Computing the Nash Equilibria of Intruder Classification Games*

Lemonia Dritsoula[1], Patrick Loiseau[2], and John Musacchio[1]

[1] UC Santa Cruz, Santa Cruz, USA
[2] EURECOM, Sophia-Antipolis, France
{lenia, johnm} @soe.ucsc.edu, patrick.loiseau@eurecom.fr

Abstract. We investigate the problem of classifying an intruder of two different types (spy or spammer). The classification is based on the number of file server and mail server attacks a network defender observes during a fixed window. The spammer naively attacks (with a known distribution) his main target: the mail server. The spy strategically selects the number of attacks on his main target: the file server. The defender strategically selects his classification policy: a threshold on the number of file server attacks. We first develop parameterized families of payoff functions for both players and analyze the Nash equilibria of the non-cooperative nonzero-sum game. We analyze the strategic interactions of the two players and the tradeoffs each one of them faces: The defender chooses a classification threshold that balances the cost of missed detections and false alarms while the spy seeks to hit the file server as much as possible while still evading detection. We give a characterization of the Nash equilibria in mixed strategies, and demonstrate how the Nash equilibria can be computed in polynomial time. We give two examples of the general model, one that involves forensics on the side of the defender and one that does not. Finally, we evaluate how investments in forensics and data logging could improve the Nash equilibrium payoff of the defender.

Keywords: Nash equilibria, intruder classification, polynomial complexity

1 Introduction

Classifying an attacker is not an easy task. In almost every network security situation, the defender has limited resources. The defender needs to distinguish between different types of attackers (spy or spammer) and decide what kind of actions should be triggered. For example, an attack on a mail server by a spammer (causing at most network congestion) should be treated differently than an attack on a file server (possibly involving identity theft). Knowing that a defender is trying to classify attackers, the strategic spy is likely to change the way he attacks in order to make it more difficult to be classified as a spy.

* This work was supported by AFOSR grant FA9550-09-1-0049.

The paper focuses on the specific situation of a network with one mail server and one file server. However, the model developed is very flexible and independent of the underlying architecture. In particular, the model can fit many situations in which a strategic attacker has an incentive to blend in with or be mistaken for other more benign kinds of attackers or even as legitimate users. We first develop a generic model that guarantees the NE computation in polynomial time, but also provides insights on how the players' NE strategies are derived. Computing the NE is a tractable process, even for larger N. We propose two characteristic examples, in which the defender has made different security investments in forensic mechanisms. The analysis of these models provide us with a qualitative and quantitative view on how changes on the network parameters affect the strategies of the players. We explore the relation between the NE strategies of the two strategic players and the behavior of the non-strategic one, and we evaluate the defender's expected gain after investing in forensics.

1.1 Related Work

There is a substantially increasing body of work in the game theory community that explores and suggests equilibrium solutions in security games (see e.g., a recent survey in [4]). Particularly relevant to the present paper is the growing body of work is on the topic of intrusion detection. In [5], Alpcan and Başar present a security game between an attacker and an intrusion detection system and address some of the basic security tradeoffs, e.g., false alarms versus undetected attackers. They also provide insightful overview on how different network parameters affect the performance of the intruder detection system. Our game-theoretic framework investigates a more complex game and provides analytic expressions for the defender's NE strategies for any network size. We also investigate the way the nonstrategic player influences the spy's strategy.

Gueye, Walrand, and Anantharam [7, 8] have investigated the structure of the Nash equilibria in a network topology game, in which attacker and defender select which links to attack and use for communication respectively. They consider a special case of nonzero-sum games, in which the different term in the players' payoffs is controlled only by the one player. In these games, one player optimizes his payoff against his opponent who has optimized his payoff as well. Such games are easier to analyze than general nonzero-sum games, and they give interesting insights on the strategies of the two players. Our work is using a similar payoff formulation in a different setting: the defender selects a threshold on file server attacks (not a set of links to use) and there are two different types of attackers.

This is a follow up of our recent work [1], where we investigated and characterized the Nash equilibria of a game similar to what we study in the current paper, but with a much more specific form assumed for the payoff functions of the players. In the current paper, we provide a comprehensive way to derive the strategies of the two players who have generalized payoffs. We further provide evaluation results between two different models, with different assumptions on the resources available to the defender.

1.2 Summary of contributions

In summary, our contributions are the following:

- We propose a generic game-theoretic model to analyze the interactions between two adversaries: a classifier (defender) and a malicious attacker when a nonstrategic spammer is present (Sec. 2).
- We show how to derive the NE strategies in polynomial time (Sec. 3).
- We develop two models for intrusion detection (Sec. 4.1 and 4.2).
- By comparing the above two models, we extract key insights on the expected gain from the defender's investment in forensic capabilities. This is an example of how our methodology can be used to evaluate how changes in the strategic situation affect the equilibrium payoffs of the players. We also investigate the impact of the different network parameters on the resulting NE strategies (Sec. 5).

2 Game Model

The network we consider consists of a defender and two servers that he monitors for potential attacks: a File Server (FS) with sensitive data and a Mail Server (MS) with contents of inferior importance. The defender observes the number of hits from an attacker to each server for a fixed classification window of N time slots. The attacker may be a spy or a spammer with probabilities p and $1-p$ respectively.

The defender is a strategic player that seeks to correctly classify the potential intruder by selecting a threshold T. When he observes T or more hits on the FS, he classifies the attacker as spy; otherwise as spammer. The spy is also a strategic player that selects the number of FS attacks H he will perform. He seeks to attack the FS as frequently as possible, while evading detection. The spammer is a non-strategic player that mostly attacks the MS and adds noise to the network. He also attacks the FS Z times (Z follows a known distribution). For example, the spammer can be modeled to follow the binomial distribution, with a small probability θ_0 to attack the FS at each time slot.

Our solution captures a more general setting than the one presented above. We only require that the attacker has some cost function if he gets detected or missed. We describe the model around the example scenario in which there are two servers, one of which is of primary interest to the strategic attacker (the file server) in order to be more concrete. However, the model we develop is quite general and applicable to many settings in which there is a target of special interest to a strategic attacker but who is incentivized to mix his attack across other targets to make classification more difficult.

Notational Conventions:
We use "min[$\boldsymbol{v}$]" to denote the minimum element of a vector $\boldsymbol{v}$ and "minimize" when we minimize a specific expression over some constraints. We use the *prime* sign ($'$) for transpose of matrices and vectors. All vectors are assumed to be

column vectors and are denoted by bold lowercase letters (e.g., $\boldsymbol{\alpha}$, $\boldsymbol{\beta}$). For matrix notation we use capital greek letters (e.g., Λ). The indicator function is denoted by $\mathbb{1}_{\text{cond}}$; it is equal to 1 if "cond" holds and is equal to 0 otherwise. The column vector of ones of length N is denoted by $\mathbf{1}_N$ and the matrix of ones of dimensions $N \times M$ is denoted by $1_{N\times M}$. The norm of a vector $\boldsymbol{x}$ of length N, denoted by $\|\boldsymbol{x}\|$, always refers to the 1-norm, i.e, $\|\boldsymbol{x}\| = |x_1| + |x_2| + \ldots + |x_N|$. An overview of our notation is shown in Table 1.

2.1 Spy's cost function

The spy cost depends on the defender's classification decision and the number of FS hits. We denote the spy cost function when the spy is detected (i.e., $T \leq H$) by $D(H)$ and when the spy is not detected (i.e., $T > H$) by $M(H)$. Thus, the overall spy cost function is expressed as follows

$$J_A(T, H) = D(H) \cdot \mathbb{1}_{T\leq H} + M(H) \cdot \mathbb{1}_{T>H},$$

or by making the appropriate simplifications

$$J_A(T, H) = [D(H) - M(H)] \cdot \mathbb{1}_{T\leq H} + M(H).$$

2.2 Defender's payoff function

We now describe how the defender's expected payoff function is constructed. We distinguish two cases:

- With probability p the defender faces a spy. If the defender correctly classifies the intruder as a spy (i.e., $T \leq H$), he gains $D(H)$. If the defender misclassifies the spy (i.e., $T > H$), he gains $M(H)$.
- With probability $1-p$ the defender faces a spammer. If the defender correctly classifies the intruder as spammer (i.e., $T \geq Z$), he does not benefit. The defender incorrectly classifies the spammer with probability $\phi(T) = \Pr\{Z \geq T\}$ and in this case there is a false alarm penalty c_{fa}.

Combining these two scenarios, the defender's expected payoff is

$$\tilde{U}_D(T, H) = p \cdot [D(H) \cdot \mathbb{1}_{T\leq H} + M(H) \cdot \mathbb{1}_{T>H}] - (1 - p) \cdot c_{\text{fa}} \cdot \phi(T). \quad (1)$$

By scaling the above function, we get

$$U_D(T, H) = D(H) \cdot \mathbb{1}_{T\leq H} + M(H) \cdot \mathbb{1}_{T>H} - \mu(T),$$

where $\mu(T) = \dfrac{1-p}{p} \cdot c_{\text{fa}} \cdot \phi(T)$. Function $\phi(T)$ is decreasing on T, and we assume that it is strictly decreasing: $\Pr\{Z \geq T\} > \Pr\{Z \geq T + 1\}$.

2.3 Players' interactions

For a classification window of N time slots, the spy has $N+1$ available actions (attack the file server $H \in \{0, \ldots, N\}$ times). The defender has $N+2$ available actions (select $T \in \{0, \ldots, N+1\}$ as the classification threshold).

We model our problem as a nonzero-sum game. However, the defender's payoff is different from the spy's cost function in only one term $\mu(T)$ that depends *only* on the defender's strategy ($U_D(T,H) = J_A(T,H) - \mu(T)$). These games are known as almost zero-sum games or quasi zero-sum games.

We are interested in Nash equilibria in mixed strategies for the following reason. In most cases the spy's best response to a threshold T is to attack the file server a number of times H just below T (unless the cost of being detected is so low that the spy prefers to attack as often as possible even while being detected). Likewise, in most cases, the defender's best response to an H is to choose the threshold T to be just equal with H in order to have the lowest false alarm penalty possible while still detecting the spy. Since each player wants to pick "lower" than the other, there is no pure strategy Nash equilibrium in most cases of interest, so we consider mixed strategies. The spy chooses a distribution vector $\boldsymbol{\alpha}$ on the allowed number of FS hits; $\boldsymbol{\alpha}$ is a vector of size $N+1$ (with non-negative elements that sum to 1). Similarly, the defender chooses a distribution vector $\boldsymbol{\beta}$ on the collection of possible thresholds T; $\boldsymbol{\beta}$ is a vector of size $N+2$ (with non-negative elements that sum to 1).

Let $\tilde{\Lambda}$ be a $(N+1) \times (N+2)$ matrix representing the spy's (pure) strategies' cost. We express the cost matrix of the attacker as

$$\tilde{\Lambda} = \begin{bmatrix} \delta(0) & & 0 \\ & \ddots & \\ 0 & & \delta(N) \end{bmatrix} \cdot \begin{pmatrix} 1 & 0 & \cdots & \cdots & 0 & 0 \\ \vdots & 1 & \ddots & & \vdots & \vdots \\ \vdots & & \ddots & \ddots & \vdots & \vdots \\ \vdots & & & \ddots & 0 & \vdots \\ 1 & \cdots & \cdots & \ldots & 1 & 0 \end{pmatrix} + \begin{bmatrix} M(0) & \cdots & M(0) \\ M(1) & \cdots & M(1) \\ \vdots & & \vdots \\ M(N-1) & \cdots & M(N-1) \\ M(N) & \cdots & M(N) \end{bmatrix},$$

where $\delta(H) = D(H) - M(H)$. Each row i of $\tilde{\Lambda}$ corresponds to one of the $N+1$ possible spy strategies. For instance, row "0" corresponds to spy attacking the FS 0 times (or $H = 0$), row "1" corresponds to spy selecting $H = 1$ and so on. Each column of $\tilde{\Lambda}$ corresponds to one of the $N+2$ possible defender strategies.

Let $\tilde{\Lambda}$ be defined as above, and $\boldsymbol{\alpha}$, $\boldsymbol{\beta}$, be the spy and defender distributions respectively. The attacker cost can be written as $\boldsymbol{\alpha}'\tilde{\Lambda}\boldsymbol{\beta}$ and the defender payoff can be written as $\boldsymbol{\alpha}'\tilde{\Lambda}\boldsymbol{\beta} - \boldsymbol{\mu}'\boldsymbol{\beta}$, where $\boldsymbol{\mu}$ is a strictly decreasing vector (component-wise) with μ_i be the i^{th} component of vector $\boldsymbol{\mu}$. Certain computations are simplified by using a matrix with only positive entries. We define

$$\Lambda = \tilde{\Lambda} + K \cdot 1_{(N+1)\times(N+2)},$$

where $K > 0$ is such that every matrix element is positive. Since $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ must each sum to 1, the expressions $\boldsymbol{\alpha}'\Lambda\boldsymbol{\beta}$ and $\boldsymbol{\alpha}'\Lambda\boldsymbol{\beta} - \boldsymbol{\mu}'\boldsymbol{\beta}$ are respectively the

Table 1. Main Notations

p	probability for spy	$\boldsymbol{\alpha}$	spy's mixed strategy
$D(H)$	detection cost	$\boldsymbol{\beta}$	def. mixed strategy
$M(H)$	missed detection cost	$\boldsymbol{\mu}$	false alarm cost vector
$\delta(H)$	$D(H) - M(H)$	$\theta(\boldsymbol{\beta})$	defendability of $\boldsymbol{\beta}$
H	spy's strategy (# FS hits)	Λ	cost matrix of spy
T	def. strategy (threshold)	s	first tight inequality
Z	# of FS hits by spammer	f	last tight inequality

attacker cost and defender payoff shifted by a constant. Adding a constant to the players' payoff does not affect their best responses, thus from here on we will consider these expressions to be the payoff functions of each player.

3 Game-Theoretic Analysis

It is known that every finite game (finite number of players with finite number of actions for each player) has a mixed-strategy Nash equilibrium [12]. Our game is finite, thus it admits a NE in mixed strategies. In a two-player game, the players' strategies $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ are a NE if each player's stategy is a best response to the other player's mixed strategy.

3.1 Best response analysis

Here is a roadmap of the subsequent analysis.

- Lemma 1 leads to the proof of Theorem 1 on the maximization of the defender's payoff in NE.
- Lemma 2 defines the simplified problem $\Lambda\boldsymbol{x} \geq \mathbf{1}_{N+1}$, $\boldsymbol{x} \geq \mathbf{0}$.
- Theorem 2 introduces the algorithm to compute the players' NE strategies, under certain conditions. To prove the validity of this algorithm, we prove a series of Lemmata (3 – 9).

Lemma 1. *A spy who plays a best response to a defender strategy $\boldsymbol{\beta}$ has a cost* $\min[\Lambda\boldsymbol{\beta}]$.

Proof. For a given defender strategy $\boldsymbol{\beta}$ the minimum attacker cost is achieved by putting positive probability only on strategies corresponding to the minimum entries of the vector $\Lambda\boldsymbol{\beta}$ (recall Λ is positive). Thus the spy's optimal cost is $\min[\Lambda\boldsymbol{\beta}]$. □

Definition 1 (Defendability). *The defendability of a mixed strategy $\boldsymbol{\beta}$ is defined as*

$$\theta(\boldsymbol{\beta}) = \min[\Lambda\boldsymbol{\beta}] - \boldsymbol{\mu}'\boldsymbol{\beta}. \tag{2}$$

It corresponds to the defender's payoff when the attacker plays a best response to $\boldsymbol{\beta}$.

The defendability is a measure of how good a strategy $\boldsymbol{\beta}$ is. The defendability is similar to the notion of vulnerability in [7]. An interesting fact that arises from the definition of the defendability is that the defender's payoff when the attacker plays a best response against $\boldsymbol{\beta}$ depends only on $\boldsymbol{\beta}$.

Theorem 1. *A defender-attacker strategy pair $(\boldsymbol{\alpha}, \boldsymbol{\beta})$ is a NE, if and only if the defendability $\theta(\boldsymbol{\beta})$ is maximal.*

Proof. The intuition behind this proof is twofold. First, we prove that if the defender does not maximize his defendability, then the attacker' optimization problem in NE (in order to make the defender indifferent among the strategies in his support and not want to use strategies outside) is infeasible. Second, we prove that the attacker's optimization problem, when the spy limits the defender to the defendability $\theta(\boldsymbol{\beta})$ yields a spy cost $\min[\Lambda\boldsymbol{\beta}]$, i.e., the same one as if the spy was not limiting the defender to the defendability (as in Lemma 1). (Sketch, see [1] for full proof). □

Definition 2. ***Polyhedron****: A polyhedron is the solution set of a finite number of linear equalities and inequalities.*
Tight constraint*: An inequality constraint is tight, if it holds as an equality; otherwise, it is loose.*
Extreme point*: A point $\boldsymbol{x}$ of a polyhedron is said to be extreme if there is no $\boldsymbol{x}'$ whose set of tight constraints is a strict superset of the set of tight constraints of $\boldsymbol{x}$.*
"Corresponds"*: We say that a point on the polyhedron $\boldsymbol{x}$ corresponds to strategy $\boldsymbol{\beta}$, if $\boldsymbol{\beta} = \boldsymbol{x}/\|\boldsymbol{x}\|$.*

Lemma 2. *The defendability is maximized amongst strategies $\boldsymbol{\beta}$ corresponding to the extreme points of the polyhedron defined by $\Lambda\boldsymbol{x} \geq \mathbf{1}_{N+1}$, $\boldsymbol{x} \geq \mathbf{0}$.*

Proof. As we proved in Theorem 1, in NE, the defender maximizes the defendability, that is, he solves the following "defendability LP"

$$\begin{aligned} \underset{\boldsymbol{\beta}, z}{\text{maximize}} \quad & -\boldsymbol{\mu}'\boldsymbol{\beta} + z \\ \text{subject to} \quad & z \cdot \mathbf{1}_{N+1} \leq \Lambda\boldsymbol{\beta} \\ & \mathbf{1}'_{N+2} \cdot \boldsymbol{\beta} = 1,\ \boldsymbol{\beta} \geq \mathbf{0}. \end{aligned} \tag{3}$$

The solution for z is $z = \min[\Lambda\boldsymbol{\beta}]$ (finite and positive since Λ positive). We can make the following transformation $\boldsymbol{x} = \frac{1}{z} \cdot \boldsymbol{\beta}$, with $\|\boldsymbol{x}\| = \frac{1}{z} \cdot 1$ and get the LP

$$\begin{aligned} \underset{\boldsymbol{x}}{\text{maximize}} \quad & -\boldsymbol{\mu}'\boldsymbol{x} + 1 \\ \text{subject to} \quad & \Lambda \cdot \boldsymbol{x} \geq 1_{N+1},\ \boldsymbol{x} \geq \mathbf{0}. \end{aligned} \tag{4}$$

The intuition behind the proof is that we can rewrite the above LP (4) in terms of $\boldsymbol{\beta}$, and then impose the equality constraint $\sum_{i=0}^{i=N+1} x_i = 1$. Then the objective is linear in $\boldsymbol{\beta}$. We prove that the extreme points of the inequalities in $\boldsymbol{x}$ correspond to the extreme points of $\boldsymbol{\beta}$ in the above LP. The formal proof can be found in [1]. □

3.2 Form of players' strategies in NE

Since the defender maximizes his defendability in NE, the defender must solve the LP given by (4). There exist polynomial-time algorithms to solve linear programming problems [10]. Our approach not only guarantees a low-complexity algorithm to compute the NE strategies of the two players, but it also provides essential intuition about how and why the defender is behaving the way he behaves.

Defender's NE strategy As we saw in Lemma 2, the best response strategy of the defender is found by looking at the extreme points of the polyhedron $\Lambda \boldsymbol{x} \geq \mathbf{1}_{N+1}, \boldsymbol{x} \geq \mathbf{0}$. We call the first type "inequality" constraints and the second type "positivity" constraints. We have $N+1$ "inequality"- and $N+2$ "positivity" constraints. We assume that $\delta(H), M(H)$ are positive functions. If they are not, we can add a constant parameter and render them positive without affecting the Nash equilibria of the game. Writing down the "inequality" constraints, we get

$$\begin{aligned}
\delta(0) \cdot x_0 + M(0)\|\boldsymbol{x}\| &\geq 1 \\
\delta(1) \cdot (x_0 + x_1) + M(1)\|\boldsymbol{x}\| &\geq 1 \\
&\vdots \\
\delta(i) \cdot (x_0 + \cdots + x_i) + M(i)\|\boldsymbol{x}\| &\geq 1 \\
&\vdots \\
\delta(N) \cdot (x_0 + x_1 + \ldots + x_N) + M(N)\|\boldsymbol{x}\| &\geq 1.
\end{aligned}$$

Our goal is to eliminate nonextreme and other points that are not selected by a defender in NE, so that we reduce the number of points we have to check. Depending on the nature of the attacker's cost functions δ and M, we are able to compute analytically the defender's NE strategies in polynomial time. We will consider the following conditions for the subsequent analysis.

Condition 1: $\forall s \in \{0, \ldots, N-1\}$, where $\Delta_k g(i) = g(i+k) - g(i)$,

1. $\Delta_1 \delta(s+1) \geq \Delta_1 \delta(s)$, and
2. $\Delta_1 M(s+1) \geq \Delta_1 M(s)$

Condition 1 suggests that the difference between the cost of the spy upon detection and his cost upon misdetection is non decreasing with respect to H. It also suggests that the marginal cost for the spy when he is not detected is smaller for smaller values of H. We use this condition to prove that the inequalities are violated, unless there is a contiguous block of tight inequalities (see Lemma 5).

Condition 2:

1. $D(H)$ is monotone with respect to the number of attacks to the FS H.
2. $M(H)$ is a decreasing function with respect to H.

Theorem 2 summarizes our results on the computation of Nash equilibria for the intruder classification games.

Theorem 2. *Under condition 1, there exists a defender NE strategy that satisfies a contiguous block (by index) of tight inequalities (indexed s through f). Under condition 2, the contiguous block will finish at index $f = N$, or we only have pure NE. When $f = N$, we search amongst different β_{N+1} for the defender strategies $\boldsymbol{\beta}$ that maximize the defendability. The remaining vector $\boldsymbol{\beta}$ is the result of the solution of the tight inequalities with the maximum allowed integer s. The attacker's strategy is the solution of the LP given by (8).*

We now develop a series of lemmata that lead to Theorem 2. The proof is provided in the Appendix.

Lemma 3. *Two points $\boldsymbol{x_1}$ and $\boldsymbol{x_2}$ on the polyhedron, with $\|\boldsymbol{x_1}\| = \|\boldsymbol{x_2}\|$, correspond to defender strategies $\boldsymbol{\beta_1}$ and $\boldsymbol{\beta_2}$ respectively with detection cost $\min[\Lambda\boldsymbol{\beta_1}] = \min[\Lambda\boldsymbol{\beta_2}]$ against a best responding attacker.*

Proof. We showed in Lemma 2 that a defender NE strategy $\boldsymbol{\beta}$ corresponds to one of the extreme points of a polyhedron defined by $\Lambda\boldsymbol{x} \geq \mathbf{1}_{N+1}$, $\boldsymbol{x} \geq \mathbf{0}$, with $\|\boldsymbol{x}\| = 1/z = 1/\min[\Lambda\boldsymbol{\beta}]$. Thus, for the same the norm $\|\boldsymbol{x}\|$, we get the same detection cost against a best responding attacker, i.e., $\min[\Lambda\boldsymbol{\beta_1}] = \min[\Lambda\boldsymbol{\beta_2}]$. □

Lemma 4. *If $\|\boldsymbol{x_1}\| = \|\boldsymbol{x_2}\|$ and $\boldsymbol{\mu}'\boldsymbol{x_1} < \boldsymbol{\mu}'\boldsymbol{x_2}$, then $\boldsymbol{x_1}$ corresponds to a defender strategy $\boldsymbol{\beta_1}$ with a better defendability, i.e., $\theta(\boldsymbol{\beta_1}) > \theta(\boldsymbol{\beta_2})$.*

Proof. From the definition of the defendability (Definition 2), we get

$$\begin{aligned}\theta(\boldsymbol{\beta_1}) - \theta(\boldsymbol{\beta_2}) &= \min[\Lambda\boldsymbol{\beta_1}] - \boldsymbol{\mu}'\boldsymbol{\beta_1} - (\min[\Lambda\boldsymbol{\beta_2}] - \boldsymbol{\mu}'\boldsymbol{\beta_2}) \\ &= \boldsymbol{\mu}'\boldsymbol{\beta_2} - \boldsymbol{\mu}'\boldsymbol{\beta_1} \quad (5) \\ &> 0, \quad (6)\end{aligned}$$

where (5) results from Lemma 3 (since $\|\boldsymbol{x_1}\| = \|\boldsymbol{x_2}\|$, $\min[\Lambda\boldsymbol{\beta_1}] = \min[\Lambda\boldsymbol{\beta_2}]$), and (6) follows the assumption $\boldsymbol{\mu}'\boldsymbol{x_1} < \boldsymbol{\mu}'\boldsymbol{x_2}$. The point $\boldsymbol{x_1}$ corresponds to a defender strategy $\boldsymbol{\beta_1}$ with a smaller false alarm cost, i.e., $\boldsymbol{\mu}'\boldsymbol{\beta_2} > \boldsymbol{\mu}'\boldsymbol{\beta_1}$. Hence $\theta(\boldsymbol{\beta_1}) > \theta(\boldsymbol{\beta_2})$. □

Lemma 5. *Under condition 1, an extreme point $\boldsymbol{x}$ corresponding to a defender NE strategy $\boldsymbol{\beta}$ satisfies exactly one contiguous set (of indices) of tight inequalities.*

The proof is provided in the Appendix. Let s, f be the indices of the first and last tight inequalities (of the contiguous block of tight ones) respectively.

Lemma 6. *Under condition 1, an extreme point $\boldsymbol{x}$ that corresponds to a defender NE strategy $\boldsymbol{\beta}$ has zeros before s and after $f + 1$, i.e.,*

$$x_i = 0, \forall i \in \{0, \ldots, s-1\} \cup \{f+2, \ldots, N+1\}.$$

Proof. We first show that $x_i = 0, \forall i < s$. If $\exists i \in \{0, \ldots, s-1\}$, s.t. $x_i > 0$, we reduce x_i to $\hat{x}_i$ until either $\hat{x}_i = 0$ or i^{th} inequality is tight, and increase x_{i+1} by the same amount. We maintain $\|\boldsymbol{x}\|$ constant, and in case that $x_{i+1} > 0$ we get one more tight constraint. Thus the original point is not extreme, as we can find another point whose tight constraints is a strict superset of those of the original. In case that $x_{i+1} = 0$, the new $\hat{\boldsymbol{x}}$ corresponds to a defender NE strategy with a better defendability.

We now show that $x_i = 0, \forall i > f+1$. If $\exists i \in \{f+2, \ldots, N+1\}$, s.t. $x_i > 0$, we reduce x_i until $\hat{x}_i = 0$ and increase x_{f+1} by the same amount. We again keep the norm $\|\boldsymbol{x}\|$ constant but $\hat{\boldsymbol{x}}$ has one more tight constraint, thus $\boldsymbol{x}$ was not extreme. □

Lemma 7. *In any Nash equilibrium, under conditions 1 and 2,*

1. *$f = N$, when D is non increasing.*
2. *$f = N$ or $s = f$, when D is increasing.*

The proof is provided in the Appendix.

Lemma 8. *Amongst different defender mixed strategies $\boldsymbol{\beta}$ with the same component β_{N+1}, the detection cost against a best responding attacker is the same, under conditions 1 and 2.*

Proof. By Lemma 7 under conditions 1 and 2, $f = N$ or we have pure strategies NE. By Lemma 3, the points with the same norm $\|\boldsymbol{x}\|$ correspond to defender strategies with the same detection cost ($\min[\Lambda\boldsymbol{\beta}] = 1/\|\boldsymbol{x}\|$). Scaling the last tight inequality N with the norm and since $\boldsymbol{\beta}$ is a distribution, we get $\delta(N)(1 - \boldsymbol{\beta}_{N+1}) + M(N) = \frac{1}{\|\boldsymbol{x}\|}$. Thus for the same β_{N+1}, the norm is the same, which results in the same detection cost against a best responding attacker. □

Lemma 9. *Under conditions 1 and 2, amongst defender mixed strategies with different s and same β_{N+1}, the defendability is maximal when s is maximal.*

The proof is provided in the Appendix.

Note 1. There might be more than one maximizers of the defendability. In this case, we have multiple NE strategies for the defender. But, by small perturbations of the game parameters we can prevent ties. For instance, let $\boldsymbol{\beta}_1$, $\boldsymbol{\beta}_2$ be two maximizers of the defendability, with different detection costs ($\min[\Lambda\boldsymbol{\beta}_1] > \min[\Lambda\boldsymbol{\beta}_2]$, and false alarm costs $\boldsymbol{\mu}'\boldsymbol{\beta}_1 < \boldsymbol{\mu}'\boldsymbol{\beta}_2$. Perturbing the $\boldsymbol{\mu}$ such that $\boldsymbol{\mu}'\boldsymbol{\beta}_1 \geq \boldsymbol{\mu}'\boldsymbol{\beta}_2$, we get a unique maximizer of the defendability. We can follow the same approach to break the ties among multiple defender strategies $\boldsymbol{\beta}$. This way, Nash's theorem of NE existence guarantees an attacker's NE strategy.

Attacker's NE strategy Having computed and analyzed the defender NE strategy, we now explore the spy's attack strategy. Let Λ_r be a reduced matrix, after keeping only the defender strategies in his support (columns). Similarly, let μ_{r} be the reduced false alarm vector. Then the payoff of the defender must be the same for all strategies in his support, and greater (or equal) with his respective

payoff fox all strategies outside his support. Thus, the attacker is solving the following optimization problem:

$$\begin{aligned} \underset{\boldsymbol{\alpha}}{\text{maximize}} \quad & 0 \\ \text{subject to} \quad & \boldsymbol{\alpha}' \cdot \Lambda \leq \theta \cdot \mathbf{1}_{N+2} + \boldsymbol{\mu}, \\ & \boldsymbol{\alpha}' \cdot \Lambda_r = \theta \cdot \mathbf{1} + \boldsymbol{\mu}_r \\ & \mathbf{1}'_{N+1} \cdot \boldsymbol{a} = 1, \\ & \boldsymbol{a} \geq \mathbf{0}. \end{aligned} \tag{7}$$

Since this is an LP, it can be solved in polynomial time [10]. Using the Big M method (with M big) we can transform the above problem into the following one (that is more robust with respect to noise and / or small perturbations).

$$\begin{aligned} \underset{\boldsymbol{\alpha}}{\text{maximize}} \quad & -M(\alpha_s + \alpha_x) \\ \text{subject to} \quad & \boldsymbol{\alpha}' \cdot \Lambda \leq \theta \cdot \mathbf{1}_{N+2} + \boldsymbol{\mu} \\ & \boldsymbol{\alpha}' \cdot \Lambda_r + \alpha'_s \cdot 1_{(N+1)\times R} \leq \theta \cdot \mathbf{1}'_R + \boldsymbol{\mu}_r \\ & \mathbf{1}'_{N+1} \cdot \boldsymbol{a} + \alpha_x \leq 1, \\ & \boldsymbol{a} \geq \mathbf{0}, \alpha_s \geq \mathbf{0}, \alpha_x \geq 0. \end{aligned} \tag{8}$$

To solve problem (8) we use `CVX`, a package for specifying and solving convex programs [13, 14]. `CVX` is using the simplex method to find the solution. From the Nash Equilibrium Theorem, we know that a solution exists, since the attacker will play a best response.

Depending on the degrees of freedom N and the number of defender NE strategies that are given nonzero probability in NE (R), the above procedure might give a unique or multiple $\boldsymbol{\alpha}$. This $\boldsymbol{\alpha}$ must be a valid probability distribution (sum to one and have nonnegative elements) for if otherwise, it would contradict Nash's existence theorem.

4 Evaluation with model examples

In this section, we present two characteristic examples of the above general problem and evaluate them in terms of the expected defender NE payoff.

4.1 Example model 1

In the first model, which is analyzed in [1], the spy's cost function in case of detection is $D(H) = c_{\mathrm{d}} - H \cdot c_{\mathrm{a}}$. There is a constant cost c_{d} associated with the detection and a benefit proportional to the number of attacks H. In case of missed detection, the spy gets the benefit from the attacks, without suffering from the detection cost, thus $M(H) = -H \cdot c_{\mathrm{a}}$, where c_{a} is the cost associated with a single FS attack. The spy cost is

$$J_A(T, H) = c_{\mathrm{d}} \cdot \mathbb{1}_{T \leq H} - c_{\mathrm{a}} \cdot H.$$

The defender's expected reward function depends on the true type of the attacker and following the general model analysis and is given by

$$U_D(T, H) = J_A(T, H) - \mu(T),$$

where $\mu(T) = \frac{1-p}{p} \cdot c_{\text{fa}} \cdot \phi(T)$. All lemmata that were proved in Sec. 3 hold since conditions 1 and 2 hold. Note that $M(H) = -H \cdot c_{\text{a}}$ is a decreasing function with respect to H and $\delta(H) = c_{\text{d}}$ is constant. Thus there is a contiguous block of tight inequalities starting at index s and finishing at index N with $x_i = 0,\ \forall i \in \{0, \ldots, s-1\}$, or we have pure NE. Furthermore, the defender's NE strategy exists amongst the two forms in Table 2. The proof is given in [1].

Table 2. Defender's strategy in NE ($\beta_m = c_{\text{a}}/c_{\text{d}}$)

#	$\ldots$	β_s	β_{s+1}	$\ldots$	β_N	β_{N+1}
1.	0	0	β_m	β_m	β_m	$1-(N-s)\beta_m$
2.	0	$1-(N-s)\beta_m$	β_m	β_m	β_m	0

4.2 Example model 2

In this second variation of the model, we assume that the defender maintains some logs on the type of occurred attacks. When a spy is detected, the defender has the appropriate tools to investigate the attacker's behavior. This way, the defender has the opportunity to learn about the spy's true intentions (which specific target/information he seeks to extract from the file server), his location or identity and his future attack pattern, in case he is not immediately expelled.

Each of the H FS hits now gives the spy a benefit of c_{a} only if he evades detection. In case he is correctly identified, each FS attack yields a cost of c_{a} for the spy, as they reveal the intentions of the spy. Thus $D(H) = c_{\text{d}} + H \cdot c_{\text{a}}$, and $M(H) = -H \cdot c_{\text{a}}$, giving the spy a cost function of

$$J_A(T, H) = (c_{\text{d}} + 2c_{\text{a}} \cdot H) \cdot \mathbb{1}_{T \leq H} - c_{\text{a}} \cdot H.$$

Following the analysis for the general model, the defender payoff function is

$$U_D(T, H) = J_A(T, H) - \mu(T).$$

All lemmata that were proved in Sec. 3 hold since conditions 1 and 2 hold. Note that $M(H) = -c_{\text{a}} \cdot H$ is a decreasing function with respect to H and $\delta(H) = D(H) - M(H) = c_{\text{d}} + 2Hc_{\text{a}}$ is increasing with respect to H. Thus there is a contiguous block of tight inequalities starting at index s and finishing at index N with $x_i = 0,\ \forall i \in \{0, \ldots, s-1\}$, or we have pure NE.

Defender's strategy in example model 2. After subtracting the two tight inequalities N and $N-1$, we get $\beta_{N+1} \geq 1/2$, because a tight inequality $(N-1)$ also suggests that $\beta_{N-1} \geq 0$. Thus in either case, it must be that $\beta_{N+1} \geq 1/2$. The upper bound for $\boldsymbol{\beta}$ is 1. But since the index of the first tight inequality is an integer, only certain values of β_{N+1} result in an optimal s, which is also an integer.

By Theorem 2, given a certain β_{N+1} for the defender NE strategy, we need to find the largest possible s such that inequality s is tight and $(s-1)^{\text{th}}$ is loose, with $\beta_0 = \ldots = \beta_{s-1} = 0$. Subtracting the tight inequality N from the loose inequality $s-1$, we get $s \leq \frac{(c_\text{a}-Nc_\text{a})+(c_\text{d}+2Nc_\text{a})\beta_{N+1}-c_\text{d}}{c_\text{a}}$. Since s must be an integer, $\beta_{N+1} = \frac{(N-1+k)c_\text{a}+c_\text{d}}{c_\text{d}+2Nc_\text{a}}$, with k integer. Thus the search over the optimal β_{N+1} has a linear complexity with respect to N, with $\beta_{N+1}^{\min} = 1/2$, $\beta_{N+1}^{\max} = 1$ and $step = \frac{(N-1+k)c_\text{a}+c_\text{d}}{c_\text{d}+2Nc_\text{a}}$. Alternatively, solving for the integer k, we get $k^{\min} = 1 - N - c_\text{d}/c_\text{a}$ and $k^{\max} = N+1$.

5 Parameter Effects in the Game

The two previously presented models have an essential difference: While in the first model, the spy benefits from the FS attacks regardless of the defender's classification decision, in the second model, the spy benefits from the FS attacks only when he is misclassified. The assumption under the second model is that the defender has invested in forensic techniques, and is able to identify, preserve and analyze attacks within the network. This way, each FS attack reveals information about the identity and the intention of the attacker.

Computer forensics is a costly investment, thus the defender needs to decide under which circumstances he should develop such tools. By comparing the two above models, and essentially the defender's expected payoff in NE, we extract key insights on the expected gain from the forensics. The crucial parameters are prevalent in both models, like p, c_fa, θ_0, c_a, c_d. The critical difference is that the spy's cost function in case of detection $D(H)$ is decreasing in the first model and increasing in the second model, with respect to H.

Some network parameters are correlated in sets of combinations, in the sense that a change in any element on the set alters the Nash equilibrium and payoffs of the players toward the same direction. For instance, looking at the defender's payoff function for the first model, we observe that a change in p or c_fa affects the false alarm penalty, thus changing p and keeping c_fa constant will provide us with the same implications as if we kept p the same and changed c_a. Thus it makes sense to investigate the impact of only a subset of the parameters.

5.1 Effect of the probability of the spy p and false alarm cost c_fa

We expect that when p is small, the defender will suffer a small cost from potential FS attacks. As p increases it becomes more difficult to distinguish between

spy and spammer and the defender's payoff will be decreasing. When p becomes larger, the defender will classify him correctly and receive a higher payoff.

Indeed, in Fig.1(a) we observe two areas of different behavior. When $p < 0,5$, the defender payoff function is decreasing whereas it is increasing as the spy's probability reaches $p = 0.5$. We also observe that as p increases, the spy's attack policy becomes more aggressive and the defender reduces his threshold to catch the more-probable-to-exist spy. In Fig. 1(b) we note that as p increases, the benefit of investing in forensics (and employ model 2) is an increasing function on p. We note here that the depicted values for the defender payoff are the unscaled and unshifted initial payoffs, as expressed in (1).

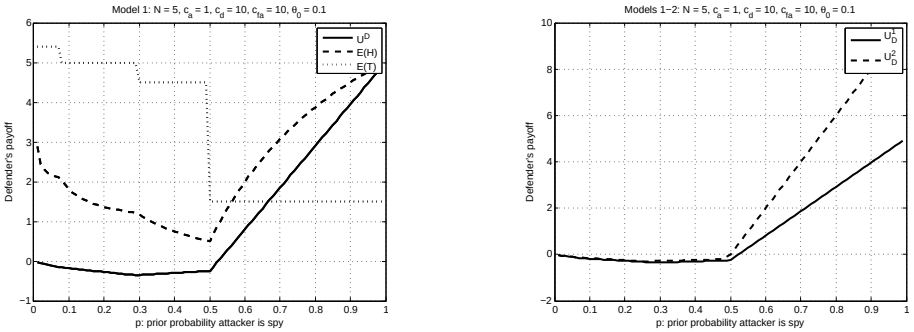

(a) NE defender payoff first decreases and then increases on p

(b) Comparison of models 1-2

Fig. 1. As p increases, the NE defender gain with model 2 increases

5.2 Effect of the detection cost c_{d}, the classification window N and single FS attack cost c_{a}

In Model 1, when the cost of detection c_{d} is small compared with the maximum achievable gain from the FS attacks ($N \cdot c_{\mathrm{a}}$), the spy does not care about getting detected and is attacking with his maximum allowed strength (N times). On the contrary, in Model 2, where the spy suffers a cost proportional to his attack aggressiveness in case of detection, the spy is more conservative with his attacks. This difference is depicted in Fig. 2(a).

As we can see, in model 1, the cost of detection is so small, than the attacker always attacks N times. On the other hand, the defender selects a threshold equal to the pure strategy of the spy and detects him. If the defender selected $T = N + 1$ or $T = N - 1$ instead of $T = N$ as his classification threshold, he would miss the spy and would have smaller payoff due to the increased false alarm, respectively. In the second model, though, the spy takes into consideration the potential benefit his FS attacks would give the defender. The spy is less

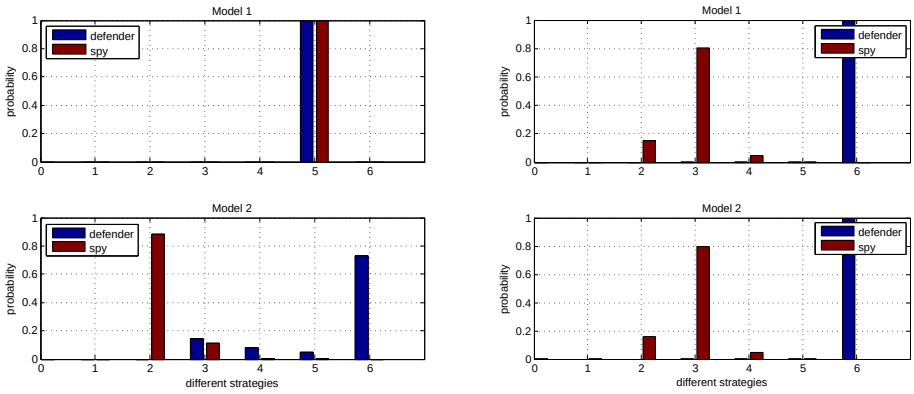

(a) small c_{d} (b) small c_{a}

Fig. 2. A Difference of the two models

aggressive, and attacks fewer times. Other parameters of the game are $N = 5$, $c_{\mathrm{d}} = 1$, $c_{\mathrm{a}} = 1$, $p = 0.1$, $c_{\mathrm{fa}} = 10$, and $\theta_0 = 0.1$. We also note here that the spy's strategy is a weighted truncated binomial distribution. Every defender's strategy in his NE support gives the defender the same payoff. Thus the difference in the false alarm penalty for the different thresholds matches the difference in the misdetection cost. For instance $\Pr\{H = 3\} = \frac{(1-p)\cdot c_{\mathrm{fa}}\cdot[\phi(3)-\phi(4)]}{c_{\mathrm{d}}+2\cdot 3\cdot c_{\mathrm{a}}} = 0.1041$.

When c_{a} is small, (or else when c_{d} is most important than $N \cdot c_{\mathrm{a}}$), we observe that the two models result in the same strategies for the two players (Fig. 2(b)). Indeed, when the spy expects not to reveal a lot of information to the defender if he gets detected, he will act as if there was not risk (as in model 1). Thus, when the defender expects to lose little from the FS hits, he will avoid investing in forensics to learn more about the intentions of the spy.

5.3 Effect of the spammer's distribution parameter θ_0

In these two models we have assumed a specific distribution on the FS attacks for the naive player, i.e., the spammer. Each time slot (period) of the available N time slots, the spammer attacks the FS with a frequency of θ_0. In the case that θ_0 is small (the spammer is mostly interested in attacking the MS instead of FS) the task of the defender to differentiate between the two types of attackers becomes easier.

On the contrary, if the spammer is attacking with a high θ_0 each period, then the defender is hurt from the false alarms, since he will be confused from the large number of FS hits and will classify the attacker as spy. We can see this difference in the defender NE payoff as θ_0 increases.

In Fig. 3 we see the effect of the spammer's strategy, essentially θ_0, on the two players' NE strategies. In both models, as θ_0 increases, the spy becomes more aggressive (to imitate the spammer's behavior). As θ_0 increases, the spam-

mer attacks the FS more frequently, and it is more difficult for the defender to distinguish the two types of attacker. The spy then exploits this uncertainty to increase his payoff (by attacking more times). When θ_0 is small, the defender sets the threshold low for spy classification. As θ_0 increases, his false alarm penalty gets smaller and the defender assigns a larger weight to the "always classify as a spammer" strategy.

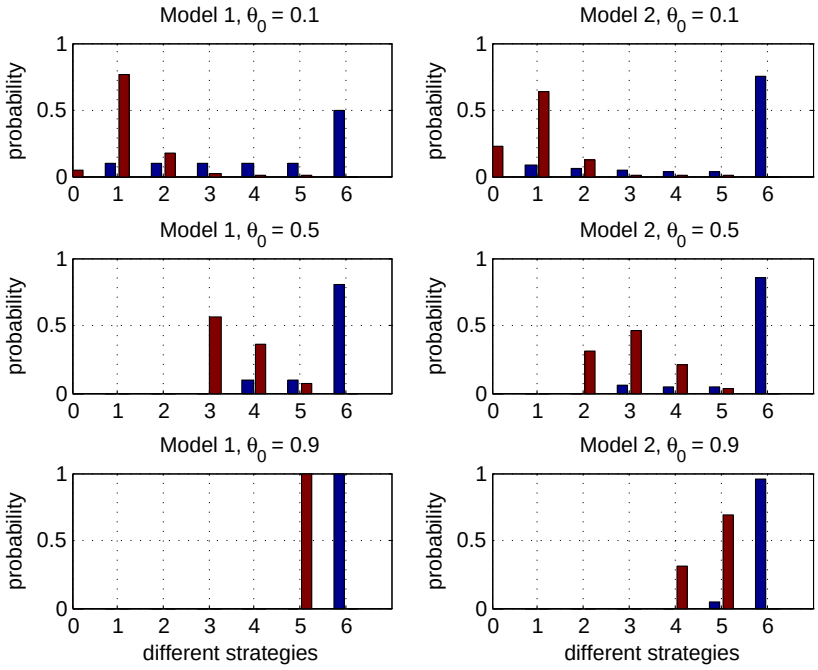

Fig. 3. Effect of θ_0 on the players' NE strategies ($N = 5$, $c_{\mathrm{a}} = 1$, $c_{\mathrm{d}} = c_{\mathrm{fa}} = 10$, $p = 0.3$). The bar left and right of numeral represents the defender and the spy respectively.

In Fig. 4 we see the effect of the θ_0 on the defender's payoff for the two models, for various values of the prior probability of the spy p. In model 1, depicted in Fig. 4(a), 4(b), we observe that as θ_0 increases, the defender's NE payoff decreases for any value of p, because higher θ_0 signifies a higher false alarm penalty for the defender. In contrast, the second model depicted in Fig. 4(b), the above rule applies only for the ranges of θ_0 below $\theta_0 = 0.5$. For $p > 0.5$, the defender will always select the same pure strategy, that yields the same payoff.

6 Conclusion

We investigate a classification game, where a network administrator (defender) seeks to classify an attacker as a strategic spy or a naive spammer. We first

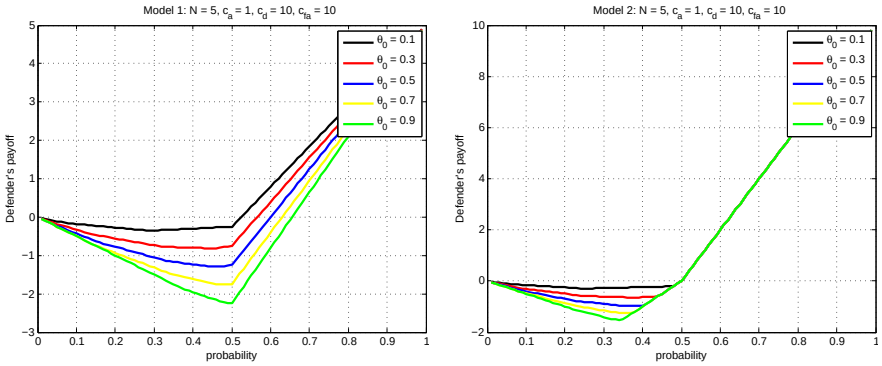

(a) Model 1 (b) Model 2

Fig. 4. The defender's NE payoff decreases as θ_0 increases for all values of p for model 1, but only for $p < 0.5$ for model 2.

prove that a nonzero-sum game with general payoff functions that satisfy some conditions can lead to a NE computation in polynomial time. Our approach characterizes the structure of the best response strategies of the two players and explains the intuition for the resulting strategies. We investigate two specific game models: model 1 is a simpler game, where the spy benefits from his attacks, regardless of the defender's classification decision. In model 2, the defender is equipped with forensic tools and the spy only benefits from his attacks upon a misclassification. By analyzing these two games, we extract important information about when the defender should invest in forensics and how the strategies of two players in NE are affected by the various control parameters of the game.

References

1. L. Dritsoula, P. Loiseau, and J. Musacchio,"A game-theoretic approach for finding optimal strategies in an intruder classification game," to appear in *Proc. of the 51th IEEE Conf. Decision and Control (CDC)*, December 2012.
2. "Cyber Security Research Report," Bit9, 2012.
3. "TMT Global Security Study Key Findings," Deloitte, 2011.
4. M. H. Manshaei, Q. Zhu, T. Alpcan, T. Basar, and J. -P. Hubaux, "Game Theory Meets Network Security and Privacy," Ecole Polytechnique Federale de Lausanne (EPFL), Tech. Rep. EPFL-REPORT-151965, April 2011.
5. T. Alpcan, and T. Başar, "A Game Theoretic Approach to Decision and Analysis in Network Intrusion Detection," in *Proc. of the 42nd IEEE Conf. Decision and Control*, December 2003, pp. 2595–2600.
6. L. Chen, J. Leneutre, "A game theoretical framework on intrusion detection in heterogeneous networks," in *IEEE Transactions on Information Forensics and Security*, v.4 n.2, pp.165-178, June 2009.

7. A. Gueye, J. C. Walrand, and V. Anantharam, "A Network Topology Design Game: How to Choose Communication Links in an Adversarial Environment?," in *GameNets*, April 2011.
8. A. Gueye, "A Game Theoretical Approach to Communication Security," PhD dissertation. University of California, Berkeley, Electrical Engineering and Computer Sciences, March 2011.
9. N. Dalvi, P. Domingos, Mausam, S. Sanghai, and D. Verma, "Adversarial classification," in *Proc. of the ACM SIGKDD*, pp. 99–108, 2004.
10. D. G. Luenberger, "Linear and Nonlinear Programming." 2nd ed. A. W., 1984.
11. Gambit, "Gambit game theory analysis software and tools," http://www.hss.caltech.edu/gambit, 2002.
12. J. Nash, "Non-Cooperative Games, " The Annals of Mathematics 54(2):286-295, 1951.
13. M. Grant and S. Boyd. "CVX: Matlab software for disciplined convex programming, version 1.21. ../../cvx", April 2011.
14. M. Grant and S. Boyd. "Graph implementations for nonsmooth convex programs," in *Recent Advances in Learning and Control (a tribute to M. Vidyasagar)*, V. Blondel, S. Boyd, and H. Kimura, editors, pages 95-110, *Lecture Notes in Control and Information Sciences*, Springer, 2008.

Appendix: Omitted Proofs from Section 3

Proof (Proof of Lemma 5). An extreme point $\boldsymbol{x}$ satisfies at least one tight inequality. If none of the inequalities are tight, we scale the vector $\boldsymbol{x}$ down until one inequality becomes tight. The new vector's set of tight inequalities is a strict superset of those of the original vector, thus the point with no tight inequalities is not extreme. Let there be two tight inequalities with indices s and $s+k$ and let all their intermediate inequalities be loose. There exist two possible cases:

1. $\exists i \in \{1, \ldots, k-1\}$, with $x_{s+i} > 0$. We make the following transformation that increases the defendability. We reduce x_{s+i} by a small amount $\epsilon > 0$ and increase x_{s+i+1} by the same amount, maintaining the same norm $\|\boldsymbol{x}\|$. All the inequalities before and after the one with index $(s+i)$ are intact, while the previously loose inequality $(s+i)$ is now tight. For the new vector $\hat{\boldsymbol{x}}$ it holds that $\boldsymbol{\mu}'\hat{\boldsymbol{x}} < \boldsymbol{\mu}'\boldsymbol{x}$, since $\boldsymbol{\mu}$ is a vector with decreasing values and we have shifted some weight from x_{s+i} to x_{s+i+1}. By Lemma 4, the new point corresponds to a defender NE strategy with a better defendability. We continue the above procedure until there are no loose inequalities between the initial tight ones, or until $x_{s+i} = 0, \forall i \in \{1, \ldots, k-1\}$.
2. $x_{s+i} = 0$, $\forall i \in \{1, \ldots, k-1\}$. Subtracting the first tight inequality (of index s) from any loose inequality of index $s+i$, with $i \in \{1, \ldots, k-1\}$, we get

$$\begin{aligned} \Delta_1\delta(s) \cdot (x_0 + \ldots + x_s) + \Delta_1 M(s)\|\boldsymbol{x}\| > 0 \\ \vdots \\ \Delta_{k-1}\delta(s) \cdot (x_0 + \ldots + x_s) + \Delta_{k-1} M(s)\|\boldsymbol{x}\| > 0. \end{aligned} \tag{9}$$

   Similarly, subtracting the last tight inequality $(s+k)$ from all the loose inequalities of index $s+i, \forall i \in \{1, \ldots, k-1\}$, we get

$$\begin{aligned}
&\Delta_{k-1}\delta(s+1)\cdot(x_0+\cdots+x_s)+\Delta_{k-1}M(s+1)\|\boldsymbol{x}\|<-\delta(s+k)x_{s+k}\\
&\qquad\qquad\vdots\\
&\Delta_1\delta(s+k-1)\cdot(x_0+\cdots+x_s)+\Delta_1M(s+k-1)\|\boldsymbol{x}\|<-\delta(s+k)x_{s+k}.
\end{aligned}\tag{10}$$

Under condition 1, the set of equations (9) and (10) cannot be satisfied simultaneously. Indeed, the last equation of (10) gives $0>\Delta_1\delta(s+k-1)\cdot(x_0+\ldots+x_s)+\Delta_1M(s+k-1)\|\boldsymbol{x}\|>\Delta_1\delta(s)\cdot(x_0+\ldots+x_s)+\Delta_1M(s)\|\boldsymbol{x}\|$, which contradicts the first equation of (9). □

Proof (Lemma 7). Suppose that $f<N$. Then the inequality of index $(f+1)$ exists, is loose and all positivity constraints are satisfied. Subtracting the tight inequality of index f from the loose inequality of index $(f+1)$, we get

$$x_{f+1}>\frac{[D(f)-D(f+1)]\cdot\|\boldsymbol{x}\|}{\delta(f)}\triangleq C.\tag{11}$$

1. If D is non increasing, since δ is positive, $C\geq 0$ and $x_{f+1}>0$. We consider the following transformation

$$\hat{x}_i=\begin{cases}x_i & \text{for } i\in\{0,\ldots,f\}\cup\{f+3,\ldots,N+1\}\\ C & \text{for } i=f+1\\ x_{f+1}-C & \text{for } i=f+2.\end{cases}\tag{12}$$

With the above transformation we get

$$\begin{aligned}\boldsymbol{\mu}'(\hat{\boldsymbol{x}}-\boldsymbol{x})&=\mu_{f+1}\cdot(\hat{x}_{f+1}-x_{f+1})+\mu_{f+2}\cdot(\hat{x}_{f+2}-x_{f+2})\\&=\mu_{f+1}\cdot(C-x_{f+1})+\mu_{f+2}\cdot(x_{f+1}-C-0)\\&=(x_{f+1}-C)\cdot(\mu_{f+2}-\mu_{f+1})\\&<0,\end{aligned}$$

since $x_{f+2}=0$, $x_{f+1}>C$, and $\boldsymbol{\mu}$ is a strictly decreasing vector ($\mu_{f+2}<\mu_{f+1}$). Hence, for the new point $\hat{\boldsymbol{x}}$, $\|\hat{\boldsymbol{x}}\|=\|\boldsymbol{x}\|$, but $\boldsymbol{\mu}'\hat{\boldsymbol{x}}<\boldsymbol{\mu}'\boldsymbol{x}$. By Lemma 4 point $\hat{\boldsymbol{x}}$ corresponds to a defender NE strategy with a better defendability. We can continue making the above transformation until $f=N$.

2. If D is increasing, then $C<0$ and $x_{f+1}\geq 0$. If $x_{f+1}>0$, while $f<N$ we can shift a small amount ϵ from x_{f+1} to x_{f+2}, keeping the same norm but getting a better defendability. We keep making the above transformation until $f=N$. If $x_{f+1}=0$, then $\|\boldsymbol{x}\|=x_s+\ldots+x_f$. Subtracting the two tight inequalities (s) and (f) and since $D(H)$ is an increasing function,

$$x_s=\frac{[D(f)-M(s)]\cdot\|\boldsymbol{x}\|}{\delta(s)}>\frac{[D(s)-M(s)]\cdot\|\boldsymbol{x}\|}{\delta(s)}=\frac{\delta(s)\cdot\|\boldsymbol{x}\|}{\delta(s)}=\|\boldsymbol{x}\|,$$

or $x_s>\|\boldsymbol{x}\|$. Contradiction, unless $s=f$. □

Proof (Lemma 9). Let $\boldsymbol{\beta}$, $\hat{\boldsymbol{\beta}}$ be two different defender NE strategies with $\beta_{N+1} = \hat{\beta}_{N+1}$. By Lemma 8, since β_{N+1} is the same for both vectors, the cost of detection is the same. Let $\hat{s} = s-1$. We will show that the false alarm penalty for the largest index s is larger, i.e., $\boldsymbol{\mu'} \cdot (\beta_s + \ldots + \beta_{N+1}) < \boldsymbol{\mu'} \cdot (\hat{\beta}_{\hat{s}} + \ldots + \beta_{N+1})$. Subtracting the tight inequalities N and $(N-1)$, results in $\beta_N = \hat{\beta}_N$. Similarly, iteratively subtracting the tight inequalities $(s+k)$ and $(s+k+1)$, $\forall k \in \{1, \ldots, N-s-1\}$ results in $\beta_{s+k} = \hat{\beta}_{s+k}$. By Lemma 6, $\beta_{s-1} = \ldots = \beta_0 = 0$ and $\hat{\beta}_{s-2} = \ldots = \hat{\beta}_0 = 0$. Thus the two different NE strategies differ only in β_{s-1}, and β_s. The remaining weight is the same for both vectors ($\beta_{s-1}+\beta_s = \hat{\beta}_{s-1}+\hat{\beta}_s = 1-\sum_{i=s+1}^{N+1} \beta_i$). In the case of the vector $\hat{\boldsymbol{\beta}}$, this weight is divided into two different components ($\hat{\beta}_{s-1}$ and $\hat{\beta}_s$) while in the case of $\boldsymbol{\beta}$ it is all assigned into the component with index s. Since $\mu_s < \mu_{s-1}$, the vector $\boldsymbol{\beta}$ with the largest index of the first tight inequality s will provide a smaller false alarm cost, and hence a greater defendability. □

Proof (Theorem 2). Depending on the nature of the cost functions, there are two potential constructions for the defender NE strategy $\boldsymbol{\beta}$. By Theorem 1). We select the one that yields the maximal defendability.

1. Mixed strategies NE with $f = N$. By Lemma 9, defender strategies $\boldsymbol{\beta}$ with the same β_{N+1} yield the maximal defendability when s is maximal. Thus, we need to find the largest possible s such that the inequalities 0 through $(s-1)$ are loose and $x_0 = \ldots = x_{s-1} = 0$. Since we are in mixed NE strategies, there exist at least two tight inequalities. Starting from the last tight inequality N and subtracting the next tight inequality $N-1$, we compute β_N. In general, subtracting the i, $(i-1)$ inequalities, we compute $\beta_i = \frac{D(i-1)-D(i)+[\delta(i)-\delta(i-1)]\cdot\sum_i^{N+1}\beta_i}{\delta(i-1)}$. In every step we check whether the previous inequality $(s-1)$ can be loose. If this is possible, then we assign all the remaining weight to $\beta_s (\beta_s = 1-\sum_{i=s+1}^{N+1}\beta_i)$. Since the block of tight inequalities that ranges from s through N (integers) is unique, only a certain number of selections on β_{N+1} will produce valid vectors $\boldsymbol{\beta}$ (with unit norm and nonnegative weights). Thus we need to solve the following equations
$$\begin{gathered}\delta(s-1)\cdot 0 + M(s-1) > 1/\|\boldsymbol{x}\| \\ \delta(s)\cdot\beta_s + M(s) = 1/\|\boldsymbol{x}\| \\ \vdots \\ \delta(N)\cdot(1-\beta_{N+1}) + M(N) = 1/\|\boldsymbol{x}\|.\end{gathered}$$
Subtracting the tight inequality N from the $(s-1)$ loose inequality we get $M(s-1) > \delta(N)\cdot(1-\beta_{N+1}) + M(N)$. Solving for the integer s, we compute the increments of β_{N+1} that give a valid distribution $\boldsymbol{\beta}$.
2. Pure NE with $s = f$. This case implies that when D is an increasing function, a pure defender strategy maximizes the defendability. For each selection of s in $\{0, \ldots, N\}$, we compute the defendability of the resulting strategy $\boldsymbol{\beta}$ ($\beta_s = 1$), and select the strategy that maximizes the defendability.

Given the defender strategy $\boldsymbol{\beta}$, the attacker is solving his LP (8) and selects his strategy $\boldsymbol{\alpha}$. Nash's existence theorem guarantees a Nash equilibrium, thus the LP will always provide a valid solution. □