

Effective Cost Allocation for Deterrence of Terrorists

Eugene Lee Quan

Susan Martonosi, Advisor

Francis Su, Reader

May, 2007

HARVEY MUDD
COLLEGE

Department of Mathematics

Copyright © 2007 Eugene Lee Quan.

The author grants Harvey Mudd College the nonexclusive right to make this work available for noncommercial, educational purposes, provided that this copyright statement appears on the reproduced materials and notice is given that the copying is by permission of the author. To disseminate otherwise or to republish requires written permission from the author.

Abstract

The attacks on the World Trade Center in New York, the subway and bus bombings in London, and the suicide bombings in Casablanca are only a few of the examples in which in recent years, terrorists have opted to attack multiple targets at once. Often, their strong determination to attack makes it impossible to completely deter terrorists from attacking altogether, and instead, counterterrorist units must consider how to defend targets effectively to minimize damages. We attempt to model a version of this scenario by presenting a two-target sequential game where two players try to attack and defend the targets respectively. The probability of successfully destroying a target is a function of resource allocations from both players, who are also subject to budget constraints. We attempt to find the defender's strategy that will minimize expected damages by first exploring the attacker's optimal strategy. We show that the attacker's decision to attack only one or both targets is dependent on the size of the attacker's allowed budget relative to other game parameters, and use that information to evaluate the defender's strategy. We also numerically determine the optimal defender security investment, as well its sensitivity to other game parameters. We conjecture that as the damage and expected reward at a target increases, the defender's allocation towards that target tends to increase, while an increase in the punishment results in the opposite effect. Such conjectures allow for the creation of a flexible defense policy in the more applicable bigger picture.

Acknowledgments

I would like to thank Professor Susan Martonosi for her guidance throughout this project, Professor Francis Su for serving as my second reader, and Professor Andrew Bernoff for organizing Senior Thesis for the 2006–2007 school year. Claire Connelly for her technical support.

Contents

Abstract	iii
Acknowledgments	v
1 Introduction	1
2 The Model	5
2.1 Rules and parameters	5
2.2 Attacker's expected benefit	8
2.3 Defender's expected damage	9
3 Attacker's Optimal Strategy	11
3.1 undefended targets	12
3.2 Defended targets	13
3.3 Other characteristics of the attacker's strategy	26
4 Defender's Optimal Strategy	29
4.1 Known reward parameters	30
4.2 Unknown reward parameters	31
5 Future Work	39
6 Conclusions	43
Bibliography	45

List of Figures

2.1	Game Tree of Model	8
3.1	Attacker's Optimal Allocation: Large Attacker Budget	17
3.2	Attacker's Optimal Allocation: Small Attacker Budget: Sym- metric Parameters	21
3.3	Attacker's Optimal Allocation: Small Attacker Budget: Asym- metric Parameters	25
4.1	Small Attacker Budget: Defender's Optimal Allocation	31
4.2	Defender's Optimal Allocation as Punishment Increases . . .	33
4.3	Defender's Optimal Allocation as Damage Increases	34
4.4	Defender's Optimal Allocation as Expected Reward Increases	35
4.5	Volatility in Defender's Expected Damage	36
4.6	Defender's Optimal Allocation as Attacker Budget Increases	37
5.1	Attacker's Optimal Allocation: Medium Attacker Budget . . .	40

Chapter 1

Introduction

On September 11, 2001, terrorists successfully destroyed the Twin Towers of the World Trade Center in New York City, USA, using hijacked airplanes. On May 16, 2003, several restaurants, a hotel, as well as a Jewish community center were attacked by suicide bombers. Numerous subway trains as well as a bus in London, England, were bombed as a result of coordinated terrorist attacks on the morning of July 5, 2005. These are only several of the examples that show, in recent years, that terrorists do continue to plan attacks against multiple targets.

Such instances have resulted in the increase of studies on terrorism. Research has ranged from understanding how to seek and destroy terrorist networks to developing effective defensive measures [Woo (2003)]. For example, the Department of Homeland Security has initiated a program to prepare and respond to acts of terrorism by financially assisting urban areas that are perceived to be at risk. However, the Department of Homeland Security has also been criticized for inadequately calculating such risk and thus, disproportionately providing financial resources [Willis (2006)]. Examples such as this call for a more organized and systematic approach to determining risk, as well as studies on deterring terrorism in general.

In conducting such research, it is important to develop formal definitions and a strong understanding of risk as well as its other components. Risk is defined to be as the product of threat, vulnerability, and consequence, where the three are defined as the probability an attack occurs, the probability an attack results in damage, and the expected damage respectively [Willis (2006)]. We attempt to model a situation which evaluates methods of minimizing risk by looking at all three aforementioned factors.

Since terrorists often simultaneously select multiple targets for attack,

terrorist deterrence strategies must consider how defense measures at one target will affect the terrorists' decisions to attack the remaining targets. Sometimes, their determination to attack is so strong that complete deterrence is impossible. In that case, the question becomes *at which* targets the attack will occur as opposed to *if* the attacks will occur.

Our situation involves the terrorists choosing at least one of two targets for attack, implying that total deterrence is impossible. The defender will invest a security allocation for protecting one or both targets. The attacker will then respond to the defender's decision by allocating resources for attacking the targets. Both the defender and the attacker are subject to budget constraints with sunk costs; in other words, both are required to use a certain amount for defense and attack respectively. The probability of a successful attack is dependent on both allocations, and there exist reward and damage parameters in the event of successful and failed attacks respectively; hence, these parameters incorporate the components of risk defined by Willis (2006).

Prior literature has suggested that this type of game theoretic model is more appropriate than a reliability theoretic one [Bier (2004)]. Game theoretic models have considered defenders and attackers as two players in a game; scenarios have included multiple targets for attack, both in parallel and in series, and both perfect and limited attacker-knowledge regarding the targets [Abhichandani and Bier (2005)]. Martonosi and Walton (2006) look at a sequential model where both the attacker and the defender pick attack and defensive allocations for a single target case. Bier et al. (2006) consider a sequential two-target model where only the defender has an allocation; the attacker can attack only one target, and makes his decision based on the defender's allocation. Sandler (2005) examine a similar model where side effects from a target being attacked are considered. Bier and Zhuang (2006) look at a two-target case with both defender and attacker allocations, and the attacker valuations of the targets are *known* to the attacker. The aforementioned models incorporate the probability of a successful attack based on the allocations as well. My model utilizes these ideas in a sequential two-target scenario, where *both* the defender and the attacker will choose allocations, and the attacker valuations are *unknown* to the defender.

While the model presented in this thesis assumes sequential turns, a similar one with simultaneous moves known as the Colonel Blotto game also exists [Shubik and Weber (1981)]. The Colonel Blotto game consists of two players and n battlefields, and each player allocates forces to the battlefields. Each player wishes to maximize the number of fields secured; many

models consider the probability of capturing a field as a function of the number of allocated forces. Other models incorporate a value function for the number of battlefields secured, as well as simultaneous generalizations of the game [Shubik and Weber (1978), Coughlin (1992)]. My work is a *sequential* version of the aforementioned models, with the two targets and the investments being analogous to the battlefields and the forces respectively.

Our model attempts to find the security investment for the defender that minimizes his total expected damage. To do this, we begin by determining the attacker's strategy that maximizes his expected benefit, and use that information to determine how the defender should act in order to force the attacker into an optimal scenario for himself.

We first find that the attacker's optimal allocation is dependent on the relative magnitude of his allowed budget, and solve for the attacker's strategy in two cases: when his allowed budget is "large", and when it is "small". We show the attacker will attack both targets in a symmetric parameter case where the attacker's budget is sufficiently large (Theorem 3.4). We use this information to demonstrate that the defender can minimize expected damages by defending both targets equally (Theorem 4.1). We also determine sufficient criteria for the attacker's optimal allocation in an asymmetric parameter case. If the attacker budget is sufficiently small, we prove that the attacker will attack only one target (Theorem 3.10). In addition, we look at how the attacker's strategy changes with alterations in the game parameters (Section 3.3.1).

We then proceed to determine the security investment that minimizes the defender's expected damage numerically (Section 4.2). In doing so, we also determine how that investment responds to changes in the game parameters. Because one must often approximate ranges for the game parameters in real-world scenarios, such sensitivity analysis is important in determining an applicable and flexible defense investment strategy. While we have not yet rigorously proven results regarding the optimal defense investment, our numerical simulations provide a strong starting point for creating the aforementioned defense policy.

Chapter 2

The Model

2.1 Rules and parameters

In our model, there exist two players: an attacker who wishes to attack two targets, and a defender who wishes to protect them. The defender begins by spending $c_1 \geq 0$ and $c_2 \geq 0$ towards defending targets 1 and 2 respectively. Note that $c_i = 0$ implies that the defender chooses to leave target i undefended. The defender is subject to a budget constraint

$$c_1 + c_2 = c_M > 0,$$

where c_M is a constant known to both the defender and the attacker from the start of the game. Note that this is a sunk cost: the defense allocations towards targets 1 and 2 must always sum to c_M . We define c_M as the **defender budget**. We will also refer to the defender allocation by $\vec{c} = (c_1, c_2) = (c_1, c_M - c_1)$.

The attacker observes the security investment \vec{c} . The attacker then responds by spending $x_1 \geq 0$ and $x_2 \geq 0$ towards attacking targets 1 and 2 respectively. If the attacker chooses not to attack target i , then $x_i = 0$. Analogous to the defender budget constraint is the attacker budget constraint

$$x_1 + x_2 = x_M > 0,$$

where x_M is also a constant known to both players from the start of the game. Therefore, the attacker budget is also a sunk cost. We define x_M as the **attacker budget**, and refer to the attack allocation as $\vec{x} = (x_1, x_2) = (x_1, x_M - x_1)$.

One observation to make is that if x_1 increases, x_2 must decrease since $x_2 = x_M - x_1$. Similarly, as c_1 increases, c_2 decreases.

We will also refer to the term **feasible value** for (c_1, c_2) or (x_1, x_2) to mean any pair of allocations $c_1, c_2 \geq 0$ or $x_1, x_2 \geq 0$ such that $c_1 + c_2 = c_M$ or $x_1 + x_2 = x_M$.

2.1.1 Probability of success

Once the attacker has attacked, each target i is, independently, successfully destroyed or not. The probability of target i being successfully destroyed depends on the defense and attack allocations towards target i ; we call this probability $p_i(x_i, c_i)$. In other words, the probability that target i is successfully destroyed given that the attacker and defender allocate x_i and c_i respectively is $p_i(x_i, c_i)$.

We now make assumptions about $p_i(x_i, c_i)$. Note that although p_1 is a function of x_1 , we can still differentiate p_1 with respect to x_2 since $x_2 = x_M - x_1$, for we shall do so later in the paper.

- $p_1(x, c) = p_2(x, c)$ for all x and c . Both targets inherently have the same probability of success function. Without loss of generality, any property that holds for p_1 with respect to x_2 also holds for p_2 with respect to x_1 .
- p_i is continuous and twice differentiable with respect to x_i and c_i when $0 < x_i < x_M$ and $0 < c_i < c_M$.
- $p_i(x_i, 0) = 1$ if $x_i > 0$. Attacking an undefended target guarantees success.
- $p_i(0, c_i) = 0$ for all possible values of c_i . A target that is not attacked cannot be destroyed.
- $\frac{dp_1}{dx_1} > 0$ and $\frac{dp_1}{dx_2} < 0$. Obviously, as the investment towards attacking target 1 increases, the probability of successfully destroying target 1 increases. As the investment towards attacking target 2 increases, the investment towards target 1 decreases, resulting in the opposite phenomenon.
- $\frac{dp_1}{dc_1} < 0$ and $\frac{dp_1}{dc_2} > 0$. The opposite occurs when the defense investment towards target 1 increases (decreases).
- $\lim_{c_i \rightarrow \infty} p_i(x_i, c_i) = 0$. When an infinite amount is invested into defending a target, that target cannot be successfully attacked.

- $\lim_{x_i \rightarrow \infty} p_i(x_i, c_i) = 1$. The opposite phenomenon occurs when an infinite amount is invested into attacking a target.
- $\frac{d^2 p_1}{dx_1^2} < 0$. p_1 is concave down with respect to x_1 . This follows the economic law of diminishing returns when the investment into attacking target 1 increases.
- $\frac{d^2 p_1}{dx_2^2} < 0$. p_1 is concave down with respect to x_2 ; this is because p_1 decreases at an increasing rate as x_1 decreases, or as x_2 increases
- $\frac{d^2 p_1}{dx_1 dx_2} > 0$. This holds since $\frac{d^2 p_1}{dx_1^2} < 0$, and $x_2 = x_M - x_1$.

Note that for all examples in this paper, we shall use the probability function

$$p_i(x_i, c_i) = \frac{1 - e^{-x_i/c_i}}{1 + e^{-x_i/c_i}}.$$

This function is rather convenient for it also incorporates the ratio of the attacker's investment to the defender's investment; as this ratio tends to infinity, the probability of success tends to 1, while the opposite is true when the ratio tends to 0.

2.1.2 Rewards and punishment

If target i is successfully destroyed, then the attacker receives a reward $a_i \geq 0$, and the defender suffers damage $d_i \geq 0$. If target i is attacked, but not destroyed, the attacker suffers punishment $f_i \geq 0$, and the defender suffers no damage as a result of target i being attacked. If target i is not attacked, neither player suffers any damage from target i . We refer to a_i, d_i , and f_i as the *game parameters*.

While d_i and f_i are known to both players from the start of the game, a_i is known only to the attacker. Although the defender does not know the value of a_i , he does know the density $g(a_i)$ from which the a_i are drawn. As in Bier et al. (2006), we assume that g is twice continuously differentiable, and that the a_i are independent. The continuity assumption allows for numerous convenient mathematical operations, while the independence assumption is consistent with that of the other game parameters. Independence also allows for easier numerical simulation, which we explore later in the paper.

Figure 2.1 shows a tree describing the rules and sequence of events in the model.

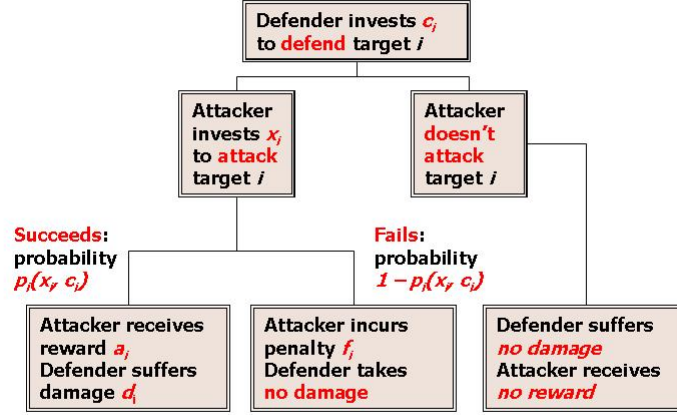


Figure 2.1: A tree describing the rules of the model.

For all examples in this paper, we shall assume that the a_i are distributed exponentially with parameter λ_i ; in other words, $g(a_i) = \lambda_i e^{-a_i \lambda_i}$. Note that this assumption is only for the examples; the theorems in this paper still hold regardless of the choice of g . Now, this density is rather convenient because it takes in only nonnegative values of a_i ; it also incorporates the often real-world phenomenon of diminishing returns for utility. In other words, the chances of receiving a larger reward decreases at an increasing rate.

2.2 Attacker's expected benefit

The attacker chooses the level at which to attack each target based on the security imposed by the defender, with the goal of maximizing his expected reward. First, let y_i represent whether or not target i is attacked. Then

$$y_i = \begin{cases} 0 & \text{if } x_i = 0, \\ 1 & \text{otherwise.} \end{cases}$$

Note that it always follows that $y_1 + y_2 \geq 1$, because $x_1 + x_2 = x_M$, so at least one target must be attacked.

Now, denote $T_i(x_i|c_i)$ to be the expected benefit to the attacker if the attacker allocates x_i towards target i , given the defender allocates c_i towards protecting target i . Thus, we define

$$\begin{aligned} T_i(x_i|c_i) &= y_i[a_i p_i(x_i, c_i) - f_i(1 - p_i(x_i, c_i))] \\ &= y_i[(a_i + f_i)p_i(x_i, c_i) - f_i]. \end{aligned}$$

Note that the first term in the first line represents the expected reward while the second represents the expected punishment. Also note that because the attacker is obligated to spend his entire budget, x_M , the cost of his investment does not appear in his net benefit function.

The total benefit for the attacker is the sum of the T_i . We define the total benefit as

$$\begin{aligned} B(\vec{x}|\vec{c}) &= y_1 T_1(x_1|c_1) + y_2 T_2(x_2|c_2) \\ &= y_1 T_1(x_1|c_1) + y_2 T_2(x_M - x_1|c_M - c_1). \end{aligned}$$

Given that the defender allocates $\vec{c} = (c_1, c_2) = (c_1, c_M - c_1)$ toward protecting targets 1 and 2, there may exist values of $\vec{x} = (x_1, x_2) = (x_1, x_M - x_1)$ that maximize the expected benefit for the attacker. We can consider these as functions of \vec{c} , and we refer to them as $\vec{x}_{opt}(\vec{c}) = (x_{opt1}(\vec{c}), x_{opt2}(\vec{c})) = (x_{opt1}(\vec{c}), x_M - x_{opt1}(\vec{c}))$. Hence,

$$B(\vec{x}_{opt}(\vec{c})|\vec{c}) \geq B(\vec{x}|\vec{c}), \quad \forall \vec{x} \neq \vec{x}_{opt}(\vec{c}).$$

From here on, we shall write \vec{x}_{opt} instead of $\vec{x}_{opt}(\vec{c})$ and x_{opt1} instead of $x_{opt1}(\vec{c})$ for brevity. As we show later in the paper, there are also scenarios where \vec{x}_{opt} does not exist.

2.3 Defender's expected damage

The defender chooses a security allocation \vec{c} to minimize the expected damages incurred in successful attacks. When the defender chooses a security allocation \vec{c} , the attacker responds by investing $\vec{x}_{opt}(\vec{c})$ (if it exists) into attacking the targets. Note that \vec{x}_{opt} depends on the values of a_i, f_i in addition to \vec{c} . As a result, the expected damage to the defender is

$$\begin{aligned} D(\vec{c}) &= d_1 p_1(x_{opt1}(\vec{c}), c_1) + d_2 p_2(x_{opt2}(\vec{c}), c_2) \\ &= d_1 p_1(x_{opt1}(\vec{c}), c_1) + d_2 p_2(x_M - x_{opt1}(\vec{c}), c_M - c_1). \end{aligned}$$

Since the a_i are unknown, however, the defender wishes to minimize his *expected damage with respect to* a_1 and a_2 :

$$\begin{aligned} E[D(\vec{c})] &= \int_0^\infty \int_0^\infty D(\vec{c}) g_1(a_1) g_2(a_2) da_1 da_2 \\ &= \int_0^\infty \int_0^\infty [d_1 p_1(x_{opt1}, c_1) + d_2 p_2(x_M - x_{opt1}, c_M - c_1)] g_1(a_1) g_2(a_2) da_1 da_2. \end{aligned} \tag{2.1}$$

Our objective is to determine the defender allocation \vec{c} that minimizes the defender's expected damage as indicated by (2.1).

Chapter 3

Attacker's Optimal Strategy

In order to determine the security investment the defender should make to minimize the expected damage, we first look at how the attacker would invest in an attack, as a function of the security in place, to maximize his expected reward. This information can then help the defender pick the appropriate allocation to force the attacker into a scenario which results in the least damage to the defender. We characterize the attacker's optimal strategy by solving for \vec{x}_{opt} in various scenarios.

First, we show that if a target having a nonzero value to the attacker is undefended, the attacker should always attack it (Theorem 3.1). Next, when the parameters are symmetric, when the attacker's budget is relatively large compared to the defender's, the attacker should attack both targets to maximize his expected benefits (Theorem 3.4). When the optimal solution is to attack both targets, the attacker should invest such that changing the investment would result in the increase in the expected reward at one target to equal the decrease at the other target (Theorem 3.5).

When the attacker's budget is relatively small, then the attacker should attack only the less heavily defended target when the parameters are symmetric (Theorem 3.7). In an asymmetric parameter case, the attacker also attacks only one target; the difference is that it is unknown analytically at which defender allocation he would switch targets (Theorem 3.10).

We also consider how the attacker's optimal allocation changed with respect to the game parameters. Under certain conditions, if the attack or punishment parameter at a target increases, then the attacker should increase the investment towards that target (Theorems 3.11 and 3.12), and if symmetric game parameters change simultaneously, the attacker should not change his strategy (Theorem 3.13).

3.1 undefended targets

We shall first look at the attacker's behavior when a target is undefended. We show that if the attacker values that target with a nonzero amount, then he should attack it.

Lemma 3.1. *If $a_1 > 0$ and $c_1 = 0$, then $\vec{x} = (0, x_M)$ is not optimal.*

Proof. If $x_1 = 0$, then $x_2 = x_M$, and $\vec{x} = (0, x_M)$. Likewise, since $c_1 = 0$, then $\vec{c} = (0, c_M)$. Therefore, the attacker's expected benefit from attacking only target 2 is

$$\begin{aligned} B((0, x_M)|(0, c_M)) &= T_2(x_M|c_M), \text{ because no benefit is received from target 1,} \\ &= (a_2 + f_2)p_2(x_M, c_M) - f_2, \text{ by definition.} \end{aligned}$$

By the continuity of p_2 , there exists $\epsilon > 0$ such that

$$(a_2 + f_2)(p_2(x_M, c_M) - p_2(x_M - \epsilon, c_M)) < a_1.$$

Rearranging yields

$$(a_2 + f_2)p_2(x_M, c_M) < a_1 + (a_2 + f_2)p_2(x_M - \epsilon, c_M). \quad (3.1)$$

Now we show that for such $\epsilon > 0$, attacking only target 2 does not maximize the expected reward. That is,

$$\begin{aligned} B((0, x_M)|(0, c_M)) &= (a_2 + f_2)p_2(x_M, c_M) \\ &< a_1 + (a_2 + f_2)p_2(x_M - \epsilon, c_M) - f_2, \text{ by (3.1).} \end{aligned}$$

Since $p_1(\epsilon, 0) = 1$, this is equal to

$$\begin{aligned} a_1 + (a_2 + f_2)p_2(x_M - \epsilon, c_M) - f_2 &= (a_1 + f_1)p_1(\epsilon, 0) - f_1 + (a_2 + f_2)p_2(x_M - \epsilon, c_M) - f_2 \\ &= T_1(\epsilon|0) + T_2(x_M - \epsilon|c_M) \\ &= B((\epsilon, x_M - \epsilon)|(0, c_M)), \text{ by definition.} \end{aligned}$$

Thus, the expected reward is higher for the attack allocation

$\vec{x} = (\epsilon, x_M - \epsilon)$, and $\vec{x} = (0, x_M)$ is not optimal. \square

Intuitively, if one target is undefended, then the attacker is guaranteed to secure the reward at that target if he attacks, no matter how small the attack allocation is. This implies that attacking that target results in a larger payoff than leaving it unattacked.

Although the attacker can increase his expected benefit by investing ϵ into attacking the undefended target, no optimal value of ϵ exists, as the next theorem shows.

Theorem 3.2. *If $\vec{c} = (0, c_M)$, and $a_1 > 0$, then there is no optimal value for \vec{x} .*

Proof. By Lemma 3.1, $\vec{x} = (0, x_M)$ is not optimal. Consider $\vec{x} = (\epsilon, x_M - \epsilon)$ for some $\epsilon > 0$. We show that there always exists some $0 < \epsilon^* < \epsilon$ such that investing $\vec{x}^* = (\epsilon^*, x_M - \epsilon^*)$ will result in a higher expected benefit than investing \vec{x} . The expected benefit from investing \vec{x} is

$$\begin{aligned} B((\epsilon, x_M - \epsilon)|(0, c_M)) &= T_1(\epsilon|0) + T_2(x_M - \epsilon|c_M) \\ &= (a_1 + f_1)p_1(\epsilon, 0) - f_1 + (a_2 + f_2)p_2(x_M - \epsilon, c_M) - f_2 \\ &= a_1 + (a_2 + f_2)p_2(x_M - \epsilon, c_M) - f_2. \end{aligned}$$

For $\epsilon^* < \epsilon$, however, since $p_2(x_M - \epsilon, c_M) < p_2(x_M - \epsilon^*, c_M)$,

$$\begin{aligned} a_1 + (a_2 + f_2)p_2(x_M - \epsilon, c_M) - f_2 &< a_1 + (a_2 + f_2)p_2(x_M - \epsilon^*, c_M) - f_2 \\ &= T_1(\epsilon^*|0) + T_2(x_M - \epsilon^*|c_M) \\ &= B((\epsilon^*, x_M - \epsilon^*)|(0, c_M)), \end{aligned}$$

contradicting the optimality of $x_1 = \epsilon$. Thus, for any value of $\epsilon > 0$ invested into attacking target 1, a better value can be found. \square

An intuitive reason for this is that any positive attack allocation towards an undefended target guarantees that the attacker will successfully destroy the target. However, there is no minimal value for such a positive allocation; hence no optimal value for \vec{x} exists.

3.2 Defended targets

Next, we show that when both targets are defended, the optimal attack allocation depends on the magnitude of the attacker's budget, x_M , relative to that of the defender's budget, c_M .

One question we will answer is whether the attacker will attack one or both targets. If the attacker were to attack both targets, the expected benefit from attacking both targets should be greater than that from solely attacking one target. This idea will be used in determining whether the attacker's optimal investment is to attack one target, or both. The following sections explain how the relative magnitude of the attacker's budget will affect whether he attacks one target or both.

3.2.1 Large attacker budget—symmetric parameters

We begin by looking at symmetric cases (a_i 's and f_i 's are equal) where the attacker has a relatively large budget. We must first, however, introduce a new expression:

$$P(x_M, \vec{x}, \vec{c}) = p_1(x_1, c_1) + p_2(x_M - x_1, c_M - c_1) - p_1(x_M, c_1)$$

where \vec{x} and \vec{c} are feasible. P represents the increase in the probability of a successful attack when the attacker chooses to attack both targets at an investment of $(x_1, x_M - x_1)$ rather than just target 1. For brevity, we will refer to P as merely the *increase in the probability of a successful attack*. We will use the relative magnitude of P to derive a sufficient condition for when the attacker will attack both targets.

We first prove a lemma that shows that given a value $e < 1$, if the attacker budget is large enough, then the attacker can choose a feasible allocation such that the increase in the probability of success is greater than e .

Lemma 3.3. *Fix c_M and let $e < 1$. There exists a threshold \underline{x}_M such that if $x_M > \underline{x}_M$, then for every security allocation \vec{c} , there exist feasible values for \vec{x} such that $P(x_M, \vec{x}, \vec{c}) = p_1(x_1, c_1) + p_2(x_M - x_1, c_M - c_1) - p_1(x_M, c_1) > e$.*

Proof. First, observe that as $x_M \rightarrow \infty$, both x_1 and x_2 are unconstrained. Since $\lim_{x_i \rightarrow \infty} p_i(x_i, c_i) = 1$, and we can pick infinitely large x_1 and x_2 , it follows that there exists x_1 and x_2 such that

$$\begin{aligned} \lim_{x_M \rightarrow \infty} P &= \lim_{x_M \rightarrow \infty} [p_1(x_1, c_1) + p_2(x_M - x_1, c_M - c_1)] - 1 \\ &= 1 + 1 - 1 \\ &> e. \end{aligned}$$

(Note that P strictly increases with respect to x_M as long as, without loss of generality, x_1 increases and x_2 stays fixed; this guarantees that once $P > e$, P does not oscillate around it).

Therefore, there exists $\underline{x}_M < \infty$ such that for all feasible \vec{c} , there exist feasible $\vec{x} = (\underline{x}_1, \underline{x}_M - \underline{x}_1)$ such that

$$P(\underline{x}_M, \vec{x}, \vec{c}) > e \tag{3.2}$$

Now, for any $x_M > \underline{x}_M$, let $k = x_M - \underline{x}_M$. For any $\underline{x}_1 < \underline{x}_M$ such that (3.2) holds, let $x_1 = \underline{x}_1 + k < x_M$. We know show that for any attacker budget x_M greater than \underline{x}_M , there always exists feasible attacker allocations

such that the increase in the probability of a successful attack, P , is greater than e .

The concavity of p_1 with respect to x_1 shows that

$$p_1(x_1, c_1) - p_1(\underline{x}_1, c_1) > p_1(x_M, c_1) - p_1(\underline{x}_M, c_1).$$

Rearranging terms and adding $p_2(\underline{x}_M - \underline{x}_1, c_M - c_1)$, this is equivalent to

$$\begin{aligned} & p_1(x_1, c_1) + p_2(\underline{x}_M - \underline{x}_1, c_M - c_1) - p_1(x_M, c_1) \\ & > p_1(\underline{x}_1, c_1) + p_2(\underline{x}_M - \underline{x}_1, c_M - c_1) - p_1(x_M, c_1). \end{aligned}$$

Note that the right hand side is just $P(\underline{x}_M, \vec{x}, \vec{c})$ which is greater than e by (3.2); thus,

$$p_1(x_1, c_1) + p_2(\underline{x}_M - \underline{x}_1, c_M - c_1) - p_1(x_M, c_1) > e. \quad (3.3)$$

Since

$$\begin{aligned} p_2(\underline{x}_M - \underline{x}_1, c_M - c_1) &= p_2(x_M - k - \underline{x}_1, c_M - c_1) \\ &= p_2(x_M - k - (x_1 - k), c_M - c_1) \\ &= p_2(x_M - x_1, c_M - c_1), \end{aligned}$$

then by (3.3),

$$\begin{aligned} p_1(x_1, c_1) + p_2(x_M - x_1, c_M - c_1) - p_1(x_M, c_1) &> e \\ &\Leftrightarrow P(x_M, \vec{x}, \vec{c}) > e. \end{aligned}$$

Therefore, for all $x_M > \underline{x}_M$, there exists feasible \vec{x} where $P(x_M, \vec{x}, \vec{c}) > e$ for all feasible \vec{c} . \square

Intuitively, if the attacker budget is large enough relative to the defender budget, then regardless of what security allocation the defender chooses, the attacker can invest such that the probability of a successful attack at each target is almost as high as that had the attacker attacked only one target. This allows for the increase in the probability of success to be greater than any $e < 1$.

We now show that if the game parameters are symmetric, and the attacker budget is sufficiently large (a precise definition provided below), then the attacker will choose to attack both targets.

Theorem 3.4. *Let $a_1 = a_2 = a > 0$ and $f_1 = f_2 = f > 0$. If x_M is sufficiently large such that for any feasible value of \vec{c} , there exists feasible values for \vec{x} such that $P(x_M, \vec{x}, \vec{c}) > \frac{f}{a+f}$, then for all feasible values of \vec{c} , $0 < x_{opti} < x_M$. We show that if the attacker's budget is large enough so that the increase in the probability of a successful attack is larger than the aforementioned ratio, then the optimal solution is to attack both targets.*

Proof. First, note that Lemma 3.3 indicates the existence of the necessary conditions. Now, assume to the contrary that there exists a security allocation \vec{c} such that $\vec{x}_{opt} = (x_M, 0)$. Then it follows that the expected benefit from attacking only target 1 is greater than any expected benefit from attacking both targets. Therefore, for all \vec{x} such that $0 < x_1 < x_M$, we have

$$B(\vec{x}_{opt}|\vec{c}) \geq B(\vec{x}|\vec{c}),$$

which implies

$$T_1(x_M|c_1) \geq T_1(x_1|c_1) + T_2(x_M - x_1|c_M - c_1),$$

and substituting yields

$$(a + f)p_1(x_M, c_1) - f \geq (a + f)[p_1(x_1, c_1) + p_2(x_M - x_1, c_M - c_1)] - 2f.$$

Rearranging terms, we have

$$\frac{f}{a + f} \geq p_1(x_1, c_1) + p_2(x_M - x_1, c_M - c_1) - p_1(x_M, c_1),$$

which is a contradiction, because we've assumed the existence of \vec{x} such that the above inequality fails to hold. This implies $x_{opt1} < x_M$ for all feasible \vec{c} . By the symmetry of the problem, $x_{opt2} < x_M$, implying $x_M > x_{opt1} > 0$ for all feasible \vec{c} . \square

In other words, if x_M is large enough so the attacker can choose an investment such that the improvement in the probability of success outweighs the cost of failure as represented by $\frac{f}{a+f}$, then the attacker should attack both targets no matter what the defender does. If the reward is small relative to the punishment, or if the attacker's budget is small relative to the defender's, the attacker has less incentive to attack both targets, since either the net benefit from or the probability of successfully attacking both targets is low; the attacker would be better off allocating everything towards attacking one target. In Theorem 3.4, a decrease in the reward results in the

increase of the cost of failure, $\frac{f}{a+f}$, implying that the requirements for a sufficiently large attacker budget become more stringent; a larger attacker budget would be necessary to ensure the attacker can optimally attack both targets.

Figure 3.1 shows a plot of x_{opt1} versus c_1 , when x_M is sufficiently small; Matlab calculated x_{opt1} numerically for every value of c_1 to generate the plot. The figure shows that $0 < x_{opt1} < x_M$ for all feasible values of c_1 .

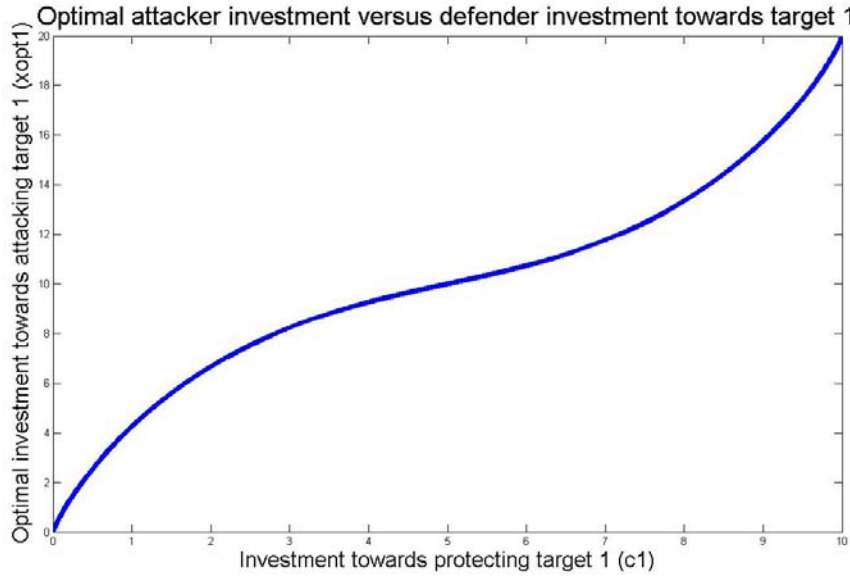


Figure 3.1: Optimal allocation towards target 1 when x_M is sufficiently large. $x_M = 20, c_M = 10, a = 6, f = 5$.

Now that we have characterized sufficient criteria for the attacker to attack both targets, we would like to determine the optimal attacker amount to invest into each target. While we have not determined sufficient criteria for when the attacker will attack both targets when the game parameters are asymmetric, the following result will hold as long as it is known the attacker will attack both targets; whether or not the game parameters are symmetric is irrelevant.

We show that the optimal investment occurs where the increase in the expected reward from one target equals the decrease in the expected reward from the other target, or in other words, where there is no increase in

the expected reward.

Theorem 3.5. *If there exists a feasible \vec{x}^* such that $\frac{dT_1}{dx_1}|_{x_1^*} = -\frac{dT_2}{dx_1}|_{x_1^*}$, then \vec{x}^* is optimal.*

Proof. First recall that since

$$\begin{aligned} B(\vec{x}|\vec{c}) &= T_1(x_1|\vec{c}) + T_2(x_M - x_1|\vec{c}) \\ &= (a_1 + f_1)p_1 - f_1 + (a_2 + f_2)p_2 - f_2, \end{aligned}$$

B is concave down with respect to x_1 . Also recall that $\frac{dT_1}{dx_1} > 0$ and $\frac{dT_2}{dx_1} < 0$.

Now let there exist \vec{x}^* such that $\frac{dT_1}{dx_1}|_{x_1^*} = -\frac{dT_2}{dx_1}|_{x_1^*}$. If \vec{x}^* were not optimal, then the attacker could either increase or decrease x_1 to increase the expected reward, B . If the attacker can increase x_1 to obtain a higher expected reward, then

$$\begin{aligned} \frac{dB}{dx_1}|_{\vec{x}^*} &> 0 \\ \Rightarrow \frac{dT_1}{dx_1}|_{x_1^*} &> -\frac{dT_2}{dx_1}|_{x_1^*}, \end{aligned}$$

yielding a contradiction. By symmetry, the attacker would be unable to decrease x_1 to increase B . Since B is concave down with respect to x_1 , then

$$\begin{aligned} \frac{dB}{dx_1}|_{\vec{x}^*} &= 0 \\ \Leftrightarrow \frac{dT_1}{dx_1}|_{x_1^*} &= -\frac{dT_2}{dx_1}|_{x_1^*} \end{aligned}$$

implies that $x_1^* = x_{opt1}$, and hence, \vec{x}^* is optimal. \square

Intuitively, when an attacker can only decrease his expected reward by increasing or decreasing his investment towards target 1, he is at the optimal investment. Only if the increase in the reward at one target outweighs the decrease at the other can he improve his reward. Recall also that this theorem holds when the parameters are asymmetric; the net increase and decrease being equal depends not on the parameters being symmetric, but only on the overall expected reward from each target.

3.2.2 Relatively small attacker budget—symmetric parameters

We now show that when the attacker budget is small relative to the defender budget, the attacker will attack only one target. We begin by proving a lemma that states for a given h where $0 < h < 1$, if the attacker budget constraint is sufficiently small, then the probability of successfully attacking the more heavily defended target is less than h regardless of what the attacker does.

Lemma 3.6. *Fix c_M, \vec{c} , and $0 < h < 1$. There exists a threshold $\bar{x}_M > 0$ such that if $x_M < \bar{x}_M$, then $p_2(x_2, c_2) < h$ for all feasible \vec{x} . In other words, we show that for any defender allocation, there exists a maximum attacker budget such that this budget will force the probability of success at a target to always be less than some constant h .*

Proof. First,

$$\begin{aligned} \lim_{x_M \rightarrow 0} p_2(x_M, c_2) &= 0 \\ &< h. \end{aligned}$$

Hence, there exists a threshold $\bar{x}_M > 0$ where $p_2(\bar{x}_M, c_2) < h$. Therefore, for any $x_2 \leq x_M \leq \bar{x}_M$,

$$\begin{aligned} p_2(x_2, c_2) &\leq p_2(x_M, c_2) \\ &\leq p_2(\bar{x}_M, c_2) \\ &< h, \end{aligned}$$

completing the proof. \square

It makes sense that if the attacker budget is sufficiently small, then the probability of successfully attacking a target is always smaller than h regardless of what the attacker allocates, since the attacker is limited by his small budget.

We proceed to examine the scenario where the game parameters are symmetric and the defender defends one target more than the other. If the attacker budget is sufficiently small (the precise definition will be stated below), then the attacker will attack only the less heavily defended target.

Theorem 3.7. *Let $a_1 = a_2 = a > 0$ and $f_1 = f_2 = f > 0$. Assume the defender invests \vec{c} where $c_2 > \frac{c_M}{2}$. If x_M is sufficiently small such that $p_2(x_2, c_2) < \frac{f}{a+f}$ for all feasible values of \vec{x} , then $\vec{x}_{opt} = (x_M, 0)$.*

Proof. First, note that Lemma 3.6 ensures the existence of the necessary conditions. For any feasible \vec{x} such that $0 < x_1 < x_M$, and feasible \vec{c} for $c_1 < \frac{c_M}{2}$, we know that $p(x_1, c_1) < p_1(x_M, c_1)$. Also, $p_2(x_M - x_1, c_M - c_1) < \frac{f}{a+f}$ by assumption. Hence,

$$\frac{f}{a+f} - p_2(x_M - x_1, c_M - c_1) > p_1(x_1, c_1) - p_1(x_M, c_1),$$

since left and right sides are positive and negative respectively. Rearranging yields

$$(a+f)p_1(x_M, c_1) - f > (a+f)p_1(x_1, c_1) - f + (a+f)p_2(x_2, c_2) - f,$$

and thus, by definition,

$$\begin{aligned} T_1(x_M|c_1) &> T_1(x_1|c_1) + T_2(x_M - x_1|c_M - c_1) \\ \Leftrightarrow B((x_M, 0)|\vec{c}) &> B(\vec{x}|\vec{c}). \end{aligned}$$

Therefore, the expected benefit to the attacker from attacking target 1 is greater than any expected benefit from attacking both targets. In other words, $\vec{x} = (x_M, 0)$ is preferable to any \vec{x} where $0 < x_1 < x_M$. Now, it also follows that

$$\begin{aligned} B((x_M, 0)|\vec{c}) &= T_1(x_M|c_1) \\ &= (a+f)p_1(x_M, c_1) - f \\ &> (a+f)p_2(x_M, c_M - c_1) - f, \text{ since } c_1 < \frac{c_M}{2} \\ &= T_2(x_M|c_M - c_1) \\ &= B((0, x_M)|c_1), \end{aligned}$$

which implies that the expected benefit from attacking only target 1 is greater than that from attacking only target 2 ($\vec{x} = (x_M, 0)$ is also preferable to $\vec{x} = (0, x_M)$). Hence, $\vec{x}_{opt} = (x_M, 0)$. \square

By symmetry, it also follows that the attacker should attack only target 2 if target 1 is more heavily defended, and if $p_1(x_1, c_1) < \frac{f}{a+f}$ for all feasible \vec{x} . Intuition states that if the attacker budget is small enough such that the probability of successfully attacking the more heavily defended target is always less than $\frac{f}{a+f}$, then the attacker should attack only one target. The target to attack should be the less defended one because the probability of successfully destroying that one is higher. As the reward increases, $\frac{f}{a+f}$

decreases, implying a lower attacker budget is needed to ensure that the optimal solution is to attack only one target.

Figure 3.2 contains a plot of x_{opt1} vs c_1 when x_M is sufficiently small. It shows that the attacker should attack only the less heavily defended target.




Figure 3.2: Optimal allocation towards target 1 when x_M is sufficiently small. $x_M = 5, c_M = 10, a = 6, f = 5$.

Figure 3.2 raises the question of what the attacker should do when $\vec{c} = (\frac{c_M}{2}, \frac{c_M}{2})$, that is, when the defender has protected both targets evenly. We proceed by showing that if the attacker budget is sufficiently small and the defender evenly defends both targets, it would be suboptimal for the attacker to attack both targets, and that his expected benefit would be maximized by attacking only one target.

Theorem 3.8. *Let $a_1 = a_2 = a > 0$ and $f_1 = f_2 = f > 0$. If $\vec{c} = (\frac{c_M}{2}, \frac{c_M}{2})$, and x_M is sufficiently small such that $p_2(x_2, \frac{c_M}{2}) < \frac{f}{a+f}$ for all feasible \vec{x} , then both $\vec{x} = (x_M, 0)$ and $\vec{x} = (0, x_M)$ are optimal.*

Proof. We first show that attacking only target 1 is preferable to attacking both targets. For all feasible \vec{x} such that $0 < x_1 < x_M$, we know that

$p_1\left(x_1, \frac{c_M}{2}\right) < p_1\left(x_M, \frac{c_M}{2}\right)$. Also, since $\frac{f}{a+f} > p_2\left(x_M - x_1, \frac{c_M}{2}\right)$:

$$\frac{f}{a+f} - p_2\left(x_M - x_1, \frac{c_M}{2}\right) > p_1\left(x_1, \frac{c_M}{2}\right) - p_1\left(x_M, \frac{c_M}{2}\right),$$

since the left and right hand sides are positive and negative respectively. Rearranging and subtracting f from both sides, we see that this is equivalent to

$$\begin{aligned} (a+f)p_1\left(x_M, \frac{c_M}{2}\right) - f &> (a+f)p_1\left(x_1, \frac{c_M}{2}\right) - f \\ &\quad + (a+f)p_2\left(x_M - x_1, \frac{c_M}{2}\right) - f \\ \Leftrightarrow T_1\left(x_M \middle| \frac{c_M}{2}\right) &> T_1\left(x_1 \middle| \frac{c_M}{2}\right) + T_2\left(x_2 \middle| \frac{c_M}{2}\right) \\ \Leftrightarrow B((x_M, 0) | (\vec{c})) &> B(\vec{x} | \vec{c}). \end{aligned}$$

This implies that the attacker prefers $\vec{x} = (x_M, 0)$ to any \vec{x} where $0 < x_1 < x_M$ because the former results in the greater expected benefit. Now we show that attacking only target 2 results in the same expected benefit as attacking only target 1. First,

$$\begin{aligned} B((x_M, 0) | \vec{c}) &= T_1\left(x_M \middle| \frac{c_M}{2}\right) \\ &= (a+f)p_1\left(x_M, \frac{c_M}{2}\right). \end{aligned}$$

Since $p_1(x, c) = p_2(x, c)$ for all x, c , then

$$\begin{aligned} (a+f)p_1\left(x_M, \frac{c_M}{2}\right) &= (a+f)p_2\left(x_M, \frac{c_M}{2}\right) \\ &= T_2\left(x_M \middle| \frac{c_M}{2}\right) \\ &= B((0, x_M) | \vec{c}). \end{aligned}$$

Hence, the allocation $\vec{x} = (0, x_M)$ is also optimal. \square

3.2.3 Relatively small attacker budget—asymmetric parameters

We now proceed with the case where the game parameters are asymmetric and the attacker's budget is relatively small.

We show that again, if the attacker's budget is relatively small, then the attacker will only attack one target. We begin by showing sufficient criteria for the existence of a defender allocation \hat{c} such that when $c_1 < \hat{c}$,

the attacker prefers attacking only target 1 to attacking only target 2. The opposite is true when $c_1 > \hat{c}$. We then proceed by showing sufficient (but not necessary) criteria where attacking target 1 is optimal when $c_1 < \hat{c}$ and attacking target 2 is optimal when $c_1 > \hat{c}$.

Theorem 3.9. Fix c_M , and let x_M be sufficiently small such that $p_i(x_M, c_M) < \min\left(\frac{a_1+f_2}{a_2+f_2}, \frac{a_2+f_1}{a_1+f_1}\right)$. Then there exists a feasible value \hat{c} such that the following hold:

- $B((x_M, 0)|(\hat{c}, c_M - \hat{c})) = B((0, x_M)|(\hat{c}, c_M - \hat{c}))$
- $B((x_M, 0)|(c_1, c_M - c_1)) > B((0, x_M)|(c_1, c_M - c_1))$ for all $c_1 < \hat{c}$
- $B((x_M, 0)|(c_1, c_M - c_1)) < B((0, x_M)|(c_1, c_M - c_1))$ for all $c_1 > \hat{c}$.

Proof. Let x_M be sufficiently small such that the aforementioned conditions hold. Let

$$\begin{aligned} R(c_1) &= (a_1 + f_1)p_1(x_M, c_1) - f_1 - (a_2 + f_2)p_2(x_M, c_M - c_1) + f_2 \\ &= B((x_M, 0)|(\hat{c}, c_M - \hat{c})) - B((0, x_M)|(\hat{c}, c_M - \hat{c})). \end{aligned}$$

Note that R denotes the difference in expected reward from attacking only target 1 versus attacking only target 2. Observe that by assumption,

$$\begin{aligned} p_2(x_M, c_M) &< \frac{a_1 + f_2}{a_2 + f_2}, \text{ and rearranging yields} \\ 0 &< (a_1 + f_1)p_1(x_M, 0) - f_1 - (a_2 + f_2)p_2(x_M, c_M) + f_2, \text{ so} \\ 0 &< R(0). \end{aligned}$$

Similarly,

$$\begin{aligned} p_1(x_M, c_M) &< \frac{a_2 + f_1}{a_1 + f_1} \text{ implies that} \\ (a_1 + f_1)p_1(x_M, c_M) - f_1 - (a_2 + f_2)p_2(x_M, 0) - f_2 &< 0, \text{ and thus} \\ R(c_M) &< 0. \end{aligned}$$

In other words, when target 1 is undefended, the difference in expected reward is positive; the opposite is true when target 2 is undefended. Now, by the continuity of R , there exists $0 < \hat{c} < c_M$ such that

$$\begin{aligned} R(\hat{c}) &= 0, \text{ which implies} \\ B((x_M, 0)|(\hat{c}, c_M - \hat{c})) &= B((0, x_M)|(\hat{c}, c_M - \hat{c})). \end{aligned}$$

Hence, there exists \hat{c} where the expected reward from attacking only one target is the same, regardless of the target.

Also, since T is decreasing in c_1 , that implies that for all $c_1 < \hat{c}$,

$$\begin{aligned} R(c_1) &> 0, \text{ so} \\ B((x_M, 0)|(c_1, c_M - c_1)) &> B((0, x_M)|(c_1, c_M - c_1)), \end{aligned}$$

and for all $c_1 > \hat{c}$,

$$\begin{aligned} R(c_1) &< 0, \text{ thus} \\ B((x_M, 0)|(c_1, c_M - c_1)) &< B((0, x_M)|(c_1, c_M - c_1)). \end{aligned}$$

□

Note that Theorem 3.9 provides sufficient criteria for when attacking only target 1 is preferable to attacking only target 2. However, that does not imply that attacking only target 1 is optimal. (The optimal solution may involve attacking both targets). The sufficient criteria for attacking only one target are provided in the following theorem.

Theorem 3.10. Fix c_M . Let x_M be sufficiently small such that the conditions of Theorem 3.9 hold. Let \hat{c} be defined as in Theorem 3.9. If the following hold:

- $(a_1 + f_1)[p_1(x_1, c_1) - p_1(x_M, c_1)] + (a_2 + f_2)p_2(x_M - x_1, c_M - c_1) < f_2$
for all feasible x_1 when $c_1 < \hat{c}$,
- $(a_2 + f_2)[p_2(x_M - x_1, c_M - c_1) - p_2(x_M, c_M - c_1)] + (a_1 + f_1)p_1(x_1, c_1) < f_1$
for all feasible x_1 when $c_1 > \hat{c}$,

then $\vec{x}_{opt} = (x_M, 0)$ if $c_1 < \hat{c}$, and $\vec{x}_{opt} = (0, x_M)$ if $c_1 > \hat{c}$.

Proof. Let the defender invest $c_1 < \hat{c}$. Then it follows that for all possible values of x_1 , by the first condition in the statement,

$$(a_1 + f_1)[p_1(x_1, c_1) - p_1(x_M, c_1)] + (a_2 + f_2)p_2(x_M - x_1, c_M - c_1) < f_2.$$

After rearranging terms, we have

$$(a_1 + f_1)p_1(x_1, c_1) - f_1 + (a_2 + f_2)p_2(x_M - x_1, c_M - c_1) - f_2 < (a_1 + f_1)p_1(x_M, c_1) - f_1,$$

and substituting definitions,

$$\begin{aligned} T_1(x_1|c_1) + T_2(x_M - x_1|c_M - c_1) &< T_1(x_M|c_1) \\ \Leftrightarrow B(\vec{x}|\vec{c}) &< B((x_M, 0)|\vec{c}). \end{aligned}$$

Hence, for $c_1 < \hat{c}$, the expected benefit from attacking only target 1 is greater than any expected benefit from attacking both targets. Now, by Theorem 3.9, we know that $B((0, x_M)|\vec{c}) < B((x_M, 0)|\vec{c})$, implying $\vec{x}_{opt} = (x_M, 0)$ when the defender invests $c_1 < \hat{c}$.

By symmetry, $\vec{x}_{opt} = (0, x_M)$ when the defender invests $c_1 > \hat{c}$. □

Note that the conditions of Theorem 3.9 guarantee that $c_1 < \hat{c}$ implies attacking only target 1 is preferable to attacking only target 2, and the new conditions introduced in Theorem 3.10 imply that attacking only target 1 is preferable to attacking both targets, which proves the optimality of only attacking target 1. Therefore, if the attacker's budget is sufficiently small by both Theorems 3.9 and 3.10, then the attacker will attack only one target.

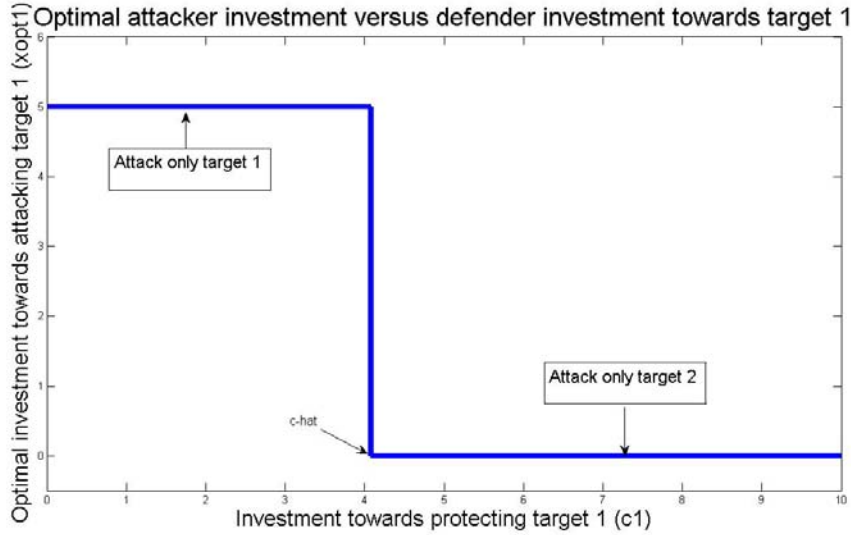


Figure 3.3: Optimal allocation towards target 1 when x_M is sufficiently small and parameters are asymmetric; \hat{c} is approximately 4.1.

Figure 3.3 shows a plot of the optimal attacker allocation versus the defender's investment towards target 1 when the attacker's budget is sufficiently small by the criteria of Theorem 3.10.

Also note that Theorem 3.7 does not fall out as a special case of Theorem 3.10, for Theorem 3.7 allows for a less stringent, yet sufficient, condition for the attacker to attack only one target.

3.3 Other characteristics of the attacker's strategy

We now examine how the optimal attacker allocation depends on the model parameters. Although in this model, the defender cannot change the model parameters to his benefit, it is still valuable from a policy standpoint to understand how the attacker's behavior depends on such parameters.

3.3.1 Optimal behavior with respect to parameters

We first show that if the reward for successfully attacking one target increases, the attacker's optimal allocation towards that target increases, given certain conditions.

Theorem 3.11. *Let $0 < x_{opt1} < x_M$, $a_i > 0$, and $c_i \neq 0$ for both i . As a_1 increases and a_2 stays fixed, if there exists a new optimal attacker allocation towards target 1, \tilde{x}_{opt1} , where $\frac{dB}{dx_1}$ is defined at \tilde{x}_{opt1} , then $\tilde{x}_{opt1} > x_{opt1}$.*

Proof. By Theorem (3.5), we know that

$$\begin{aligned} \frac{dB}{dx_1} \Big|_{\tilde{x}_{opt}} &= 0 \\ \Rightarrow \frac{dT_1}{dx_1} \Big|_{\tilde{x}_{opt}} &= -\frac{dT_2}{dx_1} \Big|_{\tilde{x}_{opt}} \\ \Rightarrow (a_1 + f_1) \frac{dp_1}{dx_1} \Big|_{x_{opt1}} &= -(a_2 + f_2) \frac{dp_2}{dx_1} \Big|_{x_M - x_{opt1}}. \end{aligned} \quad (3.4)$$

Let a_1 increase. If $\tilde{x}_{opt1} = x_{opt1}$, then the left hand side is greater than the right hand side, yielding a contradiction. If $0 < \tilde{x}_{opt1} < x_{opt1}$, then $\frac{dp_1}{dx_1} \Big|_{\tilde{x}_{opt1}} > \frac{dp_1}{dx_1} \Big|_{x_{opt1}}$, while $\frac{dp_2}{dx_1} \Big|_{x_M - x_{opt1}}$ becomes less negative, implying again that the left hand side is greater than the right hand side. Since (3.4) must be defined at the new value of \tilde{x}_{opt1} , then it follows that $x_{opt1} < \tilde{x}_{opt1}$ (since $\tilde{x}_{opt1} \neq 0$). \square

Similarly, if a_2 increases while a_1 stays fixed, then x_{opt1} would decrease and x_{opt2} would increase.

The results make intuitive sense; as the reward for one target increases, the attacker will wish to invest more to increase the probability of securing a higher reward. Also note if $c_i = 0$, then by Theorem 3.2, there is no optimal value of x_1 . If $x_{opt_i} = x_M$, then it is impossible for x_{opt_i} to increase

as a_i increases. Hence, it is necessary for the proof that $0 < c_1 < c_M$ and $0 < x_{opt1} < x_M$.

We now show that under the same conditions as Theorem 3.11, the attacker optimal investment towards a target also increases if the punishment for a failed attack at that target increases.

Theorem 3.12. *Let $0 < x_{opt1} < x_M$, $f_i > 0$, and $c_i \neq 0$ for both i . As f_1 increases and f_2 stays fixed, if there exists a new optimal attacker allocation towards target 1, \tilde{x}_{opt1} , where $\frac{dB}{dx_1}$ must be defined at \tilde{x}_{opt1} , then $\tilde{x}_{opt1} > x_{opt1}$.*

Proof. The logic for this proof is exactly the same as that of Theorem 3.11. \square

Here, as the punishment for a failed attack on a target increases, the attacker will invest more in that target to reduce the probability of failure. However, note that Theorem 3.12 requires that $\frac{dB}{dx_1}$ be defined at the new value of x_{opt1} , \tilde{x}_{opt1} , which precludes \tilde{x}_{opt1} from being 0. This does not necessarily imply that \tilde{x}_{opt1} is not 0; in fact, numerical results have shown that the attacker will not attack target 1 at all if the punishment becomes large enough that attacking is not worth the risk of failure. However, we have not yet been able to prove this, since $\frac{dB}{dx_1}$ is not defined at $x_1 = 0$.

We now show that the attacker's optimal allocation doesn't change when symmetric attack parameters change simultaneously, given certain conditions.

Theorem 3.13. *Let $a_1 = a_2 = a$ and $f_1 = f_2 = f$, and fix \vec{c} such that $0 < x_{opt1} < x_M$. Let the values of a_1 and a_2 change to $a_1 = a_2 = \hat{a}$. If the new optimal solution, \vec{x}_{opt} , is such that $\frac{dB}{dx_1}|_{\vec{x}_{opt}}$ is defined, then $\vec{x}_{opt} = \vec{x}_{opt}$.*

Proof. Let $a_1 = a_2 = a$ and $f_1 = f_2 = f$. Then (3.4) evaluated at x_{opt1} becomes

$$\begin{aligned} \frac{dB}{dx_1} \Big|_{\vec{x}_{opt}} &= (a + f) \left(\frac{dp_1}{dx_1} \Big|_{x_{opt1}} + \frac{dp_2}{dx_1} \Big|_{x_M - x_{opt1}} \right) = 0 \\ \Rightarrow \frac{dp_1}{dx_1} \Big|_{x_{opt1}} &= - \frac{dp_2}{dx_1} \Big|_{x_M - x_{opt1}}. \end{aligned} \quad (3.5)$$

For $a_1 = a_2 = \hat{a} \neq a$, (3.5) still holds. If $x_{opt1} < \tilde{x}_{opt1} < x_M$, the left hand side decreases, while the right hand side increases due to $\frac{dp_2}{dx_1}$ becoming less negative, invalidating the equality. If $x_{opt1} > \tilde{x}_{opt1} > 0$, then the reverse would occur. Since $\frac{dB}{dx_1}$ is assumed to be defined at \vec{x}_{opt} , \tilde{x}_{opt1} can neither equal 0 nor x_M . Therefore, $\vec{x}_{opt} = \vec{x}_{opt}$. \square

The attacker will not change his optimal strategy since the net expected payoff at both targets is still the same, even though it has changed in value. However, intuition also says that if the punishment rises significantly such that the net expected payoff is negative, then it would be optimal to attack only one target to minimize such losses. Numerical results do indicate that x_{opt1} can be 0 or x_M ; however, we have not been able to prove this since $\frac{dB}{dx_1}$ is not defined at those points.

Chapter 4

Defender's Optimal Strategy

Now that we have examined the attacker's optimal strategy in different scenarios, we proceed to determine the allocation that results in the least damage to the the defender. Recall that since the reward parameters are unknown to the defender, he wishes to minimize (2.1), his expected damage with respect to the rewards:

$$E[D(\vec{c})] = \int_0^\infty \int_0^\infty [d_1 p_1(x_{opt1}, c_1) + d_2 p_2(x_M - x_{opt1}, c_M - c_1)] g_1(a_1) g_2(a_2) da_1 da_2.$$

We refer to the defender's optimal allocation as the allocation that minimizes this expected value.

Analytically determining the defender allocation that minimizes (2.1), however, is rather difficult due to the uncertainty of the rewards. For simplicity, we first look, in the following section, at a symmetric parameter case when the reward parameters are known, and show that the defender minimizes his expected damage by defending both targets equally (Theorem 4.1).

Afterwards, we proceed by numerical simulation to approximate the allocation that would minimize his expected damage, and make conjectures regarding trends of the defender's allocation with respect to parameters when the rewards are unknown. One conjecture is that an increase in the expected reward or damage at a target tends to increase the optimal defense allocation towards that target (Conjectures 4.3 and 4.4). An increase in the punishment at a target results in the opposite effect (Conjecture 4.2). Finally, while we also attempt to examine trends regarding the volatility or the increase in the attacker's budget, it has been difficult to state relevant conjectures with a high degree of confidence.

Because it is often difficult to get exact values for the game parameters in real-world scenarios, it is important to be able to approximate such parameters as well as determine *how* a defense policy should change as those parameters differ within a certain range. These numerical conjectures provide a starting point for how one should plan a flexible defense strategy given the ability to determine a general range for relevant parameters.

4.1 Known reward parameters

We start with a symmetric parameter case when the attacker's budget is sufficiently small. We will assume, for simplification purposes, that the defender knows the rewards a_1, a_2 the attacker will get at each target if it is successfully attacked. As previously mentioned, in the later sections, we shall relax this assumption.

Theorem 4.1. *Let $d_1 = d_2 = d, a_1 = a_2 = a, f_1 = f_2 = f$, and assume all parameters are known to the defender. If x_M is sufficiently small such that $p_2(x_2, c_2) < \frac{f}{a+f}$ for all feasible values of \vec{c} such that $c_2 \geq \frac{c_M}{2}$ and all feasible values of \vec{x} , then $\vec{c} = (\frac{c_M}{2}, \frac{c_M}{2})$ is optimal.*

Proof. For any $c_1 < \frac{c_M}{2}$, we know that $\vec{x}_{opt} = (x_M, 0)$ by Theorem 3.7. Hence, for any $c_1 < \frac{c_M}{2}$, the expected damage to the defender is

$$\begin{aligned} D(\vec{c}) &= dp_1(x_{opt1}, c_1) + dp_2(x_M - x_{opt1}, c_M - c_1) \\ &= dp_1(x_M, c_1), \end{aligned}$$

since $x_{opt1} = x_M$. Now, because $c_1 < \frac{c_M}{2}$,

$$\begin{aligned} dp_1(x_M, c_1) &> dp_1\left(x_M, \frac{c_M}{2}\right) \\ &= D\left(\left(\frac{c_M}{2}, \frac{c_M}{2}\right)\right). \end{aligned}$$

This implies that $\vec{c} = (\frac{c_M}{2}, \frac{c_M}{2})$ results in less damage to the defender than any \vec{c} where $c_1 < \frac{c_M}{2}$. A symmetric argument holds for $c_1 > \frac{c_M}{2}$, thus proving the optimality of $\vec{c} = (\frac{c_M}{2}, \frac{c_M}{2})$. □

In other words, if all the parameters are symmetric, and the attacker budget constraint is sufficiently small such that the aforementioned criteria are satisfied, then the defender should evenly defend both targets. Figure

4.1 shows a plot of the defender's expected damage versus feasible values for \vec{c} when parameters are symmetric. The plot was generated by determining the attacker's optimal allocation towards target 1 for every feasible defender allocation towards target 1, and using those two allocations to calculate the expected damage. We see that the defender minimizes his expected damage by evenly defending both targets.




Figure 4.1: Expected damage to the defender when x_M is sufficiently small. $x_M = 5, c_M = 10, d = 7, a = 6, f = 5$. The minimum damage occurs when the defender invests $\frac{c_M}{2}$ into defending each target.

4.2 Unknown reward parameters

Since it has been difficult to prove results regarding the defender's optimal strategy when the rewards are unknown to the defender, we have instead made conjectures based on numerical simulation. Our numerical simulation consists of the following process: For every possible value of c_1 ,

1. Randomly generate a_1 and a_2 independently according to their distribution, which was assumed to be exponential.
2. Determine \vec{x}_{opt} numerically. In other words, for the generated a_i in

- step 1, determine the attacker allocation that maximizes his expected benefit.
3. Determine the defender's expected damage using that value of x_{opt1} and c_1 .
4. Repeat steps (1)–(3) 100 times.
5. Determine the defender's *average* expected damage based on the 100 iterations.

We then create a plot of the defender's expected damage versus c_1 . The minimum point marks the allocation towards defending target 1 where the defender approximately minimizes his expected damage. Note that we use only 100 iterations due to the long run-time of the algorithm; our plots, however, suggests that 100 iterations are sufficient to determine the general shape of the graph as well as the region of the minimum point in most cases.

We now propose several conjectures regarding the defender's optimal strategy.

Conjecture 4.2. *As the punishment at one target increases, and if the expected reward and attacker budget are not too small, then the defender's optimal allocation towards that target decreases.*

Intuitively, as the punishment at a target increases, the attacker has less incentive to attack that target, implying the defender could also decrease his allocation towards that same target. However, if the expected reward or attacker budget are relatively small, then the attacker may originally allocate very little to that target. Thus, if the punishment increases, chances are the attacker's, as well as the defender's strategy, changes very little. Figure 4.2 shows how the defender's optimal allocation towards target 1 decreases as the punishment at target 1 increases.

Conjecture 4.3. *Let the damage incurred by the defender at one target be neither too small nor large relative to the other parameters. As the damage increases, the defender's optimal allocation towards that target increases.*

It makes sense that the defender will more strongly defend a target that is more valuable. However, if the damage is too small at a target, then a slight increase gives the defender little incentive to increase protection to that target. If the damage is too large and he increases the allocation towards that target, the expected damage from that target is still relatively

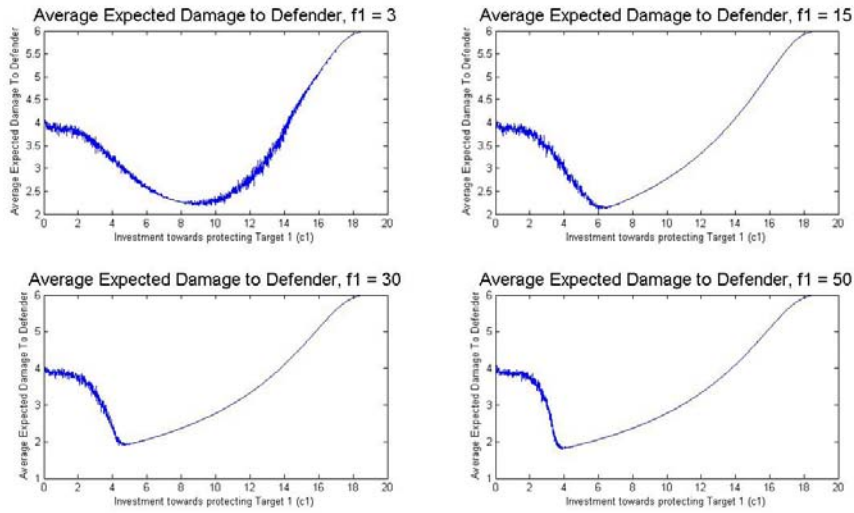


Figure 4.2: The optimal allocation towards target 1 decreases as f_1 increases. $x_M = 10, c_M = 20, d_1 = 4, d_2 = 6, E[a_1] = 1, E[a_2] = 8, f_2 = 7$. The minimum point shifts from about $c_1 = 8.3$ to $c_1 = 3.8$.

high, which again gives him little incentive to increase his allocation there. Figure 4.3 shows how the conjecture holds as d_1 increases from 4 to 40; one can see that the minimum point shifts to the right. However, when d_1 goes from 40 to 60, since d_1 is already relatively large, the defender's strategy does not differ by much.




Figure 4.3: The optimal allocation increases from about $c_1 = 10.3$ to about $c_1 = 18$ as d_1 increases from 4 to 40, but how it changes is unclear as d_1 increases from 40 to 60. $x_M = 10, c_M = 20, d_2 = 6, E[a_1] = 10, E[a_2] = 8, f_1 = 3, f_2 = 7$.

Conjecture 4.4. *Let the damage and expected reward at one target be neither too small nor too large relative to other parameters. As the expected reward increases, the defender's optimal allocation towards that target increases.*

An increase in the expected reward results in an increased attacker allocation towards that target, which implies the defender should do the same to minimize his expected damages. However, if the damage is too small or large, then the defender won't necessarily change his strategy for the same reasons. If the expected reward is too small, the attacker has little incentive to change his strategy because his expected benefit changes little anyways. If the expected reward is too large, the attacker will allocate most of his resources to secure it, so if the expected reward becomes even larger, the

attacker cannot increase his allocation towards that target by much; therefore, these two scenarios result in the defender responding the same way as before. Figure 4.4 shows how the optimal defender allocation towards target 1 increases when $E[a_1]$ increases from 1 to 50, but not from 50 to 100 as the expected reward is too large.

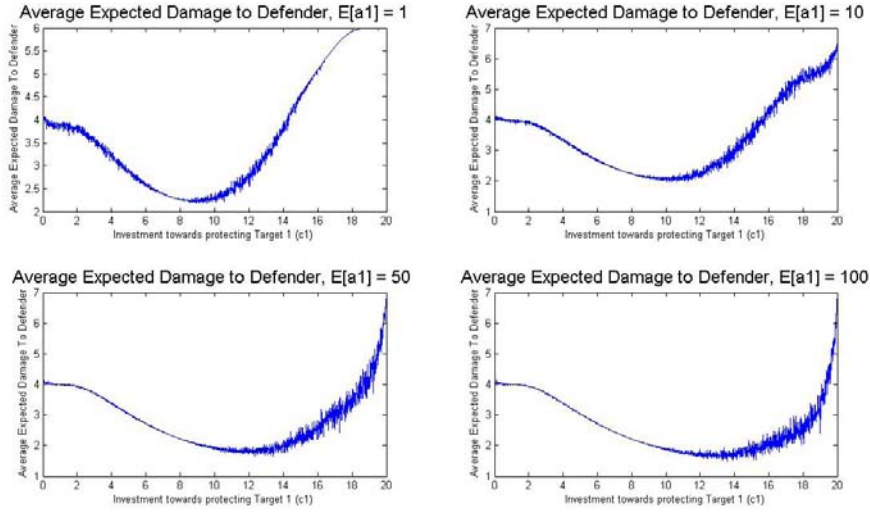


Figure 4.4: The optimal allocation increases from about $c_1 = 8.3$ to about $c_1 = 12.5$ as $E[a_1]$ increases from 1 to 50, but how it changes is unclear as $E[a_1]$ increases from 50 to 100. $x_M = 10, c_M = 20, d_1 = 4, d_2 = 6, E[a_2] = 8, f_1 = 3, f_2 = 7$.

Other characteristics regarding the defender's optimal strategy we explored include the *volatility* of the expected damage, where the volatility is defined as the vertical spread in the graph of the defender's expected damage at similar defense allocations. In addition, we also explored how his strategy changes as the attacker budget increases. Unfortunately, we were unable to find enough general trends to make any confident conjectures regarding these statements.

Figure 4.5 shows that the volatility to the left and the right of the defender's optimal allocation towards target 1 increases and decreases respectively as the punishment at target 1 increases. Figure 4.6 is an example of how difficult it is to observe general trends as the attacker budget grows; the volatility is huge, and one cannot easily determine the defender's opti-

mal allocation.




Figure 4.5: A look at the volatility of the defender's expected damage as f_1 increases from 4 to 50. $x_M = 10, c_M = 20, d_1 = 4, d_2 = 6, E[a_1] = 10, E[a_2] = 8, f_2 = 7$.

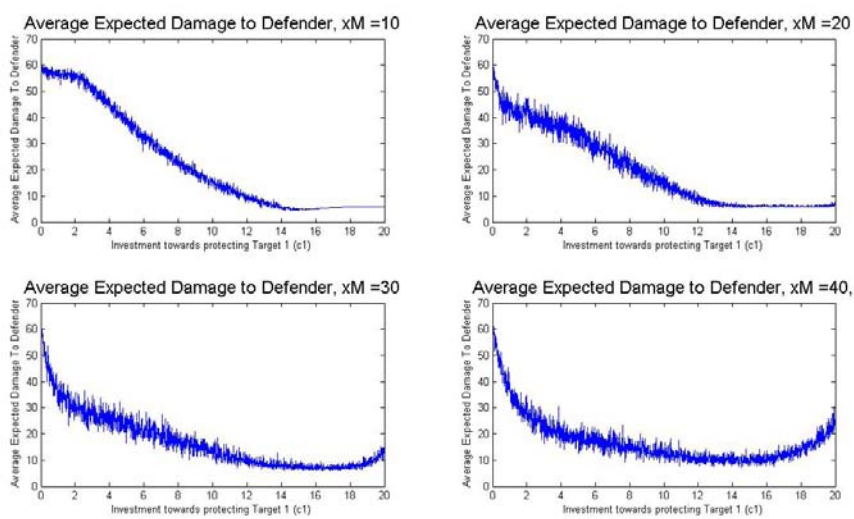


Figure 4.6: How the defender's expected damage changes as x_M increases from 10 to 40. $c_M = 20, d_1 = 60, d_2 = 6, E[a_1] = 1, E[a_2] = 8, f_1 = 3, f_2 = 7$.

Chapter 5

Future Work

We ultimately wish to solve for the defender's optimal strategy for all possible parameter values and relative magnitudes of the attacker budget constraint. The starting point for determining the defender's allocation in each case is to find the attacker's optimal investment. We have determined the attacker's optimal investment in the symmetric case where the attacker's budget is relatively small, as well as the number of targets the attacker should attack when his budget is relatively small or large. In addition, we have found sufficient conditions for the attacker's optimal investment when he will attack both targets. In regards to the defender's optimal strategy, we found that in a sufficiently small attacker budget, known symmetric parameter case, the defender should equally defend both targets. Finally, we have approximated the defender's optimal allocation numerically as well as observed general trends of how his strategy changes with respect to certain game parameters.

Short-term future work includes observing more numerical trends with the defender's optimal strategy, especially in regards to the attacker budget and the volatility of the expected damage. More long-term goals include proving results regarding such trends, as well as determining the defender's optimal strategy analytically.

Other interesting cases to explore include when the attacker's budget is neither small nor large enough to attack only one or both targets consistently for all feasible defense allocations. Currently, we have plots, like Figure 5.1, that indicate that the optimal strategy is to attack both targets for ranges of the defender allocation, and to attack only one target for the remaining ranges.

Other goals including determining criteria for when the attacker's bud-




Figure 5.1: Attacker's optimal strategy when x_M is neither sufficiently small nor large, symmetric parameter case. $x_M = 15, c_M = 10, d = 7, a = 6, f = 5$.

get is sufficiently large such that he will always attack both targets in a non-symmetric case, as well as a closed form solution for the attacker's optimal strategy when his budget is sufficiently small. Ultimately, we also hope to extend this model to the case of several targets, or budget constraints that are upper bounds on the defender's and attacker's allocations, not sunk costs.

Chapter 6

Conclusions

Due to the strong determination of terrorists to attack, complete prevention of the attacks is not always possible. In such scenarios, defenders must find strategies to minimize total damages to all the targets being attacked, while considering that defense investments at one target will affect a terrorist's decision to attack the other targets. We attempted to model this situation with two targets, a defender, and an attacker. The defender, and then the attacker, pick allocations for defense and attack respectively, each subject to a sunk cost budget constraint. The probability of a successful attack is dependent on both allocations, and there are reward, damage, and punishment parameters in the event of a successful and failed attack respectively. We first examined the attacker's optimal strategy in response to various defense allocations and used that information to determine the defender's optimal allocation.

We began by showing that an attacker should always attack an undefended target if it has a nonzero reward. Next, we found that when the attacker's budget is sufficiently small, attacking only one target is preferable to attacking both. If it is sufficiently large, then the attacker will choose to attack both targets. We also determined sufficient criteria for the attacker's optimal strategy if he were to attack both targets. Based on these findings, we then show that in a symmetric parameter case with a sufficiently small attacker budget and known reward parameters (to the defender), the defender can minimize his expected damages by equally defending both targets.

We also consider how the attacker's optimal strategy changes as the game parameters change. If either the reward or punishment parameter for a target increases, then the attacker should allocate more towards attacking

that target, given certain restrictions. Finally, if changes in the rewards in a symmetric parameter case result in new symmetric rewards, then the attacker's optimal strategy doesn't change (again, with certain restrictions).

We proceeded by determining the defender's optimal strategy numerically, as well as making conjectures regarding general trends of how the strategy changed with respect to the game parameters. We noticed that an increase in the damage or expected reward at a target tends to increase the defender's optimal allocation at that target, while an increase in the punishment results in the opposite effect. Since one must usually approximate ranges for the game parameters in real-world scenarios, such conjectures provide a strong starting point for creating a flexible defense strategy to account for such ranges.

These results, of course, provide stepping stones for future work. This includes proving results regarding the defender's optimal investment, as well as observing more numerical trends. Other work includes examining cases where the attacker budget is neither too large nor small and extending the model to multiple targets or flexible budgets.

With the recent terrorist attacks on multiple targets in numerous parts of the world, there has been demand for a systematic approach to studying risk and terrorism. The model we presented accounts for the multiple components of risk as well as incorporates elements from prior work done on this subject. In addition to proving many theorems regarding strategies, we were able to numerically calculate trends and solutions that have a direct application to real-world scenarios. These results as well as the ideas for future work form the necessary building blocks for the progression towards solving the open-ended problem of deterring terrorism.

Bibliography

- Abhichandani, V. and Bier, V. M. (2005). Optimal allocation of resources for defense and simple series and parallel systems from determined adversaries. *Reliability Engineering and System Safety*, 87:313–323.
- Arce M., D. G. and Sandler, T. (2005). Counterterrorism: A game-theoretic analysis. *Journal of Conflict Resolution*, 49(2):183–200.
- Bier, V. M. (2004). Should the model for security be game theory rather than reliability theory? *Communications of the Fourth International Conference on Mathematical Methods in Reliability: Methodology and Practice, Santa Fe, New Mexico*.
- Bier, V. M., Oliveros, S., and Samuelson, L. (2006). Choosing what to protect: Strategic defensive allocation against an unknown attacker. *Journal of Public Economic Theory*.
- Bier, V. M. and Zhuang, J. (2006). Balancing terrorism and natural disasters - defense strategy with endogenous attacker effort. *Operations Research*.
- Coughlin, P. J. (1992). Pure strategy equilibria in a class of system defense games. *International Journal of Game Theory*, 20:195–210.
- Martonosi, S. E. and Walton, D. (2006). Optimal defense allocations to deter terrorists.
- Roberson, B. (2006). The Colonel Blotto game. *Economic Theory*.
- Sandler, T. (2005). Collective versus unilateral responses to terrorism. *Public Choice*, 124:75–93.
- Sandler, T. and Lapan, H. E. (1988). The calculus of dissent: An analysis of terrorists choice of targets. *Synthese*, 76:245–261.

- Shogren, J. F. and Crocker, T. D. (1991). Cooperative and noncooperative protection against transferable and filterable externalities. *Environmental and Resource Economics*, 1.
- Shubik, M. and Weber, R. J. (1978). Competitive valuation of cooperative games. *Cowles Foundation Discussion Papers*, 482.
- Shubik, M. and Weber, R. J. (1981). System defense games: Colonel Blotto, command and control. *Naval Research Logistics Quarterly*, 28.
- Weber, R. J. (1978). Probabilistic values for games. *Cowles Foundation Discussion Papers*, 471.
- Willis, H. H. (2006). Guiding resource allocations based on terrorism risk. *Working Paper*.
- Woo, G. (2003). Insuring against Al-Qaeda. National Bureau of Economic Research Meeting.