Proceedings of the 2015 Winter Simulation Conference
L. Yilmaz, W. K. V. Chan, I. Moon, T. M. K. Roeder, C. Macal, and M. D. Rossetti, eds.

MODELING ADVERSARIAL DYNAMICS

Ignacio J. Martinez-Moyano
Rogelio Oliva

Argonne National Laboratory
9700 S. Cass Avenue
Argonne, IL 60439, USA

Donald Morrison

Department of Homeland Security—TSA
701 South 12th Street
Arlington, VA 20852, USA

David Sallach

Argonne National Laboratory
9700 S. Cass Avenue
Argonne, IL 60439, USA

ABSTRACT

This document describes the current state of the Adversary Dynamics Modeling (ADM) project currently under development. Given the dynamic nature of the terrorist threat, the purpose of this modeling effort is to increase current understanding of adversarial decision-making processes and possible behavior in order to help guide countermeasure technology decisions and deployment. The system dynamics approach is used to capture the underlying systemic structure responsible for adversarial activity.

1 INTRODUCTION

This document describes the current state of a model of adversarial processes under development as part of the Adversary Dynamics Modeling (ADM) project. Given the dynamic nature of the terrorist threat (Cragin and Daly 2004), the purpose of the modeling effort described here is to increase current understanding of adversarial decision-making processes and behavior in order to help guide countermeasure technology decisions and deployment. The system dynamics approach is used to capture the underlying systemic structure responsible for adversarial activity (Forrester 1961; Richardson and Pugh 1981; Sterman 2000). Additional model development iterations, as prescribed by the system dynamics modeling approach, will be created as additional relevant structure is added to the model.

Currently, the ADM effort is designed to investigate how motivated adversaries will adjust their strategy to implement an attack at an airport in the face of changes or challenges implemented by defenders. For the purposes of this paper, a motivated adversary is one who has the intent to engage in an attack against an airport. Although we recognize that intent is a potentially dynamic variable worth including in the model, we make the initial assumption that the adversary has the intent to attack the system (i.e., the level of intent is high). The conditions or pressures that cause adversarial change are investigated at both aggregate and prototypical levels. At this point in the development process, specific effects of countermeasures—such as the effect of increasing the use of passenger screening canines to match an adversary's attack schedule—might be explored only at a generic level. However, such an exploration could yield important insights into the specific problem being studied. The following section describes the main elements of the model structure. It also illustrates the range of model behaviors and

explains how the model structure and behavior are being used to further validate the model structure and to obtain more accurate data from subject matter experts.

2 MODEL STRUCTURE

2.1 Overview

The ADM depicts prototypical aggregate defenders and attackers whose intent is to make their functions as effective as possible. The attackers seek to inflict damage on the aviation industry, while the defenders seek to make protection measures as effective as possible. The interplay between defenders and attackers creates a reinforcing feedback mechanism that could explain the escalation of capabilities on both sides of the simulated interaction (see Figure 1). As depicted in Figure 1, as the capabilities of the attackers (readiness of attack system) become evident to the defenders, effort is exerted to improve the defense system so it is able to meet the challenges presented by the attackers. As the defenders improve the level of security, the attackers learn (at least partially) about the defenders' capabilities (readiness of the defense system), which fuels the attackers' need to improve their own capabilities in order to succeed in their attack endeavors. As the attackers improve their capabilities and the defenders identify evidence, the cycle starts anew, fueling a long-term escalation process.

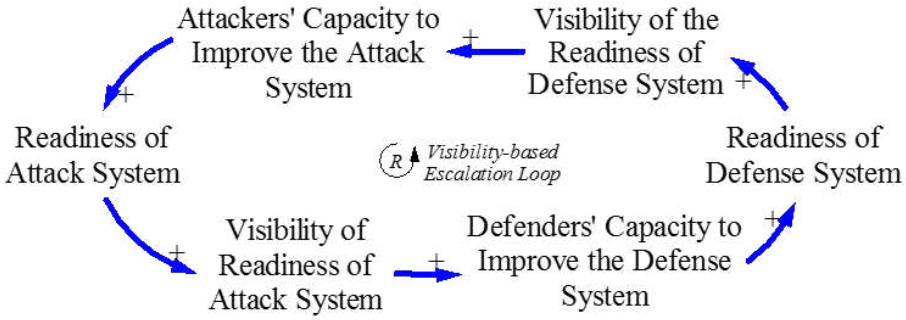

Figure 1: Adversarial escalation process.

The ADM is formulated to capture the cycle just described through understanding how the development of attack projects carried out by nine different types of motivated adversaries interact with defense-side action. The attackers develops attack projects; the defenders gather intelligence about attack development and launch defense actions when warranted.

The nine types of attackers were chosen based on two central characteristics: the size of attacker organization (individual, group, and network) and its overall capability (low, medium, and high). In the model, the attackers engage in the creation of attack projects that are initiated at a certain rate depending on the type of attacker. The different types of attacks capture the complexity and potential impacts that the nine types of attackers would have.

2.2 Simplified View of the Model

A simplified view of the model is presented next. The simplified view is broken into segments for clarity.

2.2.1 Attack Project Development

In the model, attacks are characterized as projects under development (Lyneis and Ford 2007). Figure 2 shows the structure of attack project initiation, progress, and completion. Projects are initiated with different characteristics depending on the type of attacker. The projects consist of different types of tasks that, through a completion rate, are developed and become attack tasks completed. The completion rate is

a function of the level of attacker productivity and the number of attack tasks that the attacker performs. When attackers are more productive, more tasks can be completed and more attacks can be launched. In the model, attacks are launched once an attack threshold is met. The attack threshold is measured as the percent of attack tasks completed, capturing the idea that different types of attackers have different levels of risk tolerance associated with launching an attack. The higher the attack threshold is, the lower the risk tolerance of the attacker is and, all other things being equal, the smaller number of attacks launched is because the necessary percent completion is higher. Theoretically, with low levels of attack threshold, we would see a very high attack frequency when considering motivated attackers with the resources to attack. Such attacks, however, would not necessarily be ready according to a predefined plan. One consequence of this is that the attacks would have a lower likelihood of success. Empirical evidence shows a relatively low frequency of attacks even in the face of resourceful, motivated attackers, which indicates that common attack thresholds might be high. In this conceptualization, attacks are generated when enough attack activity (preparation, planning, mobilization of resources, etc.) has been completed.

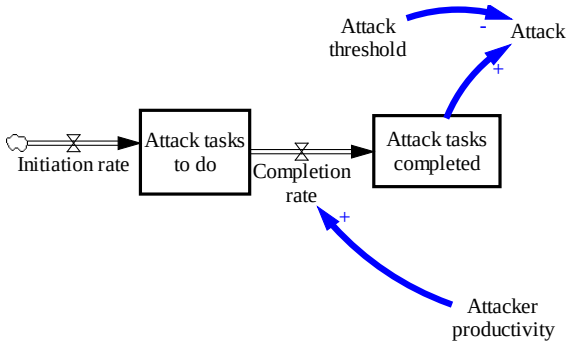

Figure 2: Basic attack projects structure.

The completion of attack activities (captured in the completion rate), in addition to advancing attack projects, creates the possibility for defenders to identify attacks that are in development. As the completion rate increases, defenders increase in knowledge about attacks. An increase in knowledge about attack activity is possible only when attackers leave behind clues about the attack during attack development and when defenders identify these clues. Consequently, the increase in knowledge is a function of the completion rate, the attackers' ability to cover progress rate, and the defenders' ability to discover progress rate. The increase in knowledge rate adds to the accumulation of evidence available to the defender, thus causing the number of information cues available to the defender to grow. As the evidence available to the defender increases, the defender prepares defense actions to prevent attacks from materializing and to prevent attacks from being successfully carried out. Once the defenders' accumulated evidence reaches a defense threshold level, a defense action is launched. Such defense actions, if successful, may change the results of attacker activity, thwart attacks, or discourage certain attack vectors. In some cases, however, the defense threshold level acts as a moving target, as new evidence becomes available and decision-makers push for more intelligence before moving forward. This effect, called the "ratchet effect" (economics) or "sliding goals," can be pervasive and has the potential of deactivating the apparatus of intelligence gathering as an actionable defense process.

2.2.2 Attack Impact

Attack impact is conceptualized as a function of the number of attacks generated, the number of defense actions, and the probability of attack success for the different attacks launched. Attack impact influences the level of perceived attacker effectiveness for attackers and their constituencies, increasing attacker motivation. Attacker motivation increases attacker productivity, which increases the completion rate,

ultimately increasing the number of attacks completed and launched. This closes an important reinforcing (also called positive feedback) mechanism in which attacks, effectiveness, and motivation are reinforced over time (R1, the attack-motivation loop shown in Figure 3). This motivation loop can be an engine of destruction as attackers launch attacks that are perceived as effective, thus increasing their motivation levels and productivity. However, the same feedback mechanism can become an engine of de-escalation if the attack impact is contained, leading to a decreased perception of effectiveness, decreasing motivation, and lessening productivity. It is important to mention that this loop is not the only one that influences motivation. Therefore, even under sustained conditions of highly constrained attack impact, determined attackers can, and will, continue their activities and will maintain their motivation to attack.

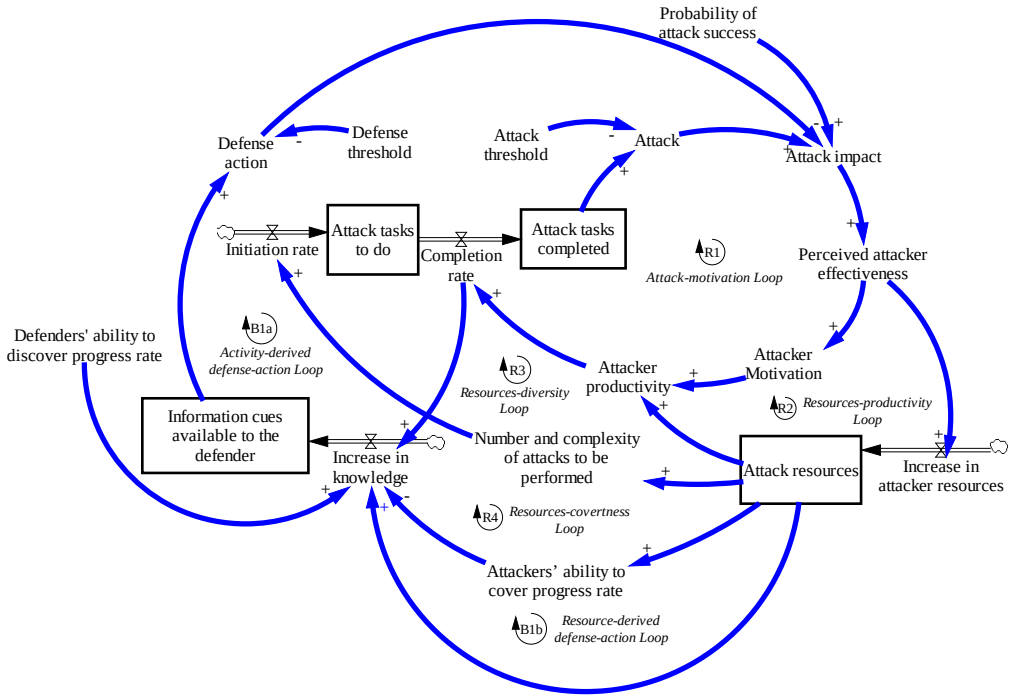

Figure 3: Influence of defense action on attacks.

2.2.3 Perceived Attacker Effectiveness

Perceived attacker effectiveness also influences the resources available to the attacker by increasing the inflow of such resources via the commitment of interested constituents. As the perception of attacker effectiveness increases, more resources are accumulated, which leads to additional attacker productivity. Attacker productivity may be increased in many ways by use of resources, including access to better-qualified attackers; by being able to train the attackers in a better way; by increasing the quality and availability of attack technology; and in other ways. Once attacker productivity increases, completion rates increase, which leads to more attacks and likely increases in attack impact. As described earlier, attack impact leads to increases in perceived attacker effectiveness, further increasing the inflow of funds to the attacker. The feedback mechanism described here (R2, the resources-productivity loop shown in Figure 3) is another reinforcing feedback mechanism that allows attackers to continue orchestrating attacks, using resources, and replenishing their resource pool as long as attack impact drives perceptions of effectiveness. This mechanism, the resources-productivity loop, describes how attackers exploit current attack scenarios by increasing the productivity of their attackers and overall plans of attack. In this sense, the expectation would be that if no other effects or mechanisms were present, as more resources became available to the attacker, more attacks of the same type and scope would be generated**.**

2.2.4 Attack Resources

The accumulation of attack resources, besides allowing attacker organizations to increase their overall effectiveness, allows attacker organizations to increase the number and/or complexity of attacks to be performed and to increase their ability to hide the progress they make in attack projects. As attack resources increase, attacker organizations have the potential to branch out, develop more attacks, and develop these attacks by using inventive attack plans that may explore new types of vulnerabilities, weapons, and targets. The creative thrust produced by the availability of resources, and the effects these resources can generate, lead to higher levels of (1) initiation of attacks, (2) number of attacks, and (3) likely attack impact. As the attack impact increases (e.g., in the case of the motivation and the resource-productivity loops), perceived attacker effectiveness increases, which leads to even more resources closing a feedback mechanism (R3, the resources-diversity loop in Figure 3) that reinforces the creation and delivery of creative and inventive attack projects as a result of past action. In general, the acquisition of resources supports adding affordances that potentially support innovative initiatives. Reinforcing processes, such as the R3, tell a (partial) story about the exploration of new types of attacks, technology, and delivery methods, which can potentially explain the existence of changes over time in attack configurations and processes. For example, when a state sponsor of terrorism, such as Iran, provides resources to a terrorist organization, such as Hamas, the resources can be used to acquire more powerful weapons (such as rockets) that allow the emergence of new strategies (such as initiating systematic rocket attacks into Israel). Although this feedback mechanism is not the only one that can help explain such changes, it adds to the overall feedback-based approach used to explain empirical evidence related to attacker activity.

The attackers' ability to conceal progress is also, in part, a function of accumulated attack resources. As attack resources increase, all activities conducted by the attackers may be better protected and conducted without leaving clues behind. Because additional resources can be applied to cover all possible cues of attack progress, a resource-rich attack process increases the likelihood of becoming invisible to the defender. As the attackers' ability to cover progress increases, the defenders' ability to discover cues and increase their knowledge about attacks and attack configurations decreases, which makes it less likely that the accumulation of usable knowledge will occur. If the defenders assemble insufficient evidence, the defense threshold may not be met, so fewer defense actions are triggered, consequently allowing for higher levels of attack impact. Higher levels of attack impact lead to increases in perceived attacker effectiveness and more attack resources; this ultimately allows for additional increases in the attackers' ability to cover their attack progress, thus closing another reinforcing feedback mechanism (R4, the resource-covertness loop in Figure 3) that further increases the growth of attacks and attack resources.

The accumulation of attack resources may consequently lead to increases in the exploitation of current attack practices, exploration of innovative attack configurations, and an increased ability to cover attack progress, causing defense actions and effectiveness to suffer. However, the accumulation of attack resources is in itself an activity, and it is a process that leaves behind clues that need to be managed by attackers. Defenders may exploit such clues in the same way they exploit attack activity clues (e.g., resource-related information cues) that are most likely precursors to attack-related activity. Consequently, the accumulation of information cues available to the defender can be conceptualized as coming from two main sources: attack activity and resource accumulation (Martinez-Moyano et al. 2008). As attack activity is identified and used to inform and launch defense actions, a balancing feedback mechanism (also called a negative feedback process) emerges (B1a, the activity-derived defense-action loop in Figure 3); this creates a counterbalance to the different reinforcing cycles that promote an ever-increasing generation of attack activity. This balancing process creates a response based on the accumulation of evidence that leads to decreases in attack impact and attacker motivation, ultimately decreasing attacker activity. The success of this mechanism is predicated on the existence of enough evidence of attack activity; when the defense threshold is met, this evidence will become the reason for launching a defensive action.

However, the success of the B1a mechanism leads to decreased levels of attack activity, which slows down the accumulation of evidence about such activity. Therefore, defenders need to be aware that highly successful defense actions have the potential to generate an evidence-starved future in the system, possibly leading to a decrease in defense activity and other consequences that may lead to decreased levels of available resources for the defender. However, when previous defense actions successfully decrease current and future attacker activity, this situation should be identified as evidence of the effectiveness of the defense system—also called deterrence—and not as evidence of a lack of need for defense. In addition, as identified earlier, clues available to the defender come not only from information related to attack activity but also from all the mechanisms and processes needed to accumulate and use attack resources (see B1b, the resource-derived defense-action loop in Figure 3). As more attack resources accumulate, more clues become available to the defender, increasing the likelihood of meeting the defense threshold for action and eventually leading to fewer attacks and less accumulation of attack resources.

2.2.5 Defense-side Consequences

The six feedback mechanisms described thus far (four reinforcing and two balancing mechanisms) interact with one another to determine the results of the different variables of interest. For example, to determine how many attacks are generated, all six feedback mechanisms are active and interacting at the same time. Depending on attacker characteristics, the different mechanisms have different relative strengths (also called "gain") at different points in time. The interaction of the various feedback mechanisms and their different strength levels shape the results and trajectories of the variables of interest.

In addition to attack-related mechanisms, defense-side processes also influence attack results. As discussed earlier, perceived attacker effectiveness increases together with increases in attack impact. Attack impact, however, has defense-side consequences as well.

As attack impact increases, perceived defender effectiveness decreases, influencing both motivation and resource generation and creating two reinforcing processes on the defense side (see dR1, the defense-motivation loop, and dR2, the defense resources-productivity loop, in Figure 4). These reinforcing processes modify defense productivity, leading to changes in defense readiness that influence the probability of attack success. As defense readiness increases, defense technology and processes produce a better system of protection, thus lowering the probability of attack success.

The probability of attack success captures the ability of an attack project to defeat the defense measures in place at a specific target. Well-developed attack projects have a higher intrinsic probability of attack success than projects that are less well suited. Consequently, improving the defense system (i.e., increasing the level of defense readiness) ultimately has the potential to lower attack impact, thereby increasing motivation and resource generation on the defense side.

Access to defense resources, in addition to increasing defense productivity and defense readiness, also increases the defenders' ability to discover progress that has been made on attack projects by the attackers. This link between the accumulation of defense resources and the ability to discover attack processes closes another reinforcing feedback mechanism (an "increasing discovery loop") that, when paired with dR1 and dR2, creates a network of processes that push for increases in defense-related activity and resources. Increases in defense resources allow the defenders to develop and deploy better (and sometimes more) technology, manpower, and processes to detect clues that are related to attack activity in progress and that are associated with the accumulation of attack-related resources that may lead to attack activity. This increased detection ability leads to an increased accumulation of evidence (i.e., clues available to the defender) to exercise defense action with the intent of lowering attack prevalence and impacts. However, as identified earlier, very high levels of defender success may lead to a very low

level of attack prevalence, potentially influencing the way in which defense activities are funded and assessed.

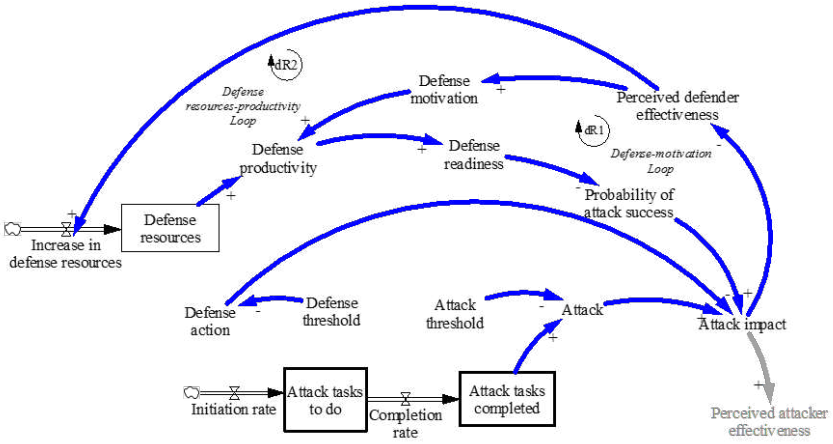

Figure 4: Defense motivation and productivity.

2.3 Model Summary

In this simplified version of the model, as identified in the literature and by subject matter experts interviewed, nine feedback mechanisms (seven reinforcing and two balancing) are identified as being crucial to adversarial processes. Of the 30 variables in the simplified model, attack impact and perceived attacker effectiveness are situated at the core of the complex set of relationships because they are part of most of the feedback mechanisms present in the model (i.e., the most central variables in the system).

In the current version of the full model, of 100 variables, progress rate (called completion rate in the simplified model)—a precursor to attack and attack impact—is at the center of the complex set of interconnections. Progress rate participates in more than 80% of the feedback mechanisms present in the model (see Martinez-Moyano 2012 for a description of the tool used to calculate centrality in the model). The centrality of the variables related to the development and delivery of attack projects is crucial to increasing the understanding of adversarial decision processes.

As an illustration of the type of information being explored the subject matter experts, the next section describes the model output under different resource allocation criteria. This output is being used as the starting point for conversations to elucidate the expected behavior of different attacker organizations under different resource constraints.

3 TRADEOFFS BETWEEN OBJECTIVES AND RESOURCES IN ATTACK PROJECTS

The process used to handle the tradeoffs between (a) the objectives and resources used in attack projects and (b) the ways in which attackers might handle such pressures lies at the core of attacker behavior and choice. Interest in this process arose as discussions with subject matter experts about attack types led to the realization that interdiction created by defense actions led, in some cases, to giving up attack patterns and, in other cases, to speeding up attacks. We have hypothesized that when attack projects are well designed and staffed, the attack projects will progress as planned, and no pressures will create the need to change objectives or resources. However, if there is a mismatch between resources and the attack project definition, or between original objectives for the attack project and current possibilities (due to defense action or other factors), pressures will emerge, and the attackers will need to decide how to deal with these pressures.

Objectives in this tradeoff space include schedule (measured in time units) and scope (measured in tasks) of attack projects; for simplification, resources are collapsed into one type of resource (attackers) of the five types identified as important in attack development (i.e., attackers, funds, logistics, intelligence,

and technology). In order to explore this tradeoff, we created a small attack project module that, once refined, will become part of the larger model.

At the core of the small project module's structure is the attack project concept, which captures the development of attack projects via a task completion rate that transforms tasks to do to complete an attack into completed tasks. The task completion rate is a function of the number of attackers assigned to the attack project, the attackers' productivity, and the number of tasks that need to be done to complete an attack.

As tasks are completed, the level of required effort is updated, which changes the schedule pressure driving the desired number of attackers in the project. In addition, as schedule pressure changes, the desired completion date is modified and the time remaining is computed. The time remaining in the project, paired with the number of attackers assigned (level of resources) also influences the schedule pressure, creating the potential to change the level of resources used, the scope of the project, the desired completion date, or all of these simultaneously, depending on attacker preferences. Attacker preferences are captured via a set of parameters that controls the assignment of adjustments over the simulated time.

3.1 Single-determinant Strategies

First, we test the use of the different levers to deal with schedule pressure (i.e., resources, completion date [also referred to as "schedule"], and scope) one at a time, as if each one was the only one available to the attackers. We simulate a prototypical (i.e., normal) attack project consisting of 100 tasks, and we introduce a change to the original definition of the scope (from 100 tasks to 200 tasks in the larger scope case and from 100 tasks to 50 tasks in the smaller scope case) to create a stressed situation in which the tradeoff becomes salient and creates a need to have it addressed. We identify the different runs to test the use of the levers by adding a word at the beginning of the name that identifies the size of the scope being tested (i.e., smaller, normal, larger) and by adding a word at the end of the name to identify the lever being used (i.e., resources, schedule, scope, all). Figure 5 shows the results for attack project tasks (tasks to do and completed tasks).

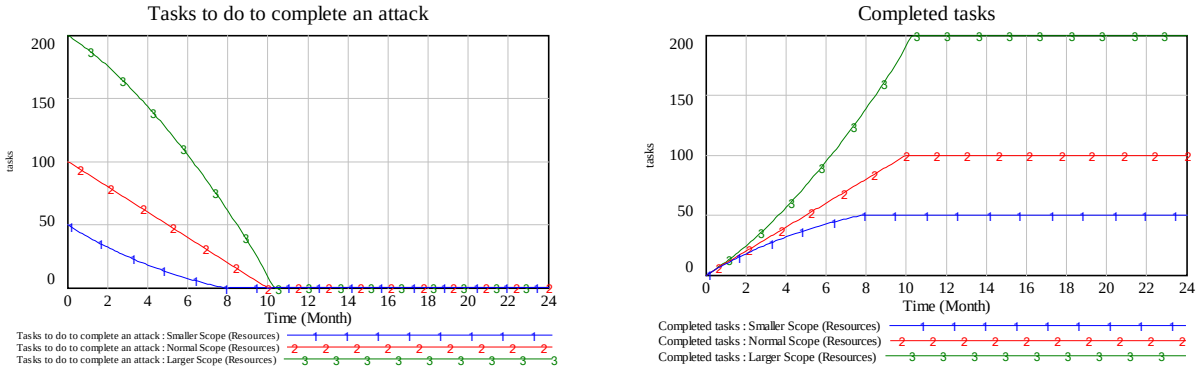

(a) (b)

Figure 5: Attack project tasks.

When using a pure resource-based strategy or a pure schedule-based strategy to correct for unplanned changes in project definition (in this case, scope), all tasks needed to complete the revised scope are performed. However, when attackers choose to use a scope-based strategy, the new scope suffers and is adjusted back to the original definition of the project. In the case of the larger scope tested (200 tasks), the pressures and choice posture of the attackers redefines the scope back to 100 tasks so that the attackers can complete it in the desired amount of time and with the desired amount of resources invested. This option would also represent a case in which the attacker is constrained by resources and time beyond the possibility of adjusting either one of the two variables to launch an attack.

In Figures 6 through 9, results for the different scope levels and strategy levers are shown. First, in the case of using a pure resource-based strategy, as expected, the overall use of resources increases as the imbalance is identified and corrected. At the peak, the resources used have tripled. Although the modified scope is double the base scope, the internal dynamics of resource allocation generate the overshoot in resources. The actual completion date of the different attack projects is shown in Figure 6a. Figure 6a shows three horizontal lines together at 10 months (vertical axis), from simulated time 0 until simulated time 8 when the first project (the smaller-scope project) is completed. A vertical drop of the "1" line is introduced, changing the completion date from 10 months (original, or designed, completion date) to 8 months (see vertical axis). The "1" line continues at 8 from that point in time on, since the project completion date will no longer change (the project is completed). In the same way, the completion dates for the project with a normal scope (10 months, represented by the "2" line) and large scope (10.38 months, represented by the "3" line) are displayed.

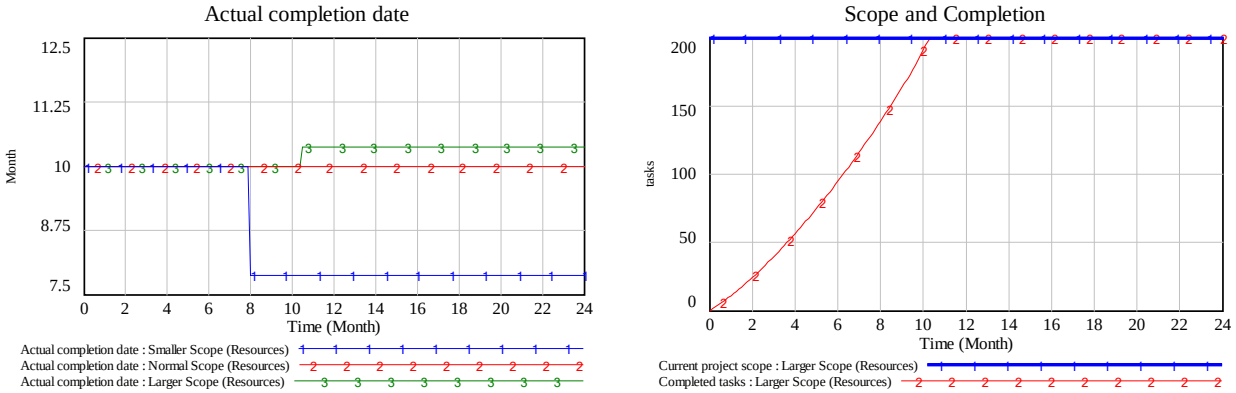

(a) (b)

Figure 6: Resources and objectives (lever used: resources).

In Figure 7, when a pure schedule-based strategy is applied, both the level of resources used and the scope of the project remain constant. However, the completion date slides out to adjust for the imbalance; in addition, the amount of time for the completion of the project doubles in order to accommodate the additional tasks (the length of the project increases from 10 to 20 months).

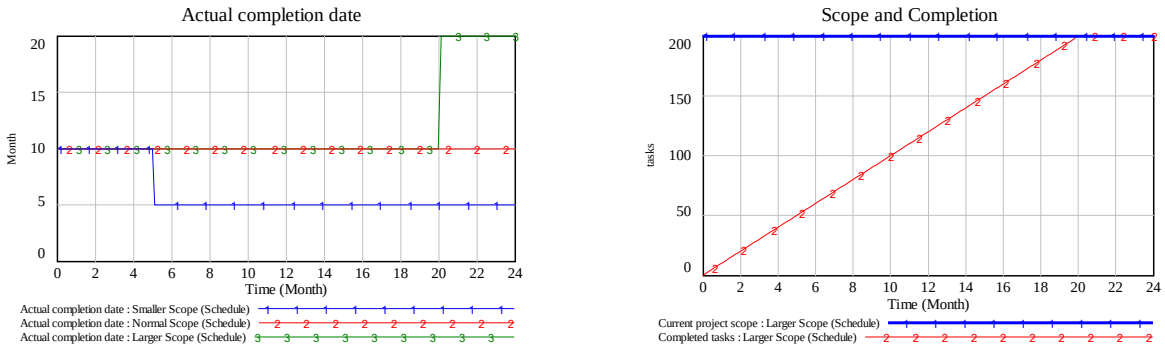

(a) (b)

Figure 7: Resources and objectives (lever used: schedule).

The last of the single-determinant strategies is to use changes to the attack project's scope as a mechanism to correct for the imbalance created by the change of scope (from 100 to 200 tasks). In this

case, both the imbalance and the correction affect the same factor, making the dynamics of the adjustment the most interesting part of the results. Figure 8 shows that the results for resource level and for completion date (schedule) are not perturbed from the baseline results. However, given the preferences tested, scope adjustment processes kick in as the schedule pressure rises due to the additional personnel required to complete the tasks. Because no adjustment in resources or schedule is allowed, the scope is adjusted rapidly (in the first 4 months) back to its original level of 100 tasks to match the original levels of attackers and allocated time to complete the project. The adjustment is quick and nonlinear due to the structure of the model.

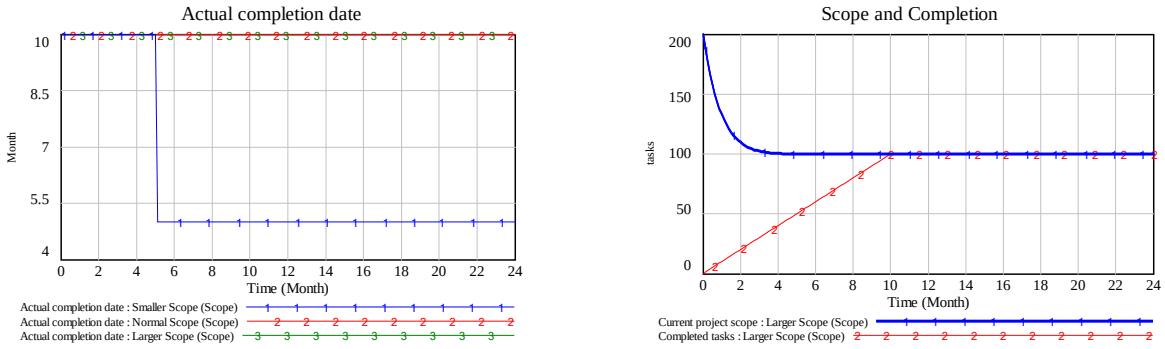

(a) (b)

Figure 8: Resources and objectives (lever used: scope).

3.2 Combined Strategy

The single-determinant strategies of intervention represent extreme choice conditions that are not necessarily realistic or available to the attackers all the time. Next, we test a combined strategy in which all the strategies are equally weighted. The results are presented in Figure 9.

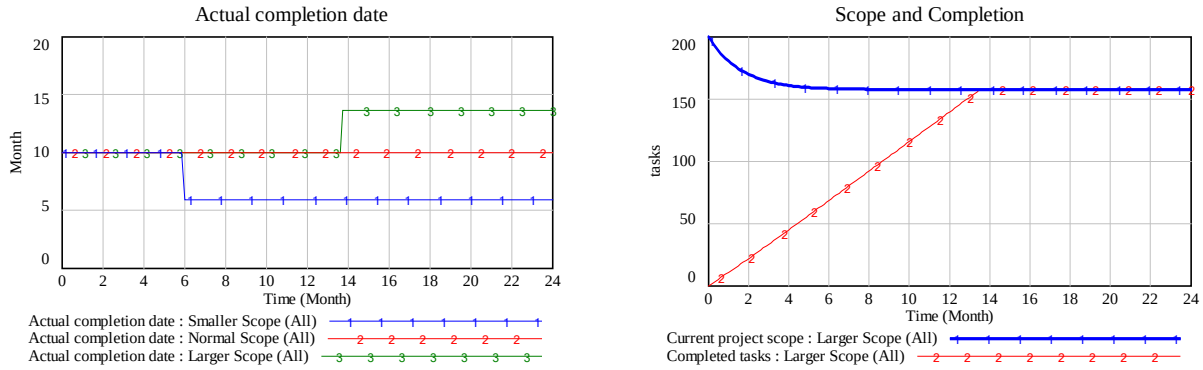

(a) (b)

Figure 9: Resources and objectives (using all change factors simultaneously).

This test is equivalent to assuming that resources and objectives are of equal importance to the attacker and are therefore interchangeable as a means of dealing with the imbalance imposed on the system. Although this is a strong assumption (it will be relaxed in subsequent versions of the model), it makes sense to present it as a prototypical hybrid case in which all current levers are used. This equally weighed option provides a boundary condition to explore and discuss with subject matter experts.

When a combined strategy is used, the use of resources does not increase as sharply, the scope adjustment is not as extreme, and the change in completion date experienced is not as large as when other levers are used. However, the use of a combined strategy—independent from the relative weights used to

balance the use of resources, schedule, and scope—assumes that the attacker's new situation allows for a certain latitude in each of the three important elements of project performance.

Table 1 summarizes the results for the larger-scope case. The rows represent the results of the different strategies (or levers) depicted in the columns. The last column (labeled "All") is the strategy in which all three levers are used at the same time and with equivalent weights (1/3 each). For example, when the strategy is to use a resource-pure approach to manage the pressures created by increasing the scope from 100 tasks (normal scope) to 200 tasks (larger scope), the use of resources changes from 10 to 32.27 people for the project, the schedule moves from the original 10 to 10.38 months, and the scope stays at 200 tasks.

Table 1: Summary of results for larger-scope case.

	Strategy			
Lever	Resources	Schedule	Scope	All
Resources (people)	32.27	0	10	11.85
Schedule (months)	10.38	20	10	13.63
Scope (tasks)	200	200	100	157.5

Understating how attackers adjust resources and objectives in order to maintain an attack project's progress (and the likelihood of perceived success) is crucial to the continued understanding of attacker proclivities and likely ways of countering their actions.

As discussed, knowing the tradeoffs between objectives and resources in attack projects is the key to understanding attacker choice. In addition, knowing how attackers might allocate resources and how they might determine a project's attractiveness in order to control project development is important in enhancing the understanding of attacker choice and the continued development of the model.

4 FUTURE WORK

We are currently in the process of identifying expected behavior for the other modules/structures described in the model structure section through subject matter elicitation and through literature review. The adversary dynamics model is an ongoing effort that has the potential to yield important results related to increasing the efficiency of security capability deployment at airports. The process followed in its development is highly iterative, providing many opportunities for reconceptualization, refinement, and incorporation of new evidence. Immediate effort will be exerted into the integration of the different modules described in this document with the overall adversarial dynamics model and in the integration of numerical data collected.

ACKNOWLEDGMENTS

The submitted manuscript has been created by UChicago Argonne, LLC, Operator of Argonne National Laboratory ("Argonne"). Argonne, a U.S. Department of Energy Office of Science laboratory, is operated under Contract No. DE-AC02-06CH11357. The U.S. Government retains for itself, and others acting on its behalf, a paid-up nonexclusive, irrevocable worldwide license in said article to reproduce, prepare derivative works, distribute copies to the public, and perform publicly and display publicly, by or on behalf of the Government.

Argonne National Laboratory's work was sponsored by the U.S. Department of Homeland Security, Transportation Security Administration under interagency agreement, through U.S. Department of Energy contract DE-AC02-06CH11357.

REFERENCES

Cragin, K., and S. A. Daly. 2004. "The Dynamic Terrorist Threat: An Assessment of Group Motivations and Capabilities in a Changing World." RAND Corporation,

Forrester, J. W. 1961. *Industrial Dynamics*. Cambridge MA: Productivity Press

Lyneis, J. M., and D. N. Ford. 2007. "System Dynamics Applied to Project Management: A Survey, Assessment, and Directions for Future Research." *System Dynamics Review* 23 (2-3):157-189.

Martinez-Moyano, I. J. 2012. "Documentation for Model Transparency." *System Dynamics Review* 28 (2):199-208.

Martinez-Moyano, I. J., E. Rich, S. Conrad, D. F. Andersen, and T. R. Stewart. 2008. "A Behavioral Theory of Insider-Threat Risks: A System Dynamics Approach." *ACM Transactions on Modeling and Computer Simulation* 18 (2):1-27.

Richardson, G. P., and A. L. Pugh, III. 1981. *Introduction to System Dynamics Modeling with Dynamo*. Cambridge MA: Productivity Press

Sterman, J. D. 2000. *Business Dynamics: Systems Thinking and Modeling for a Complex World*. Boston, MA: Irwin McGraw-Hill

AUTHOR BIOGRAPHIES

IGNACIO MARTINEZ-MOYANO is a Computational Social Scientist in the Global Security Sciences Division at Argonne National Laboratory and a Senior Fellow at the Computation Institute of The University of Chicago. Dr. Martinez-Moyano is Editor of the Notes and Insights Section of the System Dynamics Review and has published in academic journals such as *Organization Science*, *Journal of Public Administration Research and Theory*, *ACM Transactions on Modeling and Computer Simulation* (TOMACS), *Computers & Security*, *System Dynamics Review*, and *Government Information Quarterly*. His email address is imartinez@anl.gov.

ROGELIO OLIVA is a System Dynamics Specialist in the Global Security Sciences Division at Argonne National Laboratory, an Associate Professor of Information and Operation Management at the Mays Business School, and a Research Affiliate at the Massachusetts Institute of Technology (MIT) Center for Transportation and Logistics. Dr. Oliva's research work has been published in academic journals such as *Management Science*, *Organization Science*, *California Management Review*, *Production and Operations Management*, *Journal of Operations Management*, and *System Dynamics Review*. He currently serves as Associate Editor of the System Dynamics Review. His email address is roliva@anl.gov.

DONALD MORRISON is a Program Analyst in the Systems Risk Analysis Branch within the Mission Analysis Division of the Office of Security Capabilities at the Transportation Security Administration, Department of Homeland Security. His email address is Donald.Morrison@tsa.dhs.gov.

DAVID L. SALLACH is a social theorist and computational sociologist in the Global Security Sciences Division of Argonne National Laboratory who specializes in the design of rich social agent architectures. He is also a Senior Fellow in the joint Computation Institute of the University of Chicago and Argonne. From 1998 to 2003, he served as Director of Social Science Research Computing at The University of Chicago. His work has been published in a variety of journals, including *Rationality and Society*, *Communications of the ACM*, and *Social Science Computer Review*. His email address is sallach@anl.gov.