far. Note that the current belief of each agent is represented by the color of their respective node symbols, which is based on the color spectrum shown in Figure 3.1.

For the convenience of the reader, Appendix A contains a list of all variable notation and their definitions introduced in this chapter.



*Figure 3.2 – Example Network Diagram and Symbol Notation*

### 3.2.4 Pairwise Interactions

We assume all agents have an identical Poisson process with rate 1. Thus, conditioned on the fact that an arrival has occurred in the network, it has an equal probability of having arrived in any agent's Poisson process.

$$\begin{aligned}
 &P(\text{arrival occurred for agent } i | \text{an arrival occurred in the network}) \\
 &= P(\text{arrival occurred for agent } j | \text{an arrival occurred in the network}) \quad \forall i, j \\
 \therefore P(\text{arrival occurred for agent } i | \text{an arrival occurred in the network}) &= \frac{1}{|V|} = \frac{1}{n} \\
 &\text{where } n \text{ is the number of agents (nodes) in the network}
 \end{aligned}$$

When an arrival occurs in a particular agent's Poisson process, it becomes the *active* agent in the network and will subsequently search for one of their neighbors to interact with. We assume that all neighbors of the active agent have an equal probability of being selected for a pairwise interaction with the active agent.

$$\begin{aligned}
 &P(\text{agent } j \text{ is chosen for pairwise interaction with agent } i | \text{agent } i \text{ is the active agent}) \\
 &= P(\text{agent } k \text{ is chosen for pairwise interaction with agent } i | \text{agent } i \text{ is the active agent}) \\
 &\quad \forall j, k \in N_i \\
 &\text{where } N_i \text{ is the set of adjacent neighbors (edges) to agent } i
 \end{aligned}$$

$$\therefore P(\text{agent } j \text{ is chosen for pairwise interaction} | \text{agent } i \text{ is active agent}) = \begin{cases} \frac{1}{|N_i|} & \forall j \in N_i \\ 0 & \text{else} \end{cases}$$

Although some agents may be more inclined to interact with a particular neighbor(s) due to prior good relationships and/or proximity to one another, we do not allow a non-uniform distribution for the selection of neighbors for the active agent to choose from. We rely on the use of a uniform distribution in order to simplify the model. Also, due to a lack of sufficient information and data in order to revise this neighbor selection distribution, we feel a uniform distribution is the best alternative.

Furthermore, we denote the belief of agent  $i$  at time step  $t$  as  $X_i(t) \in [-0.5, 0.5]$ . A time step is defined as the occurrence of one pairwise interaction in the network. Thus, time step  $t$  is defined as the  $t$ -th interaction in the network, and time step 0 is defined as the initial state of the network. The vector  $X(t) \in \mathbb{R}^{n \times 1}$  denotes the beliefs of all agents in the network at time step  $t$ :

$$X(t) = [X_1(t), X_2(t), \dots, X_n(t)]^T$$

After the active agent (agent  $i$ ) selects one of his neighbors (agent  $j$ ) at random, one of three possible interaction types will occur:

**I. Forceful Interaction ( $\alpha$ -type interaction)**

With probability  $\alpha_{ij}$ , agent  $i$  will *attempt* to ‘forcefully’ impart  $(1 - \tau_j)$  of his belief on agent  $j$ , where  $\tau_j$  represents the threshold value of agent  $j$ . Several outcomes can occur during a forceful interaction which will determine how agent  $j$  will update his belief. We will explain this and the meaning of an agent’s threshold value later in more detail when we describe the formulation of the proposed threshold model.

**II. Averaging Interaction ( $\beta$ -type interaction)**


With probability  $\beta_{ij}$ , agent  $i$  and agent  $j$  will each *attempt* to reach a consensus equal to the average of their prior beliefs. Either one, both, or neither agent will update their belief to the average of their prior beliefs. Such outcomes depend on the threshold values of both agents as well as the prior beliefs of their neighbors (again, to be discussed formally later).

**III. Identity Interaction ( $\gamma$ -type interaction)**

Lastly, with probability  $\gamma_{ij} = (1 - \alpha_{ij} - \beta_{ij})$ , both agents exhibit no change in their prior beliefs. This can best be describe as two individuals interacting with each other, yet they both disagree with the opinions of each other and are unwilling to change their own belief, or perhaps they both have identical beliefs prior to the interaction, and thus, there is no reason to change their beliefs.

### Specifying the Alpha, Beta, and Gamma ( $\alpha$ , $\beta$ , and $\gamma$ ) Values

We previously mentioned the four levels of influence that can be assigned to an agent in the network (regular, forceful, forceful+, or stubborn). The idea of implementing the different levels of forcefulness in the network is to enable more powerful leaders in society to have a greater probability of spreading their ideas in the network, while at the same time making it more likely that regular agents (and other less forceful agents) will have a greater probability of adopting the idea of a more influential person rather than the idea of a lesser influential person. Figure 3.3 (below) shows the  $\alpha_{ij}$ ,  $\beta_{ij}$ , and  $\gamma_{ij}$  values for all possible interactions between agents based on their level of influence—represented by the symbol notation which was discussed earlier. In total, there are 16 different sets of  $\alpha_{ij}$ ,  $\beta_{ij}$ , and  $\gamma_{ij}$  values—i.e. the active agent (agent  $i$ ) can be one of four possible levels of influence, and he interacts with agent  $j$ , who can also be one of four different influence levels.



*Figure 3.3 – Alpha, Beta, and Gamma Default Values*

For instance, looking at the first entry of the top-left table in Figure 3.3, we see that if a regular (square) agent is the active agent and interacts with another regular (square) agent, then they have a 100% chance of having a  $\beta$ -type (averaging) interaction. Meanwhile, the next entry shows that if a regular agent is the active agent and interacts with a forceful (upside-down triangle) agent then they have a 100% chance of having a  $\gamma$ -type (identity) interaction due to the fact that a regular agent is assumed not to be able to forcefully influence the belief of an agent with a higher level of influence. Note that order matters (i.e. who is the active agent and who is the neighboring agent that is selected by the active agent) when determining the specific  $\alpha_{ij}$ ,  $\beta_{ij}$ , and  $\gamma_{ij}$  values. Additionally, we make the assumption that all agents of the same level of influence have the same  $\alpha_{ij}$ ,  $\beta_{ij}$ , and  $\gamma_{ij}$  values. Although we agree agents of the same influence level may have different  $\alpha_{ij}$ ,  $\beta_{ij}$ , and  $\gamma_{ij}$  values (which might be more accurately obtained through interpersonal relationships with such individuals), due to lack of information we use the same values for all agents of the same influence level in order to simplify the model.

The default set of alpha, beta, and gamma values (Figure 3.3) were originally designed by Hung [8] to realistically mimic the types of interactions which occur between individuals in Pashtun societies in Afghanistan. The experiments in Chapter 4 will use these default values, but we will also experiment with different values in the parameter sensitivity analysis section for other types of societies, such as western, hierarchical, and consensus societies. Now that we have established the general network characteristics, agent characteristics, and types of pairwise interactions, we are ready to discuss the explicit formulation of the proposed threshold model.

### 3.2.5 Formulation of the Threshold Model

In this section, we detail the modeling approach and formulation of the threshold model which is designed to incorporate the concept of peer pressure into the interactions between agents. Thus, not only will the forcefulness (influence) level and belief of each agent determine the propagation of beliefs in the network during a pairwise interaction, but the beliefs of the neighbors of both agents engaging in a pairwise interaction will also influence the outcome (change in beliefs, if any) of such interactions. The threshold model naturally implements the effect peer influence has on interactions between agents by assigning all agents a specific threshold value,  $\tau_i \in [0,1], \forall i \in V$ . The threshold value,  $\tau_i$ , is the proportion of agent  $i$ 's

neighbors who must have beliefs either greater than (case 1) or less than (case 2) agent  $i$ 's prior belief in order for agent  $i$  to *agree* to change his belief *toward the desired direction* (either greater than or less than agent  $i$ 's prior belief). Therefore, we can interpret  $\tau_i$  as agent  $i$ 's willingness to be persuaded by the beliefs of his peers. The higher an agent's threshold value, the less likely he will be influenced during a pairwise interaction to change his belief away from the beliefs of his peers. For example, a threshold value of 0.1 means at least 10% of the agent's neighbors must have an attitude in the desired direction (greater than or less than) in order for the agent to be persuaded to update his belief toward the desired direction of pairwise communication.

We establish the threshold values for each agent depending on their individual level of influence. Thus, the higher an agent's level of influence is, the higher their threshold value will be. The assumption behind this stems from the notion that the more influential an agent is in the network, the more unlikely they will be to succumb to peer pressure from their neighbors, especially from those neighbors of lesser influence level. Thus, we establish the following threshold values for the agents in the network:

$$\tau_i = \begin{cases} 0.5 & \forall i \in V_R \\ 0.75 & \forall i \in V_F \cup V_{F+} \\ 1 & \forall i \in V_S \end{cases}$$

Note that the threshold value for the stubborn agents is only intended to be a placeholder, as stubborn agents have immutable beliefs which never change, and thus, peer pressure does not affect their beliefs. In Chapter 4, we experiment with different threshold values, including random assignment of threshold values, to determine the impact (if any) on the targeting strategies used by the stubborn agents.

### 3.2.5.1 Variable Notation and Definitions

Before we explain the formulation of the threshold model simulation procedure, we define the key variable notation and give their definitions (see Table 3.1).

<u>Notation</u>	<u>Description</u>
$X_i(t)$	<i>belief of agent i (active agent) at interaction t</i>
$X_j(t)$	<i>belief of agent j (selected neighboring agent) at interaction t</i>
$\tau_i$	<i>threshold value of agent i</i>
$\tau_j$	<i>threshold value of agent j</i>
$\eta_{i \setminus j}(t)$	<i>set of i's neighbors at interaction t (excluding j)</i>
$\eta_{j \setminus i}(t)$	<i>set of j's neighbors at interaction t (excluding i)</i>
$\eta_{i \setminus j, A}(t)$	<i>set of i's neighbors at interaction t (excluding j) who have beliefs <math>&gt; X_i(t)</math></i>
$\eta_{i \setminus j, B}(t)$	<i>set of i's neighbors at interaction t (excluding j) who have beliefs <math>&lt; X_i(t)</math></i>
$\eta_{j \setminus i, A}(t)$	<i>set of j's neighbors at interaction t (excluding i) who have beliefs <math>&gt; X_j(t)</math></i>
$\eta_{j \setminus i, B}(t)$	<i>set of j's neighbors at interaction t (excluding i) who have beliefs <math>&lt; X_j(t)</math></i>
$n$	<i>number of agents (nodes) in the network</i>
$\alpha_{ij}$	<i>prob. that agent i will attempt to forcefully impart <math>(1 - \tau_j)</math> of its attitude on agent j</i>
$\beta_{ij}$	<i>prob. that agent i will attempt to reach a consensus with agent j equal to the avg. of their prior beliefs</i>
$\gamma_{ij}$	<i>prob. that agents i and j will undergo an identity interaction and exhibit no change in their prior beliefs</i>
<b>Reminder:</b> $\alpha_{ij} + \beta_{ij} + \gamma_{ij} = 1, \forall i, j$ where $i \neq j$ (since agents cannot interact with themselves)	

Table 3.1 – Variable Notation Used In Threshold Model Simulation

### 3.2.5.2 Threshold Simulation Procedure

We now detail the threshold simulation procedure. More specifically, we describe the possible belief exchanges which can occur during the three types of pairwise interactions. As previously stated, once an agent is selected as the *active* agent (agent  $i$ ), he randomly chooses one of his neighbors (agent  $j$ ). Next, the two agents will randomly, yet probabilistically, engage

in one of the three types of interactions (either forceful, averaging, or identity) depending on the specific  $\alpha_{ij}$ ,  $\beta_{ij}$ , and  $\gamma_{ij}$  values.

Once the type of interaction has been selected, we determine the new beliefs for agents  $i$  and  $j$ . The previous models proposed by Hung [8] and Howard [13] are based on the same types of interactions, namely forceful (alpha), averaging (beta), and identity (gamma). However, in both of their modeling approaches, the belief exchange outcomes (formulas) which determine the interacting agents' new beliefs were slightly different in comparison to those used in the proposed threshold model. We first present Hung and Howard's pairwise interaction belief exchange formulas below:

**Forceful ( $\alpha$ -type interaction):**

With probability  $\alpha_{ij}$ , agent  $i$  'forcefully' imparts  $(1 - \varepsilon_{ij})$  of its attitude on agent  $j$ :

$$\begin{aligned} X_i(t+1) &= X_i(t) \\ X_j(t+1) &= \varepsilon_{ij} \cdot X_j(t) + (1 - \varepsilon_{ij}) \cdot X_i(t) \quad 0 \leq \varepsilon_{ij} \leq 0.5 \end{aligned}$$

where  $\varepsilon_{ij}$  represents the stubbornness for agent  $j$  interacting with agent  $i$

In 'forceful' interactions the parameter  $\varepsilon_{ij}$  is a rating of stubbornness for each agent. This parameter represents the amount of their own belief an agent will retain after being forcefully influenced by another agent. For simplicity this is assumed to be identical for all agent pairs.

$$\varepsilon_{ij} = \varepsilon \quad \forall i, j \in \{1, 2, \dots, n\}$$

**Averaging ( $\beta$ -type interaction):**

With probability  $\beta_{ij}$ , they reach a consensus equal to the average of their prior attitudes:

$$X_i(t+1) = X_j(t+1) = \frac{X_i(t) + X_j(t)}{2}$$

**Identity ( $\gamma$ -type interaction):**

With probability  $\gamma_{ij}$ , both agents exhibit no change in attitude:

$$\begin{aligned} X_i(t+1) &= X_i(t) \\ X_j(t+1) &= X_j(t) \end{aligned}$$

Due to the properties of the threshold model, modifications are required to the previous belief exchange outcomes resulting from pairwise interactions. In the threshold model, belief exchanges during pairwise interactions are determined not only by the prior beliefs of agents  $i$  and  $j$ , but also by their individual threshold values as well as the beliefs of their neighbors. Thus, more possible outcomes may occur during the interactions which we must account for. We detail the specific belief exchange outcomes (formulas) for each type of interaction in the proposed threshold model below:

I. ***ALPHA INTERACTION***: If an alpha (forceful) interaction occurs (which happens with probability  $\alpha_{ij}$ ) the following belief outcomes are possible:

A. If  $X_i(t) < X_j(t)$  and  $\frac{|\eta_{j \setminus i, B}(t)|}{|\eta_{j \setminus i}(t)|} \geq \tau_j$ , then agent  $i$  forcefully imparts  $(1 - \tau_j)$  of his belief on agent  $j$ , while agent  $j$ 's neighbors (excluding agent  $i$ ) impart  $\tau_j$  of their average belief on agent  $j$ :

$$\begin{aligned} X_i(t+1) &= X_i(t) \\ X_j(t+1) &= \tau_j * \frac{\sum_{k \in \eta_{j \setminus i, B}(t)} X_k(t)}{|\eta_{j \setminus i, B}(t)|} + (1 - \tau_j) * X_i(t) \end{aligned} \quad (3.1)$$

B. If condition 'A' does not hold and  $X_i(t) < X_j(t)$ , then with probability  $(1 - \tau_j)$  agent  $i$  forcefully imparts  $(1 - \tau_j)$  of his belief on agent  $j$ :

$$\begin{aligned} X_i(t+1) &= X_i(t) \\ X_j(t+1) &= \tau_j * X_j(t) + (1 - \tau_j) * X_i(t) \end{aligned} \quad (3.2)$$

C. If  $X_i(t) > X_j(t)$  and  $\frac{|\eta_{j \setminus i, A}(t)|}{|\eta_{j \setminus i}(t)|} \geq \tau_j$ , then agent  $i$  forcefully imparts  $(1 - \tau_j)$  of his belief on agent  $j$ , while agent  $j$ 's neighbors (excluding agent  $i$ ) impart  $\tau_j$  of their average belief on agent  $j$ :

$$\begin{aligned} X_i(t+1) &= X_i(t) \\ X_j(t+1) &= \tau_j * \frac{\sum_{k \in \eta_{j \setminus i, A}(t)} X_k(t)}{|\eta_{j \setminus i, A}(t)|} + (1 - \tau_j) * X_i(t) \end{aligned} \quad (3.3)$$

D. If condition ‘C’ does not hold and  $X_i(t) > X_j(t)$ , then with probability  $(1 - \tau_j)$  agent  $i$  forcefully imparts  $(1 - \tau_j)$  of his belief on agent  $j$ :

$$\begin{aligned} X_i(t+1) &= X_i(t) \\ X_j(t+1) &= \tau_j * X_j(t) + (1 - \tau_j) * X_i(t) \end{aligned} \quad (3.4)$$

E. If no previous conditions hold, then no change in belief occurs:

$$\begin{aligned} X_i(t+1) &= X_i(t) \\ X_j(t+1) &= X_j(t) \end{aligned} \quad (3.5)$$

II. ***BETA INTERACTION***: If a beta (averaging) interaction occurs (which happens with probability  $\beta_{ij}$ ), the following belief outcomes are possible:

First, we determine if agent  $i$  will change his belief to the consensus average of the two agents’ prior beliefs, or if no change in belief will occur for agent  $i$ .

A. If  $X_i(t) < X_j(t)$  and  $\frac{|\eta_{i \setminus j, A}(t)|}{|\eta_{i \setminus j}(t)|} \geq \tau_i$ , then agent  $i$ ’s new belief is the average of the two agents’ prior beliefs:

$$X_i(t+1) = \frac{X_i(t) + X_j(t)}{2} \quad (3.6)$$

B. If  $X_i(t) > X_j(t)$  and  $\frac{|\eta_{i \setminus j, B}(t)|}{|\eta_{i \setminus j}(t)|} \geq \tau_i$ , then agent  $i$ ’s new belief is the average of the two agents’ prior beliefs:

$$X_i(t+1) = \frac{X_i(t) + X_j(t)}{2} \quad (3.7)$$

C. If neither conditions ‘A’ or ‘B’ hold, then no change in belief occurs for agent  $i$ :

$$X_i(t+1) = X_i(t) \quad (3.8)$$

Next, we determine if agent  $j$  will change his belief to the consensus of the two agents' prior beliefs, or if no change in belief will occur for agent  $j$ .

D. If  $X_j(t) < X_i(t)$  and  $\frac{|\eta_{j \setminus i, A}(t)|}{|\eta_{j \setminus i}(t)|} \geq \tau_j$ , then agent  $j$ 's new belief is the average of the two agents' prior beliefs:

$$X_j(t+1) = \frac{X_i(t) + X_j(t)}{2} \quad (3.9)$$

E. If  $X_j(t) > X_i(t)$  and  $\frac{|\eta_{j \setminus i, B}(t)|}{|\eta_{j \setminus i}(t)|} \geq \tau_j$ , then agent  $j$ 's new belief is the average of the two agents' prior beliefs:

$$X_j(t+1) = \frac{X_i(t) + X_j(t)}{2} \quad (3.10)$$

F. If neither conditions 'D' or 'E' hold, then no change in belief occurs for agent  $j$ :

$$X_j(t+1) = X_j(t) \quad (3.11)$$

III. **GAMMA INTERACTION:** Lastly, if neither an alpha (forceful) or beta (averaging) interaction occurs, then agents  $i$  and  $j$  will engage in a gamma (identity) interaction since we define:  $\alpha_{ij} + \beta_{ij} + \gamma_{ij} = 1$ . The only outcome of the gamma interaction is presented below as both agents maintain their prior beliefs:

$$\begin{aligned} X_i(t+1) &= X_i(t) \\ X_j(t+1) &= X_j(t) \end{aligned} \quad (3.12)$$

Finally, once the new beliefs of agents  $i$  and  $j$  are determined, we update the network belief vector for the next interaction time step,  $X(t+1) \in \mathbb{R}^{n \times 1}$ . Note that only the indices for the two interacting agents (agents  $i$  and  $j$ ) can possibly change when we compare  $X(t)$  to  $X(t+1)$ . The beliefs of all other agents will remain the same. Once the belief vector  $X(t+1)$  has been determined, a new *active* agent is randomly selected and the same process is repeated until the desired number of pairwise interactions have occurred in the network.

Thus, we see from the models proposed by Hung [8] and Howard [13] that regardless of the characteristics of the two interacting agents and the beliefs of their neighbors, the new beliefs of the interacting agents are solely dependent on the type of interaction and the prior beliefs of the interacting agents. Meanwhile, the presence of the threshold model, in which peer pressure affects agents' decisions on whether or not to update their beliefs, forces us to implement additional outcome possibilities for the alpha (forceful) and beta (averaging) interactions since these interactions not only involve threshold values, but also the beliefs of the neighbors of each interacting agent. During all three types of interactions in the threshold model, if the criteria is not satisfied to enable an agent to update his/her belief, then no change in belief is seen from interaction  $t$  to  $t + 1$ . We see examples of such outcomes in equations (3.5), (3.8), and (3.11).

### 3.3 Transient, Dynamic Two-Player Game Description

With the detailed simulation procedure of the threshold model, we describe below the implementation of the two-player transient, dynamic game involving the two stubborn agents—the U.S. and Taliban. We discuss the main assumptions about the model with respect to the formulation of the two-player game, followed by the objective of the game and its parameters.

#### 3.3.1 Assumptions

In formulating the two-player game on the social network, we make the following key assumptions:

1. ‘No man is an island’ [7] in the networks which are created by the network generator tool, which means that each network is connected such that there exists a direct or indirect path from every node to all other nodes in the network. Moreover, this assumption means that every agent in the network communicates with and is capable of being influenced by someone else in the network (with the important exception that stubborn agents have immutable beliefs).
2. The social network has static edges such that all agents and their locations are known throughout the entirety of the game. We make one exception to this assumption in that stubborn agents may update their connection strategies during

the game as they see fit, which allows for the edges from stubborn agents to be dynamic in this sense.

3. The same number of connections by each stubborn agent to mutable agents is predetermined at the beginning of the game and will remain the same throughout the entire game.
4. All agents in the network have an equal, uniform probability of ‘being selected’ as the *active* agent for pairwise interactions.
5. The influence-type probabilities are known and remain constant among all agents of the same influence level.
6. Lastly, we use Monte Carlo simulations for determining the propagation of beliefs in the network over time.

### **3.3.2 Dynamic Game Parameters**

Next, we define the various parameter inputs and settings for the dynamic game which are determined by the user.

#### **3.3.2.1 Number of Realizations**

The number of realizations refers to the number of replications or iterations of the Monte Carlo simulation for the pairwise interactions of the threshold model. We have found that generally 30 (sometimes even less) realizations of the Monte Carlo simulation give an accurate picture of the evolution of beliefs in the network over time.

#### **3.3.2.2 Number of Steps**

The number of steps can be interpreted as the total number of ‘turns’ or ‘moves’ to be allowed by the stubborn agents during the game. During a step, only one stubborn agent (either US or TB—whoever’s turn it is) may evaluate their current strategy (which includes looking at their opponent’s strategy since we assume a perfect information game) and subsequently decide what their new strategy is during the current turn based on this information. For instance, if the number of steps in the game is 10, each stubborn agent will be given five steps (or turns) during the game during which they may update their strategy in order to maximize their payoff. The

steps during the game follow an alternating pattern between the two opponents such that no stubborn agent is allowed to make two moves in a row.

### **3.3.2.3 Expected Number of Interactions per Step ( $\mu$ )**

The expected number of interactions per step ( $\mu$ ) is simply the number of pairwise interactions which are expected to occur during each step of the game. This can be interpreted as the number of interactions we expect to occur during each step before the opposing stubborn agent wishes to update his strategy, and thus, end the current step (turn) of the other stubborn agent. The number of interactions per step follows a normal distribution with a mean ( $\mu$ ) and standard deviation ( $\sigma$ ). By allowing variability to exist in the number of interactions per step, we add more flexibility to, and increase the realism of, the two-player game. We discuss the standard deviation parameter,  $\sigma$ , next.

### **3.3.2.4 Standard Deviation of the Number of Interactions per Step ( $\sigma$ )**

The standard deviation of the number of interactions per step ( $\sigma$ ) is an input into the game and determines how much variability and uncertainty exists in the number of interactions per each step in the game. During the experiment analysis in Chapter 4 we analyze a more symmetric game, whereby each stubborn agent ‘receives’ an equal amount of interactions per step, and thus, we set  $\sigma = 0$ .


### **3.3.2.5 Number of Connections per Stubborn Agent ( $z$ )**

As previously discussed, each stubborn agent has the same number of connections (strategies), which is predetermined at the start of the game. Once a stubborn agent selects a strategy at the beginning of one of his steps (turns) in the game, the next opportunity to update his strategy will be after the subsequent completion of the opposing stubborn agent’s next step.

### **3.3.2.6 Initial Connection Strategy for One Stubborn Agent**

In order to initiate the start of the game, we must set an initial strategy for one stubborn agent. We let step 0 represent the initial state of the network with an initial strategy established for one stubborn agent. The game begins at step 1, whereby the opposing stubborn agent

analyzes the initial network from step 0 and subsequently chooses his best strategy based on this information. Figure 3.4 below shows an example network with the initial strategy for the TB agent set during step 0 and the first strategy chosen by the US agent during step 1 in order to maximize his payoff given the TB agent’s initial strategy from step 0.



*Figure 3.4 – Step 0 (Initial Strategy) and Step 1 (First Strategy) Example*

### 3.3.2.7 Penalty for Identical Strategy ( $\lambda$ )

We enforce a ‘penalty’ for a stubborn agent who connects to a mutable agent who is also currently connected to the opposing stubborn agent. The motivation behind implementing this  $\lambda$  penalty factor is two-fold. First, in a realistic setting, if a person is talking to (and previously succumbing to pressure from) an influential person such as a stubborn agent, then another influential person of opposite belief trying to subsequently persuade that same person toward his own belief will have a more difficult time persuading such a person. Thus, it is realistic to implement a penalty on the stubborn agent in situations such as these in order to mimic reality, and therefore, decrease that stubborn agent’s effectiveness. Second, in Iraq or Afghanistan, it is important for leaders to gain steady support over time for the U.S. cause, while staying away from high risk strategies that can be extremely variable. Because instances where a US and TB stubborn agent connect to the same agent result in higher variances in terms of network beliefs, the US would most likely prefer to stay away from such strategies.

For instance, if the US agent,  $V_{US}$ , connects to a mutable agent who is currently a strategy connection for the TB agent,  $V_{TB}$ , then there is a penalty inflicted which impacts the alpha and



function called ‘risk averse mean belief’ which was an attempt to prevent pure Nash equilibrium strategy profiles with identical strategies for both players. The ‘risk averse mean belief’ payoff function,  $h(S_{US}, S_{TB})$ , depends on both the predicted standard deviation and the equilibrium mean belief seen below:

$$h(S_{US}, S_{TB}) = \text{Mean Belief Payoff} - \lambda * \text{Predicted Std Dev of Network}$$

$$h_i(S_{US}, S_{TB}) = f_i(S_{US}, S_{TB}) - \lambda * \sigma(X^*(S_{US}, S_{TB})) \quad i \in \{US, TB\}$$

The penalty parameter ( $\lambda$ ) is set to a positive value (for instance, 1). Although this alternative ‘penalty’ implementation may be useful for some situations, there are a couple problems which limit its usefulness for the transient-threshold, two-player game. First, the predicted standard deviation is in the form of a bound (not an actual number), but more importantly, the standard deviation is for the mean, long run equilibrium belief of the network for the non-threshold model. Thus, in order to use a similar approach, we obtain the standard deviation of the network beliefs through the use of simulations. Overall, our penalty implementation appears to not only be simpler to implement, but also more easily understood from the intuitive sense of using the penalty to affect the likelihood of the types of interactions and thus indirectly affecting the payoffs, rather than directly affecting the payoffs.

### 3.3.3 Payoff Function

The last input for the game is selecting the payoff function for each stubborn agent. The goal of each stubborn agent is to select a set of strategy connections throughout the game which will allow them to maximally influence the mutable agents in the network. However, the notion of what ‘maximal influence’ is could depend on the specific payoff function selected by the stubborn agents. Below we introduced two different payoffs for the two-player game—Mean Belief and Number of Agents (Nodes) Won.

#### Mean Belief

If the ‘mean belief’ payoff function is selected, each player will maximally influence the network by selecting the mutable agents as strategies which will move the mean belief of all mutable agents in the network toward their side. For example, if the US agent has higher

influence in the network compared to the TB agent, then the mean belief of the mutable agents in the network will move towards +0.5. In the reverse case, if the strategies chosen by the TB agent are more influential than those chosen by the US agent, then the mean belief of the mutable agents will move towards -0.5. Thus, the US agent wants the mean belief of the network to be as high as possible (close to +0.5), while the TB agent wants the network mean belief to be as low as possible (close to -0.5).

The ‘mean belief’ payoff function, denoted  $f(\cdot)$ , is defined as:

$$f_{US}(S_{US}(t), S_{TB}(t)) = \frac{1}{|V_M|} \sum_{k=1}^{|V_M|} X_k(t) = -f_{TB}(S_{US}(t), S_{TB}(t)) \quad (3.13)$$

### Number of Agents (Nodes) Won

If the ‘number of agents (nodes) won’ payoff function is chosen, each player will maximally influence the network by choosing the mutable agents as strategies which will yield the highest number of mutable agents toward supporting their side. Each player is awarded 1 point for every mutable agent who is deemed to be in support of their side. A mutable agent is deemed to be in support of a player’s side if their belief is above/below a predetermined buffer threshold (B), where  $B \in [0, 0.5)$ . We present the point system for the Number of Agents Won payoff function below in Table 3.2.

US Points Awarded	Mutable Agent Belief	TB Points Awarded
+1	> B	-1
0	[-B,B]	0
-1	< -B	+1

*Table 3.2 – Number of Agents Won Point System*

For instance, if  $B = 0.1$ , then the US and TB agents are awarded +1 and -1 points, respectively, for all mutable agents with beliefs above +0.1. Subsequently, the TB and US agents are awarded +1 and -1 points, respectively, for all mutable agents with beliefs below -0.1. No points are awarded to either player for indecisive mutable agents with beliefs between [-0.1, +0.1]. Thus, we see that the range of possible discrete values for this payoff function is on the interval  $[-|V_M|, |V_M|]$ , while the range of possible continuous values for the ‘mean belief’ payoff is on the interval  $[-0.5, 0.5]$ .

More specifically, the ‘number of agents (nodes) won’ payoff function, denoted  $g(\cdot)$ , is defined as:

$$h_{US}(x) \begin{cases} +1 & \text{if } x > B \\ 0 & \text{if } x \in [-B, B] \\ -1 & \text{if } x < -B \end{cases} \quad (3.14)$$

$$h_{TB}(x) \begin{cases} -1 & \text{if } x > B \\ 0 & \text{if } x \in [-B, B] \\ +1 & \text{if } x < -B \end{cases} \quad (3.15)$$

$$g_i(S_{US}(t), S_{TB}(t)) = \sum_{k=1}^{|V_M|} h_i(X_k(t)) \quad i \in \{US, TB\} \quad (3.16)$$

Note:  $g_{US}(S_{US}(t), S_{TB}(t)) = -g_{TB}(S_{US}(t), S_{TB}(t))$

### Zero Sum vs. Non-Zero Sum Games

If both players choose the same payoff function, either  $f_{US}(\cdot)$  and  $f_{TB}(\cdot)$  or  $g_{US}(\cdot)$  and  $g_{TB}(\cdot)$ , then both players are playing a zero sum game. A zero sum game occurs whenever the payoff for one player is exactly the opposite payoff for the other player. We call these games zero sum games because summing up both players payoffs at any particular time during the game will result in the summation equaling zero:

$$f_{US}(S_{US}(t), S_{TB}(t)) + f_{TB}(S_{US}(t), S_{TB}(t)) = 0 \quad \forall t$$


$$g_{US}(S_{US}(t), S_{TB}(t)) + g_{TB}(S_{US}(t), S_{TB}(t)) = 0 \quad \forall t$$

However, any combination of different payoff functions will not result in a zero sum game. Therefore, different combinations of  $f(\cdot)$  and  $g(\cdot)$ , such as  $f_{US}(\cdot)$  and  $g_{TB}(\cdot)$ , result in non-zero sum games.

### Absolute Payoff vs. Relative Payoff

There are two ways we can report the payoffs for each player—(1) the true absolute value of the payoff which is directly calculated by each payoff function, or (2) the relative value which takes into account the initial starting conditions of the network. The reason we are interested in

knowing the relative payoff, besides solely the absolute payoff, is to gauge how much the payoff for each player has changed relative to the initial payoff. The relative payoff allows us to better analyze the progress a player has made, which can sometimes be unapparent when strictly viewing the absolute payoff. We emphasize this point through the use of the following example shown below in Figure 3.6:



*Figure 3.6 – Absolute Payoff vs. Relative Payoff Example*

In the above example, both players used the ‘mean belief’ payoff,  $f(\cdot)$ , and a total of 6 steps occurred during the game (3 for each player, starting with step 1 for the US agent) with 100

pairwise interactions per step. The initial mean belief of the network was -0.5, meaning that all mutable agents very strongly favored the Taliban. The absolute ‘mean belief’ payoff at the end of the game with respect to the US agent is -0.3832, which means the average belief of the mutable agents in the network is still strongly in favor of the Taliban. Therefore, from the absolute perspective of the network beliefs, the Taliban agent won the game. However, the relative ‘mean belief’ payoff at the end of the game with respect to the US agent is +0.1168, which indicates that the US agent won the game from the relative perspective since on average, each mutable agent in the network moved +0.1168 in their mean belief toward the US side from the start of the game. Thus, we see that both payoff measures (absolute and relative) are necessary since they serve different purposes for how we can interpret the outcome of the game.

### Strategy Profile and Payoff Table

Finally, we show an example strategy profile and payoff table below (Table 3.3), which illustrates the progression of the two-player, dynamic game over 4 steps based on the initial network shown in Figure 3.4. Because there are 4 steps in this example game, each player is given 2 steps (opportunities) to update their strategies, where step 0 does not count. The initial step, step 0, is the initial connection strategy {2} of the TB stubborn agent which is used to initiate the start of the game. The payoff used for both players is the mean belief payoff from (3.13). Because the initial mean belief of the network is 0, the relative payoff is identical to the absolute payoff.

	Step #	TB Strategy	US Strategy	Absolute Payoff	Relative Payoff
<b>Initial TB Strategy (no US Strategy)</b>	0	2	--	0.0000	0.0000
<b>1st US step (TB Strategy Fixed from Step 0)</b>	1	2	7	0.0156	0.0156
<b>1st TB step (US Strategy Fixed from Step 1)</b>	2	12	7	0.0045	0.0045
<b>2nd US step (TB Strategy Fixed from Step 2)</b>	3	12	2	-0.0056	-0.0056
<b>2nd TB step (US Strategy Fixed from Step 3)</b>	4	7	2	-0.0342	-0.0342

*Table 3.3 – Example Strategy Profile and Payoff Table*

### 3.3.4 Finding Strategies to Improve Payoff

Lastly, we describe the approach both players take in locating strategies to improve their payoffs during the game. Due to the nature of the game, whereby both players know the other player's strategy, the network topology (meaning the influence level of all agents and their connections), and the beliefs of all agents throughout the entirety of the game, no player has an advantage from an intelligence viewpoint when selecting their strategies. However, given all this information, the task of determining which agent(s) to choose as strategies at any particular time during the game can be quite overwhelming. Below we discuss the strategy space which is searched by the players in order to select strategies to improve their payoff, and later, we propose heuristics which drastically reduce the time it takes to find solutions. As the complexity of the network and two-player game increases, we find that the use of heuristics are necessary as the computational times become too burdensome.

#### 3.3.4.1 Players and Strategies

In the two-player game there are two opposing stubborn agents of opposite, immutable beliefs (refer to the belief spectrum shown earlier in Figure 3.1). The United States (US) agent has a belief of +0.5, while the Taliban (TB) agent has a belief of -0.5. The *objective* of the game for each stubborn agent is to select mutable agents to communicate with in order to maximally spread their influence (belief) throughout the network. We measure a stubborn agent's success in the game by determining their specific payoff, which will be discussed in a later section.

Although there is only one stubborn agent per side in this two-player game, each stubborn agent has  $z > 0$  connections to mutable agents in the network. Each stubborn agent will have the same number of connections they will be allowed to make during the game. Thus, if  $z = 3$ , for example, the US agent and TB agent will each have three connections they can make to mutable agents in the network throughout the course of the game. We will refer to these connections as strategies, and they represent communication links between the targeted mutable agents in the network and the stubborn agent. We use the following notion:

$S(t)$  = set of all possible strategies available at time  $t$  (identical for both players)

$$|S(t)| = \binom{|V_M| + z - 1}{z} = \frac{(|V_M| + z - 1)!}{z! (|V_M| + z - 1 - z)!} = \frac{(|V_M| + z - 1)!}{z! (|V_M| - 1)!}$$

$S_{US}(t)$  = strategy chosen by US agent at time  $t$

$S_{TB}(t)$  = strategy chosen by TB agent at time  $t$

$$|S_{US}(t)| = |S_{TB}(t)| = z \quad \forall t$$

where time  $t$  refers to the total number of interactions that have occurred in the network

The number of possible strategies is on the order of  $|V_M|^z$ , but more precisely, the set is determined by using a combination with repetition because (1) we do not care about order (i.e. the strategy  $\{1, 2\}$  is the same as the strategy  $\{2, 1\}$ ), and (2) we allow repetition of strategies (i.e. the strategy  $\{1, 1\}$  is allowed). We allow repetition of the same targeted agent in strategies in order to mimic the possibility for a stubborn agent wishing to place more of his time and effort on a particular agent, rather than spreading his efforts to multiple agents. Also, although we mention in our assumptions that we only consider a static network concerning the number of agents (nodes), we prefer the notation  $S(t)$  instead of simply,  $S$ , to describe the set of all strategies during the game in order to allow the possibility of future work dealing with dynamic, evolving networks in which the list of possible strategies changes over time.

For instance, in the example network diagram (see Figure 3.2), given we know  $t = 1$ , then we get the following:

$$S(1) = \{1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15, 16\}$$

$$|S(1)| = \binom{16 + 1 - 1}{1} = \binom{16}{1} = 16$$

$$S_{US}(1) = \{12\}$$

$$S_{TB}(1) = \{2\}$$

$$|S_{US}(1)| = |S_{TB}(1)| = 1$$

Throughout the course of the game, each stubborn agent will be given opportunities in which they may update (change) their current connection strategy to a more favorable strategy. In this sense, we view the two-player game as a dynamic “chess game” such that the stubborn agents will constantly reevaluate their current strategy (position) during the game in relation to their opponent’s strategy and then make subsequent strategy changes based on these evaluations. However, we view the game as a perfect information game, such that during all times of the game, each stubborn agent knows the opposing stubborn agent’s strategy. We leave the notion of imperfect information games, whereby stubborn agents are unsure of the exact strategies of their opponent, for future work.

### **3.3.4.2 Exhaustive Enumeration**


Locating strategies through the use of exhaustive enumeration means we calculate all possible strategies available to a player and choose the strategy which yields the highest payoff. The clear advantage of using exhaustive enumeration is the guarantee that the best empirical strategies will be found because we have searched the entire strategy space. However, due to the size of the strategy space, which can exponentially increase with the number of connections each player makes, the time it takes to locate such strategies can become quite cumbersome. Circumstances involving very large networks with multiple connections per stubborn agent present an obvious tradeoff scenario for the players between the degree of optimality they are willing to accept and the run time needed for evaluating the best strategies. Therefore, we must explore additional measures each player can take in order to significantly reduce simulation run times without significantly jeopardizing the quality of the chosen strategies.

### **3.3.4.3 ‘Selective Search’ Heuristic**

Heuristics are “problem-solving methods which tend to produce efficient solutions to difficult problems by restricting the search through the space of possible solutions” [20]. Similar to the way humans utilize heuristics in order to keep the information processing demands of a task within the bounds of limited cognitive capacity [21], we also implement heuristics in order to decrease the size of the strategy space that is searched. A simple method a player can implement in order to decrease the time necessary to evaluate the best available strategy is by first determining which strategies are consistently observed as being empirically suboptimal in

comparison to other available strategies. If a player can determine what characteristics inherently make a strategy worse than other strategies, then a heuristic can be used by the player during the game which reduces the strategy space that is searched by not including strategies which are deemed suboptimal. Due to the presence of more influential agents in the network, both Hung and Howard repeatedly found that ‘regular’ agents were very rarely selected as strategies. Only in networks strictly dominated by regular agents and lacking agents of higher influence level did regular agents become chosen strategies. Thus, because regular agents are the least influential agents in the networks, connecting to them does not typically maximize a player’s payoff because regular agents do not have the ability to forcefully influence other agents.

Similar to our predecessors, we also find that players in the two-player game on the threshold model rarely choose regular agents as strategies to maximize their payoffs, given the presence of agents who are more influential in the network. Only very special network structures, which do not necessarily mimic a realistic society, do we find cases involving regular agents being selected as strategies. Figure 3.7 shows a line graph example where a regular agent is chosen as a strategy by the U.S. agent.



*Figure 3.7 – Line Graph Example with Regular Agent Chosen as Strategy*

Based on the consistent observation that regular agents are rarely chosen as strategies by the players, we implement the ‘Selective Search’ Heuristic for both players in order to reduce the computational time needed to find strategies to improve a player’s payoff.

## Description

The ‘Selective Search’ Heuristic removes all strategy combinations involving at least one regular agent from the list of possible strategies in the strategy space, and subsequently uses exhaustive enumeration on the remaining strategies (all combinations with repetition involving only forceful and/or forceful+ agents) to determine the best available strategy for a player. Thus, this heuristic is designed to exhaustively enumerate only a subset of the entire strategy space since it removes all strategies containing regular agents. Although using the ‘Selective Search’ Heuristic removes the guarantee of finding the best empirical solutions when compared to using exhaustive enumeration, we find that the strategies found by the heuristic almost exclusively match the strategies found using exhaustive enumeration, with the few exceptions involving special network cases.

### DEFINITION:

The ‘Selective Search’ Heuristic reduces the size of the strategy space at time  $t$ ,  $|S(t)|$ , from  $|S(t)| = \binom{|V_M|+z-1}{z}$  under exhaustive enumeration to  $|S(t)| = \binom{|V_{F+U}|+z-1}{z}$ .

### Example

In order to better illustrate the power of this simple, yet quite effective heuristic, we return to an example involving the network shown in Figure 3.2. Table 3.4 below lists all possible strategies that are searched using both exhaustive enumeration and the ‘Selective Search’ Heuristic, given each player only has one connection ( $z = 1$ ). As we see, the simple heuristic is able to reduce the entire strategy space from 16 to 4 strategies since the network contains 12 regular agents {3-6, 8-11, 13-16}, 3 forceful agents {2, 7, 12}, and 1 forceful+ agent {1}.

	Exhaustive Enumeration	Selective Search
Strategy Space:	{1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15, 16}	{1, 2, 7, 12}

*Table 3.4 – Strategy Space Comparison: Exhaustive Enumeration vs. ‘Selective Search’*

More importantly, for scenarios in which we have multiple connections per stubborn agent ( $z > 1$ ), the strategy space can be exponentially reduced depending on the number of regular agents which are in the network. For example, using the same network in Figure 3.2, if both players are allowed two connections ( $z = 2$ ) then the size of the strategy space is reduced from  $|S(t)| = \binom{|V_M|+z-1}{z} = \binom{16+2-1}{2} = 136$  (exhaustive enumeration) to  $|S(t)| = \binom{4+2-1}{2} = 10$  ('Selective Search'). Thus, we see how powerful this heuristic can be at significantly reducing the run time by reducing the size of the strategy space.

#### 3.3.4.4 'Sequential Search' Greedy Algorithm

The 'Selective Search' Heuristic offers us one way to significantly reduce the size of the strategy space, of which the degree in reduction depends on the number of regular agents in the network. However, for large networks with many influential agents in which each stubborn agent has multiple connections, the run time can still be quite burdensome, even after removing all combinations of regular agents from the strategy space. Thus, we implement another heuristic 'Sequential Search' which is a greedy algorithm specifically designed for reducing the run time for cases involving  $z > 1$  connections per stubborn agent. The use of a greedy algorithm to improve the computational time for finding strategies which maximize belief propagation in social networks has been previously proposed by Kempe et al. [22], and we similarly propose the use of a greedy algorithm for selecting strategies for complex network problems.

#### Description

The 'Sequential Search' Greedy Algorithm further reduces the size of the strategy space after the 'Selective Search' Heuristic only if the number of connections per player (stubborn agent) is greater than one ( $z > 1$ ). The idea behind this greedy approach is to solve the best connections for a player 'sequentially' rather than simultaneously. In other words, the greedy algorithm implements the idea of finding and adding the single best connection (one at a time) to the strategy profile for a player that yields the maximum marginal return to the player's payoff function. Thus, rather than simultaneously searching all available combinations of size 'z', we find the best available connection one at a time until the entire strategy profile is of length 'z'.

DEFINITION:

The ‘Sequential Search’ Greedy Algorithm reduces the size of the strategy space at time  $t$  from  $|S(t)| = \binom{|V_{F+} \cup V_F| + z - 1}{z}$  under ‘Selective Search’ to  $|S(t)| = (|V_{F+} \cup V_F| * z)$ . If  $z = 1$ , the two strategy spaces become equivalent.

### Greedy Algorithm Summary

Below we detail the steps performed by the greedy algorithm, ‘Sequential Search’:


- Given  $step = k - 1$  at time  $t - 1$  and  $step = k$  at time  $t$
- Given payoff function for each player:  $Payoff \in \{f(\cdot), g(\cdot)\}$
- The player which updates their strategy at step  $k$  and time  $t$  is designated as:  $P \in \{US, TB\}$
- The opposing player which maintains a fixed strategy when player  $P$  is updating their strategy is designated as:  $F \in \{US, TB\}$

1.  $S_P(t) \leftarrow null, best_{strat} \leftarrow null, current\_time \leftarrow t$
2. **WHILE**  $|S_P(t + 1)| < z$
3. Given  $S_P(t)$ , the opposing player’s strategy,  $S_F(t)$ , which is fixed and of magnitude ‘ $z$ ’, and the strategy space,  $S(t)$ , find the single best connection strategy,  $best_{strat}$ , which yields the highest payoff for player  $P$  after time  $t + \mu$
4.  $S_P(t) \leftarrow [S_P(t) \text{ } best_{strat}]$
5.  $best_{strat} \leftarrow null, t \leftarrow current\_time$
6. **END**

### Example

In order to get a better understanding of the ability of this greedy algorithm to reduce the strategy space that is searched, we explore the following example. Figure 3.8 shows a 73 node network (excluding the two players), of which there are 35 agents who have either a forceful+ or forceful level of influence (i.e.  $|V_{F+} \cup V_F| = 35$ ). We assume each player is allowed  $z = 3$  strategy connections. Thus, once the ‘Selective Search’ Heuristic is used, the size of the strategy

space is reduced from  $|S(t)| = \binom{|V_M|+z-1}{z} = \binom{73+3-1}{3} = 67,525$  to  $|S(t)| = \binom{35+3-1}{3} = 7,770$ . Although we have considerably reduced the size of the strategy space through the ‘Selective Search’ Heuristic, exhaustively searching all 7,770 different strategies during each realization of the Monte Carlo simulation still presents a sizeable run time problem.



*Figure 3.8 – Example 73 Agent Network*

However, upon using the greedy algorithm, after first reducing the strategy space through the use of the ‘Selective Search’ Heuristic, we are able to further reduce the strategy space from:

$$|S(t)| = \binom{|V_{F+} \cup V_F| + z - 1}{z} = \binom{35 + 3 - 1}{3} = 7,770$$

To

$$|S(t)| = (|V_{F+} \cup V_F| * z) = (35 * 3) = 105$$

We provide further detail below on how a player uses the ‘Sequential Search’ Greedy Algorithm during the two-player game to significantly reduce the size of the strategy space that is searched during a given step k of the game.

This example shows how efficient the ‘Sequential Search’ Greedy Algorithm can be at significantly reducing the strategy space for cases where  $z > 1$ . For cases where  $z = 1$ , we see that no further reductions in the strategy space are made by the greedy algorithm. One pertinent question that needs answering in relation to the greedy algorithm is how the potentially drastic reduction in strategy space affects the quality of the solutions in comparison to exhaustive enumeration. Such quality consequences of the ‘Sequential Search’ Greedy Algorithm will be analyzed in Chapter 4.

### 3.4 The Necessity of Simulations in Calculating Belief Propagation

The previous social influence network models used by Hung [8] and Howard [13] have one clear advantage over our proposed threshold model—the existence of an analytic expression for calculating the long-term expected beliefs of the agents in the network. Due to the nature of the threshold model, in which the diffusion of beliefs depends on the current beliefs of the agents at every interaction, an analytic expression which calculates the expected belief of the network  $k > 1$  interactions into the future has not been found to exist. Thus, we must employ the use of Monte Carlo simulations in order to achieve a reasonable understanding of the behavior of the threshold model.

### 3.5 Modeling Formulation Summary

In this chapter, we described the formulation of a proposed threshold model which is a stochastic, pairwise-interaction model with mutable agents whose beliefs can change and with immutable (stubborn) agents whose beliefs do not change. The purpose of this formulation was to incorporate more realistic assumptions in previously proposed models [8,13] such as:

1. The long-term beliefs of the mutable agents in the network should be *dependent* on the initial beliefs of those mutable agents, despite the presence of immutable stubborn agents.
2. Introduce a dynamic, two-player game whereby each player controls a predetermined number of connections to mutable agents in the network, and the objective of the game for each player is to maximize their respective payoff function.

3. We shifted the previous focus from analyzing the long-term asymptotic equilibrium behavior of the network to instead analyzing the transient behavior of the network during the course of the two-player game. We will see in Chapter 4 that both players continually evaluate, change, and reevaluate their strategies throughout the course of the game, while not being content on settling on identical Nash Equilibria strategies for the long-run.
4. Peer pressure should affect the outcome of belief exchanges during pairwise interactions.
5. Lastly, we implemented the ability to penalize the payoffs for identical strategies among opposing stubborn agents (U.S. and Taliban) in order to disincentivize their appeal during the two-player game, which thus creates a more realistic setting for the game.

Furthermore, we defined how the players locate strategies during the game and introduced two heuristic methods, the ‘Selective Search’ Heuristic and the ‘Sequential Search’ Greedy Algorithm which are used to significantly reduce the run time needed to find solutions for complex problems where the use of exhaustive enumeration becomes too cumbersome. In the next chapter, we explore the characteristics of the strategies chosen by the players and discuss experimental results designed to better understand the properties of the threshold model under the two-player game.

[This Page Intentionally Left Blank]

## 4 Experiments and Analysis

This chapter discusses the results of simulation experiments aimed at understanding the dynamics of the two-player game. The first set of experiments in Section 4.1 details the general characteristics of the strategies chosen by the players to maximize their payoffs in the transient setting. Next, we present experiments which analyze how different payoff functions influence the strategies that are chosen by the players followed by sensitivity analyses of additional model parameters. In Section 4.2 we examine the run time performance and compare the strategies and payoffs found by exhaustive enumeration, ‘Selective Search’, and ‘Sequential Search’. Lastly, in Section 4.3 we introduce population-focused actions—stimulus projects and assassinations—which are available for use during the two-player game by the US agent and TB agent, respectively, in order further aid in spreading their influence throughout the network.

### Experimental Scope

We remind the reader that all observations and conclusions drawn from the following experiments are based off of extensive empirical evidence obtained through the use of Monte Carlo simulations. While lacking formal theoretical proof, we do find some very interesting results from the simulations concerning the characteristics of the strategies that are chosen by the players, and we provide the intuitive reasons which help explain why we observe such results.

### Experimental Parameters

In Chapter 3 we described the various parameters for the threshold model, which are used for the two-player game. During the experiments presented in this chapter, unless otherwise stated, we used the following default experimental parameters found in Table 4.1 below. Also, recall the following set notation defined in Chapter 3:

$V_R = \text{set of all regular agents}$

$V_F = \text{set of all forceful agents}$

$V_{F+} = \text{set of all forceful + agents}$

$V_S = \text{set of all stubborn (immutable) agents} = V_{US} \cup V_{TB}$

Threshold ( $\tau_i$ ) Values:	$\tau_i = \begin{cases} 0.5 & \forall i \in V_R \\ 0.75 & \forall i \in V_F \cup V_{F+} \\ 1 & \forall i \in V_S \end{cases}$
Interaction-type Probabilities ( $\alpha, \beta, \gamma$ ):	$\alpha_{ij} = \begin{matrix} & & & j \in \\ & & & V_R & V_F & V_{F+} & V_S \\ i \in V_R & [0 & 0 & 0 & 0] \\ i \in V_F & [1 & 0 & 0 & 0] \\ i \in V_{F+} & [1 & 0.4 & 0 & 0] \\ i \in V_S & [1 & 1 & 1 & 0] \end{matrix}$
	$\beta_{ij} = \begin{matrix} & & & j \in \\ & & & V_R & V_F & V_{F+} & V_S \\ i \in V_R & [1 & 0 & 0 & 0] \\ i \in V_F & [0 & 0.1 & 0.1 & 0] \\ i \in V_{F+} & [0 & 0.1 & 0.1 & 0] \\ i \in V_S & [0 & 0 & 0 & 0] \end{matrix}$
	$\gamma_{ij} = \begin{matrix} & & & j \in \\ & & & V_R & V_F & V_{F+} & V_S \\ i \in V_R & [0 & 1 & 1 & 1] \\ i \in V_F & [0 & 0.9 & 0.9 & 1] \\ i \in V_{F+} & [0 & 0.5 & 0.9 & 1] \\ i \in V_S & [0 & 0 & 0 & 1] \end{matrix}$
Penalty for Identical Strategy ( $\lambda$ ):	$\lambda = 0$ (maximum penalty imposed)
Standard Deviation of the Number of Interactions per Step ( $\sigma$ ):	$\sigma = 0$

*Table 4.1 – Default Experimental Parameters*

## Experimental Setup

All experiments were performed on a Dell Studio 1558 with an Intel® Core™ i3-350M Dual-Core processor running at 2.27 GHz. The laptop has 4 GB of RAM, and runs 64-bit Windows 7 Home Premium. All of our code is written and run in 64-bit MATLAB version 7.11.0.584.

## 4.1 Results

### 4.1.1 Characteristics of Strategies

The purpose of this set of experiments is to (1) show the general characteristics of the strategies chosen by the stubborn agents in the two-player game and (2) provide the intuitive reasoning why such strategies are targeted by the stubborn agents. For these experiments, we exhaustively enumerate the payoffs for every strategy to ensure the best strategies are chosen only after examining the payoffs from the entire strategy space and not a subset which would decrease our confidence in the quality of the chosen strategies. Although we limit ourselves to one connection per stubborn agent ( $z = 1$ ) for the examples in this experiment set, since the computational time for exhaustive enumeration can become very slow for multiple connections per stubborn agent, we do compare the results of multiple connection scenarios between the two heuristic methods and exhaustive enumeration in Section 4.2.


Table 4.2 below shows the experimental design parameters used for all experiments in this experiment set in which we conduct the two-player game in the transient setting on the threshold model.

<b>Experiment Set Description</b>	
Network	Large (73 agents)
# Realizations	30
# Steps	10
Interactions per Step	300
Connections per Player ( $z$ )	1
Initial, Fixed Strategy	TB [70]


*Table 4.2 – Characteristics of Strategies: Experiment Set Description*

We tested the large, 73 agent Pashtun network created through the use of Hung’s network generator tool with 9 different initial starting beliefs for the mutable agents in the network. The network diagrams for these 9 cases are seen below in Figure 4.1. For further detail on the initial beliefs of all agents in these networks refer to Appendix B. Throughout the remainder of this chapter we will refer to these network configurations in various experiments by the titles shown in Figure 4.1 to identify the initial beliefs of the agents in the networks.


### 1. All Neutral




### 2. Village Mix 0




### 3. Village Mix 1




### 4. Village Mix 2




### 5. Village Mix 3




### 6. Village Mix 4




### 7. Village Mix 5



### 8. Village Mix 6



## 9. All Random



*Figure 4.1 – Experimental Network Diagrams*

As discussed in Chapter 3, the strategies targeted by the stubborn agents in the models used by both Hung and Howard were independent of the beliefs of the mutable agents in the network. The proposed threshold model and two-player game incorporated assumptions aimed at formulating a more realistic model, such as:

1. The long-term beliefs of the mutable agents in the network should be *dependent* on the initial beliefs of those mutable agents, despite the presence of immutable stubborn agents.
2. We shifted the previous focus from analyzing the long-term asymptotic equilibrium behavior of the network to instead analyzing the transient behavior of the network during the course of the two-player game.
3. Introduce a dynamic, two-player game whereby both players continually evaluate, change, and reevaluate their strategies throughout the course of the game, while not being content on settling on identical Nash Equilibria strategies for the long-run.
4. Peer pressure should affect the outcome of belief exchanges during pairwise interactions.
5. Lastly, we implemented the ability to penalize the payoffs for identical strategies among opposing stubborn agents (U.S. and Taliban) in order to disincentivize their appeal during the two-player game, which thus creates a more realistic setting for the game.

We first compare the strategies chosen by Howard’s stubborn agents during the two-player game on his network model with the strategies chosen by the stubborn agents on the proposed threshold model. The results are presented for all 9 test cases. Table 4.3 below shows the strategies and ‘mean belief’ payoffs for the TB and US agents for the 9 network diagrams shown in Figure 4.1 using Howard’s simulated annealing heuristic to find the Nash equilibria, long-term strategies for the two-player game. Note that regardless of initial starting conditions (beliefs), both stubborn agents choose fixed, identical strategies (agent 70) to connect to during the entire game. Also, the ‘mean belief’ payoff of the network is 0 in all cases, as the long-term, expected equilibrium beliefs of the network are independent of initial beliefs, but rather determined by the beliefs of the stubborn agents, the chosen connection strategies, and the topology of the network.

		Strategy Profile and Payoff Table Long-Term ( $10^{12}$ ) Interactions		
	Starting Beliefs	TB Strategy	US Strategy	Mean Belief Payoff
Experiment 1:	All Neutral	70	70	0
Experiment 2:	Village Mix 0	70	70	0
Experiment 3:	Village Mix 1	70	70	0
Experiment 4:	Village Mix 2	70	70	0
Experiment 5:	Village Mix 3	70	70	0
Experiment 6:	Village Mix 4	70	70	0
Experiment 7:	Village Mix 5	70	70	0
Experiment 8:	Village Mix 6	70	70	0
Experiment 9:	All Random	70	70	0

*Table 4.3 – Characteristics of Strategies: Strategy Profile and Payoff Table Using Howard’s Model*

By comparison, the same 9 experiments for the dynamic, two-player game on the threshold model in the transient setting yield quite different results. The strategy profile and payoff tables for experiments 5 and 6 are shown in Table 4.4. Appendix B shows the strategy profile and payoff tables for all 9 experiments.

Experiment 5

Step #	TB Strategy	US Strategy	Absolute Payoff	Relative Payoff
1	70	31	-0.0065	0.0017
2	48	31	-0.0025	0.0058
3	48	39	0.0074	0.0156
4	49	39	0.0068	0.0150
5	49	40	0.0155	0.0237
6	39	40	0.0151	0.0233
7	39	23	0.0214	0.0296
8	12	23	0.0199	0.0281
9	12	39	0.0269	0.0351
10	13	39	0.0221	0.0304

Experiment 6

Step #	TB Strategy	US Strategy	Absolute Payoff	Relative Payoff
1	70	23	-0.0135	0.0015
2	48	23	-0.0174	-0.0023
3	48	21	-0.0041	0.0110
4	49	21	-0.0113	0.0038
5	49	48	-0.0062	0.0089
6	13	48	-0.0059	0.0091
7	13	49	0.0019	0.0169
8	12	49	-0.0011	0.0140
9	12	13	0.0065	0.0216
10	49	13	0.0048	0.0199

*Table 4.4 – Characteristics of Strategies: Experiments 5 and 6 Strategy Profile and Payoff Tables Using the Proposed Threshold Model*

In all 9 cases, the strategy profile and payoff tables show that not only do the initial beliefs dictate the strategies chosen during each step of the game, but the ending ‘mean belief’ payoff of the network is dependent upon the initial starting beliefs as well. Although the experiments conducted on the proposed threshold model focus on the transient behavior of the network, the asymptotic equilibrium state of the network is also dependent on the initial beliefs of the agents in the network. Not only do the beliefs of the mutable agents in the threshold network dictate the equilibrium state of the network due to peer pressure influencing changes in beliefs, but we find that the beliefs of the mutable agents also influence the strategy decisions of the stubborn agents.

## KEY OBSERVATIONS

We present four key observations below concerning the strategies that are chosen to improve a stubborn agent’s payoff during the dynamic, two-player game for the threshold model in the transient setting. As previously discussed in Chapter 3, we use Monte Carlo simulations to observe the propagation of beliefs in the social network during the two-player game. The observations presented below are thus based on empirical evidence obtained through Monte Carlo simulation experiments and are not obtained through the use of theoretical proofs. Furthermore, these findings are consistently observed among various different network

scenarios; however, for simplicity we limit ourselves to the same network example when presenting these observations.

**OBSERVATION #1:** Stubborn agents should target influential agents of opposite belief.


We illustrate the point made in observation #1 by viewing the strategy profile for experiment 5 found in Table 4.3. During step 1 of the game the U.S. agent chooses agent 31, who is a forceful agent with an initial belief of -0.3. There are no other influential agents of lesser initial belief. Thus, the U.S. strategy makes sense intuitively because a stubborn agent has more potential to improve his payoff by selecting agents who hold opinions opposite their own in an attempt to persuade them to change their belief toward their own side. A stubborn agent stands little to gain by selecting agents as strategies who are already in favor of their opinion. Although we mention the characteristic for stubborn agents to choose strategies that are of opposite belief, we find that choosing strategies is far more complex than simply finding the most influential leader of opposite belief from the stubborn agent.

**OBSERVATION #2:** Stubborn agents should target influential agents who have many neighbors (large  $N_i$ 's) of lesser influence level, and whose neighbors live in small neighborhoods (small  $N_j$ 's).

Influential agents who have the qualities discussed in observation #2, who have opposite belief from the stubborn agent (observation #1), and who additionally have neighbors who are also of opposite belief, tend to be the most beneficial strategies for stubborn agents. The intuitive reason why influential leaders whose neighbors have small neighborhoods (where an agent's neighborhood is defined as the number of adjacency connections they have) are chosen strategies is that small neighborhoods means the neighbors of the targeted agent are as isolated as possible from other agents in the network which helps to prevent outside influence from affecting them. Thus, this maximizes the probability that the stubborn agent's influence will diffuse throughout the entire neighborhood or village of the targeted agent unhindered by outside influence. Empirical evidence also suggests that stubborn agents target agents whose neighborhoods are also dominated by opinions which are in opposition to the stubborn agent. Studies have shown that people who are faced with opinions contrary to their own belief will often shift their opinions in the "direction of the views of the majorities or the experts" [23].

Subsequently, the effect that group pressure can have on individuals to conform to the belief held by the majority of their neighbors may hint at the notion that people living in the same village or vicinity often have the same or similar beliefs.

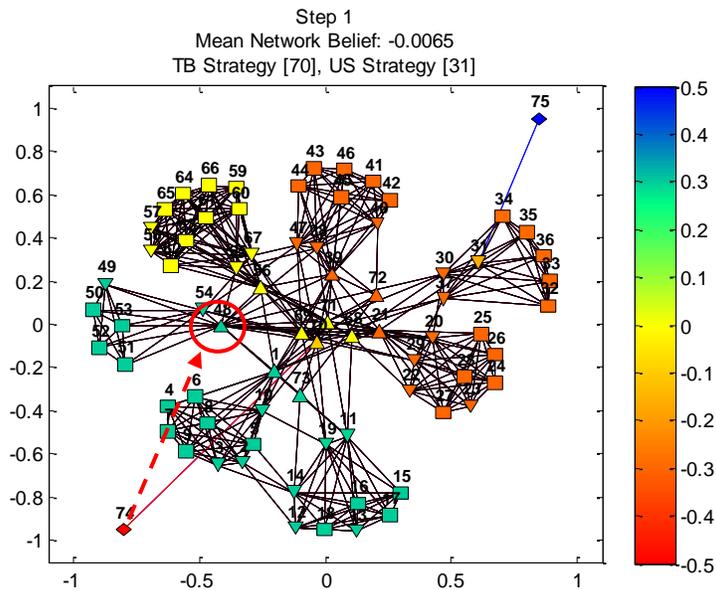
A convincing example of observation #2 is found by looking at experiment 5 (network shown below). The Taliban agent initially has a starting connection to agent 70 at step 0, and the U.S. agent is given the first opportunity to select its best strategy given this information. Thus, agent 31 is selected by the U.S. agent as its first strategy (see Figure 4.2). There are a couple reasons why agent 31 was targeted by the U.S. agent. First, agent 31 and his neighbors all have strongly negative (anti-U.S.) beliefs (-0.3). Second, out of the three villages which have strongly negative beliefs, the village containing agent 31 has the smallest neighborhoods—which means less influence from other influential agents who are neighbors to agent 31’s neighbors will be able to affect the beliefs of the neighbors of agent 31.



*Figure 4.2 – Experiment 5: First Chosen U.S. Strategy*

Next, upon the conclusion of step 1, where the US agent selects agent 31, the TB agent updates his strategy (previously agent 70) at step 2 given the US agent is now connected to agent 31. Using the same logic, the TB agent determines that agent 48 is the best strategy (see Figure 4.3 below) due to—(1) agent 48 is the most influential leader with direct access to regular agents in that village, (2) all the agents in the village are strongly pro-US (anti-TB) with beliefs +0.3,

and (3) agent 48 lives in an isolated village—the village neighbors of agent 48 have small neighborhoods, and thus, outside influence from more power agents is minimal in the transient state compared to the other pro-US villages. We next discuss a consistent observation concerning strategies involving villages which contain more than one influential leader.




*Figure 4.3 – Experiment 5: First Chosen Taliban Strategy*


**OBSERVATION #3:** The targeting strategy for villages containing more than one influential leader of similar belief is for stubborn agents to sequentially target some, or all, of those village leaders.

The intuitive reason behind observation #3 is (1) the targeted village most likely has agents of homogenous beliefs which are in opposition to the stubborn agent, which means that (2) the stubborn agent must target multiple village leaders in order to gradually influence the beliefs of all members in the village. Simply targeting one influential leader within a village containing multiple leaders is typically not observed because the targeted agent will quickly succumb to the peer pressure of the other village leaders upon the withdrawal of the stubborn agent’s influence on the targeted agent. Thus, due to the peer pressure effect of the threshold model, which affects how easily agents are persuaded to change their beliefs, a stubborn agent must sequentially target leaders within the same village in order to gradually change the beliefs of the entire village.

We demonstrate the principle of observation #3 by continuing with the same example (experiment 5). We previously saw that the TB agent targeted agent 48 as his first strategy. Upon his next step, the TB agent selects agent 49, which is a forceful leader located in the same village as agent 48. We see the dramatic effect this sequential targeting technique can have on the beliefs in a village when we compare Figures 4.3 and 4.4, and later, 4.5 and 4.6. Once a stubborn agent is able to influence the beliefs of the influential leaders in a village, it becomes easier for the rest of the agents of lesser influence level (regulars) in the village to be influenced since the influential agents in a village have a higher probability of spreading their beliefs to the agents of lower level of influence in that village.




*Figure 4.4 – Experiment 5: Before Second Chosen Taliban Strategy*



*Figure 4.5 – Experiment 5: After Second Chosen Taliban Strategy*

The resulting network below (Figure 4.6) shows the beliefs of the agents in the network at the conclusion of the 10-step, 3000 pairwise interaction game. Note that the presence of peer pressure in the threshold model impacts the diffusion of beliefs in the network, such that even after 3000 interactions the initial beliefs of the agents are still recognizable—i.e. the initial beliefs of the agents in the network affect the end state beliefs. The presence of the threshold model, and the ability it has to affect belief propagation through peer pressure influence, has significantly changed the behavior of the stubborn agents in choosing strategies during the two-player game when viewed in the transient setting compared to our predecessors’ models.



*Figure 4.6 – Experiment 5: End Game Network Beliefs*

During the previous analysis above, we used the default threshold values shown below:


$$\begin{cases} \tau_i = 0.5 & \forall i \in V_R \\ \tau_i = 0.75 & \forall i \in V_F \cup V_{F+} \\ \tau_i = 1 & \forall i \in V_S \end{cases}$$

However, we also observe how sensitive the chosen strategies in the two-player game are to the specific threshold values of the agents in the network. In general, we observe that low threshold values in the network result in a rapid diffusion of beliefs as stubborn agents wield greater influence ability, while high threshold values result in a low diffusion of beliefs as the initial beliefs of the agents will prevail. We illustrate this point by presenting the two extreme cases for the threshold values—(1) all mutable agents have threshold values of 0, and (2) all mutable agents have threshold values of 1. In both experiments, we use the same network under the same initial beliefs for the agents in the network. This network is shown as experiment 5 in Figure 4.1.

In the first experiment, all mutable agents have a threshold value equal to 0:

$$\tau_i = \begin{cases} 0 & \forall i \in V_M \\ 1 & \forall i \in V_S \end{cases}$$

Figure 4.7 below shows the ending network beliefs using these threshold values for a 10-step game with 300 interactions per step, and where each stubborn agent has one connection ( $z = 1$ ). With the threshold values for all mutable agents set to 0, we observe that the stubborn agents are able to more quickly propagate their beliefs through the network, which enables huge changes in agent beliefs from their initial values. The reason why the stubborn agents are able to quickly diffuse their beliefs to the mutable agents in the network when all mutable threshold values are 0 is that the peer pressure aspect of the threshold model has essentially been ‘turned off’ in this case. Thus, the most influential agents in the network—the two stubborn agents—are able to exert the most influence in the network as mutable agents are no longer apprehensive to change their beliefs in opposition to their neighbors’ beliefs.




*Figure 4.7 – Ending Network Diagram: Mutable Thresholds = 0 Example*

Next, we compare the results of this first experiment where all mutable threshold values are zero with the following experiment on the same network whereby all mutable agents have the other extreme threshold value of 1:

$$\tau_i = 1 \quad \forall i \in V$$

Figure 4.8 is the network diagram which shows the beliefs of the agents at the end of the 10-step game with 300 interactions per step when all threshold values are 1. In this extreme

case, no change in beliefs from the initial beliefs occurs because each mutable agent has effectively become stubborn in their own beliefs and is unwilling to adopt a new belief.



*Figure 4.8 – Ending Network Diagram: Mutable Thresholds = 1 Example*

Due to the sensitive nature of the individual threshold values, which can have a tremendous impact on the strategies chosen by the stubborn agents, we present the final key observation for strategies chosen by the players to maximize their payoffs.

**OBSERVATION #4:** Stubborn agents should target influential agents with low threshold values who also have neighbors with low threshold values.

We have observed that not only do stubborn agents consider the threshold value of each individual when choosing strategies, but also the threshold values of the neighbors of each individual. In other words, the appeal of an agent as a targeted strategy depends not only on the agent’s individual threshold value, but also the threshold values of those agents in his neighborhood. We present two examples below which illustrate this observation. Both examples use the experiment 5 network shown in Figure 4.1; however, we slightly modify the threshold values of the agents to demonstrate how sensitive the threshold values are in determining targeted strategies.

In the first experiment, we assume the following threshold values:

$$\tau_i = \begin{cases} 0.5 & \forall i \in V_R \\ 0.75 & \forall i \in V_F \cup V_{F+\setminus 1} \\ 0.5 & i \in 1 \\ 1 & \forall i \in V_S \end{cases}$$

Thus, the only difference from the default threshold values is that agent 1, a forceful+ leader, has a threshold value that is reduced from 0.75 to 0.5. The purpose of this experiment is to show that players target agents with low threshold values. Table 4.5 shows the strategy profile and payoff table for this experiment. When we compare the chosen strategies between this table and Table 4.3 (the default baseline case), we note that agent 1 is chosen three different times (twice by the TB agent and once by the US agent) in Table 4.5 when his threshold value has been reduced to 0.5, while agent 1 is not previously chosen as a strategy in the baseline scenario in Table 4.3. Therefore, we have shown that an influential agent with a low threshold value, relative to the other influential agents in the network, strengthens his position as a candidate to be chosen as a strategy by the players.

Step #	TB Strategy	US Strategy	Absolute Payoff	Relative Payoff
1	70	31	-0.0065	0.0017
2	1	31	-0.0046	0.0036
3	1	39	0.0027	0.0109
4	48	39	0.0015	0.0097
5	48	1	0.0140	0.0222
6	49	1	0.0098	0.0180
7	49	23	0.0142	0.0225
8	1	23	0.0069	0.0151
9	1	21	0.0098	0.0181
10	3	21	0.0045	0.0127

*Table 4.5 – Observation #4: Experiment 1 Strategy Profile and Payoff Table*

Although we just saw that stubborn agents target agents with low threshold values, the purpose of the second experiment (shown below) is to show that the not only does an individual's threshold value influence strategy decisions, but also the threshold values of his neighbors. In this experiment, we maintain the same default threshold values with the exception

that regular agents 50, 51, 52, and 53, increase their threshold values from 0.5 to 0.95. Thus, we assume the following threshold values for this experiment:

$$\tau_i = \begin{cases} 0.5 & \forall i \in V_R \setminus \{50, 51, 52, 53\} \\ 0.75 & \forall i \in V_F \cup V_{F+} \\ 0.95 & i \in \{50, 51, 52, 53\} \\ 1 & \forall i \in V_S \end{cases}$$

Table 4.6 shows the strategy profile and payoff table for this experiment. Upon comparing the strategy profile and payoff table from the baseline scenario (Table 4.3) to Table 4.6, we see how the threshold values of an agent’s neighbors also factors into determining the strength of that agent as a connection strategy. Previously, both agents 48 and 49 (two village leaders located in the same village as regular agents 50-53, were deemed as payoff-maximizing strategies by the TB agent. However, due to the drastic increase in threshold values of the regular agents in their village, which subsequently reduces the probability that a stubborn agent can forcefully influence the entire village, both agents 48 and 49 are no longer selected as strategies by the TB agent.

Step #	TB Strategy	US Strategy	Absolute Payoff	Relative Payoff
1	70	30	-0.0066	0.0016
2	68	30	-0.0031	0.0051
3	68	21	0.0117	0.0199
4	30	21	0.0205	0.0287
5	30	22	0.0320	0.0402
6	57	22	0.0411	0.0493
7	57	23	0.0518	0.0601
8	56	23	0.0540	0.0622
9	56	57	0.0613	0.0695
10	23	57	0.0617	0.0699

*Table 4.6 – Observation #4: Experiment 2 Strategy Profile and Payoff Table*

## 4.1.2 How Different Payoff Functions Influence Strategies

Next, we present a series of experiments which analyze the characteristics of strategies under different payoff functions. Because stubborn agents can have different perspectives on their success during the game depending on the specific payoff function they decide to use, the strategies that are chosen by a stubborn agent can vary in relation to the payoff function. In this section, we compare the characteristics of the strategies chosen by the players using the ‘number of agents won’ payoff in comparison to the strategies chosen by the ‘mean belief’ payoff.

### 4.1.2.1 Mean Belief

We first conduct two experiments where both players use the ‘mean belief’ payoff function. The results from these experiments will serve as the baseline scenarios from which we will compare the results to subsequent experiments involving another payoff function described earlier—the number of agents won. Table 4.7 describes the parameter conditions of the two experiments.

	<b>Experiment 1</b>	<b>Experiment 2</b>
Network	Large (73 agents)	Large (73 agents)
# Realizations	30	30
# Steps	8	8
Interactions per Step	300	300
Connections per Player (z)	1	1
Initial, Fixed Strategy	TB [70]	TB [70]
Initial Agent Beliefs	Village Mix 1	Village Mix 3

*Table 4.7 – Payoff Function Comparison: Experiment Descriptions*

Table 4.8 shows the ‘mean belief’ strategy profile and payoff tables for both experiments described in Table 4.7. As discussed earlier in Section 4.1.1 concerning the characteristics of strategies, we observe that the strategies chosen by the players in each experiment are dependent on the beliefs of the agents in the network. Thus, each stubborn agent considers the belief of each mutable agent when determining which strategy will yield the highest mean belief payoff during the current step in the game.

Experiment 1

Step #	TB Strategy	US Strategy	Absolute Payoff	Relative Payoff
1	70	21	-0.0123	0.0014
2	48	21	-0.0155	-0.0018
3	48	39	-0.0057	0.0080
4	49	39	-0.0081	0.0056
5	49	40	0.0001	0.0138
6	12	40	-0.0010	0.0127
7	12	23	0.0051	0.0188
8	13	23	0.0008	0.0145

Experiment 2

Step #	TB Strategy	US Strategy	Absolute Payoff	Relative Payoff
1	70	31	-0.0065	0.0017
2	48	31	-0.0025	0.0058
3	48	39	0.0074	0.0156
4	49	39	0.0068	0.0150
5	49	40	0.0155	0.0237
6	39	40	0.0151	0.0233
7	39	23	0.0214	0.0296
8	12	23	0.0199	0.0281

*Table 4.8 – Payoff Function Comparison: Mean Belief Strategy Profile and Payoff Tables*

#### 4.1.2.2 Number of Agents Won

Next, we determine the strategies found during the two-player game for the same two experiments described in Table 4.7 with the exception that both stubborn agents now use the ‘number of agents won’ payoff function instead of the ‘mean belief’ payoff function. By conducting the same two experiments with a different payoff function, we gain a better understanding of how different payoff functions can alter the strategies that the players pursue during the game. In order to use the ‘number of agents won’ payoff function, we must determine the buffer threshold value ( $B$ ), where  $B \in [0, 0.5)$ , which will be used in determining the points awarded to each player based on each mutable agent’s belief. For these experiments we set  $B = 0.1$ , which yields the following point system for the two players:

US Points Awarded	Mutable Agent Belief	TB Points Awarded
+1	$> 0.1$	-1
0	$[-0.1, 0.1]$	0
-1	$< -0.1$	+1

*Table 4.9 – Payoff Function Comparison: Number of Agents Won Point System*

The following table shows the ‘number of agents won’ strategy profile and payoff tables for both experiments. Upon comparing the strategy profile and payoff tables for two different payoff functions for each experiment, we notice drastic differences in the strategies that are chosen by the two players. Because points are only awarded to (or subtracted from) a player’s payoff as mutable agents change in belief by crossing above or below  $\pm B$ , we find that the

incentive for players to target strategies which are near  $\pm B$  in belief is very high when the payoff is number of agents won. On the other hand, when the ‘mean belief’ payoff function is used by the players, the incentive to target mutable agents of opposite belief from the player is very high as the player’s attempt to influence the mean network belief toward their own belief.

Experiment 1					Experiment 2				
Step #	TB Strategy	US Strategy	Absolute Payoff	Relative Payoff	Step #	TB Strategy	US Strategy	Absolute Payoff	Relative Payoff
1	70	31	5	9	1	70	1	-3	-1
2	1	31	-2	2	2	48	1	-4	-2
3	1	3	-1	3	3	48	20	-3	-1
4	12	3	-5	-1	4	1	20	-3	-1
5	12	1	2	6	5	1	49	-7	-5
6	3	1	-4	0	6	2	49	-10	-8
7	3	12	-5	-1	7	2	48	-6	-4
8	14	12	-8	-4	8	49	48	-15	-13


*Table 4.10 – Payoff Function Comparison: Number of Agents Won Strategy Profile and Payoff Tables*

When viewing the strategy profile and payoff tables in Table 4.10, we find that the stubborn agents select mutable agents as strategies who have beliefs close to the buffer threshold value ( $\pm B$ ). For instance, in experiment 1, agent 31 has an initial belief of -0.1, while agents 1 and 3 have initial beliefs of 0.1. Thus, we find that these influential agents are heavily targeted by both players since small deviations in the targeted agents’ beliefs result in noticeable payoff changes when using the number of agents won payoff function.

#### 4.1.2.3 Mean Belief vs. Number of Agents Won


Now that we have explored two different payoff functions and how a stubborn agent’s strategy can differ depending on the specific payoff function that is used, we draw some conclusions about which payoff function a player should consider using when viewing their own payoff during the game. Figure 4.9 shows the absolute payoffs over time with respect to the U.S. agent during experiments 1 and 2 previously described in Table 4.7. The plot on the left shows the ‘mean belief’ and ‘number of agents won’ payoffs for experiment 1, while the plot on the right shows both payoffs for experiment 2. Upon visual comparison of the payoff plots and subsequent empirical evidence, we notice that the ‘number of agents won’ payoff function can

lead to high variations in a player’s payoff over time, while the ‘mean belief’ payoff function yields payoffs for a player which are less subject to high deviations.



*Figure 4.9 – Payoff Function Comparison: Mean Belief and Number of Agents Won Payoff Plots*

A great example of the potential for high variations in a player’s payoff over time when using the ‘number of agents won’ payoff function is seen in Figure 4.10 below which shows the absolute payoff plots for an identical experiment as experiment 2 with the exception that the game is carried out for 20 steps rather than 8 steps.



*Figure 4.10 – Payoff Function Comparison: Payoff Plots for 20-Step Experiment 2*

Thus, as we observe from the payoff plots for the 20-step game for experiment 2, the number of agents (nodes) won payoff is highly variable for the players. After step 8, the payoff for the U.S. agent is -15 (which means that the Taliban agent has ‘won’ 15 more agents in the network to supporting the Taliban than the U.S. agent has supporting the United States cause). However, after only 5 more steps, the payoff has changed by 22 points such that the U.S. agent now has the advantage in the game with a payoff of +7. Therefore, if a stubborn agent is risk averse, and therefore wants to minimize the possibility of huge changes in payoff over time, the ‘mean belief’ payoff should be used since the ‘number of agents won’ payoff is more subject to higher variations. Lastly, we note from experimental evidence that the degree of variability in the ‘number of agents won’ payoff is largely determined by the buffer threshold value (B) that is determined, and specifically, how many agents in the network have beliefs which are near the buffer threshold value—i.e. the more agents which are close to the buffer allows for bigger changes in payoff over time as only slight changes in those agents’ beliefs equate to +/-1 changes in ‘number of agents won’ payoff for a player.

### **4.1.3 Sensitivity to Model Parameters**

We previously discussed that threshold values are extremely sensitive parameters in determining strategies for the two-player, dynamic game. However, we present a few additional model parameters which can also influence the strategies chosen by the players during the game and the rate of belief diffusion in the network.

#### **4.1.3.1 Alpha, Beta, and Gamma Values ( $\alpha$ , $\beta$ , and $\gamma$ )**

##### **Different Types of Societies**

The default  $\alpha$ ,  $\beta$ , and  $\gamma$  values which were developed by our predecessor’s modeling approach [8] in conjunction with the MIT Political Science Department are primarily used in our experiments since they were created to accurately represent Pashtun society in Afghanistan. However, our modeling approach can easily be extended to analyze opinion dynamics in other societal types, such as Western (democratic), consensus, and hierarchical societies, by simply changing the interaction-type probabilities ( $\alpha$ ,  $\beta$ , and  $\gamma$  values) for the agents in the network. Table 4.11 below shows different alpha, beta, and gamma values that are assigned depending on

the type of society that is modeled by the network. Our motivation is to determine how the propagation of beliefs in the threshold model is affected by different society types, and more specifically, different  $\alpha$ ,  $\beta$ , and  $\gamma$  values.

<b>Society Type</b>	<b>Description</b>	<b>alpha matrix</b>				<b>beta matrix</b>			
Default	normal values for Pashtun society	0	0	0	0	1	0	0	0
		1	0	0	0	0	0.1	0.1	0
		1	0.4	0	0	0	0.1	0.1	0
		1	1	1	0	0	0	0	0
Consensus	everyone averages, except stubborn agents	0	0	0	0	1	1	1	0
		0	0	0	0	1	1	1	0
		0	0	0	0	1	1	1	0
		1	1	1	0	0	0	0	0
Hierarchical	forcefully influence all lower nodes, avg. with same level	0	0	0	0	1	0	0	0
		1	0	0	0	0	1	0	0
		1	1	0	0	0	0	1	0
		1	1	1	0	0	0	0	0
Western	realistic democratic society	0.1	0	0	0	0.9	0.1	0.1	0
		0.5	0.1	0	0	0.1	0.5	0.1	0
		0.9	0.5	0.1	0	0.1	0.1	0.5	0
		1	1	1	0	0	0	0	0

Note: Gamma matrix values are found by the following equation:  $\gamma_{ij} = 1 - \alpha_{ij} - \beta_{ij}$

*Table 4.11 – Alpha, Beta, Gamma Values for Different Society Types*


### Society Type Experiments

We conducted experiments with identical conditions with the exception of the  $\alpha$ ,  $\beta$ , and  $\gamma$  values that represent the four society types in the table above. By doing so, we can better gauge the effect that different  $\alpha$ ,  $\beta$ , and  $\gamma$  values have on the evolution of beliefs in the network over time. A description of the experiment set is found in the table below.

<b>Experiment Set Description</b>	
Network	Toy (16 agents)
# Realizations	30
# Steps	50
Interactions per Step	100
Connections per Player (z)	1
Initial, Fixed Strategy	TB [1]
Initial Agent Beliefs	All -0.5
Society Type	1. Default
	2. Consensus
	3. Hierarchical
	4. Western

*Table 4.12 – Different Society Types ( $\alpha$ ,  $\beta$ , and  $\gamma$  values): Experiment Set Description*

The plot below (Figure 4.11) shows the ‘mean belief’ payoff with respect to the U.S. agent during the two-player game conducted for the four types of societies described in Table 4.12. A significant difference exists between the payoffs of the U.S. agent for the consensus society compared to the payoffs of the other three types of societies, all three of which yielded similar payoffs as the network beliefs evolved over time. The main difference between the consensus society and the other three societies is that the probability of averaging ( $\beta - type$ ) interactions occurring in the consensus society between two agents is much higher than the other societies, whereas the probability of forceful ( $\alpha - type$ ) interactions occurring in the other three societies is more prominent than in the consensus society. Moreover, upon viewing the payoff plots over time in Figure 4.11 for the default, hierarchical, and Western societies, we recognize that the payoff results are highly robust to considerable variations in the  $\alpha$ ,  $\beta$ , and  $\gamma$  values. Meanwhile, the consensus society is considered radically different in its interaction-type characteristics compared to the other three societies due to the absence of forceful ( $\alpha - type$ ) interactions occurring between interacting mutable agents. Thus, the robustness of the  $\alpha$ ,  $\beta$ , and  $\gamma$  values gives us a higher degree of confidence in our results, even if there are inaccuracies in these values.



*Figure 4.11 – Different Society Types ( $\alpha$ ,  $\beta$ , and  $\gamma$  values): Mean Belief Payoff Plot*

As we have seen by this example, and subsequent empirical evidence, the evolution of beliefs in the networks depends upon the interaction-type probabilities, which are specified by the  $\alpha$ ,  $\beta$ , and  $\gamma$  values. The table below summarizes the general effect that these values have on the propagation of beliefs in the networks under the threshold model. Thus, we find that the diffusion of beliefs in different societies can be very different depending on the probabilities interaction-type probabilities.

	<b>Effect on Evolution of Beliefs</b>
Increased Probability of $\alpha$ – type interactions	Increased deviations in network beliefs away from initial network beliefs will result as the probability of forceful interactions increases.
Increased Probability of $\beta$ – type interactions	Initial starting beliefs will more likely prevail with increased probability of averaging interactions, which results in small deviations away from initial network beliefs.
Increased Probability of $\gamma$ – type interactions	Initial starting beliefs will more likely prevail with increased probability of identity interactions as agents will exhibit no change in belief when identity interactions occur, which thus results in small deviations away from initial network beliefs.

*Table 4.13 – Different Society Types ( $\alpha$ ,  $\beta$ , and  $\gamma$  values): Effect of Evolution of Beliefs*

#### 4.1.3.2 Expected Number of Interactions per Step ( $\mu$ )

During the two-player game, a player locates a strategy that is expected to yield the best payoff after  $\mu$  pairwise interactions occur in the network, after which the opposing player will reevaluate their current strategy and subsequently find their next strategy. We often fix the number of pairwise interactions that occur between steps in the experiments in order to create an equal and more symmetric game for the two opposing stubborn agents. However, we pose the following question concerning the number of interactions per step: “If given the opportunity, would stubborn agents prefer more opportunities to reevaluate their current strategy and choose a new strategy (i.e. more steps) but with fewer interactions per step, or would they rather prefer fewer steps but with more interactions per step, where both scenarios result in the same total number of pairwise interactions occurring during the game?”

The following table details the conditions of the experiments that were conducted to analyze this question:

<b>Experiment Set Description</b>	
Network	Large (73 agents)
# Realizations	30
# Steps	1. 10 2. 30
Interactions per Step	1. 300 2. 100
Connections per Player (z)	2
Initial, Fixed Strategy	TB [69, 70]
Initial Agent Beliefs	All Random
Payoff Function for TB	A. Mean Belief B. Random Selection (excluding regulars) C. Random Selection (including regulars)
Payoff Function for US	A. Mean Belief B. Mean Belief C. Mean Belief

*Table 4.14 –  $\mu$  Sensitivity: Experiment Set Description*

In total, we performed six experiments: two different scenarios (10 steps with 300 interactions per step vs. 30 steps with 100 interactions per step) each on three different payoff

functions for the TB agent—(1) mean belief, (2) choosing a random strategy excluding all regular agents, and (3) choosing a random strategy including regular agents. The US agent maintained the mean belief payoff function during all experiments. Table 4.15 shows the mean belief payoff with respect to the US agent at the end of the game for each experiment.

	Payoff: A	Payoff: B	Payoff: C
<b>Exp 1 (10 steps, 300 interactions per step)</b>	0.0091	0.0461	0.0574
<b>Exp 2 (30 steps, 100 interactions per step)</b>	0.0100	0.0529	0.0826

*Table 4.15 –  $\mu$  Sensitivity: Experiment Set Payoff Table*

Thus, based on the experimental evidence, we conclude that stubborn agents would prefer more opportunities (steps) to update their strategies, but with fewer interactions per step since the potential for higher payoffs can be obtained when a stubborn agent is given more opportunities to reevaluate his current strategy during the two-player game. In a realistic environment, whereby opposing players do not necessarily know the strategies of one another with certainty, the players would benefit most by reevaluating their strategies as soon as possible upon receiving intelligence about the opposing stubborn agent’s current strategy. Moreover, upon analyzing the strategy profile and payoff tables from the experiments, we conclude that the strategies chosen by the players to maximize their payoffs are sensitive to the specific  $\mu$  parameter (see Table B.3 in Appendix B). Table B.3 shows that the strategies which maximize a player’s payoff are subject to change depending on how many interactions occur per step in the network.

#### **4.1.3.3 Initial (Step 0) Strategy for One Stubborn Agent**

Another parameter which can subsequently impact the strategies that are selected by the stubborn agents is the strategy that is fixed for one stubborn agent during step 0, which is used to initiate the start of the two-player game. The strategy that is initially selected by a stubborn agent can alter the strategy profiles that are chosen by the stubborn agents during the subsequent steps of the game. We present the following example detailed in Table 4.16 below whereby we examine the strategy profile and payoff tables for two experiments with the same initial conditions, with the exception that the initial, fixed strategy for the TB agent during step 0 is

different. The motivation behind this experiment is to show that the strategy chosen by the stubborn agent during step 0 as the initial strategy can subsequently impact the strategies and payoffs during the rest of the game.

<b>Experiment Set Description</b>	
Network	Large (73 agents)
# Realizations	30
# Steps	10
Interactions per Step	300
Connections per Player (z)	2
Initial, Fixed Strategy	1. TB [1] 2. TB [70]
Initial Agent Beliefs	Village Mix 0

*Table 4.16 – Initial (Step 0) Strategy for Stubborn Agent: Experiment Set Description*

As the following strategy profile and payoff tables for the two experiments reveal (see Table 4.17 below), the initial strategy chosen for the TB agent during step 0 (agent 1 during experiment 1 and agent 70 during experiment 2) subsequently impacts the payoffs and successive strategies that are chosen by the players throughout the entire game. However, we do observe similarities between several of the strategies that are chosen by the players in both experiments, but the order in which they are selected as strategies can be affected by the initial (step 0) strategy that is selected for a stubborn agent.

Experiment 1				Experiment 2			
Step #	TB Strategy	US Strategy	Absolute Payoff	Step #	TB Strategy	US Strategy	Absolute Payoff
1	1	72	-0.0036	1	70	72	-0.0027
2	69	72	-0.0102	2	68	72	-0.0074
3	69	12	-0.0049	3	68	39	0.0006
4	21	12	-0.0096	4	21	39	-0.0024
5	21	13	-0.0071	5	21	12	0.0042
6	23	13	-0.0141	6	23	12	0.0001
7	23	39	-0.0130	7	23	13	0.0057
8	13	39	-0.0181	8	12	13	0.0008
9	13	48	-0.0140	9	12	49	0.0057
10	39	48	-0.0157	10	13	49	0.0029

*Table 4.17 – Initial (Step 0) Strategy for Stubborn Agent: Strategy Profile and Payoff Tables*

#### 4.1.3.4 Penalty for Identical Strategy ( $\lambda$ )

In Chapter 3 we discussed the motivation behind implementing a penalty for identical connections among stubborn agents. Now, we explore cases where the  $\lambda$  penalty factor may or may not be necessary for limiting identical connection strategies. Table 4.18 shows the experiment set description for the experiments analyzing how stubborn agents select strategies under different initial network beliefs when no penalty is inflicted on their ‘mean belief’ payoffs for identical strategies.

<b>Experiment Set Description</b>	
Network	Large (73 agents)
# Realizations	30
# Steps	10
Interactions per Step	300
Connections per Player ( $z$ )	1
Initial, Fixed Strategy	TB [70]
Penalty for Identical Strategy ( $\lambda$ )	$\lambda = 1$ (none)
Initial Agent Beliefs	<ol style="list-style-type: none"> <li>1. Village Mix 0</li> <li>2. All +0.2</li> <li>3. Village Mix 3</li> <li>4. All Random</li> <li>5. Village Mix 5</li> </ol>

*Table 4.18 – Penalty for Identical Strategy ( $\lambda$ ): Experiment Set Description*

Table 4.19 shows the strategy profile and payoff tables for experiments 1 and 2 described in the table above. We indicate strategies which are identical during a particular step with red text. Experiment 1 (Village Mix 0) contains villages of various beliefs (both pro-US and pro-TB) which are clustered homogenously by village, whereas experiment 2 (All +0.2) contains a completely homogenous network whereby all agents have initial beliefs of +0.2 (i.e. marginally pro-US). Upon inspection of the strategy profile and payoff tables for these two experiments and the other three experiments listed in the table above, we observe that the penalty for identical strategies is not always necessary for preventing identical connection strategies among stubborn agents.

Experiment 1

Step #	TB Strategy	US Strategy	Absolute Payoff
1	70	72	-0.0027
2	68	72	-0.0076
3	68	39	0.0003
4	21	39	-0.0029
5	21	49	0.0030
6	23	49	-0.0015
7	23	48	0.0042
8	49	48	0.0014
9	49	13	0.0064
10	48	13	0.0041

Experiment 2

Step #	TB Strategy	US Strategy	Absolute Payoff
1	70	70	0.1995
2	67	70	0.1977
3	67	67	0.1977
4	57	67	0.1941
5	57	57	0.1947
6	67	57	0.1924
7	67	67	0.1945
8	57	67	0.1909
9	57	57	0.1933
10	67	57	0.1915

*Table 4.19 – Penalty for Identical Strategy ( $\lambda$ ): Experiments 1 and 2 Strategy Profile and Payoff Tables*

We first observe that networks which contain initial network beliefs representing both sides of the belief spectrum (both pro-US and pro-TB agents) do not necessarily require the penalty for identical strategies as stubborn agents are motivated to target agents of highly opposite belief (when such strategies are available in the network) since these strategies offer higher payoffs. Meanwhile, for networks which contain mostly, or completely, homogenous beliefs (such as experiment 2), or networks which may have different beliefs clustered homogeneously by village, yet the majority of all agents in the network are still highly supportive of one stubborn agent (such as experiment 5), identical strategies are prone to exist without the use of the  $\lambda$  penalty. Moreover, for networks which are prone to identical strategies, the stubborn agent who is ‘winning’ the game based on the absolute payoff from the initial network beliefs has the incentive to choose the identical strategy as the opposing stubborn agent, while the ‘losing’ stubborn agent is motivated not to choose the same strategy as the ‘winning’ stubborn agent. We observe this interesting finding by viewing the strategy profile for experiment 2 in Table 4.19 and subsequently provide the rationale behind this observation.

In experiment 2, the initial mean network belief is +0.2, indicating that the U.S. agent is initially ‘winning’ the game. Upon each U.S. step during the 10-step game, the U.S. agent chooses the same strategy as the TB agent in order to minimize the potential payoff loss due to the TB agent’s influence in the network (i.e. the U.S. and TB agents exert equal influence on the

same targeted strategy for identical connection strategies). From an intuitive perspective, it makes sense that the U.S. agent is motivated to choose identical strategies as the TB agent since the U.S. agent is initially ‘winning’ the beliefs of the agents in the network, and thus, attempts to thwart the negative influence from the TB agent by choosing the same strategies. Meanwhile, during each TB step, the TB agent selects a strategy not presently occupied by the U.S. agent in order to have full (and not shared) influence over the targeted agent(s) in order to have the greatest chance of increasing his payoff.

## **4.2 Analysis of ‘Selective Search’ and ‘Sequential Search’ Heuristics**

In Chapter 3 we introduced the ‘Selective Search’ and ‘Sequential Search’ heuristics as methods used to significantly reduce the computational time (compared to exhaustive enumeration) necessary to find strategies in large, complex networks with multiple connections per stubborn agent. In Section 4.2.1 we present the run time approximations for both heuristic methods as well as exhaustive enumeration followed by example network cases whereby the use of the ‘Sequential Search’ Greedy Algorithm is necessary for finding solutions as the run times for exhaustive enumeration and ‘Selective Search’ become intractable. Later, in Section 4.2.2, we present a payoff maximization comparison between the three methods on identical scenarios. The purpose of this comparison is to analyze how the quality of the solutions obtained by each heuristic compare to the solutions found using exhaustive enumeration.

### **4.2.1 Run Time Comparison: Exhaustive Enumeration vs. ‘Selective Search’ vs. ‘Sequential Search’**

Before we analyze the quality of the solutions obtained by the ‘Sequential Search’ Greedy Algorithm, we first present the run time approximation equations for the exhaustive enumeration, ‘Selective Search’, and ‘Sequential Search’. Moreover, we provide example cases and their run time approximations to give the reader a better understanding of the computational savings afforded by the use of the two heuristic methods. The run time approximation for exhaustively enumerating all possible strategies without the use of ‘Selective Search’ is:

$$Run\ Time\ (sec) = \theta * \#\ steps * \mu * \#\ Realizations * (|V_M|)^z \quad (4.1)$$

The run time approximation for the ‘Selective Search’ Heuristic is:

$$Run\ Time\ (sec) = \theta * \#\ steps * \mu * \#\ Realizations * (|V_F \cup V_{F+}|)^z \quad (4.2)$$

Meanwhile, the run time approximation for the ‘Sequential Search’ Greedy Algorithm (after the ‘Selective Search’ Heuristic is applied to remove all combinations of regular agents from the strategy space) is:

$$Run\ Time\ (sec) = \theta * \#\ steps * \mu * \#\ Realizations * (|V_F \cup V_{F+}|) * z \quad (4.3)$$

where  $\theta$  is the network parameter associated with the specific network

The  $\theta$  network parameter is on the order of  $10^{-4}$  for the networks used in this thesis. For the small, 16 node network seen in Figure 3.7,  $\theta = 1.58 * 10^{-4}$  while for large, 73 node network seen in Figure 3.9,  $\theta = 1.39 * 10^{-4}$ . It is important to note that the computational times for both the exhaustive enumeration method (4.1) and the ‘Selective Search’ Heuristic (4.2) increase linearly with the number of steps in the game, the expected number of interactions per step, and the number of realizations of the Monte Carlo simulation, while they increase exponentially with the number of connections per stubborn agent. The difference between these two methods is the size of the strategy space that is searched:  $|V_M|$  vs.  $|V_F \cup V_{F+}|$  which is due to the removal of all strategies containing regular agents under the ‘Selective Search’ Heuristic. On the other hand, the run time for the ‘Sequential Search’ Greedy Algorithm (4.3) increases linearly with all inputs, including the number of connections per stubborn agent. Therefore, as the number of connections per stubborn agent increases, the run time for the greedy algorithm scales very well as opposed to the run times for exhaustive enumeration and the ‘Selective Search’ Heuristic.

Table 4.20 shows the run time approximations for the exhaustive enumeration method, the ‘Selective Search’ Heuristic and the ‘Sequential Search’ Greedy Algorithm using the large, 73 node network under different conditions for the number of connections per stubborn agent. For all cases we assume a 10-step game, with 300 interactions per step, and 30 realizations of the Monte Carlo simulation.

	<b>Exhaustive Enumeration</b>	<b>'Selective Search'</b>	<b>'Sequential Search'</b>
$z = 1$	15.2 min	7.3 min	7.3 min
$z = 2$	18 hr 30 min	4 hr 15 min	14.6 min
$z = 3$	56 days 6 hr	6 days 5 hr	21.9 min
$z = 10$	$1.70 * 10^{12}$ years	$1.09 * 10^9$ years	1 hr 13 min

*Table 4.20 – Run Time Approximation Comparison Table*

Table 4.20 shows the significant advantage offered by the ‘Sequential Search’ Greedy Algorithm in terms of run time scalability for scenarios where the number of connections per stubborn agent is greater than one ( $z > 1$ ). For a realistically small example of  $z = 3$  connections per stubborn agent, the run times for the exhaustive enumeration method (56 days 6 hours) and ‘Selective Search’ Heuristic (6 days 5 hours) become quite cumbersome, while the ‘Sequential Search’ run time of approximately 22 minutes is very acceptable. As the number of connections increases further, for instance  $z = 10$ , the computational times for both exhaustive enumeration and ‘Selective Search’ become intractable, while the ‘Sequential Search’ method takes just over one hour.

#### **4.2.2 Payoff Maximization: Exhaustive Enumeration vs. ‘Selective Search’ vs. ‘Sequential Search’**

Although we have seen the ability of the ‘Sequential Search’ Greedy Algorithm to drastically reduce the computational time for cases where the number of connections per stubborn agent is greater than one, there is an inherent tradeoff in the quality of the obtained solutions since the greedy algorithm significantly reduces the strategy space that is searched by each player. Thus, the question we must determine is to what extent the greedy algorithm jeopardizes the quality of the solutions. If ‘Sequential Search’ is able to obtain solutions of similar quality compared to exhaustive enumeration, then we increase our willingness to accept the tradeoff between having lower quality solutions with dramatically reduced run times instead of guaranteeing solutions of no worse quality but with exceptionally longer computational times under exhaustive enumeration or ‘Selective Search’.

We present three experiments which analyze the quality of the solutions that are found by the three methods. In all experiments, the large, 73 agent network was used under different

conditions. The first key conclusion from the solution comparisons is that the ‘Selective Search’ Heuristic found precisely the same strategies selected by the exhaustive enumeration method. This finding seems intuitively obvious given the presence of numerous influential agents in the large network, which subsequently makes connecting to regular agents unappealing. Details of the first experiment and its results are found in the tables below.

<b>Experiment 1 Description</b>	
Network	Large (73 agents)
# Realizations	10
# Steps	8
Interactions per Step	100
Connections per Player (z)	2
Initial, Fixed Strategy	TB [1, 70]
Initial Agent Beliefs	0

*Table 4.21 – Payoff Maximization: Experiment 1 Description*

Step #	TB Strategy	US Strategy	Absolute Payoff	Relative Payoff
1	[1, 70]	[1, 39]	0.0001	0.0001
2	[13, 39]	[1, 39]	-0.0004	-0.0004
3	[13, 39]	[13, 38]	0.0002	0.0002
4	[1, 38]	[13, 38]	-0.002	-0.002
5	[1, 38]	[38, 39]	-0.0007	-0.0007
6	[13, 39]	[38, 39]	-0.0026	-0.0026
7	[13, 39]	[1, 13]	-0.0014	-0.0014
8	[13, 40]	[1, 13]	-0.0021	-0.0021

*Table 4.22 – Payoff Maximization: Experiment 1 Exhaustive Enumeration and ‘Selective Search’ Results*

Step #	TB Strategy	US Strategy	Absolute Payoff	Relative Payoff
1	[1, 70]	[1, 39]	0.0001	0.0001
2	[13, 39]	[1, 39]	-0.0004	-0.0004
3	[13, 39]	[13, 38]	0.0002	0.0002
4	[1, 1 (38)]	[13, 38]	-0.0016	-0.0016
5	[1, 1 (38)]	[39 (38), 39]	-0.0004	-0.0004
6	[13, 38 (39)]	[39 (38), 39]	-0.0024	-0.0024
7	[13, 38 (39)]	[1, 13]	-0.0009	-0.0009
8	[39 (38), 39 (40)]	[1, 13]	-0.0018	-0.0018

*Table 4.23 – Payoff Maximization: Experiment 1 ‘Sequential Search’ Results*

For the results of the ‘Sequential Search’ method, we designated a strategy connection as incorrectly matching the exhaustive enumeration solution with red text and include the exhaustive enumeration solution in parenthesis next to the incorrect strategy. Thus, when viewing the results from experiment 1, we note that the first three steps for ‘Sequential Search’ identically match those strategies found by exhaustive enumeration. Also, for steps 4 thru 8, 12 out of 20 strategy connections were identical. What is important to note, though, is of the 8 strategy connections that did not match the exhaustive enumeration solutions, all but two of them were agents located within the same village as the agents in the exhaustive enumeration strategies. Agents 38, 39, and 40 are all influential agents located in the same village, and thus offer very similar payoffs to the players. Thus, although the greedy algorithm incorrectly swapped these strategies, the detriment to each player’s payoff was insignificant.

The second experiment was conducted using the same conditions as experiment 1 but with a different initial, fixed strategy for the Taliban agent. Details of experiment 2 and its results are found in the following tables below:

<b>Experiment 2 Description</b>	
Network	Large (73 agents)
# Realizations	10
# Steps	8
Interactions per Step	100
Connections per Player (z)	2
Initial, Fixed Strategy	TB [48, 49]
Initial Agent Beliefs	0

*Table 4.24 – Payoff Maximization: Experiment 2 Description*

Step #	TB Strategy	US Strategy	Absolute Payoff	Relative Payoff
1	[48, 49]	[1, 39]	0.0001	0.0001
2	[13, 48]	[1, 39]	0.0000	0.0000
3	[13, 48]	[40, 47]	0.0003	0.0003
4	[39, 39]	[40, 47]	-0.0012	-0.0012
5	[39, 39]	[13, 13]	0.0000	0.0000
6	[40, 47]	[13, 13]	-0.0029	-0.0029
7	[40, 47]	[39, 39]	-0.0025	-0.0025
8	[13, 13]	[39, 39]	-0.0056	-0.0056

*Table 4.25 – Payoff Maximization: Experiment 2 Exhaustive Enumeration and ‘Selective Search’ Results*

Step #	TB Strategy	US Strategy	Absolute Payoff	Relative Payoff
1	[48, 49]	[1, 39]	0.0001	0.0001
2	[13, 38 (48)]	[1, 39]	0.0000	0.0000
3	[13, 38 (48)]	[48 (40), 48 (47)]	-0.0002	-0.0002
4	[39, 39]	[48 (40), 48 (47)]	-0.0020	-0.0020
5	[39, 39]	[13, 13]	-0.0008	-0.0008
6	[40, 48 (47)]	[13, 13]	-0.0039	-0.0039
7	[40, 48 (47)]	[38 (39), 39]	-0.0031	-0.0031
8	[13, 13]	[38 (39), 39]	-0.0059	-0.0059

*Table 4.26 – Payoff Maximization: Experiment 2 ‘Sequential Search’ Results*

The strategy profile of experiment 2 for the greedy algorithm did not fare quite as well as experiment 1 as the first incorrect strategy was chosen during step 2. In total, 10 strategy connections did not match the exhaustive enumeration solution compared to 8 from experiment 1. If a mistake is made earlier in the game by the greedy algorithm, we expect more mistakes to occur during the game as the initial mistake will influence the subsequent decisions made by the greedy algorithm during the rest of the game. We point out that the follow-on mistakes made by the greedy algorithm in subsequent steps could be attributable to the fact that the greedy algorithm is trying to optimize the strategy for a player based on the opponent’s strategy. If the opponent’s strategy did not match the exhaustive enumeration solution then optimizing the strategy for the other player based on this incorrect strategy may only further lead to mismatches with the exhaustive enumeration strategy profile. For instance, during step 2 the TB agent was supposed to select agent 48 as a connection strategy, but instead chose agent 38. Upon the next

step, the US agent selects agent 48 rather than agents 40 and 47. Thus, although the greedy algorithm incorrectly chose agent 38 in step 2 for the TB agent instead of agent 48, we notice that during step 3, because agent 48 was not previously chosen, the greedy algorithm now recognized agent 48 as a quality strategy and the US agent subsequently chooses this strategy.

The last experiment was performed using the same network but with different initial agent beliefs and  $z = 3$  connections per stubborn agent. The description of experiment 3 and its results are found in the tables below.

<b>Experiment 3 Description</b>	
Network	Large (73 agents)
# Realizations	10
# Steps	6
Interactions per Step	100
Connections per Player (z)	3
Initial, Fixed Strategy	TB [48, 49, 54]
Initial Agent Beliefs	Village Mix 4

*Table 4.27 – Payoff Maximization: Experiment 3 Description*

Step #	TB Strategy	US Strategy	Absolute Payoff	Relative Payoff
1	[48, 49, 54]	[1, 1, 39]	-0.0150	0.0001
2	[49, 49, 73]	[1, 1, 39]	-0.0151	-0.0001
3	[49, 49, 73]	[47, 49, 54]	-0.0154	-0.0003
4	[48, 48, 48]	[47, 49, 54]	-0.0227	-0.0077
5	[48, 48, 48]	[38, 40, 47]	-0.0208	-0.0058
6	[49, 49, 54]	[38, 40, 47]	-0.0290	-0.0140

*Table 4.28 – Payoff Maximization: Experiment 3 Exhaustive Enumeration and ‘Selective Search’ Results*

Step #	TB Strategy	US Strategy	Absolute Payoff	Relative Payoff
1	[48, 49, 54]	[1, 1, 39]	-0.0150	0.0001
2	[13 (49), 40 (49), 73]	[1, 1, 39]	-0.0151	0.0000
3	[13 (49), 40 (49), 73]	[47, 69 (49), 72 (54)]	-0.0144	0.0007
4	[12 (48), 14 (48), 19 (48)]	[47, 69 (49), 72 (54)]	-0.0199	-0.0048
5	[12 (48), 14 (48), 19 (48)]	[14 (38), 40, 47]	-0.0187	-0.0036
6	[13 (49), 13 (49), 13 (54)]	[14 (38), 40, 47]	-0.0213	-0.0063

*Table 4.29 – Payoff Maximization: Experiment 3 ‘Sequential Search’ Results*

As the number of connections per stubborn agent increases, the proportion of the entire strategy space that is not searched by the greedy algorithm also increases. Therefore, in the third experiment the greedy algorithm had more errors in finding the same strategies as exhaustive enumeration, as expected, than the previous two experiments. However, there were instances where some of the strategies found by the greedy algorithm identically matched those found through exhaustive enumeration, such as the strategies selected in step 1. Even though exhaustive enumeration can guarantee solutions of no worse quality compared to the greedy algorithm, the computational time becomes impractical very quickly. Table 4.30 below shows the run time comparisons for all three experiments for each method.

<b>Run Time Comparison</b>			
	Exhaustive Enumeration	'Selective Search'	'Sequential Search'
Experiment 1	1 hr 42 min	23 min	1.3 min
Experiment 2	1 hr 42 min	23 min	1.3 min
Experiment 3	3 days 18 hr 13 min	9 hr 57 min	1.5 min

*Table 4.30 – Payoff Maximization: Run Times*

As we see, even for the relatively small number of total pairwise interactions (600) and 10 realizations which occurred during experiment 3, almost four days were needed days to find the solutions through exhaustive enumeration and nearly 10 hours using the ‘Selective Search’ Heuristic. On the other hand, the ‘Sequential Search’ Greedy Algorithm took just 90 seconds to find its solutions. Thus, even though exhaustive enumeration and ‘Selective Search’ provide no worse (and often better) solutions than the greedy algorithm , such huge reductions in run times provided by ‘Sequential Search’, especially for highly complex network scenarios, increase our willingness to accept lower quality solutions. For instance, given the time constraints in a realistic setting such as military leaders choosing which village leaders to communicate with in Afghanistan, rapidly determining decent strategies in a fraction of the time is likely to be more attractive than spending days or even months to precisely determine incrementally better strategies, especially since the landscape of the battlefield may have already changed by then. Therefore, the greedy algorithm is a good tool that should provide acceptable solutions whenever the computational time becomes too burdensome to handle.

## 4.3 Population-Focused Actions Which Affect Belief Propagation

In this section we introduce population-focused actions which the players can use in order to influence the belief propagation throughout the network. In a realistic environment, such as Afghanistan, the U.S. military and Taliban insurgents will unlikely limit themselves to strictly direct communications with the local populace in order to win over the minds of the people toward supporting their cause. Although much emphasis is placed on face-to-face communications with the local population, research on the attitudes of the local populace has shown that Afghans also “subscribe to the axiom that deeds speak louder than words” [24]. Subsequently, the U.S. military also seeks to increase popular support by aiding the Afghani people with stimulus projects. These stimulus projects may include but are not limited to—building a school, well, or road for a village and/or giving money for economic projects. On the other hand, the Taliban insurgents may give death threats or assassinate villagers who are deemed to be in direct support of the counterinsurgency effort. Because these population-focused actions by both the U.S. military and Taliban insurgents will affect the opinions of the populace, we expand the two-player game to incorporate the ability for the US and TB agents to strategically use these additional resources.

### 4.3.1 U.S. Stimulus Projects

#### Description

We implement the ability for the US agent to conduct stimulus projects during the two-player game as follows:

- The number of stimulus projects that will occur during the game is predetermined.
- A limit of one stimulus project can occur per U.S. step during the game.
- No stimulus projects may occur during a TB step.
- We designate four different types of stimulus projects (build a school, well, or road, or give money), yet we can easily allow for more types.
- Given the predetermined number of stimulus projects set to occur during the game, we either randomly or selectively choose the type of project to occur as well as the U.S. step for which it will occur.

- Given a stimulus project is to occur during a U.S. step, the U.S. agent updates his strategy per one of the specified payoff-maximization approaches discussed in section 3.3.4, but while simultaneously taking into consideration the impact the stimulus project will have on the beliefs of the agents in the network due to the stimulus project in an attempt to yield the highest payoff.
- The effect of the stimulus project:
  - The targeted agent, which is selected as the U.S. connection strategy during the step, and all of the mutable agents (nodes) that are adjacent to the agent that is targeted for the stimulus project will receive a one-time ‘belief shock’ from the stimulus project. All other mutable agents receive no belief shock from the stimulus project.
  - The belief shock follows a sinusoidal function (see below), and can either be positive or negative depending on the prior belief of the individual before the stimulus project occurs

DEFINITION:

We define the set of mutable agents who receive a belief shock from a U.S. stimulus project that occurs during step k of the two-player game as:

$$V_{k,Stimulus} \quad \text{where } Stimulus \in \{School, Well, Road, Money\}$$

The belief of agent  $i$  is updated due to the stimulus project occurring during step k and time t as follows:

$$X_i(t + 1) = \begin{cases} X_i(t) + A_{Stimulus} * \sin(Period * (X_i(t) - Phaseshift)) & \forall i \in V_{k,Stimulus} \\ X_i(t) & \text{else} \end{cases}$$


where the default amplitude ( $A_{Stimulus}$ ),  $Period$ , and  $Phaseshift$  values are:

$$A_{Stimulus} = \begin{cases} 0.15 & Stimulus = school \\ 0.1 & Stimulus = well \\ 0.05 & Stimulus = road \\ 0.025 & Stimulus = money \end{cases}$$

$$Period = 2$$

$$Phaseshift = -0.2$$

Figure 4.12 depicts the belief shock curves for the four types of stimulus projects using the default parameters discussed above. Given an agent's current belief,  $X_i(t)$  on the x-axis, the value of the belief shock that is added to the agent's belief is found on the y-axis. We see that a stimulus project can have either a positive or negative impact on an agent's belief depending on the agent's belief before the project occurs and the phaseshift value. The phaseshift value represents the prior belief of an agent in which the stimulus project has no effect on the subsequent belief of the agent. For example, in Figure 4.12, the phaseshift value is -0.2, which means any agent with a belief of -0.2 before the stimulus project will remain with the same belief of -0.2 after the project. Meanwhile, all agents with a belief greater than the phaseshift value will increase their beliefs of varying degrees due to the project, while all agents with a belief less than the phaseshift value will decrease their beliefs because of the stimulus project.



*Figure 4.12 – Stimulus Project Default Belief Shocks*

The motivation behind this implementation is that not all villagers will react the same way to a stimulus project by U.S. forces. For instance, agents who hold strongly negative opinions of the United States may only be further enraged by the stimulus project affecting their village, and therefore, may subsequently strengthen their hatred of the U.S. (i.e. the stimulus project would decrease their prior belief). Note, however, that agents strictly have beliefs between  $[-0.5, 0.5]$ . An agent who theoretically would have a belief below -0.5 or above +0.5

due to a stimulus project will subsequently have their beliefs readjusted to the extreme value of -0.5 or +0.5, respectively, in order to stay consistent with the modeling approach.

### Sensitivity to Parameters

In order to find the appropriate default values of the parameters which determine the belief shock values of the agents due to stimulus projects, we conducted a series of experiments aimed at finding reasonable empirical bounds for the parameters. The goal was to find realistic amplitude and phaseshift values given a fixed value of 2 for the period. All experiments were conducted using the same conditions exception for variations in the amplitude and phaseshift values. The details of the experiments are found in the Table 4.31.

<b>Phaseshift and Amplitude Sensitivity Experiment Description</b>	
Network	Large (73 agents)
# Realizations	10
# Steps	10
# Stimulus Projects	3 (steps 5, 7, 9)
Connections per Player (z)	1
Initial, Fixed Strategy	TB [1]
Initial Agent Beliefs	Village Mix 1

*Table 4.31 – Phaseshift and Amplitude Sensitivity Experiment Description*

Table 4.32 below shows the payoff for the U.S. agent at the end of the 10-step game for each experiment using the ‘mean belief’ payoff function. Three stimulus projects were used by the U.S. agent during step 5 (well), step 7 (school), and step 9 (school). All three projects were the same type and occurred during the same step in each experiment for equivalent comparison between experiments with different values for amplitude and phaseshift. The amplitude values ranged from 25% to 200% of the default values shown earlier, while the phaseshift value ranged from 0 to -0.3. As we have seen in previous experiments, because the U.S. and TB agents use the same payoff-maximization methods for choosing their strategies, the relative payoffs during the game are very small and close to zero. Thus, the payoff gain from stimulus projects should also be small in comparison and not too overpowering such that communications with village leaders become obsolete.

The red box in the Table 4.32 shows the empirical bounds for the amplitude and phaseshift values which are deemed realistic payoff gains for the U.S., and the bold number corresponds to the chosen default amplitude and phaseshift values. In general, as the amplitude value increases and the phaseshift value decreases, the payoff increases as expected. However, setting the amplitude too high and the phaseshift too low will result in unrealistically high payoff gains from the stimulus projects (i.e. amplitudes 200% of the default value and a phaseshift of -0.3 results in 0.2289 relative payoff). Meanwhile, setting the amplitude too low and the phaseshift too high will result in pessimistically low payoff gains from the stimulus projects which would make their use in the game negligible. The reader should keep in mind that a total of 3 stimulus projects were used to produce these payoff gains. Thus, the relative payoff gain from a single project would be much less than the payoffs seen below.

		Phaseshift			
		0	-0.1	-0.2	-0.3
Amplitude	25%	0.0141	0.0165	0.0235	0.0225
	50%	0.0085	0.0128	0.039	0.054
	original	0.0199	0.0396	<b>0.0434</b>	0.1003
	150%	0.0366	0.0657	0.096	0.1691
	200%	0.0462	0.0716	0.129	0.2289

*Table 4.32 – Relative U.S. ‘Mean Belief’ Payoffs for Various Phaseshift and Amplitude Values*

### Belief Propagation with Stimulus Projects

We wish to understand how the use of stimulus projects can affect the evolution of beliefs in the network during the course of the two-player game. In Afghanistan, for instance, U.S. military commanders would like to know what effect stimulus projects are having on the opinions of the Afghanistan people, and whether or not such projects are deemed cost effective. After subsequent experiments of analyzing how the stimulus projects affect the propagation of beliefs in the network we have observed the following:

## OBSERVATION

U.S. stimulus projects have the ability to cause long-lasting changes in the beliefs of mutable agents in the network and can subsequently be referred to as game-changing actions available to the U.S. agent.


### Example

We present an example below (see Table 4.33) which shows how the use of stimulus projects can result in long-lasting ‘shocks’ to the beliefs of the agents in the network and leads to an advantageous payoff for the U.S. agent. The initial belief of all mutable agents in the network is neutral. Three different stimulus projects occur during the beginning of the 20-step game—step 1 (well), step 3 (road), and step 5 (school). Moreover, the default values for the amplitude of each stimulus project, the phaseshift (-0.2), and the period (2) were used in this experiment.

<b>Stimulus Project Belief Propagation Experiment Description</b>	
Network	Large (73 agents)
# Realizations	30
# Steps	20
Interactions per Step	300
# Stimulus Projects	3 (steps 1, 3, 5)
Connections per Player (z)	1
Initial, Fixed Strategy	TB [48]
Initial Agent Beliefs	All Neutral (0)


*Table 4.33 – Stimulus Project Belief Propagation Experiment Description*

Figure 4.13 shows the beliefs of the agents in the network at the end of the 20-step game. We notice the drastic difference the use of three stimulus projects has had on the diffusion of beliefs in the network as the mean network belief increased from neutral (0) at the start of the game to slightly in favor of the United States (0.0744) by the end of the game. We conducted the same experiment as described in Table 4.33 with the exception that no stimulus projects occurred during the game. By doing so, we can compare the payoffs and end-game network diagrams to see the impact stimulus projects can have on the diffusion of beliefs.




*Figure 4.13 – Ending Network Diagram: Stimulus Project Belief Propagation Example*

Figure 4.14 shows the end-game network diagram for the experiment where no stimulus projects were used by the U.S. agent under the same conditions as the previous experiment. Upon visual comparison of Figures 4.13 and 4.14 we observe the beneficial impact that stimulus projects can have on the beliefs of the mutable agents in the network.



*Figure 4.14 – Ending Network Diagram: No Stimulus Project Example*

Although the end-game network beliefs for the experiment using the stimulus projects provides a better payoff for the U.S. agent than when no stimulus projects are used, we would also like to observe a plot comparison of the payoffs over time throughout the experiment. More specifically, we would like to determine if the use of stimulus projects can result in long-lasting ‘shocks’ in the beliefs of the agents in the network, or if rather, such belief ‘shocks’ are temporary as the beliefs will eventually return to prior values. Figure 4.15 shows the plot of the relative ‘mean belief’ payoff for the U.S. agent throughout the course of the 20-step game for both experiments. Thus, we see how game-changing the use of stimulus projects can be on the propagation of beliefs as the belief ‘shocks’ resulted in long-lasting changes in the mean belief of the mutable agents in the network.



*Figure 4.15 – Relative Payoff Comparison over Time: Stimulus vs. No Stimulus*

The reason why the belief shocks from the stimulus projects are long-lasting relates to the characteristics of the threshold model, and specifically the peer pressure aspect of the model, which makes it difficult to change an individual’s belief away from the beliefs of the individual’s neighbors. Because a stimulus project causes a belief shock to an agent and all of his neighboring adjacency connections (excluding stubborn agents), the resulting beliefs can be long-lasting since the shock simultaneously affects a cluster of neighboring agents. For instance, if all agents who are affected by the stimulus project have the same belief prior to the stimulus project, then they will all increase (or decrease) their beliefs by the same value, and subsequently have a tendency to remain at the new belief due to the peer pressure aspect of the threshold model.

## Strategic Planning with Limited Resources

A typical question which may arise by the U.S. military is: “Given we have limited resources to fund stimulus projects, when should we use the projects in order to produce the greatest benefit toward our cause?” More specifically, one may wonder whether it is better to spread out the use of stimulus projects over time, or rather concentrate their use during a certain time interval. Moreover, the decision on when to use a stimulus project may depend on the current landscape of the network (i.e. current beliefs of the mutable agents). Thus, there could be instances where it is more advantageous to use a stimulus project sooner rather than later, and vice versa. Below we explore scenarios involving limited resources and subsequently provide general rules for when stimulus projects should be used during the two-player game.

**Question 1:** Should the U.S. agent use stimulus projects sooner rather than later, vice versa, or perhaps not at all?

The first set of experiments is designed to determine the general characteristics of networks, specifically in relation to the beliefs of the agents, which suggest a trend for when to use stimulus projects in the context of Question 1. The default values for the amplitude of each stimulus project and the period (2) were used in this experiment set, while the phaseshift was set to -0.1. The details of the experiments are shown below:

<b>Experiment Set 1 Description</b>	
Network	Large (73 agents)
# Realizations	10
# Steps	10
Interactions per Step	300
# Stimulus Projects	1 (either step 1 or step 9) or None
Stimulus Project Type	School
Connections per Player (z)	1
Initial, Fixed Strategy	TB [70]
Initial Agent Beliefs	Exp 1. Village Mix 1
	Exp 2. Village Mix 2
	Exp 3. All Neutral (0)
	Exp 4. All +0.4
	Exp 5. All -0.4
	Exp 6. Village Mix 6

*Table 4.34 – Strategic Planning with Limited Resources: Experiment Set 1 Description*

For each of the six experiments, there were three different cases tested concerning the usage of stimulus projects. The description of the three cases is seen in the table below.

<b>Experiment Set 1 Case Description</b>	
Case A	1 stimulus project (step 1)
Case B	1 stimulus project (step 9)
Case C	No stimulus project

*Table 4.35 – Strategic Planning with Limited Resources: Experiment Set 1 Case Description*

We compare the ‘mean belief’ payoff function for the U.S. at the end of the game as well as the average ‘mean belief’ payoff during the course of the game (found by averaging the ‘mean belief’ payoff at the end of each step in the 10-step game). Table 4.36 shows the results for each of the six experiments, where the case for each experiment with the highest payoff is highlighted.

	End Absolute Payoff	End Relative Payoff	Average Absolute Payoff	Average Relative Payoff
Exp 1A	0.0027	0.0164	0.0032	0.0169
Exp 1B	0.0230	0.0367	0.0062	0.0199
Exp 1C	0.0159	0.0296	0.0043	0.0180
Exp 2A	0.0083	0.0261	0.0075	0.0253
Exp 2B	0.0422	0.0600	0.0119	0.0297
Exp 2C	0.0321	0.0499	0.0097	0.0275
Exp 3A	0.0149	0.0149	0.0120	0.0120
Exp 3B	0.0086	0.0086	0.0016	0.0016
Exp 3C	-0.0031	-0.0031	-0.0005	-0.0005
Exp 4A	0.4184	0.0184	0.4184	0.0184
Exp 4B	0.3984	-0.0016	0.3913	-0.0087
Exp 4C	0.3732	-0.0268	0.3860	-0.0140
Exp 5C	-0.3694	0.0306	-0.3771	0.0229
Exp 5C	-0.3753	0.0247	-0.3818	0.0182
Exp 5C	-0.3718	0.0282	-0.3807	0.0193
Exp 6C	-0.1676	0.0345	-0.1833	0.0188
Exp 6C	-0.1573	0.0447	-0.1764	0.0257
Exp 6C	-0.1550	0.0470	-0.1761	0.0260

*Table 4.36 – Strategic Planning with Limited Resources: Experiment Set 1 Payoff Table*

The results of these experiments in conjunction with consistent observations from other experiments analyzing the use of stimulus projects have led us to propose the general guidelines

(see below) concerning the timing/use of stimulus projects during the two-player game if resources are limited to the U.S. agent.

### **Guidelines for answering Question 1**

- If the network has completely (or semi) homogenous initial agent beliefs
  - If homogeneously positive (or at least above the phaseshift value):
    - Empirical evidence suggests it is better to conduct stimulus projects as soon as possible during the game
  - If homogeneously negative (or at least below the phaseshift value):
    - Empirical evidence suggests it is better to conduct stimulus projects as soon as possible during the game or perhaps not at all
- If the network has many mixed beliefs, clustered homogeneously by village
  - Empirical evidence shows that it is better to wait until later in the game to use stimulus projects once the beliefs have ‘settled down’

Experiments 1 and 2 show that the U.S. agent will attain the highest mean belief payoff (both the average payoff over time and the end-game payoff) if the stimulus project is used toward the end of the game, rather than the beginning. The initial network beliefs of experiments 1 and 2 are Village Mix 1 and Village Mix 2, respectively, which both contain beliefs that are mixed and clustered homogeneously by village. Thus, the findings from experiments 1 and 2 coincide with the guideline that the U.S. agent should wait until later in the game to use a stimulus project given the initial network beliefs are mixed and clustered homogeneously by village. Experiments 3, 4, and 5 contain initial network beliefs of all neutral (0), all +0.4, and all -0.4, respectively. Based on the guidelines presented above, we expect that the U.S. agent will obtain a higher payoff if the stimulus project is used at the beginning of the game rather than towards the end of the game. As expected, we find that the highest mean belief payoff for the U.S. agent occurs when the stimulus project is used during step 1 of the 10-step game. One key observation to point out, however, is that a higher payoff is obtained for the U.S. agent in experiment 5 when no stimulus project is used, rather than the stimulus project being used at the end of the game during step 9. This finding shows that for networks which are mostly (or completely) anti-US in their initial agent beliefs that the use of stimulus projects may actual do

more harm than good for the U.S. payoff. Experiment 6 is a good example of this finding, as the initial network is mostly anti-US in their sentiments, and empirical evidence shows us that the best payoff for the U.S. agent occurs when no stimulus project is used during the game.

**Question 2:** Should the U.S. agent spread out the use of stimulus projects, or rather concentrate their use consecutively during a specific time period?

Another realistic question which may arise for the U.S. player is: “Given the use of stimulus projects will improve my payoff during the game, should the projects be spread out during the game or should I instead concentrate their use during a specific time interval?” We present the experiment and case descriptions which analyzes this question in the tables shown below.

<b>Experiment Set 2 Description</b>	
Network	Large (73 agents)
# Realizations	10
# Steps	20
Interactions per Step	300
# Stimulus Projects	5
Stimulus Project Type	School
Connections per Player (z)	1
Initial, Fixed Strategy	TB [1]
Initial Agent Beliefs	All Neutral (0)

*Table 4.37 – Strategic Planning with Limited Resources: Experiment Set 2 Description*

<b>Experiment Set 2 Case Description</b>	
Exp A	5 stimulus projects (concentrated—used during steps 1, 3, 5, 7 and 9)
Exp B	5 stimulus projects (concentrated—used during steps 11, 13, 15, 17, and 19)
Exp C	5 stimulus projects (spread out—used during steps 1, 5, 9, 13, and 17)


*Table 4.38 – Strategic Planning with Limited Resources: Experiment Set 2 Case Description*

Similar to previous experiment set, we compare the ‘mean belief’ payoff for the U.S. agent at both the end of the two-player game as well as the average ‘mean belief’ payoff during the course of the game. The results of this experiment are presented in Table 4.39 where we highlight the experiment case which yields the highest payoff for the U.S. agent.

	End Absolute Payoff	End Relative Payoff	Average Absolute Payoff	Average Relative Payoff
Exp A	0.1440	0.1440	0.0954	0.0954
Exp B	0.1005	0.1005	0.0241	0.0241
Exp C	0.0535	0.0535	0.0336	0.0336

*Table 4.39 – Strategic Planning with Limited Resources: Experiment Set 2 Payoff Table*

Upon viewing the payoff table above as well as the ‘mean belief’ payoff plots of the three cases in Figure 4.16 below, we observe the noticeable advantage the U.S. agent obtains in ‘mean belief’ payoff when the stimulus projects are used consecutively, rather than spread out evenly over the course of the entire game.



*Figure 4.16 – Strategic Planning with Limited Resources: Experiment Set 2 Payoff Plots*

Both cases where the stimulus projects are used consecutively (either at the beginning or end of the game) yields a higher payoff than spreading out the use of the projects over the course of the entire game. Moreover, the highest payoff occurs when the stimulus projects are concentrated at the beginning of the game. Empirical evidence, including the results of this experiment, supports the notion that stimulus projects can have a compounding effect in payoff gains when they are used in a consecutive manner which results in higher payoffs than distributing the stimulus projects evenly over time. Thus, given the U.S. agent has limited resources to devote to stimulus projects, we recommend the projects be used in a consecutive manner during the game in order to produce the greatest benefit for the U.S. cause.

### **4.3.2 Assassinations**

While the U.S. agent has additional actions, such as stimulus projects, which can affect belief propagation in the network, the Taliban agent is also given the ability to target mutable agents for assassinations in order to instill negative ‘belief shocks’ in the network. We provide a description of the Taliban agent’s ability to carry out assassinations below:

#### **Description**

We implement the ability for the Taliban agent to conduct assassinations during the two-player game as follows:

- The number of assassinations that will occur during the game is predetermined.
- A limit of one assassination can occur per TB step during the game.
- No assassinations may occur during a U.S. step.
- Given the predetermined number of assassinations set to occur during the game, we either randomly or selectively choose the TB steps for which assassinations will occur.
- Given an assassination is set to occur during a TB step, the agent that is assassinated is an agent that is presently communicating with the U.S. agent (i.e. is a chosen strategy by the U.S. agent). The reason for this implementation is that the Taliban is more likely to assassinate those people who show favor to the U.S.

cause, and thus, someone who is communicating with the U.S. is a likely target for an assassination.

- The effect of the assassination:
  - The influence level of the assassinated agent drops by one level. For instance, if the agent was previously forceful+ he is now forceful, and if the agent was previously forceful he is now regular. This is to mimic the effect that the heir to the position will most likely be the agent's son and is less likely to be as influential as the father. Thus, assassinated agents are not removed from the network, but rather replaced by the next of kin.
  - The influence level of the agent will return to the previous level before the assassination occurred at the conclusion of the current step in the game.
  - All agents which are adjacent to the agent that is assassinated (as well as the next of kin replacing the assassinated agent) will undergo a negative 'belief shock' whereby the assassination will decrease their current beliefs based on their influence level as described below:

DEFINTION:

We define the set of mutable agents who receive a negative belief shock from an assassination by the TB agent that occurs during step k of the two-player game as:

$$V_{k,Assassination} \quad \text{where } Assassination = \text{index of assassinated agent}$$

The belief of agent  $i$  is updated due to the assassination occurring during step k and time t as follows:

$$X_i(t+1) = \begin{cases} X_i(t) + shock & \forall i \in V_{k,Assassination} \cap V_R \\ X_i(t) + \frac{1}{2} * shock & \forall i \in V_{k,Assassination} \cap V_F \\ X_i(t) + \frac{1}{4} * shock & \forall i \in V_{k,Assassination} \cap V_{F+} \\ X_i(t) & \text{else} \end{cases}$$

where  $shock \leq 0$  and the default value of  $shock = -0.1$

Regular agents are the least influential agents in the network, and therefore, more likely to follow or comply in their opinions due to fear tactics imposed by the use (or threat) of an assassination. We reason that the more influential an agent is, the less likely they are to succumb to the threat or use of an assassination. Thus, we reflect this notion by decreasing the negative belief shock that results from an assassination as the influence level of the affected agent increases.

### **Belief Propagation with Assassinations**

We wish to explore how the use of assassinations can affect the dissemination of beliefs in the network during the course of the two-player game. In Afghanistan, for instance, U.S. military commanders may be interested in determining how the beliefs of the local populace within a village reacts to the news that a leader in their village was assassinated by the Taliban for speaking with and showing favor to the United States. Would the fear of subsequent assassination attempts detrimentally affect the progress the U.S. has made with the opinions of the people in a targeted village? After conducting various experiments which examine how assassinations affect the propagation of beliefs in the network we have observed the following:

### **OBSERVATION**

Taliban assassinations have the ability to cause long-lasting changes in the beliefs of mutable agents in the network and can subsequently be referred to as game-changing actions available to the TB agent.

### **Example**


We present an example below which shows how the use of assassinations can result in long-lasting negative ‘shocks’ to the beliefs of the agents in the network and leads to an advantageous payoff for the Taliban agent. These long-lasting ‘shocks’ are analogous to the long-lasting shocks brought on by the use of stimulus projects by the U.S. agent, with the exception that the shocks from assassinations are intended to promote more favorable beliefs for the Taliban cause. For the experiment presented below, the initial belief of all mutable agents in the network is neutral, and the Taliban agent carries out three assassinations during the beginning of the 20-step game—steps 2, 4 and 6. Moreover, the default value for the belief shock from an

assassination (-0.1) was used during the experiment. Table 4.40 provides further details about the parameters of the experiment.

Assassination Belief Propagation Experiment Description	
Network	Large (73 agents)
# Realizations	30
# Steps	20
Interactions per Step	300
# Stimulus Projects	3 (steps 2, 4, 6)
Connections per Player (z)	1
Initial, Fixed Strategy	TB [48]
Initial Agent Beliefs	All Neutral (0)

*Table 4.40 – Assassination Belief Propagation Experiment Description*


Figure 4.17 shows the beliefs of the agents in the network at the end of the 20-step game. During the experiment, the three agents which were assassinated were agents 49 (step 2), 48 (step 4), and 54 (step 6).



*Figure 4.17 – Ending Network Diagram: Assassination Belief Propagation Example*


In order to observe the effect that the assassinations had on the evolution of beliefs in the network over time, we conducted the same experiment as described in Table 4.40 with the

exception that no assassinations occurred during the game. By doing so, we can compare the payoffs and end-game network diagrams to see the impact assassinations can have on the diffusion of beliefs. Figure 4.18 shows the end-game network diagram for the experiment where no assassinations were used by the TB agent.



*Figure 4.18 – Ending Network Diagram: No Assassination Project Example*

Although the end-game network beliefs for the experiment using the three assassinations yields a better payoff for the Taliban agent than when no assassinations were used, we would also like to observe a plot comparison of the payoffs over time throughout the experiment. More specifically, we would like to determine if the use of assassinations can result in long-lasting ‘shocks’ in the beliefs of the agents in the network, or if rather, such belief ‘shocks’ are temporary as the beliefs will eventually return to prior values. Figure 4.19 shows the plot of the relative ‘mean belief’ payoff for the U.S. agent (the payoff for the Taliban agent is simply the negative value of the U.S. agent’s payoff since this is a zero-sum game) throughout the course of the 20-step game for both experiments. Thus, we see how game-changing the use of assassinations can be on the propagation of beliefs as the negative belief ‘shocks’ resulted in long-lasting changes in the mean belief of the mutable agents in the network.



*Figure 4.19 – Relative Payoff Comparison over Time: Assassination vs. No Assassination*

Much like the long-lasting shocks brought upon the beliefs in the network due to stimulus projects, the reason why the belief shocks from the assassinations are long-lasting relates to the characteristics of the threshold model, and specifically the peer pressure aspect of the model, which makes it difficult to change an individual’s belief away from the beliefs of the individual’s neighbors. Because an assassination causes a belief shock to the targeted agent and all of his neighboring adjacency connections (excluding stubborn agents) as well as temporarily degrades the level of influence of the assassinated agent by one level, the resulting beliefs can be long-lasting since the shock simultaneously affects a cluster of neighboring agents. Thus, once the agents who are affected by the assassination decrease their belief, they will subsequently have a tendency to remain at the new belief due to the peer pressure aspect of the threshold model.

## 4.4 Summary of Experiments

The formulation of the proposed threshold model and subsequent experimentation and analysis of the dynamic, two-player game on the network model have resulted in very interesting and surprising conclusions, such as:

- The long-term beliefs of the mutable agents in the network are *dependent* on the initial beliefs, which is due to the peer pressure aspect of the threshold model which affects the belief exchanges during pairwise interactions.
- The two-player game operates in a transient, “chess game” setting which allows for the players to analyze their opponent’s strategy as well as the current state of the network and then subsequently make dynamic changes in their strategies in order to improve their payoffs.
- **SURPRISING RESULT:** The population-focused actions (stimulus projects and assassinations) introduced in Section 4.3 cause *long-lasting changes* in the beliefs of mutable agents in the network and can subsequently be deemed as game-changing actions available to the players. This conclusion gives credibility to the notion that “actions speak louder than words” when trying to influence the opinions of others!
- Finally, we implemented the ability to penalize the payoffs for identical strategies among opposing stubborn agents (U.S. and Taliban) in order to disincentivize their appeal during the two-player game, which thus creates a more realistic setting for the game. However, as experiments have shown, the  $\lambda$  penalty factor is not always necessary to prevent identical strategies as stubborn agents ideally target different mutable agents of opposite beliefs.

Moreover, in Section 4.1 we discussed the characteristics of strategies chosen by the players during the dynamic, two-player game on the proposed threshold model in the transient setting. Based on empirical evidence, we presented four key observations:

1. Stubborn agents target influential agents of opposite belief.
2. In the transient setting, stubborn agents target influential agents who have many neighbors (large  $N_i$ 's) of lesser influence level, but whose neighbors live in small neighborhoods (small  $N_j$ 's).
3. For villages containing more than one influential leader of similar beliefs, stubborn agents should sequentially target some, or all, of those village leaders.
4. Stubborn agents target influential agents with low threshold values who also have neighbors with low threshold values.

We also showed how different payoff functions influence the strategies targeted by the players, as well as the sensitivity of the strategies to other model parameters, such as the interaction-type probabilities ( $\alpha, \beta, and \gamma$  values), the initial strategy chosen by a stubborn agent during step 0 of the game, and the imposed penalty for restricting identical strategies ( $\lambda$ ). We showed that the threshold model is fairly robust to changes in the  $\alpha, \beta,$  and  $\gamma$  values, which is important because our estimates of these values do not need to be highly accurate. Moreover, we showed that the strategies chosen by the players to maximize their payoffs are sensitive to the specific  $\mu$  parameter as well as the initial strategy chosen by a stubborn agent during step. Additionally, the penalty for limiting identical strategies ( $\lambda$ ) is not always necessary, as stubborn agents find agents of opposite belief as desirable strategies, and thus, networks representing a wide array of different agent beliefs enable stubborn agents to focus their strategy selection on influencing different agents.

Later, in Section 4.2, we analyzed the run time performance of three methods used to find strategies during the two-player game: exhaustive enumeration, ‘Selective Search’, and ‘Sequential Search’. Scenarios in which the stubborn agents have  $z > 1$  connections during the game can be computationally burdensome for exhaustive enumeration and even ‘Selective Search’. However, the ‘Sequential Search’ greedy algorithm scales extremely well with increasing complexity without tremendously sacrificing the quality of the solutions compared to exhaustive enumeration and ‘Selective Search’.

Finally, Section 4.3 introduced population-focused actions—stimulus projects and assassinations—which can be used by the US and TB agents, respectively, to further affect the

propagation of beliefs during the two-player game. The belief shocks which occurred as a result of the stimulus projects and assassinations affected the beliefs of a neighborhood of mutable agents. Thus, due to the nature of the peer pressure influence which is built into the threshold model, these population-focused actions surprisingly resulted in *long-lasting changes* in the beliefs of the mutable agents in the network. Subsequently, this astounding result gives credence to the saying that “actions speak louder than words” when attempting to influence the opinions of other people!

[This Page Intentionally Left Blank]

## 5 Future Work Recommendations

In this chapter we explore several interesting topics as enhancements to the current peer-influence, social network and two-player game. First, we discuss the possibility of considering a dynamic network which evolves over time. Next, we discuss modifying the two-player game as an imperfect information game whereby opposing players are limited in the information they know about each other's strategies, the beliefs and influence levels of the agents in the network, and/or when their opponents will update their strategies. Moreover, we discuss the implications of 'tipping points' in the two-player game, whereby future research could provide better insights into the underlying parameters which factor into the strategy choices made by the players. Later, we suggest the introduction of additional population-focused actions which affect belief propagation, such as using a surge of troops to improve security in a village, spreading propaganda, or organizing meetings with the populace. Finally, we introduce the concept of using iterated reasoning in order to predict opponent strategies with the objective of finding better strategies for the stubborn agents in the two-player game.

### 5.1 Dynamic Networks

During the modeling approach and implementation of the two-player game, we made the assumption of a static (fixed) network whereby the location of all agents and their adjacency connections were known and constant. However, in a realistic setting, in which the analysis of the network occurs for an extended period of time, this assumption may no longer be valid as the social network evolves over time with respect to the number of agents, the adjacency connections between agents, and the influence level of the agents. Thus, modifying the social network to incorporate the idea of a dynamic network that can change during the course of the two-player game unveils a new area of interesting analysis concerning this work.

Moreover, although the primary focus of this research was on analyzing the belief diffusion and characteristics of strategies chosen by the stubborn agents on a realistic Pashtun network, future work on a wide variety of different network structures could provide additional insights into the characteristics of strategies chosen by the stubborn agents during the two-player game. While different network types, such as line, ring (circle), or star, may not be realistic

depictions of societal networks, their unique topologies may result in notable differences in the behavior of the stubborn agents when choosing strategies in order to maximize their payoffs during the game.

## **5.2 Imperfect Information Games**

The two-player game implementation assumes a symmetric, perfect information game whereby both players know each other's strategies, the exact beliefs and levels of influence of all agents in the network, as well as when each player will update their strategies. These assumptions are not valid in a realistic setting whereby players do not necessarily know who their opponents are trying to influence, let alone when their opponents will switch their strategies to targeting new agents. Moreover, having accurate estimates of the beliefs of all agents in the network and their influence levels at any given time is a difficult task as these values may be subject to change on a regular basis.

Thus, future work could analyze the two-player game as an imperfect information game whereby we limit a player's knowledge about their opponent's strategy, the beliefs and influence level of the agents in the network, and/or the times during the game when the opponent will update their strategy. By reformulating the game as such, players would attempt to make educated guesses about their opponent's strategy through observation methods, such as viewing the percentage change in agent beliefs over time. Therefore, as the network beliefs evolve over the course of the game, a player would continually update their confidence in the strategy they believe their opponent has chosen. Moreover, if a player is restricted in their knowledge of the beliefs of particular agents in the network, they may consider inferring their beliefs by looking at the beliefs of their neighboring agents.


## **5.3 Tipping Points**

In Chapter 4 we observed how different parameters weigh into the decisions made by stubborn agents on choosing their targeting strategies. The topology of the network (the number of agents, their level of influence, and their connectivity), the threshold values, and the current beliefs of the agents all factor into the quality of strategies chosen by the players to maximize their payoffs. However, the topic of tipping points presents an interesting area for future

research with the ultimate aim at producing generalized relationships (formulas) between the various parameters which affect the strategies chosen by the stubborn agents. For instance, we present the following example below which gives insight into the relationship between the threshold values and beliefs of the agents. We wish to determine the ‘tipping point’ at which a stubborn agent will change his strategy based on changes in threshold values and agent beliefs.

### Example

The network diagram in Figure 5.1 contains three symmetric villages, each containing 1 forceful+ agent, 3 forceful agents, and 4 regular agents. During the two-player game on this network, given the initial TB strategy (agent 4), the default threshold values, and ‘mean belief’ payoff function, the U.S. agent’s first strategy selection (step 1) for maximizing his payoff 100 interactions into the future will be to target Village 3 (specifically agent 3) due to the high payoff obtained by targeting a village of opposite belief (-0.5). However, we wish to determine the ‘tipping point’ for the threshold values of the forceful+ and forceful agents in Village 3 which would cause the US agent to instead target Village 2 during step 1. There is an inherent tradeoff between the threshold values and the beliefs of agents since agents of opposite belief yield higher payoffs for stubborn agents, while agents with higher threshold values are harder to influence and subsequently result in lower payoffs.




*Figure 5.1 – Tipping Point: Network Diagram*

For this simple experiment set, we conduct five experiments where the beliefs of all agents in Village 2 differ between experiments ranging from 0 to -0.4, while the beliefs of all agents in Villages 1 and 3 remain the same (+0.5 and -0.5, respectively). The network diagram in Figure 5.1 shows the initial network for experiment case #5. We subsequently determine at which point (i.e. tipping point) the U.S. agent will change his strategy during step 1 from targeting Village 3 to Village 2 by increasing the threshold values of the forceful+ and forceful agents in Village 3 until this change in strategy occurs. Table 5.1 details the experiment set and shows the resulting threshold value tipping points for the five experiment cases.

				Threshold Tipping Point of Village 3 Leaders	
	Village 1 Belief	Village 2 Belief	Village 3 Belief	US Agent Targets Village 3 for all Threshold Values Between:	US Agent Targets Village 2 for all Threshold Values Between:
<b>Case #1:</b>	+0.5	-0.4	-0.5	[0, 0.7661]	[0.7662, 1]
<b>Case #2:</b>	+0.5	-0.3	-0.5	[0, 0.7831]	[0.7832, 1]
<b>Case #3:</b>	+0.5	-0.2	-0.5	[0, 0.7892]	[0.7893, 1]
<b>Case #4:</b>	+0.5	-0.1	-0.5	[0, 0.8061]	[0.8062, 1]
<b>Case #5:</b>	+0.5	0	-0.5	[0, 0.8320]	[0.8321, 1]

*Table 5.1 – Tipping Point: Results*

Moreover, Figure 5.2 shows the linear relationship which exists for the tipping points between the threshold values and beliefs of the agents for the five experiment cases. Thus, we conclude that the strategies chosen (and subsequent payoffs received) by a stubborn agent depend linearly on changes in threshold values and the beliefs of the agents in the network. Hopefully, future research into this topic would produce a generalized relationship between all interacting variables which affect the strategies that are chosen by the stubborn agents.



*Figure 5.2 – Tipping Point: Linear Relationship Plot*

## 5.4 Applying Iterated Reasoning to Predict Opponent Strategies

During the course of the dynamic, two-player game, the two stubborn agents have perfect information about the current state of the network, including their opponent’s strategy. Although the game operates in the transient setting with continual updates in strategies for both players, the implementation of a smarter algorithm which uses iterated reasoning to predict opponent strategies would be useful. Currently, in the two-player game, the players select the best strategies given the current state of the network and their opponent’s current strategy, without looking ahead into the future to predict what the opponent’s strategies may be in response to their newly updated strategies. Much like in the game of chess, whereby players ideally want to ‘checkmate’ their opponent as a result of analyzing the game several moves into the future in order to make the best possible move, the stubborn agents in the two-player game would be better served in their strategy selections if they applied the same iterated reasoning to predict their opponent’s strategies.

## 5.5 Additional Population-Focused Actions Which Affect Belief Propagation

In Section 4.3 we introduced population-focused actions (stimulus projects and assassinations) which gave the US agent and TB agent additional capabilities for affecting the beliefs in the network, besides choosing connection strategies. Experiments demonstrated that the belief shocks resulting from these population-focused actions produced long-lasting changes on the beliefs of the agents in the network. However, besides pairwise interactions (all agents in the network), stimulus projects (US agent), and assassinations (TB agent), no other actions are available to the agents in order to impact the diffusion of beliefs in the network. This presents an area for future research by enhancing the model to allow for even more population-focused actions which can affect the propagation of beliefs. The implementation of the following strategy options could provide valuable insights into how other actions affect the attitudes of the agents in the network:

- ***Surge of Troops***: If the counterinsurgents significantly increase the security of a particular village or district by increasing the number of troops, how will this impact the attitudes of the agents in the network?
- ***Propaganda***: How will propaganda—posters, pamphlets, radio, internet, loud speakers, etc.—supporting either the insurgents or counterinsurgents affect agent beliefs?
- ***Meetings and Negotiations***: Can organized village meetings or negotiations prove as valuable options for the insurgents or counterinsurgents in achieving significant gains in popular support of the people?

The strategy options presented above as well as many others could be formulated in the current model for use by the agents during the two-player game. We have already seen the surprising ability for stimulus projects and assassinations to produce long-lasting changes in the network beliefs, and it would be of further interest in determining the affects of additional population-focused actions on the network attitudes.

[This Page Intentionally Left Blank]

## 6 Summary and Conclusions

In this chapter we summarize the work presented in this thesis and give conclusions obtained through experiments and analysis of the proposed threshold model and dynamic, two-player game. This thesis considered the problem of finding the best strategies for maximally influencing the beliefs of individuals in social networks toward supporting the counterinsurgency efforts of the United States in Afghanistan.

In Chapter 2 we described the nature of insurgencies and counterinsurgencies and the inherent struggle between both sides for the popular support of the people. The U.S. military has limited resources and personnel, and therefore, critical decisions must be made by ground forces on whom to selectively target with the end goal of winning the ‘hearts and minds’ of the most people. Because the success of any counterinsurgency ultimately hinges on the ability to win the ‘hearts and minds’ of the populace, this thesis focused on the extremely difficult task of determining which individuals to engage through nonlethal means with the objective of winning the most support of the populace.

In Chapter 3 we described the formulation of a proposed threshold model which is a stochastic, pairwise-interaction model with mutable agents whose beliefs can change and with immutable (stubborn) agents whose beliefs do not change. The U.S. military forces and Taliban insurgents are represented by a stubborn agent (US and TB agent), and they have opposite, immutable beliefs of  $+0.5$  and  $-0.5$ , respectively. We introduced a dynamic, two-player game whereby each player controls a predetermined number of connections from their respective stubborn agent to mutable agents in the network, and the objective of the game for each player is to maximize their respective payoff function. Two payoff functions (‘mean belief’ and ‘number of agents won’) are defined for the players. The two-player game operates in a transient, “chess game” setting which allows for the players to analyze their opponent’s strategy as well as the current state of the network and then subsequently make dynamic changes in their strategies in order to improve their payoffs. Furthermore, we defined how the players locate strategies during the game and introduced two heuristic methods, the ‘Selective Search’ Heuristic and the ‘Sequential Search’ Greedy Algorithm which are used to significantly reduce the run time

needed to find solutions for complex problems where the use of exhaustive enumeration becomes too cumbersome.

In Chapter 4 we presented the results and analysis of the proposed threshold model and dynamic, two-player game. The dynamic, two player game allows players to update their strategies in the transient setting which creates a more realistic game environment than previous modeling approaches [8, 13]. Furthermore, we showed that the long-term beliefs of the mutable agents in the network are *dependent* on the initial beliefs, which is due to the peer pressure aspect of the threshold model which affects the belief exchanges during pairwise interactions. The population-focused actions (stimulus projects and assassinations) introduced in Section 4.3 cause *long-lasting changes* in the beliefs of mutable agents in the network and can subsequently be deemed as game-changing actions available to the players. This conclusion gives credibility to the notion that “actions speak louder than words” when trying to influence the opinions of others! Even though we implemented the ability to penalize the payoffs for identical strategies among opposing stubborn agents (U.S. and Taliban) in order to create a more realistic setting for the game, we showed that the  $\lambda$  penalty factor is not always necessary to prevent identical strategies as stubborn agents are incentivized on their own to target different mutable agents of opposite beliefs.

Moreover, in Section 4.1 we discussed the characteristics of strategies chosen by the players during the dynamic, two-player game on the proposed threshold model in the transient setting. Based on empirical evidence, we presented four key observations:

1. Stubborn agents target influential agents of opposite belief.
2. In the transient setting, stubborn agents target influential agents who have many neighbors (large  $N_i$ 's) of lesser influence level, but whose neighbors live in small neighborhoods (small  $N_j$ 's).
3. For villages containing more than one influential leader of similar beliefs, stubborn agents should sequentially target some, or all, of those village leaders.
4. Stubborn agents target influential agents with low threshold values who also have neighbors with low threshold values.

We also showed how different payoff functions influence the strategies targeted by the players, as well as the sensitivity of the strategies to other model parameters, such as the interaction-type probabilities ( $\alpha$ ,  $\beta$ , and  $\gamma$  values), the initial strategy chosen by a stubborn agent during step 0 of the game, and the imposed penalty for restricting identical strategies ( $\lambda$ ). We demonstrated that the threshold model is fairly robust to changes in the  $\alpha$ ,  $\beta$ , and  $\gamma$  values, which is important because our estimates of these values do not need to be highly accurate. Moreover, we showed that the strategies chosen by the players to maximize their payoffs are sensitive to the specific  $\mu$  parameter as well as the initial strategy chosen by a stubborn agent during step. Additionally, the penalty for limiting identical strategies ( $\lambda$ ) is not always necessary, as stubborn agents find agents of opposite belief as desirable strategies, and thus, networks representing a wide array of different agent beliefs enable stubborn agents to focus their strategy selection on influencing different agents.

Later, in Section 4.2, we analyzed the run time performance of three methods used to find strategies during the two-player game: exhaustive enumeration, ‘Selective Search’, and ‘Sequential Search’. Scenarios in which the stubborn agents have  $z > 1$  connections during the game can be computationally burdensome for exhaustive enumeration and even ‘Selective Search’. However, the ‘Sequential Search’ greedy algorithm scales extremely well with increasing complexity without tremendously sacrificing the quality of the solutions compared to exhaustive enumeration and ‘Selective Search’.

In Section 4.3 we introduced population-focused actions—stimulus projects and assassinations—which can be used by the US and TB agents, respectively, to further affect the propagation of beliefs during the two-player game. The belief shocks which occurred as a result of the stimulus projects and assassinations affected the beliefs of a neighborhood of mutable agents. Thus, due to the nature of the peer pressure influence which is built into the threshold model, these population-focused actions surprisingly resulted in *long-lasting changes* in the beliefs of the mutable agents in the network. Subsequently, this astounding result gives credence to the saying that “actions speak louder than words” when attempting to influence the opinions of other people!

Finally, in Chapter 5, we identified interesting areas for future research. The five areas for future research were: (1) dynamic networks, (2) imperfect information games, (3) tipping points, (4) additional population-focused actions which affect belief propagation, and (5) applying iterated reasoning to predict opponent strategies. These five areas, among others, offer new opportunities to enhance the current modeling approach and will provide new insights into the opinion dynamics of the two-player game.

This thesis focused on the extremely difficult task of determining which individuals to engage through nonlethal means in order to gain the most support of the populace during counterinsurgency efforts. Although the task of winning the most popular support during counterinsurgency efforts is tremendously difficult and complex due to the multitude of variables that need to be considered at any given time, the modeling formulation and game theoretic approach provides a systematic method for the U.S. military to use in their evaluation of potential nonlethal targeting strategies. Thus, rather than solely relying on intuition and limited intelligence, military commanders are capable of assigning quantitative values to their list of high priority nonlethal targets based on the predicted effects on the opinions of the populace. Hopefully, the decision methods, analysis, and insights presented in this thesis will serve as an important step towards efficiently and effectively winning the ‘hearts and minds’ of the population in Afghanistan as well as future areas of concern for the United States military.

[This Page Intentionally Left Blank]

## Appendix A: Variable Notation and Definitions

<u>Notation</u>	<u>Description</u>
$V_R$	<i>set of all regular agents</i>
$V_F$	<i>set of all forceful agents</i>
$V_{F+}$	<i>set of all forceful + agents</i>
$V_{US}$	<i>set of all US agents</i>
$V_{TB}$	<i>set of all TB agents</i>
$V_M$	<i>set of all mutable agents = <math>V_R \cup V_F \cup V_{F+}</math></i>
$V_S$	<i>set of all stubborn (immutable) agents = <math>V_{US} \cup V_{TB}</math></i>
$V$	<i>set of all agents = <math>V_R \cup V_F \cup V_{F+} \cup V_{US} \cup V_{TB} = V_M \cup V_S</math></i>
$ V  = n$	<i>total number of agents (nodes) in the network</i>
$X_i(t)$	<i>belief of agent <math>i</math> (active agent) at interaction <math>t</math></i>
$X_j(t)$	<i>belief of agent <math>j</math> (selected neighboring agent) at interaction <math>t</math></i>
$\tau_i$	<i>threshold value of agent <math>i</math></i>
$\tau_j$	<i>threshold value of agent <math>j</math></i>
$\eta_{i \setminus j}(t)$	<i>set of <math>i</math>'s neighbors at interaction <math>t</math> (excluding <math>j</math>)</i>
$\eta_{j \setminus i}(t)$	<i>set of <math>j</math>'s neighbors at interaction <math>t</math> (excluding <math>i</math>)</i>
$\eta_{i \setminus j, A}(t)$	<i>set of <math>i</math>'s neighbors at interaction <math>t</math> (excluding <math>j</math>) who have beliefs <math>&gt; X_i(t)</math></i>
$\eta_{i \setminus j, B}(t)$	<i>set of <math>i</math>'s neighbors at interaction <math>t</math> (excluding <math>j</math>) who have beliefs <math>&lt; X_i(t)</math></i>
$\eta_{j \setminus i, A}(t)$	<i>set of <math>j</math>'s neighbors at interaction <math>t</math> (excluding <math>i</math>) who have beliefs <math>&gt; X_j(t)</math></i>
$\eta_{j \setminus i, B}(t)$	<i>set of <math>j</math>'s neighbors at interaction <math>t</math> (excluding <math>i</math>) who have beliefs <math>&lt; X_j(t)</math></i>
$\alpha_{ij}$	<i>prob. that agent <math>i</math> will attempt to forcefully impart <math>(1 - \tau_j)</math> of its attitude on agent <math>j</math></i>

$\beta_{ij}$	<i>prob. that agent <math>i</math> will attempt to reach a consensus with agent <math>j</math> equal to the avg. of their prior beliefs</i>
$\gamma_{ij}$	<i>prob. that agents <math>i</math> and <math>j</math> will undergo an identity interaction and exhibit no change in their prior beliefs</i>
<b>Reminder:</b> $\alpha_{ij} + \beta_{ij} + \gamma_{ij} = 1, \forall i, j$ where $i \neq j$ (since agents cannot interact with themselves)	

## Appendix B: Experimental Data and Results

Agent (Node)	1. All Neutral	2. Village Mix 0	3. Village Mix 1	4. Village Mix 2	5. Village Mix 3	6. Village Mix 4	7. Village Mix 5	8. Village Mix 6	9. All Random
1	0	0	0.1	0.15	0.3	0.2	0.2	-0.2	-0.3680
2	0	0	0.1	0.15	0.3	0.2	0.2	-0.2	0.3873
3	0	0	0.1	0.15	0.3	0.2	0.2	-0.2	-0.1430
4	0	0	0.1	0.15	0.3	0.2	0.2	-0.2	0.0959
5	0	0	0.1	0.15	0.3	0.2	0.2	-0.2	-0.0801
6	0	0	0.1	0.15	0.3	0.2	0.2	-0.2	-0.4745
7	0	0	0.1	0.15	0.3	0.2	0.2	-0.2	-0.1469
8	0	0	0.1	0.15	0.3	0.2	0.2	-0.2	0.2324
9	0	0	0.1	0.15	0.3	0.2	0.2	-0.2	0.0924
10	0	0	0.1	0.15	0.3	0.2	0.2	-0.2	0.1268
11	0	-0.3	0.3	0.4	0.3	0.3	0.3	-0.3	0.1698
12	0	-0.3	0.3	0.4	0.3	0.3	0.3	-0.3	0.3033
13	0	-0.3	0.3	0.4	0.3	0.3	0.3	-0.3	-0.3440
14	0	-0.3	0.3	0.4	0.3	0.3	0.3	-0.3	0.2122
15	0	-0.3	0.3	0.4	0.3	0.3	0.3	-0.3	-0.2767
16	0	-0.3	0.3	0.4	0.3	0.3	0.3	-0.3	0.2781
17	0	-0.3	0.3	0.4	0.3	0.3	0.3	-0.3	0.1173
18	0	-0.3	0.3	0.4	0.3	0.3	0.3	-0.3	-0.2890
19	0	-0.3	0.3	0.4	0.3	0.3	0.3	-0.3	0.0842
20	0	0.3	-0.3	-0.4	-0.3	-0.4	0.4	-0.4	-0.1861
21	0	0.3	-0.3	-0.4	-0.3	-0.4	0.4	-0.4	-0.3452
22	0	0.3	-0.3	-0.4	-0.3	-0.4	0.4	-0.4	0.0741
23	0	0.3	-0.3	-0.4	-0.3	-0.4	0.4	-0.4	-0.4569
24	0	0.3	-0.3	-0.4	-0.3	-0.4	0.4	-0.4	-0.2875
25	0	0.3	-0.3	-0.4	-0.3	-0.4	0.4	-0.4	0.1266
26	0	0.3	-0.3	-0.4	-0.3	-0.4	0.4	-0.4	-0.0244
27	0	0.3	-0.3	-0.4	-0.3	-0.4	0.4	-0.4	-0.2511
28	0	0.3	-0.3	-0.4	-0.3	-0.4	0.4	-0.4	0.2522
29	0	0.3	-0.3	-0.4	-0.3	-0.4	0.4	-0.4	0.2466
30	0	0	-0.1	-0.15	-0.3	-0.2	0.1	-0.1	-0.0397
31	0	0	-0.1	-0.15	-0.3	-0.2	0.1	-0.1	0.3640
32	0	0	-0.1	-0.15	-0.3	-0.2	0.1	-0.1	-0.2392
33	0	0	-0.1	-0.15	-0.3	-0.2	0.1	-0.1	-0.2083
34	0	0	-0.1	-0.15	-0.3	-0.2	0.1	-0.1	-0.0910
35	0	0	-0.1	-0.15	-0.3	-0.2	0.1	-0.1	0.0419
36	0	0	-0.1	-0.15	-0.3	-0.2	0.1	-0.1	0.4305
37	0	0	-0.1	-0.15	-0.3	-0.2	0.1	-0.1	0.3285
38	0	-0.2	-0.3	-0.4	-0.3	-0.3	0.25	-0.25	0.4787
39	0	-0.2	-0.3	-0.4	-0.3	-0.3	0.25	-0.25	0.4970
40	0	-0.2	-0.3	-0.4	-0.3	-0.3	0.25	-0.25	-0.2466
41	0	-0.2	-0.3	-0.4	-0.3	-0.3	0.25	-0.25	0.2415
42	0	-0.2	-0.3	-0.4	-0.3	-0.3	0.25	-0.25	-0.3148
43	0	-0.2	-0.3	-0.4	-0.3	-0.3	0.25	-0.25	-0.4339
44	0	-0.2	-0.3	-0.4	-0.3	-0.3	0.25	-0.25	-0.2194
45	0	-0.2	-0.3	-0.4	-0.3	-0.3	0.25	-0.25	0.0443
46	0	-0.2	-0.3	-0.4	-0.3	-0.3	0.25	-0.25	-0.2964
47	0	-0.2	-0.3	-0.4	-0.3	-0.3	0.25	-0.25	-0.0669
48	0	-0.2	0.3	0.4	0.3	0.4	0.35	-0.35	-0.2797
49	0	-0.2	0.3	0.4	0.3	0.4	0.35	-0.35	0.3398
50	0	-0.2	0.3	0.4	0.3	0.4	0.35	-0.35	-0.1593

51	0	-0.2	0.3	0.4	0.3	0.4	0.35	-0.35	0.1182
52	0	-0.2	0.3	0.4	0.3	0.4	0.35	-0.35	0.1789
53	0	-0.2	0.3	0.4	0.3	0.4	0.35	-0.35	-0.0581
54	0	-0.2	0.3	0.4	0.3	0.4	0.35	-0.35	-0.2257
55	0	0.2	0	0	0	0	0	0	-0.4347
56	0	0.2	0	0	0	0	0	0	0.2045
57	0	0.2	0	0	0	0	0	0	0.4007
58	0	0.2	0	0	0	0	0	0	-0.0925
59	0	0.2	0	0	0	0	0	0	-0.4868
60	0	0.2	0	0	0	0	0	0	-0.1742
61	0	0.2	0	0	0	0	0	0	-0.2058
62	0	0.2	0	0	0	0	0	0	-0.2956
63	0	0.2	0	0	0	0	0	0	0.3381
64	0	0.2	0	0	0	0	0	0	-0.0257
65	0	0.2	0	0	0	0	0	0	-0.2102
66	0	0.2	0	0	0	0	0	0	-0.0652
67	0	0.2	0	0	0	0	0	0	-0.0842
68	0	0.3	0	0	0	0	0	0	0.2341
69	0	0.3	0	0	0	0	0	0	0.4328
70	0	-0.3	0	0	0	0	0	0	0.1333
71	0	0.3	0	0	0	0	0	0	-0.1003
72	0	-0.4	-0.3	-0.4	-0.3	-0.2	0.1	-0.1	-0.1223
73	0	-0.4	0.3	0.4	0.3	0.2	0.2	-0.2	-0.3322

*Table B.1 – Characteristics of Strategies: Initial Agent Beliefs for Experimental Network Diagrams*

Experiment 1

Step #	TB Strategy	US Strategy	Absolute Payoff	Relative Payoff
1	70	2	0.0004	0.0004
2	22	2	0.0005	0.0005
3	22	70	0.0010	0.0010
4	2	70	-0.0027	-0.0027
5	2	22	-0.0001	-0.0001
6	70	22	-0.0021	-0.0021
7	70	2	0.0015	0.0015
8	22	2	-0.0009	-0.0009
9	22	37	0.0006	0.0006
10	2	37	-0.0031	-0.0031

Experiment 2

Step #	TB Strategy	US Strategy	Absolute Payoff	Relative Payoff
1	70	72	-0.0041	0.0055
2	68	72	-0.0085	0.0011
3	68	14	-0.0029	0.0067
4	21	14	-0.0073	0.0023
5	21	49	-0.0055	0.0041
6	56	49	-0.0080	0.0016
7	56	70	-0.0066	0.0030
8	57	70	-0.0128	-0.0032
9	57	13	-0.0096	0.0000
10	58	13	-0.0153	-0.0058

Experiment 3

Step #	TB Strategy	US Strategy	Absolute Payoff	Relative Payoff
1	70	39	-0.0124	0.0013
2	48	39	-0.0088	0.0049
3	48	21	0.0022	0.0159
4	49	21	-0.0026	0.0111
5	49	40	0.0041	0.0178
6	21	40	0.0050	0.0187
7	21	49	0.0106	0.0243
8	57	49	0.0109	0.0246
9	57	21	0.0181	0.0318
10	49	21	0.0159	0.0296

Experiment 4

Step #	TB Strategy	US Strategy	Absolute Payoff	Relative Payoff
1	70	39	-0.0161	0.0017
2	48	39	-0.0102	0.0076
3	48	21	0.0030	0.0208
4	49	21	-0.0014	0.0164
5	49	72	0.0066	0.0244
6	54	72	0.0098	0.0276
7	54	23	0.0193	0.0371
8	72	23	0.0221	0.0399
9	72	29	0.0320	0.0498
10	23	29	0.0321	0.0499

Experiment 5

Step #	TB Strategy	US Strategy	Absolute Payoff	Relative Payoff
1	70	31	-0.0065	0.0017
2	48	31	-0.0025	0.0058
3	48	39	0.0074	0.0156
4	49	39	0.0068	0.0150
5	49	40	0.0155	0.0237
6	39	40	0.0151	0.0233
7	39	23	0.0214	0.0296
8	12	23	0.0199	0.0281
9	12	39	0.0269	0.0351
10	13	39	0.0221	0.0304

Experiment 6

Step #	TB Strategy	US Strategy	Absolute Payoff	Relative Payoff
1	70	23	-0.0135	0.0015
2	48	23	-0.0174	-0.0023
3	48	21	-0.0041	0.0110
4	49	21	-0.0113	0.0038
5	49	48	-0.0062	0.0089
6	13	48	-0.0059	0.0091
7	13	49	0.0019	0.0169
8	12	49	-0.0011	0.0140
9	12	13	0.0065	0.0216
10	49	13	0.0048	0.0199

Experiment 7

Step #	TB Strategy	US Strategy	Absolute Payoff	Relative Payoff
1	70	55	0.2023	0.0003
2	48	55	0.2000	-0.0021
3	48	68	0.1952	-0.0068
4	49	68	0.1797	-0.0224
5	49	48	0.1738	-0.0282
6	21	48	0.1603	-0.0418
7	21	49	0.1601	-0.0419
8	48	49	0.1505	-0.0515
9	48	21	0.1483	-0.0537
10	49	21	0.1371	-0.0650

Experiment 8

Step #	TB Strategy	US Strategy	Absolute Payoff	Relative Payoff
1	70	70	-0.2024	-0.0003
2	69	70	-0.2058	-0.0037
3	69	21	-0.1928	0.0093
4	56	21	-0.1862	0.0158
5	56	49	-0.1734	0.0287
6	21	49	-0.1692	0.0328
7	21	48	-0.1612	0.0409
8	49	48	-0.1608	0.0413
9	49	21	-0.1540	0.0480
10	57	21	-0.1550	0.0470

Experiment 9

Step #	TB Strategy	US Strategy	Absolute Payoff	Relative Payoff
1	70	21	0.0197	0.0413
2	57	21	0.0167	0.0383
3	57	48	0.0251	0.0467
4	49	48	0.0217	0.0433
5	49	57	0.0269	0.0485
6	48	57	0.0223	0.0439
7	48	49	0.0265	0.0481
8	57	49	0.0226	0.0442
9	57	48	0.0261	0.0477
10	49	48	0.0215	0.0431

*Table B.2 – Characteristics of Strategies: Strategy Profile and Payoff Tables for Experimental Network Diagrams*

Experiment 1A

Step #	TB Strategy	US Strategy	Absolute Payoff	Relative Payoff
1	[69, 70]	[23, 1]	0.0223	0.0439
2	[57, 39]	[23, 1]	0.0138	0.0354
3	[57, 39]	[13, 48]	0.0232	0.0448
4	[49, 2]	[13, 48]	0.0142	0.0358
5	[49, 2]	[57, 39]	0.0166	0.0382
6	[48, 13]	[57, 39]	0.0098	0.0314
7	[48, 13]	[49,48]	0.0146	0.0362
8	[12, 57]	[49,48]	0.0095	0.0311
9	[12, 57]	[40, 72]	0.0144	0.0360
10	[49, 48]	[40, 72]	0.0091	0.0307

Experiment 2A

Step #	TB Strategy	US Strategy	Absolute Payoff	Relative Payoff
1	[69, 70]	[23, 73]	0.0014	0.0230
2	[39, 49]	[23, 73]	0.0023	0.0239
3	[39, 49]	[13, 1]	0.0083	0.0299
4	[57, 68]	[13, 1]	0.0060	0.0276
5	[57, 68]	[3, 21]	0.0102	0.0318
6	[31, 2]	[3, 21]	0.0074	0.0290
7	[31, 2]	[20, 72]	0.0097	0.0313
8	[12, 10]	[20, 72]	0.0068	0.0284
9	[12, 10]	[23, 31]	0.0088	0.0304
10	[29, 14]	[23, 31]	0.0062	0.0277
11	[29, 14]	[49, 40]	0.0083	0.0299
12	[13, 31]	[49, 40]	0.0060	0.0276
13	[13, 31]	[31, 29]	0.0076	0.0292
14	[49, 40]	[31, 29]	0.0060	0.0276
15	[49, 40]	[58, 67]	0.0079	0.0295
16	[31, 29]	[58, 67]	0.0056	0.0272
17	[31, 29]	[49, 57]	0.0076	0.0292
18	[23, 3]	[49, 57]	0.0060	0.0276
19	[23, 3]	[31, 48]	0.0088	0.0304
20	[57, 58]	[31, 48]	0.0075	0.0291
21	[57, 58]	[23, 54]	0.0103	0.0319
22	[31, 10]	[23, 54]	0.0090	0.0306
23	[31, 10]	[58, 68]	0.0112	0.0328
24	[23, 54]	[58, 68]	0.0088	0.0304
25	[23, 54]	[31, 10]	0.0110	0.0326
26	[68, 58]	[31, 10]	0.0094	0.0310
27	[68, 58]	[54, 23]	0.0120	0.0336
28	[31, 10]	[54, 23]	0.0100	0.0316
29	[31, 10]	[68, 49]	0.0120	0.0336
30	[23, 2]	[68, 49]	0.0100	0.0316

*Table B.3 –  $\mu$  Sensitivity: Strategy Profile and Payoff Tables for Experiments 1A and 2A*

## References

- [1] Kalev I. Sepp, "Best Practices in Counterinsurgency," *Military Review*, pp. 8-12, May-June 2005.
- [2] *Field Manual 3-24: Counterinsurgency*. Washington D.C.: Headquarters, Department of the Army, 2006.
- [3] William B. Caldwell, IV, and Steven M. Leonard, "Field Manual 3-07, Stability Operations: Upshifting the Engine of Change," *Military Review*, vol. 88, no. 4, Jul-Aug 2008.
- [4] Mao Tse-tung. *On Guerrilla Warfare*. London: Cassell, 1965.
- [5] *Field Manual 3-24.2: Tactics in Counterinsurgency*. Washington D.C.: Headquarters, Department of the Army, 2009.
- [6] Lane V. Packwood, "Popular Support as the Objective in Counterinsurgency: What Are We Really After?" *Military Review*, vol. 89, no. 3, May-June 2009.
- [7] D. Acemoglu A. Ozdaglar and A. ParandehGheibi, "Spread of (Mis)Information in Social Networks," *Games and Economic Behavior*, vol. 70, no. 2, pp. 194-227, Nov. 2010.
- [8] Benjamin Hung, "Optimization-Based Selection of Influential Agents in a Rural Afghan Social Network," Masters, Massachusetts Institute of Technology, 2010.
- [9] Elihu Katz and Paul Lazarsfeld, *Personal influence: the Part Played by People in the Flow of Mass Communications*. Glencoe, IL: Free Press, 1955.
- [10] Duncan J. Watts and Peter S. Dodds, "Influentials, Networks, and Public Opinion Formation," *Journal of Consumer Research*, vol. 34, Dec. 2007.
- [11] Robert P. Abelson, "Mathematical Models of the Distribution of Attitudes Under Controversy," in *Contributions to Mathematical Psychology*, New York: Holt, Rinehart, and Winston, Inc., 1964, pp. Ch 6, p 141-160.
- [12] Morris H. DeGroot, "Reaching a Consensus," *Journal of American Statistical Association*, vol. 69, no. 345, pp. 118-121, Mar. 1974.
- [13] Nicholas Howard, "Finding Optimal Strategies for Influencing Social Networks in Two Player Games," Masters, Massachusetts Institute of Technology, 2011.
- [14] Thomas Schelling, "Models of Segregation," *American Economic Review*, vol. 59, no. 2, pp. 488-493, May 1969.
- [15] Thomas Schelling, "Dynamic Models of Segregation," *Journal of Mathematical Sociology*, vol. 1, pp. 143-186, 1971.

- [16] Mark Granovetter, "Threshold Models of Collective Behavior," *American Journal of Sociology*, vol. 83, no. 6, pp. 1420-1443, May 1978.
- [17] David Kempe, Jon Kleinberg, and Éva Tardos, "Influential Nodes in a Diffusion Model for Social Networks," in *ICALP, LNCS 3580*, 2005.
- [18] Asuman Ozdaglar and Ercan Yildiz, "Optimal Stubborn Agent Placement Problem - Transient Case," *Unpublished Manuscript*, 2011.
- [19] A.-S. Zadran, *Socio-economic and legal-political processes in a Pashtun village, Southeastern Afghanistan*. Buffalo, NY: PhD Thesis, 1977.
- [20] Myron L. Braunstein, *Depth Perception Through Motion*. New York: Academic Press, 1976.
- [21] John W. Payne, "Heuristic Search Processes in Decision Making," *Advances in Consumer Research*, vol. 3, pp. 321-327, 1976.
- [22] David Kempe, Jon Kleinberg, and Éva Tardos, "Maximizing the Spread of Influence through a Social Network," in *KDD*, 2003.
- [23] Solomon E. Asch, "Opinions and Social Pressure," *Scientific American*, vol. 193, no. 5, pp. 31-35, Nov. 1955.
- [24] William Flavin, *Civil Military Operations: Afghanistan: Observations on Civil Military Operations During the First Year of Operation Enduring Freedom*. U.S. Army Peacekeeping and Stability Operations Institute, 2004.