

Stackelberg vs. Nash in Security Games: Interchangeability, Equivalence, and Uniqueness*

Zhengyu Yin¹, Dmytro Korzhyk², Christopher Kiekintveld¹,
Vincent Conitzer², and Milind Tambe¹

¹University of Southern California, Los Angeles, CA 90089, USA

{zhengyu, kiekintv, tambe}@usc.edu

²Duke University, Durham, NC 27708, USA

{dima, conitzer}@cs.duke.edu

ABSTRACT

There has been significant recent interest in game theoretic approaches to security, with much of the recent research focused on utilizing the leader-follower Stackelberg game model; for example, these games are at the heart of major applications such as the ARMOR program deployed for security at the LAX airport since 2007 and the IRIS program in use by the US Federal Air Marshals (FAMS). The foundational assumption for using Stackelberg games is that security forces (leaders), acting first, commit to a randomized strategy; while their adversaries (followers) choose their best response *after* surveillance of this randomized strategy. Yet, in many situations, the followers may act without observation of the leader's strategy, essentially converting the game into a simultaneous-move game model. Previous work fails to address how a leader should compute her strategy given this fundamental uncertainty about the type of game faced.

Focusing on the complex games that are directly inspired by real-world security applications, the paper provides four contributions in the context of a general class of security games. First, exploiting the structure of these security games, the paper shows that the Nash equilibria in security games are interchangeable, thus alleviating the equilibrium selection problem. Second, resolving the leader's dilemma, it shows that under a natural restriction on security games, any Stackelberg strategy is also a Nash equilibrium strategy; and furthermore, the solution is unique in a class of real-world security games of which ARMOR is a key exemplar. Third, when faced with a follower that can attack multiple targets, many of these properties no longer hold. Fourth, our experimental results emphasize positive properties of games that do not fit our restrictions. Our contributions have major implications for the real-world applications.

Categories and Subject Descriptors

I.2.11 [Artificial Intelligence]: Distributed Artificial Intelligence

General Terms

Theory

*Zhengyu Yin and Dmytro Korzhyk are both first authors of this paper.

Cite as: Stackelberg vs. Nash in Security Games: Interchangeability, Equivalence, and Uniqueness, Zhengyu Yin, Dmytro Korzhyk, Christopher Kiekintveld, Vincent Conitzer, and Milind Tambe, *Proc. of 9th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2010)*, van der Hoek, Kaminka, Lespérance, Luck and Sen (eds.), May, 10–14, 2010, Toronto, Canada, pp. XXX-XXX.

Copyright © 2010, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

Keywords

Stackelberg Equilibrium, Nash Equilibrium, Security

1. INTRODUCTION

There has been significant recent research interest in game-theoretic approaches to security at airports, ports, transportation, shipping and other infrastructure [12, 3, 4, 7]. Much of this work has used a *Stackelberg* game framework to model interactions between the security forces and attackers. That is, the defender (i.e., the security forces) acts first by committing to a patrolling or inspection strategy, and the attacker chooses where to attack after observing the defender's choice. The typical solution concept applied to these games is Strong Stackelberg Equilibrium (SSE), which assumes that the defender will choose an optimal mixed (randomized) strategy based on the assumption that the attacker will observe this strategy and choose an optimal response. This leader-follower paradigm appears to fit many real-world security situations. Indeed, Stackelberg games are at the heart of two major decision-support applications: the ARMOR program in use at the Los Angeles International Airport since 2007 to randomize allocation of checkpoints and canine patrols [12], and the IRIS program in use by the US Federal Air Marshals to randomize assignments of air marshals to flights [13].

However, there are legitimate concerns about whether the Stackelberg model is appropriate in all cases. In some situations attackers may choose to act without acquiring costly information about the security strategy, especially if security measures are difficult to observe (e.g., undercover officers) and insiders are unavailable. In such cases, a simultaneous-move game model may be a better reflection of the real situation. The defender faces an unclear choice about which strategy to adopt: the recommendation of the Stackelberg model, or of the simultaneous-move model, or something else entirely? In general settings, the equilibrium strategy can in fact differ between these models. Consider the following game in normal form:

	c	d
a	2,1	4,0
b	1,0	3,1

Table 1: Example game where the Stackelberg Equilibrium is not a Nash Equilibrium

If the row player has the ability to commit, the SSE strategy is to play *a* with .5 and *b* with .5, so that the best response for the column player is to play *d*, which gives the row player an expected utility

of 3.5². On the other hand, if the players move simultaneously the only Nash Equilibrium (NE) of this game is for the row player to play a and the column player c . This can be seen by noticing that b is strictly dominated for the row player. Previous work has failed to resolve the defender's dilemma of which strategy to select when the attacker's observation capability is unclear.

We conduct theoretical and experimental analysis of the leader's dilemma, focusing on *security games* [7]. These are non-zero-sum games motivated by real-world security domains, and are at the heart of applications such as ARMOR and IRIS [7, 12, 13]. We make four primary contributions. First, we show that Nash equilibria are interchangeable in security games, avoiding equilibrium selection problems. Second, if the game satisfies the SSAS (Subsets of Schedules Are Schedules) property, the defender's set of SSE strategies is a subset of her NE strategies. In this case, the defender is always playing a best response by using an SSE regardless of whether the attacker observes the defender's strategy or not. Third, we provide counter-examples to this (partial) equivalence in two cases: (1) when the SSAS property does not hold for defender schedules, and (2) when the attacker can attack multiple targets simultaneously. In these cases, the defender's SSE strategy may not be part of any NE profile. Finally, our experimental tests show that the fraction of games where the SSE strategy played is not part of any NE profile is vanishingly small. However, when attackers can attack multiple targets a relatively large number of games have distinct SSE and NE strategies.

2. MOTIVATING DOMAINS

We study quite general classes of security games in this work, but with assumptions motivated by two real-world applications. The first is the ARMOR security system deployed at the Los Angeles International Airport (LAX) [12]. In this domain police are able to set up checkpoints on roads leading to particular terminals, and assign canine units (bomb-sniffing dogs) to patrol terminals. Police resources in this domain are homogeneous, and do not have significant scheduling constraints.

IRIS is a similar application deployed by the Federal Air Marshals Service (FAMS) [13]. Armed marshals are assigned to commercial flights to deter and defeat terrorist attacks. This domain has more complex constraints. In particular, marshals are assigned to tours of flights that return to the same destination, and the tours on which any given marshal is available to fly are limited by the marshal's current location and timing constraints. The types of scheduling and resource constraints we consider in this work are motivated by those necessary to represent this domain.

Additionally, there are many other potential security applications, e.g., the Los Angeles Port domain, where port police patrol docks to ensure the safety and security of all passenger, cargo, and vessel operations.

3. DEFINITIONS AND NOTATION

A security game [7] is a two-player game between a defender and an attacker. The attacker may choose to attack any target from the set $T = \{t_1, t_2, \dots, t_n\}$. The defender tries to prevent attacks by covering targets using resources from the set $R = \{r_1, r_2, \dots, r_K\}$. As shown in Figure 1, $U_d^c(t_i)$ is the defender's utility if t_i is attacked while t_i is covered by some defender resource. If t_i is not covered, the defender gets $U_d^u(t_i)$. The attacker's utility is denoted similarly by $U_a^c(t_i)$ and $U_a^u(t_i)$. We use $\Delta U_d(t_i) = U_d^c(t_i) -$

$U_d^u(t_i)$ to denote the difference between defender's covered and uncovered utilities. Similarly, $\Delta U_a(t_i) = U_a^c(t_i) - U_a^u(t_i)$. As a key property of security games, we assume $\Delta U_d(t_i) > 0$ and $\Delta U_a(t_i) > 0$. In words, adding resources to cover a target helps the defender and hurts the attacker.




Figure 1: Payoff structure of security games.

Motivated by FAMS and similar real-world domains, we introduce resource and scheduling constraints for the defender. Resources may be assigned to *schedules* covering multiple targets, $s \subseteq T$. For each resource r_i , there is a subset S_i of the schedules S that resource r_i can potentially cover. That is, r_i can cover any $s \in S_i$. In the FAMS domain, flights are targets and air marshals are resources. Schedules capture the idea that air marshals fly tours, and must return to a particular starting point. Heterogeneous resources can express additional timing and location constraints that limit the tours on which any particular marshal can be assigned to fly. An important subset of the FAMS domain can be modeled using fixed schedules of size 2 (i.e., a pair of departing and returning flights). The LAX domain is also a subclass of security games as defined here, with schedules of size 1 and homogeneous resources.

A security game described above can be represented as a normal form game, as follows. The attacker's pure strategy space \mathcal{A} is the set of targets. The attacker's mixed strategy $\mathbf{a} = \langle a_i \rangle$ is a vector where a_i represents the probability of attacking t_i . The defender's pure strategy is a feasible assignment of resources to schedules, i.e., $\langle s_i \rangle \in \prod_{i=1}^K S_i$. Since covering a target with one resource is essentially the same as covering it with any positive number of resources, the defender's pure strategy can also be represented by a coverage vector $\mathbf{d} = \langle d_i \rangle \in \{0, 1\}^n$ where d_i represents whether t_i is covered or not. For example, $\langle \{t_1, t_4\}, \{t_2\} \rangle$ can be a possible assignment, and the corresponding coverage vector is $\langle 1, 1, 0, 1 \rangle$. However, not all the coverage vectors are feasible due to resource and schedule constraints. We denote the set of feasible coverage vectors by $\mathcal{D} \subseteq \{0, 1\}^n$.

The defender's mixed strategy \mathbf{C} specifies the probabilities of playing each $\mathbf{d} \in \mathcal{D}$, where each individual probability is denoted by $C_{\mathbf{d}}$. Let $\mathbf{c} = \langle c_i \rangle$ be the vector of coverage probabilities corresponding to \mathbf{C} , where $c_i = \sum_{\mathbf{d} \in \mathcal{D}} d_i C_{\mathbf{d}}$ is the marginal probability of covering t_i . For example, suppose the defender has two coverage vectors: $\mathbf{d}_1 = \langle 1, 1, 0 \rangle$ and $\mathbf{d}_2 = \langle 0, 1, 1 \rangle$. Then $\mathbf{C} = \langle .5, .5 \rangle$ is one defender's mixed strategy, and the corresponding $\mathbf{c} = \langle .5, 1, .5 \rangle$. Denote the mapping from \mathbf{C} to \mathbf{c} by φ , so that $\mathbf{c} = \varphi(\mathbf{C})$.

If strategy profile $\langle \mathbf{C}, \mathbf{a} \rangle$ is played, the defender's utility is

$$U_d(\mathbf{C}, \mathbf{a}) = \sum_i^n a_i (c_i U_d^c(t_i) + (1 - c_i) U_d^u(t_i)),$$

while the attacker's utility is

$$U_a(\mathbf{C}, \mathbf{a}) = \sum_i^n a_i (c_i U_a^c(t_i) + (1 - c_i) U_a^u(t_i)).$$

²In these games it is assumed that if the follower is indifferent, he breaks the tie in the leader's favor (otherwise, the optimal solution is not well defined).

If the players move simultaneously, the standard solution concept is Nash equilibrium.

DEFINITION 1. A pair of strategies $\langle \mathbf{C}, \mathbf{a} \rangle$ forms a Nash Equilibrium (NE) if they satisfy the following:

1. The defender plays a best-response:
 $U_d(\mathbf{C}, \mathbf{a}) \geq U_d(\mathbf{C}', \mathbf{a}) \forall \mathbf{C}'.$
2. The attacker plays a best-response:
 $U_a(\mathbf{C}, \mathbf{a}) \geq U_a(\mathbf{C}, \mathbf{a}') \forall \mathbf{a}'.$

In our Stackelberg model, the defender chooses a mixed strategy first, and the attacker chooses a strategy after observing the defender's choice. The attacker's response function is $g(\mathbf{C}) : \mathbf{C} \rightarrow \mathbf{a}$. In this case, the standard solution concept is Strong Stackelberg Equilibrium [8, 16].

DEFINITION 2. A pair of strategies $\langle \mathbf{C}, g \rangle$ forms a Strong Stackelberg Equilibrium (SSE) if they satisfy the following:

1. The leader (defender) plays a best-response:
 $U_d(\mathbf{C}, g(\mathbf{C})) \geq U_d(\mathbf{C}', g(\mathbf{C}')),$ for all $\mathbf{C}'.$
2. The follower (attacker) plays a best-response:
 $U_a(\mathbf{C}, g(\mathbf{C})) \geq U_a(\mathbf{C}, g'(\mathbf{C})),$ for all $\mathbf{C}, g'.$
3. The follower breaks ties optimally for the leader:
 $U_d(\mathbf{C}, g(\mathbf{C})) \geq U_d(\mathbf{C}, \tau(\mathbf{C})),$ for all $\mathbf{C},$ where $\tau(\mathbf{C})$ is the set of follower best-responses to $\mathbf{C}.$

We denote the set of mixed strategies for the defender that are played in some Nash Equilibrium by Ω_{NE} , and the corresponding set for Strong Stackelberg Equilibrium by Ω_{SSE} .

4. EQUILIBRIA IN SECURITY GAMES

The challenge for us is to understand the fundamental relationships between the SSE and NE strategies in security games. A special case is zero-sum security games, where the defender's utility is the exact opposite of the attacker's utility. For finite two-person zero-sum games, it is known that the different game theoretic solution concepts of NE, minimax, maximin and SSE all give the same answer. In addition, Nash equilibrium strategies of zero-sum games have a very useful property in that they are *interchangeable*: an equilibrium strategy for one player can be paired with the other player's strategy from *any* equilibrium profile, and the result is an equilibrium, and the payoffs for both players remain the same.

Unfortunately, security games are not necessarily zero-sum (and are not zero-sum in deployed applications). Many properties of zero-sum games do not hold in security games. For instance, a minimax strategy in a security game may not be a maximin strategy. Consider the example in Table 2, in which there are 3 targets and one defender resource. The defender has three actions; each of defender's actions can only cover one target at a time, leaving the other targets uncovered. While all three targets are equally appealing to the attacker, the defender has varying utilities of capturing the attacker at different targets. For the defender, the unique minimax strategy, $\langle 1/3, 1/3, 1/3 \rangle$, is different from the unique maximin strategy, $\langle 6/11, 3/11, 2/11 \rangle$.

Strategically zero-sum games [10] are a natural and strict superset of zero-sum games for which most of the desirable properties of zero-sum games still hold. This is exactly the class of games for which no completely mixed Nash equilibrium can be improved upon. Moulin and Vial proved a game (A, B) is strategically zero-sum if and only if there exist $u > 0$ and $v > 0$ such that

	t_1		t_2		t_3	
	C	U	C	U	C	U
Def	1	0	2	0	3	0
Att	0	1	0	1	0	1

Table 2: Security game which is not strategically zero-sum

$uA + vB = U + V$, where U is a matrix with identical columns and V is a matrix with identical rows [10]. Unfortunately, security games are not even strategically zero-sum. The game in Table 2 is a counterexample, because otherwise there must exist $u, v > 0$ such that,

$$u \begin{pmatrix} 1 & 0 & 0 \\ 0 & 2 & 0 \\ 0 & 0 & 3 \end{pmatrix} + v \begin{pmatrix} 0 & 1 & 1 \\ 1 & 0 & 1 \\ 1 & 1 & 0 \end{pmatrix} \\ = \begin{pmatrix} a & a & a \\ b & b & b \\ c & c & c \end{pmatrix} + \begin{pmatrix} x & y & z \\ x & y & z \\ x & y & z \end{pmatrix}$$

From these equations, $a + y = a + z = b + x = b + z = c + x = c + y = v$, which implies $x = y = z$ and $a = b = c$. We also know $a + x = u$, $b + y = 2u$, $c + z = 3u$. However since $a + x = b + y = c + z$, u must be 0, which contradicts the assumption $u > 0$.

Nevertheless, we show in the rest of this section that security games still have some important properties. We start by establishing equivalence between the set of defender's minimax strategies and the set of defender's NE strategies. Second, we show Nash equilibria in security games are interchangeable, resolving the defender's equilibrium strategy selection problem in simultaneous-move games. Third, we show that under a natural restriction on schedules, any SSE strategy for the defender is also a minimax strategy and hence an NE strategy. This resolves the defender's dilemma about whether to play according to SSE or NE when there is uncertainty about attacker's ability to observe the strategy. Finally, for a restricted class of games (including the games from the LAX domain), we find that there is a unique SSE/NE defender strategy and a unique attacker NE strategy.

4.1 Equivalence of NE and Minimax

We first prove that any defender's NE strategy is also a minimax strategy. Then for every defender's minimax strategy \mathbf{C} we construct a strategy \mathbf{a} for the attacker such that $\langle \mathbf{C}, \mathbf{a} \rangle$ is an NE profile.

DEFINITION 3. For a defender's mixed strategy \mathbf{C} , define the attacker's best response utility by $E(\mathbf{C}) = \max_{i=1}^n U_a(\mathbf{C}, t_i)$. Denote the minimum of the attacker's best response utilities over all defender's strategies by $E^* = \min_{\mathbf{C}} E(\mathbf{C})$. The set of defender's minimax strategies is defined as:

$$\Omega_M = \{\mathbf{C} | E(\mathbf{C}) = E^*\}.$$

We define the function f as follows. If \mathbf{a} is an attacker's strategy in which target t_i is attacked with probability a_i , then $f(\mathbf{a}) = \bar{\mathbf{a}}$ is an attacker's strategy such that

$$\bar{a}_i = \lambda a_i \frac{\Delta U_d(t_i)}{\Delta U_a(t_i)}$$

where $\lambda > 0$ is a normalizing constant such that $\sum_i^n \bar{a}_i = 1$. The inverse function $f^{-1}(\bar{\mathbf{a}}) = \mathbf{a}$ is given by the following equation.

$$a_i = \frac{1}{\lambda} \bar{a}_i \frac{\Delta U_a(t_i)}{\Delta U_d(t_i)} \quad (1)$$

LEMMA 4.1. Consider a security game \mathcal{G} . Construct the corresponding zero-sum security game $\bar{\mathcal{G}}$ in which the defender's utilities are re-defined as follows.

$$\begin{aligned} U_d^c(t) &= -U_a^c(t) \\ U_d^u(t) &= -U_a^u(t) \end{aligned}$$

Then $\langle \mathbf{C}, \mathbf{a} \rangle$ is an NE profile in \mathcal{G} if and only if $\langle \mathbf{C}, f(\mathbf{a}) \rangle$ is an NE profile in $\bar{\mathcal{G}}$.

PROOF. Note that the supports of strategies \mathbf{a} and $\bar{\mathbf{a}}$ are the same, and also that the attacker's utility function is the same in games \mathcal{G} and $\bar{\mathcal{G}}$. Thus \mathbf{a} is a best response to \mathbf{C} in \mathcal{G} if and only if $\bar{\mathbf{a}}$ is a best response to \mathbf{C} in $\bar{\mathcal{G}}$.

Denote the utility that the defender gets if profile $\langle \mathbf{C}, \mathbf{a} \rangle$ is played in game \mathcal{G} by $U_d^{\mathcal{G}}(\mathbf{C}, \mathbf{a})$. To show that \mathbf{C} is a best response to \mathbf{a} in game \mathcal{G} if and only if \mathbf{C} is a best response to $\bar{\mathbf{a}}$ in $\bar{\mathcal{G}}$, it is sufficient to show equivalence of the following two inequalities.

$$\begin{aligned} U_d^{\mathcal{G}}(\mathbf{C}, \mathbf{a}) - U_d^{\mathcal{G}}(\mathbf{C}', \mathbf{a}) &\geq 0 \\ \Leftrightarrow U_d^{\bar{\mathcal{G}}}(\mathbf{C}, \bar{\mathbf{a}}) - U_d^{\bar{\mathcal{G}}}(\mathbf{C}', \bar{\mathbf{a}}) &\geq 0 \end{aligned}$$

We will prove the equivalence by starting from the first inequality and transforming it into the second one. On the one hand, we have,

$$U_d^{\mathcal{G}}(\mathbf{C}, \mathbf{a}) - U_d^{\mathcal{G}}(\mathbf{C}', \mathbf{a}) = \sum_i^n a_i(c_i - c'_i)\Delta U_d(t_i).$$

Similarly, on the other hand, we have,

$$U_d^{\bar{\mathcal{G}}}(\mathbf{C}, \bar{\mathbf{a}}) - U_d^{\bar{\mathcal{G}}}(\mathbf{C}', \bar{\mathbf{a}}) = \sum_i^n \bar{a}_i(c_i - c'_i)\Delta U_a(t_i).$$

Given Equation (1) and $\lambda > 0$, we have,

$$\begin{aligned} U_d^{\mathcal{G}}(\mathbf{C}, \mathbf{a}) - U_d^{\mathcal{G}}(\mathbf{C}', \mathbf{a}) &\geq 0 \\ \Leftrightarrow \sum_i^n a_i(c_i - c'_i)\Delta U_d(t_i) &\geq 0 \\ \Leftrightarrow \sum_i^n \frac{1}{\lambda} \bar{a}_i \frac{\Delta U_a(t_i)}{\Delta U_d(t_i)} (c_i - c'_i) \Delta U_d(t_i) &\geq 0 \\ \Leftrightarrow \frac{1}{\lambda} \sum_i^n \bar{a}_i (c_i - c'_i) \Delta U_a(t_i) &\geq 0 \\ \Leftrightarrow \frac{1}{\lambda} \left(U_d^{\bar{\mathcal{G}}}(\mathbf{C}, \bar{\mathbf{a}}) - U_d^{\bar{\mathcal{G}}}(\mathbf{C}', \bar{\mathbf{a}}) \right) &\geq 0 \\ \Leftrightarrow U_d^{\bar{\mathcal{G}}}(\mathbf{C}, \bar{\mathbf{a}}) - U_d^{\bar{\mathcal{G}}}(\mathbf{C}', \bar{\mathbf{a}}) &\geq 0 \end{aligned}$$

□

LEMMA 4.2. Suppose \mathbf{C} is a defender NE strategy in a security game. Then $E(\mathbf{C}) = E^*$, i.e., $\Omega_{NE} \subseteq \Omega_M$.

PROOF. Suppose $\langle \mathbf{C}, \mathbf{a} \rangle$ is an NE profile in the security game \mathcal{G} . According to Lemma 4.1, $\langle \mathbf{C}, f(\mathbf{a}) \rangle$ must be an NE profile in the corresponding zero-sum security game $\bar{\mathcal{G}}$. Since \mathbf{C} is an NE strategy in a zero-sum game, it must also be a minimax strategy [5]. Thus $E(\mathbf{C}) = E^*$. □

LEMMA 4.3. In a security game \mathcal{G} , any defender's strategy \mathbf{C} such that $E(\mathbf{C}) = E^*$ is an NE strategy, i.e., $\Omega_M \subseteq \Omega_{NE}$.

PROOF. \mathbf{C} is a minimax strategy in both \mathcal{G} and the corresponding zero-sum game $\bar{\mathcal{G}}$. Any minimax strategy is also an NE strategy in a zero-sum game [5]. Then there must exist an NE profile $\langle \mathbf{C}, \bar{\mathbf{a}} \rangle$ in $\bar{\mathcal{G}}$. By Lemma 4.1, $\langle \mathbf{C}, f^{-1}(\bar{\mathbf{a}}) \rangle$ is an NE profile in \mathcal{G} . Thus \mathbf{C} is an NE strategy in \mathcal{G} . □

THEOREM 4.4. In a security game, the set of defender's minimax strategies is equal to the set of defender's NE strategies, i.e., $\Omega_M = \Omega_{NE}$.

PROOF. Lemma 4.2 shows that every defender's NE strategy is a minimax strategy, and Lemma 4.3 shows that every defender's minimax strategy is an NE strategy. Thus the sets of defender's NE and minimax strategies must be equal. □

4.2 Interchangeability of Nash Equilibria

We now show that Nash Equilibria in security games are interchangeable.

THEOREM 4.5. Suppose $\langle \mathbf{C}, \mathbf{a} \rangle$ and $\langle \mathbf{C}', \mathbf{a}' \rangle$ are two NE profiles in a security game \mathcal{G} . Then $\langle \mathbf{C}, \mathbf{a}' \rangle$ and $\langle \mathbf{C}', \mathbf{a} \rangle$ are also NE profiles in \mathcal{G} .

PROOF. Consider the corresponding zero-sum game $\bar{\mathcal{G}}$. From Lemma 4.1, both $\langle \mathbf{C}, f(\mathbf{a}) \rangle$ and $\langle \mathbf{C}', f(\mathbf{a}') \rangle$ must be NE profiles in $\bar{\mathcal{G}}$. By the interchange property of NE in zero-sum games [5], $\langle \mathbf{C}, f(\mathbf{a}') \rangle$ and $\langle \mathbf{C}', f(\mathbf{a}) \rangle$ must also be NE profiles in $\bar{\mathcal{G}}$. Applying Lemma 4.1 again in the other direction, we get that $\langle \mathbf{C}, \mathbf{a}' \rangle$ and $\langle \mathbf{C}', \mathbf{a} \rangle$ must be NE profiles in \mathcal{G} . □

By Theorem 4.5, the defender's equilibrium selection problem in a simultaneous-move security game is resolved. The reason is that given the attacker's NE strategy \mathbf{a} , the defender must get the same utility by responding with any NE strategy. Next, we give some insights on expected utilities in NE profiles. We first show the attacker's expected utility is the same in all NE profiles, followed by an example demonstrating that the defender may have varying expected utilities corresponding to different attacker's strategies.

THEOREM 4.6. Suppose $\langle \mathbf{C}, \mathbf{a} \rangle$ is an NE profile in a security game. Then, $U_a(\mathbf{C}, \mathbf{a}) = E^*$.

PROOF. From Lemma 4.2, \mathbf{C} is a minimax strategy and $E(\mathbf{C}) = E^*$. On the one hand,

$$U_a(\mathbf{C}, \mathbf{a}) = \sum_i^n a_i U_a(\mathbf{C}, t_i) \leq \sum_i^n a_i E(\mathbf{C}) = E^*.$$

On the other hand, because \mathbf{a} is a best response to \mathbf{C} , it should be at least as good as the strategy of attacking $t^* \in \arg \max_t U_a(\mathbf{C}, t)$ with probability 1, that is,

$$U_a(\mathbf{C}, \mathbf{a}) \geq U_a(\mathbf{C}, t^*) = E(\mathbf{C}) = E^*.$$

Therefore we know $U_a(\mathbf{C}, \mathbf{a}) = E^*$. □

Unlike the attacker who gets the same utility in all NE profiles, the defender may get varying expected utilities depending on the attacker's strategy selection. Consider the game shown in Table 3. The defender can choose to cover one of the two targets at a time. The only defender's NE strategy is to cover t_1 with 100% probability, making the attacker indifferent between attacking t_1 and t_2 . One attacker's NE response is always attacking t_1 , which gives the defender an expected utility of 1. Another attacker's NE strategy is $\langle 2/3, 1/3 \rangle$, given which the defender is indifferent between defending t_1 and t_2 . In this case, the defender's utility decreases to $2/3$ because she captures the attacker with a lower probability.

4.3 SSE and Minimax / NE

We have already shown that the set of defender's NE strategies coincides with her minimax strategies. If every defender's SSE strategy is also a minimax strategy, then SSE strategies must also be NE strategies. The defender can then safely commit to an SSE

	t_1		t_2	
	C	U	C	U
Def	1	0	2	0
Att	1	2	0	1

Table 3: A security game where the defender's expected utility varies in different NE profiles

strategy; there is no selection problem for the defender. Unfortunately, if a security game has arbitrary scheduling constraints, then an SSE strategy may not be part of any NE profile. For example, consider the game in Table 4 with 4 targets $\{t_1, \dots, t_4\}$, 2 schedules $s_1 = \{t_1, t_2\}$, $s_2 = \{t_3, t_4\}$, and a single defender resource. The defender always prefers that t_1 is attacked, and t_3 and t_4 are never appealing to the attacker.

	t_1		t_2		t_3		t_4	
	C	U	C	U	C	U	C	U
Def	10	9	-2	-3	1	0	1	0
Att	2	5	3	4	0	1	0	1

Table 4: A schedule-constrained security game where the defender's SSE strategy is not an NE strategy.

There is a unique SSE strategy for the defender, which places as much coverage probability on s_1 as possible without making t_2 more appealing to the attacker than t_1 . The rest of the coverage probability is placed on s_2 . The result is that s_1 and s_2 are both covered with probability 0.5. In contrast, in a simultaneous-move game, t_3 and t_4 are dominated for the attacker. Thus, there is no reason for the defender to place resources on targets that are never attacked, so the defender's unique NE strategy covers s_1 with probability 1. That is, the defender's SSE strategy is different from the NE strategy. The difference between the defender's payoffs in these cases can also be arbitrarily large because t_1 is always attacked in an SSE and t_2 is always attacked in a NE.

The above example restricts the defender to protect t_1 and t_2 together, which makes it impossible for the defender to put more coverage on t_2 without making t_1 less appealing. If the defender could assign resources to any subset of a schedule, this difficulty is resolved. More formally, we assume that for any resource r_i , any subset of a schedule in S_i is also a possible schedule in S_i :

$$\forall 1 \leq i \leq K : s' \subseteq s \in S_i \Rightarrow s' \in S_i. \quad (2)$$

If a security game satisfies Equation (2), we say it has the SSAS property. This is natural in many security domains, since it is often possible to cover fewer targets than the maximum number that a resource could possible cover in a schedule. We find that this property is sufficient to ensure that the defender's SSE strategy must also be an NE strategy.

LEMMA 4.7. *Suppose \mathbf{C} is a defender strategy in a security game which satisfies the SSAS property and $\mathbf{c} = \varphi(\mathbf{C})$ is the corresponding vector of marginal probabilities. Then for any \mathbf{c}' such that $0 \leq c'_i \leq c_i$ for all $t_i \in T$, there must exist a defender strategy \mathbf{C}' such that $\varphi(\mathbf{C}') = \mathbf{c}'$.*

PROOF. The proof is by induction on the number of t_i where $c'_i \neq c_i$, as denoted by $\delta(\mathbf{c}, \mathbf{c}')$. As the base case, if there is no i such that $c'_i \neq c_i$, the existence trivially holds because $\varphi(\mathbf{C}) = \mathbf{c}'$. Suppose the existence holds for all \mathbf{c}, \mathbf{c}' such that $\delta(\mathbf{c}, \mathbf{c}') = k$, where $0 \leq k \leq n - 1$. We consider any \mathbf{c}, \mathbf{c}' where $\delta(\mathbf{c}, \mathbf{c}') =$

$k + 1$. Then for some j , $c'_j \neq c_j$. Since $c'_j \geq 0$ and $c'_j < c_j$, we have $c_j > 0$. There must be a nonempty set of coverage vectors \mathcal{D}_j that cover t_j and receive positive probability in \mathbf{C} . Because the security game satisfies the SSAS property, for every $\mathbf{d} \in \mathcal{D}_j$, there is a valid \mathbf{d}^- which covers all targets in \mathbf{d} except for t_j . From the defender strategy \mathbf{C} , by shifting $\frac{C_{\mathbf{d}}(c_j - c'_j)}{c_j}$ probability from every $\mathbf{d} \in \mathcal{D}_j$ to the corresponding \mathbf{d}^- , we get a defender strategy \mathbf{C}^\dagger where $c_i^\dagger = c_i$ for $i \neq j$, and $c_j^\dagger = c'_j$ for $i = j$. Hence $\delta(\mathbf{c}^\dagger, \mathbf{c}') = k$, implying there exists a \mathbf{C}' such that $\varphi(\mathbf{C}') = \mathbf{c}'$ by the induction assumption. By induction, the existence holds for any \mathbf{c}, \mathbf{c}' . \square

THEOREM 4.8. *Suppose \mathbf{C} is a defender SSE strategy in a security game which satisfies the SSAS property. Then $E(\mathbf{C}) = E^*$, i.e., $\Omega_{SSE} \subseteq \Omega_M = \Omega_{NE}$.*

PROOF. The proof is by contradiction. Suppose $\langle \mathbf{C}, g \rangle$ is an SSE profile in a security game which satisfies the SSAS property, and $E(\mathbf{C}) > E^*$. Let $T_a = \{t_i | U_a(\mathbf{C}, t_i) = E(\mathbf{C})\}$ be the set of targets that give the attacker the maximum utility given the defender strategy \mathbf{C} . By the definition of SSE, we have

$$U_d(\mathbf{C}, g(\mathbf{C})) = \max_{t_i \in T_a} U_d(\mathbf{C}, t_i).$$

Consider a defender mixed strategy \mathbf{C}^* such that $E(\mathbf{C}^*) = E^*$. Then for any $t_i \in T_a$, $U_a(\mathbf{C}^*, t_i) \leq E^*$. Consider a vector \mathbf{c}' :

$$c'_i = \begin{cases} c_i^* - \frac{E^* - U_a(\mathbf{C}^*, t_i) + \epsilon}{U_a^u(t_i) - U_a^l(t_i)}, & t_i \in T_a, \\ c_i^*, & t_i \notin T_a, \end{cases} \quad (3a)$$

where ϵ is an infinitesimal positive number. Since $E^* - U_a(\mathbf{C}^*, t_i) + \epsilon > 0$, we have $c'_i < c_i^*$ for all $t_i \in T_a$. On the other hand, since for all $t_i \in T_a$,

$$U_a(\mathbf{c}', t_i) = E^* + \epsilon < E(\mathbf{C}) = U_a(\mathbf{C}, t_i),$$

we have $c'_i > c_i \geq 0$. Then for any $t_i \in T$, we have $0 \leq c'_i \leq c_i^*$. From Lemma 4.7, there exists a defender strategy \mathbf{C}' corresponding to \mathbf{c}' . The attacker's utility of attacking each target is as follows:

$$U_a(\mathbf{C}', t_i) = \begin{cases} E^* + \epsilon, & t_i \in T_a, \\ U_a(\mathbf{C}^*, t_i) \leq E^*, & t_i \notin T_a. \end{cases} \quad (4a)$$

Thus, the attacker's best responses to \mathbf{C}' are still T_a . For all $t_i \in T_a$, since $c'_i > c_i$, it must be the case that $U_d(\mathbf{C}, t_i) < U_d(\mathbf{C}', t_i)$. By definition of attacker's SSE response g , we have,

$$\begin{aligned} U_d(\mathbf{C}', g(\mathbf{C}')) &= \max_{t_i \in T_a} U_d(\mathbf{C}', t_i) \\ &> \max_{t_i \in T_a} U_d(\mathbf{C}, t_i) = U_d(\mathbf{C}, g(\mathbf{C})). \end{aligned}$$

It follows that the defender is better off using \mathbf{C}' , which contradicts the assumption \mathbf{C} is an SSE strategy of the defender. \square

Theorem 4.4 and 4.8 together imply the following corollary.

COROLLARY 4.9. *In security games with the SSAS property, any defender's SSE strategy is also an NE strategy.*

We can now answer the original question posed in this paper: when there is uncertainty over the type of game played, should the defender choose an SSE strategy or a mixed strategy Nash equilibrium or some combination of the two? For domains that satisfy the SSAS property, we have proven that any of the defender's SSE strategies is also an NE strategy.

Among our motivating domains, the LAX domain satisfies the SSAS property since all schedules are of size 1. Other patrolling domains, such as patrolling a port, also satisfy the SSAS property. In such domains, the defender could thus commit to an SSE strategy, which is also now known to be an NE strategy. The defender retains the ability to commit, but is still playing a best-response to an attacker in a simultaneous-move setting (assuming the attacker plays an equilibrium strategy – it does not matter which one, due to the interchange property shown above). However, the FAMS domain does not naturally satisfy the SSAS property because marshals must fly complete tours (though in principle they could fly as civilians on some legs of a tour). The question of selecting SSE vs. NE strategies in this case is addressed experimentally in Section 6.

4.4 Uniqueness in Restricted Games

The previous sections show that SSE strategies are NE strategies in many cases. However, there may still be multiple equilibria to select from (though this difficulty is alleviated by the interchange property). Here we prove an even stronger uniqueness result for an important restricted class of security domains, which includes the LAX domain. In particular, we consider security games where the defender has homogeneous resources that can cover any single target. The SSAS property is trivially satisfied, since all schedules are of size 1. Any vector of coverage probabilities $\mathbf{c} = \langle c_i \rangle$ such that $\sum_i^n c_i \leq K$ is a feasible strategy for the defender, so we can represent the defender strategy by marginal coverage probabilities. With a minor restriction on the attacker's payoff matrix, the defender always has a unique minimax strategy which is also the unique SSE and NE strategy. Furthermore, the attacker also has a unique NE response to this strategy.

THEOREM 4.10. *In a security game with homogeneous resources that can cover any single target, if for every target $t_i \in T$, $U_a^c(t_i) \neq E^*$, then the defender has a unique minimax, NE, and SSE strategy.*

PROOF. We first show the defender has a unique minimax strategy. Let $T^* = \{t | U_a^u(t) \geq E^*\}$. Define $\mathbf{c}^* = \langle c_i^* \rangle$ as

$$c_i^* = \begin{cases} \frac{U_a^u(t_i) - E^*}{U_a^u(t_i) - U_a^c(t_i)}, & t_i \in T^*, \\ 0, & t_i \notin T^*. \end{cases} \quad (5a)$$

$$(5b)$$

Note that E^* cannot be less than any $U_a^c(t_i)$ – otherwise, regardless of the defender's strategy, the attacker could always get at least $U_a^c(t_i) > E^*$ by attacking t_i , which contradicts the fact that E^* is the attacker's best response utility to a defender's minimax strategy. Since $E^* \geq U_a^c(t_i)$ and we assume $E^* \neq U_a^c(t_i)$,

$$1 - c_i^* = \frac{E^* - U_a^c(t_i)}{U_a^u(t_i) - U_a^c(t_i)} > 0 \Rightarrow c_i^* < 1.$$

Next, we will prove $\sum_i^n c_i^* \geq K$. For the sake of contradiction, suppose $\sum_i^n c_i^* < K$. Let $\mathbf{c}' = \langle c'_i \rangle$, where $c'_i = c_i^* + \epsilon$. Since $c_i^* < 1$ and $\sum_i^n c_i^* < K$, we can find $\epsilon > 0$ such that $c'_i < 1$ and $\sum_i^n c'_i < K$. Then every target has strictly higher coverage in \mathbf{c}' than in \mathbf{c}^* , hence $E(\mathbf{c}') < E(\mathbf{c}^*) = E^*$, which contradicts the fact that E^* is the minimum of all $E(\mathbf{c})$.

Next, we show that if \mathbf{c} is a minimax strategy, then $\mathbf{c} = \mathbf{c}^*$. By the definition of a minimax strategy, $E(\mathbf{c}) = E^*$. Hence, $U_a(\mathbf{c}, t_i) \leq E^* \Rightarrow c_i \geq c_i^*$. On the one hand $\sum_i^n c_i \leq K$ and on the other hand $\sum_i^n c_i \geq \sum_i^n c_i^* \geq K$. Therefore it must be the case that $c_i = c_i^*$ for any i . Hence, \mathbf{c}^* is the unique minimax strategy of the defender.

Furthermore, by Theorem 4.4, we have that \mathbf{c}^* is the unique defender's NE strategy. By Theorem 4.8 and the existence of SSE [2], we have that \mathbf{c}^* is the unique defender's SSE strategy. \square

THEOREM 4.11. *In a security game with homogeneous resources that can cover any one target, if for every target $t_i \in T$, $U_a^c(t_i) \neq E^*$ and $U_a^u(t_i) \neq E^*$, then the attacker has a unique NE strategy.*

PROOF. \mathbf{c}^* and T^* are the same as in the proof of Theorem 4.10. Given the defender's unique NE strategy \mathbf{c}^* , in any attacker's best response, only $t_i \in T^*$ can be attacked with positive probability, because,

$$U_a(\mathbf{c}^*, t_i) = \begin{cases} E^* & t_i \in T^* \\ U_a^u(t_i) < E^* & t_i \notin T^* \end{cases} \quad (6a)$$

$$(6b)$$

Suppose $\langle \mathbf{c}^*, \mathbf{a} \rangle$ forms an NE profile. We have

$$\sum_{t_i \in T^*} a_i = 1 \quad (7)$$

For any $t_i \in T^*$, we know from the proof of Theorem 4.10 that $c_i^* < 1$. In addition, because $U_a^u(t) \neq E^*$, we have $c_i^* \neq 0$. Thus we have $0 < c_i^* < 1$ for any $t_i \in T^*$. For any $t_i, t_j \in T^*$, necessarily $a_i \Delta U_d(t_i) = a_j \Delta U_d(t_j)$. Otherwise, assume $a_i \Delta U_d(t_i) > a_j \Delta U_d(t_j)$. Consider another defender's strategy \mathbf{c}' where $c'_i = c_i^* + \epsilon < 1$, $c'_j = c_j^* - \epsilon > 0$, and $c'_k = c_k^*$ for any $k \neq i, j$.

$$U_d(\mathbf{c}', \mathbf{a}) - U_d(\mathbf{c}^*, \mathbf{a}) = a_i \epsilon \Delta U_d(t_i) - a_j \epsilon \Delta U_d(t_j) > 0$$

Hence, \mathbf{c}^* is not a best response to \mathbf{a} , which contradicts the assumption that $\langle \mathbf{c}^*, \mathbf{a} \rangle$ is an NE profile. Therefore, there exists $\beta > 0$ such that, for any $t_i \in T^*$, $a_i \Delta U_d(t_i) = \beta$. Substituting a_i with $\beta / \Delta U_d(t_i)$ in Equation (7), we have

$$\beta = \frac{1}{\sum_{t_i \in T^*} \frac{1}{\Delta U_d(t_i)}}$$

Then we can explicitly write down \mathbf{a} as

$$a_i = \begin{cases} \frac{\beta}{\Delta U_d(t_i)}, & t_i \in T^*, \\ 0, & t_i \notin T^*. \end{cases} \quad (8a)$$

$$(8b)$$

As we can see, \mathbf{a} defined by (8a) and (8b) is the unique attacker NE strategy. \square

The implication of Theorem 4.10 and Theorem 4.11 is that in the simultaneous-move game, both the defender and the attacker have a unique NE strategy, which gives each player a unique expected utility as a result.

5. MULTIPLE ATTACKER RESOURCES

To this point we have assumed that the attacker will attack exactly one target. We now extend our security game definition to allow the attacker to use multiple resources to attack multiple targets simultaneously. To keep the model simple, we assume homogeneous resources (for both players) and schedules of size 1. The defender has $K < n$ resources which can be assigned to protect any target, and the attacker has $L < n$ resources which can be used to attack any target. Attacking the same target with multiple resources is equivalent to attacking with a single resource. The defender's pure strategy is a coverage vector $\mathbf{d} = \langle d_i \rangle \in \mathcal{D}$, where $d_i \in \{0, 1\}$ represents whether t_i is covered or not. Similarly, the attacker's pure strategy is an attack vector $\mathbf{q} = \langle q_i \rangle \in \mathcal{Q}$. We have $\sum_i^n d_i = K$ and $\sum_i^n q_i = L$. If pure strategies $\langle \mathbf{d}, \mathbf{q} \rangle$ are played, the attacker gets a utility of

$$U_a(\mathbf{d}, \mathbf{q}) = \sum_i^n q_i (d_i U_a^c(t_i) + (1 - d_i) U_a^u(t_i))$$

while the defender’s utility is given by

$$U_d(\mathbf{d}, \mathbf{q}) = \sum_i^n q_i (d_i U_d^c(t_i) + (1 - d_i) U_d^u(t_i))$$

The defender’s mixed strategy is a vector \mathbf{C} which specifies the probability of playing each $\mathbf{d} \in \mathcal{D}$. Similarly, the attacker’s mixed strategy \mathbf{A} is a vector of probabilities corresponding to all $\mathbf{q} \in \mathcal{Q}$.

In security games with multiple attacker resources, the defender’s SSE strategy may not be part of any NE profile, even if there are no scheduling constraints. Consider the game shown in Table 5.

	t_1		t_2		t_3	
	C	U	C	U	C	U
Def	0	-1	-100	-100 - ϵ	0	0 - ϵ
Att	100 - ϵ	100	0	10	5 - ϵ	5

Table 5: A security game with multiple attacker resources where the defender’s SSE strategy is not an NE strategy.

There are 3 targets t_1, t_2, t_3 . The defender has 1 resource, and the attacker has 2 resources. Therefore the defender’s pure strategy space is the set of targets to protect: $\{t_1, t_2, t_3\}$, while the attacker’s pure strategy space consists of the pairs of targets: $\{\langle t_1, t_2 \rangle, \langle t_1, t_3 \rangle, \langle t_2, t_3 \rangle\}$. If the defender protects t_1 and the attacker attacks $\langle t_1, t_2 \rangle$, the defender’s utility is $U_d^c(t_1) + U_d^u(t_2) = -100 - \epsilon$ and the attacker’s utility is $U_a^c(t_1) + U_a^u(t_2) = 110 - \epsilon$. In this example, t_1 is very appealing to the attacker no matter if it is covered or not, so t_1 is always attacked. If t_2 is attacked, the defender gets a very low utility, even if t_2 is defended. So in the SSE, the defender wants to make sure that t_2 is not attacked. The defender’s SSE strategy places at least .5 probability on t_2 , so that t_1 and t_3 are attacked instead of t_2 (recall that the attacker breaks ties in the defender’s favor in an SSE). The attacker’s SSE response is $\mathbf{A} = \langle 0, 1, 0 \rangle$, i.e., to always attack t_1 and t_3 . The other .5 defense probability will be placed on t_1 because $\Delta U_d(t_1) > \Delta U_d(t_3)$. So, the SSE profile is $\langle \mathbf{C}, \mathbf{A} \rangle$, where $\mathbf{C} = \langle .5, .5, 0 \rangle$.

Next, we show that there is no NE in which the defender plays \mathbf{C} . Suppose there is an NE profile $\langle \mathbf{C}, \mathbf{A}' \rangle$. Given \mathbf{C} , the attacker’s utility for attacking t_1 is higher than the utility for attacking t_2 , so it must be that t_1 is always attacked in this NE. Therefore, the attacker never plays $\langle t_2, t_3 \rangle$. However, this implies that t_1 is the most appealing target for the defender to cover, because $U_d(t_1, \mathbf{A}) > U_d(t_i, \mathbf{A}), i \in \{2, 3\}$. So, to be a best response the coverage of t_1 would need to be 1 instead of 0.5, contradicting the assumption that \mathbf{C} is an equilibrium strategy for the defender.

6. EXPERIMENTAL RESULTS

While our theoretical results resolve the leader’s dilemma for many interesting classes of security games, as we have seen, there are still some cases where SSE strategies are distinct from NE strategies for the defender. One case is when security games do not satisfy the SSAS property, and another is when the attacker has multiple resources. We conduct experiments to further investigate these two cases, offering evidence about the frequency with which SSE strategies differ from all NE strategies across randomly generated games using 36 different parameter settings.

For a particular game instance we first compute an SSE strategy \mathbf{C} using the DOBSS mixed-integer linear program [12]. We then use the linear feasibility program below to determine whether or not this SSE strategy is part of some NE profile by attempting to

find an appropriate attacker response strategy.

$$A_{\mathbf{q}} \in [0, 1], \text{ for all } \mathbf{q} \in \mathcal{Q} \quad (9)$$

$$\sum_{\mathbf{q} \in \mathcal{Q}} A_{\mathbf{q}} = 1 \quad (10)$$

$$A_{\mathbf{q}} = 0, \text{ for all } U_a(\mathbf{q}, \mathbf{C}) < E(\mathbf{C}) \quad (11)$$

$$\sum_{\mathbf{q} \in \mathcal{Q}} A_{\mathbf{q}} U_d(\mathbf{d}, \mathbf{q}) \leq Z, \text{ for all } \mathbf{d} \in \mathcal{D} \quad (12)$$

$$\sum_{\mathbf{q} \in \mathcal{Q}} A_{\mathbf{q}} U_d(\mathbf{d}, \mathbf{q}) = Z, \text{ for all } d \in D \text{ with } C_d > 0 \quad (13)$$

Here \mathcal{Q} is the set of attacker pure strategies, which is just the set of targets when there is only 1 attacker resource. The probability that the attacker plays \mathbf{q} is denoted by $A_{\mathbf{q}}$ which must be between 0 and 1 (Constraint (9)). Constraint (10) forces the probabilities to sum to 1. Constraint (11) prevents the attacker from placing positive probabilities on pure strategies which give the attacker a utility less than the best response utility $E(\mathbf{C})$. In constraints (12) and (13), Z is a variable which represents the maximum expected utility the defender can get among all pure strategies given the attacker’s strategy \mathbf{A} , and C_d denotes the probability of playing \mathbf{d} in \mathbf{C} . These two constraints require the defender’s strategy \mathbf{C} to be a best response to the attacker’s mixed strategy. Therefore a feasible solution \mathbf{A} is an NE strategy for the attacker. Conversely, if $\langle \mathbf{C}, \mathbf{A} \rangle$ is an NE profile, \mathbf{A} must satisfy all of the LP constraints.

We first test single-attacker games, fixing the number of targets at 10 and the number of defender resources at 3. We vary the number of schedules, the size of the schedules, and the number of resource types. Each test set consists of 10000 games, with payoffs drawn from $U[-100, 0]$ for $U_d^u(t_i)$ and $U_a^c(t_i)$, and $U[0, 100]$ for $U_d^c(t_i)$ and $U_a^u(t_i)$. Table 6 summarizes our results. A column represents a number of schedules, and a row represents a pair of schedule size and number of resource types. For example, looking at row 2 and column 2, we see that among 10000 games with 5 schedules of size 2 and 1 resource type, there are 316 cases where the defender’s SSE strategy is not an NE strategy. The number of cases where the defender’s SSE strategy is not an NE strategy is never more than 10.5% in any of the 36 settings we tested. This number decreases as we increase the number of schedules. With 20 available schedules, the number is less than 2%. The main implication of these results is that in practice, committing to an SSE strategy is likely to be a good approach in almost all cases. This is particularly true in domains like FAMS where schedule sizes are relatively small (2 in most cases) and the number of possible schedules is large relative to the number of targets.

	5	10	15	20
2S / 1R	316	103	27	3
2S / 2R	313	82	22	2
2S / 3R	297	101	18	3
3S / 1R	933	555	165	32
3S / 2R	858	494	172	35
3S / 3R	867	551	155	35
4S / 1R	990	912	515	183
4S / 2R	1029	950	492	190
4S / 3R	1005	927	483	173

Table 6: Number of instances out of 10000 single-attacker security games where SSE is not NE

Table 7 shows the results for varying numbers of attacker re-

sources. Again, each set has 10000 games. As the number of attacker resources increases, the number of cases where the defender's SSE strategy is not an NE strategy increases. With 2, 3, and 4 attacker resources, the numbers are 27%, 54%, and 69% respectively, which implies the defender cannot simply play an SSE strategy when there are multiple attacker resources. This result poses an interesting direction for future work, since it is unclear how a defender should play in these games if the attacker's ability to observe the mixed strategy is uncertain.

	2	3	4
#SSE \neq NE	2692	5368	6873

Table 7: Number of instances out of 10000 multiple-attacker security games where SSE is not NE

7. SUMMARY AND RELATED WORK

There has been significant interest in understanding the interaction of observability and commitment in general Stackelberg games. Bagwell's early work [1] questions the value of commitment to pure strategies given noisy observations by followers; but the ensuing and on-going debate illustrated that the leader retains her advantage in case of commitment to mixed strategies [14, 6]. The value of commitment for the leader when observations are costly is also studied in [9]. In contrast with this research, our work focuses on real-world security games, illustrating subset, equivalence, interchangeability, and uniqueness properties that are non-existent in general Stackelberg games studied previously.

Pita et al. [11] provide experimental results on observability in Stackelberg games: they test a variety of defender strategies against human players (attackers) who choose their optimal attack when provided with limited observations of defender strategy. Results show the superiority of a defender's strategy computed assuming human "anchoring bias" in attributing probability distribution over the defender's actions. This research complements ours, which provides new mathematical foundations. Testing the insights of our research with the experimental paradigm of [11], with expert players is an interesting topic for future research.

Going back to the foundations of game theory, Von Neumann and Morgenstern [15] provided a key result on interchangeability: for two-player zero-sum games, any combination of players' maximin strategies is in equilibrium. However, our security games are neither zero-sum nor strategically zero-sum (as seen earlier).

To summarize, this paper is focused on a general class of defender-attacker Stackelberg games that are directly inspired by real-world security applications. The paper confronts fundamental questions of how a defender should compute her mixed strategy. In this context, this paper provides four key contributions. First, exploiting the structure of these security games, the paper shows that the Nash equilibria in security games are interchangeable, thus alleviating the defender's equilibrium selection problem for simultaneous-move games. Second, resolving the defender's dilemma, it shows that under the SSAS restriction on security games, any Stackelberg strategy is also a Nash equilibrium strategy; and furthermore, this strategy is unique in a class of real-world security games of which ARMOR is a key exemplar. Third, when faced with a follower that can attack multiple targets, many of these properties no longer hold, providing a key direction for future research. Fourth, our experimental results emphasize positive properties of security games that do not fit the SSAS property. In practical terms, these contributions imply that defenders in applications such as ARMOR [12]

and IRIS [13] can simply commit to SSE strategies, thus helping to resolve a major dilemma in real-world security applications.

8. ACKNOWLEDGEMENTS

This research was supported by the United States Department of Homeland Security through the Center for Risk and Economic Analysis of Terrorism Events (CREATE). Korzhyk and Conitzer are supported by NSF IIS-0812113, ARO 56698-CI, and an Alfred P. Sloan Research Fellowship. However, any opinions, conclusions or recommendations herein are solely those of the authors and do not necessarily reflect views of the funding agencies. We thank Ronald Parr for detailed comments and discussions.

9. REFERENCES

- [1] K. Bagwell. Commitment and observability in games. *Games and Economic Behavior*, 8:271–280, 1995.
- [2] T. Basar and G. J. Olsder. *Dynamic Noncooperative Game Theory*. Academic Press, San Diego, CA, 2nd edition, 1995.
- [3] N. Basilico, N. Gatti, and F. Amigoni. Leader-follower strategies for robotic patrolling in environments with arbitrary topologies. In *AAMAS-09*, 2009.
- [4] V. Conitzer and T. Sandholm. Computing the optimal strategy to commit to. In *ACM EC-06*, pages 82–90, 2006.
- [5] D. Fudenberg and J. Tirole. *Game Theory*. MIT Press, October 1991.
- [6] S. Huck and W. Müller. Perfect versus imperfect observability—an experimental test of Bagwell's result. *Games and Economic Behavior*, 31(2):174 – 190, 2000.
- [7] C. Kiekintveld, M. Jain, J. Tsai, J. Pita, M. Tambe, and F. Ordonez. Computing optimal randomized resource allocations for massive security games. In *AAMAS-09*, 2009.
- [8] G. Leitmann. On generalized Stackelberg strategies. *Optimization Theory and Applications*, 26(4):637–643, 1978.
- [9] J. Morgan and F. Vardy. The value of commitment in contests and tournaments when observation is costly. *Games and Economic Behavior*, 60(2):326–338, August 2007.
- [10] H. Moulin and J. P. Vial. Strategically zero-sum games: The class of games whose completely mixed equilibria cannot be improved upon. *International Journal of Game Theory*, 7(3-4):201–221, September 1978.
- [11] J. Pita, M. Jain, F. Ordóñez, M. Tambe, S. Kraus, and R. Magori-Cohen. Effective solutions for real-world Stackelberg games: When agents must deal with human uncertainties. In *AAMAS-09*, 2009.
- [12] J. Pita, M. Jain, C. Western, C. Portway, M. Tambe, F. Ordonez, S. Kraus, and P. Parachuri. Deployed ARMOR protection: The application of a game-theoretic model for security at the Los Angeles International Airport. In *AAMAS-08 (Industry Track)*, 2008.
- [13] J. Tsai, S. Rathi, C. Kiekintveld, F. Ordóñez, and M. Tambe. IRIS a tool for strategic security allocation in transportation networks. In *AAMAS-09 (Industry Track)*, 2009.
- [14] E. van Damme and S. Hurkens. Games with imperfectly observable commitment. *Games and Economic Behavior*, 21(1-2):282 – 308, 1997.
- [15] J. von Neumann and O. Morgenstern. *Theory of Games and Economic Behavior*. Princeton University Press, May 2004.
- [16] B. von Stengel and S. Zamir. Leadership with commitment to mixed strategies. Technical Report LSE-CDAM-2004-01, CDAM Research Report, 2004.