Cooperation and Learning in Cyber Security Training Exercises

SAND 2013-2198C

31st International Conference of the System Dynamics Society
July 21th – 25th, 2013 – Cambridge, Massachusetts

Asmeret Bier and Ben Anderson
Sandia National Laboratories

Abstract

Cyber attacks pose a major threat to many organizations, and cooperation within and between organizations has potential to improve defensive capabilities. Tracer FIRE, a training program for cyber security incident responders, has begun to explore whether cooperation during training exercises can enhance learning. A system dynamics model was created using the Behavioral Influence Assessment framework, which uses well-established psychological, social, and economic theory to simulate cognition and interactions between people and their environments. The model was used to understand the relationship between cooperation and learning during Tracer FIRE, and to explore methods, using scenario exploration and sensitivity analysis, of increasing participants' learning.

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 1) 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.

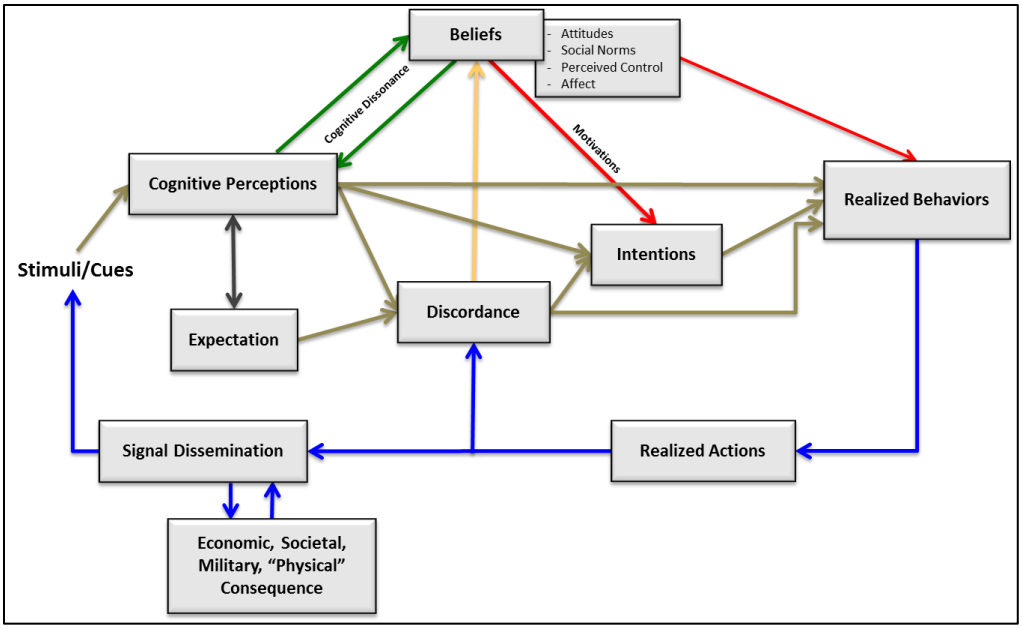

Figure 1: Computational structure of the BIA framework

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 (Appendix 1) 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.

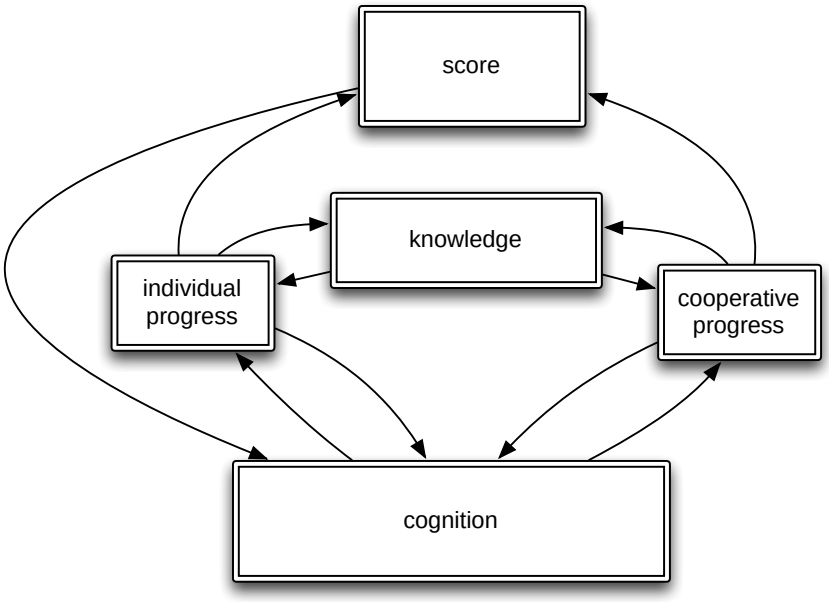

Figure 2: Model structure overview

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.

	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			-		

Table 1: Cues, cognitive perceptions, and potential behaviors

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 3. 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 challenges left to complete 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.

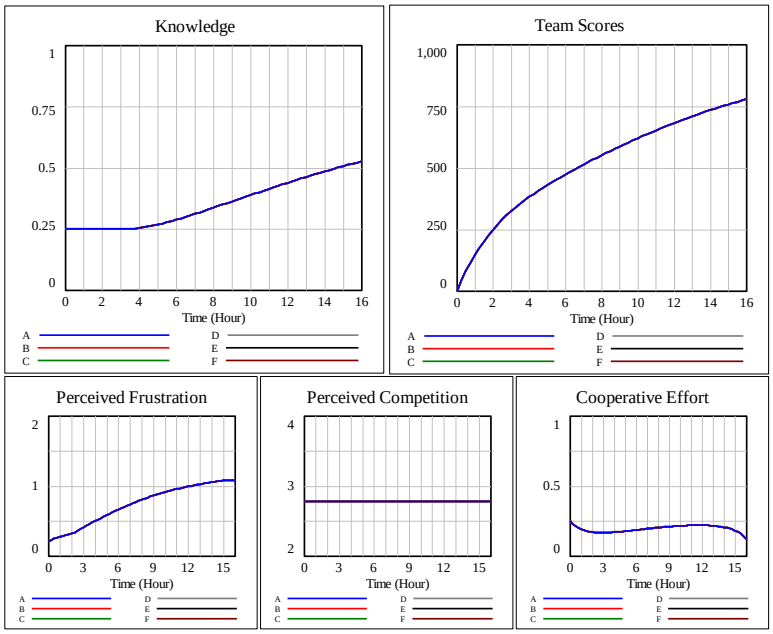

Figure 3: Base case simulation (init knowledge = 0.25, baseline cooperation = 0.25)

Figures 4a and 4b 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 4a), and if barriers to cooperation are removed to make work required to cooperate 5% (rather than 25%), knowledge reaches 70% (figure 4b).

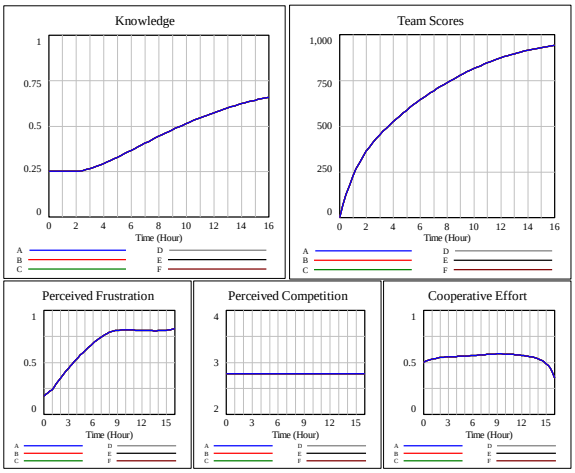

Figure 4a: Baseline cooperation = 50%; work required to cooperate = 25%

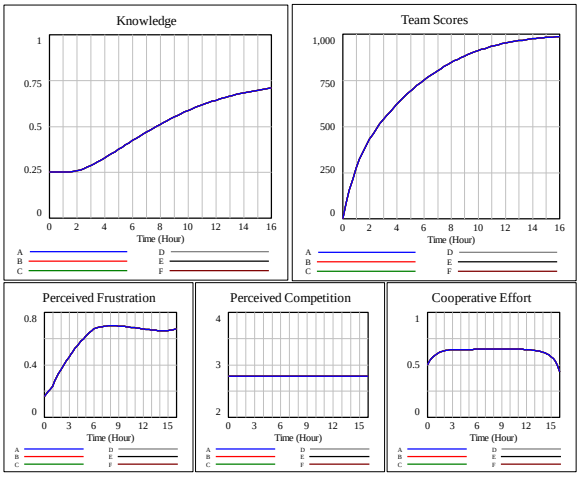

Figure 4b: Baseline cooperation = 50%; work required to cooperate = 5%

Learning can be further improved by increasing the difficulty of tasks, as in the scenario shown in figure 5. This scenario is the same as the one shown in figure 4a, 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.

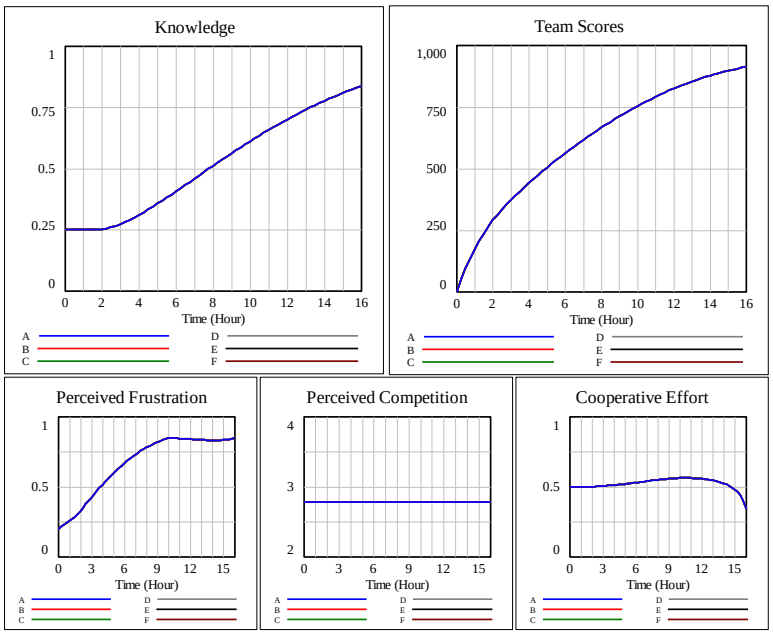

Figure 5: Baseline cooperation = 50%, task difficulty = 1

It is also likely that different teams will have different baseline cooperation levels. Figures 6a 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 6b 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.

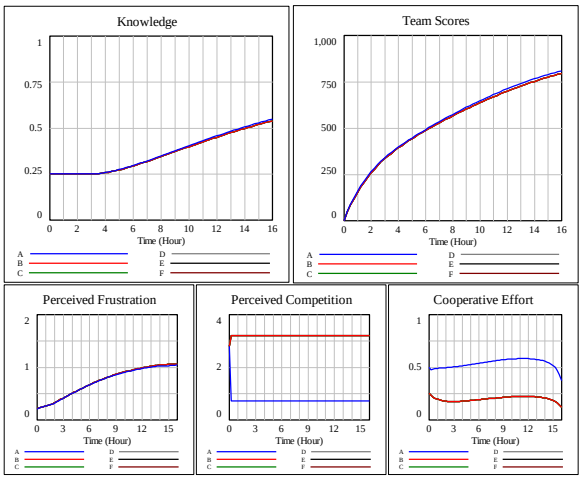

Figure 6a: One team with baseline cooperation = 50%

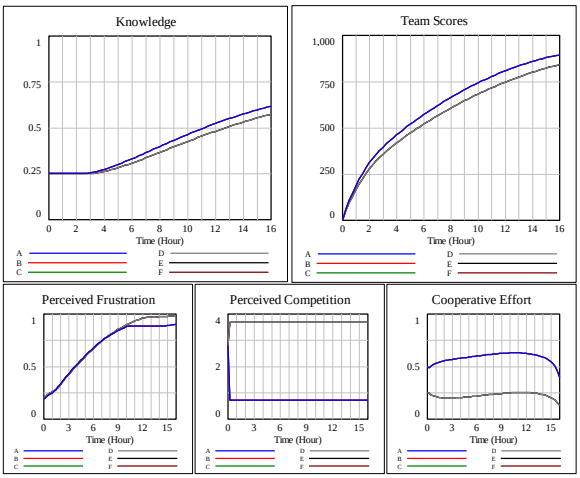

Figure 6b: Half of teams with baseline cooperation = 50%

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.

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

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

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.

References

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

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

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.

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).

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

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

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

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. *2010 6th International Conference on Collaborative Computing: Networking, Applications and Worksharing (CollaborateCom).* (pp. 1–6).

Acknowledgements

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.

Appendix 1: Model structure