SANDIA REPORT

SAND2014-0442
Unlimited Release
January 2014

A Cognitive and Economic Decision Theory for Examining Cyber Defense Strategies

Asmeret B. Bier

Prepared by
Sandia National Laboratories
Albuquerque, New Mexico 87185 and Livermore, California 94550

Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

Issued by Sandia National Laboratories, operated for the United States Department of Energy by Sandia Corporation.

Available to DOE and DOE contractors from
U.S. Department of Energy
Office of Scientific and Technical Information
P.O. Box 62
Oak Ridge, TN 37831

Telephone: (865) 576-8401
Facsimile: (865) 576-5728
E-Mail: reports@adonis.osti.gov
Online ordering: http://www.osti.gov/bridge

Available to the public from
U.S. Department of Commerce
National Technical Information Service
5285 Port Royal Rd.
Springfield, VA 22161

Telephone: (800) 553-6847
Facsimile: (703) 605-6900
E-Mail: orders@ntis.fedworld.gov
Online order: http://www.ntis.gov/help/ordermethods.asp?loc=7-4-0#online

SAND2014-0442
Unlimited Release
January 2014

A Cognitive and Economic Decision Theory for Examining Cyber Defense Strategies

Asmeret B. Bier
Cognitive Modeling Department
Sandia National Laboratories
P.O. Box 5800
Albuquerque, New Mexico 87185-MS1327

Abstract

Cyber attacks pose a major threat to modern organizations. Little is known about the social aspects of decision making among organizations that face cyber threats, nor do we have empirically-grounded models of the dynamics of cooperative behavior among vulnerable organizations. The effectiveness of cyber defense can likely be enhanced if information and resources are shared among organizations that face similar threats. Three models were created to begin to understand the cognitive and social aspects of cyber cooperation. The first simulated a cooperative cyber security program between two organizations. The second focused on a cyber security training program in which participants interact (and potentially cooperate) to solve problems. The third built upon the first two models and simulates cooperation between organizations in an information-sharing program.

CONTENTS

FIGURES

TABLES

NOMENCLATURE

BIA	Behavioral Influence Assessment
CSIRT	Cyber Security Incident Response Team
DOE	Department of Energy
ISAC	Information Sharing and Analysis Center
JC3	Joint Cybersecurity Coordination Center
TF-BIA	Tracer FIRE Behavioral Influence Assessment

1. INTRODUCTION

Cyber attacks pose a major threat to modern organizations. These attacks can have nefarious aims and serious consequences, including disruption of operations, espionage, identity theft, and attacks on critical infrastructure. Organizations must put substantial resources into protecting themselves and their customers, clients, and others against cyber attacks. Even with a substantial investment in cyber defense resources, however, the risk of harm from a cyber attack is significant for many organizations.

The effectiveness of cyber defense can likely be enhanced if programs are implemented that allow organizations that face similar cyber threats to share information and resources. The threats faced by different organizations may be similar or identical (figure 1), and much of the work done by cyber defenders at these organizations may be redundant (Hui et al. 2010). By sharing information about cyber attacks, effective defense strategies, and personnel with specific expertise, organizations may better protect themselves against cyber threats while maintaining or even reducing the resources dedicated to cyber security.

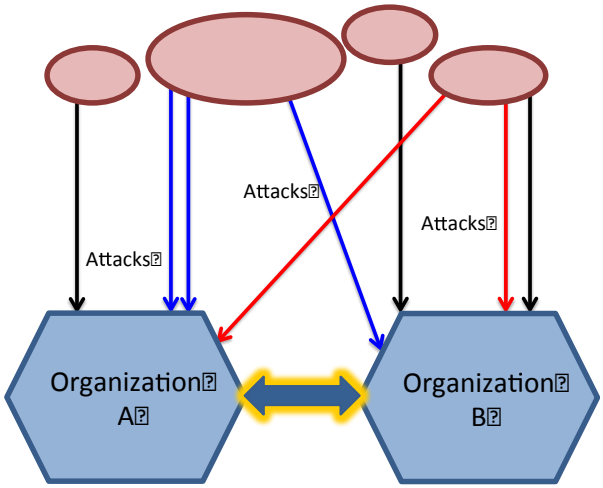

Figure 1: Cooperation can guard against attacks with similar traits or sources

Despite these potential benefits, cooperative cyber defense strategies are not common. Cyber defense teams must balance the potential benefits of cooperation against motivations not to cooperate. For example, if its vulnerabilities are made publicly known, an organization might become more susceptible to cyber attacks and might face damage to its reputation. Trust in cooperating organizations is therefore necessary for successful cooperative cyber security programs. Since organizations that are likely to cooperate with each other are those that face similar threats, they might also be in similar industries and have competitive relationships. Competition for customers, clients, or funding may raise concerns about motive and competitive advantage, making organizations less likely to trust each other. Finally, group inertia is a significant factor to overcome, and individual habits may be even more difficult to change than organizational strategy.

The potential for cooperation to improve defense and reduce resources may outweigh the obstacles. Three models were created to begin to understand the cognitive and social aspects of cyber cooperation. The first simulated a cooperative cyber security program between two organizations. This model provided insight into some of the potential dynamics that might be seen when organizations cooperate with each other. The model also brought to light a lack of data for model validation. To address this issue, a second model was built, focusing on a cyber security training program in which participants interact (and potentially cooperate) to solve problems. This training program provided substantial validation data, and lessons from the second model were used to build a third and final model. The third model simulates and information-sharing program between six organizations, and was used to understand how free riding behavior might impact the success of cooperative cyber security.

2. A TWO-ORGANIZATION MODEL OF COOPERATION FOR CYBER SECURITY

Model Overview

An organization must consider many different factors when making decisions about participation in a cooperative cyber security program. The risks and benefits of such a program must be weighed against each other, which is a difficult task when such programs are not widespread and potential outcomes are thus not readily apparent. A system dynamics model might be useful in understanding how the dynamics of such a program might unfold, which could help potential participants to understand the potential costs and benefits of cooperation.

This model depicts a simple system in which two organizations face similar cyber threats and are considering sharing their cyber defense resources. Each organization does some amount of cyber defense work that is redundant with work done by the other organization. In other words, there is some amount of cyber defense work that must be done separately for each organization, but the rest could be shared, rather than completed by each organization separately.

Figure 2 shows the basic feedback structure of the resource allocation decisions faced by the two organizations. Each organization has some amount of resources that it devotes to cyber security, and allocates those resources between two types of tasks. The first type of task is non-redundant, and must be done separately for each organization. The second type of task is redundant. Redundant tasks are those that can be done once, by either organization, and results of the tasks can be shared with the other organization to reduce workload. Each organization uses the fraction of tasks (both non-redundant and redundant) being completed to decide whether more or fewer resources should be allocated to cyber security. Each organization attempts to minimize the resources it allocates to cyber security while ensuring that the cyber tasks are completed to the maximum possible extent. This minimizes (but does not eliminate) the risk of a successful cyber attack, while maximizing the resources available for non-cyber-related organizational activities.

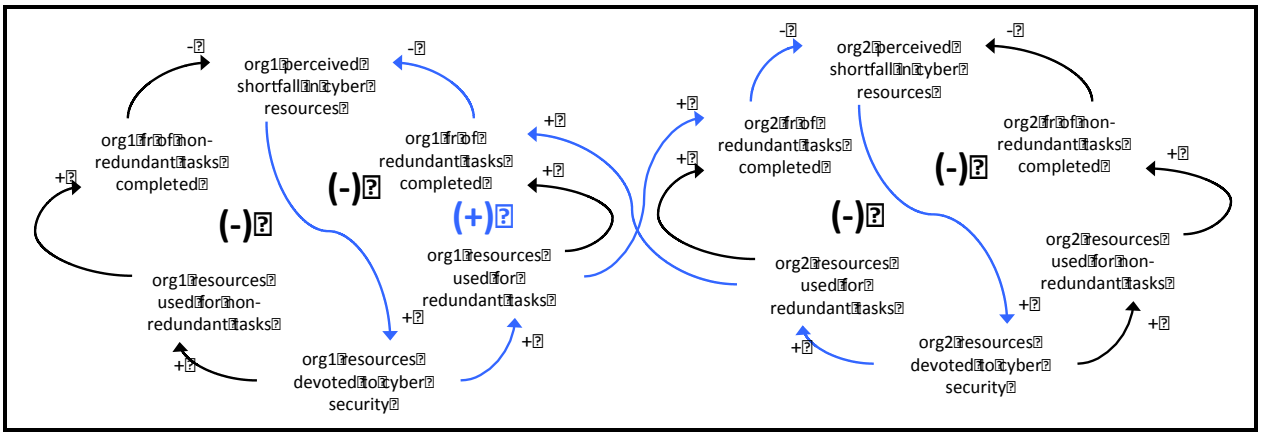

Figure 2: Feedback structure of resource sharing between organizations

A new feature of the causal structure is formed when cooperation becomes viable. In this case, resources allocated to cyber defense by one organization can augment the completion of redundant tasks for the other organization, allowing the second organization to reduce the resources it devotes to cyber security without losing effectiveness of cyber defense. If both organizations agree to cooperate to complete redundant tasks, both organizations may be able to devote fewer resources to cyber defense without sacrificing effectiveness.

The resource allocation structure shown in figure 2 addresses the potential benefits of cooperation in cyber security, which are weighed against risks to determine whether such a program should be established. Figure 3 shows the feedback structure of the decision-making process for a single organization. This portion of the model determines the strength of the cooperative agreement between the two organizations. The first feedback loop in figure 3, shown in blue, includes a simplification of the structure shown in figure 2. This loop represents how the benefits of cooperation, especially the increase in efficiency when resources are shared for redundant tasks, encourage an organization to strengthen its cooperative agreements. If benefits of cooperation have been realized in the past, then the organization is more likely to support cooperation in the future.

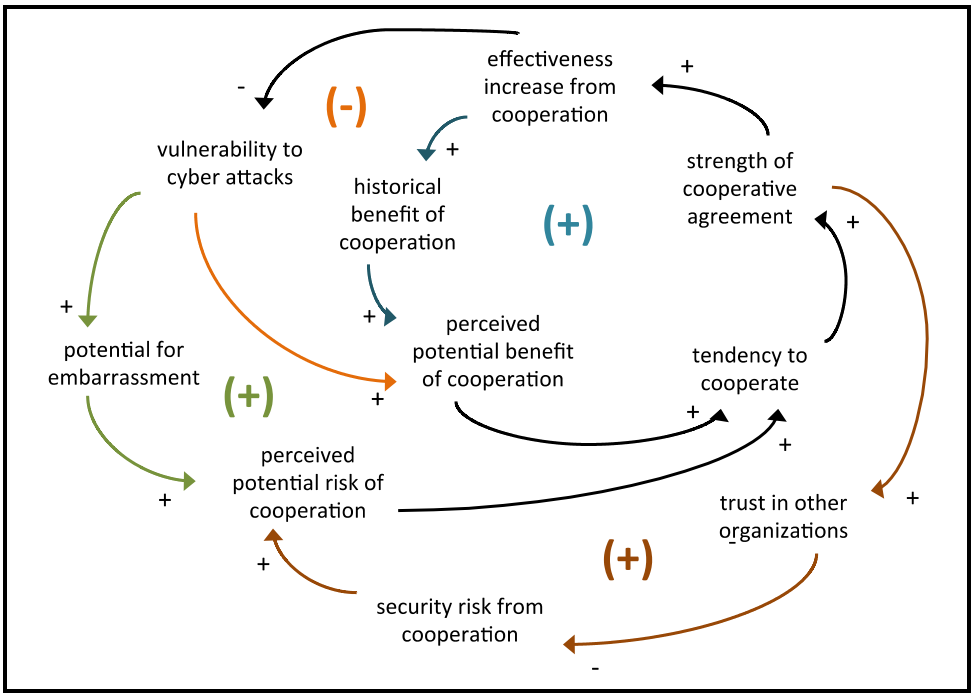

Figure 3: Feedback structure of decision making model for one organization

Three feedback loops might counteract the benefit loop. First (shown in orange), if the cyber security of an organization is strengthened then it may feel less vulnerable to cyber attacks. This would encourage the organization to reduce its support for a cooperative agreement, since perceived vulnerability encourages cooperation. There are also two feedback loops in this system that concern the risks involved in cooperation. The first (shown in green) addresses the potential for embarrassment if it becomes known that the organization is vulnerable to cyber attacks. This could mean lost business, reduced trust from customers, or lost reputation for security practices, any of which could cause serious damage to the organization. However, if cooperation improves security, the risk of embarrassment from cyber attacks decreases.

The other risk-based loop (shown in brown) addresses the possibility that cooperating organizations may not fully trust one other. Cooperative agreements may involve sharing sensitive information, such as details of organizational structure, vulnerabilities, and information about cyber attacks and strategies for counteracting those attacks. This information could be dangerous if used for the wrong purposes. Furthermore, organizations that are likely to cooperate with each other are those that face similar threats, and are thus likely to be in similar industries and perhaps have competitive relationships. Trust may be difficult to build in these situations. This model assumes that trust between organizations is stronger when cooperative agreements have existed and produced benefits over some period of time. If trust grows, organizations become more likely to promote cooperation.

The model described here uses the same decision making structure to represent each of the two organizations in the system (future work will include more detailed and varied structures). Each organization determines its desire to cooperate, and the two desires govern the

strength of the cooperative agreement. The strength of that agreement and the risks and benefits that it produces then support future decision-making processes for each organization.

Results

The model was used to simulate two scenarios, where the primary difference was the intensity of cyber attacks experienced by the two organizations. This intensity is an important driver of the system because it helps to determine the organizations' perceived vulnerabilities to cyber attacks. In the base case scenario, both organizations face similar threats, and the intensity of attacks faced by the two organizations is equal. The second scenario involves uneven threats; in this simulation organization 2 faces a substantially more intense threat than organization 1. This alters the risk/benefit calculations for the two organizations as described below, changing the organizations' desires to participate in a cooperative agreement.

Figure 4 shows the strength of the cooperative agreement that results from each scenario. The simulation begins with no cooperative agreement in place. In the similar threats (base) case, the strength of the agreement builds slowly over the first year and a half. This growth depends on both organizations having some baseline belief that cooperation is likely to help with the effectiveness of cyber defense. After the first year and a half, both organizations begin to see significant benefits resulting from the cooperative agreement. The perceived benefits of cooperation encourage more cooperation, and the strength of the cooperative agreement grows more quickly in the next few years before leveling off with a strong agreement.

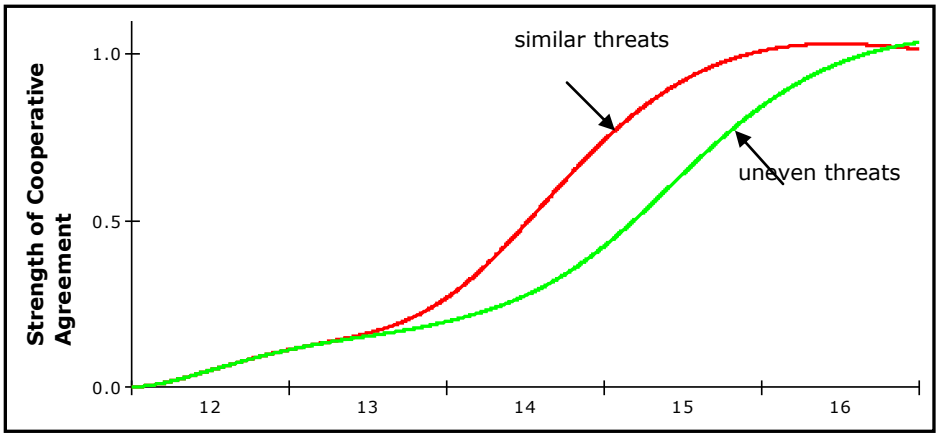

Figure 4: Strength of cooperative agreement for base case and uneven threat scenarios

The uneven threats case exhibits similar behavior to the base case at the beginning of the time horizon. For the first two years of the simulation, the cooperative agreement grows slowly based on a pre-existing belief that cooperation may help cyber defense. In the uneven threats case, the organization that faces a smaller cyber threat has less to gain from cooperation. This

organization is less enthusiastic about strengthening the cooperative agreement, and the agreement grows much more slowly than in the similar threats case.

The benefits of cooperation play a large role in decision-making, particularly in the later part of the simulations. These benefits result from the fact that cooperation allows organizations to achieve strong cyber defense while significantly reducing the resources they dedicate to cyber security. Figure 5 shows the resources dedicated to cyber security and used for cyber security by organization 1 for the similar threats (base) case. The results for organization 2 are identical. In this scenario, both organizations begin with a baseline level of cyber resources. As the cooperative agreement is strengthened, much of the redundant work is eliminated. This allows both organizations to achieve the same level of cyber security they would without cooperation, but at a reduced investment. Even though fewer resources are now allocated by organization 1 for cyber defense, more resources are actually used for the cyber defense of organization 1, because organization 2 contributes resources through the cooperative agreement. Since the tasks being eliminated are redundant, both organizations can reduce their investments in cyber defense resources, yet see more cyber defense work being done.

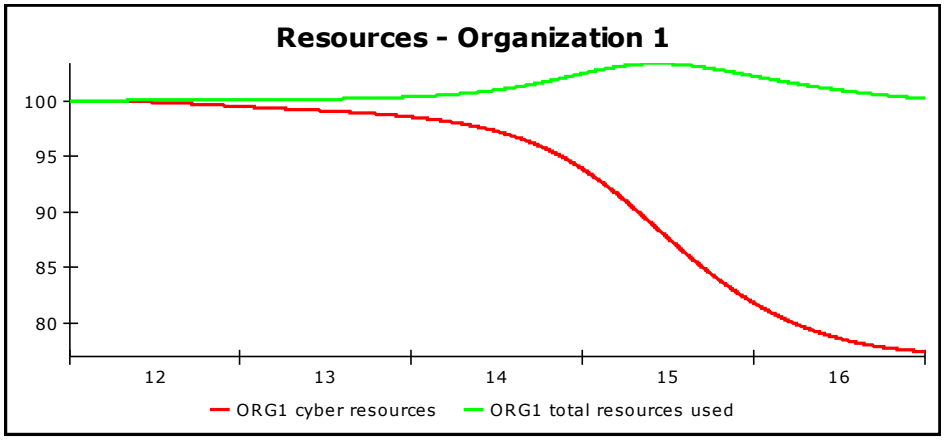

Figure 5: Resources contributed and used by one organization

When the risks faced by the two organizations are uneven, the risks and benefits of cooperation that each perceives (figure 6) also differ. In the uneven threats scenario, organization 2 faces a substantially more intense cyber threat than organization 1. Both organizations begin with low perceived benefits of cooperation; since no benefits of cooperation have yet been realized, these are based on a pre-existing belief that cooperation may be helpful. When benefits from cooperation do become apparent, organization 2 realizes that cooperation could provide a very large benefit. This perception also relies on the intensity of the cyber threat. Since organization 1 faces a less intense threat than organization 2, its perception of the potential benefits of cooperation is smaller. The intensity of the cyber threat also directly impacts each organization's perception of the potential risks involved in cooperation. Organization 2 sees a stronger threat, and thus considers itself more vulnerable and understands that the risks it faces

(from security or embarrassment) are quite large. Since it faces a less intense threat, organization 1 perceives a smaller risk of cooperation than organization 2.

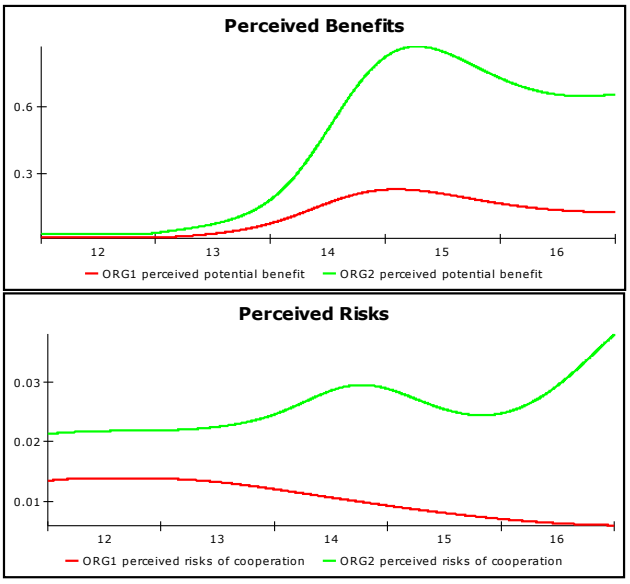

Figure 6: Perceived benefits and risks for each organization in the uneven threats case

For both organizations, the potential benefits of cooperation are substantially larger in magnitude than the risks. Organization 2 is therefore much more eager to strengthen the cooperative agreement than organization 1. Both parties must agree in order for the agreement to be strengthened, so diminished interest from organization 1 in the uneven threats scenario (as compared to the similar threats scenario) results in a weaker agreement.

Conclusions

This model indicates that in a simple system where redundant cyber security work can be reduced through cooperation, the benefits of a cooperative agreement can be substantial. Rather than duplicating work to detect, understand, and defend against cyber threats, energy can be deferred into more useful defensive strategies or other organizational goals. Stronger defense can be realized without increasing the resources dedicated to cyber security.

These results also suggest that cooperative cyber agreements are likely to work best when participating groups face threats at similar intensities. An organization that faces fewer threats is likely to be less interested in a cooperative agreement than an organization that faces many serious cyber threats. Differences in the intensity of threats to cooperating organizations could cause distrust and a high perceived risk of cooperation.

In the first few years of a program of cooperation, organizations are likely to participate minimally. They might declare support for a cooperative program, but substantial resources will likely not be contributed until the benefits of cooperation are apparent. The success of these programs is thus likely to depend on whether benefits are realized before the organizations

involved lose interest. Once benefits are apparent, participation will likely be influenced by the threats faced by each organization. The success of an agreement will depend on there being sufficient threat to make cooperation attractive. Full participation is also likely to depend on trust between the organizations; low-trust or competitive relationships will make a cooperative agreement less successful.

This model simulates the potential outcomes and decision-making processes involved in cooperative cyber security agreements designed to reduce redundant work. It is the first step in a project designed to understand the potential for organizational cooperation to improve cyber defense. A substantial amount of work remains to be done to understand this problem. Future adaptations of this model will incorporate cognitive models of the individuals and groups involved in decision-making about cooperation in cyber defense. The model will be used to explore likely outcomes of these systems when the organizations involved have different characteristics and tendencies. We will also explore cooperative agreements with more than two participating organizations. Validation data will be collected from cyber security training exercises, historical data, and subject matter experts. Further psychological and economic theory, including cognitive dissonance (Festinger 1957), the theory of planned behavior (Ajzen 1991), bounded rationality (Simon 1957), qualitative choice theory (McFadden 1982), and prospect theory (Tversky & Kahneman 1974) will be incorporated to enhance the decision-making model. Cooperative agreements in contexts other than redundant work will be analyzed, and potential program designs will be studied. We will also explore likely changes in attitudes toward these programs as they become widespread, including tipping points that affect whether an organization will be willing to participate. We hope that this work will lead to a better understanding of the decision-making processes involved in cooperative agreements between organizations for cyber security, and will contribute to successful design of these programs.

3. COOPERATION AND LEARNING IN CYBER SECURITY TRAINING EXERCISES

Introduction

Cyber attacks pose a major threat to modern organizations. The consequences of these attacks include disruption of operations, espionage, identity theft, and attacks on critical infrastructure. Organizations put substantial resources into protecting themselves and their customers against cyber attacks, but even with considerable investment in cyber defense resources the risk of harm from a cyber attack is significant for many organizations.

Sandia and Los Alamos National Laboratories, realizing the increasing threat from cyber attacks, created a training program called Tracer FIRE (Forensic and Incident Response Exercise) to increase the effectiveness of cyber security incident response teams (CSIRTs). Tracer FIRE combines traditional classroom and hands-on training with a competitive game forum. In the classroom portion, students cover incident response topics and are given hands-on training with tools commonly used by CSIRT personnel. In the game portion of the exercise, the students form teams and use these tools to solve a series of challenges based on real-world incidents. The challenges cover a variety of cyber defense topics, and the number of points awarded is based on the difficulty of the challenge. The size of the teams varies from 4-10 players, and an effort is made to ensure that each team has a balanced skill set, and that all teams have roughly the same skill level. Tracer FIRE has been used to train almost 1000 incident responders from DOE, US Government, critical infrastructure and academia. In fact, the most recent Tracer FIRE event was held online, and had hundreds of participants from over 10 countries around the world.

Tracer FIRE also presents an opportunity for human-focused research on cyber security and training. The exercise offers a controlled environment with a variety of challenges and an opportunity for data collection that does not often exist in traditional security environments. A variety of research projects have used Tracer FIRE to study individual and group characteristics in relation to effectiveness of cyber defense and training.

Tracer FIRE has begun to explore incorporating challenges that encourage cooperation between players. By cooperating with other organizations (sharing information about cyber attacks, effective defense strategies, and personnel with specific expertise), cyber defenders might increase the resources and information available for solving a particular cyber problem and thus better protect their organizations. Researchers have begun to explore the possibility of organizational cooperation in cyber defense (Hui et al. 2010; Sandhu et al. 2010; Luna-Reyes 2006; Ring and Van de Ven 1994; Oliver 1990; Luna-Reyes et al. 2008), and the Tracer FIRE team is exploring methods for enhancing cooperation both during and after the exercise. The current design of Tracer FIRE encourages cooperation within teams (points are rewarded by team) and does not prohibit cooperation between teams. Some teams do cooperate with each other to solve challenges, but the point structure, combined with a tendency toward a culture of

individualistic work in cyber security (Gates and Whalen 2004), does not always encourage high levels of cooperation.

This paper presents a model that was created to explore the potential for enhancing cooperation during Tracer FIRE. The model uses a decision-making framework based on psychological, social, and economic theory that was designed to dynamically simulate and allow exploration of cognition, including learning. The model was used to explore whether cooperation can improve learning in an exercise like Tracer FIRE, and how the characteristics of the exercise and of the participants and teams would likely affect the benefit (or cost) of cooperation. The model proved useful for understanding how the exercise might be tuned to encourage cooperation and enhance learning.

The Tracer FIRE Behavioral Influence Assessment (TF-BIA) Model

In order to study the dynamics of cooperation in Tracer FIRE, the Tracer FIRE Behavioral Influence Assessment (TF-BIA) model was created. The model was populated based on interviews with subject matter experts, who were past participants in the Tracer FIRE program and also cyber security professionals, and was calibrated using data collected during Tracer FIRE exercises. The model is based on the Behavioral Influence Assessment (BIA) framework, which was designed to model decision making using well-established psychological, social, and economic theories, all within a system dynamics structure.

Behavioral Influence Assessment (BIA)

Behavioral Influence Assessment (BIA) is a system dynamics-based modeling framework for simulating systems that involve human behavior and decision making. The theoretical framework of the BIA is based on well-established psychological, social, and economic theories that have been incorporated into a single structure (figure 7) that is both self-consistent and dynamic. BIA uses a hybrid cognitive-system dynamics architecture. Cognitive models are implemented using system dynamics and embedded into an encompassing system dynamics model, which simulates interactions between people, groups, and physical, economic, or other system components.

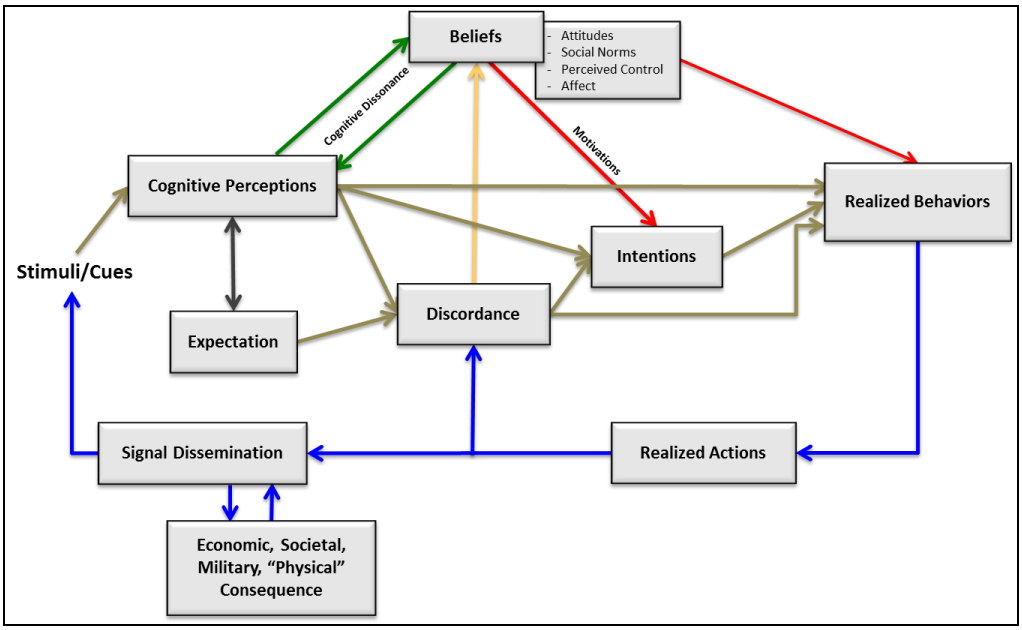

Figure 7: Computational structure of the BIA framework

The cognitive portion of the BIA begins with individuals or groups being exposed to cues (stimuli relevant to the decision-maker). These cues are processed to create cognitive perceptions, the decision-maker's assessment of the world or situation. Over time, cognitive perceptions become expectations, which are compared to cognitive perceptions to determine discordance with the current situation. Discordance and cognitive perception affect beliefs, a category of cognitive processes that includes the components of the theory of planned behavior (attitudes, social norms, perceived behavioral control) (Ajzen 1991) and affect. Intentions are calculated using utility functions. A multinomial logit function (McFadden 1982) compares intentions to determine realized behaviors, and over time those behaviors become physical realized actions.

One of these cognitive models is populated for each individual or group being included in the system. These cognitive models are connected to each other and to a world model sector using system dynamics. The world model sector includes all of the non-cognitive components of the system of interest, including physical systems, economics, etc. Outputs from the world model and the cognitive models act as inputs, or stimuli, for the cognitive model in subsequent time steps. Theoretical and mathematical details of the BIA are discussed by Backus et al. (2010).

Tracer FIRE BIA (TF-BIA)

The Tracer FIRE BIA (TF-BIA) model uses the BIA framework to simulate behaviors of participants in Tracer FIRE. The model simulates six teams, each with the same basic cognitive structure (cognitive parameters can vary between teams). Each team determines the amount of effort it spends working individually versus working cooperatively with other teams.

Considering the difficulty of the remaining challenges, individual and cooperative progress are calculated. Cooperative progress also takes into account the amount of work required to cooperate with other teams and shared knowledge available through cooperation. Shared knowledge available depends on the amount of knowledge that each team has and the effort that each team puts toward cooperation.

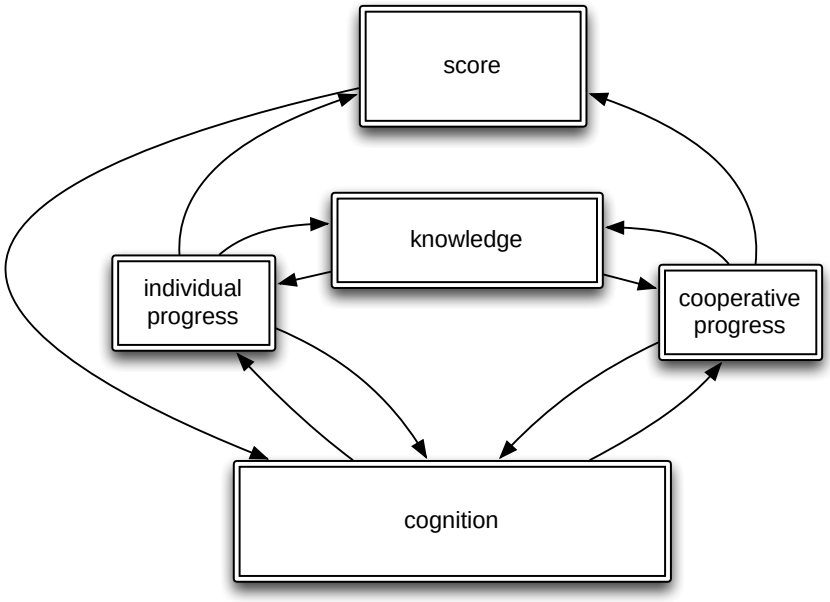

Figure 8: Model structure overview

Individual and cooperative progress for each team are combined to determine the increase in overall score. As teams solve more challenges, remaining challenges become more difficult. Increase in score and challenge difficulty are used as indicators to determine learning for each team. As knowledge increases, teams become more efficient at solving problems and have more to contribute to cooperative efforts if they choose to do so.

Both behavioral and non-behavioral portions of the model feed into the cognitive models as cues. Interviews with subject matter experts (SMEs) were held to determine how decisions are made during Tracer FIRE. The SMEs were previous participants in the exercise and also work as cyber security professionals. These interviews were used to determine the structure of the decision process (which cues and perceptions are considered, how cues determine perceptions, etc.) and to understand the relative importance of each input for model parameterization. The cues and cognitive perceptions that feed into each potential behavior are shown in table 1.

Table 1: Cues, cognitive perceptions, and potential behaviors

	potential behaviors ->	Work Individually			Work Cooperatively	
	cognitive perceptions ->	Competition	Benefit of indiv. work	Time pressure	Benefit of cooperation	Frustration
	effect on behavior ->	+	+	+	+	+
cues	Score difference from nearest competitor	-				
	Team rank	+				
	Recent individual progress		+			
	Recent cooperative progress				+	
	Recent total progress					-
	Difficulty of remaining tasks					+
	Time remaining in game			-		

Each team determines how much effort it puts into individual versus cooperative work. Teams tend to increase individual work when they feel time pressure or competition (based on team rank and having competitors close in score), or when individual work has increased the team's score in the recent past. They tend to work cooperatively when they are frustrated (due to lack of progress or high task difficulty), or when cooperation has recently produced benefits. These factors are compared to determine the effort that goes toward each type of work (individual and cooperative), which then affects score and knowledge, as described above.

Results

A key goal of Tracer FIRE participants is to win the game (by generating a higher score than any other team), but the primary goal of Tracer FIRE is to increase participants' knowledge about cyber security incident handling. Cooperation allows teams to learn from others, but requires effort and may give competitors an advantage. Teams must decide how much effort to put into cooperation versus individual work, and this decision affects both learning and scores.

There are four adjustable inputs in the TF-BIA model. The first two, initial knowledge (for each team) and baseline cooperation (for each team) are characteristics of the teams but can be altered by the Tracer FIRE designers. In the simulations discussed here, we assume that all teams have the same initial knowledge and baseline cooperation unless otherwise indicated. The other two variables of interest can be directly manipulated by the white cell (the people running Tracer FIRE). The white cell can modify the difficulty of the challenges, which is represented in the model by a maximum task difficulty variable. It can also make it easier or more difficult for teams to cooperate with each other. This might involve changes to communication infrastructure (instant messaging, shared message boards, etc.), locating players in the same room, challenges

that encourage cooperation between teams, verbal encouragement to cooperate from the white cell, or other strategies.

The base case simulation is shown in figure 9. In the base case, each team begins with 25% of the knowledge necessary to complete all of the Tracer FIRE challenges. Work required to cooperate is 25% (in other words, only 75% of the effort put into cooperation actually goes toward progress in the exercises). Challenge difficulty is .75 (of a maximum of 1), and each team begins the exercises with a baseline 25% of effort going toward cooperation. The teams end up with about 78% of the maximum score and about 52% of the total knowledge that can be gained from the exercises, doubling their knowledge over the course of the exercise. Cooperative effort starts out at 25% (the baseline), but declines after the beginning of the exercise. Since all the teams have similar, relatively low levels of initial knowledge, not much can be gained from cooperation and teams put more focus into individual work. Competition remains stable in this scenario because the teams' scores are equal. Near the middle of the time horizon, learning and frustration encourage more cooperation. All teams are gaining knowledge, so the potential benefit of cooperation is increasing. The remaining challenges are getting harder (teams tend to solve the easiest challenges first), so frustration is also increasing. At the end of the exercises, time pressure causes teams to focus more on individual work.

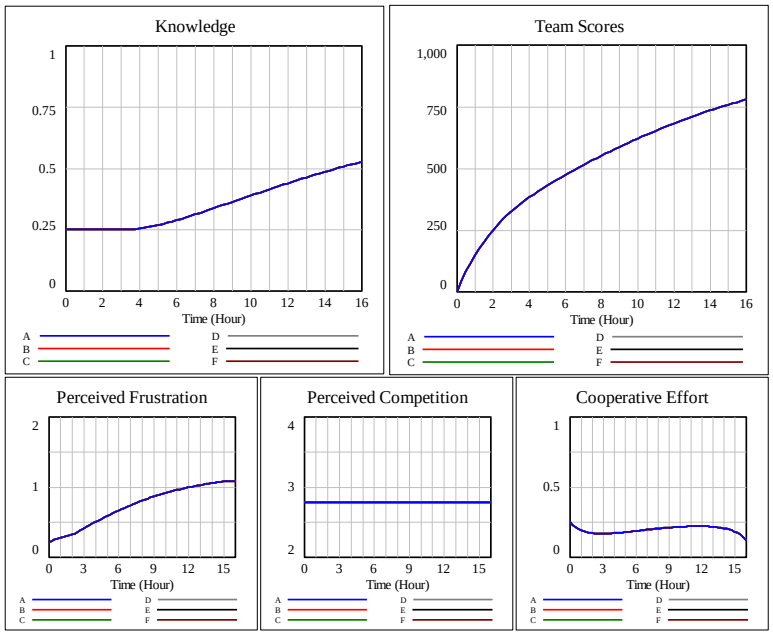

Figure 9: Base case simulation (init knowledge = 0.25, baseline cooperation = 0.25)

Figures 10 and 11 show scenarios where teams have a higher baseline rate of cooperation (50%) than in the base case (25%). This could represent a situation where teams or participants were chosen specifically for characteristics (personality traits, familiarity with other players, etc.) that encourage cooperation. It could also represent an exercise where teams are encouraged to cooperate before the game starts, or where challenges are designed to encourage cooperation between teams. Both scenarios show that learning increases from the base case. The final

knowledge variable for each team nears 66% when baseline cooperation increases to 50% (figure 10), and if barriers to cooperation are removed to make work required to cooperate 5% (rather than 25%), knowledge reaches 70% (figure 11).

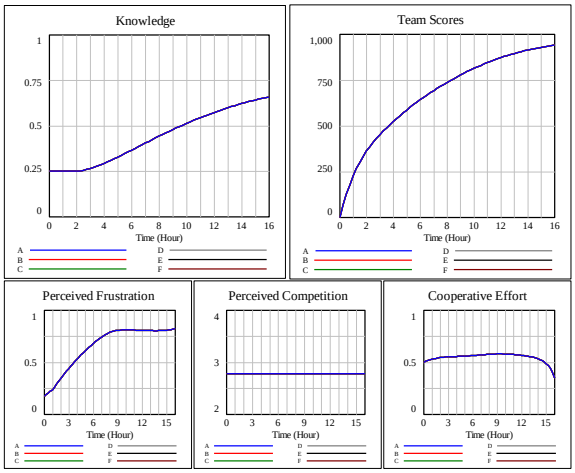

Figure 10: Baseline cooperation = 50%; work required to cooperate = 25%

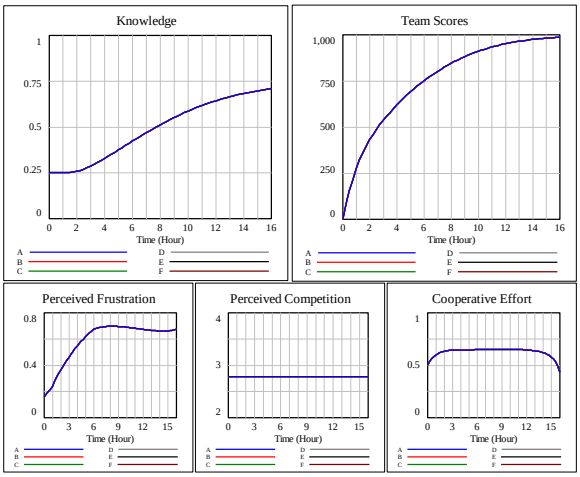

Figure 11: Baseline cooperation = 50%; work required to cooperate = 5%

Learning can be further improved by increasing the difficulty of tasks, as in the scenario shown in figure 12. This scenario is the same as the one shown in figure 10, except that the task difficulty is at its maximum. Participants learn more with higher task difficulty in this scenario, but frustration is also higher. This could cause participants to reduce overall effort levels or to dislike the Tracer FIRE program, discouraging their colleagues from participating in the future. While this model does not consider distraction or future participation in the program, it is a consideration for exercise design and implementation.

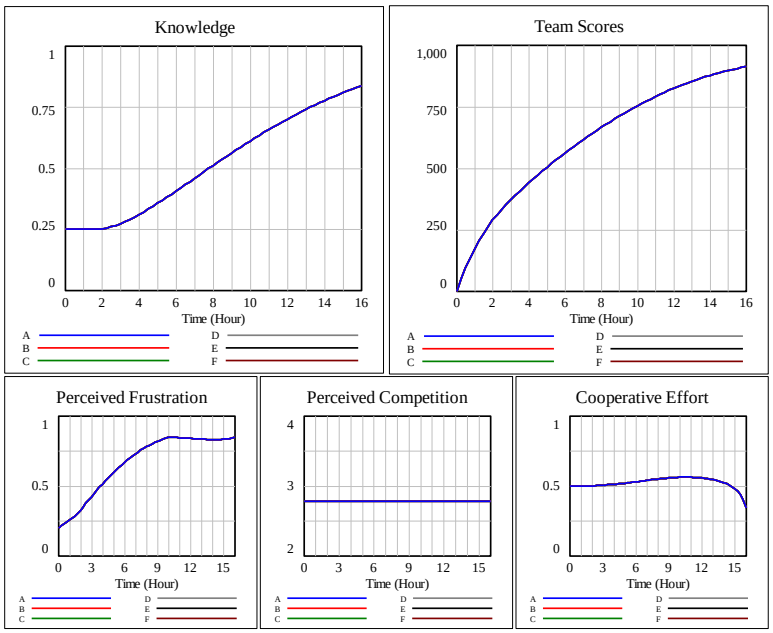

Figure 12: Baseline cooperation = 50%, task difficulty = 1

It is also likely that different teams will have different baseline cooperation levels. Figures 13 shows a scenario in which five teams have baseline cooperation of 25% and one team has a higher level of baseline cooperation (50%). Learning and score both increase a small amount for the team that cooperates more than the others. Figure 14 shows a scenario in which three of the six teams have the higher (50%) baseline level of cooperation. Because more teams are more willing to cooperate, the pool of shared knowledge increases and these teams see an even higher increase in score and knowledge than the others. These scenarios assume that work required to cooperate is the same as in the base case. As barriers to cooperation increase, benefits of cooperation will decrease, at some point (around 50% work required for cooperation in this scenario) creating a negative incentive to cooperate.

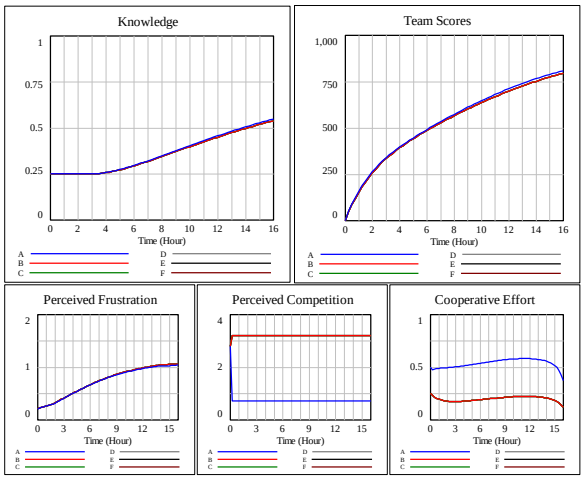

Figure 13: One team with baseline cooperation = 50%

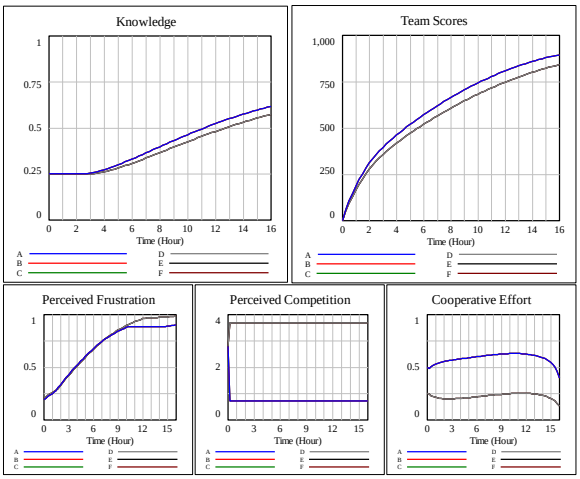

Figure 14: Half of teams with baseline cooperation = 50%

The goal of Tracer FIRE is to increase the participants' knowledge about cyber security incident response. Sensitivity analysis was conducted to indicate which of the four adjustable inputs to this model were most important in determining the teams' average knowledge at the end of the simulation. Partial correlation coefficients are shown in table 2. All of the inputs have high correlation with the knowledge output with high confidence. The maximum task difficulty has the highest (negative) correlation, but the others are also important.

Table 2: Partial correlation coefficients for average knowledge at end of simulation

variable	partial correlation coefficient	p-value
Maximum task difficulty	-0.93516	7.8392e-90
Work required to cooperate	-0.92539	4.2709e-84
Average initial knowledge	0.81709	1.5894e-48
Average baseline cooperation	0.75821	4.5148e-38

Conclusions

The model described above represents learning and cooperation in the cyber security training program Tracer FIRE, using a scenario in which six teams compete against each other for points. The model was used to indicate how the exercises might be designed to best improve participants' knowledge of the subject area. The four inputs to the model that are adjustable by the white cell are maximum task difficulty, work required to cooperate, initial knowledge, and baseline cooperation. All of these proved to be highly correlated with learning.

These results suggest various strategies that the white cell might try to improve learning during Tracer FIRE. They might make challenges more difficult (but not so much that frustration causes participants to dislike the exercise, which we plan to explore in future implementations of

this model). They might also remove barriers to cooperation by improving communication infrastructure, locating participants in the same room, verbally encouraging cooperation, incorporating challenges that require cooperation, or other methods. They might increase the initial knowledge of participants by including more classroom-style lessons before the exercise begins. Finally, they might increase participants' baseline levels of cooperation. This could be accomplished based on personality types of participants, composition of teams, familiarity of players with each other, structure of the game, or other strategies.

The Behavioral Influence Assessment (BIA) framework proved useful for modeling this problem. Because the framework includes an explicit cognitive model, we can use the model to understand intermediate phases in participants' decision-making process, such as cognitive perceptions, affect, and motivations. This might be more useful for understanding problems like learning than the decision rule method most common in system dynamics models. The BIA framework shows promise for modeling human behavior, especially in situations where details of cognition may be important.

This model was useful for indicating factors that could increase learning during Tracer FIRE, but there are aspects of the model that should be improved in future phases of this project. We would like to incorporate an extra behavioral variable that allows participants to take breaks from working during Tracer FIRE, which would allow assessment of frustration versus progress. Incorporation of the types of challenges and knowledge that would be useful for solving them would be also be useful. Finally, we would like to understand how other characteristics of an exercise, such as the number of teams, number and expertise of participants on each team, and challenge design might affect the success of Tracer FIRE.

4. COOPERATION AND FREE RIDING IN CYBER SECURITY INFORMATION SHARING PROGRAMS

Organizational Cooperation in Cyber Security

Cyber attacks pose a major threat to modern organizations. These attacks can have nefarious aims and serious consequences, including disruption of operations, espionage, identity theft, and attacks on critical infrastructure. The ubiquity of interconnected machines and advances in hacking techniques lead organizations to allocate substantial resources to protecting themselves and their customers, clients, and others against cyber attacks. Even with substantial investment in cyber defense resources the risk of harm from a cyber attack is significant for many organizations.

The effectiveness of cyber defense can likely be enhanced through programs that allow organizations that face similar cyber threats to share information about vulnerabilities, attacks, and defense strategies (ENISA 2010, MITRE Corporation 2012). Threats faced by different organizations are often similar, and much of the cyber defense work done may be redundant (Hui et al. 2010). Sharing information might allow organizations to better protect themselves while maintaining or even reducing the resources they dedicate to cyber security.

Despite the potential benefits of sharing information, cooperative cyber defense programs are not widespread. Cyber defense teams must balance the potential benefits of cooperation against motivations not to cooperate. For example, if an organization's vulnerabilities are leaked, that organization might become more susceptible to cyber attacks and face damage to its reputation. Trust in other organizations is therefore necessary for successful cooperative cyber security programs. Since organizations that are likely to cooperate with each other are those that face similar threats, they might also have competitive relationships. Competition for customers, clients, or funding may raise concerns about motive and competitive advantage, making organizations less likely to trust each other. Finally, group inertia is a significant factor to overcome, and both individual habits and organizational strategy may need to change to establish a successful program.

Increased recognition of the potential benefits of information sharing has led to various programs and proposals for cooperative cyber defense programs. A presidential executive order in the United States (The White House, 2013) lays the framework to create policy to increase the security and resilience of the nation's critical infrastructures. A major component of the U.S. strategy is increased communication, including information sharing between the public and private sectors (Raduege, 2013). This aspect of cyber security regulation has proven controversial, given the potential for privacy breaches (The Economist, 2013).

The United States Department of Energy (DOE) recently created the Joint Cybersecurity Coordination Center (JC3), and requires DOE-related entities to report cyber security incidents to the JC3 (US DOE 2013). Information Sharing and Analysis Center (ISAC) and Information Exchange (IE) models (ENISA 2010; ISAC Council 2004; The MITRE Corporation 2012) have been used in various critical infrastructure sectors in the U.S. and Europe, including financial services, electricity, public transportation, and health care, to allow entities within these sectors

to share information about cyber and other threats to critical infrastructure. The ISACs have had varied but limited success, due to hesitancies about sharing sensitive information and lags in data sharing when compared to direct relationships between organizations (MITRE Corporation 2012). More recently, the U.S. House of Representatives passed the Cyber Intelligence Sharing and Protection Act (CISPA). The legislation would have allowed and encouraged the U.S. government to share information about cyber threats with the private sector, but CISPA was not passed by the U.S. Senate and did not become law. The European Network and Information Security Agency published a document outlining that the key to security is cooperation across citizens, industry, and government (ENISA, 2010), and the European Commission is in the process of designing cyber security legislation with an information sharing component (The Economist, 2013).

The potential for cooperation to improve defense and reduce costs may outweigh the obstacles and potential pitfalls for collaborating organizations. This work is a first step toward understanding the social and operational issues involved in implementing a program of cooperative cyber defense between organizations. The model described here simulates an information sharing program involving six generic organizations. The model describes the social and organizational dimensions of a potential cooperative relationship, focusing on decisions about whether and how much an organization should participate in cooperative behaviors. The simulations described here are intended to assess how free riding behavior might affect participation in an information-sharing program.

Previous Research

Most research into cooperation in cyber security has focused on the technical issues involved and the effectiveness of tools and emerging technologies to combat cyber criminals (Hui et al. 2010), as well as potential program designs in cyber (Sandhu et al. 2010, Krishnan et al. 2011) and other information sharing (Luna-Reyes 2006) applications. Human factors in cyber security, including public awareness, advisory profiling, and cognitive factors of the cyber defender have begun to be considered (Forsythe et al., 2012). Several education and training initiatives have incorporated team collaboration to help teach the skills of information sharing and team work in the cyber domain (Reed et al., 2013; Hill et al., 2001). However, these programs are limited to the university level or weeklong courses in virtual training environments. The collaborative skill transfer has yet to be documented, and future research might assess whether those who participate continue to share information with teammates, coworkers, and other organizations they competed with after the game has finished.

Information sharing between organizations for cyber security has been explored in the economics literature, using game theory and non-dynamic modeling, as well as descriptive methods. Anderson and Moore (2006) give an overview of economic issues in cyber security. Gordon et al. (2003) examined the relationship between information sharing and security and the incentive structure that affects that relationship. Hausken (2007) used game theory to investigate information sharing and security investment tendencies between two firms. Gal-Or and Ghose (2004) used game theory to assess how sharing information and investment strategies might

affect competition. Liu et al. (2013) used game theory to assess free riding behavior in FS-ISAC, the information sharing program for the financial services sector in the United States. Rich et al. (2006) outlined some of the factors that are likely to influence information sharing programs and built a model of a basic system.

System dynamics and systems thinking has also been used to a small extent to look at cooperation and information sharing between organizations. Luna-Reyes and Andersen (2007) presented causal maps of a theory of interagency collaboration for information system development, and Luna-Reyes et al. modeled information sharing for requirements analysis in information technology projects. Dutta and Roy (2006) built and discussed the most substantial work on cyber security in the system dynamics literature to date, a model looking at compliance with information security protocols, as well as investment in training and technology. The system dynamics community has studied a limited number of cyber security topics, including insider threat (Rich et al. 2005, Moore et al. 2006, Martinez-Moyano et al. 2006), system vulnerability (Radianti et al. 2009, Goldsmith and Siegel 2010, Goldsmith and Siegel 2012), general theories of human factors and security (Gonzalez and Sawicka 2003, Hillen et al. 2006)

We created a model that uses system dynamics to simulate organizational decisions about whether and how much to participation in cooperative cyber programs. The model allows explicit representation of decision-making strategies, and enables exploration of different scenarios. Cooperative relationships between organizations have been examined (Ring and Van de Ven 1994; Oliver 1990; Luna-Reyes et al. 2008), but these relationships may be substantially different when their purpose is cyber security rather than for commercial purposes. This work looks at social and organizational aspects that will likely play a major role in cooperative dynamics but have not been sufficiently analyzed in a dynamic capacity.

Information Sharing Model: Conceptual Design

We created a model that simulates information sharing between multiple organizations. The model focuses on how organizations make decisions about whether and how much to participate in the program. In each organization, management and cyber defense staff both weight the risks and benefits of participation, and their desires to contribute determine the organizations' participation rates.

Figure 15 shows the basic feedback structure of the information sharing. The causal loop diagram is shown at the top of figure 15, and the small diagrams (a-g) highlight the important feedback loops in the system. The red variables in the causal loop diagram are the cyber threat and information sharing aspects of the system. The overall cyber threat will influence the threats faced by any particular organization. Some fraction of these threats will be new to the organization, while others will be known by the organization through experience, research, or information sharing. When threats are known, there is some background on how to deal with them and the time needed to process the threats is smaller. The number of cyber security staff in

an organization will increase as the number of hours needed for cyber defense increases, but there is a delay involved in the hiring process and a delay for training new cyber defense staff.

Two variables determine the effort that any organization puts into contributing to the information sharing program: management support for information sharing and cyber defense staff support for information sharing. Inputs to decisions made by cyber defense staff are shown in blue in figure 15, and highlighted in loops a, b, c, and d. As information sharing is utilized, cyber defenders will tend to see more benefit from previous participation in the program, and will be more likely to contribute information (figure 15a). However, they are also concerned about security risk from information sharing (figure 15b). When information is shared with another organization, it can reveal vulnerabilities, defense strategies, or other information that might increase vulnerability. The level of trust in other participating organizations is important in determining willingness to participate. Baseline levels of trust between organizations are considered, as are any breaches of trust that might occur outside of this system. For example, if an organization has a substantial cyber security problem, or if they have neglected to share useful or important information in the past, trust may be damaged. Trust can also be affected by perceived free riding in the information sharing program. If a particular organization seems to utilize shared information but does not contribute as expected, other organizations may lose trust in them. Free riding can also directly affect cyber defenders' desire to share information (figure 15c), since effort put into the program may not be worth the benefits received if other organizations are not sharing information at similar levels. Finally, as the workload for cyber defenders increases they are more likely to support information sharing, since having knowledge about a threat decreases the work required to fix it (figure 15d).

The green variables in figure 15 represent inputs to management's decision about whether and how much to support an information sharing program. The cost savings from the information sharing program (loops e and f in figure 15) is a key driver of this decision, and is affected by the cost of cyber security and the amount of information that has been shared. Management also considers the perceived potential cost of the cyber threat (figure 15g). This cost takes into account direct costs (loss of productivity, loss of intellectual property, theft, etc.) as well as indirect costs (loss of reputation, embarrassment, loss of future clients or customers, etc.) that may occur if a breach of cyber security occurs. As these potential costs increase or as the perceived likelihood of a breach of cyber security increases, management will tend to support information sharing more due to its potential to improve cyber security within the organization.

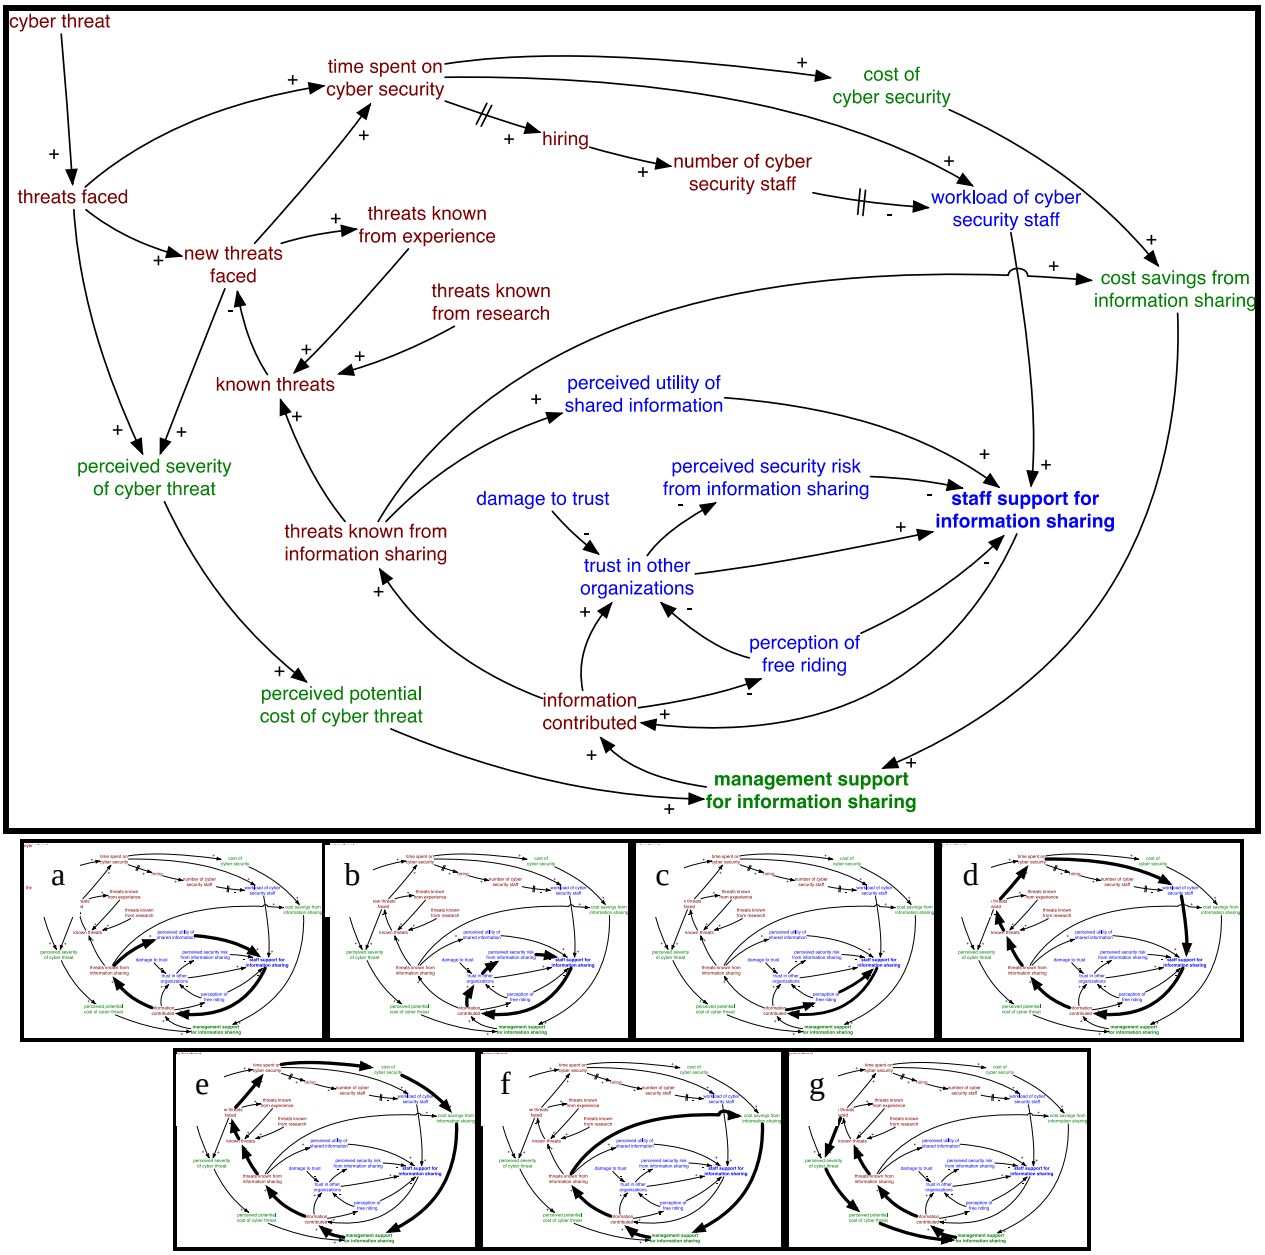

Figure 15: Feedback structure of the organizational cooperation model; important feedback loops are highlighted in a-g

Information Sharing Model: Model Formulation

The information sharing model contains three main sectors: Threats and Information, Management Decisions, and Cyber Defender Decisions. In the Threats and Information sector , a stock of existing threats represents potential threats to the simulated organizations that have been created but not necessarily seen by any of the organizations. This stock grows linearly over time, at a rate that approximately matches data on new cyber attacks. Some fraction of new threats and existing threats are faced by each organization, and existing threats faced might be known by the

organization or might not be known. Unknown existing and new threats faced by an organization must be processed and dealt with, but after this process they are added to the organizations' pool of known threats. This distinction is important, because if information about threats is known, the time needed to address those threats decreases. Threats can also be known through research or through the information sharing program.

As an organization learns about threats, it has the option to share that information with other organizations. Information contributed to the sharing program is determined by the fraction of information that is considered sharable (an exogenous variable in this model, which accounts for classified or secret information, information not considered of use to other organizations, etc.) and the effort put into the information sharing by the organization. This effort is determined by multiplying management support for information sharing by cyber defenders' desire to share information, both determined in separate model sectors. In this way, each organization determines how much information to contribute to the sharing program. Since some of the information contributed by an organization is likely to already be included in the pool of shared information, we use a modified version of the equation for calculating the probability of a union. If we assume that any given bit of information has equal probability of being in the shared information pool or in the information being shared by an organization, we can say that the probability that shared information will be new is the same as the fraction of information that has not yet been shared. In a similar way, we calculate the fraction of threats known by each organization through information sharing, and add that to the known threats for each organization.

The other two model sectors describe how management and cyber defenders in each organization make decisions. Each of these sectors depends on array variables, with each organization in the model represented separately. The Management Decisions sector determines management support for sharing information as well as changes to the cyber defense staff. Management's perceived threat is based on the known and unknown threats faced, how severe they perceive each type of threat to be, and the perceived likelihood of a successful attack. The severity and perceived likelihood values are exogenous and based on conversations with subject matter experts. Perceived potential costs of direct and indirect threats are determined and feed into the managements' perceived utility of information sharing.

The Management Decisions section also depends heavily on cost calculations. Hours spend on known and unknown threats are determined based on average values and multiplied by the cost per hour of cyber defense work to determine the cost of dealing with cyber threats. By assessing the threats known from information sharing, the cost savings from information sharing is also assessed, and feeds into managements' perceived utility of information sharing. The delayed cost of dealing with threats also determines the cyber security budget, which determines the desired number of cyber defenders in each organization. Cyber defenders can work entirely on the cyber security job, or can work partially in cyber security while they also spend time elsewhere in the organization. A delay is present in the staffing process, and new staff spend time in training. While in training, new staff members are able to deal with threats at a slower rate

than experienced staff, and experienced staff must spend time training the staff members in training. There is also a baseline loss of cyber defenders, as people move on to different organizations or different jobs within the same organization. Functional cyber defense staff is determined by the number of cyber defenders who are both experienced and in training, how efficient each group is, and how much time goes into training newcomers. Hours spent dealing with threats is divided by the functional cyber defense staff to calculate the workload of cyber defenders.

The final sector of the model, Staff Decisions, determines cyber defense staffs' desire to participate in the information sharing program. Four main variables contribute to the staffs' perceived utility of information sharing. First, as the workload of cyber defenders increases, staff are more likely to want to share information, since doing so can reduce their workload. As the information sharing program progresses and staff have seen a benefit from participating in the program, their perceived utility of shared information can increase. If staff perceive a high level of free-riding in the program, they are less willing to participate. Free riding is determined by perceived discrepancies in the amount of information contributed by each organization. The duration of participation by each organization can increase trust between participants, which is also affected by the initial strength of relationships between organizations and by any damage to trust. As trust increases, the perceived potential security risk from sharing information about cyber attacks will decreases, improving cyber defenders' desires to share information.

Results and Analysis

For this study, the model was used to assess how free riding behavior can affect information sharing in a cooperative setting. In the base case scenario, all six organizations have the same initial affinity toward sharing information. Cyber defenders in each organization had a baseline desire to participate of 0.9 (where 1 is the maximum), the level of desire that our subject matter experts suggested as realistic. The results of the base case simulation are shown in figure 16. Since all of the organizations have the same behavior in this scenario, when appropriate, only results for organization A are shown. Cyber defenders start with a .9 desire to participate in the information sharing program, and management supports participation at a 0.5 level. When the information sharing program begins after week 2, each organization begins to share information based on the product of these desires (figure 16a). Total shared information (figure 16b) grows steadily over the time horizon. As information sharing grows, the fraction of known threats that are known because of information sharing (figure 16c) also grows steadily.

While information sharing cuts the average time needed to defend against a cyber attack, the cyber threat still grows, and organizations are hit with more attacks as the time horizon progresses. Cyber staff retire or move on to new jobs (the average time in the job is five years in this model), and training must occur before new staff are fully up to speed. Cyber defense staffing dynamics, including staff in training, fully trained staff, and functional staff (taking into account reduced effectiveness of staff in training and experienced staff's time needed to train

new staff) are shown in figure 16d. Figures 16e and 16f show the dynamics of the variables that go into decisions about how strongly management (figure 16e) and staff (figure 16f) want to share information. Management consider potential direct and indirect costs of a cyber attack, both of which grow steadily, as well as the cost savings realized through information sharing. Cyber defense staff consider their workload (which depends on the staffing dynamics shown in figure 16d), the utility of shared information (which grows as more information is shared), free riding behaviors and pressure from other organizations to participate (zero in this scenario since all organizations participate at the same rate), and perceived security risk from sharing information (which decreases as the program goes on, since trust in participating organizations builds).

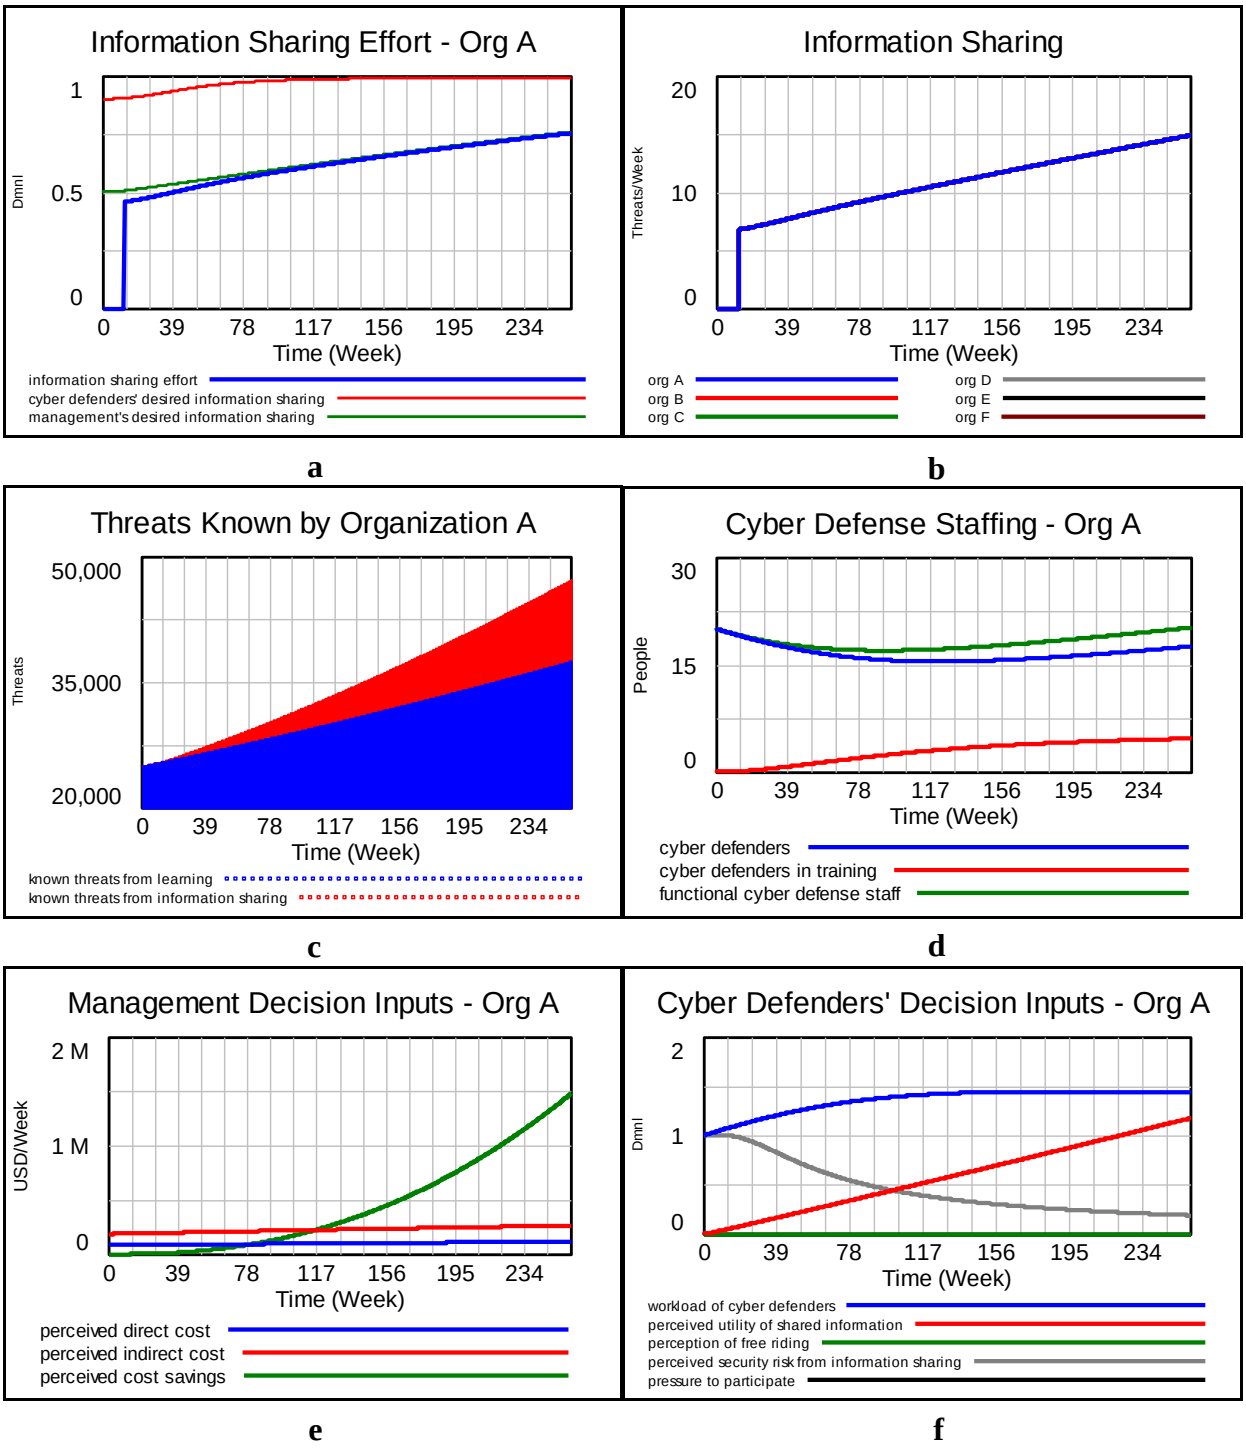

Figure 16: Results of the base case simulation (all six organizations' cyber defenders have 0.9 initial desire to share information)

Results of the second scenario, in which organization F acts as a low-grade free rider, are shown in figure 17. When one organization participates less than the others, it still has access to information shared by the other organizations but does not put as much effort into the program.

Liu et al. (2013) found that real-world information sharing programs offer incentives to free ride, and discussed policy designs to counteract this behavior. In this simulation, five organizations' (figure 17a) cyber defenders have a baseline desire to participate of 0.9, about the level indicated by our subject matter experts. The last organizations' cyber defenders' baseline participation (figure 17b) is 0.5, so that the organization begins at just over half the participation rate of the other organizations. In the first year of the program, the free riding behavior causes a noticeable dip in participation rates for the fully-participating organizations.

As the free riding behavior is noticed, the other organizations have to main actions. They reduce their participation rates, and they pressure the free riding organization to share more information. This pressure is effective, and the free riding organization steadily increases its participation (figure 17c) over the time horizon. While all participating organizations see substantial cost savings from the information sharing program (figure 17d), the free-riding organization gains slightly more savings, since it has access to information shared by five high-participating organizations, while the others have access to information shared by four fully-participating organizations and one low-participating organization.

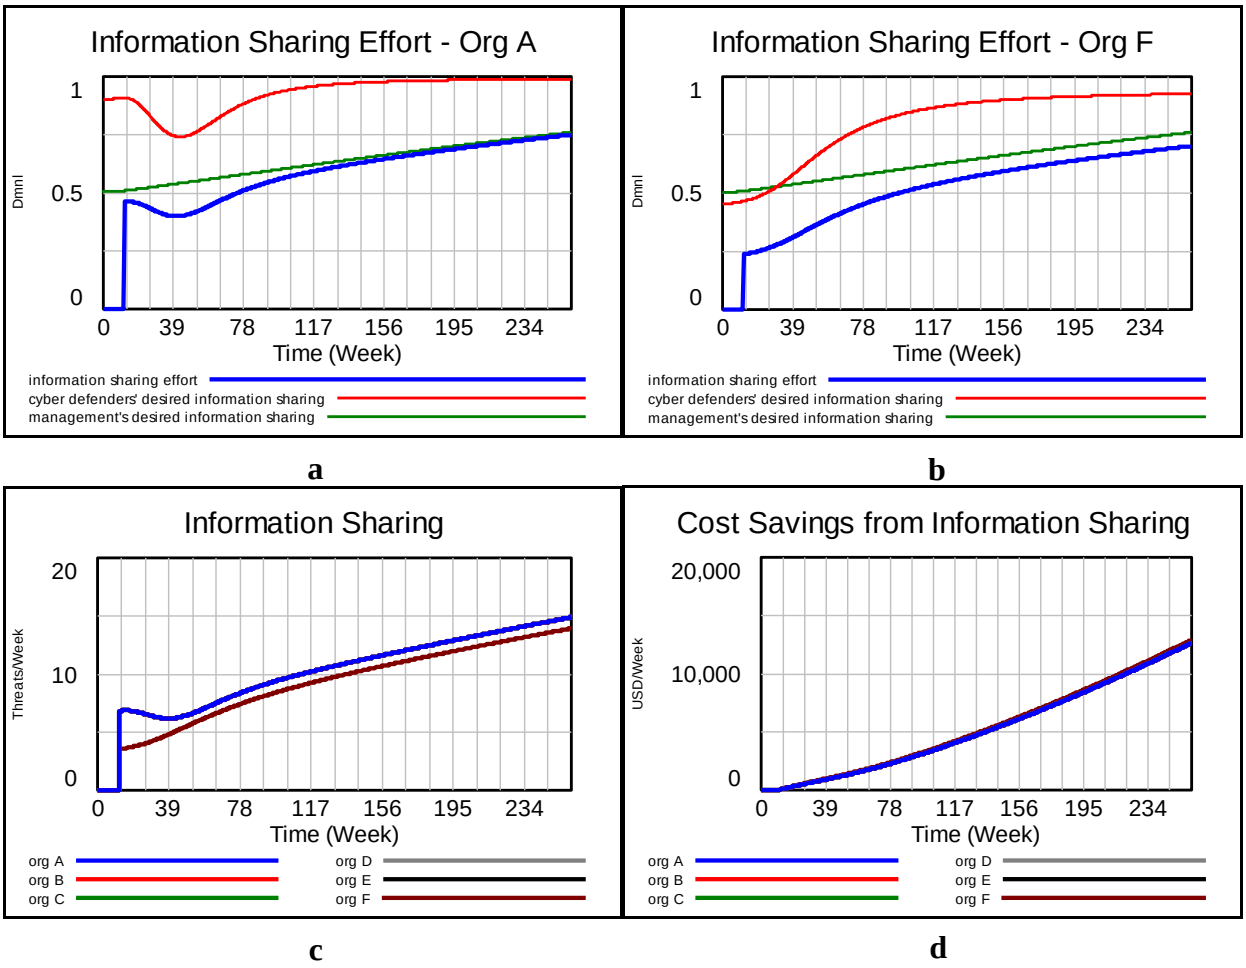

Figure 17: Results of the low free-rider simulation (cyber defenders in organization F have 0.5 baseline desire to share information, while other organizations begin with 0.9 desire)

In the final scenario, cyber defenders from organization F have a very low baseline desire to participate. In the first six months of the program, the other organizations resent this free riding, and their participation drops dramatically (figure 18a). They also pressure organization F to participate more, which is effective until approximately week 40 when organization F realizes that it is now sharing more information than the other organizations (figure 18b). Organizations F then decreases it's participation, the other organizations follow, and a dampening oscillation pattern is seen (figure 18c). As in the previous scenario, organization F shares less information than the others overall but sees the most cost savings from the program, since it has access to information shared by the other organizations. Even with this free riding behavior, the information sharing program is quite stable and successful by the end of the time horizon. Volatile behaviors, driven by free riding, at the beginning of the simulation are eventually overridden when organizations realize major benefits from the program, including cost savings, a major input to management decisions, and perceived utility of previously shared information and growing trust in the other organizations, which cyber defenders take into account.

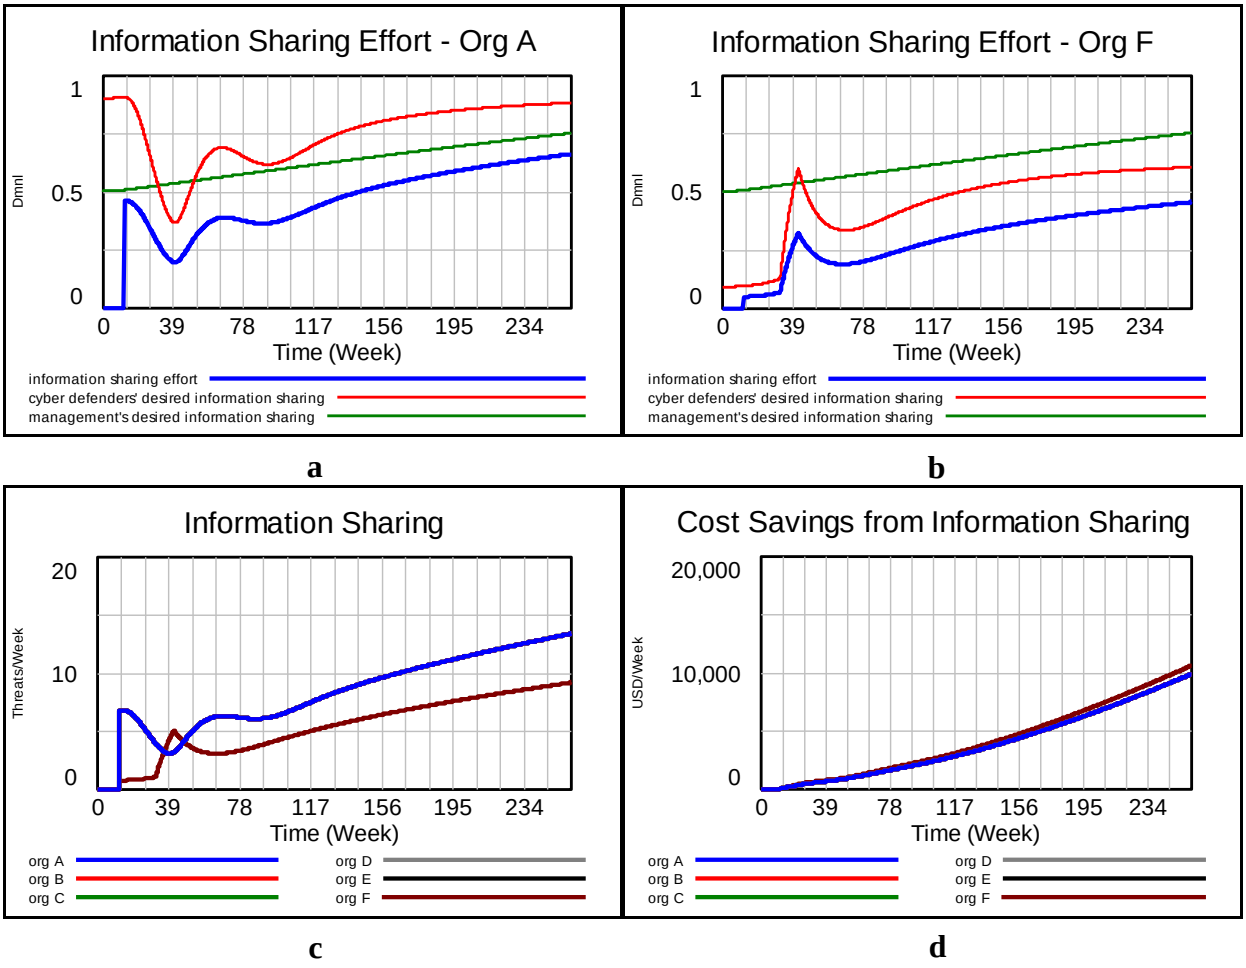

Figure 18: Results of the high free-rider simulation (organization F has 0.1 initial desire to share information, while all other organizations begin with 0.9 desire to share information)

Conclusions

Cyber security is a large and growing problem, and one that many organizations would like to deal with more effectively. Information sharing programs have the potential to improve cyber security at relatively low cost, but the human aspect of such programs can add uncertainty and volatility that may not be expected by program designers. Including human decision making in a simulation of an information sharing program can shed light on potential problems that likely wouldn't be found in typical simulations that avoid modeling cognitive processes.

The model described here simulates an information-sharing program between six organizations. The decision-making structure for management and cyber defenders in each organization is identical, and based on interviews with subject matter experts about their decision-making strategies. We simulated three scenarios to shed light on how free riding behavior would be likely to affect a cooperative program such as this. While free riding and

resulting pressure to participate caused volatility at the beginning of the time horizon, it was eventually overridden by other factors and the program was successful. Management groups took info account the cost savings gained through the information-sharing program, which grew as more information was shared and the cyber threat became more intense. Cyber defenders considered the perceived utility of shared information, which grew as the program continued and the pool of shared information grew, as well as the perceived security risk from information sharing. As the information-sharing program progressed, trust in other organizations grew and this perceived security risk diminished.

5. REFERENCES

Ajzen, I. (1991). The theory of planned behavior. *Organizational behavior and human decision processes, 50*(2), 179–211.

Anderson, R. & Moore, T. (2006). The Economics of Information Security. *Science* **314,** 610–613.

Backus, G., Bernard, M., Verzi, S., Bier, A., and M. Glickman. (2010). Foundations to the Unified Psycho-Cognitive Engine. SAND2010-6974, Sandia National Laboratories.

Dutta, A. & Roy, R. (2008). Dynamics of organizational information security. *System Dynamics Review* **24,** 349–375.

The Economist. (2013). Cyber-Security: To the Barricades. Feb 16, 2013.

European Network and Information Security Agency (ENISA). (2010). Incentives and challenges for information sharing in the context of network and information security.

Festinger, L. (1957). *A Theory Of Cognitive Dissonance*. Stanford University Press.
Hui, P., Bruce, J., Fink, G., Gregory, M., Best, D., McGrath, L., & Endert, A. (2010). Towards efficient collaboration in cyber security. *Collaborative Technologies and Systems (CTS), 2010 International Symposium on* (pp. 489–498).

Gates, C. and T. Whalen. (2004). Profiling the defenders. *Proceedings of the 2004 workshop on New security paradigms (NSPW '04)*. ACM, New York, NY, USA, 107-114.

Goldsmith, D., and Siegel, M. (2010). Understanding Cyber Complexity: Systems Modeling in the Financial Services Sector. ECIR Working Paper, February 2010.

Goldsmith, D., and Siegel, M. (2012). Systematic Approaches to Cyber Insecurity. ECIR Working Paper, January 2012.

Gonzalez, J. J. & Sawicka, A. (2003). The role of learning and risk perception in compliance. in *Proceedings of the 21st International Conference of the System Dynamics Society, New York*.

Gordon, L.A., Loeb, M.P., & Lucyshyn, W. (2003). Sharing Information on Computer Systems Security: An Economic Analysis. *Journal of Accounting and Public Policy*, 22(6), 461-485.

Hillen, S., Sveen, F. O. & Gonzalez, J. J. (2006). Using Dynamic Stories to Communicate Information Security. in *International System Dynamics Conference, Nijmegen, The Netherlands*.

Hui, P., J. Bruce, G. Fink, M. Gregory, D. Best, L. McGrath, and A. Endert. (2010). Towards Efficient Collaboration in Cyber Security. *2010 International Symposium on Collaborative Technologies and Systems (CTS).* (pp. 489–498).

ISAC Council. (2013). "A Functional Model for Critical Infrastructure Information Sharing and Analysis: Maturing and Expanding Efforts." 31 Jan. 2004. Web. 26 Aug. 2013.

Krishnan, R., Niu, J., Sandhu, R. & Winsborough, W. H. (2011). Group-Centric Secure Information-Sharing Models for Isolated Groups. *ACM Transactions on Information and System Security* **14,** 1–29.

Liu, C. Z., Zafar, H. & Au, Y. A. (2013). Rethinking FS-ISAC: An IT security information sharing model for the financial services sector. University of Texas at San Antonia, College of Business Working Paper Series, WP # 0023IS-673-2013.

Luna-Reyes, L. F. (2006). Trust and Collaboration in Interagency Information Technology Projects. *Proceedings of 2006 International Conference of the System Dynamics Society, Nijmegen, The Netherlands*.

Luna-Reyes, L. F., Black, L. J., Cresswell, A. M., & Pardo, T. A. (2008). Knowledge sharing and trust in collaborative requirements analysis. *System Dynamics Review, 24*(3), 265-297. doi:10.1002/sdr.404

Luna-Reyes, L. F., & Andersen, D. F. (2007). Towards a Theory of Interorganizational Collaboration: Generic Structures of Cross-Boundary Requirements Analysis. *Proceedings of 2007 International Conference of the System Dynamics Society, Nijmegen, Boston MA*.

Martinez-Moyano, I. J., Conrad, S. H., Rich, E. H. & Andersen, D. F. (2006). Modeling the emergence of insider threat vulnerabilities. in *Simulation Conference, 2006. WSC 06. Proceedings of the Winter* 562–568.

McFadden, D. (1982). "Qualitative Response Models," in Advances in Econometrics, Ed. Werner Hildenbrand, Cambridge University Press, New York.

The MITRE Corporation. (2012). Cyber information-sharing models: An overview.

Moore, A. P., Cappelli, D. M., Joseph, H. & Trzeciak, R. F. (2006). An Experience Using System Dynamics to Facilitate an Insider Threat Workshop. *Unpublished Paper, Carnegie Mellon University CERT Software Engineering Institute*.

Oliver, C. (1990). Determinants of interorganizational relationships: Integration and future directions. *Academy of management review*, 241–265.

Ponemon Institute. (2013). 2013 Cost of Data Breach Study: Global Analysis. Ponemon Institute Research Report.

Radianti, J., Gonzalez, J. J. & Rich, E. (2009). A quest for a framework to improve software security: Vulnerability black markets scenario. in *Proceedings of the the 27th International Conference of the System Dynamics Society*.

Rich, E., Martinez-Moyano, I. J., Conrad, S., Cappelli, D. M., Moore, A. P., Shimeall, T. J., Andersen, D. F., Gonzalez, J. J., Ellison, R. J., Lipson, H. F., Mundie, D., Sarriegui, J. M., Sawicka, A., Stewart, T. R., Torres, J. M., Weaver, E. A., & J. Wiik. (2005). Simulating Insider Cyber-Threat Risks: A Model-Based Case and a Case-Based Model. in *Proceedings of the 23rd International Conference of the System dynamics Society* 17–21.

Rich, E., Sveen, F. O. & Jager, M. (2006). Overcoming organizational challenges to secure knowledge management. in *Second Secure Knowledge Management Workshop (SKM)*.

Ring, P. S., & Van de Ven, A. H. (1994). Developmental processes of cooperative interorganizational relationships. *Academy of management review*, 90–118.

Sandhu, R., Krishnan, R., & White, G. B. (2010). Towards secure information sharing models for community cyber security. *Collaborative Computing: Networking, Applications and Worksharing (CollaborateCom), 2010 6th International Conference on* (pp. 1–6).

Simon, H.A. (1957). Administrative Behavior (2nd ed.). New York, NY: Macmillan

Tversky, A. & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. Science, 185, 1124-1131.

U.S. Department of Energy (DOE). (2013). "DOE O 205.1B." 11 Mar. 2013. Web. 26 Aug. 2013.

U.S. House of Representatives. (2013). "HR3523: Cyber Intelligence Sharing and Protection Act." 2012. Web. 26 Aug. 2013.

DISTRIBUTION

1	MS0899	Technical Library	9536 (electronic copy)
1	MS0359	D. Chavez, LDRD Office	1911


Sandia National Laboratories