ABSTRACT

In this paper, we investigate the dynamics between a defender and an attacker when considering the issue of technology in a multiperiod sequential game with uncertainty. In this setting, defenders can improve their chances of defending against an attack by investing in technology, whereas attackers can forego attacking by using their time to accumulate knowledge, resources, or technology to improve their future chances of success. Because dynamic games of this type are generally difficult to solve, we examine a simple modified dynamic programming algorithm that could be used to computationally analyze problems in this framework. We study how parameters behave in this model in order to understand how they affect the optimal behavior of each player and later compare simple heuristics for each player to the optimal solution to this model. We show that there could be gross inefficiencies when the optimal timing of technology adoption and accumulation is not considered.

INTRODUCTION

One of the key challenges in maintaining security is keeping up with developments in technology. In recent years, government agencies have faced many potential threats to its security, from deterring the entry of terrorists to more nontraditional forms of threats ranging from shoe bombers to contraptions involving printer cartridges, liquids, and gels. With attackers becoming more creative and technologically sophisticated, an important challenge for government agencies is to find ways to respond to these challenges in a timely manner, keeping in mind their budget and resource limitations.

In the context of warfare, the choice of technology levels can also be critical in determining the ability to defend a country or to invade another. Nations that face potential threats continually invest in upgrading and maintaining their defense mechanisms by allocating significant portions of their annual budgets to this end, while other states or terrorist groups are also investing, purchasing, and utilizing more sophisticated weaponry, devices, and technologies.

The timing then of the new technologies to be adopted by both parties becomes an important issue to consider, especially because a single successful attack can have a tremendous impact on one or both parties, may it be with respect to human lives lost, damage to property, impact on the economy, hysteria that it may cause the citizens of a country that was attacked, as well as the significant costs associated in defending and attacking a state. As new technologies to maintain security, safety, and defense emerge, implementation and adoption may be viewed as a necessary step for a defender since these changes should lead to improved levels of security. However, given the state of the economy, the adoption of certain technologies may not be practical or even feasible due to the costs associated with adopting and operating these technologies. Similarly, an attacker neeeds to consider the issue of timing. He may need to balance the decision of whether to attack at any period with the decision of whether to allocate the resources for a certain period toward improving their technology to ensure that future attacks have a higher probability of success.

In this paper, we study technology adoption in the context of strategic multiperiod security games. Although technology accumulation and adoption has been widely studied in the economics and decision analysis communities, this topic has not yet been extensively studied in the defense studies literature. This paper tries to fill in this gap. One novel feature of this paper is that the dynamics between an attacker and defender are fully analyzed in a multiperiod setting. This multiperiod setting is a tacit requirement in analyzing technology adoption issues, since the tradeoffs between immediately adopting a technology and delaying such action can only be meaningful in a dynamic setting. We recommend a framework that could be used in analyzing such a complex decision where interdependence between periods plays a critical role. Moreover, we illustrate the importance of these timing decisions for the defender by looking at the significant improvements in costs savings that can be achieved compared to simple reasonable heuristics that could be employed.

LITERATURE REVIEW

The use of games in modeling the relationship between an attacker and a defender

Technology Adoption, Accumulation, and Competition in Multiperiod Attacker-Defender Games

Victor Richmond R. Jose

Georgetown University
vrj2@georgetown.edu

Jun Zhuang

University at Buffalo
jzhuang@buffalo.edu

APPLICATION AREAS: Strategic Operations, Arms Control and Proliferation, Countermeasures, Protection, Analysis of Alternatives, Decision Analysis

OR METHODOLOGIES: Dynamic Programming, Decision Analysis

has a long history in economics, starting with the work of Dresher (1961). The variety of applications and research related to issues in military operations research and defense studies is rich, and is highlighted in recent surveys by Kardes and Hall (2005) and Sandler and Arce (2003). However, most of these applications simply focus on single-period games or repeated games where the significant information from previous periods is ignored.

In modeling the timing of when to adopt new technologies, single-stage games are no longer adequate in showing the dynamic nature of accumulating and investing resources to obtain new and better technologies. Past decisions and information become relevant to decision making. Several studies consider multistage games in the context of homeland security. For example, Bandyopadhyay and Sandler (2011) consider the interaction between preemption and defense in a two-stage game, where two defenders (home and foreign) face potential threats from a common attacker. Faria (2003) considers a continuous time game between a terrorist that maximizes his damage and a government that maximizes national security. His dynamic system-based model shows that equilibrium for such a relationship is a cyclical pattern of terror. Feichtinger and Novak (2008) also conclude that a persistent oscillatory behavior is justifiable in the long run using an open-loop Nash solution concept of dynamic game theory.

Zhuang et al. (2010) model secrecy and deception in a multiperiod signaling game, where defenders can achieve cost savings in terms of defense costs when secrecy and deception are considered in a game. Hausken and Zhuang (2011a) consider a multiperiod game in which defenders can allocate their resources to both defend these resources and attack attackers in hopes of downgrading their resources for future attacks. Similarly, attackers can allocate their resources to launch attacks and defend their own resources. Their game is designed by repeating a single-period game and linking them by the end of period resources to the start of period resources at the next period. In another paper, Hausken and Zhuang (2011b) consider a government facing a terrorist who can stockpile in a two-period game. In particular, they see that in this two-period game the first-period resources, cost functions, and resource growth factors can generate scenarios where stockpiling can occur.

With respect to the literature on technology adoption, accumulation, and competition, most of the research in this area has focused on a single decision maker or firm who plays a game against nature; however, a few papers have discussed the notion of competition in technology adoption typically involving a single common technology. For example, Reiganum (1981) studied the timing of adoption of a new technology in a game highlighting the phenomenon of a leader and a follower even among identical firms when competition is introduced. Chambers and Kouvelis (2003) use a Cournot-style game between firms in a two-period setting to study the interaction between learning and investment effects for a single-product-producing firm. The focus of these papers, however, differs. They tend to focus on the adoption of a common technology and competition for market share in a common market pool, making the insights somewhat distant and irrelevant to the context of security and defense games. In addition, these papers generally consider a zero-sum game where the share of the customer base is the prize, which may not be appropriate for the attacker-defender game.

In this paper, we connect the literatures on technology adoption and attacker-defender games by using this modeling framework in understanding the problem of optimal timing and purchase of newer and potentially costlier technologies for attacking and defending. Although some papers in the attacker-defender games literature discuss some elements or notions of technology adoption, this paper also differs from these papers in several respects. First, almost all models in the attacker-defender games are single- or two-period games, which may not fully capture the issue of timing investments in newer technologies. One exception to this is Hausken and Zhuang (2011a), who consider a *T*-period game. In their model, however, decisions are made myopically, without considering the entire event horizon. Given that investments in technology can be viewed as investments for the future, this myopic approach that led to analytically tractable solutions

will be inappropriate for our context since we would like to capture the tradeoffs each player makes at every period, which are dependent on the history of decisions that were made. To resolve this issue, we use a dynamic program approach that searches for a global optimum, which a myopic approach cannot provide.

The model we propose for technology adoption is a multiperiod attacker-defender game in which both players are strategic or forward-looking—i.e., they consider not only the benefits they will attain at the present period but also how their choices would impact latter periods in the game. Given that these problems tend to be analytically intractable, we also propose a dynamic programming framework that can be used to solve this problem. The advantage of the dynamic programming approach is that it provides a global optimum that may not necessarily be the same as the local optima for each period that the myopic approach provides.

We show through some numerical illustrations, sensitivity analyses, and benchmarking studies of certain heuristics some of the potential benefits this framework provides.

MODEL

Model Setup and Notation

We consider a sequential, finite-horizon game with uncertainty between a defender and an attacker, where the basic sequence and dynamics of the game is summarized in Figure 1. In this game, the defender, at the beginning of each period t, can choose to maintain her current level of technology for defense or adopt a higher and costlier level of technology. The attacker observes the level of technology of the defender and can choose to launch an attack or use the time to accumulate resources to improve his chances of success in future attacks.

If the attacker chooses to launch an attack, the attack's success is determined probabilistically by the current levels of capability by the attacker to attack and the defender to defend.

This process continues until the end of the horizon differing only with the level of technology available for each player at the start of each period, which depends on the sequence of acts that each player makes in the earlier periods.

Denote

t = time period, where $t = 1, \ldots, T \leq \infty$;
M = maximum level of technology feasible for the defender;
D_t^d = decision of the defender at time t, where $D_t^d \in \{0, 1, \ldots, M\}$, which represents the additional levels in technology that the defender chooses to add at time t;
D_t^a = decision of the attacker at time t, where $D_t^a \in \{0, 1\}$, where 1 and 0 represent the decision to accumulate and to attack, respectively;
$\psi_t = \sum_{j=1}^{t} D_t^d$ is the level of technology for the defender at time t;
$\lambda_t = \sum_{j=1}^{t} D_t^a$ is the level of accumulation for the attacker at time t;
$\mathbb{P}(\psi, \lambda)$ = probability of a successful attack when the technology level of the attacker and defender are ψ and λ, respectively;
v_a, v_d = damage valuation associated with a successful attack, to attacker and defender, respectively;
δ_a, δ_d = discount factor of the attacker and defender, respectively.

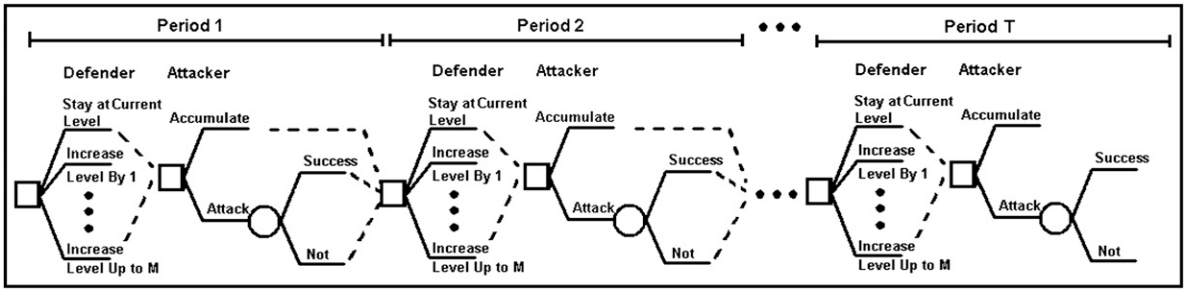

Figure 1. Attacker-defender game.

Although many factors could potentially affect the probability of a successful attack, in this paper we model this success probability primarily as a function of the technology level ψ of the defender and the accumulation level λ of the attacker. Without loss of generality, we order ψ and λ such that higher values are considered to be of higher quality—i.e., the higher ψ (λ), the lower (higher) the chance an attack will be successful. In addition, we typically should choose functions that demonstrate some notion of a diminishing relative effectiveness of technology as the technology levels of both players increase.

To clarify what we mean by level of technology (and accumulation), we associate this level to the set of technologies (resources) that are adapted by the defender (attacker) to maintain (compromise) security of the target. It does not necessarily refer to a specific technology but to a portfolio of techniques that are used and employed, although it may be possible that the main difference between a certain level and next level available for adoption is a single technology. This implies that the state space is generated by an ordering of the set of possible actions, which typically is comprised by some small set of specific technologies together with some appropriate combinations of these individual technologies. Thus the probability associated with ψ is not attributable to a specific technology but to the joint implementation of the tools in that portfolio for a specific chosen level. In the end, the technology set employed by the defender (attacker) indexed from 0 to M (0 to T) is an ordered set in which the portfolio of technologies associated with a higher index leads to a lower (higher) chance of a successful attack. We also assume that once a higher level of technology is implemented, a defender will not choose a lower level of technology in the future.

As a simple but realistic example, consider a Transportation Security Administration (TSA) point of entry trying to deter the entrant of a potential attacker into a secured area such as an airport. They can employ a variety of approaches such as physical search, metal detector, and x-ray body scanner. One way to create a technology set is to have the following $\{\psi = 1 =$ physical search, $\psi = 2 =$ physical search and metal detector, $\psi = 3 =$ x-ray full body scanner$\}$. Here, as ψ increases, the probability of defense (i.e., deterring the entrance of an attacker) increases because of the quality of the technology. Additional refinements can be made to have a larger set of alternatives (e.g., machines of differing qualities, newer methods for detection, more combinations of alternatives). In this instance, we see that the increments may be discrete and the difference between each step is ''equal,'' but the implications that they have in terms of effectiveness and costs are not necessarily ''equal.''

Related to better chances for defense, we associate a fixed cost for adopting a new technology for the defender at period t, denoted by $C_d(\psi, t)$. Aside from this one-time cost, we also include a period-dependent operating cost $C_o(\psi, t)$ incurred every period when technology level ψ is implemented. Similarly, the attacker has some associated cost $C_a(\lambda, t)$, which is dependent on the attacker's accumulation level. Although time is not explicit in this notation, we note that λ is associated with time since D_t^a in $\{0, 1\}$ implies that $\lambda \leq t + \lambda_0$ for every period t.

Problem Formulation and Equilibrium Concept

In each period t, the defender is faced with a decision D_t^d whether to maintain the current level of technology ($D_t^d = 0$) or invest in a higher level of technology by increasing from the current level ψ_{t-1} up to $M(D_t^d \in \{1, \ldots, M - \psi_{t-1}\})$. The defender incurs an adoption cost if she chooses to invest and an operating cost dependent on the ending level of technology. In terms of optimization, the defender tries to minimize the value function V_t^d representing the expected costs by weighing the costs associated with these investments and the expected savings in the reduced expected damage of a successful attack if the attacker chooses to launch an attack. After the defender sets up a defense, the attacker chooses whether to launch an attack in a manner that maximizes his value function V_t^a. The advantage of not launching an attack is that the attacker can further develop capabilities

or technology to improve the success of future attacks.

In analyzing this problem, we employ the traditional Nash equilibrium concept. Moreover, similar to Powell (2007), Zhuang et al. (2010), and Golalikhani and Zhuang (2011), we employ the subgame perfect Nash equilibrium (SPNE) refinement. This means that the solution at every subgame is a Nash equilibrium itself, i.e.,

$$D_t^{d*} = \arg\min_{D_t^d} V_t^d(D_t^d \mid \psi_{t-1}, \lambda_{t-1})$$

$$D_t^{a*} = \arg\max_{D_t^a} V_t^a(D_t^a \mid \psi_{t-1}, \lambda_{t-1}, D_t^{d*})$$

for $t = 1, \ldots, T$. This can usually be solved through backward induction. Specifically, in period T, the attacker decides to maximizes his or her value by comparing the expected benefit of attacking (expected value of target less the cost of attacking) versus doing nothing, i.e., the optimization problem for choosing between attacking or doing nothing is

$$V_T^a(D_T^a \mid \psi_T, \lambda_{T-1}) = \max_{D_T^a}\{\mathbb{P}(\psi_T, \lambda_{T-1})v_a - C_a(\lambda_{T-1}, T), 0\}. \quad (1)$$

Equation 1 implies that if the costs are not significantly high, the attacker would choose to attack. Otherwise, the attacker would simply choose to accumulate, which is paramount to doing nothing since action is no longer permissible after time T. The defender has the following optimization problem:

$$\begin{aligned} V_T^d(D_T^d \mid \psi_{T-1}, \lambda_{T-1}) = & \min_{D_T^d : D_T^d + \psi_{T-1} \le M} \{C_d(\psi_{T-1} + D_T^d, T)\mathbf{1}_{[D_T^d \neq 0]} \\ & + C_o(\psi_{T-1} + D_T^d, T) \\ & + \mathbb{P}(\psi_{T-1} + D_T^d, \lambda_{T-1})v_d\mathbf{1}_{[D_T^a = 0]}\}. \end{aligned} \quad (2)$$

Repeating this, we have the following Bellman equation for any intermediate period $t < T$

$$\begin{aligned} V_t^a(D_t^a \mid \psi_t, \lambda_{t-1}) = & \max_{D_t^a}\{[\mathbb{P}(\psi_t, \lambda_{t-1})v_a - C_a(\lambda_{t-1}, t)]\mathbf{1}_{[D_t^a = 1]} \\ & + \delta_a V_{t+1}^a(\psi_{t+1}(D_t^a), \lambda_{t-1} + (1 - D_t^a))\} \\ V_t^d(D_t^d \mid \psi_{t-1}, \lambda_{t-1}) = & \min_{D_t^d : D_t^d + \psi_{t-1} \le M} \{C_d(\psi_{t-1} + D_t^d, t)\mathbf{1}_{[D_t^d \neq 0]} \\ & + C_o(\psi_{t-1} + D_t^d, t) \\ & + \mathbb{P}(\psi_{t-1} + D_t^d, \lambda_{t-1})v_d\mathbf{1}_{[D_t^a = 0]} \\ & + \delta_d V_{t+1}(\psi_{t-1} + D_t^d, \lambda_t(D_t^d))\} \end{aligned} \quad (3)$$

Here, we note that the value function becomes complicated quickly because the next period's state variable λ_t for the defender is dependent on the decision D_t^d since this will dictate what the attacker will do in the intermediate stage. Similarly, this is true for ψ_{t+1} and D_t^a for the attacker since D_t^a will dictate the defender's optimal strategy in the next period.

Figure 2 summarizes the steps in the dynamic programming algorithm that we use to solve this game for any set of cost functions C_a, C_d, and C_o. The recovery of the optimal solution often yields a clear decision for a player. In the rare event that a player is indifferent between two strategies, we note that an alternative optimal solution exists and that the Nash equilibrium would then involve a mixed strategy between alternatives. For example, in the case of an attacker, it may be possible that the attacker would instead choose to attack with a certain probability a^*, and with probability $1 - a^*$ choose to accumulate.

In addition, if we choose reasonable functions that are bounded, compactness guarantees that an optimal solution exists. One thing to note is that the algorithm performs better than brute force enumeration of all possible paths that each player can take, but like most dynamic programming models, the algorithm still would be prone to some computational issues when T and M are high. Fortunately, this does not pose a significant problem since the expansion is on the action space and horizon rather than the state space, where issues related to the curse of dimensionality may occur.

Algorithm:

Initialize:
Set $T, M, v_a, v_d, \delta_a, \delta_d, \psi_0, \lambda_0$, parameters of C_a, C_d, C_o

Last Period Solution:
for $(a,b) \in \{\psi_0, \ldots, \mathtt{M}\} \times \{\lambda_0, \ldots, T+\lambda_0\}$
Compute and Store
$V_T^a(D_T^a|a,b) := \max_{D_T^a \in \{0,1\}} [\mathbb{P}(a,b)v_a - C_a(b,T)]\mathbf{1}_{[D_T^a=0]}$
$D_T^{a*}(a,b) := \arg\max V_T^a(D_T^a|a,b)$
$V_T^d(D_T^d|a,b) := \min_{D_T^d: D_T^d + a \le M} C_d(a+D_T^d, T)\mathbf{1}_{[D_T^d \neq 0]} + C_o(a+D_T^d, T)$
$\quad + \mathbb{P}(a+D_T^d, b)v_a \mathbf{1}_{[D_T^{a*}=0]}$
$D_T^{d*}(a,b) := \arg\min V_T^d(D_T^d|a,b)$

Iteration Steps:
Repeat for $t = T-1, T-2, \ldots, 1$
for $(a,b) \in \{\psi_0, \ldots, M\} \times \{\lambda_0, \ldots, T+\lambda_0\}$
Compute and Store
$V_t^a(D_t^a|a,b) = \max_{D_t^a \in \{0,1\}} [\mathbb{P}(a,b)v^a - C_a(b,t)]\mathbf{1}_{[D_t^a=0]}$
$\quad + \delta_a V_{t+1}^a(a + D_{t+1}^{d*}(a, b+1-D_t^a), b+1-D_t^a)$
$D_t^{a*}(a,b) = \arg\max V_t^a(D_t^a|a,b)$
$V_t^d(D_t^d|a,b) = \min_{D_t^d: D_t^d + a \le M} C_d(a+D_t^d, t)\mathbf{1}_{[D_t^d \neq 0]} + C_o(a+D_t^d, t) + \mathbb{P}(a+D_t^d, b)v_a \mathbf{1}_{[D_t^{a*}=0]}$
$\quad + \delta_d V_{t+1}^d(a+D_t^d, b + [1 - D_t^a(a+D_t^d, b+1-D_t^{a*}(a+D_t^d, b))])$
$D_t^{d*}(a,b) = \arg\min V_t^d(D_t^d|a,b)$

Recovering the Optimal Solution:
for $t = 1, 2, \ldots, T$
$\psi_t = \psi_{t-1} + D_t^{d*}(\psi_{t-1}, \lambda_{t-1})$
$\lambda_t = \lambda_{t-1} + D_t^{a*}(\psi_t, \lambda_{t-1})$

Figure 2. Summary of the dynamic programming algorithm.

NUMERICAL ILLUSTRATION

Base Case Model

To illustrate how the solution for this game behaves, we consider a numerical illustration. For this purpose, we have to choose specific functional forms for the probability success and cost functions. In particular, we select

$$\mathbb{P}(\psi,\lambda) = \frac{\lambda}{\lambda+\psi}, \quad C_d(\psi,t) = \frac{b_d\psi}{t},$$
$$C_o(\psi,t) = \frac{b_o\psi}{t}, \quad C_a(\lambda,t) = \frac{b_a}{t\lambda},$$

where $\mathbb{P}$ is decreasing in ψ and increasing in λ, C_d and C_o are increasing in ψ and decreasing in t, C_a is decreasing in t and λ, and the scaling constants b_a, b_d, and b_o are positive. For our baseline values, we select: $\psi_0 = \lambda_0 = 1$, $v_a = v_d = v = 10$, $T = M = 20$, $\delta_a = \delta_d = 0.9$, $b_a = b_d = b_o = 1$.

The selection of the appropriate cost functions can be at times tricky and in many instances, they may not necessarily be well-behaved. For example, these functions may not necessarily be strictly monotonic or continuous in their parameters. In addition, several realistic alternatives may also be possible for functions such as the probability success function $\mathbb{P}$. In selecting these functions, one important criterion should be the ability of these functional forms to capture many essential features in the dynamics between attacker and defender. For example, the notion of an ''adaptive adversary'' dilemma should not be neglected in the choice of the probability function. The adaptive adversary phenomenon refers to the fact that the relative effectiveness of a technology may diminish over time as the opponent accumulates or invests in technology as well. By choosing a function that has the appropriate concavity that we mentioned earlier, we can capture part of this reality, where a diminishing level of effectiveness can be observed as the parameters λ and ψ increase.

The optimal solution for the baseline case is provided in the first plot (subplot a1) in

Figure 3, where over time (x-axis), we plot the optimal level of technology (y-axis) for each player. Here, changes in y represent technological improvements for a player, whereas a flat region represents retention of technology for the defender, and an attack for the attacker. In this plot, we see that the optimal Nash strategy for the attacker is to attack in almost every period (except period 15, when the attacker should accumulate to increase the probability of success), whereas the defender invests only in certain periods (periods 2, 3, 9, 15 and 16).

Under closer scrutiny, the decision of how much to invest in a technology is not as surprising. Because the cost functions chosen are steep in the sense that they have exponential tails, intuition suggests that the most advanced technologies will be costly relative to the benefits at early stages. Thus, one could predict that the decision on how much to invest should be relatively small, i.e., the defender should invest small amounts over time. In the unlikely case that the cost functions are not that steep and the overall costs are relatively low, the amount to invest in new technologies should be relatively large at some early period. However, it is difficult for a decision maker to predict the optimal time to invest without performing calculations. One possible reason for this is that most individuals would have a hard time reconciling the tradeoffs between the cost of adopting new technologies at different periods of time and the net benefits the new technologies would provide if an attack were made—a phenomenon that is supported by research

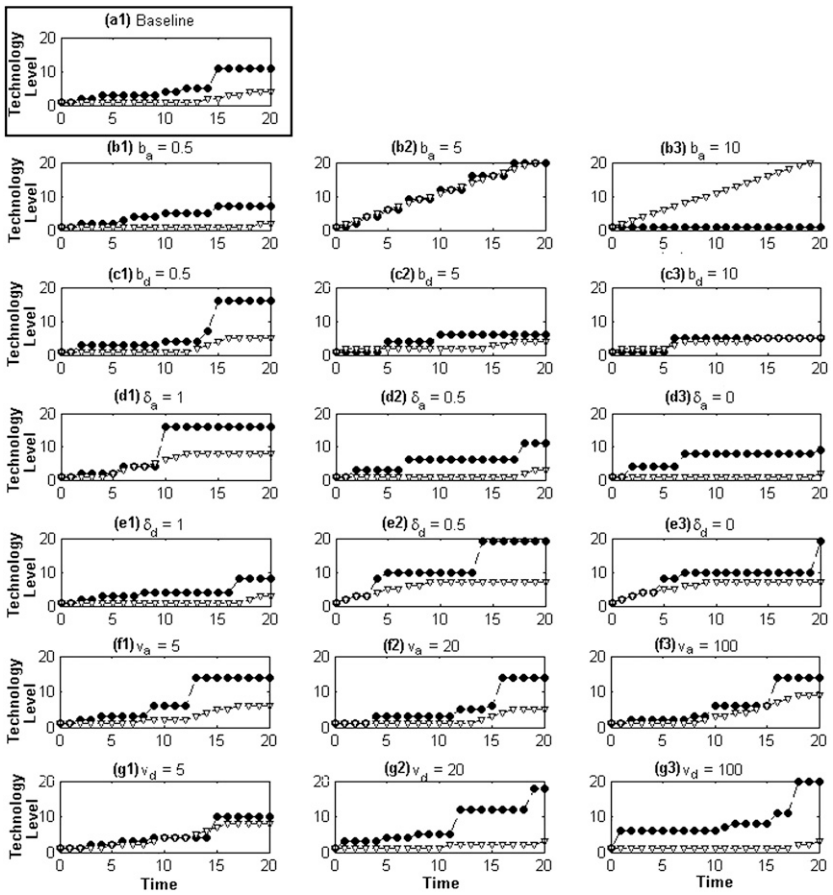

Figure 3. Sensitivity analysis for $T = 20$ when parameters for the defender (∇) and attacker ($\bullet$) are made to vary.

on behavioral decision making on time inconsistency and discounting (e.g., Shane et al., 2002).

Figure 4 shows the policy function D_t^a and D_t^d for this problem at several time periods. We note that the region on the x-axis from $t + 1$ to T for various periods t is empty because these are unreachable states during the first t periods of the game. In addition, we note that the policy function takes on much more values for the attacker than for the defender given the size of their action spaces. The defender can improve up to level M at any period, whereas the attacker can only choose to attack or accumulate at every period.

One quick observation here is that the policy function does not seem to be well-behaved even with the use of simple functional forms for the cost and probability success function. This is driven by the discrete nature of the action space but more so by the discontinuous nature of the objective function, making it difficult to predict over what regions it would be necessary to invest in a dynamic game. When the appropriate functions for the cost and success functions are less well-behaved (e.g., noncontinuous, nonmonotonic), which arises naturally in settings where the effectiveness of a technology set is not necessarily strictly increasing or perfectly correlated with respect to the level and the cost, the prediction of the optimal behavior becomes more difficult. For this reason, we believe, here lies the contribution of using such a model.

Moreover, we show in the next section that using even simple reasonable heuristics for either party could lead to gross inefficiencies, which further strengthens the argument for considering such models in policy formulation, analysis, and decision making.

Sensitivity Analysis

Compared to a single-period model, changes in parameters for a multiperiod model could potentially reverberate throughout the entire time horizon. To have a better understanding, we illustrate how the optimal state space values (ψ_t^*, λ_t^*) change for several parameter values. The optimal solution (D_t^{d*}, D_t^{a*}) can easily be recovered by looking at the difference in the value of the state space variables at every time period. Subfigures in rows (b) to (g) of Figure 3 show three instances for the parameters b_a, b_d, δ_a, δ_d, v_a, and v_d, which can be compared to the baseline case located at the top row of the same figure.

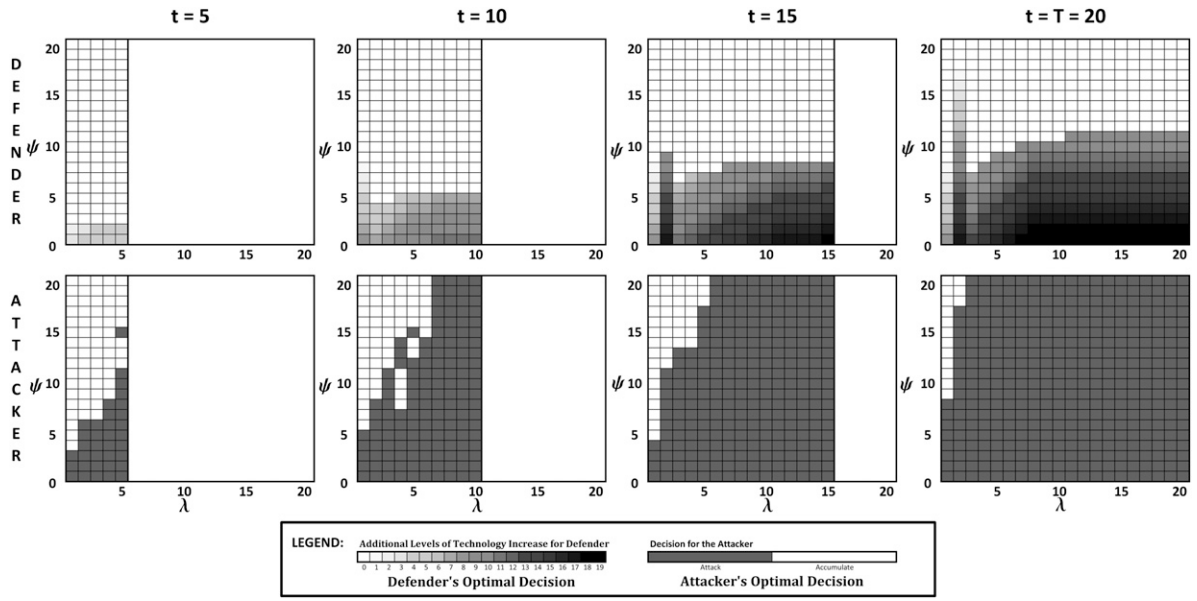

Figure 4. Policy function for the attacker and defender at different periods ($t = 5, 10, 15, 20$).

- *Cost of an attack* (b_a). By adjusting a multiplicative constant associated with the cost of implementing an attack, we analyze how the attacker's (and consequently the defender's) optimal behavior changes. As the cost of launching attacks increases, the attacker will reduce the frequency of attacks. In the baseline case, it is optimal for the attacker to attack in almost every period. As the cost increases, the number of attacks decreases in favor of having a higher level of accumulation to improve the success rate of an attack. In particular, we note that most of the accumulation occurs early on and the attacker behaves similarly to the baseline case in the latter stages. On the side of the defender, the impact of such a change is that the defender is forced to invest in technology to cope with the attacker's level of accumulated resources/technology.
- *Cost of implementing a better technology* (b_d). When the cost of adopting a new technology increases, we observe the defender to be more careful in choosing when to adopt a higher level of technology. We note that when the cost becomes too high, the defender simply chooses to maintain the current level of technology and hopes that the attacks become unsuccessful. Here, the defender expects to absorb the expected cost of an attack since the cost of implementing a higher technology outweighs the benefits from the reduction in the expected damage from an attack.
- *Discount factor for the attacker and defender* (δ_a and δ_d). When the discount factor for the attacker increases (i.e., less discounting), we notice that the attacker would choose to accumulate more in early periods. This happens because the value of future attacks will be potentially higher since the level of accumulation would increase and the net contribution of these future attacks would be higher at the present. Conversely, we observe that if the discount factor is low enough (i.e., discounting is significant), the attacker would choose not to accumulate and simply keep on attacking at every possible instance, because foregoing an attack for a higher chance of success in a future period is less attractive since the discounted value of a future attack is now significantly reduced.

  For the defender, the timing for the adoption of newer technology also depends on the discount factor. The defender would delay investing in newer technologies when δ_d is high and adopts immediately when δ_d is low. In the latter case, the attacker correspondingly acts by accumulating more than the case when discounting is low since he knows that the defender can easily move to a higher level of technology in latter periods since large jumps in technology levels are significantly discounted.
- *Value of a successful attack for the attacker and defender* (v_a and v_d). When a player values an attack much more than the other, the optimal decision would be to accumulate more (invest more in technology) since a single successful attack now has a greater effect on the value function. Interestingly, since the valuation of the opponent is not changed, we see in plots (f1) to (g3) in Figure 3 that the opponent's change in strategy is not significantly different in form from the baseline. The rationale here is that the player with an increased valuation would do whatever it takes to have or avoid a successful attack.

Aside from varying the parameters, we note that certain scenarios are interesting in that they lead to certain phenomena studied in the technology innovation literature. For example, one interesting phenomenon that occurs in certain settings is the notion of a technology competition or arms race (see Anderton (1989) for a review). For example, in Figure 3 (b2) the two players accumulate technology to improve their chances of successfully defending and attacking for intermediate values of b_a. This phenomenon ceases when the cost of attacking is either too low or too high because in these two extremes the attacker will clearly either just attack or just accumulate. It is in the intermittent case of accumulating and attacking that the defender needs to improve resources. Despite the fact that an attacker tries to accumulate resources in this setting, the option to launch an attack is open to the attacker if it seems that there is a reasonable chance of success even if the cost of an attack is not that low.

With respect to the other parameters, we see that the technology race is not really clearly seen

because what usually happens is that the for a significant change in parameters, only one of the parties would significantly try to improve its level of technology/accumulation. For example, when v_d is high, the defender will invest more in technology as the expected cost of an attack starts to outweigh the cost of adopting and implementing a newer technology. The same is true for v_a with the caveat that at much higher levels of v_a the defender will also somewhat increase his level of technology. The attacker would then start to gradually accumulate resources since the value of a successful attack is higher. He tries to achieve it at higher levels of v_a when multiple tries as well as improving his chances through accumulation.

We note that the results in this example are specific to the functional forms employed but the contribution of this paper lies in the use of the suggested framework and the illustration of some types of analysis that could be done. In addition, we believe that many of the insights that we discover from these numerical examples will hold for a wide range of function forms that have similar monotonic properties with respect to model parameters ψ, λ, and t.

Analysis of Some Attacker and Defender Heuristics

The importance of timing investment decisions manifests most clearly when the savings are substantial relative to the decision-making effort. We perform some benchmark analysis by comparing the optimal strategy to some potentially reasonable heuristic strategies that a defender may employ. We measure the difference in the optimal costs when the defender employs these heuristics and the attacker responds optimally to the defender's choices in investment.

First, we consider the case in which the defender chooses to invest immediately in the most effective albeit most expensive technology, which we refer to as Policy D1. Typically, this type of behavior can be justified when a decision maker who cannot correctly value the costs and benefits of adopting a technology is mistakenly led to believe that the costs of adopting the best safeguards would always outweigh the cost associated with a successful attack.

A second reasonable strategy (hereafter referred to as Policy D2) is for the defender to simply invest slowly by choosing to invest in one level of technology at a time, i.e., increase the level of technology by one unit per period. Unlike Policy D1, this one yields a more modest approach to investing, which could be viewed by the public as an act where the defender is constantly doing something to protect its internal security.

Figure 5 shows the percentage increase in the optimal solution when comparing these two policies when the time horizon is made to vary. For example, here we see that when the players are engaged in a one-period game, Policy D1 yields a 566% increase in expected costs for the defender while Policy D2 yields

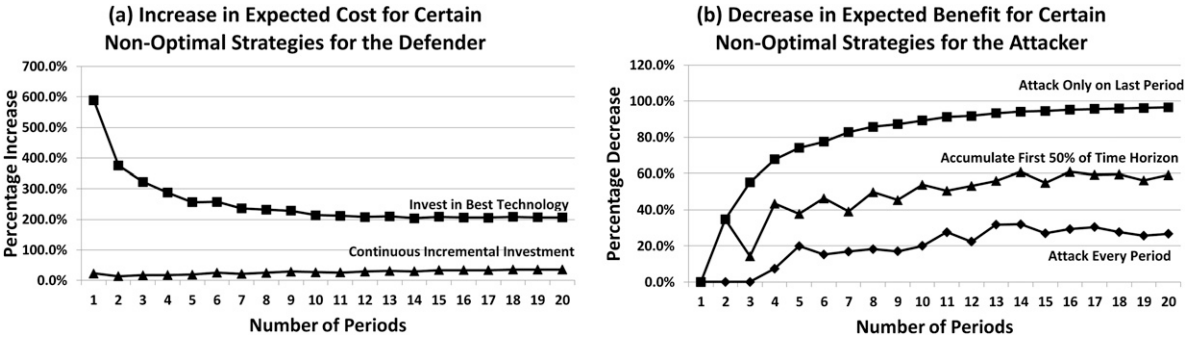

Figure 5. (Left) Percentage increase in expected cost under varying time horizon (T) when (a) investing one unit of technology per period and (b) choosing to invest in the most expensive technology immediately. (Right) Percentage decrease in expected benefit under varying time horizon (T) when attacker accumulates nonoptimally for varying amounts of time and attacks every period thereafter.

a 22% increase. In both instances, these two strategies provide substantially higher costs for the defender.

As the time horizon increases, we observe that Policy D1's inefficiency decreases, because for longer time horizons, the cost of adopting a higher level of technology early on can be ''spread out'' over the longer horizon since the defender can reap additional benefits in the reduced chance of having a successful attack over time. On the other hand, Policy D2 slightly increases over time because there are more periods for potential discrepancies between the optimal solution and the solution that Policy D2 provides. Recalling the structure of the optimal solution in the baseline case, we see that there are now more periods in which a moderate investment is no longer considered optimal. In fact, for several periods, the defender should simply stay at its current level rather than continually investing in newer technologies. Although this may be slightly increasing over time, Policy D2 still yields a better solution than Policy D1 for all values of T considered, which yields an increase in expected costs of somewhere between 170% and 570%. For Policy D2, the increase in expected costs is more modest with a range of 13% to 36%. When discussing large sums of money for implementing this defensive strategy, we see that using such a model could still yield substantial cost savings compared to implementing some nonoptimal heuristics that can be thought to be natural or reasonable. Similarly, we can examine how the attacker can potentially be at a disadvantage when acting in a nonoptimal fashion while the defender reacts optimally to the attacker's choices. In particular, we examine the case in which the attacker accumulates for the first few periods and then simply attacks until the end. A potential justification for this is that the attacker may want to first attain a certain level of technology before attacking to improve the chances of success.

Figure 5 also shows the effect on the optimal value function for the attacker for different levels/periods of accumulation (0%, 50%, and almost 100% of the time horizon and hereafter referring to them as Policies A1, A2, and A3, respectively). Although the plots are nonmonotonic due to the discrete nature of the time horizon, we see that their general pattern indicates that as the time horizon increases, the percentage decrease generally increases over time. We also note that the case where the attacker attacks every period is best among the heuristics, which is not surprising since the optimal SPNE strategy in the baseline case is for him to attack in almost every period. Using the numeric example that we have, we can see that for the longer time horizons, the decrease in the objective value function can be substantial (~20–30%). We note that these patterns remain persistent when both players play the heuristic strategies discussed. In this case, both players will experience significant changes in their optimal cost or benefit value that is similar in pattern and generally worse in value than the plots in Figure 5.

Finally, we conduct some sensitivity analyses on the parameters and see how well the heuristics presented perform relative to the optimal strategy. For the defender, we see from Figure 6 that the difference between the performance of the heuristics is relatively stable for most parameters (v_a, b_a, δ_a, and δ_d). However, for v_d and b_d, we observe that the performance is more or less the same except for extremely large values. This tells us that for most reasonable values of the parameters, there is a substantial cost savings for the defender when using the optimal strategy compared to the heuristics analyzed. Similar plots are provided for the attacker in Figure 7. Here, we see that the general ranking between the heuristics remains the same for all parameters with a few minor exceptions—e.g., see Policies A2 and A3 in plots (a), (c), and (e) of Figure 7.

CONCLUSION

Conventional wisdom teaches that in many aspects of life, timing is everything. This definitely holds in the area of multiperiod attacker-defender games where improper timing in investments toward defense, preparation, and security can have a serious economic consequence for both attacker and defender. Effective planning requires the careful discernment of when and how much to invest at every period, taking into account that the interaction

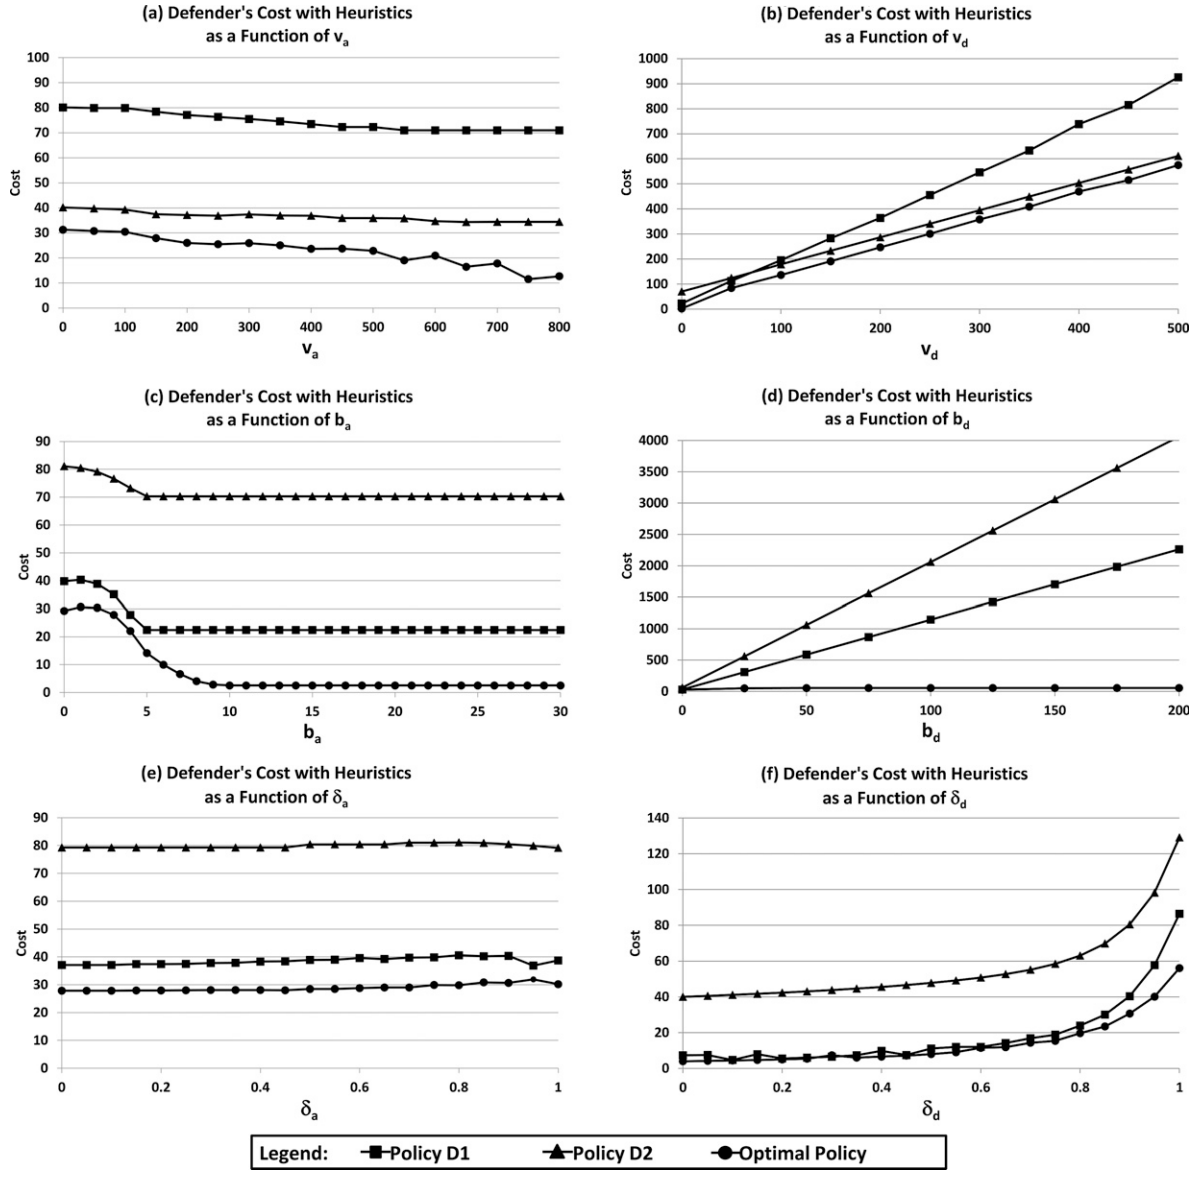

Figure 6. Sensitivity analysis for the defender heuristics.

between an attacker and defender goes beyond a single period.

The research in the area of technology adoption and competition has been well-established in the economics and decision analysis literatures. However, this is not true in the security games and military operations research literature. This paper fills that gap by presenting a multiperiod dynamic game between an attacker and defender, which captures the tradeoffs between immediately implementing important technology improvements and the cost and risk associated with potential threats and attacks. We provide a modeling framework that could be used for planning and understanding the costs and benefits of adopting new technologies in the context of homeland security. Although these models can capture the complex interactions between the two parties, the price for this is the lack of analytic tractability even for simple nontrivial functional forms, which nowadays is easily addressed

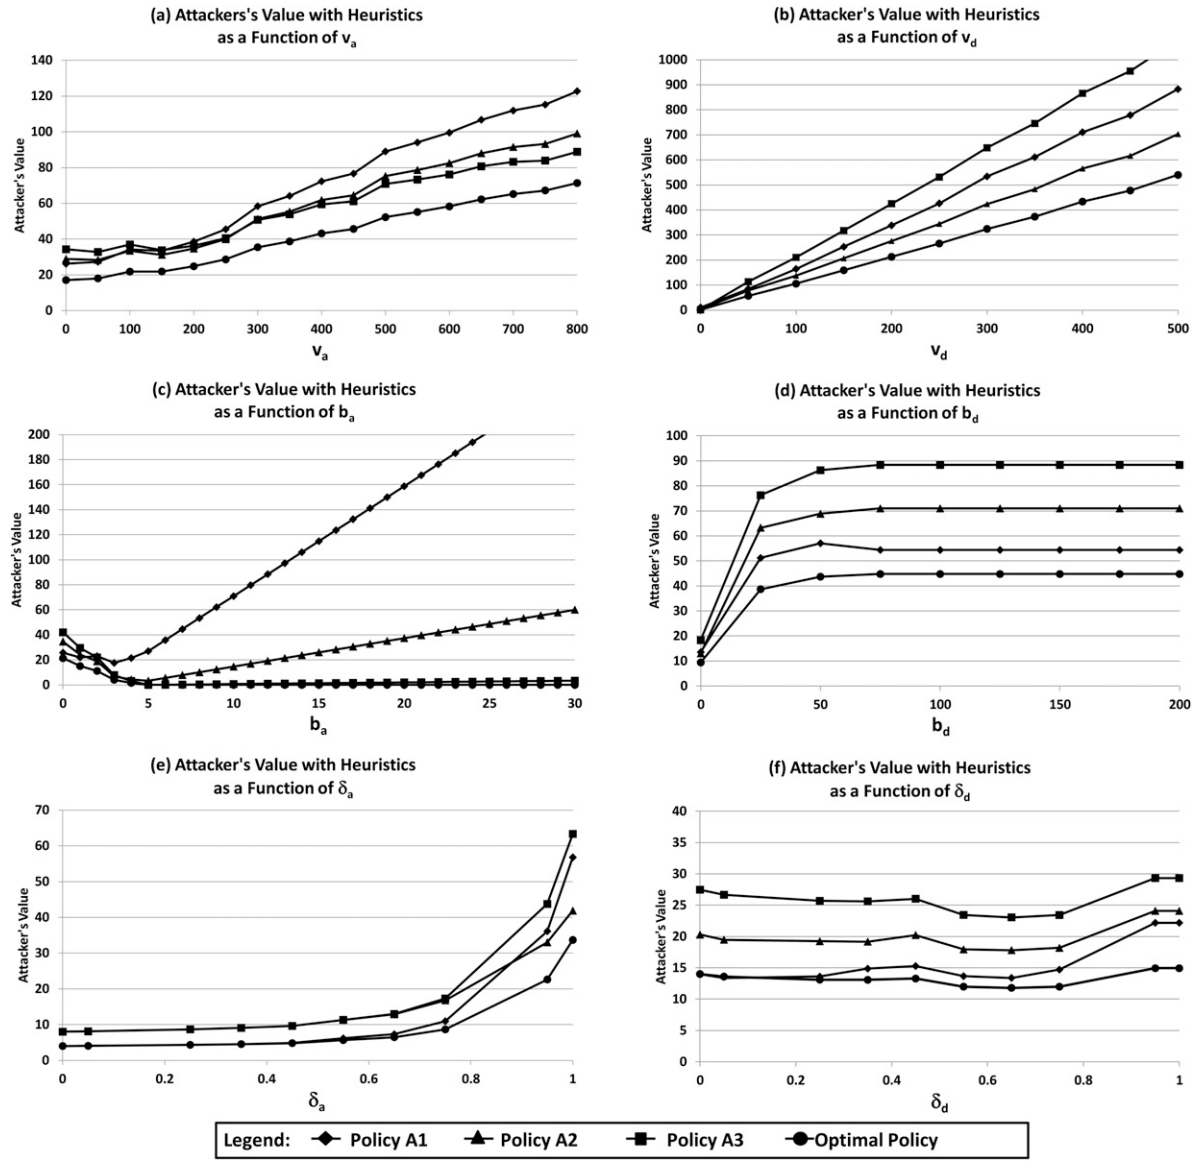

Figure 7. Sensitivity analysis for the attacker heuristics.

by the rapid development in computational methods.

By using some simple examples, we see that despite the fact that some simple heuristics seem to perform reasonably well, the costs savings in correctly timing investments in security could be significant considering the current spending levels of government agencies. Although the paper focuses on a specific set of functional forms and numerical examples, we believe that the idea and framework is still general enough to be applied to specific technology adoption problems of a government agency or an international police/security organization. We hope that this type of analysis can provide operations research professionals in the military and homeland security areas insightful justification to provide evidence either for or against proposals to policy and decision makers related to the planning, scheduling, and acquisition of portfolios of defense technologies.

Although we have addressed an important problem, there are other challenges and interesting questions that can be analyzed going forward. First, the incorporation of incomplete information in this setting may be applicable in many settings. For example, each player may not have full knowledge of the technologies being adopted by each party. Instead, each party may have only partial information, similar to what was done by Zhuang et al. (2010). Another interesting avenue to pursue is the incorporation of investment opportunities for technology with uncertain quality or payoff. This issue deals with the common problem that many technologies are being proposed for use and need funding, but the effectiveness and potential cost of developing them are typically uncertain. We believe that embedding this additional layer of uncertainty will yield interesting insights with respect to the potential tradeoffs a defender would make to allocate scarce resources between investing them for protection in the immediate versus the distant future.

Another area of future work is to see how the dynamics of an attacker-defender game would change when multiple parties work against an attacker (e.g., several nations or government agencies) or perhaps multiple parties working against a defender (e.g., multiple terrorist groups), or both. Here, one can investigate how the sharing of technology could improve the system or make it worse for one or both parties given the asymmetry in information, resources, and technology of parties involved.

ACKNOWLEDGEMENTS

This research was partially supported by the United States Department of Homeland Security (DHS) through the National Center for Risk and Economic Analysis of Terrorism Events (CREATE) under award number 2010-ST-061-RE0001. This research was also partially supported by the United States National Science Foundation (NSF) under award number 1200899. However, any opinions, findings, and conclusions or recommendations in this document are those of the authors and do not necessarily reflect views of the DHS, CREATE, or NSF. We thank three anonymous referees for their helpful comments and Marie Catalano for editorial help. Corresponding authors are Victor Jose and Jun Zhuang. The authors assume responsibility for any errors.

REFERENCES

Anderton, C. H. 1989. Arms Race Modeling: Problems and Prospects, *Journal of Conflict Resolution*, Vol 33, No 2, 346–387.

Bandyopadhyay, S., and Sandler, T. 2011. The Interplay between Preemptive and Defensive Counterterrorism Measures: A Two-Stage Game, *Economica*, Vol 78, No 311, 546–564.

Chambers, C., and Kouvelis, P. 2003. Competition, Learning and Investment in New Technology, *IIE Transactions*, Vol 35, No 9, 863–878.

Dresher, M. 1961. *Games of Strategy: Theory and Applications.* Prentice-Hall.

Faria, J. R. 2003. Terror Cycles, *Studies in Nonlinear Dynamics and Econometrics*, Vol 7, No 1, 1–9.

Feichtinger, G., Novak, A. J. 2008. Terror and Counterterror Operations: Differential Game with Cyclical Nash Solution, *Journal of Optimization Theory and Applications*, Vol 139, No 3, 541–556.

Golalikhani, M., Zhuang, J. 2011. Modeling Arbitrary Layers of Continuous-Level Defenses in Facing with Strategic Attackers, *Risk Analysis*, Vol 31, No 4, 533–547.

Hausken, K., and Zhuang, J. 2011a. Governments' and Terrorists' Defense and Attack in a T-period Game, *Decision Analysis*, Vol 8, No 1, 46–70.

Hausken, K., and Zhuang, J. 2011b. Defending against a Terrorist who Accumulates Resources, *Military Operations Research*, Vol 16, No 1, 21–39.

Kardes, E., and Hall, R. 2005. Survey of Literature on Strategic Decision Making in the Presence of Adversaries. CREATE Working Paper 05-006, University of Southern California.

Powell, R. 2007. Defending against Terrorist Attacks with Limited Resources, *American*

Political Science Review, Vol 101, No 3, 527–541.

Reinganum, J. F. 1981. On the Diffusion of New Technology: A Game Theoretic Approach, *The Review of Economic Studies*, Vol 48, No 3, 395–405.

Sandler, T., and Arce, D. G. 2003. Terrorism & Game Theory, *Simulation & Gaming*, Vol 34, No 3, 319–337.

Shane, F., Loewenstein, G., and O'Donoghue, T. 2002. Time Discounting and Time Preference: A Critical Review, *Journal of Economic Literature*, Vol 40, No 2, 351–401.

Zhuang, J., Bier, V. M., and Alagoz, O. 2010. Modeling Secrecy and Deception in a Multiple-Period Attacker-Defender Signaling Game, *European Journal of Operational Research*, Vol 203, No 2, 409–418.