A Game Theoretic Model for Information Security Incorporating Success and Evasion

Mike O'Leary
Department of Mathematics, Towson University
Towson, MD 21252, USA
molearytowson.edu

Andrew Engel
SAS Institute
10188 Telesis Court, Suite 200
San Diego, CA 92121, USA
andrew.engelsas.com

Abstract

We present a game theoretic model of the interaction between an attacker and a system administrator. The model is general and can be applied at both the strategic as well as the operational level. It accounts separately for the possibility of an attack succeeded as well as the possibility that the attack evades detection by the system administrator. The main result is that, for a general class of models, the optimal strategy for the defender can be chosen independently of the attacker.

Keywords

Game theory, Security, Theory

1 Introduction

What is the "optimal" method to defend a computer information system from an attack? An information system in this context can represent a single computer, a database, or even a geographically dispersed network. There are many technical things that can be done to improve security of an information system. Possibilities include firewalling systems, strengthening access controls, or implementing and reviewing the data provided by an intrusion detection system. However each of these comes at a cost, in terms of money, personnel, time, and training. The "optimal" solution then gives the best balance of risk versus cost. This is our question: What is the right level of information security for an enterprise or organization?

In this paper, we present a general framework to analyze these decisions using ideas from game theory. Game theory has long been used to analyze strategies for classical games like chess [37], and it is also successful in analyzing games with incomplete information like poker [7]. Game theory has been used to solve diverse problems of practical interest which include increasing revenue from the sale of the spectrum for the Federal Communications Commission [29], modeling military applications from disarmament [6] and national security [10] to warfare [23] and air combat [19], and determining pricing for the telecommunications industry [39].

In this paper, we create game theoretic models that describe the interaction between an attacker and a defender. These models can be interpreted at two basic levels: a strategic level and at an operational level. At the strategic level, the main question for an organization is to determine the optimal investment in information security. Too little, and one is open to loss and liability; too much and money has been wasted. We model this as a game between an attacker and a defender. The attacker wants to compromise the security of the information system while the defender attempts to construct policies and deploy equipment, people, and resources to prevent such a breach. One can then ask what is the optimal balance between cost and capability for the defender.

A similar set of questions obtains at the operational level. Here a network or security administrator needs to protect a particular resource; for example a database system or a web server. How should the administrator respond to the alerts of the intrusion detection system? Different responses have different effectiveness and costs.

Despite the different scales of these questions, they have a number of common elements, and these common elements are the focus of this work. We propose a general model where each attack and each defensive posture have three associated elements

- the chance an attack succeeds,
- the chance an attack evades detection, and
- the costs necessary to either launch the attack or to adopt that defensive posture.

To keep the model as generally applicable as possible, these are the only three features that we model.

The chance of success for a given attack depends on both the particular attack and the particular defense chosen. We assume that each attack has a strength which represents the likelihood that the attack would succeed against an average target. Similarly, we also assume that each defensive posture has a strength which gives the likelihood that an average attack would succeed against that posture. Then the probability that a given (non-average) attack would succeed agains a given (non-average) defensive posture is modeled. We consider both linear and non-linear models for this process, and we show that, for attacks and defensive postures near the average, that the linear model is a good approximation for a large class of nonlinear models.

We handle the chances that an attack evades detection in the same fashion; we assume that each attack has an evasiveness that describes the probability that the attack would go undetected against an average target. Similarly, we associate an evasiveness to each defensive posture that gives the likelihood that this defense would allow an average attack to evade detection.

We also assume that each attack or defensive posture has an associated cost. This cost can be financial- modeling costs of time and equipment, for example. It also includes costs to an organization in lost productivity caused by more stringent security requirements.

We explicitly assume that both the attacker and the defender have multiple objectives in this game. The attacker wants the attack to succeed, to evade detection and to do so at minimal cost. Similarly, the defender wants the attack to fail, to detect the attack, and to do so at minimal cost. As a consequence, we model the interaction as a multi-objective game. Each player is trying to maximize a payoff function that is a linear combination of the success probability, evasion probability and cost. The precise linear combination is chosen by each player to reflect the relative weight that they place on each component. This enables our model to examine how an organization might balance investments in protective tools, like firewalls and encryption, versus investments in detection tools, like intrusion detection systems.

Our main result is that, for attacks and defenses near the average, the optimal strategy for each player can be determined by analyzing the relative weights assigned to the various components by that player; the objectives or capabilities play no role. As a corollary, we also show that investments in detection tools do not deter an attacker bent on attacking an organization into using more evasive but less powerful attacks.

2 Previous Work

Some of the first efforts to apply game theory to computer security were [20, 21] who took a very high level approach to the question. Alpcan and Basar in [3] modeled the interaction between an attacker and an intrusion detection system. In their work, the attacker chooses either a system to attack or to not attack at all, while the defending intrusion detection system chooses either to set off an alarm or not. In particular, this model explicitly allows for the possibility of false alarms. The model does not account for attacks of differing likelihoods of success however. The authors found that, in this case, the Nash equilibrium solutions depend only on the cost function of their opponent. This work was extended in [4], including allowing multiple attackers.

Lye and Wing [28] present a model at the operational level. They examine attacks on a web server and on an ftp server, and include detailed maps of the possible states of both the attacker and the administrator, together with estimates of the rewards and transition probabilities for each state. The authors then compute the corresponding Nash equilibria of the game. This approach has its share of difficulties however. First is the difficulty of determining what are the rewards and transition probabilities for a given state. There is also the problem of completeness; since this low level model is meant to account for all of the possible actions of both the defender and the administrator, any state that can occur but is not modelled will be a source of error.

Liu, Zang and Yu [27] focus more on analyzing the attackers intent, objective and strategies, rather than on finding optimal defender strategies. They assume that the attacker's behavior is rational, and that it can be modeled by examining incentives and costs. Incentives are measured, for example, by comparing the degraded state of a network undergoing a DoS attack to the undegraded state. Costs include financial costs as well as constraints and measures of risk. They too use a low level model of attacks; for example they split the Mstream DDoS attack into five phases for each zombie. Their emphasis is on techniques to learn what the attacker's strategy might be. Because of the low level nature of the model, it is difficult to draw general conclusions. For example, they present a nice application of their method to a DDoS attack, but find 42 different Nash equilibria.

A number of different authors have examined the difficult problem of modeling attacks at an operational level. These include [31, 32, 33], who examine alerts to construct potential attacker strategies. Sheyner *et.al.* [36] use symbolic model checking algorithms to generate attack graphs. Cuppens and Miége [13] describe a method to correlate different alerts from different intrusion detection systems. Templeton and Levitt [38] describe an attack language called JIGSAW that describes and models attacks.

Bistarelli, Fioravanti, and Peretti [8] examine "defense trees" as the defensive counterpart to attack trees and describe how to analyze the financial costs of various countermeasures, including examining the resulting return on investment of different defensive techniques.

Kotenko and Stepashkin [25] model attacks with the goal of improving vulnerability assessment; we also mention Liu and Li [26] who use game theory to predict potential attacks.

At the strategic level, we have the analysis of Gordon and Loeb [16]. In their work, they construct a model to determine the optimal investment in security to protect a given set of information. They use two different models to calculate the expected benefit of an investment in information security. Additional models were studied by Hausken [22] and the results extended by Willemson [40]. For additional analysis of the economic costs of security at the strategic level, see also [9, 11, 12, 15, 17, 18].

Neubauer, Stummer and Weippl [30] have constructed a mathematical model to evaluate the costs and expected benefits of investments in security. In particular, when examining losses they look at both monetary losses as well as losses to the corporate image, while when examining defensive postures, they account for acceptance costs, setup manpower, setup time, setup costs, and running costs.

Other applications of game theory to computer security include [1, 2, 35].

3 Review of Game Theory

Because we rely heavily on game theory in what follows, we shall briefly recap some of the salient features of the theory that we require in what follows. Solid introductions to the subject can be found in [5, 14, 34].

Consider a two player game between players A and D. It has two components; first is a pair of sets S_A, S_D, called the *pure strategy* sets for each player. They represent the collection of actions each player is allowed to make. The second component is a pair of functions π_A, π_D, called *payoff functions*. Given a pair of strategy choices, one for each player, the payoff function π_A represents the value of that pair of choices to player A; the function π_D does does the same thing for player D.

If the strategy sets S_A and S_D of both players are finite, then the result is called a bimatrix game. Indeed, if player A has strategies $\{a_1, a_2, \dots, a_m\}$ while player D has strategies $\{d_1, d_2, \dots, d_n\}$, then we can construct the pair of matrices

$$\pi_A = \begin{bmatrix} \pi_A(a_1,d_1) & \pi_A(a_1,d_2) & \cdots & \pi_A(a_1,d_n) \\ \pi_A(a_2,d_1) & \pi_A(a_2,d_2) & \cdots & \pi_A(a_2,d_n) \\ \vdots & \vdots & \ddots & \vdots \\ \pi_A(a_m,d_1) & \pi_A(a_m,d_2) & \cdots & \pi_A(a_m,d_n) \end{bmatrix}$$

$$\pi_D = \begin{bmatrix} \pi_D(a_1,d_1) & \pi_D(a_1,d_2) & \cdots & \pi_D(a_1,d_n) \\ \pi_D(a_2,d_1) & \pi_D(a_2,d_2) & \cdots & \pi_D(a_2,d_n) \\ \vdots & \vdots & \ddots & \vdots \\ \pi_D(a_m,d_1) & \pi_D(a_m,d_2) & \cdots & \pi_D(a_m,d_n) \end{bmatrix}$$

that describe the payoffs for the game.

In the game, both players act to maximize their own payoff function. However, they do not necessarily try to minimize the payoff of their opponent.

In many games, players do not want to act in a deterministic fashion, as this behavior may be anticipated by their opponent. Instead, many times the best choice is to choose different pure strategies randomly. A mixed strategy is a probability density function defined on a pure strategy set. Let x be a mixed strategy for player A, and y a mixed strategy for player D; we then let $\pi_A(x,y)$ and $\pi_D(x,y)$ represent the expected payoffs to players A and D. In the simple case of a bimatrix game, the mixed strategies are vectors $\mathbf{x}$ and $\mathbf{y}$ with $\sum x_i = \sum y_j = 1$, where x_i is the probability of player A choosing strategy i and y_j is the probability that player D chooses strategy j. Further, the payoff functions for the mixed strategies can be found by matrix multiplication

$$\pi_A(\mathbf{x},\mathbf{y}) = \mathbf{x}^T \pi_A \mathbf{y},$$
$$\pi_D(\mathbf{x},\mathbf{y}) = \mathbf{x}^T \pi_D \mathbf{y}.$$

A Nash equilibrium is a pair of mixed strategies $(\hat{x},\hat{y})$ so that neither side can unilaterally improve their own expected payoff by choosing a different strategy. In particular, for any different strategy x for player A we have

$$\pi_A(\hat{x},\hat{y}) \geq \pi_A(x,\hat{y}).$$

Similarly, for any different strategy y for player D, we have

$$\pi_D(\hat{x},\hat{y}) \geq \pi_D(\hat{x},y).$$

Every bimatrix game has a Nash equilibrium in mixed strategies, as do more general games with some reasonable assumptions on the strategy sets S_A, S_D, and the payoff functions π_A, π_D [14, Thm. 1.1, 1.2]. Although the existence of a Nash equilibrium is guaranteed, the Nash equilibrium is not necessarily unique. It should also be noted that other notions of equilibrium solution exist; c.f. [34, Chp. VII].

4 The Model

4.1 Features

As described in the introduction, we would like to create an abstract model with a wide range of validity. To that end, we suppose only that each attack $\mathbf{a}$ can be characterized by exactly two components $\mathbf{a} = (a_s, a_e)$. The number a_s, the attack's strength, represents the probability that the attack, when launched against an "average" target will succeed. The number a_e, the attack's evasiveness, represents the probability that the attack, when launched against an "average" system, will evade detection.

We also assume that there is a function $C_A(\mathbf{a})$ that represents the relative cost to the attacker for performing the attack $\mathbf{a}$. This represents the time, effort, and energy that the attacker needs to expend to implement the attack $\mathbf{a}$.

We also assume that each posture $\mathbf{d}$ that the defender can take can be characterized by two components $\mathbf{d} = (d_s, d_e)$. Here d_s is the strength of the defense, which represents the probability that an "average" attack will succeed while d_e is the probability that an "average" attack will evade detection. Note that increasing values of d_s correspond to increasing chances of success for the attacker, and increasing values of d_e correspond to increasing chances of evasion for the attacker; thus the defender wants these values to be as small as possible.

We also assume that there is a cost function $C_D(\mathbf{d})$ that represents the relative cost to the defender of adopting posture d. This cost includes direct financial costs of equipment as well as the value of the time and work of the network staff. In addition, this also includes any decrease in the usability of the system caused by the use of more stringent security.

Both the attacker and defender have three different objectives. The attacker wants the attack to succeed, to evade detection, and to do so at minimal cost. On the other hand, the defender wants to prevent the success of the attack, to detect the attack, and to do so at minimal cost.

For a given attack $\mathbf{a}$, and defense $\mathbf{d}$, let $P_s(\mathbf{a}, \mathbf{d})$ give the probability that this attack will succeed against this defense; similarly let $P_e(\mathbf{a}, \mathbf{d})$ give the probability that this attack will avoid detection against this defense. The goal of the attacker is to maximize the payoff function

$$\pi_A(\mathbf{a}, \mathbf{d}) = A_s P_s(\mathbf{a}, \mathbf{d}) + A_e P_e(\mathbf{a}, \mathbf{d}) - A_c C_A(\mathbf{a})$$

where A_s, A_e, and A_c are positive weights that give the relative importance of success, evasion, and cost. Similarly, the goal of the defender is to maximize the payoff function

$$\pi_D(\mathbf{a}, \mathbf{d}) = -D_s P_s(\mathbf{a}, \mathbf{d}) - D_e P_e(\mathbf{a}, \mathbf{d}) - D_c C_D(\mathbf{d})$$

where again D_s, D_e and D_c are positive weights.

Because we are interested in the maxima of π_A and π_D rather than their precise values, we can assume without loss of generality that $A_c = D_c = 1$. Then if the cost functions $C_A(\mathbf{a})$ and $C_D(\mathbf{d})$ return the monetary cost of a strategy, then A_s represents the monetary benefit to the attacker of successful attack while D_s is the corresponding monetary loss to the defender of a successful attack. Similarly, A_e represents the monetary benefit of the attack remaining undetected while D_e is the monetary loss of an attack remaining undetected. Of course, we need not assume that either $C_a(\mathbf{a})$ or $C_d(d)$ represent monetary values; in fact for an attacker motivated by notoriety, bravado, or boredom, the use of a monetary measures for $C_a(\mathbf{a})$ and $\pi_A(\mathbf{a}, \mathbf{d})$ are probably inapproptiate.

4.2 Success probabilities

To proceed, we need to construct a model for the probability functions P_s and P_e. To begin, we assume that the success probability depends only on the attack strength and the defenders strength; in other words that it is independent of the corresponding costs and evasiveness. Thus $P_s(\mathbf{a}, \mathbf{d}) = P_s(a_s, d_s)$. Similarly, we assume that the probability of evading detection depends only on a_e and d_e so that $P_e(\mathbf{a}, \mathbf{d}) = P_e(a_e, d_e)$.

Temporarily dropping subscripts for readability, let $P(a, d)$ be the probability that an attack with strength a against a target with defense d will succeed. The definition of a and d tells us the probability of success when either a is the strength of an average attack, or d is the strength of the average defense. However, we do not know what those values are. If we let $\bar{a}$ and $\bar{d}$ denote the strength of the average attack and the average defense, then the definition of a and d then implies that

$$P(a, \bar{d}) = a$$
$$P(\bar{a}, d) = d$$

for any a and d. In particular, setting $a = \bar{a}$ in the above, we see that $\bar{a} = \bar{d}$.

To continue, let us begin by choosing the simplest form for P; namely that it is linear. Suppose that

$$\ell(a, d) = \bar{a} + (a - \bar{a}) + (d - \bar{a});$$

we would like to model P by

$$P(a, d) = \ell(a, d)$$

for all $0 \leq a \leq 1$ and $0 \leq d \leq 1$. However, this is impossible, as this gives us values of $P(a, d) > 1$ for $a + d - \bar{a} > 1$ and $P(a, d) < 0$ for $a + d - \bar{a} < 0$.

Rather than change the simple algebraic form of P, we instead restrict the domain of allowable choices for (a, d). We start by insisting that $0 \leq a + d - \bar{a} \leq 1$. In particular, we are eliminating from consideration scenarios where the attacker applies an attack much stronger than average to a defense that is much weaker than average.

Next, we note that the attacker and the defender choose their strategies independently. As a consequence, if a^* is an allowable attack strategy against a defense d^*, then it must be an allowable attack strategy against any other allowable defense d. In particular, this implies that the set of allowable strategy pairs (a, d) must be rectangular. Combining this with the restriction $0 \leq a + d - \bar{a} \leq 1$, we use the model

$$P(a, d) = \ell(a, d) = \bar{a} + (a - \bar{a}) + (d - \bar{a}) \qquad (1)$$

in the region

$$\mathcal{D} = \{(a, d) : \tfrac{1}{2}\bar{a} \leq a, d \leq \tfrac{1}{2} + \tfrac{1}{2}\bar{a}\}$$

This choice of domain $\mathcal{D}$ implies that almost every choice of attack and defensive posture is neither guaranteed to succeed nor to fail. In fact only the combination $a = d = \bar{a}/2$ is guaranteed to fail and only the combination $a = d = \bar{a}/2 + 1/2$, is guaranteed to succeed.

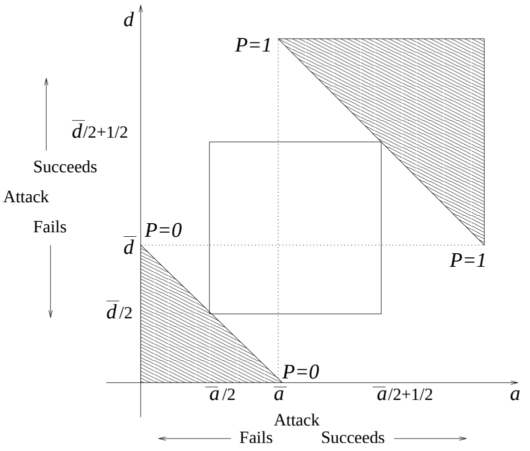

Moreover, $\mathcal{D}$ is the largest rectangular domain where almost every choice of (a, d) is neither assured of success or failure.

Although this description shows how to find the success probability of an attack with strength a against a defense with strength d, the same argument can be applied to find the probability that an attacker will evade detection. As a consequence, we obtain the linear model

$$P_s(a_s, d_s) = \bar{a}_s + (a_s - \bar{a}_s) + (d_s - \bar{a}_s)$$
$$P_e(a_e, d_e) = \bar{a}_e + (a_e - \bar{a}_e) + (d_e - \bar{a}_e)$$

for

$$(a_s, d_s) \in \mathcal{D}_s = \left\{(a, d) : \tfrac{1}{2}\bar{a}_s \leq a, d \leq \tfrac{1}{2} + \tfrac{1}{2}\bar{a}_s\right\} \quad \text{(2a)}$$
$$(a_e, d_e) \in \mathcal{D}_e = \left\{(a, d) : \tfrac{1}{2}\bar{a}_e \leq a, d \leq \tfrac{1}{2} + \tfrac{1}{2}\bar{a}_e\right\} \quad \text{(2b)}$$

where $\bar{a}_s$ is the strength of the average attack and $\bar{a}_e$ is the evasiveness of the average attack.

Before we explore the consequences of this linear model for the probabilities $P(a, d)$, we would like to briefly consider the possibility of a nonlinear form for P. We can, for example, look at choices of P of the form

$$P(a, d) = \ell(a, d) + \sigma(a - \bar{a}, d - \bar{d})$$

where σ is a nonlinear correction term. However, this correction cannot be chosen arbitrarily; there are three significant conditions that a model, linear or nonlinear, must satisfy.

First, because of the definition of a and d, we know that $P(a, d) = \ell(a, d)$ when either $a = \bar{a}$ or $d = \bar{a}$. Further, because P is to represent a probability, we must have $0 \leq P(a, d) \leq 1$ for all (a, d) in the domain of definition of P. Finally, $P(a, d)$ must be monotone non-decreasing in a and in d. Indeed, since $P(a, d)$ is the probability of success, increasing the strength of the attack and keeping the defense constant must not decrease the success probability. As a consequence we must have $\frac{\partial P}{\partial a} \geq 0$ and $\frac{\partial P}{\partial d} \geq 0$.

As an example of the effect of these conditions, we have proven (Proposition 1) that there is no nonlinear quadratic choice of P that meets all of these conditions on the same domain $\mathcal{D}$ that we have used in the linear model. To find a nonlinear quadratic model, we must reduce the size of the domain of definition of P. The domain $\mathcal{D}$ is a square of side length $\frac{1}{2}$, so consider $\mathcal{D}^*_\mu = \{(a, d) : \bar{a} - \frac{\mu}{2}\bar{a} \leq a, d \leq \bar{a} + \frac{\mu}{2}(1 - \bar{a})\}$ which is a square where the side length has been multiplied by the factor μ. We have shown (Proposition 2) that, in this case,

$$|P(a, d) - \ell(a, d)| \leq (1 - \mu)\min\{\bar{a}, 1 - \bar{a}\}.$$

As a consequence, if, for example, we reduce the side length of the domain of definition by 10%, then the difference between the linear model and any nonlinear quadratic model is no more than 10% of the value of an average attack. In particular, any quadratic model for P is either close to the linear model, or has a markedly smaller domain of definition.

Although quadratic choices of P need to be close to the linear model, there are non-trivial choices of P that satisfy all of our conditions (Corollary 3).

4.3 Cost functions

The form of the cost function depends on whether we are using this game to model the operational problem or the strategic problem. In the operational case, the form of the cost function can be based on the actual costs incurred to adopt a given defensive posture or to make a given attack.

The strategic case is more difficult. We can begin with the cost functions of Gordon and Loeb [24] as starting points. That paper contained two models of the relationship between the cost of a defensive posture and the likelihood that an attack on that defensive posture would succeed; they labelled these as type I and type II. Their type

I model yields a cost function of the general form

$$C^I(d) = \frac{1}{\alpha}\left[\left(\frac{\bar{d}}{d}\right)^{1/\beta} - 1\right] + C_0 \tag{3}$$

where $\alpha > 0$ and $\beta \geq 1$. Their type II model yields a cost function of the general form

$$C^{II}(d) = \frac{1}{\alpha}\left[\frac{\ln d}{\ln \bar{d}} - 1\right] + C_0 \tag{4}$$

for $\alpha > 0$; in both cases C_0 represents the cost needed to obtain an "average" level of security. There are other models that could also be used; see also [22, 40]. These cost models assume that the entire benefit of increased investment in security only decreases the likelihood of a successful attack; the question of detection is omitted. Thus, our model requires more sophisticated cost models. One could construct a cost function that is simply the sum of terms like (3) and/or (4) for both success and evasion. However it is likely that investments in protection would improve detection and vice-versa, which indicates that a model that separates these costs and then sums them is deficient.

As we will see, our results are not terribly sensitive to the precise form of the cost functions $C_A(\mathbf{a})$ and $C_D(\mathbf{d})$.

5 The Games

5.1 Simple bimatrix game

To begin our analysis of these games, we start by assuming that both the attacker and defender only have a finite number of attacks / defensive postures. In particular, let us assume that the attacker can only choose from the attacks $\{\mathbf{a}_1, \mathbf{a}_2, \ldots, \mathbf{a}_m\}$, while the defender can choose from $\{\mathbf{d}_1, \mathbf{d}_2, \ldots, \mathbf{d}_n\}$. A mixed strategy for the attacker is a choice $\mathbf{x} \in \mathbf{R}^n$ with $\sum x_i = 1$, while a mixed strategy for the defender is a choice $\mathbf{y} \in \mathbf{R}^m$ with $\sum y_j = 1$. Here x_i is the probability that the attacker chooses $\mathbf{a}_i$, while y_j is the probability that the defender chooses $\mathbf{d}_j$.

In this case the game always has Nash equilibria in pure strategies. In particular, the pure strategy $(\mathbf{a}_i, \mathbf{d}_j)$ is a Nash equilibrium provided the attacker chooses i so that

$$A_s a_{i,s} + A_e a_{i,e} - C_A(\mathbf{a}_i) \tag{5}$$

is as large as possible and the defender chooses j so that

$$-D_s d_{j,s} - D_e d_{j,e} - C_D(\mathbf{d}_j) \tag{6}$$

is as large as possible. Further, the problem has more than one Nash equilibrium if and only if there is more than one pair (i, j) at which the maxima of (5) and (6) obtain. In this case, we also obtain Nash equilibria for the mixed strategies composed of linear combinations of optimal pure strategies; together these are all of the Nash equilibria. The precise details are found in Proposition 4.

A number of practical conclusions follow from these game theoretic facts. First, the choice of optimum solution for one side can be determined without any knowledge of the behavior of the opponent. In particular, a defender need not hypothesize about the relative importance of cost, likelihood of success, and likelihood of detection for a potential attacker. Instead, the optimal solution for the defender uses only the relative weights and costs associated with the defender. It is interesting to note that this also means that the defender's optimal strategy is not affected by the existence of attacks that are unknown to the defender.

Secondly, a defender might feel that putting increased effort into detection might deter an attacker into using stealthier but less effective techniques, thus reducing the likelihood of a successful attack. However, this is not the case. Indeed, the attacker's optimal solution also only depends on parameters associated with the attacker. In particular, the attacker's optimal choice can be made without knowledge of the possible choices for the defender's posture.

There are some practical limitations to the significance of this result however. In the model under consideration, we have assumed that the attacker will definitely attack the defender, and that there is only one defender. The significance of our result then, is that, if the defender has definitely decided to attack a particular target, then the defender cannot deter the attacker into using a stealthier but weaker attack by placing more emphasis on attack detection.

An interesting generalization of this model would be to allow for an attacker to choose from multiple defenders, and then analyze the effect of possible deterrence when the attacker does not have a predetermined target.

Finally, we note that all of these results are proven using linear models for P_s and P_e. However, even if the functions P_s and P_e are nonlinear, then the payoff from the strategies chosen by using the linear model will be close to the payoffs for the optimal strategies for the nonlinear model- even though the actual strategies may be quite different.

Indeed, suppose that

$$P_s(a_s, d_s) = \ell_s(a_s, d_s) + \sigma_s(a_s - \bar{a}_s, d_s - \bar{a}_s)$$
$$P_e(a_s, d_s) = \ell_e(a_s, d_s) + \sigma_e(a_s - \bar{a}_e, d_s - \bar{a}_e)$$

for nonlinear choices σ_s and σ_e. Then the payoff functions for a pure strategy for the nonlinear game have the form

$$\pi_A(\mathbf{a}, \mathbf{d}) = \pi_A^{(\ell)}(\mathbf{a}, \mathbf{d}) + N_A(\mathbf{a}, \mathbf{d}) \tag{7a}$$
$$\pi_D(\mathbf{a}, \mathbf{d}) = \pi_D^{(\ell)}(\mathbf{a}, \mathbf{d}) + N_D(\mathbf{a}, \mathbf{d}) \tag{7b}$$

where $\pi_A^{(\ell)}(\mathbf{a}, \mathbf{d})$ and $\pi_D^{(\ell)}(\mathbf{a}, \mathbf{d})$ are the payoffs from the linear model

$$\begin{aligned}\pi_A^{(\ell)}(\mathbf{a}, \mathbf{d}) = A_s\ell_s(a_s, d_s) \\ + A_e\ell_e(a_s, d_s) - C_A(\mathbf{a})\end{aligned} \tag{8a}$$
$$\begin{aligned}\pi_D^{(\ell)}(\mathbf{a}, \mathbf{d}) = -D_s\ell_s(a_s, d_s) \\ - D_e\ell_e(a_s, d_s) - C_D(\mathbf{a})\end{aligned} \tag{8b}$$

while the nonlinear terms have the form

$$\begin{aligned}N_A(\mathbf{a}, \mathbf{d}) &= A_s\sigma_s(a_s - \bar{a}_s, d_s - \bar{a}_s) \\ &\quad + A_e\sigma_e(a_s - \bar{a}_e, d_s - \bar{a}_e) \\ N_D(\mathbf{a}, \mathbf{d}) &= -D_s\sigma_s(a_s - \bar{a}_s, d_s - \bar{a}_s) \\ &\quad - D_e\sigma_e(a_s - \bar{a}_e, d_s - \bar{a}_e)\end{aligned}$$

We have proven (Proposition 5) that the difference in the payoff the defender receives by using an equilibrium strategy constructed as a solution of the linear game rather than as an equilibrium of the full nonlinear game is no larger that the value of $|N_D|$. Since the functions σ_s, σ_e have been shown to be small if P_s and P_e are quadratic and the domains of definition of P_s and P_e are close to $\mathcal{D}_s$ and $\mathcal{D}_e$, we know that N_A and N_D are both small; thus the payoff to the defender obtained by using the equilibrium from the linear problem is close to the payoff they would obtain if they used the equilibrium solution to the full problem.

5.2 The general game

In the previous section, we assumed that both the attacker and the defender have only a finite number of possible choices of attacks and defensive postures; this assumption was made only for mathematical simplicity. All of the results of that section remain valid if we remove this assumption.

Indeed, let us instead assume that the attacker can choose any attack $\mathbf{a} = (a_s, a_e)$ in some set S_A, while the defender can choose any defensive posture $\mathbf{d} = (d_s, d_e)$ in the set S_D. Given the restriction on the choices for $\mathbf{a}$ and $\mathbf{d}$ from (2), we know that

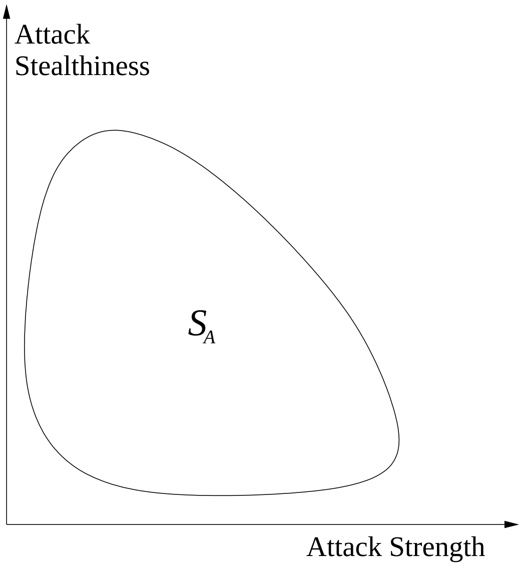

$$S_A, S_D \subseteq (\bar{a}_s/2, \bar{a}_s/2 + 1/2) \times (\bar{a}_e/2, \bar{a}_e/2 + 1/2)$$

Then, the pure strategy of always choosing $(\mathbf{a}, \mathbf{d})$ is a Nash equilibrium if and only if

$$A_s a_s + A_e a_e - C_A(\mathbf{a})$$

and

$$-D_s d_s - D_e d_e - C_D(\mathbf{d})$$

are as large as possible. Once again, the optimal strategy for each side can be determined without any knowledge of the relative weights assigned by the opponent. A precise statement of this result is given in Proposition 6.

6 Theorems and Proofs

Proposition 1 *There are no non-trivial quadratic choices of P that satisfy $0 \le P(a,d) \le 1$ on the domain $\mathcal{D} = \{(a,d) : \frac{1}{2}\bar{a} \le a, d \le \frac{1}{2}\bar{a} + \frac{1}{2}\}$.*

Proof: A general quadratic form for P is

$$\begin{aligned} P(a,d) &= \ell(a,d) \\ &\quad + A(a-\bar{a})^2 + B(a-\bar{a})(d-\bar{a}) + C(d-\bar{a})^2. \end{aligned}$$

The definition of a and d implies that $P(\bar{a},d) = \ell(\bar{a},d)$ and $P(a,\bar{d}) = \ell(a,\bar{d})$; thus $A = C = 0$ and the general form is

$$P(a,d) = \ell(a,d) + B(a-\bar{a})(d-\bar{a}).$$

If P has the domain $\mathcal{D}$, then $0 \le P(a,d) \le 1$ for all $(a,d) \in \mathcal{D}$. However, because $\ell(\frac{1}{2}\bar{a}, \frac{1}{2}\bar{a}) = 0$,

$$P(\tfrac{1}{2}\bar{a}, \tfrac{1}{2}\bar{a}) = \tfrac{1}{4}B\bar{a}^2$$

so $P \ge 0$ implies $B \ge 0$. On the other hand, because $\ell(\frac{1}{2}\bar{a} + \frac{1}{2}, \frac{1}{2}\bar{a} + \frac{1}{2}) = 1$,

$$P(\tfrac{1}{2}\bar{a} + \tfrac{1}{2}, \tfrac{1}{2}\bar{a} + \tfrac{1}{2}) = 1 + \tfrac{1}{4}B(\bar{a}-1)^2$$

so that the requirement $P \le 1$ implies $B \le 0$. ■

Proposition 2 *Let P be any quadratic choice of P satisfying $0 \le P(a,d) \le 1$ on the domain $\mathcal{D}^*_\mu$. Then*

$$|P(a,d) - \ell(a,d)| \le (1-\mu)\min\{\bar{a}, 1-\bar{a}\}.$$

Proof: The general quadratic form of P is

$$P(a,d) = \ell(a,d) + B(a-\bar{a})(d-\bar{a})$$

for some choice of B; however B is not arbitrary. Indeed, because $0 \le P \le 1$ on $\mathcal{D}^*_\mu$, we know

$$P(\bar{a} - \tfrac{\mu}{2}\bar{a}, \bar{a} - \tfrac{\mu}{2}\bar{a}) = \bar{a} - \mu\bar{a} + \tfrac{1}{4}B\mu^2\bar{a}^2 \ge 0$$

so that

$$B \ge -\frac{4(1-\mu)}{\mu^2\bar{a}}.$$

Similarly

$$\begin{aligned} &P(\bar{a} + \tfrac{\mu}{2}(1-\bar{a}), \bar{a} + \tfrac{\mu}{2}(1-\bar{a})) \\ &\qquad = \bar{a} + \mu(1-\bar{a}) + \tfrac{1}{4}B\mu^2(1-\bar{a})^2 \le 1 \end{aligned}$$

so that

$$B \le \frac{4(1-\mu)}{\mu^2(1-\bar{a})}.$$

For any choice $(a,d) \in \mathcal{D}^*_\mu$,

$$|P(a,d) - \ell(a,d)| \le |B|\,|a-\bar{a}|\,|d-\bar{a}|.$$

But $|a-\bar{a}|, |d-\bar{a}| \le \min\{\frac{\mu}{2}\bar{a}, \frac{\mu}{2}(1-\bar{a})\}$ so that

$$\begin{aligned} &|P(a,d) - \ell(a,d)| \\ &\le \frac{4(1-\mu)}{\mu^2}\max\left\{\frac{1}{\bar{a}}, \frac{1}{1-\bar{a}}\right\}\left[\frac{\mu}{2}\min\{\bar{a}, 1-\bar{a}\}\right]^2. \end{aligned}$$

■

Corollary 3 *The quadratic function $P(a,d) = \ell(a,d) + B(a-\bar{a})(d-\bar{a})$ satisfies*

1. *$P(a,d) = \ell(a,d)$ for $a = \bar{a}$ or $d = \bar{a}$,*
2. *$0 \le P(a,d) \le 1$ for all $(a,d) \in D^*_\mu$, and*
3. *$\frac{\partial P}{\partial a} > 0$ and $\frac{\partial P}{\partial d} > 0$ for all $(a,d) \in D^*_\mu$*

for any choice of B with

$$\begin{aligned} &\max\left\{\frac{-4(1-\mu)}{\mu^2\bar{a}}, \frac{-2}{\mu(1-\bar{a})}\right\} \\ &\qquad\qquad \le B \le \min\left\{\frac{4(1-\mu)}{\mu^2(1-\bar{a})}, \frac{2}{\mu\bar{a}}\right\}. \end{aligned}$$

Proof: Clearly (1) follows immediately. Direct calculation shows us that

$$\frac{\partial P}{\partial a} = 1 + B(d-\bar{a}), \qquad \frac{\partial P}{\partial d} = 1 + B(a-\bar{a}).$$

If $B > 0$, then because we have $a-\bar{a}, d-\bar{a} > (-\mu/2)\bar{a}$, we see that

$$\frac{\partial P}{\partial a} > 1 - B\frac{\mu}{2}\bar{a}, \qquad \frac{\partial P}{\partial d} > 1 - B\frac{\mu}{2}\bar{a}$$

so that the condition $B < 2/(\mu\bar{a})$ implies both $\frac{\partial P}{\partial a} > 0$ and $\frac{\partial P}{\partial d} > 0$. Similarly, if $B < 0$, then because $a-\bar{a}, d-\bar{a} < \mu/2(1-\bar{a})$, we see that

$$\frac{\partial P}{\partial a} > 1 + B\frac{\mu}{2}(1-\bar{a}), \qquad \frac{\partial P}{\partial d} > 1 - B\frac{\mu}{2}(1-\bar{a})$$

so that the condition $B > -2/(\mu(1-\bar{a}))$ implies both $\frac{\partial P}{\partial a} > 0$ and $\frac{\partial P}{\partial d} > 0$; thus (3) is proven. With this in hand, we see that, for any $(a, d) \in D^*_\mu$

$$\begin{aligned} &P(\bar{a} + \tfrac{\mu}{2}(1-\bar{a}), \bar{a} + \tfrac{\mu}{2}(1-\bar{a})) \\ &\qquad \leq P(a, d) \leq P(\bar{a} - \tfrac{\mu}{2}\bar{a}, \bar{a} - \tfrac{\mu}{2}\bar{a}) \end{aligned}$$

so that (2) follows from the proof of the previous result. ■

Proposition 4 *In the linear bimatrix game, the pure strategy formed by choosing* $(\mathbf{a}, \mathbf{d})$ *is a Nash equilibrium if and only if*

$$A_s a_s + A_e a_e - C_A(\mathbf{a})$$

and

$$-D_s d_s - D_e d_e - C_D(\mathbf{d})$$

are as large as possible.

A mixed strategy is a Nash equilibrium solution if and only if it is a linear combination of pure strategies that are themselves Nash equilibrium strategies.

Proof: Suppose that the attacker has m possible attacks $\mathbf{a}_i$ for $i \in \{1, 2, \ldots, m\}$, while the defender has n possible defenses $\mathbf{d}_j$ for $j \in \{1, 2, \ldots, n\}$.

Let $\mathbf{S}$ be the matrix of success probabilities, so that

$$S_{ij} = P(a_{i,s}, d_{j,s}), \tag{9}$$

and let $\mathbf{E}$ be the matrix of evasion probabilities, so that

$$E_{ij} = P(a_{i,e}, d_{j,e}). \tag{10}$$

In each case, the attacker chooses the row and the defender chooses the column; thus if $\mathbf{x} \in \mathbf{R}^m$ is a mixed strategy for the attacker and $\mathbf{y} \in \mathbf{R}^n$ is a mixed strategy for the defender then the expected success probability is

$$S = \sum_{i,j} S_{ij} x_i y_j = \mathbf{x}^T \mathbf{S} \mathbf{y}$$

and the expected evasion probability is

$$E = \sum_{i,j} E_{ij} x_i y_j = \mathbf{x}^T \mathbf{E} \mathbf{y}$$

The payoff matrix for the attacker then is

$$\pi_A = A_s \mathbf{S} + A_e \mathbf{E} - \mathbf{C}_A$$

where $\mathbf{C}$ is the cost matrix

$$\mathbf{C}_A = \begin{pmatrix} C_A(\mathbf{a}_1) & C_A(\mathbf{a}_1) & \ldots & C_A(\mathbf{a}_1) \\ C_A(\mathbf{a}_2) & C_A(\mathbf{a}_2) & \ldots & C_A(\mathbf{a}_2) \\ \ldots & \ldots & \ldots & \ldots \\ C_A(\mathbf{a}_m) & C_A(\mathbf{a}_m) & \ldots & C_A(\mathbf{a}_m) \end{pmatrix}.$$

Similarly, the payoff matrix for the defender is

$$\pi_D = -D_s \mathbf{S} - D_e \mathbf{E} - \mathbf{C}_D$$

where $\mathbf{C}_D$ is the cost matrix

$$\mathbf{C}_D = \begin{pmatrix} C_D(\mathbf{d}_1) & C_D(\mathbf{d}_2) & \ldots & C_D(\mathbf{d}_n) \\ C_D(\mathbf{d}_1) & C_D(\mathbf{d}_2) & \ldots & C_D(\mathbf{d}_n) \\ \ldots & \ldots & \ldots & \ldots \\ C_D(\mathbf{d}_1) & C_D(\mathbf{d}_2) & \ldots & C_D(\mathbf{d}_n) \end{pmatrix}.$$

The expected payoff for the attacker is

$$\begin{aligned} \mathbf{x}^T \pi_A \mathbf{y} &= \mathbf{x}^T (A_s \mathbf{S} + A_e \mathbf{E} - \mathbf{C}_A) \mathbf{y} \\ &= \sum_{i,j} (A_s S_{ij} + A_e E_{ij}) x_i y_j - \sum_i x_i C_A(\mathbf{a}_i) \end{aligned}$$

so that if we now use the model (1) to substitute for the matrices $\mathbf{S}$ and $\mathbf{E}$ in (9) and (10) we find that the payoff is

$$\begin{aligned} \mathbf{x}^T \pi_A \mathbf{y} = &\sum_i (A_s a_{i,s} + A_e a_{i,e} - C_A(\mathbf{a}_i)) x_i \\ &+ \sum_j (A_s d_{j,s} + A_e d_{j,e}) y_j \\ &- A_s \bar{a}_s - A_e \bar{a}_e. \end{aligned} \tag{11}$$

Similarly, the expected payoff for the defender is

$$\begin{aligned} \mathbf{x}^T \pi_D \mathbf{y} &= \mathbf{x}^T (-D_s \mathbf{S} - D_e \mathbf{E} - \mathbf{C}_D) \mathbf{y} \\ &= \sum_{i,j} (-D_s S_{ij} - D_e E_{ij}) x_i y_j - \sum_j y_j C_D(\mathbf{d}_j) \end{aligned}$$

and thus

$$\begin{aligned} \mathbf{x}^T \pi_D \mathbf{y} = &- \sum_i (D_s a_{i,s} + D_e a_{i,e}) x_i \\ &- \sum_j (D_s d_{j,s} + D_e d_{j,e} + C_D(\mathbf{d}_j) y_j \\ &+ D_s \bar{a}_s + D_e \bar{a}_e. \end{aligned} \tag{12}$$

Examining (11), we see that we can maximize the attacker's expected payoff by choosing i so that

$$A_s a_{i,s} + A_e a_{i,e} - C_A(\mathbf{a}_i)$$

is as large as possible and using the pure strategy that always chooses $\mathbf{a}_i$. Moreover, this choice is independent of the defender's choice of strategy. Similarly, examining (12), the defender wants to choose j so that

$$D_s d_{j,s} + D_e d_{j,e} + C_D(\mathbf{d}_j)$$

is as large as possible.

If there is more than one choice of i or j at which these maxima are obtained, then any mixed strategy consisting of linear combinations of these choices will also form a Nash equilibrium. ■

Proposition 5 *Let (x, y) be a mixed strategy that forms a Nash equilibrium for the nonlinear game with payoff functions* (7)*, and let $(\hat{x}, \hat{y})$ be a strategy that forms a Nash equilibrium for the linear game with payoff functions* (8)*. Then the difference in payoff functions satisfies*

$$\left|\mathbf{x}^T \pi_A \mathbf{y} - \hat{\mathbf{x}}^T \pi_A^{(\ell)} \mathbf{y}\right| \le \max |N_A|,$$
$$\left|\mathbf{x}^T \pi_D \mathbf{y} - \mathbf{x}^T \pi_D^{(\ell)} \hat{\mathbf{y}}\right| \le \max |N_D|.$$

Proof: Let K be the set of indices so that

$$A_s a_{k,s} + A_e a_{k,e} - C_A(\mathbf{a}_k)$$

is as large as possible for any $k \in K$. Clearly, the value of $A_s a_{k,s} + A_e a_{k,e} - C_A(\mathbf{a}_k)$ is a constant independent of k. Further, Proposition 4 implies that $\hat{x}_i = 0$ for all $i \notin K$. We also note that

$$\pi_A^{(\ell)}(\mathbf{a}_i, \mathbf{d}_j) \le \pi_A^{(\ell)}(\mathbf{a}_k, \mathbf{d}_j) \tag{13}$$

for any $k \in K$. Indeed

$$\begin{aligned}\pi_A^{(\ell)}(\mathbf{a}_k, \mathbf{d}_j) &- \pi_A^{(\ell)}(\mathbf{a}_i, \mathbf{d}_j)\\ &= [A_s a_{k,s} + A_e a_{k,e} - C_A(\mathbf{a}_k)]\\ &\quad - [A_s a_{i,s} + A_e a_{i,e} - C_A(\mathbf{a}_i)] \ge 0.\end{aligned}$$

Similarly, we have $\pi_A^{(\ell)}(\mathbf{a}_k, \mathbf{d}_j) = \pi_A^{(\ell)}(\mathbf{a}_\kappa, \mathbf{d}_j)$ for all $k, \kappa \in K$.

With these preliminaries completed, we start by noticing that because $(\mathbf{x}, \mathbf{y})$ is a Nash equilibrium for the nonlinear problem, that

$$\hat{\mathbf{x}}^T \pi_A \mathbf{y} \le \mathbf{x}^T \pi_A \mathbf{y}.$$

Now

$$\hat{\mathbf{x}}^T \pi_A \mathbf{y} = \hat{\mathbf{x}}^T \pi_A^{(\ell)} \mathbf{y} + \hat{\mathbf{x}}^T N_A \mathbf{y},$$

so that

$$\mathbf{x}^T \pi_A \mathbf{y} \ge \hat{\mathbf{x}}^T \pi_A^{(\ell)} \mathbf{y} + \hat{\mathbf{x}}^T N_A \mathbf{y}.$$

On the other hand,

$$\begin{aligned}\mathbf{x}^T \pi_A \mathbf{y} &= \mathbf{x}^T \pi_A^{(\ell)} \mathbf{y} + \mathbf{x}^T N_A \mathbf{y}\\ &= \sum_{i,j} x_i y_j \pi_A^{(\ell)}(\mathbf{a}_i, \mathbf{d}_j) + \mathbf{x}^T N_A \mathbf{y}.\end{aligned}$$

Thus, using (13), we see that for any $k \in K$ we have

$$\begin{aligned}\mathbf{x}^T \pi_A \mathbf{y} &\le \sum_{i,j} x_i y_j \pi_A^{(\ell)}(\mathbf{a}_k, \mathbf{d}_j) + \mathbf{x}^T N_A \mathbf{y}\\ &\le \sum_{j} y_j \pi_A^{(\ell)}(\mathbf{a}_k, \mathbf{d}_j) + \mathbf{x}^T N_A \mathbf{y}.\end{aligned}$$

Since $\hat{x}_i = 0$ for $i \notin K$, and $\sum \hat{x}_i = 1$ and $\pi_A^{(\ell)}(\mathbf{a}_k, \mathbf{d}_j) = \pi_A^{(\ell)}(\mathbf{a}_\kappa, \mathbf{d}_j)$ for all $k, \kappa \in K$, we have

$$\begin{aligned}\mathbf{x}^T \pi_A \mathbf{y} &\le \sum_{i,j} \hat{x}_i y_j \pi_A^{(\ell)}(\mathbf{a}_i, \mathbf{d}_j) + \mathbf{x}^T N_A \mathbf{y}\\ &\le \hat{\mathbf{x}}^T \pi_A^{(\ell)} \mathbf{y} + \mathbf{x}^T N_A \mathbf{y}.\end{aligned}$$

Thus

$$\hat{\mathbf{x}}^T N_A \mathbf{y} \le \mathbf{x}^T \pi_A \mathbf{y} - \hat{\mathbf{x}}^T \pi_A^{(\ell)} \mathbf{y} \le \mathbf{x}^T N_A \mathbf{y},$$

giving us our result, at least for the attacker. The analysis for the defender follows in the same fashion. ■

Proposition 6 *In the general linear game, the pure strategy formed by choosing $(\mathbf{a}, \mathbf{d}) \in S_A \times S_D$ is a Nash equilibrium if and only if*

$$A_s a_s + A_e a_e - C_A(\mathbf{a})$$

and

$$-D_s d_s - D_e d_e - C_D(\mathbf{d})$$

are as large as possible.

Proof: This follows immediately from the fact that the payoff functions for the attacker and defender are

$$\begin{aligned}\pi_A(\mathbf{a},\mathbf{d}) &= [A_s a_s + A_e a_e - C_A(\mathbf{a})] \\ &\quad + [A_s d_s + A_e d_e] \\ &\quad - A_s \bar{a}_s - A_e \bar{a}_e \\ \pi_D(\mathbf{a},\mathbf{d}) &= -[D_s d_s + D_e d_e + C_D(\mathbf{d})] \\ &\quad - [D_s a_s + D_e a_e] \\ &\quad + D_s \bar{a}_s + D_e \bar{a}_e\end{aligned}$$

■

References

[1] A. Agah, S.K. Das, and K. Basu, “A non-cooperative game approach for intrusion detection in sensor networks,” in *Vehicular Technology Conference, 2004*, 2902- 2906.

[2] A. Agah, S. Das, K. Basu, and M. Asadi, “Intrusion detection in sensor networks: A non-cooperative game approach,” in *3rd IEEE International Symposium on Network Computing and Applications*, (NCA 2004), Boston, MA, August 2004, pp. 343-346.

[3] T. Alpcan and T. Basar, “A game theoretic approach to decision and analysis in network intrusion detection,” in *Proceedings of the 42nd IEEE Conference on Decision and Control*, Maui, HI, 2595-2600, December 2003.

[4] T. Alpcan and T. Basar, “A game theoretic analysis of intrusion detection in access control systems,” in *Proc. of the 43rd IEEE Conference on Decision and Control*, Paradise Island, Bahamas, December 2004, pp. 1568-1573.

[5] R.J. Aumann and S. Hart (eds.), *Handbook of Game Theory*, Elsevier Press, 1992.

[6] R.J. Aumann and M. Maschler, “Game-theoretic aspects of gradual disarmament”, Chapter V in *Report to the U.S. Arms Control and Disarmament Agency ST-80*. Princeton: Mathematica, 1966.

[7] D. Billings, N. Burch, A. Davidson, *et. al.*, “Approximating game-theoretic optimal strategies for full-scale poker” in *Proceedings of the 2003 International Joint Conference on Artificial Intelligence*.

[8] S. Bistarelli, F. Fioravanti and P. Peretti, “Defense trees for economic evaluation of security investments”, in *First International Conference on Availability, Reliability and Security* (ARES’06), pp. 416-423 2006.

[9] L.D. Bodin, L.A. Gordon and M.P. Loeb, “Evaluating information security investments using the analytic hierarchy process,” *Comm. ACM*, **48**(2), 79-83, 2005.

[10] S.J. Brams and D.M. Kilgour, *Game Theory and National Security*, Blackwell Publishers, 1988.

[11] H. Cavusoglu, B. Mishra and S. Raghunathan, “A model for evaluating IT security investments,” *Comm. ACM*, **47**(7), 87-92, 2004.

[12] H. Cavusoglu, B. Mishra and S. Raghunathan, “The value of intrusion detection systems in information technology security architecture,” *Information Systems Research*, **16**(1), 28-46, March 2005.

[13] F. uppens, and A. Miége, “Alert Correlation in a Cooperative Intrusion Detection Framework,” *Proceedings of the 2002 IEEE Symposium on Security and Privacy*, 202-215, May 12-15, 2002

[14] D. Fudenberg, and J. Tirole, *Game Theory*, MIT Press, 1991.

[15] L.A. Gordon and M.P. Loeb, “Using Information Security as a Response to Competitor Analysis Systems,” *Comm. ACM*, **44**(9), 70-75, 2001.

[16] L.A. Gordon and M.P. Loeb, “The Economics of Information Security Investment,” *ACM Transactions on Information and System Security*, **5**(4), 438-457, 2002.

[17] L.A. Gordon, M.P. Loeb and T. Sohail, “A Framework for using insurance for cyber-risk management,” *Comm. ACM*, **46**(3), 81-85, 2003.

[18] L.A. Gordon and M.P. Loeb, “Budgeting process for information security expenditures,” *Comm. ACM*, **49**(1), 121-125, 2006.

[19] W. Grimm and K.H. Well, “Modeling air combat as a differential game - recent approaches and future requirements,” in *Proceedings of the 4th International Symposium on Differential Games and Applications*, 1990.

[20] S.N. Hamilton, W.L. Miller, A. Ott and O.S. Saydjari, “The role of game theory ininformation warfare,” in *The Fourth Information Survivability Workshop (ISW-2001/2002)*, Vancouver, BC, Canada, March 2002.

[21] S.N. Hamilton, W.L. Miller, A. Ott and O.S. Saydjari, "Challenges in applying game theory to the domain of information warfare," in *The Fourth Information Survivability Workshop (ISW-2001/2002)*, Vancouver, BC, Canada, March 2002.

[22] K. Hausken, "Returns to information security investment: The effect of alternative information security breach functions on optimal investment and sensitivity to vulnerability," *Information Systems Frontiers*, (to appear).

[23] R. Isaacs, *Differential Games: A Mathematical Theory with Applications to Warfare and Pursuit, Control and Optimization*. New York: Wiley, 1967.

[24] M. Kodialam and T.V. Lakshman, "Detecting network intrusions via sampling: A game theoretic approach," in *Proceedings of the 22nd Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM 2003)*, San Francisco CA, USA, March 30 - April 3, 2003.

[25] I. Kotenko and M. Stepashkin, "Analyzing vulnerabilities and measuring security level at design and exploitation stages of computer life cycle," in *Computer Network Security*, Proceedings of the Third International Workshop on Mathematical Methods, Models, and Architectures for Computer Network Security, MMM-ACNS 2005, St. Petersburg, Russia, September 24-28, 2005, V. Gorodetsky, I. Kotenko, and V. Skormin (Eds.), 311-324, 2005.

[26] P. Liu, L. Li, "A Game Theoretic Approach to Attack Prediction," Technical Report, Cyber Security Group, 2002. Online at http://ist.psu.edu/s2/paper/predict.pdf

[27] P. Liu, W. Zang and M. Yu, "Incentive-based modeling and inference of attacker intent, objectives, and strategies," *ACM Trans. Inf. Syst. Secur.* **8**(1) (Feb. 2005), 78-118.

[28] K. Lye and J. Wing, "Game strategies in network security," *Int. J. Inf. Secur.*, **4**, 71-86, 2005.

[29] J. McMillan, "Selling spectrum rights," *The Journal of Economic Perspectives*, **8**(3), 145–162, 1994.

[30] T. Neubauer, C. Stummer, and E. Weippl, "Workshop-based Multiobjective Security Safeguard Selection, ares," in *First International Conference on Availability, Reliability and Security* (ARES'06), 366-373, 2006.

[31] P. Ning, Y. Cui and D.S. Reeves, "Constructing attack scenarios through correlation of intrusion alerts," In *Proceedings of the 9th ACM Conference on Computer and Communications Security* (Washington, DC, USA, November 18 - 22, 2002). V. Atluri, Ed. CCS '02. ACM Press, New York, NY, 245-254.

[32] P. Ning and D. Xu, "Learning attack strategies from intrusion alerts," In *Proceedings of the 10th ACM Conference on Computer and Communications Security* (Washington D.C., USA, October 27 - 30, 2003). CCS '03. ACM Press, New York, NY, 200-209.

[33] P. Ning and D. Xu, "Hypothesizing and reasoning about attacks missed by intrusion detection systems," *ACM Trans. Inf. Syst. Secur.* **7**, 4 (Nov. 2004), 591-627.

[34] G. Owen, *Game Theory*, Academic Press, 1995.

[35] P. Patcha and J.-M. Park, "A game theoretic formulation for intrusion detection in mobile ad-hoc networks," *Int. J. Net. Sec.*, **2**(2) (Mar. 2006), 131-137.

[36] O. Sheyner, J. Haines, S. Jha, R. Lippmann and J.M. Wing, "Automated generation and analysis of attack graphs," in *Security and Privacy, 2002. Proceedings. 2002 IEEE Symposium on Security and Privacy*, 2002, 273-284.

[37] H.A. Simon and J. Schaeffer, "The Game of Chess" in *Handbook of Game Theory, Volume 1*, R.J. Aumann and S. Hart (eds.), Elsevier Science Publishers, 1992, pp. 1-17

[38] S.J. Templeton and K. Levitt, "A requires/provides model for computer attacks," In *Proceedings of the 2000 Workshop on New Security Paradigms* (Ballycotton, County Cork, Ireland, September 18 - 21, 2000). NSPW '00. ACM Press, New York, NY, 31-38.

[39] A. van den Nouweland, P. Borm, W. van Golstein Brouwers, R. Groot Bruinderink, and S. Tijs, "A game theoretic approach to problems in telecommunication," *Management Science*, **42**(2), 294–303, February 1996.

[40] J. Willemson, "On the Gordon and Loeb model for information security investment", Fifth Workshop on the Economics of Information Security (WEIS 2006), Cambridge, England, June 2006.