More Targets to Attack Makes a Better Defense

Undergraduate Thesis
University of California, Davis

Denise Cerna
Supervised by Dr. Burkhard Schipper

June, 2023

Abstract

In recent years, Stackelberg security games have been used for the improvement of resource allocation in areas like counterterrorism in airport security and the protection of wildlife against poachers. In a Stackelberg security game, the defender optimally uses a mixed strategy to defend a set of targets against the attacker. The attacker then observes this strategy and decides which target to attack. In the current paper, we consider not only optimal strategies with respect to a given set of targets, but also the defender's ability to create or raise awareness of additional targets. We show that raising awareness/creating additional targets can be beneficial because the defender may introduce lower value targets that divert the adversary's attention away from other higher value targets.

Contents

1 Introduction

Game theory has been increasingly used to study security resource allocation problems relating to infrastructure and the enforcement of legal behavior, such as arms control policies. For instance, starting in the sixties, inspection games have been used for this type of application by modelling an inspector who must verify whether or not an inspectee obeys certain legal rules (Avenhaus et al, 2002). In equilibrium, the inspector's strategy is a statistical test with a false alarm probability, while the inspectee's strategy is a violation procedure, determining when he behaves illegally. Notably, this type of formulation has been used to represent situations like inspections led by the International Atomic Energy Agency's under the Nuclear Non-Proliferation Treaty. A key result of these inspection games determined in Avenhaus et al. 2002 has been that, as the leading player, the inspector benefits from committing to and revealing his strategy in comparison to a simultaneous-move game.

This feature inspires the use of Stackelberg games, which were first introduced as a way to model the impact of leadership and commitment to strategies in a duopoly setting (von Stackelberg 1934). In the standard model, a leader commits to a strategy first and the follower then observes this strategy before committing to his own. In recent years, new formulations of this game have been introduced specifically for security applications (Kar et al., 2018). The Stackelberg format is able to portray a worst case scenario for the leader, where information about her defense strategy is fully leaked or exposed to the attacker. Modelling the defense in terms of a leadership game can therefore help the defender identify an optimal strategy that she may even benefit from announcing. In contrast with inspection games, Stackelberg security games model the leading player as a defender who is trying to cover a set of targets. The follower is an attacker who chooses which target to attack after observing the defender's mixed strategy. Some notable implementations of these games include the scheduling of checkpoints and canine patrol routes at the LAX airport and the dispatching of US Federal Air Marshalls (Jain et al. 2010). Other recent applications include "Green Security Games" with emphasis on helping conservation agencies protect the environment. Since adversaries like poachers and illegal fishermen often engage in smaller attacks repeatedly, in these games, the defender is able to replan her mixed strategy after observing multiple attacks under the current strategy (Kar et al., 2018). Throughout these various domains, the ultimate goal is to optimally distribute a limited amount of resources to best protect the given assets in the event of an attack. These games can also serve to deter criminal behavior in general by making it more difficult to stage an effective attack.

In this paper, we analyze how a defender can utilize her power to expose certain targets and consequently influence the information available to the attacker. We first describe the general model and equilibria present in current literature on Stackelberg Security Games. Next we provide an example of a two-target Stackelberg security game illustrating the tie-breaking assumption necessary for the existence and uniqueness of an equilibrium. Further, we consider the defender's ability to create and raise awareness of additional targets. In both a leadership and simultaneous move setting, we are subsequently able to define conditions for a low-value target such that revealing low-value targets is beneficial to the defender.

2 Notation and Baseline Leadership Model with Full Awareness

To model a Stackelberg security game, we define a sequential two-stage extensive-form game as follows:

There are two players, a defender (player 1) and an attacker (player 2). The defender moves before the attacker.

There is a nonempty finite set of n targets, T where $T = \{t_1, t_2, ..., t_n\}$. The defender allocates a finite amount of defense resources over targets in T. W.l.o.g. let the total amount of defense resource be normalized to 1. Thus, we can understand the defender's strategy as a probability distribution τ over T. The set of the defender's strategies $\Delta(T)$ is the set of probability distributions over T.

The attacker is assumed to observe the defender's strategy. In other words, each of the defender's strategies corresponds to an information set of the attacker. Conditional on the defender's strategy observed, the attacker chooses whether to attack and, if yes, which target to attack. We assume that the attacker chooses to attack a target with probability 1 or no targets at all. Consequently, the attacker's set of actions is $T \cup \{\emptyset\}$, where $\emptyset$ denotes "no target attacked". The attacker chooses a strategy denoted by $s : \Delta(T) \longrightarrow T \cup \{\emptyset\}$ that assigns to each information set $\tau \in \Delta(T)$ an action $s(\tau) \in T \cup \{\emptyset\}$ of the attacker. We let S_2 denote the set of all attacker's strategies.

We say that a target $t \in T$ is *covered* if $\tau(t) > 0$. There are four outcomes for each target: covered and attacked, (c, a), covered and not attacked, $(c, \neg a)$, uncovered and attacked, (u, a), as well as uncovered and not attacked, $(u, \neg a)$.

We denote the utility to player $i \in \{1, 2\}$ from target t that is covered and attacked by $u_i^{c,a}(t) \in \mathbb{R}$. Analogous notation is used for all other outcomes.

Reasonable assumptions on payoffs in Stackelberg security games are as follows:

[1] $u_i^{c,\neg a}(t) = u_i^{u,\neg a}(t) := u_i^{\neg a}(t) \quad \forall i \in \{1, 2\}, \forall t \in T$

[2] $u_1^{c,a}(t) > u_1^{u,a}(t) \quad \forall t \in T$

[3] $u_2^{u,a}(t) > u_2^{c,a}(t) \quad \forall t \in T$

To elaborate, [1] follows from the commonly accepted assumption that, unless a target is attacked, the payoffs do not depend on whether a target is covered or not (Kiekintveld et al, 2009). This is used to model the game in a compact format, allowing for easier computation in algorithm development.

Requirements [2] and [3] describe the opposing interests of each player. For the defender, an attack on a covered target is always better than an attack on an uncovered target. Conversely, for the attacker, an attack on an uncovered target is always better than a covered target. This is less-restrictive than the zero-sum assumption. Security games are not explicitly zero-sum for reasons such as the varying cost of protecting different targets or the possibility of the attacker profiting from the media attention received by the attack even if they are caught (Pita, 2009).

The defender's expected payoff from an attack on a single target t' and strategy τ is

$$U_1(t', \tau) = \tau(t')u_1^{c,a}(t') + (1 - \tau(t'))u_1^{u,a}(t') + \sum_{t \in T \setminus \{t'\}} u_1^{\neg a}(t).$$

To define the attacker's expected payoff, the belief systems of the attacker are trivial as he observes the strategy of the defender. Given the defender's strategy τ, the attacker's expected payoff conditional on τ from attacking target t' is

$$U_2(t', \tau) = \tau(t')u_2^{c,a}(t') + (1 - \tau(t'))u_2^{u,a}(t') + \sum_{t \in T \setminus \{t'\}} u_2^{\neg a}(t).$$

In the case that no target is attacked, i.e., $s(\tau) = \{\emptyset\}$ the expected payoffs are simply the sum of each players payoff for a target that is not attacked.

$$U_1(\tau, \{\emptyset\}) = \sum_{t \in T} u_1^{\neg a}(t) = nu_1^{\neg a}(t).$$

$$U_2(\tau, \{\emptyset\}) = \sum_{t \in T} u_2^{\neg a}(t) = nu_2^{\neg a}(t).$$

2.1 Solution Concept

A Stackelberg Equilibrium is a refinement of Nash Equilibrium similar to a subgame perfect equilibrium. However, this alone does not guarantee a unique solution, since the follower can be indifferent among multiple strategies, which are in this case perceived as targets to attack (see next section). In order to break ties, two major rules have been defined in literature (Leitmann, 1978). These are generally referred to as "strong" and "weak" Stackelberg equilibria and are adaptations of the "pessimistic" and "optimistic" strategies described in von Stengel and Zamir (2004). In a strong Stackelberg equilibrium (SSE), the attacker breaks ties by choosing the best strategy for the defender, while in a weak Stackelberg equilibrium (WSE) he chooses the worst option for the defender. In Basar and Olsder (1995) it has been shown that a strong Stackelberg equilibrium will always exist while a weak Stackelberg equilibrium may not, due the defender's incentive to deviate from the equilibrium to influence the follower to play the preferred strategy. More recently, an "Inducible Stackelberg Equilibrium" has been defined as well to model instances where a defender cannot infinitesimally deviate from her strategy due to resource assignment constraints (Guo et al., 2018). However, it still shares the optimistic tie-breaking assumption of a SSE. Since it is the most widely accepted equilibrium concept in current literature, we will consider a strong Stackelberg equilibrium in our discussions of the leadership game. In the next section, we provide an example illustrating the importance of a strong Stackelberg Equilibrium for existence in a sequential game.

For simplicity, we define an attack set A containing all targets that provide the highest payoff to the attacker in response to a defender's strategy τ (Kiekintveld et al., 2009).

$$A(\tau) = \{t : U_2(t, \tau) \geq U_2(t', \tau), \forall t' \in T\}.$$

A formal definition of a Strong Stackelberg Equilibrium is as follows:

Definition 1 *A pair of strategies (s^*, τ^*) are a Strong Stackelberg Equilibrium (SSE) if:*

1. *The defender's action, τ^* is a best response to the defenders belief $b_1 \in \Delta(S_2)$.*

$$\sum_{s \in S_2} U_1(s(\tau^*), \tau^*) b_1(s) \geq \sum_{s \in S_2} U_1(s^*(\tau), \tau) b_1(s) \quad \forall \tau \in \Delta(T).$$

2. *With information set τ^*, the attacker's strategy, s^* is a best response to τ^*.*

$$U_2(s^*(\tau^*), \tau^*) \geq U_2(s(\tau^*), \tau^*) \quad \forall s \in S_2.$$

3. *The attacker breaks ties in favor of the defender.*

$$U_1(s^*(\tau^*), \tau^*) \geq U_1(t, \tau^*) \quad \forall t \in A(\tau^*).$$

2.2 Existence and Uniqueness of a Strong Stackelberg Equilibrium

In order to illustrate the importance of the tie-breaking assumption utilized in Stackelberg security games, consider a simple two-target example described in Figure 1 with payoffs for each player referenced in Tables 1 and 2. The leader is the row player and the follower is the column player. Note that the dashed line at the player 2 level represents both the set of probability distributions for player 1's mixed strategy between the two targets as well as the information set of player 2.

To find a mixed strategy for the defender, we use the principle of indifference. This means that the defender chooses τ such that the attacker receives the same expected payoff between pure strategies, which can be interpreted as an attack on a specific target. This is done so that neither player has any incentive to deviate from their strategy and so that the defender can benefit from the assumption that the attacker will break ties by choosing the target that gives her the greatest payoff. Since playing $\emptyset$ for the attacker

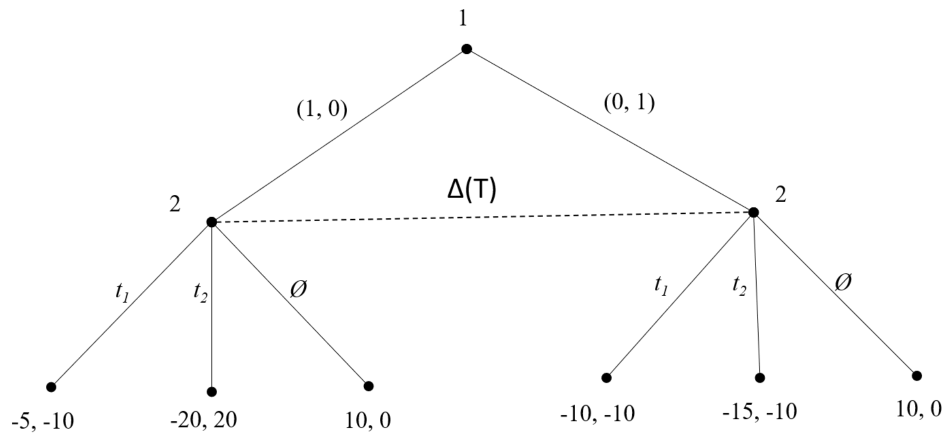

Figure 1: Two-Target Example

Defender Payoff

	a	¬ a
c	-10	5
u	-15	5

Target 1

	a	¬ a
c	-20	5
u	-25	5

Target 2

Attacker Payoff

	a	¬ a
c	-10	0
u	-10	0

Target 1

	a	¬ a
c	-10	0
u	20	0

Target 2

Table 1: Two-Target Game Payoffs

	t_1	t_2	Ø
t_1	-5, -10	-20, 20	10, 0
t_2	-10, 10	-15, -10	10, 0

Table 2: Summary Payoff Matrix

is dominated by a mixed strategy, such as $(\frac{1}{4}, \frac{3}{4})$, we can immediately eliminate it from subsequent calculations.

We denote the coverage of each target as $p = \tau(t_1)$ and $(1 - p) = \tau(t_2)$

Figure 2 displays the best response correspondence of the attacker, with each player's expected utility given in terms of the coverage probability p and the target that is attacked. Each red line represents payoffs given an attack on t_1, t_2, or $\emptyset$ as a function of p.

Notice that the attacker is indifferent between attacking t_1 and t_2 when $p = \frac{2}{5}$.

However, if the attacker chooses the worst option for the defender (weak Stackelberg Equilibrium), in this case t_2, the defender will deviate from her strategy and choose some $\varepsilon > 0$ so that $p = \frac{2}{5} - \varepsilon$. The attacker would then be incentivized to choose to attack t_1 instead. Since the defender will always be able to get arbitrarily close to $p = \frac{2}{5}$ based on the choice of ε, this weak Stackelberg Equilibrium does not exist. On the other hand, if

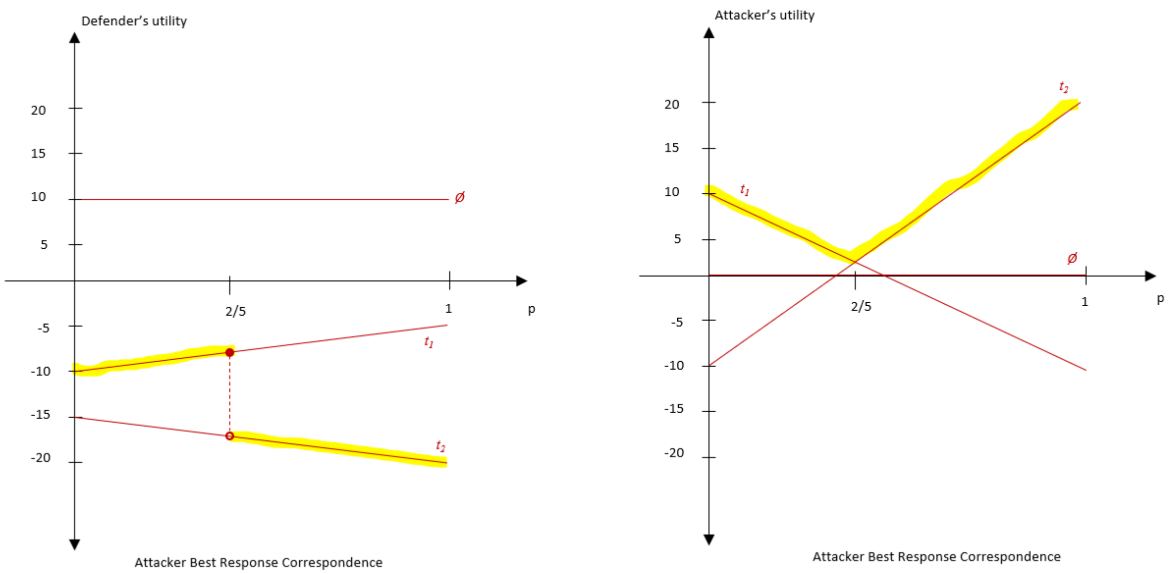

Figure 2: Illustration of Strong Stackelberg Equilibrium

we assume that the attacker chooses the best option for the defender, t_1, right away, this results in a Strong Stackelberg equilibrium that is unique and always exists since neither player has any incentive to deviate.

3 Model with Unawareness

In order to examine behavior in this game under unawareness, we can consider a similar problem and examine the results when a new target is added or revealed. Here, the defender is considering revealing more targets to the attacker. In the current formulation, we model this by assuming that the attacker will not know that a target exists unless the defender chooses to protect it with some probability greater than zero. This puts the defender in a situation where they can either reveal this target or leave it undefended under the belief that it is fully absent from the information set of the attacker. We use the following example to motivate the idea that the defender will have an incentive to reveal targets if their payoffs make them "low-value" in comparison to others.

3.1 Three-Target Example with Unawareness

Consider the following example of a three target game with two equal-value targets and one with parameterized payoff values. We assume that the attacker is initially unaware of the third target. In the case that the attacker remains unaware of the third target, he will observe and believe he is playing the game with only t_1 and t_2. This corresponds to

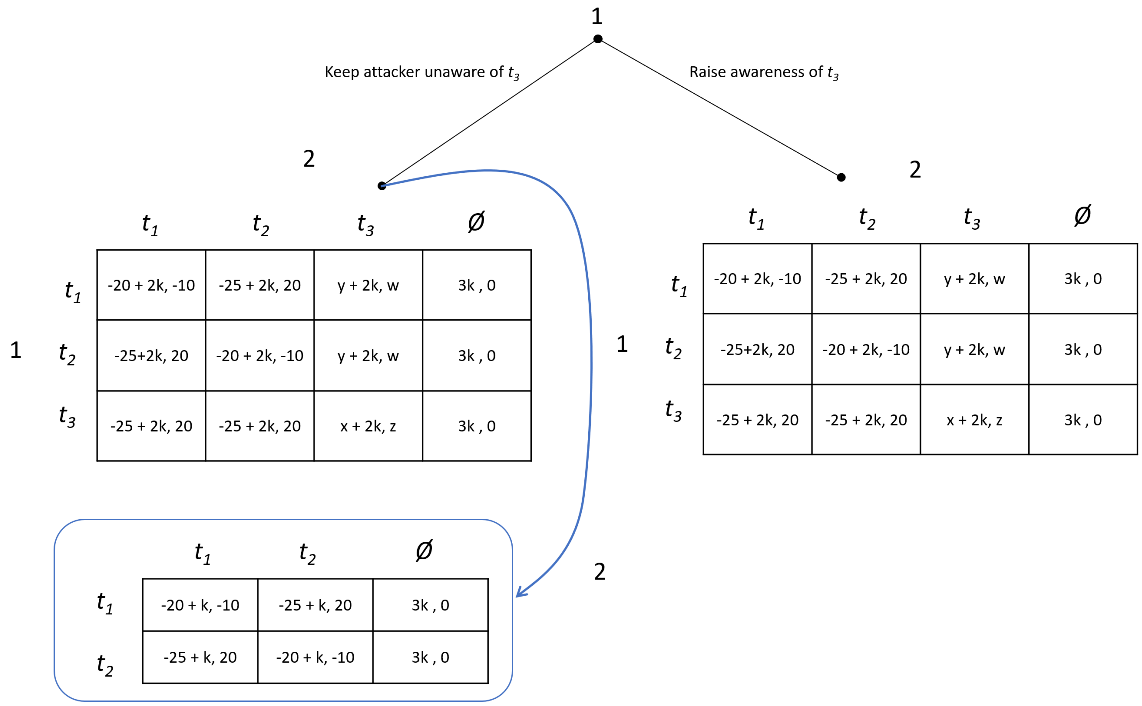

Figure 3: Three Target Example

the game circled in blue in Figure 3. In this case, the result is trivial since the targets are equal and the defender will choose to cover each target with probability $\frac{1}{2}$ and an attack on either target will provide her with the same payoff.

We parameterize the payoffs for t_3 to make it easier to discuss the game with targets of different values. Since we assume that the payoffs should not depend on targets that are not attacked, we let $u_1(t)^{\neg a}$ be some constant k such that $k > u_1^{c,a}(t) \quad \forall t \in T$ and $u_2^{\neg a}(t) = 0$.

Let

$$\begin{aligned} p &= \tau(t_1) \\ q &= \tau(t_2) \\ 1 - p - q &= \tau(t_3) \end{aligned}$$

For the Defender, when raising awareness of all three targets, the expected payoff of an attack on a target t_j is:

$$\begin{aligned} U_1(t_1) &= -25 + 2k + 5p \\ U_1(t_2) &= -25 + 2k + 5q \\ U_1(t_3) &= (p+q)(y-x) + x + 2k \\ U_1(\emptyset) &= 3k \end{aligned}$$

Table 3: Three-Target Game Payoffs

Defender Payoff

	a	¬ a
c	-20	k
u	-25	k

Target 1

	a	¬ a
c	-20	k
u	-25	k

Target 2

	a	¬ a
c	x	k
u	y	k

Target 3

Attacker Payoff

	a	¬ a
c	-10	0
u	20	0

Target 1

	a	¬ a
c	-10	0
u	20	0

Target 2

	a	¬ a
c	z	0
u	w	0

Target 3

Table 4: Summary Payoff Matrix

	t_1	t_2	t_3	∅
t_1	-20+2k, -10	-25+2k, 20	y+2k, w	3k, 0
t_2	-25+2k, 20	-20+2k, -10	y+2k, w	3k, 0
t_3	-25+2k, 20	-25+2k, 20	x+2k, z	3k, 0

For the Attacker, the expected payoff of an attack on a target t_j is:

$$\begin{aligned} U_2(t_1) &= 20 - 30p \\ U_2(t_2) &= 20 - 30q \\ U_2(t_3) &= z + (p+q)(w-z) \\ U_2(\emptyset) &= 0 \end{aligned}$$

Figure 3 describes the game tree where the defender must first choose whether or not to raise awareness of t_3. If yes, the game for both players is pictured to the right. However, if not, the defender plays the game with all three targets while the attacker on observes t_1 and t_2, circled in blue.

Note that if the payoffs for t_3 are identical to those of t_1 and t_2, the game is trivial and both the attacker and defender are indifferent between an attack on any of the three targets.

However, consider the case where t_3 seems less valuable than the other two targets. In other words, an attack on t_3 is not as damaging to the defender. For instance, let $x = -19$, $y = -24$, $z = -11$, and $w = 19$. Note that these payoffs still satisfy the general payoff requirements mentioned earlier.

A comparison of expected payoffs will show that, when p and q are chosen such that the attacker is indifferent between the three targets, the defender plays a mixed strategy of $(\frac{31}{90}, \frac{31}{90}, \frac{28}{90})$ and the attacker chooses the best option for the defender, in this case, t_3. This game results in an expected payoff of $-22.44 + 2k$ for the defender. This is slightly improved upon the $-22.5 + 2k$ equilibrium payoff that occurs when the attacker is unaware of t_3.

In contrast, if we consider a slightly higher value target where $x = -21$, $y = -26$, $z = -9$, and $w = 21$, we get an expected payoff of $-23.39 + 2k$, which is worse for the defender.

Although we have focused on the leadership setting, it may not always be feasible or of interest for the attacker to perfectly observe the defender's strategy at the beginning of a game. For a more robust discussion of behavior under the exposure of "low-value" targets, we will also consider the simultaneous move game where both players use a mixed strategy to make each other indifferent between pure outcomes.

In the same example described above, due to the symmetrical payoffs of t_1 and t_2, the results of the simultaneous move game when the attacker is unaware of the third target are the same as before. If we introduce the lower value target defined above, we get a mixture of strategies that yields an expected payoff of $-23 + 2k$ to the defender. This makes the defender worse off than before. However, if we let $x = -19$, $y = -22$, $z = -11$, and $w = 19$, we get a mixed strategy equilibrium that yields an expected payoff of $-22.27 + 2k$ to the defender. Thus past a certain point, adding a low value-target may be beneficial to the defender even if the attacker does not act as a follower. In both the sequential and simultaneous-move setting, exposing a low value target has the potential for yielding a higher payoff for the defender.

Further analysis requires a more specific characterization of a low-value target, which will be described in the next section.

3.2 Simultaneous Move Game: General Case

Despite the general use of the Stackelberg model for these security games, the fact that the attacker chooses to observe the defender's strategy before taking action or that the defender successfully announces her strategy is a strong assertion. Not every adversary is willing to plan thoroughly and, especially with small-scale attacks, they may take action more impulsively. To make sure adding low-value targets in a security setting is robust to more solution concepts, we first analyze this result in the case of a simultaneous-move game where both players choose a mixed strategy. We consider a game where the defender has $n + 1$ targets and believes that the attacker is aware of n of them. We will refer to the "hidden" target as t_{n+1}.

The attacker this time chooses a mixed strategy denoted by $s \in \Delta(T)$. The probability of an attack on some target t_j is denoted as $s(t_j)$.

The defender and attacker's expected payoff given a defender's mixed strategy τ and an attacker's mixed strategy s is now

$$U_1(s,\tau) = \sum_{j=1}^{n+1}[s(t_j)[\tau(t_j)u_1^{c,a}(t_j) + (1-\tau(t_j))u_1^{u,a}(t_j)] + (1-s(t_j))[u_1^{\neg a}(t_j)]]$$

$$U_2(s,\tau) = \sum_{j=1}^{n+1}[s(t_j)[\tau(t_j)u_2^{c,a}(t_j) + (1-\tau(t_j))u_2^{u,a}(t_j)] + (1-s(t_j))[u_2^{\neg a}(t_j)]]$$

Definition 2 *In a simultaneous move game, the mixed strategies τ^* and s^* form a Nash equilibrium if:*

1. *The defender plays a best response. So, $U_1(s^*,\tau^*) \geq U_1(s^*,\tau) \quad \forall \tau \in (\Delta T)$.*
2. *The attacker plays a best response. So, $U_2(s^*,\tau^*) \geq U_2(s,\tau^*) \quad \forall s \in (\Delta T)$.*

Within this simultaneous move game, we are able to incorporate the systematic payoff structure used in a Stackelberg security game to explicitly define an attacker's strategy s after removing a target from the information set.

Lemma 1 *Let (τ^*, s^*) be the Nash Equilibrium for a security game with n targets and let (τ^{**}, s^{**}) be the Nash Equilibrium for the game with $n+1$ targets.*

Under the payoff assumptions of a Stackelberg security game, the probability of an attack placed on target t_j in an attacker's strategy s^* can be defined as

$$s^*(t_j) = \frac{s^{**}(t_j)}{\sum_{i=1}^{n} s^{**}(t_i)}$$

PROOF. By the construction of a Stackelberg security game, there are only two possible payoff options for each player at each pure strategy outcome. Either, the target is uncovered and attacked, or, it is covered and attacked.

Table 5 provides a generalization of the defender's payoff matrix with $u_1(t_j, t_k)$ being the defender's utility when covering t_k with probability 1 given an attack on t_j with probability 1. Here, t_k is the row target and t_j is the column target. Notice that the covered and attacked utility for each column target lies on the diagonal while the uncovered and attacked utility for the column target fills in the remaining column values. This is because, following from payoff assumption [1], we assume that each player's utility only depends on the target that is attacked. For instance $u_1(t_1, t_2) = u_1^{u,a}(t_1)$ because even though t_2 is covered, it is t_1 that is attacked.

Since (τ^*, s^*) and (τ^{**}, s^{**}) are Nash equilibria, the attacker chooses a strategy such that the defender is indifferent, always receiving the same payoff no matter which target is covered.

Table 5: Generalized Defender Payoff Matrix $(u_1(t_j, t_k))$

	t_1	t_2	...	t_{n+1}
t_1	$u_1^{c,a}(t_1)$	$u_1^{u,a}(t_2)$	...	$u_1^{u,a}(t_{n+1})$
t_2	$u_1^{u,a}(t_1)$	$u_1^{c,a}(t_2)$	...	$u_1^{u,a}(t_{n+1})$
...	...	...	...	...
t_{n+1}	$u_1^{u,a}(t_1)$	$u_1^{u,a}(t_2)$	...	$u_1^{c,a}(t_{n+1})$

In other words, the defender's expected payoff in the larger game is $\sum_{j=1}^{n+1} s^{**}(t_j)u_1(t_j, t_k) = U_1(s^{**}, t_k) = U_1(s^{**}, \tau^{**})$ for every row target $k \in \{1, 2, ..., n+1\}$.

Notice that $\sum_{j=1}^{n} s^{**}(t_j)u_1(t_j, t_k) = U_1(s^{**}, \tau^{**}) - s^{**}(t_{n+1})u_1(t_{n+1}, t_k)$.

By the payoff structure of the security game $s^{**}(t_{n+1})u_1(t_{n+1}, t_k) = u_1^{u,a}(t_{n+1})$ is constant for every row $k \in \{1, 2, ..., n\}$.

Thus, $\sum_{j=1}^{n} s^{**}(t_j)u_1(t_j, t_k) = U_1(s^{**}, \tau^{**}) - s^{**}(t_{n+1})u_1(t_{n+1}, t_k)$ is also a constant in every row $k \in \{1, 2, ..., n\}$, even if the t_{n+1} column was eliminated.

Since, we need the defender's expected payoff in the smaller game $\sum_{j=1}^{n} s^{*}(t_j)u_1(t_j, t_k) = U_1(s^{*}, t_k) = U_1(s^{*}, \tau^{*})$ to be constant for every $k \in \{1, 2, ..., n\}$, for the mixed strategy of the smaller game, we can consider s^* to be s^{**} with $s^{**}(t_{n+1}) = 0$.

So, conditioning on the attacker only being aware of the smaller game, we can define s^* such that,

$$s^*(t_j) = \frac{s^{**}(t_j)}{\sum_{i=1}^{n} s^{**}(t_i)}.$$

$\square$

This result helps in the comparison of expected payoffs when defining a condition for a low value target in Proposition 1.

Proposition 1 *$U_1(s^{**}, \tau^{**}) > U_1(s^*, \tau^*)$ if and only if $u_1^{u,a}(t_{n+1}) > U_1(s^*, \tau^*)$*

In words, for the defender to receive a higher payoff by revealing a target t_{n+1} in a simultaneous move game, it is both necessary and sufficient for $u_1^{u,a}(t_{n+1}) > U_1(s^*, \tau^*)$, where s^* and τ^* are the mixed strategies that form a Nash equilibrium in the game with n targets.

If this condition is satisfied, we will refer to t_{n+1} as a low-value target.

The proof is provided in section *A.1* of the Appendix. The argument stems from writing out $U_1(s^{**}, \tau^{**}) > U_1(s^*, \tau^*)$ in terms of s^{**} using Lemma 1. Intuitively, if $u_1^{u,a}(t_{n+1}) > U_1(s^*, \tau^*)$ this will increase the "average" payoff across targets and improve the outcome for the defender.

3.3 Strong Stackelberg Equilibrium: General Case

Now we consider a low-value target characterization in a Stackelberg security game where the defender has $n+1$ targets and believes that the attacker is aware of n of them. We will refer to the "hidden" target as t_{n+1}.

In the game where the defender does not raise awareness of t_{n+1}, let the attacker's pure strategy in equilibrium be t_n and the defender's mixed strategy be τ^*. Similarly, in the game where the defender does raise awareness of t_{n+1}, let the defender's mixed strategy be τ^{**}

By the requirements for a Strong Stackelberg Equilibrium, it follows that in the smaller game,

$$U_1(t_n, \tau^*) \geq U_1(t, \tau^*) \quad \forall t \in A(\tau^*) \quad \text{and} \quad U_2(t_n, \tau^*) = U_2(t, \tau^*) \quad \forall t \in A(\tau^*).$$

This indicates that the defender receives the highest payoff possible from the targets in the attack set and the attacker is indifferent between targets.

In order for raising awareness of t_{n+1} to be beneficial to the defender in the larger game, we similarly need to satisfy the requirements for a Strong Stackelberg Equilibrium and make sure that the defender gets a better payoff:

[i] $U_2(t_{n+1}, \tau^{**}) = U_2(t, \tau^{**}), \forall t \in A(\tau^{**})$

[ii] $U_1(t_{n+1}, \tau^{**}) > U_1(t_n, \tau^{**}), \forall t \in A(\tau^{**})$

[iii] $U_1(t_{n+1}, \tau^{**}) > U_1(t_n, \tau^*)$

Here, [i] guarantees that the attacker is still indifferent between all targets.

Additionally, [ii] states that the attacker chooses to attack t_{n+1} for the larger game with $n+1$ targets. Since the attacker breaks ties in favor of the defender, this means that attacking t_{n+1} gives the defender the highest payoff.

Finally, [iii] guarantees that the defender's payoff in the game with $n+1$ targets is better than in the smaller game with n targets.

Since the coverage of the set of targets depends on the defender's mixed strategy, we first establish a way to write τ in terms of the basic payoffs for each target, given that the attacker is indifferent between targets.

Lemma 2 *[i] holds if and only if* $\tau^{**}(t_{n+1}) = \dfrac{1-\sum_{k=1}^{n} \frac{u_2^{u,a}(t_{n+1})-u_2^{u,a}(t_k)}{u_2^{c,a}(t_k)-u_2^{u,a}(t_k)}}{1+\sum_{k=1}^{n} \frac{u_2^{c,a}(t_{n+1})-u_2^{u,a}(t_{n+1})}{u_2^{c,a}(t_k)-u_2^{u,a}(t_k)}} \forall t \in T$

The proof can be found in section *A.2* of the Appendix and follows from the definition of expected payoff and the payoff structure described in Table 5.

Since the strategy τ^{**} and number of targets $n+1$ are arbitrary, this formula can be used to find τ for any target in a game of any size.

Proposition 2 *$U_1(s^{**}, \tau^{**}) > U_1(s^*, \tau^*)$ if and only if $\tau^{**}(t_{n+1}) > \frac{U_1(s^*,\tau^*) - u_1^{u,a}(t_{n+1})}{u_1^{c,a}(t_{n+1}) - u_1^{u,a}(t_{n+1})}$*

where $\tau^{**}(t_{n+1})$ is defined by Lemma 2 and $s^*(\tau^*) = t_n$ is the attacker's strategy in the game with n targets and $s^{**}(\tau^{**}) = t_{n+1}$ is the attacker's strategy in the game with $n+1$ targets. Equivalently, $U_1(s^{**}, \tau^{**}) > U_1(s^*, \tau^*)$ is true if and only if [i], [ii]. [iii] hold.

If this condition is satisfied, we will refer to t_{n+1} as a low-value target.

The proof is provided in section *A.3* of the Appendix and an alternative condition written fully in terms of primitives is derived in section *A.4*. We can interpret this characterization by noticing that the right hand side must be less than one since $\tau^{**}(t_{n+1})$ is a probability. Furthermore, $u_1^{c,a}(t_{n+1}) > u_1^{u,a}(t_{n+1})$ by the payoff assumptions of a security game and $U_1(s^*, \tau^*) > u_1^{u,a}(t_{n+1})$ otherwise the defender would strictly prefer an attack on an uncovered target. Thus, the ratio must be positive and as $u_1^{c,a}(t_{n+1})$ increases, the right hand side goes to zero. It follows that if $u_1^{c,a}(t_{n+1})$ is sufficiently large, the inequality will hold and the new target can be considered low-value.

4 Discussion

The analysis in this paper provides support for a defender's incentive to expose reasonably-defined low value targets. By discovering the underlying payoff structure of security games, we were able to characterize specific conditions that indicate when a defender would be better off revealing hidden targets in both a simultaneous move game and the original Stackelberg game. In practice, this type of behavior can potentially be seen through the creation of decoy targets. If the defender has a strong belief that an attacker is unaware of a high value target, she could keep it uncovered to prevent drawing attention and raising awareness of it. However, she could also benefit from actively creating and raising awareness of less valuable targets, subject to the assumed solution concept, effectively spreading out the risk of an attack. This would make it much harder for an attacker to plan and commit to an effective action.

Future work could be done to examine low-value target characterizations for additional solution concepts and extensions incorporating more resource and scheduling constraints. This could mean determining a condition for a low-value target that is robust across multiple solution concepts at once. Additional work could also be done in reformulating a weak Stackelberg Equilibrium to allow for the attacker to have a pessimistic tie breaking assumption, which intuitively seems more realistic. It would also be interesting to examine revealing targets in the context of repeated games where the defender would be able to re-evaluate her strategy after observing the attacker's behavior.

5 Acknowledgements

I would like to thank Professor Burkhard Schipper for providing me with incredible guidance and support throughout this research and helping me gain a deeper understanding of game theory. This project has inspired me to keep studying the intersection of mathematics and economics and I am very grateful to the UC Davis Mathematics department for this opportunity.

Appendix A

A.1 Proof of Proposition 1

For the game with full awareness to benefit the defender, we need $U_1(s^{**}, \tau^{**}) > U_1(s^*, \tau^*)$. By Lemma 1, we can write s^* in terms of s^{**} and use the defender payoffs written in Table 5.

We show the first direction by assuming that $U_1(s^{**}, \tau^{**}) > U_1(s^*, \tau^*)$ for all rows $k \in \{1, 2, ..., n\}$. Let $u_1(t_j, t_k)$ represent the defenders utility when covering target t_k given an attack on target t_j.

$$
\begin{aligned}
U_1(s^{**}, \tau^{**}) &> U_1(s^*, \tau^*) \\
s^{**}(t_1)u(t_1, t_k) + ... + s^{**}(t_{n+1})u(t_{n+1}, t_k) &> \frac{s^{**}(t_1)}{\sum_{k=1}^{n} s^{**}(t_j)}u(t_1, t_k) + ... + \frac{s^{**}(t_n)}{\sum_{k=1}^{n} s^{**}(t_j)}u(t_n, t_k)) \\
s^{**}(t_{n+1})u(t_{n+1}, t_k) &> \left(\frac{s^{**}(t_1)}{\sum_{k=1}^{n} s^{**}(t_j)} - s^{**}(t_1)\right)u(t_1, t_k) + ... + \left(\frac{s^{**}(t_n)}{\sum_{k=1}^{n} s^{**}(t_j)} - s^{**}(t_n)\right)u(t_n, t_k) \\
(1 - \sum_{k=1}^{n} s^{**}(t_j))u(t_{n+1}, t_k) &> u(t_1, t_k)s^{**}(t_1)\left(\frac{1 - \sum_{k=1}^{n} s^{**}(t_j)}{\sum_{k=1}^{n} s^{**}(t_j)}\right) + ... + u(t_n, t_k)s^{**}(t_n)\left(\frac{1 - \sum_{k=1}^{n} s^{**}(t_j)}{\sum_{k=1}^{n} s^{**}(t_j)}\right) \\
u(t_{n+1}, t_k) &> \frac{s^{**}(t_1)}{\sum_{k=1}^{n} s^{**}(t_j)}u(t_1, t_k) + ... + \frac{s^{**}(t_n)}{\sum_{k=1}^{n} s^{**}(t_j)}u(t_n, t_k) \\
u_1^{u,a}(t_{n+1}) &> U_1(s^*, \tau^*)
\end{aligned}
$$

Next, we show the second direction by assuming that $u_1^{u,a}(t_{n+1}) > U_1(s^*, \tau^*)$.

$$
\begin{aligned}
u_1^{u,a}(t_{n+1}) &> U_1(s^*, \tau^*) \\
u(t_{n+1}, t_k) &> \frac{s^{**}(t_1)}{\sum_{k=1}^{n} s^{**}(t_j)}u(t_1, t_k) + ... + \frac{s^{**}(t_n)}{\sum_{k=1}^{n} s^{**}(t_j)}u(t_n, t_k) \\
(1 - \sum_{k=1}^{n} s^{**}(t_j))u(t_{n+1}, t_k) &> u(t_1, t_k)s^{**}(t_1)\left(\frac{1 - \sum_{k=1}^{n} s^{**}(t_j)}{\sum_{k=1}^{n} s^{**}(t_j)}\right) + ... + u(t_n, t_k)s^{**}(t_n)\left(\frac{1 - \sum_{k=1}^{n} s^{**}(t_j)}{\sum_{k=1}^{n} s^{**}(t_j)}\right) \\
s^{**}(t_{n+1})u(t_{n+1}, t_k) &> \left(\frac{s^{**}(t_1)}{\sum_{k=1}^{n} s^{**}(t_j)} - s^{**}(t_1)\right)u(t_1, t_k) + ... + \left(\frac{s^{**}(t_n)}{\sum_{k=1}^{n} s^{**}(t_j)} - s^{**}(t_n)\right)u(t_n, t_k) \\
s^{**}(t_1)u(t_1, t_k) + ... + s^{**}(t_{n+1})u(t_{n+1}, t_k) &> \frac{s^{**}(t_1)}{\sum_{k=1}^{n} s^{**}(t_j)}u(t_1, t_k) + ... + \frac{s^{**}(t_n)}{\sum_{k=1}^{n} s^{**}(t_j)}u(t_n, t_k)) \\
U_1(s^{**}, \tau^{**}) &> U_1(s^*, \tau^*)
\end{aligned}
$$

Thus, $U_1(s^{**}, \tau^{**}) > U_1(s^*, \tau^*) \iff u_1^{u,a}(t_{n+1}) > U_1(s^*, \tau^*)$

This condition can also be written fully in terms of primitives where

$$U_1(s^*, \tau^*) = \sum_{j=1}^{n} [s^*(t_j)[\tau^*(t_j)u_1^{c,a}(t_j) + (1 - \tau^*(t_j))u_1^{u,a}(t_j)] + (1 - s^*(t_j))[u_1^{\neg a}(t_j)]]$$

and $\tau^*(t_j)$ is written in terms of primitives using Lemma 2 as follows:

$$\tau^*(t_j) = \frac{1 - \sum_{t_k \in T, t_k \neq t_j} \frac{u_2^{u,a}(t_j) - u_2^{u,a}(t_k)}{u_2^{c,a}(t_k) - u_2^{u,a}(t_k)}}{1 + \sum_{t_k \in T, t_k \neq t_j} \frac{u_2^{c,a}(t_j) - u_2^{u,a}(t_j)}{u_2^{c,a}(t_k) - u_2^{u,a}(t_k)}}$$

□

A.2 Proof of Lemma 2

By the definition of a strong Stackelberg Equilibrium, the attacker must maintain indifference between targets. Thus, we can compare the expected payoffs of an attack on each target $t_1, t_2, ..., t_n$, or more generally $t \neq t_{n+1}$ to the expected payoff of an attack on t_{n+1}.

$$\begin{aligned}
\tau^{**}(t_{n+1})u_2^{c,a}(t_{n+1}) + (1 - \tau^{**}(t_{n+1}))u_2^{u,a}(t_{n+1}) &= \tau^{**}(t_1)u_2^{c,a}(t_1) + (1 - \tau^{**}(t_1))u_2^{u,a}(t_1) \\
&= \tau^{**}(t_2)u_2^{c,a}(t_2) + (1 - \tau^{**}(t_2))u_2^{u,a}(t_2) \\
&\vdots \\
&= \tau^{**}(t_n)u_2^{c,a}(t_n) + (1 - \tau^{**}(t_n))u_2^{u,a}(t_n)
\end{aligned}$$

By re-arranging, we get,

$$\begin{aligned}
\tau^{**}(t_1) &= \frac{\tau^{**}(t_{n+1})[u_2^{c,a}(t_{n+1}) - u_2^{u,a}(t_{n+1})] + u_2^{u,a}(t_{n+1}) - u_2^{u,a}(t_1)}{u_2^{c,a}(t_1) - u_2^{u,a}(t_1)} \\
\tau^{**}(t_2) &= \frac{\tau^{**}(t_{n+1})[u_2^{c,a}(t_{n+1}) - u_2^{u,a}(t_{n+1})] + u_2^{u,a}(t_{n+1}) - u_2^{u,a}(t_2)}{u_2^{c,a}(t_2) - u_2^{u,a}(t_2)} \\
&\vdots \\
\tau^{**}(t_n) &= \frac{\tau^{**}(t_{n+1})[u_2^{c,a}(t_{n+1}) - u_2^{u,a}(t_{n+1})] + u_2^{u,a}(t_{n+1}) - u_2^{u,a}(t_n)}{u_2^{c,a}(t_n) - u_2^{u,a}(t_n)}
\end{aligned}$$

Since τ^{**} is a probability distribution, we must have that $\tau^{**}(t_1) + \tau^{**}(t_2) + ... + \tau^{**}(t_n) + \tau^{**}(t_{n+1}) = 1$.

By substituting in the equations above, we get:

$$\tau^{**}(t_{n+1}) + \sum_{k=1}^{n} \frac{\tau^{**}(t_{n+1})[u_2^{c,a}(t_{n+1}) - u_2^{u,a}(t_{n+1})] + u_2^{u,a}(t_{n+1}) - u_2^{u,a}(t_k)}{u_2^{c,a}(t_k) - u_2^{u,a}(t_k)} = 1$$

$$\tau^{**}(t_{n+1}) + \tau^{**}(t_{n+1}) \sum_{k=1}^{n} \frac{u_2^{c,a}(t_{n+1}) - u_2^{u,a}(t_{n+1})}{u_2^{c,a}(t_k) - u_2^{u,a}(t_k)} + \sum_{k=1}^{n} \frac{u_2^{u,a}(t_{n+1}) - u_2^{u,a}(t_k)}{u_2^{c,a}(t_k) - u_2^{u,a}(t_k)} = 1$$

$$\tau^{**}(t_{n+1}) \left[1 + \sum_{k=1}^{n} \frac{u_2^{c,a}(t_{n+1}) - u_2^{u,a}(t_{n+1})}{u_2^{c,a}(t_k) - u_2^{u,a}(t_k)} \right] = 1 - \sum_{k=1}^{n} \frac{u_2^{u,a}(t_{n+1}) - u_2^{u,a}(t_k)}{u_2^{c,a}(t_k) - u_2^{u,a}(t_k)}$$

Thus,

$$\tau^{**}(t_{n+1}) = \frac{1 - \sum_{k=1}^{n} \frac{u_2^{u,a}(t_{n+1}) - u_2^{u,a}(t_k)}{u_2^{c,a}(t_k) - u_2^{u,a}(t_k)}}{1 + \sum_{k=1}^{n} \frac{u_2^{c,a}(t_{n+1}) - u_2^{u,a}(t_{n+1})}{u_2^{c,a}(t_k) - u_2^{u,a}(t_k)}}$$

Since t_{n+1} can technically be any target and τ^{**} is simply representing the defender's mixed strategy in a game of size $n+1$, this formula can be used to find the mixed strategy of the defender in a game of any size. We can generalize the formula for the probability placed on target $t_j \in T$ in a mixed strategy τ^* to be $\tau^*(t_j) = \frac{1 - \sum_{t_k \in T, t_k \neq t_j} \frac{u_2^{u,a}(t_j) - u_2^{u,a}(t_k)}{u_2^{c,a}(t_k) - u_2^{u,a}(t_k)}}{1 + \sum_{t_k \in T, t_k \neq t_j} \frac{u_2^{c,a}(t_j) - u_2^{u,a}(t_j)}{u_2^{c,a}(t_k) - u_2^{u,a}(t_k)}}$

$\square$

A.3 Proof of Proposition 2

We need statements [ii] and [iii] to hold to ensure t_{n+1} is a SSE and the defender receives a better payoff when t_{n+1} is revealed and attacked, meaning $U_1(s^{**}, \tau^{**}) > U_1(s^*, \tau^*)$

From [ii], we get that

$$U_1(t_{n+1}, \tau^{**}) > U_1(t_n, \tau^{**}) \iff$$

$$\tau^{**}(t_{n+1}) u_1^{c,a}(t_{n+1}) + (1 - \tau^{**}(t_{n+1})) u_1^{u,a}(t_{n+1}) > \tau^{**}(t_n) u_1^{c,a}(t_n) + (1 - \tau^{**}(t_n)) u_1^{u,a}(t_n)$$

$$\tau^{**}(t_{n+1})[u_1^{c,a}(t_{n+1}) - u_1^{u,a}(t_{n+1})] + u_1^{c,a}(t_{n+1}) > \tau^{**}(t_n)[u_1^{c,a}(t_n) - u_1^{u,a}(t_n)] + u_1^{c,a}(t_n)$$

$$\tau^{**}(t_{n+1}) > \frac{\tau^{**}(t_n)[u_1^{c,a}(t_n) - u_1^{u,a}(t_n)] + u_1^{u,a}(t_n) - u_1^{u,a}(t_{n+1})}{u_1^{c,a}(t_{n+1}) - u_1^{u,a}(t_{n+1})}$$

Similarly, from [iii], we get that

$$\begin{aligned} U_1(t_{n+1},\tau^{**}) > U_1(t_n,\tau^*) \iff \\ \tau^{**}(t_{n+1})u_1^{c,a}(t_{n+1}) + (1-\tau^{**}(t_{n+1}))u_1^{u,a}(t_{n+1}) &> U_1(t_n,\tau^*) \\ \tau^{**}(t_{n+1})[u_1^{c,a}(t_{n+1}) - u_1^{u,a}(t_{n+1})] + u_1^{c,a}(t_{n+1}) &> U_1(t_n,\tau^*) \\ \tau^{**}(t_{n+1}) &> \frac{U_1(t_n,\tau^*) - u_1^{u,a}(t_{n+1})}{u_1^{c,a}(t_{n+1}) - u_1^{u,a}(t_{n+1})} \end{aligned}$$

However, notice that since the distribution τ^{**} includes an additional target that has to be covered, $\tau^*(t_n) > \tau^{**}(t_n)$. By the following comparison, [2] must follow from [3]:

$$\begin{aligned} \frac{\tau^*(t_n)[u_1^{c,a}(t_n) - u_1^{u,a}(t_n)] + u_1^{u,a}(t_n) - u_1^{u,a}(t_{n+1})}{u_2^{c,a}(t_{n+1}) - u_1^{u,a}(t_{n+1})} &> \frac{\tau^{**}(t_n)[u_1^{c,a}(t_n) - u_1^{u,a}(t_n)] + u_1^{u,a}(t_n) - u_1^{u,a}(t_{n+1})}{u_1^{c,a}(t_{n+1}) - u_1^{u,a}(t_{n+1})} \\ \frac{U_1(t_n,\tau^*) - u_1^{u,a}(t_{n+1})}{u_1^{c,a}(t_{n+1}) - u_1^{u,a}(t_{n+1})} &> \frac{\tau^{**}(t_n)[u_1^{c,a}(t_n) - u_1^{u,a}(t_n)] + u_1^{u,a}(t_n) - u_1^{u,a}(t_{n+1})}{u_1^{c,a}(t_{n+1}) - u_1^{u,a}(t_{n+1})} \end{aligned}$$

Hence, for the defender to increase their payoff by revealing t_{n+1}, it must be the case that

$$\tau^{**}(t_{n+1}) > \frac{U_1(t_n,\tau^*) - u_1^{u,a}(t_{n+1})}{u_1^{c,a}(t_{n+1}) - u_1^{u,a}(t_{n+1})}$$

where $\tau^{**}(t_{n+1})$ is defined by Lemma 2

□

A.4 Alternative Form of Proposition 2

An alternative formulation can be found by rewriting $U_1(t_n,\tau^*)$ in terms of primitives.

Since Lemma 2 defines τ for any possible target in a game of given size, we can see that

$$\tau^*(t_n) = \frac{1 - \sum_{k=1}^{n-1} \frac{u_2^{u,a}(t_n) - u_2^{u,a}(t_k)}{u_2^{c,a}(t_k) - u_2^{u,a}(t_k)}}{1 + \sum_{k=1}^{n-1} \frac{u_2^{c,a}(t_n) - u_2^{u,a}(t_n)}{u_2^{c,a}(t_k) - u_2^{u,a}(t_k)}}$$

$$\text{Similarly, } 1-\tau^*(t_n) = \frac{1+\sum\limits_{k=1}^{n-1}\frac{u_2^{c,a}(t_n)-u_2^{u,a}(t_n)}{u_2^{c,a}(t_k)-u_2^{u,a}(t_k)}}{1+\sum\limits_{k=1}^{n-1}\frac{u_2^{c,a}(t_n)-u_2^{u,a}(t_n)}{u_2^{c,a}(t_k)-u_2^{u,a}(t_k)}} - \frac{1-\sum\limits_{k=1}^{n-1}\frac{u_2^{u,a}(t_n)-u_2^{u,a}(t_k)}{u_2^{c,a}(t_k)-u_2^{u,a}(t_k)}}{1+\sum\limits_{k=1}^{n-1}\frac{u_2^{c,a}(t_n)-u_2^{u,a}(t_n)}{u_2^{c,a}(t_k)-u_2^{u,a}(t_k)}}$$

$$= \frac{\sum\limits_{k=1}^{n-1}\frac{u_2^{c,a}(t_n)-u_2^{u,a}(t_k)}{u_2^{c,a}(t_k)-u_2^{u,a}(t_k)}}{1+\sum\limits_{k=1}^{n-1}\frac{u_2^{c,a}(t_n)-u_2^{u,a}(t_n)}{u_2^{c,a}(t_k)-u_2^{u,a}(t_k)}}$$

By finding the expected utility of the smaller game we get,

$$U_1(s^*,\tau^*) = \frac{1-\sum\limits_{k=1}^{n-1}\frac{u_2^{u,a}(t_n)-u_2^{u,a}(t_k)}{u_2^{c,a}(t_k)-u_2^{u,a}(t_k)}}{1+\sum\limits_{k=1}^{n-1}\frac{u_2^{c,a}(t_n)-u_2^{u,a}(t_n)}{u_2^{c,a}(t_k)-u_2^{u,a}(t_k)}} u_1^{c,a}(t_n) + \frac{\sum\limits_{k=1}^{n-1}\frac{u_2^{c,a}(t_n)-u_2^{u,a}(t_k)}{u_2^{c,a}(t_k)-u_2^{u,a}(t_k)}}{1+\sum\limits_{k=1}^{n-1}\frac{u_2^{c,a}(t_n)-u_2^{u,a}(t_n)}{u_2^{c,a}(t_k)-u_2^{u,a}(t_k)}} u_1^{u,a}(t_n)$$

Finally, by substituting $\tau^{**}(t_{n+1})$ and $U_1(s^*,\tau^*)$ with primitives, we get the following equivalent condition for a low-value target:

$$\frac{1-\sum\limits_{k=1}^{n}\frac{u_2^{u,a}(t_{n+1})-u_2^{u,a}(t_k)}{u_2^{c,a}(t_k)-u_2^{u,a}(t_k)}}{1+\sum\limits_{k=1}^{n}\frac{u_2^{c,a}(t_{n+1})-u_2^{u,a}(t_{n+1})}{u_2^{c,a}(t_k)-u_2^{u,a}(t_k)}} > \frac{\frac{1-\sum_{k=1}^{n-1}\frac{u_2^{u,a}(t_n)-u_2^{u,a}(t_k)}{u_2^{c,a}(t_k)-u_2^{u,a}(t_k)}}{1+\sum_{k=1}^{n-1}\frac{u_2^{c,a}(t_n)-u_2^{u,a}(t_n)}{u_2^{c,a}(t_k)-u_2^{u,a}(t_k)}} u_1^{c,a}(t_n) + \frac{\sum_{k=1}^{n-1}\frac{u_2^{c,a}(t_n)-u_2^{u,a}(t_k)}{u_2^{c,a}(t_k)-u_2^{u,a}(t_k)}}{1+\sum_{k=1}^{n-1}\frac{u_2^{c,a}(t_n)-u_2^{u,a}(t_n)}{u_2^{c,a}(t_k)-u_2^{u,a}(t_k)}} u_1^{u,a}(t_n) - u_1^{u,a}(t_{n+1})}{u_1^{c,a}(t_{n+1}) - u_1^{u,a}(t_{n+1})}$$

References

[1] Avenhaus, R., von Stengel, B., Zamir, S. (2002). Inspection games, ch. 51, p. 1947-1987 in Aumann, R.J. and Hart, S. eds., *Handbook of Game Theory with Economic Applications*, vol. 3, Elsevier.

[2] Conitzer, V., Sandholm, T. (2006). Computing the optimal strategy to commit to. In *ACM Conference on Economics and Computation*. 82-90.

[3] Guo, Q., Gan, J., Fang, F., Tran-Thanh, L., Tambe, M., An, B. (2018). On the Inducibility of Stackelberg Equilibrium for Security Games. *AAAI Conference on Artificial Intelligence*. Vol. 33 (2019). 2020–2028.

[4] Heifetz, A., Martin Meier, M., Schipper, B. C., Dynamic unawareness and rationalizable behavior. *Games and Economic Behavior*. 81. 50-68.

[5] Kar, D., Nguyen, T., Fang, F., Brown, M., Sinha, A., Tambe, M., Jiang, A. X. (2018). Trends and applications in stackelberg security games. *Handbook of Dynamic Game Theory*. 1223-1269. Springer.

[6] Kiekintveld, C., Jain, Ma., Tsai, J., Pita, J., Ordóñez, F., Tambe, M. (2009). Computing optimal randomized resource allocations for massive security games. *AMAAS*-2009 Conf. 2. 689-696.

[7] Leitmann, G.. (1978). On generalized Stackelberg strategies. *Journal of Optimization Theory and Applications*. 26. 637-643.

[8] Pita, J., Jain, M., Ordóñez, F., Portway, C., Tambe, M., Western, C., Paruchuri, P., Kraus, S. (2009). Using Game Theory for Los Angeles Airport Security. *AI Magazine*. 30. 43-57.

[9] Sinha, A., Fang, F., An, B., Kiekintveld, C., Tambe, M. (2018). Stackelberg Security Games: Looking Beyond a Decade of Success. Proceedings of the *Twenty-Seventh International Joint Conference on Artificial Intelligence*. 5494-5501.

[10] von Stackelberg, H. (1934). Marktform und Gleichgewicht. Vienna and Berlin: Springer Verlag.

[11] von Stengel, B., Zamir, S. (2004). Leadership with commitment to mixed strategies. Technical Report LSE-CDAM2004-01, *CDAM* Research Report.

[12] von Stengel, B., Zamir, S. (2010). Leadership games with convex strategy sets. *Games and Economic Behavior*. 69 (2). 446-457.