

The Power of Initial Investigation in Audit Games

Ren Liu¹, Weiran Shen^{1, 2, 3*}

¹Gaoling School of Artificial Intelligence, Renmin University of China

²Beijing Key Laboratory of Research on Large Models and Intelligent Governance

³Engineering Research Center of Next-Generation Intelligent Search and Recommendation, MOE
{liuren201438, shenweiran}@ruc.edu.cn

Abstract

Audit games are an important variant of the Stackelberg security game, a widely studied game-theoretic model over the past years. It has been acknowledged that a pre-audit phase can notably enhance the audit's efficiency by informing and directing the following audit procedures.

In this paper, we model the above process with a two-stage audit game. The game encompasses two stages: an investigation stage where the auditor gathers information about potential policy breaches, and an audit stage where the auditor allocates the audit resources based on the investigation results. We formulate the problem as a set of mathematical programs. Due to the non-convexity of the programs, we consider a restricted strategy space and show that the optimal strategy in the restricted space can be determined by solving a polynomial number of convex optimization problems. Finally, we conduct extensive experiments to evaluate the effect of introducing the initial investigation stage and our algorithm. Our experiments show that even a small budget for the initial investigations can significantly enhance the defender's utility.

Introduction

As a standard theoretical tool in enhancing public security, the Stackelberg security game (SSG) has attracted much research attention in recent years. The applications of the SSG have spread across various domains that are crucial to global development and human well-being (Pita et al. 2008; Tsai et al. 2009; Jain et al. 2010). Besides urban crime prevention, wildlife protection, and cybersecurity, audit games are another important application of the SSG that draws inspiration from the audit process. Audit mechanisms are common in modern organizations (e.g., hospitals, tech companies, and financial institutions) to help them detect policy violations (e.g., unauthorized access to sensitive data).

A standard audit game features two players: an auditor and an auditee. Existing work on audit games mainly follows the framework of the standard SSG, where the auditor announces the audit strategy first and then the auditee responds by choosing a target to attack. Following the literature convention, we call the auditor the “defender” and the auditee the “attacker”.

*To whom correspondence should be addressed.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

However, such a framework may not be able to fully leverage the defender's leadership advantage as the audit resources are usually limited. To circumvent this limitation, we introduce an initial investigation stage for the defender, transforming their strategy into a two-stage process. The defender can make use of the first investigation stage to collect “clues” about the attacker's target choice, and then allocate audit resources to the targets in the subsequent stage based on the observed clues.

The initial stage of the investigation draws inspiration from the concept of a “pre-audit”. Pre-audits are a widely used and essential component in various auditing processes. Their primary purpose is to identify potential issues early on and to ensure that financial records and internal controls are in place (Hatfield, Jackson, and Vandervelde 2011; Loebbecke and Steinbart 1987). Pre-audits usually include a review of financial statements, examination of documentation, interviewing key employees, and identifying unusual adjustments in budget allocations, allowing auditors to gain a preliminary understanding of the auditee and thereby identify early issues.

The investigation stage has been confirmed by numerous researches to be helpful in guiding the audit direction (Furnham and Gunter 2015; Houck 2003; Singleton and Singleton 2010). Intuitively, the clues collected in the investigation stage reveal information about the attacker's choice, thereby enhancing the utilization efficiency of audit resources and improving the overall protection level.

In this paper, we propose a two-stage audit game. The defender's strategy has two components: the investigation strategy and the audit strategy. In the first stage, the defender allocates investigation resources to collect clues about the attacker's target. In the second stage, the defender allocates audit resources to each target to formally audit the targets. The defender's audit resource allocation is based on the clues observed in the investigation stage to take advantage of the information. Our model allows probabilistic observations of the clues, and takes into account both “false alarms” (false positives or type I errors) and “missed detections” (false negatives or type II errors).

Since each target may have its own clue observations, the total number of clue combinations can be exponential. This leads to exponentially many audit strategies, as the audit strategy depends on the observed clues in the first stage,

making the problem intractable. To address this problem, we focus on a simplified defense strategy, which we call the *independent defense strategy*. We analyze the structure of the optimal independent defense strategy and prove that the minimum budget consumption in the second stage decreases monotonically with respect to the minimum utility the attacker can obtain. Based on this property, we propose a fast algorithm for the optimal independent defense strategy using binary search. In each iteration, we only need to solve a second-order cone program, which is a convex optimization problem and can be efficiently solved.

We also conduct extensive experiments to evaluate the effectiveness of our new game framework. Our experiments demonstrate that the additional investigation stage can significantly enhance the defender’s utility with only small resource investments. Our results also provide a useful guideline for the defenders in general security games who are capable of performing simple initial investigations.

Related Work

Stackelberg security games and audit games. Stackelberg security games have been widely studied in computer science (Conitzer and Sandholm 2006; Kiekintveld et al. 2009). Since its inception, security games have found many practical applications in real life (Pita et al. 2008; Jain et al. 2010). To the best of our knowledge, Fellingham and Newman (1985) was the first to model the audit problem in a game-theoretic framework. On top of that, Blocki et al. (2013, 2015) applied the SSG model to auditing scenarios and proposed the concept of *audit games*. They introduced an additional penalty term to model the attacker’s punishment when the attack was caught by the defender. Different from their model, we add a pre-audit phase and consider the defender’s overall strategy in both the pre-audit and formal audit phases.

Multi-stage security games. Our work is also related to the literature on multi-stage security games. As solely relying on the SSG strategy is insufficient for the defender in many scenarios (Xu et al. 2015), some previous work adds additional stages to the defender’s strategy. One approach to exploit the game structures by adding stages to the defender’s strategy is through information design. The defender can send signals to the attacker, thereby inducing the attacker to change actions (e.g., (Bondi et al. 2018, 2020; Guo et al. 2017; Xu et al. 2015)).

Information asymmetry in security games. Enhancing the defender’s information collection channel is another approach that is technically relevant to our work. Recently, there is a line of work that focuses on acquiring information through informants, who will provide information about the attacker’s action (Huang et al. 2020; Shen et al. 2020, 2024). Similar ideas include introducing alarm systems, drones, and mobile sensors to monitor the attacker’s action (Bondi et al. 2018; Ma et al. 2018; Shi et al. 2020; Xu et al. 2018; Zhang et al. 2019), whereas we directly model the defender’s information-gathering behavior as part of their strategy.

Preliminaries

We consider an audit game setting with two players: the defender (the auditor) and the attacker (the auditee). The defender has only limited resources and wants to protect n different targets, while the attacker wants to attack them. We formulate the problem as a Stackelberg security game.

In the standard security game, the defender first commits to a strategy that allocates the resources to the targets in a randomized way. Knowing the defender’s strategy, the attacker then chooses a target to attack. Different from the standard version, we consider a setting where the defender performs an initial investigation (e.g., interviewing key employees or checking financial documents) to collect “clues” about potential policy violations. We model the audit process with two stages, where the two stages require different defensive resources. In the first stage, the defender allocates investigation resources to perform initial investigations. In the second stage, based on the results of the investigation, the defender uses audit resources to audit the targets.

Let $[n] = \{1, 2, \dots, n\}$ be the set of all targets. For any target $i \in [n]$, let c_i^1 be the cost of performing an initial investigation on the target in the first stage and c_i^2 the cost of auditing the target in the second stage. Denote the defender’s strategy by $s = (x, y)$, where $x, y \in [0, 1]^n$ are n -dimensional vectors with each element being the probability of investigating or auditing the corresponding target in stage 1 and 2, respectively.

Let $o \in \{0, 1\}^n$ be a binary vector, where o_i indicates whether a clue is found on target i in the investigation stage. Since the auditing strategy in stage 2 may depend on the investigation results in stage 1, the strategy y of stage 2 is actually a function of o .

Formally, we define the game as follows:

Definition 1 (Audit Game with Initial Investigations). *The audit game with initial investigations proceeds as follows:*

- The defender announces their strategy $s = (x, y)$;
- Based on the defender’s strategy s , the attacker chooses a target to attack;
- The defender allocates resources according to x to perform initial investigations and obtains observation o ;
- Based on the observation o , the defender uses strategy $y(o)$ to audit the targets.

If the defender audits a target i attacked by the attacker, the attacker gets a penalty P_i^a while the defender gets a reward R_i^d . Similarly, if the attacked target is not audited by the defender, the attacker gets a reward R_i^a , and the defender gets a penalty P_i^d . We make the standard assumption that both the defender’s and the attacker’s rewards are always larger than their corresponding penalties, i.e., $R_i^d > P_i^d$ and $R_i^a > P_i^a$ for all i .

An initial investigation may not accurately reveal whether an attack happens or not, so we view o_i as a random variable and assume that o_i and o_j are independent, no matter which target is attacked. We define the probability of getting a clue on an investigated target as follows:

$$p_i^a = \Pr\{o_i = 1 \mid \text{the investigated target } i \text{ is attacked}\},$$

$$p_i^u = \Pr\{o_i = 1 \mid \text{the investigated target } i \text{ is not attacked}\}.$$

Here, we assume that $p_i^a > p_i^u$ for all target i .

Suppose that target t is attacked. Clearly, for any target i , the result of the investigation in stage 1 gives an observation with $o_i = 1$ only if the defender allocates resources to investigate target i in stage 1 and gets a clue. Thus the probability of a positive observation for target i is:

$$\Pr\{o_i = 1 \mid t\} = \begin{cases} x_i p_i^a & \text{if } i = t \\ x_i p_i^u & \text{otherwise} \end{cases}.$$

The probability of having $o_i = 0$ is then $1 - \Pr\{o_i = 1 \mid t\}$, and the probability of vector o is:

$$\Pr\{o \mid t\} = \prod_{i=1}^n \Pr\{o_i \mid t\}.$$

Therefore, if target t is attacked by the attacker, the utilities of the defender and the attacker are:

$$\begin{aligned} u_d(s, t) &= \sum_o [y_t(o) R_t^d + (1 - y_t(o)) P_t^d] \prod_{i=1}^n \Pr\{o_i \mid t\}, \\ u_a(s, t) &= \sum_o [y_t(o) P_t^a + (1 - y_t(o)) R_t^a] \prod_{i=1}^n \Pr\{o_i \mid t\}. \end{aligned}$$

Similar to the standard SSG, the defender has only a limited amount of resources, which is usually much less than the total cost of auditing all the targets. Let B_1 and B_2 denote the total amount of investigation resources and audit resources, respectively. Thus, the defender's strategy (x, y) has to satisfy the following budget constraints:

$$\begin{aligned} \sum_{i=1}^n c_i^1 x_i &\leq B_1, \\ \sum_o \left(\prod_{i \in [n]} \Pr\{o_i \mid t\} \right) \sum_{i=1}^n c_i^2 y_i(o) &\leq B_2, \forall t \in [n]. \end{aligned}$$

The standard solution concept for a two-player Stackelberg game is strong Stackelberg equilibrium (SSE). In our model, we also adopt such a solution concept and assume that after observing the defender's strategy s , the attacker chooses a best target t , breaking ties in favor of the defender. I.e., given any defender's strategy s , the attacker's strategy $t(s)$ satisfies the following:

$$u_a(s, t(s)) \geq u_a(s, t'), \forall t' \in [n], \quad (1)$$

$$u_d(s, t(s)) \geq u_d(s, t'), \forall t' \in \arg \max_t u_a(s, t). \quad (2)$$

Definition 2 (Strong Stackelberg Equilibrium (SSE)). A strategy profile (s, t) is a strong Stackelberg equilibrium if it satisfies the following:

1. The attacker plays a best-response against the defender strategy s (Equation (1));
2. The attacker breaks ties in favor of the defender (Equation (2));
3. Strategy s is optimal for the defender:

$$u_d(s, t(s)) \geq u_d(s', t(s')).$$

To find the optimal strategy for the defender, we follow (Conitzer and Sandholm 2006) and solve the following mathematical program for each target t , and then choose the best solution among them:

$$\begin{aligned} \max_{x, y} \quad & u_d(s, t) \\ \text{s.t.} \quad & u_a(s, t) \geq u_a(s, t') \quad \forall t', \\ & \sum_{i=1}^n c_i^1 x_i \leq B_1, \\ & \sum_o \sum_{i=1}^n y_i(o) c_i^2 \prod_{i=1}^n \Pr\{o_i \mid t\} \leq B_2, \\ & 0 \leq x_i \leq 1 \quad \forall i, \\ & 0 \leq y_i(o) \leq 1 \quad \forall i, \forall o \end{aligned} \quad (3)$$

Theoretical Analysis

The programs presented above have both a non-convex objective function and non-convex constraints. Also, the second-stage solution $y(o)$ depends on the clues gathered in the first stage, which can have 2^n different possibilities. Therefore, even representing $y(o)$ takes exponential space. In this section, we analyze the problem and derive structural results about the solution to the problem. The following result shows that we can, without loss of generality, focus on solutions where y_i only depends on o_i and o_t , i.e., $y_i(o) = y_i(o_i, o_t)$.

Lemma 1. *There exists an optimal solution with y_i being a function of only o_i and o_t , for any $i \neq t$, and y_t a function of only o_t , where t is the chosen target of the attacker under the solution.*

Proof. We prove the statement by showing that, for any optimal solution $s^* = (x^*, y^*)$, we can construct a new solution that satisfies the condition in the statement of the lemma and achieves the same utility as s^* . Let t be the target chosen by the attacker under solution s^* . Construct another solution $s' = (x', y')$ as follows:

$$\begin{aligned} x'_i &= x_i^*, \quad \forall i \in \{n\}, \\ y'_i(o_i, o_t) &= \begin{cases} \mathbb{E}_{o_{-t}}[y_i^*(o_t, o_{-t})] & \text{if } i = t, \\ \mathbb{E}_{o_{-(i,t)}}[y_i^*(o_i, o_t, o_{-(i,t)})] & \text{otherwise.} \end{cases} \end{aligned}$$

Note that, as the defender's observation on each target is independent, no matter which target is attacked, we have

$$\mathbb{E}_{o_{-(i,t)}}[y_i^*(o_i, o_t, o_{-(i,t)})] = \mathbb{E}_{o_{-(i,t)}}[y_i^*(o_i, o_t)] \quad \forall i \neq t.$$

According to the construction of s' , it is straightforward that

$$0 \leq x'_i \leq 1, 0 \leq y'_i(o_i, o_t) \leq 1 \quad \forall i \in [n], \forall o_i, o_t.$$

And

$$\mathbb{E}_o[y_i^*(o)] = \mathbb{E}_{o_{i,t}}[y'_i(o_i, o_t)]. \quad (4)$$

With Equation (4), one can easily check that the new strategy s' achieves the same defender utility:

$$\begin{aligned}
& u_d(s^*, t) \\
&= \sum_o [y_t^*(o)R_t^d + (1 - y_t^*(o))P_t^d] \prod_{i=1}^n \Pr\{o_i \mid t\} \\
&= \mathbb{E}_o [y_t^*(o)R_t^d + (1 - y_t^*(o))P_t^d \mid t] \\
&= \mathbb{E}_o [P_t^d + y_t^*(o)(R_t^d - P_t^d) \mid t] \quad (5) \\
&= \mathbb{E}_o [P_t^d \mid t] + (R_t^d - P_t^d) \mathbb{E}_o [y_t^*(o_t) \mid t] \\
&= P_t^d + (R_t^d - P_t^d) \mathbb{E}_{o_t} [y_t'(o_t) \mid t] \\
&= u_d(s', t).
\end{aligned}$$

Similarly, the new strategy also gives the same utility to the attacker:

$$\begin{aligned}
u_a(s^*, i) &= R_i^a - (R_i^a - P_i^a) \mathbb{E}_{o_{i,t}} [y_i'(o_i, o_t) \mid i] \quad (6) \\
&= u_a(s', i) \quad \forall i,
\end{aligned}$$

Furthermore, the new strategy satisfies the budget constraint:

$$\begin{aligned}
\sum_{i=1}^n c_i^1 x_i' &= \sum_{i=1}^n c_i^1 x_i^* \leq B_1, \\
\sum_{i=1}^n c_i^2 \mathbb{E}_{o_{i,t}} [y_i'(o_i, o_t) \mid t] &= \sum_{i=1}^n c_i^2 \mathbb{E}_o [y_i^*(o) \mid t] \leq B_2. \quad (7)
\end{aligned}$$

Equation (5) - (7) shows that the defender's utility, the attacker's utility, and the budget consumption with s' are all the same as those with s^* , respectively, which indicates that s' satisfies the constraints while yielding the same value as the optimal solution s^* . Namely, we can construct an optimal solution where y_i is a function of o_i and o_t via the method presented above. \square

Although the original strategy space is too large to handle, Lemma 1 indicates that we can only consider a much smaller sub-space without sacrificing the solution quality. This result significantly simplifies the representation of the optimal solution. From now on, we will only consider solutions where $y_i = y_i(o_i, o_t)$ and $y_t = y_t(o_t)$. With Lemma 1, the program can be simplified as follows:

$$\begin{aligned}
& \max_{x, y} u_d(s, t) \\
& \text{s.t. } u_a(s, t) \geq u_a(s, t') \quad \forall t', \\
& \quad \sum_{i=1}^n c_i^1 x_i \leq B_1, \\
& \quad \sum_{i=1}^n c_i^2 \mathbb{E}_{o_{i,t}} [y_i(o_i, o_t) \mid t] \leq B_2, \quad (8) \\
& \quad 0 \leq x_i \leq 1 \quad \forall i, \\
& \quad 0 \leq y_i(o_i, o_t) \leq 1 \quad \forall i, \forall o_i
\end{aligned}$$

Independent Defense Strategy

According to Lemma 1, in the optimal strategy, y_i may depend on both o_i and o_t . However, bilinear terms still appear in the first and third constraints of Program (8), making the problem non-convex. Thus, we consider an alternative, simplified audit strategy y_i for the defender, with y_i being a function of only o_i , i.e., $y_i = y_i(o_i)$. Such a simplification makes y_i independent of the observations of other targets, hence called *independent defense strategy*. This not only simplifies the computation but also makes the implementation much easier.

With such a restriction on the strategy space, Program (8) becomes the following:

$$\begin{aligned}
& \max_{x, y} u_d(s, t) \\
& \text{s.t. } u_a(s, t) \geq u_a(s, t') \quad \forall t', \\
& \quad \sum_{i=1}^n c_i^1 x_i \leq B_1, \\
& \quad \sum_{i=1}^n c_i^2 \mathbb{E}_{o_i} [y_i(o_i) \mid t] \leq B_2, \quad (9) \\
& \quad 0 \leq x_i \leq 1 \quad \forall i, \\
& \quad 0 \leq y_i(o_i) \leq 1 \quad \forall i, \forall o_i
\end{aligned}$$

Denote the above program by $P(t)$. We still assume that the attacker breaks ties in favor of the defender when the attacker is indifferent among multiple targets that lead to the same maximum utility, as the defender still can slightly perturb the strategy to enforce such an outcome.

We first consider the following lemma, which will be useful for designing our algorithm.

Lemma 2. *There exists an optimal solution to $P(t)$ where the attacker's utility $u_a(i) = \min\{R_i^a, u_a(t)\}$, where t is the target chosen by the attacker in the solution.*

This result is straightforward and also holds in the standard Stackelberg security setting. So we omit its proof here.

Recall that the observation o_i is a binary random variable indicating whether a clue is discovered in the first stage. Intuitively, if the defender observes a clue ($o_i = 1$), target i is more likely to be attacked as $p_i^a > p_i^u$. Therefore, a natural idea is to set $y_i = 1$ whenever $o_i = 1$. This idea is confirmed theoretically by our next result. For ease of representation, we slightly abuse notation and write

$$y_i(1) = y_i(o_i = 1) \quad \text{and} \quad y_i(0) = y_i(o_i = 0)$$

to denote the defender's strategy in the second stage under different observations from the first stage, respectively. We also denote by d_i the expected probability of auditing target i in the second stage conditioned on target i being attacked, that is,

$$d_i = \mathbb{E}_{o_i} [y_i(o_i) \mid i] = x_i p_i^a y_i(1) + (1 - x_i p_i^a) y_i(0). \quad (10)$$

Lemma 3. *It is without loss of generality to assume that the defender will audit target i with probability 1 once a clue is observed on i in the first stage. In other words, there is an optimal solution to $P(t)$ where $y_i(1) = 1$ for all i .*

Note that Program $P(t)$ still involves non-convex constraints. To solve the program, we introduce a parameter m and consider the following program $Q(t; m)$:

$$\begin{aligned} \min_{x, y} \quad & c_t^2 d_t + \sum_{i \neq t} c_i^2 \left(\frac{p_i^u}{p_i^a} d_i + \frac{p_i^a - p_i^u}{p_i^a} y_i(0) \right) \\ \text{s.t.} \quad & u_a(t) = m, \\ & u_a(i) \leq m \quad \forall i, \\ & \sum_{i=1}^n c_i^1 x_i \leq B_1, \\ & 0 \leq x_i \leq 1 \quad \forall i, \\ & 0 \leq y_i(0) \leq 1 \quad \forall i \end{aligned} \quad (11)$$

The above program fixes the attacker's best possible utility to m and minimizes the budget of the second stage that can enforce such a utility. We fix $y_i(1)$ to 1 in the objective function according to Lemma 3 and minimize the second-stage budget consumption by optimizing x and $y(0)$. The constraints serve multiple purposes: they ensure that attacking target t is the attacker's best response, keep the first stage resource consumption within the investigation budget constraint B_1 , and specify the feasible ranges of x and $y_i(0)$.

Definition 3 (The Minimum Attacker Utility). *The minimum attacker utility $m^*(t)$ is the smallest attacker's utility on target t achievable by a feasible solution in $P(t)$.*

Theorem 1. *Any optimal solution to Program $Q(t; m^*(t))$ is also an optimal solution to Program $P(t)$.*

Proof. According to Definition 3, $P(t)$ has a feasible solution, denoted by s which leads to the attacker's utility $u_a(s, t) = m^*(t)$ and satisfies the second budget constraint. We first show that s is feasible for Program $Q(t; m^*(t))$. It is clear that the two programs share all constraints except that $P(t)$ includes an additional audit budget constraint. Therefore, it is quite intuitive that s is a feasible solution for $Q(t; m^*(t))$ as it is feasible for $P(t)$.

Let s^* denote any optimal solution of $Q(t, m^*(t))$. The optimality of s^* in Program $Q(t; m^*(t))$ leads to the following inequality:

$$\begin{aligned} & c_t^2 d_t + \sum_{i \neq t} c_i^2 \left(\frac{p_i^u}{p_i^a} d_i + \frac{p_i^a - p_i^u}{p_i^a} y_i^*(0) \right) \\ & \leq c_t^2 d_t + \sum_{i \neq t} c_i^2 \left(\frac{p_i^u}{p_i^a} d_i + \frac{p_i^a - p_i^u}{p_i^a} y_i(0) \right) \leq B_2. \end{aligned}$$

The second inequality holds since s is feasible for $P(t)$ and thus satisfies its budget constraint. Again, since $P(t)$ and $Q(t; m^*(t))$ share all the constraints except for the audit budget constraint B_2 , s^* is clearly feasible for $P(t)$.

We prove that it is also optimal for Program $P(t)$ by contradiction. Assume that s^* is not optimal for $P(t)$ and denote the optimal solution to $P(t)$ by s' . It follows that $u_d(s', t) > u_d(s^*, t)$. By definition, we know that

$$\begin{aligned} u_d(s, t) &= P_t^d + (R_t^d - P_t^d)d_t, \\ u_a(s, t) &= R_t^a + (P_t^a - R_t^a)d_t. \end{aligned}$$

On target t , the attacker's utility satisfies $m' = u_a(s', t) < u_a(s^*, t) = m^*(t)$. In other words, there is an attacker's utility $m' < m^*(t)$ on target t , which is achievable by a feasible solution to Program $P(t)$. This contradicts the definition of $m^*(t)$. Therefore, the contradiction proves that s^* is also optimal for Program $P(t)$. \square

Theorem 1 offers us an alternative plan for optimizing Program $P(t)$. By solving $Q(t; m^*(t))$, we can obtain a solution that is optimal for both $Q(t; m^*(t))$ and $P(t)$. A natural idea is to find the minimum attacker utility $m^*(t)$ and solve for the optimal solution of $P(t)$ by optimizing $Q(t; m^*(t))$. Before presenting our optimization algorithm, we first prove the following useful results.

Lemma 4. *For any m , there is an optimal solution s to $Q(t; m)$, where the attacker's utility is*

$$u_a(s, i) = \min\{R_i^a, m\}, \forall i.$$

According to Lemma 4, we can, without loss of generality, focus only on solutions with $u_a(s, i) = \min\{R_i^a, m\}, \forall i$. For any such solution, when m is given, the expected probability of covering each attacked target becomes a constant:

$$d_i = \begin{cases} \frac{R_i^a - m}{R_i^a - P_i^a} & \text{if } R_i^a > m \\ 0 & \text{if } R_i^a \leq m \end{cases}.$$

And the relationship between $y_i(0)$ and x_i can be derived from Equation (10):

$$y_i(0) = \frac{d_i - x_i p_i^a}{1 - x_i p_i^a}. \quad (12)$$

Lemma 5. *For any m , there is an optimal solution to Program $Q(t; m)$ where $x_t = 0$.*

Proof. Let s^* be an optimal solution to $Q(t; m)$. Consider solution s where

$$x_i = \begin{cases} 0 & \text{if } i = t \\ x_i^* & \text{if } i \neq t \end{cases}, \quad y_i(0) = \begin{cases} d_t^* & \text{if } i = t \\ y_i^*(0) & \text{if } i \neq t \end{cases},$$

in which the defender allocates no resources to investigate target t . We prove by showing that s is also an optimal solution. Note that the only difference between s and s^* lies in the strategy on target t . As the players' utilities are independent on each target, solely letting $x_t = 0$ in any strategy will not change their utilities on target $i \neq t$. Since $d_t = x_t p_t^a + (1 - x_t p_t^a)y_t(0) = d_t^*$, s results in the same attacker utility c on target t .

Now we check the budget constraints. Intuitively, the total cost of the first step is reduced by changing from s^* to s as the $x_t = 0 \leq x_t^*$. In the second step, the total cost is

$$\begin{aligned} & c_t^2 [x_t p_t^a + (1 - x_t p_t^a)y_t(0)] \\ & + \sum_{i \neq t} c_i^2 [x_i p_i^u + (1 - x_i p_i^u)y_i(0)] \\ & = c_t^2 d_t^* + \sum_{i \neq t} c_i^2 [x_i^* p_i^u + (1 - x_i^* p_i^u)y_i^*(0, 0)] \\ & = c_t^2 [x_t^* p_t^a + (1 - x_t^* p_t^a)] \\ & + \sum_{i \neq t} c_i^2 [x_i^* p_i^u + (1 - x_i^* p_i^u)y_i^*(0, 0)], \end{aligned}$$

which is still optimal. Therefore, s is an optimal solution with $x_t = 0$. \square

Remark 1. Lemma 5 implies that the defender will adopt different investigation strategies on different targets. When the defender believes that target t will be attacked, there is simply no need to investigate before auditing. Yet on the targets that will not be attacked from their perspective, the defender needs to employ the initial investigation to reduce the budget expenditure on them. Another way of understanding the costs on these targets is that, once the defender has allocated enough resources to cover other targets, the attacker's best choice is already t regardless of whether the defender investigates t in the first stage.

Lemma 6. The optimal value of Program $Q(t; m)$ is a monotone decreasing function of m .

Now we show that when m is given, $Q(t; m)$ can be reformulated as a convex optimization problem. We eliminate variable $y_i(0)$ and represent it by x_i according to Equation (12). We introduce u_i to represent $\frac{1}{1-x_i p_i^a}$ and v_i to represent $1 - x_i p_i^a$. We also use $\alpha = (\alpha_1, \alpha_2, \dots, \alpha_n)$, $\beta = (\beta_1, \beta_2, \dots, \beta_n)$ to represent the vector of α_i and β_i . We reformulate Program $Q(t; m)$ as the following $\tilde{Q}(t; m)$:

$$\begin{aligned} \min_{\alpha, \beta, w, x} \quad & \sum_{i \neq t} c_i^2 \left\{ \frac{p_i^u}{p_i^a} d_i + \frac{p_i^a - p_i^u}{p_i^a} [1 - (1 - d_i) \alpha_i] \right\} \\ & + c_t^2 d_t \\ \text{s.t.} \quad & \sum_{i=1}^n c_i^1 x_i \leq B_1, \\ & \beta_i = 1 - x_i p_i^a \quad \forall i, \\ & w = 1, \\ & \alpha_i \beta_i = w^2 \quad \forall i, \\ & 0 \leq x_i \leq \min \left\{ 1, \frac{d_i}{p_i^a} \right\} \quad \forall i, \\ & \alpha_i, \beta_i \geq 0 \quad \forall i \end{aligned} \quad (13)$$

Notice that given m , d_i becomes a constant for each i . In $\tilde{Q}(t; m)$, the objective function and all constraints are linear except for the fourth constraint, which is a rotated second-order cone constraint. Therefore, $\tilde{Q}(t; m)$ is a convex optimization problem.

Lemma 6 indicates that for any feasible Program $Q(t; m)$ with parameter $m \geq m^*(t)$, its optimal value will not exceed the total cost of the second stage under $P(t)$. Hence, its optimal solution is feasible for $P(t)$. According to Definition 3, for any $m < m^*(t)$, the solution to $Q(t; m)$ is infeasible for $P(t)$. Therefore, we can use binary search to find $m^*(t)$ with a feasible corresponding Program $Q(t; m)$, and according to Theorem 1, any optimal solution to $Q(t; m^*(t))$ is also an optimal solution to $P(t)$. In each iteration of the binary search, the algorithm solves Program $\tilde{Q}(t; m)$, which is a convex program and can be solved efficiently. The details are shown in Algorithm 1.

Algorithm 1: Finding the optimal independent defender strategy

Input : Payoffs $\{R^d, P^d, R^a, P^a\}$, budget B , costs $\{c^1, c^2\}$, alarm probabilities $\{p^a, p^u\}$, precision ϵ .

Output: Solution s .

```

1  $u_{max} \leftarrow -\infty;$ 
2 for  $t = 1$  to  $n$  do
3   Initialize  $m_{low} \leftarrow P_t^a$ ,  $m_{high} \leftarrow R_t^a$ ;
4   if  $\tilde{Q}(t; m_{high})$  is infeasible then
5     continue;
6   while  $m_{high} - m_{low} > \epsilon$  do
7     if  $\tilde{Q}(t; \frac{m_{high} - m_{low}}{2})$  is infeasible then
8        $m_{low} \leftarrow \frac{m_{high} - m_{low}}{2}$ ;
9     else
10       $m_{high} \leftarrow \frac{m_{high} - m_{low}}{2}$ ;
11    $u_d \leftarrow \frac{(m_{high} - R_t^a)(R_t^d - P_t^d)}{P_t^a - R_t^a} + P_t^d$ ;
12   if  $u_d > u_{max}$  then
13      $u_{max} \leftarrow u_d$ ;
14   Set  $s$  to be the solution to  $\tilde{Q}(t; m_{high})$ ;
15 return  $s$ ;

```

Experiments

We report our experiment results in this section. The results are averaged over 100 game instances. Following standard practice, player rewards and penalties are drawn from $U[0, 1]$ and $U[-1, 0]$, respectively. Audit costs c_i^2 and investigation true positive rates p_i^a are sampled from $U[0, 1]$. Let k_i and l_i represent the ratios of c_i^1 to c_i^2 and p_i^u to p_i^a , respectively, i.e.,

$$k_i = \frac{c_i^1}{c_i^2} \quad \text{and} \quad l_i = \frac{p_i^u}{p_i^a}.$$

For ease of presentation, let $k = \{k_1, k_2, \dots, k_n\}$, $l = \{l_1, l_2, \dots, l_n\}$.

We use the canonical SSG without the investigation phase as the baseline. Both models share identical audit resources and marginal audit costs. To solve for the optimal SSG solution, the following linear program is applied to each target t , where vector z represents the defender's strategy, and each element z_i indicates the probability of covering target i .

$$\begin{aligned} \max_z \quad & u_d(t) = z_t R_t^d + (1 - z_t) P_t^d \\ \text{s.t.} \quad & z_t P_t^a + (1 - z_t) R_t^a \geq z_{t'} P_{t'}^a + (1 - z_{t'}) R_{t'}^a \quad \forall t', \\ & \sum_{i=1}^n c_i^2 z_i \leq B_2, \\ & 0 \leq z_i \leq 1 \quad \forall i \end{aligned} \quad (14)$$

In the experiment results presented in this section, all programs are solved using Gurobi 11.0 (Gurobi Optimization, LLC 2023).

Effect of Price Ratio and Investigation Accuracy

Figure 1 shows the results when all the targets share the same k and l . The number of targets is 50 in all the game instances for this experiment. We set $B_2 = 5$. Additionally, in the two-stage model, the defender has investigation resources of $B_1 = 0.1$. For each budget B_2 , we compute the corresponding utilities for both the defender and the attacker with $k \in \{0.1, 0.2, \dots, 0.5\}$ and $l \in \{0.1, 0.2, \dots, 0.5\}$. Different curves on each subgraph represent the utility of the defender (attacker) with different l . The black dashed line represents the payoff in a standard Stackelberg security game without the investigation phase. The result implies that with high cost-effectiveness, the investigation phase can provide a significant utility boost to the defender.

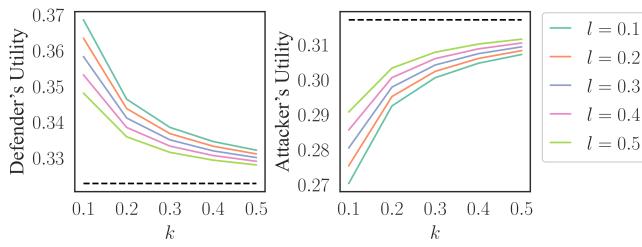


Figure 1: The defender and attacker’s utilities with different budgets, price ratios, and investigation accuracy.

Effect of Initial Investigation

In this experiment, we analyze the effect of the resources invested in the first stage on the defender’s utility. In the experiment, k and l are both randomly drawn from $U[0, 1]$. We set $B_2 \in \{1, 3\}$ and $n \in \{30, 40, 50\}$. We compare the improvement in defensive effectiveness provided by initial investigations with $B_1 \in \{0.05, 0.1, \dots, 0.5\}$ under different numbers of targets and audit resources. Figure 2 shows that when audit resources are limited, investing a small amount of resources in the initial investigation phase can significantly enhance defensive effectiveness. The defender’s utility increases with the consumption of resources in the first phase. Therefore, allocating resources to the initial investigation is highly cost-effective, providing a decent improvement in protection at a relatively low cost.

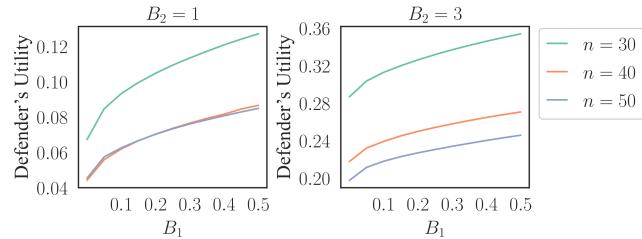


Figure 2: The defender’s utility with different investigation resources. Different curves represent different targets in the game.

The Resource-Saving Effect of Initial Investigation

Beyond improving defender’s utility, the introduction of the first-stage investigation can also reduce the amount of defensive resources required, thereby enhancing resource efficiency. To evaluate this effect, we fix the defender’s second-stage budget at 1 and vary the investigation budget B_1 in the first stage. We compare this setting with a standard SSG that uses the same total amount of defensive resources. The experimental results illustrate how much defensive resource a standard SSG must expend in order to achieve the same level of defensive performance as the independent defense strategy.

In this experiment, the parameter k is randomly sampled from $U[0, 1]$. We compare different values of $l \in \{0.1, 0.2, 0.3, 0.4, 0.5\}$ and first-stage budget levels $B_1 \in \{0.01, 0.02, 0.03, 0.04, 0.05\}$, and measure how much defensive resource a standard SSG must consume to achieve the same defensive performance as an independent defense strategy.

The experimental results indicate that the introduction of an initial investigation phase significantly reduces the amount of defensive resources required. Moreover, for all investigation accuracy levels considered, the defensive resources saved are significantly greater than those spent on the initial investigation.

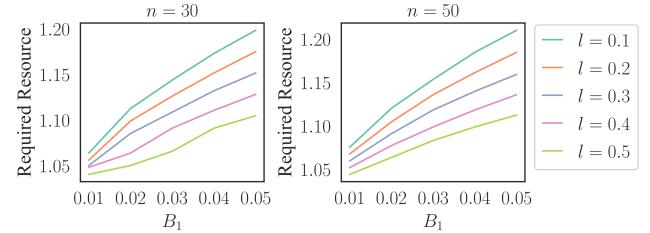


Figure 3: The defensive resources required by a standard SSG to achieve the same performance as the independent defense strategy. Different curves correspond to different levels of initial investigation accuracy.

Conclusion

In this paper, we propose a two-stage audit game model that includes an initial investigation phase. Due to the non-convex nature of the problem, we consider simpler strategies, called independent defender strategies. Such a strategy is both computationally efficient and much easier to implement in reality. We provide an efficient algorithm to find an optimal independent defender strategy. Through a series of experiments, we illustrate the impact of the initial investigation phase. The experiment results show that even with very limited resources, the investigation phase considerably improves the defender’s utility.

Acknowledgments

We thank all the anonymous reviewers for their helpful comments. This work is supported in part by the National Natural Science Foundation of China (No. 72192805), Public

Computing Cloud, Renmin University of China, the fund for building world-class universities (disciplines) of Renmin University of China.

References

- Blocki, J.; Christin, N.; Datta, A.; Procaccia, A.; and Sinha, A. 2013. Audit games. *Proceedings of the 23rd International Joint Conference on Artificial Intelligence (IJCAI)*, 41–47.
- Blocki, J.; Christin, N.; Datta, A.; Procaccia, A.; and Sinha, A. 2015. Audit games with multiple defender resources. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 29.
- Bondi, E.; Oh, H.; Xu, H.; Fang, F.; Dilkina, B.; and Tambe, M. 2018. Biodiversity Conservation with Drones: Using Uncertain Real-Time Information in Signaling Games to Prevent Poaching. Working paper at the International Conference on Machine Learning AI for Social Good Workshop, Long Beach, United States, 2018.
- Bondi, E.; Oh, H.; Xu, H.; Fang, F.; Dilkina, B.; and Tambe, M. 2020. To signal or not to signal: Exploiting uncertain real-time information in signaling games for security and sustainability. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, 1369–1377.
- Conitzer, V.; and Sandholm, T. 2006. Computing the optimal strategy to commit to. In *Proceedings of the 7th ACM conference on Electronic commerce*, 82–90.
- Fellingham, J. C.; and Newman, D. P. 1985. Strategic considerations in auditing. *Accounting Review*, 634–650.
- Furnham, A.; and Gunter, B. 2015. *Corporate assessment (Routledge Revivals): auditing a company's personality*. Routledge.
- Guo, Q.; An, B.; Bosansky, B.; and Kiekintveld, C. 2017. Comparing Strategic Secrecy and Stackelberg Commitment in Security Games. In *IJCAI*, 3691–3699.
- Gurobi Optimization, LLC. 2023. Gurobi Optimizer Reference Manual.
- Hatfield, R. C.; Jackson, S. B.; and Vandervelde, S. D. 2011. The effects of prior auditor involvement and client pressure on proposed audit adjustments. *Behavioral Research in Accounting*, 23(2): 117–130.
- Houck, T. P. 2003. *Why and how audits must change: practical guidance to improve your audits*. John Wiley & Sons.
- Huang, T.; Shen, W.; Zeng, D.; Gu, T.; Singh, R.; and Fang, F. 2020. Green security game with community engagement. 529–537.
- Jain, M.; Tsai, J.; Pita, J.; Kiekintveld, C.; Rathi, S.; Tambe, M.; and Ordóñez, F. 2010. Software assistants for randomized patrol planning for the lax airport police and the federal air marshal service. *Interfaces*, 40(4): 267–290.
- Kiekintveld, C.; Jain, M.; Tsai, J.; Pita, J.; Ordóñez, F.; and Tambe, M. 2009. Computing optimal randomized resource allocations for massive security games.
- Loebbecke, J. K.; and Steinbart, P. J. 1987. An Investigation of the Use of Preliminary Analytical Review to Provide Substantive Audit Evidence. *Auditing: A Journal of Practice & Theory*, 6(2).
- Ma, X.; He, Y.; Luo, X.; Li, J.; Zhao, M.; An, B.; and Guan, X. 2018. Camera placement based on vehicle traffic for better city security surveillance. *IEEE Intelligent Systems*, 33(4): 49–61.
- Pita, J.; Jain, M.; Marecki, J.; Ordóñez, F.; Portway, C.; Tambe, M.; Western, C.; Paruchuri, P.; and Kraus, S. 2008. Deployed armor protection: the application of a game theoretic model for security at the los angeles international airport. In *Proceedings of the 7th international joint conference on Autonomous agents and multiagent systems: industrial track*, 125–132.
- Shen, W.; Chen, W.; Huang, T.; Singh, R.; and Fang, F. 2020. When to follow the tip: security games with strategic informants. In *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence*.
- Shen, W.; Han, M.; Chen, W.; Huang, T.; Singh, R.; Xu, H.; and Fang, F. 2024. An Extensive Study of Security Games with Strategic Informants. *Artificial Intelligence*, 104162.
- Shi, Z. R.; Schlenker, A.; Hay, B.; Bittleston, D.; Gao, S.; Peterson, E.; Trezza, J.; and Fang, F. 2020. Draining the water hole: Mitigating social engineering attacks with cyber-tweak. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, 13363–13368.
- Singleton, T. W.; and Singleton, A. J. 2010. *Fraud auditing and forensic accounting*, volume 11. John Wiley & Sons.
- Tsai, J.; Rathi, S.; Kiekintveld, C.; Ordóñez, F.; and Tambe, M. 2009. IRIS-a tool for strategic security allocation in transportation networks. *AAMAS (Industry Track)*, 37–44.
- Xu, H.; Rabinovich, Z.; Dughmi, S.; and Tambe, M. 2015. Exploring information asymmetry in two-stage security games. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 29.
- Xu, H.; Wang, K.; Vayanos, P.; and Tambe, M. 2018. Strategic coordination of human patrollers and mobile sensors with signaling for security games. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32.
- Zhang, Y.; Guo, Q.; An, B.; Tran-Thanh, L.; and Jennings, N. R. 2019. Optimal interdiction of urban criminals with the aid of real-time information. In *Proceedings of the AAAI conference on artificial intelligence*, volume 33, 1262–1269.