Inspection Games with Incomplete Information and Heterogeneous Resources

Bobak McCann, Mathieu Dahan
H. Milton Stewart School for Industrial and Systems Engineering, Georgia Institute of Technology, Atlanta, GA 30332
{bmccann6@gatech.edu, mathieu.dahan@isye.gatech.edu}

Abstract. We study a two-player zero-sum inspection game with incomplete information, where an inspector deploys resources to maximize the expected damage value of detected illegal items hidden by an adversary across capacitated locations. Inspection and illegal resources differ in their detection capabilities and damage values. Both players face uncertainty regarding each other's available resources, modeled via stochastic player types. To solve this large-scale game, we first characterize the locations' marginal detection probabilities and expected damage values in equilibrium. We then design combinatorial algorithms that coordinate each player type's resources to match these marginal values while satisfying best-response conditions. Our approach computes a Nash equilibrium with linear support in polynomial time. Notably, equilibrium strategies are independent of the inspector's detection capabilities, implying no strategic advantage from concealing them. We extend our analysis to a two-stage problem, computing in pseudo-polynomial time the optimal acquisition of inspection resources given subsequent interactions with the adversary. A case study on drug interdiction at U.S. seaports shows that reducing uncertainty about drug shipments enables the inspector to detect an additional $20.04 million worth of narcotics annually, quantifying the value of intelligence. Our results provide actionable guidance for security agencies acquiring and deploying heterogeneous inspection resources under uncertainty.

Key words: Strategic inspection; zero-sum game; imperfect information; heterogeneous resources; resource acquisition
Subject classifications: Military: search/surveillance; Games/group decisions: noncooperative; Programming: linear: large scale systems

1. Introduction

1.1. Motivation

Smuggling of illicit goods remains a pervasive and pressing threat worldwide. From narcotics and weapons to human trafficking and endangered wildlife, illicit flows fuel organized crime and undermine security. The scale of the problem is enormous—the global economy loses more than $2 trillion each year to smuggling, counterfeiting, and related black-market trade (United Nations Conference on Trade and Development 2020). These illegal activities inflict not only economic damage but also grave social harm: drug smuggling, for example, has exacerbated public health crises and violence in many communities. Despite enforcement efforts, large volumes of contraband still slip through borders. In the U.S. alone, border agents seized over 418,000 pounds of illegal drugs (including more than 15,000 pounds of fentanyl) in just a nine-month period (U.S. Customs and Border Protection 2024), a figure that hints at the even greater quantity of dangerous goods that evade detection. Such statistics underscore the high stakes and societal importance of improving inspection and interdiction strategies against smuggling.

Current inspection operations in domains like port security and airport screening are extensive, yet they face the challenge of balancing security with efficiency. Modern ports of entry handle immense volumes of cargo and travelers—for example, more than 100,000 commercial cargo trucks cross U.S. land borders on a typical day (U.S. Department of Homeland Security, Science & Technology Directorate 2023)—which makes high-throughput screening essential. To cope with this, agencies employ a range of tools and protocols. Non-intrusive inspection technologies, particularly X-ray and gamma-ray scanners mounted on mobile vehicles, have become indispensable for examining freight and luggage without disrupting commerce (Port of Barcelona 2020). At airports, every passenger bag passes through X-ray scanners, and travelers themselves go through metal detectors or body scanners, supplemented by random manual checks. Canine units (sniffer dogs) and handheld detectors are also regularly deployed to detect drugs, explosives, or other contraband.

Despite such advances, operational challenges in real-world inspection remain formidable. Coordination of heterogeneous resources is a key difficulty: security agencies must deploy different types of inspection assets—from large-scale X-ray scanners and radiation detectors to specialized canine teams and human inspectors—in a coordinated fashion. Deciding how to allocate these diverse resources across numerous entry points and vast numbers of shipments is a complex task, especially under limited time and budget constraints. Multiple contraband types present another challenge. Inspectors are tasked with intercepting a broad spectrum of illegal goods, each of which can vary in the amount of damage they can cause to society. Uncertainty and adversarial adaptation further complicate inspections. Those trying to smuggle contraband are strategic, constantly probing for weaknesses in enforcement and altering their tactics to avoid detection. Inspectors, for their part, often operate with incomplete information—they do not know which travelers or shipments conceal illegal goods, and they cannot feasibly search every single one without causing significant economic disruptions (Jorgic et al. 2024). This intrinsic uncertainty about where, when, and how adversaries will act makes optimal inspection planning exceedingly difficult in practice.

These real-world challenges have spurred a growing body of analytical research on inspection and security screening. Inspection games model the strategic contest between an inspector and an adversary and have long informed applications from arms-control monitoring to airport screening (Avenhaus and Canty 2009). Yet most formulations that randomize the allocation of multiple resources assume homogeneous inspection and adversary resources, and complete information (Dziubiński and Roy 2018, Bahamondes and Dahan 2024, Boone and Dahan 2025). As such, they also cannot be used to recommend which inspection resources to purchase. In light of these gaps, we posit the following research question: *How should an inspector optimally acquire and coordinate heterogeneous inspection resources to detect illegal activity from partially-known adversaries?*

1.2. Contributions

– We formulate the first inspection game in which both players randomize the allocation of multiple resources with heterogeneous characteristics across capacitated locations, while facing uncertainty regarding their opponent's resources. Modeling such features is critical for addressing realistic scenarios that arise in security domains and improving the relevance of recommended inspection schedules. This framework builds on previous models in the literature on hide-and-seek games (Gal and Casas 2014, Dziubiński and Roy 2018) and contributes more generally to the literature on search games (Kikuta and Ruckle 1997, 2002, Zoroa et al. 2004, Alpern et al. 2010, Lidbetter and Lin 2019, Clarkson et al. 2022).

– To solve this large-scale game of incomplete information with combinatorial action sets, we develop a three-step approach for efficiently computing Nash equilibria. First, we develop an algorithm (Algorithm 1) to partition the locations into zones and determine target expected total damage values of hidden items to be achieved within each zone in equilibrium. Second, we design a combinatorial algorithm (Algorithm 2) that coordinates the illegal resources of each adversary type across the capacitated locations to meet the target expected damage values within each zone. Using properties of the target expected damage values (Lemma 1), we show that Algorithm 2 feasibly constructs in cubic time a strategy for the adversary with linear support that satisfies best-response conditions (Proposition 1). Third, we construct in quadratic time an inspection strategy by cycling inspection resources for each inspector type (Algorithm 3), equalizing the locations' marginal detection probabilities within each zone (Proposition 2). Our approach yields in polynomial time a Nash equilibrium with linear support, easily implementable by players in practice (Theorem 1). To the best of our knowledge, Algorithm 2 is the first combinatorial algorithm developed in security games for coordinating the allocation of heterogeneous resources in locations with heterogeneous capacities to construct equilibrium strategies, with applications beyond the studied inspection game.

– Our equilibrium analysis provides valuable insights to security agencies that monitor critical systems being utilized by adversaries to carry out illegal activities: The analytical characterization of the locations' detection probabilities and expected damage values in equilibrium determines the locations that are the most vulnerable, as well as the ones that must not be inspected. Remarkably, we show that the players' equilibrium strategies are independent of the inspection resources (Corollary 1), implying that the inspector has no strategic benefit from concealing their resource detection capabilities from the adversary.

– We investigate a two-stage tactical-level problem where the inspector first acquires inspection resources under cost constraints, and then interacts against the adversary according to the inspection game. Using the equilibrium properties of the inspection game, we show that the resulting mathematical program with equilibrium constraints reduces to a bounded generalized multiple-choice knapsack problem, providing optimal acquisition plans in pseudo-polynomial time using dynamic programming (Theorem 2).

– We conduct a case study on maritime port inspections for drug interdiction, using U.N. drug seizure records and container throughput data from ports in California and Florida. We find that reducing the inspector's uncertainty regarding drug shipment volumes—measured via the normalized Shannon entropy—can increase the monetary value of seized narcotics by up to \$20.04M (resp. \$7.38M) in California (resp. Florida) annually, thereby quantifying the value of improved intelligence. Furthermore, we derive qualitative insights into how the dispersion of smuggling capacities critically influences optimal inspection resource acquisition decisions, providing practical guidance to security agencies tasked with inspecting critical systems using heterogeneous resources against partially-known strategic adversaries.

1.3. Related Work

A variety of search game models have been proposed, involving settings where an adversary hides multiple items, either simultaneously or sequentially, in different locations. The inspector's goal is to optimally locate these items. This concept has been explored in numerous studies, including works by Kikuta and Ruckle (1997), Kikuta and Ruckle (2002), Zoroa et al. (2004), Alpern et al. (2010), Lidbetter and Lin (2019), and Clarkson et al. (2022). Commonly, these papers assume that the inspector systematically moves from one potential hiding location to another in pursuit of the concealed items. Our game diverges from these existing models in several key aspects. First, in our model, the inspector randomizes their search across the potential hiding locations. Second, we account for the inspector's simultaneous use of multiple imperfect and heterogeneous inspection resources, which may not always detect a hidden item, even when inspecting the correct location. Finally, our game models settings where multiple items, each with heterogeneous values, are hidden simultaneously. This means certain items hold greater value for the inspector to find than others. Such nuances are not addressed in the current body of literature.

Our inspection problem builds upon hide-and-seek formulations. The hide-and-seek game is a two-person zero-sum game that was first introduced by Von Neumann (1953). In its original formulation, the hider and seeker interact on a square matrix of non-negative entries a_{ij} in which the hider selects an entry of the matrix and the seeker selects either a row or column. If the seeker's choice contains the entry that the hider picked, then the hider will pay the seeker a_{ij}. Otherwise, the seeker will pay the hider a_{ij}. Subsequent studies incorporated imperfect detection in discrete locations, most notably the pursuit–evasion formulation of Gal and Casas (2014), and the multi–item extension of Dziubiński and Roy (2018), in which the hider distributes several identical objects while the inspector chooses a subset of boxes to open. Our work builds upon these models by allowing multiple preys to coordinately hide in heterogeneously capacitated locations, and by allowing there to be multiple coordinated predators who differ in their capabilities.

Our inspection game is closely related to the simultaneous hide-and-seek game of perfect information studied by Bahamondes and Dahan (2024), in which each location has a finite (potentially heterogeneous) hiding capacity and an associated detection probability. In their model, the seeker deploys multiple identical inspection resources, and the hider allocates multiple identical items. When a location is inspected, each hidden item is detected independently with the location's detection probability. In contrast, our model accounts for the uncertainty faced by the players and assumes that detection depends on the specific inspection resource deployed to a location. Inspection resources vary in their detection capabilities, and items differ in their damage values.

A sizable stream of game-theoretic work studies inspection and search problems in which *each player knows their own resources but only a distribution over the opponent's*. Classic treaty-verification models examine single- or multi-stage inspection under such uncertainty (Avenhaus and Zamir 2002, Avenhaus and von Stengel 2004, von Stengel 2014), while early hide-and-seek variants consider priors over the hider's location or payload (Dresher 1962, Gal and Kohn 2016). In security game applications, one-sided uncertainty (only the attacker's type is private) led to algorithms such as DOBSS and ERASER (Kiekintveld et al. 2011, Letchford and Conitzer 2013), later generalized to continuous types and learning updates (Yang et al. 2014, Nguyen et al. 2019). Work that does feature *two-sided* uncertainty—for example in patrol scheduling (Basilico et al. 2016), network interdiction (Borrero et al. 2019), or cyber-deception sensor placement (Pirani et al. 2021)—assumes either homogeneous sensor efficacy or ignores location capacities. We allow heterogeneous sensors and explicit capacity constraints.

In McCarthy et al. (2016), the authors study a security game, and examine both the tactical selection of heterogeneous resources under a finite budget and their subsequent operational allocation. They model an interdiction problem where a single illegal logger traverses a network from a source to a target node, while a defender selects and deploys resources that patrol a chosen subset of edges to intercept the logger. In contrast, we consider a setting where a defender simultaneously positions several heterogeneous inspection resources at discrete locations, each with varying capacities for concealing contraband. As far as we are aware, the dual focus in McCarthy et al. (2016) on both tactical and operational facets of resource selection and allocation is a distinctive feature not found in prior work on security games. As such, we believe ours is the first work to consider both the tactical and operational level questions of resource selection and allocation in security games where *both* players wield heterogeneous resources. Crucially, our equilibrium analysis and solution approach enable us to reduce the tactical resource selection question into an efficiently solvable knapsack problem, suitable for a wide array of real-world situations.

Preliminary results of our work appeared in McCann and Dahan (2022), which considered an inspection game with *perfect* information involving heterogeneous inspection resources and *homogeneous* illegal resources. However, the equilibrium characterization in that simpler setting does not extend to our framework, which introduces fundamental complexities due to player uncertainty and heterogeneity in the adversary's illegal items. Addressing these challenges requires a novel combinatorial algorithm for randomizing the allocation of heterogeneous resources across capacitated locations—a methodological advance not present in the security game literature. Moreover, we leverage this equilibrium analysis to address a distinct level of decision-making: the tactical selection of the inspector's resource portfolio prior to adversarial interaction. This two-stage extension significantly expands the scope and applicability of inspection models under uncertainty.

2. Problem Description

We consider an inspection problem involving an inspector tasked with monitoring a system being utilized by an adversary to carry out illegal activities. The inspector coordinates a set of inspection resources across the system to detect illegal resources allocated by the adversary. The inspection resources may consist for instance of police units equipped with sensors or K-9s that may differ in their detection capabilities. Similarly, the illegal resources may represent drugs or weapons that differ in their monetary values or adverse impacts on a nation. Furthermore, the decision-makers face uncertainty regarding the resources available to their opponent. Therefore, we investigate a game of incomplete information, where the inspector (resp. adversary) aims to maximize (resp. minimize) the expected damage value of detected illegal resources.

Formally, we consider an adversary A interested in allocating illegal resources across a system consisting of a set of n locations (or regions) $L = [\![1, n]\!]$. Depending on the setting, a location may represent a geographical area around a country border that must be crossed by smugglers, or a container used to transport illegal goods. Each location $\ell \in L$ is capacitated and can receive at most $c_\ell \in \mathbb{Z}_{>0}$ illegal resources. We denote by $m \coloneqq \sum_{\ell \in L} c_\ell$ the total capacity of the system. To prevent the adversary from carrying out their illegal activities, an inspector I is tasked with detecting illegal resources within each location by coordinating inspection resources across the system. To reduce sensing interference, we assume that each location can be inspected by no more than one inspection resource at a time.

We model the strategic interactions between the inspector and adversary by formulating a two-player zero-sum incomplete-information game $\Gamma \coloneqq \langle \{I, A\}, (\Theta_I, \Theta_A), (\pi_I, \pi_A), (\mathcal{A}_I, \mathcal{A}_A), (u, -u) \rangle$. In this game, we account for the fact that each player faces uncertainty pertaining to their opponent's resources. From the inspector's perspective, the adversary can be categorized into a finite set Θ_A of types, with each adversary type $\theta_A \in \Theta_A$ occurring with probability $\pi_A(\theta_A)$. Similarly, the adversary believes that the inspector's type belongs to a finite set Θ_I, with each type $\theta_I \in \Theta_I$ occurring with probability $\pi_I(\theta_I)$. We assume that occurrence probabilities are independent across players and are common knowledge, along with the sets of player types and the location capacities.

Each player type is characterized by the resources they can allocate: Given an adversary type $\theta_A \in \Theta_A$, we denote by $v_j^{\theta_A} \in \mathbb{R}_{\geq 0}$ the *damage value* of the adversary's j-th illegal resource. Such values account for the importance of each illegal resource, based on their monetary or informational value, and overall adverse impact to society if not detected by the inspector. Similarly, given an inspector type $\theta_I \in \Theta_I$, we let $d_i^{\theta_I} \in [0, 1]$ denote the *detection rate* of the inspector's i-th inspection resource. An inspection resource's detection rate represents the probability of detecting each illegal resource within the inspected location. This accounts for the detection capabilities of humans, K-9s, and novel sensing technologies. Overall, player types capture various uncertainties, such as the inspector's lack of knowledge regarding the number of illegal resources allocated by the adversary (which may vary over the planning horizon) and the adversary's uncertainty regarding the exact detection rates of the inspector's detection technology.

For notational convenience and without loss of generality, we assume that any inspector type allocates n inspection resources (some possibly with zero detection rate) and any adversary type allocates m illegal resources (some possibly with zero damage value). The set of actions of any inspector type is then given by $\mathcal{A}_I := \mathfrak{S}_n$ (i.e., the set of permutations of size n). Note that for each inspection plan $\varphi \in \mathcal{A}_I$ and each location $\ell \in L$, $\varphi(\ell) \in [\![1,n]\!]$ represents the inspection resource allocated to location ℓ. Similarly, we represent an allocation plan for the adversary as a mapping ψ from L to $\mathcal{P}([\![1,m]\!])$ (i.e., the power set of $[\![1,m]\!]$), where $\psi(\ell)$ determines the subset of illegal resources allocated to location $\ell \in L$ without exceeding its capacity. The set of feasible allocations for the adversary is given as follows:

$$\mathcal{A}_A := \{\psi \in \mathcal{P}([\![1,m]\!])^L \mid \bigcup_{\ell \in L} \psi(\ell) = [\![1,m]\!],\ |\psi(\ell)| \leq c_\ell \ \ \forall\, \ell \in L, \text{ and } \psi(\ell) \cap \psi(k) = \emptyset \ \ \forall\, \ell \neq k \in L\}.$$

In this model, we assume detection is independent across resources, and if an illegal resource is detected by the inspector, it is immediately captured. Thus, the inspector (resp. attacker) seeks to maximize (resp. minimize) the total expected damage value of detected illegal resources, which for any player type profile $(\theta_I, \theta_A) \in \Theta_I \times \Theta_A$ and any action profile $(\varphi, \psi) \in \mathcal{A}_I \times \mathcal{A}_A$, is given by

$$u(\varphi, \psi \mid \theta_I, \theta_A) := \sum_{\ell \in L} d^{\theta_I}_{\varphi(\ell)} \cdot \sum_{j \in \psi(\ell)} v^{\theta_A}_j. \tag{1}$$

We illustrate a realization of the game with the following example:

EXAMPLE 1. Consider a system consisting of 5 locations of capacities $(c_1, c_2, c_3, c_4, c_5) = (5,4,4,2,1)$. The inspector of type θ_I has 5 inspection resources with detection rates given by $(d^{\theta_I}_1, d^{\theta_I}_2, d^{\theta_I}_3, d^{\theta_I}_4, d^{\theta_I}_5) = (0.9, 0.8, 0.8, 0.55, 0)$ (i.e., the fifth resource is fictitious). On the other hand, the adversary of type θ_A has 16 illegal resources of damage values $v^{\theta_A}_1 = v^{\theta_A}_2 = v^{\theta_A}_3 = 1$, $v^{\theta_A}_4 = \cdots = v^{\theta_A}_7 = 4$, $v^{\theta_A}_8 = \cdots = v^{\theta_A}_{12} = 9$, $v^{\theta_A}_{13} = v^{\theta_A}_{14} = 10$, and $v^{\theta_A}_{15} = v^{\theta_A}_{16} = 12$. We suppose the players allocate their resources as in Figure 1.

In this example, the inspector's action $\varphi \in \mathcal{A}_I$ is given by $\varphi(1) = 1$, $\varphi(2) = 4$, $\varphi(3) = 2$, $\varphi(4) = 3$, and $\varphi(5) = 5$, and the adversary's action $\psi \in \mathcal{A}_A$ is given by $\psi(1) = \{1,8,9,10,15\}$, $\psi(2) = \{2,3,4,5\}$, $\psi(3) = \{11,13,14,16\}$, $\psi(4) = \{6,12\}$, and $\psi(5) = \{7\}$. We observe that the expected value of the detected illegal resources is given by

$$\begin{aligned} u(\varphi, \psi \mid \theta_I, \theta_A) &= 0.9 \cdot (1+9+9+9+12) + 0.55 \cdot (1+1+4+4) \\ &\quad + 0.8 \cdot (9+10+10+12) + 0.8 \cdot (4+9) + 0 \cdot 4 \\ &= 84.7. \end{aligned}$$

△

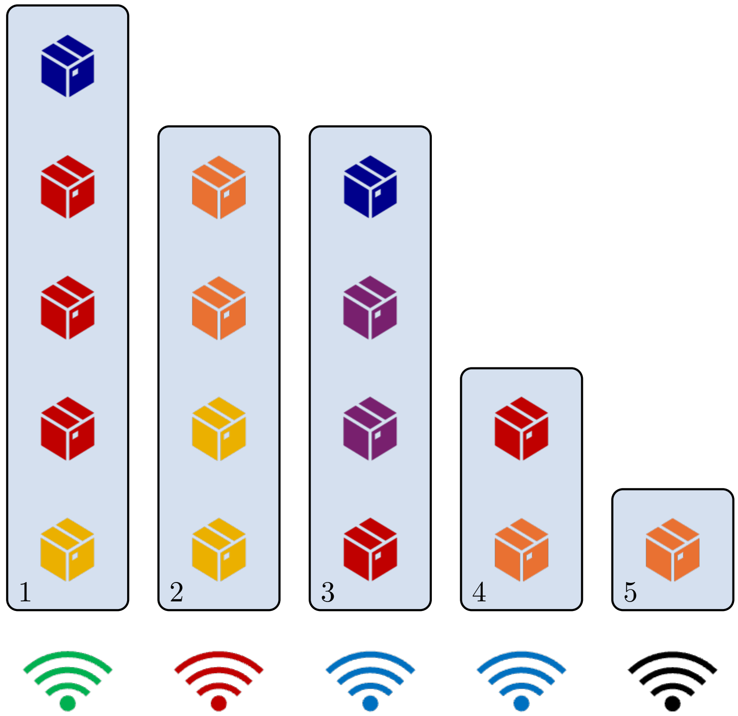

Figure 1 Game instance. The adversary assigns illegal resources as follows: one with damage value 1 (yellow), three with damage value 9 (red) and one with damage value 12 to location 1 (blue), two with damage value 1 and two with damage value 4 to location 2, one with damage value 9, two with damage value 10 (purple) and one with damage value 12 to location 3, one with damage value 4 and one with damage value 9 to location 4, and one with damage value 4 to location 5. The inspector deploys an inspection resource with detection rate 0.9 (green) at location 1, one with detection rate 0.55 (red) at location 2, ones with detection rates 0.8 (blue) at locations 3 and 4, and a fictitious resource with detection rate 0 (black) at location 5.

In such inspection problems, players often benefit from randomizing their actions (Pita et al. 2008), which, in practice, is achieved through a randomized schedule of operations. Therefore, players select *mixed strategies*, defined as type-dependent probability distributions over their action sets (Myerson 1991). Let $\Delta(\mathcal{A}_I)$ and $\Delta(\mathcal{A}_A)$ denote the sets of probability distributions over the players' respective action sets. Then, the sets of mixed strategies for the inspector and adversary are respectively given by

$$\Delta_I = \{\sigma_I \in [0,1]^{\mathcal{A}_I \times \Theta_I} \mid \sigma_I(\cdot \mid \theta_I) \in \Delta(\mathcal{A}_I) \;\; \forall\, \theta_I \in \Theta_I\}$$
$$\Delta_A = \{\sigma_A \in [0,1]^{\mathcal{A}_A \times \Theta_A} \mid \sigma_A(\cdot \mid \theta_A) \in \Delta(\mathcal{A}_A) \;\; \forall\, \theta_A \in \Theta_A\},$$

where $\sigma_I(\cdot \mid \theta_I)$ represents the probability that the inspector selects each action $\varphi \in \mathcal{A}_I$ when of type $\theta_I \in \Theta_I$. Similarly, $\sigma_A(\cdot \mid \theta_A)$ denotes the probability that the adversary selects each action $\psi \in \mathcal{A}_A$ when their type is $\theta_A \in \Theta_A$. Given a strategy profile $(\sigma_I, \sigma_A) \in \Delta_I \times \Delta_A$, we extend the notation in (1) to define the players' expected payoff as follows:

$$u(\sigma_I, \sigma_A \mid \pi_I, \pi_A) := \mathbb{E}_{(\theta_I,\theta_A)\sim(\pi_I,\pi_A)}[\mathbb{E}_{(\varphi,\psi)\sim(\sigma_I(\cdot\mid\theta_I),\sigma_A(\cdot\mid\theta_A))}[u(\varphi,\psi \mid \theta_I,\theta_A)]].$$

In inspection settings, players typically cannot observe their opponent's realized action before selecting their own strategy. For instance, when the adversary hides their smuggled goods within containers, they rarely know the inspection plan implemented the day the containers go through customs (De Wulf and Sokol 2005). This results in "simultaneous" selections of player strategies, and the solution concept is given by *Nash equilibria*. For the game of incomplete information Γ, a strategy profile $(\sigma_I^*, \sigma_A^*) \in \Delta_I \times \Delta_A$ is a Nash equilibrium (NE) if

$$\forall\, \theta_I \in \Theta_I,\ \sigma_I^*(\cdot\,|\,\theta_I) \in \operatorname*{arg\,max}_{\sigma_I(\cdot\,|\,\theta_I)\in\Delta(\mathcal{A}_I)} \mathbb{E}_{\theta_A\sim\pi_A}[\mathbb{E}_{(\varphi,\psi)\sim(\sigma_I(\cdot\,|\,\theta_I),\sigma_A^*(\cdot\,|\,\theta_A))}[u(\varphi,\psi\,|\,\theta_I,\theta_A)]]$$
$$\forall\, \theta_A \in \Theta_A,\ \sigma_A^*(\cdot\,|\,\theta_A) \in \operatorname*{arg\,min}_{\sigma_A(\cdot\,|\,\theta_A)\in\Delta(\mathcal{A}_A)} \mathbb{E}_{\theta_I\sim\pi_I}[\mathbb{E}_{(\varphi,\psi)\sim(\sigma_I^*(\cdot\,|\,\theta_I),\sigma_A(\cdot\,|\,\theta_A))}[u(\varphi,\psi\,|\,\theta_I,\theta_A)]].$$

In other words, in equilibrium, each inspector type θ_I must select a probability distribution $\sigma_I^*(\cdot\,|\,\theta_I)$ that is a best response to σ_A^*, given the incomplete adversary information given by the probability distribution π_A. Similarly in equilibrium, each adversary type θ_A selects a probability distribution $\sigma_A^*(\cdot\,|\,\theta_A)$ that is a best response to the inspector's mixed strategy σ_I^*, given the incomplete inspector information given by the probability distribution π_I. We refer to $u(\sigma_I^*, \sigma_A^*\,|\,\pi_I, \pi_A)$ as the *value* of the game Γ.

Since Γ is a zero-sum game with a finite number of player types and actions, an NE is guaranteed to exist and the value of the game is unique. Furthermore, NE are given by strategy profiles $(\sigma_I^*, \sigma_A^*) \in \Delta_I \times \Delta_A$ where σ_I^* and σ_A^* are optimal primal and dual solutions to the following linear program:

$$\max_{\sigma_I\in\Delta_I} \mathbb{E}_{\theta_A\sim\pi_A}[\min_{\psi_{\theta_A}\in\mathcal{A}_A} u(\sigma_I, \psi_{\theta_A}\,|\,\pi_I, \theta_A)]. \tag{LP$_\Gamma$}$$

Notably, (LP$_\Gamma$) represents the equivalent Stackelberg of incomplete information, where the inspector first selects their inspection strategy $\sigma_I \in \Delta_I$, and then the adversary selects an allocation plan $\psi_{\theta_A} \in \mathcal{A}_A$ based on their type θ_A after observing σ_I. This is particularly relevant in settings where the adversary can record historical inspection operations.

However, the combinatorial structure of the players' action sets renders (LP$_\Gamma$) computationally challenging, even for very small instances. Similarly, algorithms for computing approximate equilibria (e.g. Freund and Schapire 1999, Lipton et al. 2003, Hellerstein et al. 2019) are also inapplicable for realistic instances of the game Γ. Instead, we develop three algorithms for (i) computing the equilibrium expected damage value in each location; (ii) coordinating the adversary's heterogeneous illegal resources to reach these target values; and (iii) constructing an equilibrium inspection strategy in response to the adversary's allocation strategy. This results in a polynomial-time approach for solving Γ.

3. Characterization and Computation of Equilibrium Strategies

3.1. Preliminary Analysis

To facilitate the equilibrium analysis of the incomplete-information game Γ, we introduce the following notations: For any inspection strategy $\sigma_I \in \Delta_I$, we denote by $\zeta_\ell(\sigma_I) := \mathbb{E}_{\theta_I\sim\pi_I}[\mathbb{E}_{\varphi\sim\sigma_I(\cdot\,|\,\theta_I)}[d^{\theta_I}_{\varphi(\ell)}]]$ the *detection probability* of each location $\ell \in [\![1, n]\!]$. This quantity is relevant to the adversary, as their incentive is to allocate illegal resources with higher damage values to locations with lower detection probabilities. Analogously, for any allocation strategy $\sigma_A \in \Delta_A$, we denote by $\xi_\ell(\sigma_A) = \mathbb{E}_{\theta_A \sim \pi_A}[\mathbb{E}_{\psi \sim \sigma_A(\cdot\,|\,\theta_A)}[\sum_{j \in \psi(\ell)} v_j^{\theta_A}]]$ the *expected damage value* in each location $\ell \in [\![1, n]\!]$. This quantity is important for the inspector, as they must allocate inspection resources with higher detection rates to locations with higher expected damage values to maximize their expected payoff.

Henceforth, we assume that the locations are indexed in nonincreasing order of their capacities ($c_1 \geq \cdots \geq c_n$), the inspection resources of each inspector type $\theta_I \in \Theta_I$ are indexed in nonincreasing order of their detection rates ($d_1^{\theta_I} \geq \cdots \geq d_n^{\theta_I}$), and the illegal resources of each adversary type $\theta_A \in \Theta_A$ are indexed in nondecreasing order of their damage values ($v_1^{\theta_A} \leq \cdots \leq v_m^{\theta_A}$). Furthermore, we denote as $\varphi^0 \in \mathcal{A}_I$ the inspection plan defined by $\varphi^0(\ell) = \ell$ for every $\ell \in [\![1, n]\!]$, that is, that allocates the ℓ-th inspection resource to the ℓ-th location. We also denote as $\psi^0 \in \mathcal{A}_A$ the allocation plan defined by $\psi^0(\ell) = \{j + \sum_{i=1}^{\ell-1} c_i,\ \forall\, j \in [\![1, c_\ell]\!]\}$ for every $\ell \in [\![1, n]\!]$, that is, that allocates the first c_1 illegal resources to location 1, the next c_2 illegal resources to location 2, and so on.

To build the intuition regarding the equilibria of the game Γ, we first consider the special case where players have full information (i.e., $|\Theta_I| = |\Theta_A| = 1$) and the adversary's resources satisfy $\sum_{j \in \psi^0(1)} v_j^{\theta_A} > \cdots > \sum_{j \in \psi^0(n)} v_j^{\theta_A}$. Namely, if the adversary selects the allocation plan ψ^0, then the expected damage value at each location decreases with the location index. In this scenario, we can show that (φ^0, ψ^0) is an NE of Γ. Indeed, the inspector's best inspection resources (i.e., with highest detection rates) are allocated to locations containing the largest expected damage values, and swapping any two inspection resources will result in a decrease of expected detected damage value. Furthermore, since the adversary's most valuable resources (i.e., with highest damage values) are allocated to locations with lowest detection probabilities, swapping any pair of illegal resources will increase the expected detected damage value.

From the analysis of this special case, we observe the equilibrium is ensured from the decreasing expected damage values—achieved with an allocation plan that assigns illegal resources with highest damage values to locations with smallest capacities—and the decreasing detection probabilities—achieved with an inspection plan that assigns inspection resources with highest detection rates to locations with highest capacities. However in general, these conditions do not hold with the inspection plan φ^0 and the allocation plan ψ^0, and the players must account for the uncertainty they face pertaining to their opponent. Still, we will determine how we can recover similar conditions by carefully randomizing the resources for each player (and each player type), while managing the challenges arising from the heterogeneous location capacities, detection rates, and damage values.

3.2. Equilibrium Properties and Computation

We first focus on the adversary's equilibrium behavior in any instance of the incomplete-information game Γ. Our preliminary analysis in Section 3.1 suggests that the adversary should select an allocation strategy that randomizes their illegal resources in such a way that the expected damage values of the locations with higher capacities are larger, while allocating the resources with higher damage values within locations of smaller capacities. However, it is unclear if such a randomization is feasible, given the locations' capacity constraints and the illegal resources' damage values for each adversary type.

In fact, we find that the adversary's equilibrium strategy requires partitioning the set of locations based on the average expected damage value that can be achieved within each subset. Specifically, the first subset of locations is determined by the largest location $k \in [\![1,n]\!]$ that maximizes $\frac{1}{k}\sum_{\ell=1}^{k}\sum_{s\in\psi^0(\ell)}\mathbb{E}_{\theta_A\sim\pi_A}[v_s^{\theta_A}]$, namely, the average expected damage value within $[\![1,k]\!]$ if each adversary type allocates illegal resources according to ψ^0. The resulting location is denoted k_1^*, and referred to as a *demarcation location*. Furthermore, we will interpret $\bar{\xi}_0 := \frac{1}{k_1^*}\sum_{\ell=1}^{k_1^*}\sum_{s\in\psi^0(\ell)}\mathbb{E}_{\theta_A\sim\pi_A}[v_s^{\theta_A}]$ as the target expected damage value to be achieved by the adversary's strategy at each location in *zone* $[\![1,k_1^*]\!]$. We then repeat the same process with the remaining locations $[\![k_1^*+1,n]\!]$, until a partition is created. The overall partitioning is detailed in Algorithm 1.

Algorithm 1: Computation of Demarcation Locations and Target Expected Damage Values

1 Initialize $k_0^* \leftarrow 0, \quad i \leftarrow 0$

2 **while** $k_i^* < n$ **do**

3 $\quad k_{i+1}^* \leftarrow \max\{\arg\max\{\frac{1}{k-k_i^*}\sum_{\ell=k_i^*+1}^{k}\sum_{s\in\psi^0(\ell)}\mathbb{E}_{\theta_A\sim\pi_A}[v_s^{\theta_A}],\ \forall\, k\in[\![k_i^*+1,n]\!]\}\}$

4 $\quad \bar{\xi}_i \leftarrow \frac{1}{k_{i+1}^*-k_i^*}\sum_{\ell=k_i^*+1}^{k_{i+1}^*}\sum_{s\in\psi^0(\ell)}\mathbb{E}_{\theta_A\sim\pi_A}[v_s^{\theta_A}]$

5 $\quad i \leftarrow i+1$

6 **return** $(k_0^*,\ldots,k_i^*)$, $(\bar{\xi}_0,\ldots,\bar{\xi}_{i-1})$

At termination, Algorithm 1 returns $r+1$ demarcation locations that satisfy $0 = k_0^* < \cdots < k_r^* = n$ and partition the set of locations into r zones $[\![k_0^*+1,k_1^*]\!],\ldots,[\![k_{r-1}^*+1,k_r^*]\!]$. The algorithm also associates each zone $[\![k_i^*+1,k_{i+1}^*]\!]$ (for $i\in[\![0,r-1]\!]$) with a target expected damage value $\bar{\xi}_i$, which captures the interplay between the location capacities and the distribution of the illegal resources' damage values across adversary types. In the next lemma, we show key properties satisfied by the demarcation locations and target expected damage values:

LEMMA 1. *The target expected damage values are decreasing:*

$$\forall\, i\in[\![1,r-1]\!],\ \bar{\xi}_{i-1} > \bar{\xi}_i. \tag{2}$$

The distribution of the illegal resources' damage values satisfies the following:

$$\forall\, i\in[\![0,r-1]\!],\ \forall\, k\in[\![k_i^*+1,k_{i+1}^*-1]\!],\ \bar{\xi}_i \leq \frac{1}{k_{i+1}^*-k}\sum_{\ell=k+1}^{k_{i+1}^*}\sum_{s\in\psi^0(\ell)}\mathbb{E}_{\theta_A\sim\pi_A}[v_s^{\theta_A}]. \tag{3}$$

From Lemma 1, we find that the target expected damage values associated with each zone are decreasing. Furthermore, (3) proves that in each zone $[\![k_i^*+1, k_{i+1}^*]\!]$, and for every $k \in [\![k_i^*+1, k_{i+1}^*-1]\!]$, the expected damage value, averaged over the last k_{i+1}^*-k locations in that zone, is no less than the zone's target expected damage value. This will play a critical role in proving that the distributions of location capacities and illegal resources' damage values permit the computation of an equilibrium mixed strategy for the adversary that achieves the target expected damage values determined by Algorithm 1.

Building upon the game-theoretic intuition from the preliminary analysis in Section 3.1, we aim to answer the following question: *Does there exist an allocation strategy that achieves the target expected damage values returned by Algorithm 1 by only allocating the least valuable illegal resources in the first zone, the next least valuable resources in the second zone, and so on?* We design a combinatorial algorithm to positively resolve this question.

We initialize the algorithm with a strategy $\sigma_A \in \Delta_A$ that selects the allocation plan ψ^0 for each adversary type. Then, at each iteration of the algorithm, we identify two locations j^- and j^+ within a same zone $[\![k_i^*+1, k_{i+1}^*]\!]$ such that $\xi_{j^-}(\sigma_A) < \bar{\xi}_i$ and $\xi_{j^+}(\sigma_A) > \bar{\xi}_i$. We also identify an adversary type $\theta'_A \in \Theta_A$ and an allocation plan ψ' in the support of $\sigma_A(\cdot \,|\, \theta'_A)$, denoted as $\text{supp}(\sigma_A(\cdot \,|\, \theta'_A))$, such that swapping a subset of illegal resources allocated in locations j^- and j^+ by ψ' increases $\xi_{j^-}(\sigma_A)$ and decreases $\xi_{j^+}(\sigma_A)$. The algorithm terminates when the locations' expected damage values meet their target values. However, the algorithm must be carefully designed to ensure that at each iteration, there is always a possibility to swap resources allocated by an allocation plan selected with positive probability by an adversary type, until termination. We refer the reader to Algorithm 2 for the detailed implementation. We also extend the support notation to represent the support of a mixed strategy $\sigma_A \in \Delta_A$, defined as $\text{supp}(\sigma_A) = \cup_{\theta_A \in \Theta_A} \{\psi \in \mathcal{A}_A \mid \sigma_A(\psi \,|\, \theta_A) > 0\}$.

From Algorithm 2, we observe that for each zone $[\![k_i^*+1, k_{i+1}^*]\!]$ and at each iteration of the outer while loop 3-10, j^- is selected to be the largest location in that zone with an expected damage value strictly less than its target value, and j^+ is selected to be the lowest location that is greater than j^- with an expected damage value strictly larger than its target value. Then, the inner while loop 6-10 iteratively transfers probabilities between allocation plans until the target expected damage value is reached at j^- or j^+. Specifically, the algorithm selects an adversary type θ'_A, an allocation plan ψ' in the support of $\sigma_A(\cdot \,|\, \theta'_A)$, and a subset S of c_{j^+} illegal resources allocated by ψ' to location j^- such that the total damage value in S is strictly less than the damage value allocated by ψ' to location j^+—we note that location capacities satisfy $c_{j^-} \geq c_{j^+}$ since $j^- < j^+$. The algorithm then creates a new allocation plan $\psi^\dagger$ by interchanging the resources in S allocated by ψ' in location j^- with the ones allocated in location j^+. A new mixed strategy is then constructed by reassigning as much probability p from ψ' to $\psi^\dagger$ until ψ' leaves the support of $\sigma_A(\cdot \,|\, \theta'_A)$ when p reaches $\sigma_A(\psi' \,|\, \theta'_A)$, or until the target expected damage value $\bar{\xi}_i$ is achieved at location j^- (resp. j^+) when p reaches

Algorithm 2: Computation of Adversary's Equilibrium Strategy

1 Initialize $\sigma_A \in \Delta_A$ by setting $\sigma_A(\psi^0 \,|\, \theta_A) \leftarrow 1 \;\; \forall\, \theta_A \in \Theta_A$

2 **for** $i \in [\![0, r-1]\!]$ **do**

3   **while** $\exists\, \ell \in [\![k_i^* + 1, k_{i+1}^*]\!] \mid \xi_\ell(\sigma_A) \neq \bar{\xi}_i$ **do**

4         $j^- \leftarrow \max\{\ell \in [\![k_i^* + 1, k_{i+1}^*]\!] \mid \xi_\ell(\sigma_A) < \bar{\xi}_i\}$

5         $j^+ \leftarrow \min\{\ell \in [\![j^- + 1, k_{i+1}^*]\!] \mid \xi_\ell(\sigma_A) > \bar{\xi}_i\}$

6   **while** $\xi_{j^-}(\sigma_A) < \bar{\xi}_i$ **and** $\xi_{j^+}(\sigma_A) > \bar{\xi}_i$ **do**

7   Select $\theta'_A \in \Theta_A$, $\psi' \in \operatorname{supp}(\sigma_A(\cdot \,|\, \theta'_A))$, and $S \subseteq \psi'(j^-)$, such that $|S| = c_{j^+}$ and $\sum_{s \in S} v_s^{\theta'_A} < \sum_{s \in \psi'(j^+)} v_s^{\theta'_A}$

8   Construct $\psi^\dagger \in \mathcal{A}_A$ defined by

$$\psi^\dagger(\ell) = \begin{cases} \psi'(j^-) \setminus S \cup \psi'(j^+) & \text{if } \ell = j^- \\ S & \text{if } \ell = j^+ \\ \psi'(\ell) & \text{otherwise} \end{cases}$$

9             $p \leftarrow \min\left\{ \sigma_A(\psi' \,|\, \theta'_A), \dfrac{\bar{\xi}_i - \xi_{j^-}(\sigma_A)}{\pi_A(\theta'_A) \cdot (\sum_{s \in \psi'(j^+)} v_s^{\theta'_A} - \sum_{s \in S} v_s^{\theta'_A})}, \dfrac{\xi_{j^+}(\sigma_A) - \bar{\xi}_i}{\pi_A(\theta'_A) \cdot (\sum_{s \in \psi'(j^+)} v_s^{\theta'_A} - \sum_{s \in S} v_s^{\theta'_A})} \right\}$

10            $\sigma_A(\psi' \,|\, \theta'_A) \leftarrow \sigma_A(\psi' \,|\, \theta'_A) - p, \quad \sigma_A(\psi^\dagger \,|\, \theta'_A) \leftarrow p$

11 **return** σ_A

$(\bar{\xi}_i - \xi_{j^-}(\sigma_A))/(\pi_A(\theta'_A) \cdot (\sum_{s \in \psi'(j^+)} v_s^{\theta'_A} - \sum_{s \in S} v_s^{\theta'_A}))$ (resp. $(\xi_{j^+}(\sigma_A) - \bar{\xi}_i)/(\pi_A(\theta'_A) \cdot (\sum_{s \in \psi'(j^+)} v_s^{\theta'_A} - \sum_{s \in S} v_s^{\theta'_A}))$). Notably, the careful definition of j^- and j^+ maintains some ordering of the illegal resources' damage values, initially allocated in a nondecreasing order by ψ^0 for each adversary type. This is critical to guaranteeing the validity and termination of the algorithm, as shown in the next proposition:

PROPOSITION 1. *Algorithm 2 is well-defined. In particular, at the start of each iteration of the outer while loop 3-10, $\sigma_A \in \Delta_A$ and satisfies the following properties:*

(i) Each allocation plan selected with positive probability by any adversary type allocates the illegal resources with lowest damage values to the first zone, the ones with the next lowest damage values to the second zone, and so on:

$$\forall\, \psi \in \operatorname{supp}(\sigma_A),\; \forall\, i \in [\![0, r-1]\!], \bigcup_{\ell \in [\![k_i^*+1, k_{i+1}^*]\!]} \psi(\ell) = [\![1 + \sum_{j=1}^{k_i^*} c_j, \sum_{j=1}^{k_{i+1}^*} c_j]\!]. \tag{4}$$

(ii) Considering the subset of locations for which the target damage values are not met, each allocation plan selected with positive probability by any adversary type allocates illegal resources with damage

values that are nondecreasing with the location index. Specifically, for every $i \in [\![0, r-1]\!]$ *and every* $(\ell_1, \ell_2) \in [\![k_i^* + 1, k_{i+1}^*]\!]^2$ *such that* $\ell_1 < \ell_2$, $\xi_{\ell_1}(\sigma_A) \neq \bar{\xi}_i$, *and* $\xi_{\ell_2}(\sigma_A) \neq \bar{\xi}_i$, *we have the following:*

$$\forall\, \theta_A \in \Theta_A,\ \forall\, \psi \in \mathrm{supp}(\sigma_A(\cdot \,|\, \theta_A)),\ \forall\, s_1 \in \psi(\ell_1),\ \forall\, s_2 \in \psi(\ell_2),\ v_{s_1}^{\theta_A} \leq v_{s_2}^{\theta_A}. \tag{5}$$

(iii) The expected damage value, averaged over the last locations of a given zone, is at least as large as the target expected damage value of that zone:

$$\forall\, i \in [\![0, r-1]\!],\ \forall\, k \in [\![k_i^* + 1, k_{i+1}^*]\!],\ \bar{\xi}_i \leq \frac{1}{k_{i+1}^* - k} \sum_{\ell=k+1}^{k_{i+1}^*} \xi_\ell(\sigma_A). \tag{6}$$

(iv) The locations j^- *and* j^+ *are well-defined, and at the start of each iteration of the inner while loop 6-10, line 7 is well defined, namely:*

$$\exists\, \theta'_A \in \Theta_A,\ \exists\, \psi' \in \mathrm{supp}(\sigma_A(\cdot \,|\, \theta'_A)),\ \exists\, S \subseteq \psi'(j^-) \mid |S| = c_{j^+} \text{ and } \sum_{s \in S} v_s^{\theta'_A} < \sum_{s \in \psi'(j^+)} v_s^{\theta'_A}. \tag{7}$$

Furthermore, Algorithm 2 terminates in $O(mn(n + |\Theta_A|))$ *time, and returns a mixed strategy satisfying the target expected damage values from Algorithm 1 with support of size at most* $n - r + |\Theta_A|$.

To show that Algorithm 2 is well-defined, Proposition 1 requires proving by induction that properties $(i)-(iv)$ are satisfied at the start of each iteration. Importantly, the proof shows that by selecting j^- and j^+ as in Lines 4-5, the illegal resources' damage values allocated by each allocation plan with positive probability satisfy the nondecreasing condition (5). This in turn guarantees the validity of the inner while loop 6-10, namely, the algorithm will be able to swap illegal resources between j^- and j^+ until the target expected damage value is met at either location. Furthermore, this process maintains condition (6), which ensures that j^+ is well-defined in the next iteration.

From Proposition 1, we obtain that Algorithm 2 feasibly constructs an allocation strategy that achieves the target expected damage values returned by Algorithm 1 by only allocating the least valuable illegal resources in the first zone, the next least valuable resources in the second zone, and so on. The algorithm runs in polynomial time, and returns a mixed strategy with linear support. This small-support characteristic is important, as it implies that such a mixed strategy can be practically implemented by an adversary.

EXAMPLE 2. We illustrate Algorithm 2 using the game instance defined in Example 1, consisting of 5 locations, 5 inspection resources (including one that is fictitious), 16 illegal resources, and 1 type for each player. Upon running Algorithm 1, the demarcation locations are calculated as $k_0^* = 0$, $k_1^* = 3$, $k_2^* = 4$, and $k_3^* = 5$, with the target expected damage values being $\bar{\xi}_0 = 24.\bar{6}$ in the first zone $[\![1, 3]\!]$, $\bar{\xi}_1 = 22$ in the second zone $\{4\}$, and $\bar{\xi}_2 = 12$ in the third zone $\{5\}$. Figure 2 displays the adversary's mixed strategy resulting from Algorithm 2. We observe that its support contains 3 allocation plans, which is consistent with the bound established in Proposition 1. △

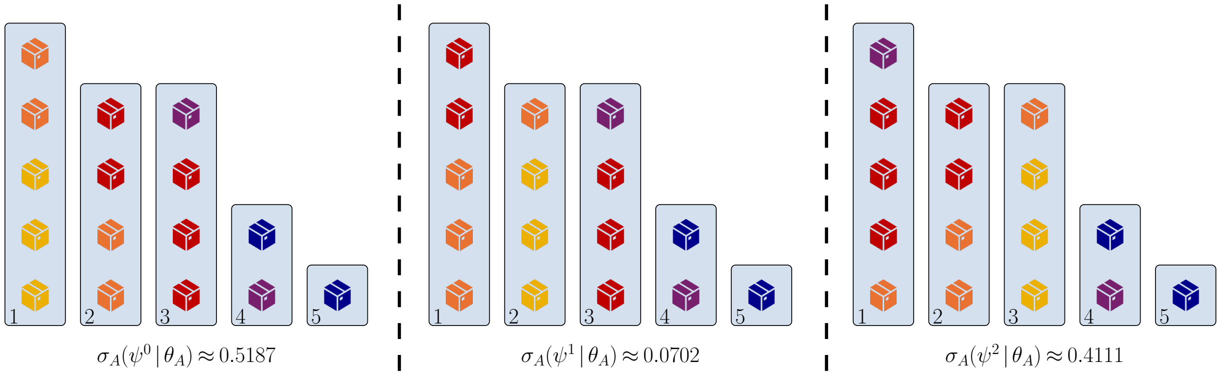

Figure 2 Adversary's mixed strategy as computed by Algorithm 2. Each section illustrates an allocation plan in its support.

Next, we aim to build the inspector's equilibrium strategy, which will be a best response to the adversary's allocation strategy constructed by Algorithm 2. Given the preliminary analysis in Section 3.1, (2) implies that each inspector type must allocate their best inspection resources to the first zones. Furthermore, to ensure that the adversary has no incentive to deviate from their strategy, the inspector must allocate their resources so that the detection probability of every location within each zone is identical.

Thus, we design an algorithm to create an inspection strategy that satisfies such properties. The algorithm cycles the resources of each inspector type over locations within each zone. Consequently in each zone $[\![k_i^*+1, k_{i+1}^*]\!]$ (for $i \in [\![0, r-1]\!]$), each resource $\ell \in [\![k_i^*+1, k_{i+1}^*]\!]$ is allocated uniformly across all locations within that zone by every inspector type. The challenge lies in efficiently constructing an inspection strategy with small support that cycles resources for each zone. In Algorithm 3, we achieve this feat by sorting in advance the cumulative probabilities associated with each cycle within each zone.

Algorithm 3 constructs a mixed strategy such that every inspector type implements the same probability distribution ρ. Lines 2-3 gather and order the cumulative distribution function values λ_j of the desired distribution ρ. In addition, Line 4 keeps track of the zones in which the cycling of resources changes at each value of the cumulative distribution function. Then, the algorithm iteratively constructs the distribution ρ: It first constructs the inspection plan φ^0, which allocates the ℓ-th inspection resource to the ℓ-th location, and assigns probability $\lambda_2 - \lambda_1$. Then, the algorithm shifts the allocation of resources within the zones characterized in I_2, and assigns probability $\lambda_3 - \lambda_2$ to the new inspection plan. The algorithm terminates when ρ is a probability distribution, returning the inspector's mixed strategy σ_I. In the next proposition, we derive the runtime of Algorithm 3, as well as the properties of the inspection strategy it constructs.

PROPOSITION 2. *Algorithm 3 terminates in $O(n(n+|\Theta_I|))$ time and returns a mixed strategy $\sigma_I \in \Delta_I$ that satisfies the following properties:*

Algorithm 3: Computation of Inspector's Equilibrium Strategy

Initialize $\rho \in \mathbf{0}_{\mathcal{A}_I}$

$\Lambda \leftarrow \bigcup_{i \in [\![0,r-1]\!]} \bigcup_{k \in [\![0,k^*_{i+1}-k^*_i]\!]} \{\frac{k}{k^*_{i+1}-k^*_i}\}$

Sort $\Lambda = (\lambda_1, \ldots, \lambda_{|\Lambda|})$ in increasing order of its values

$I_j \leftarrow \{i \in [\![0,r-1]\!] \mid \lambda_j = \frac{k}{k^*_{i+1}-k^*_i} \text{ for some } k \in [\![0, k^*_{i+1}-k^*_i]\!]\},\ \forall\, j \in [\![1,|\Lambda|]\!]$

$s_i \leftarrow 0,\ \forall\, i \in [\![0,r-1]\!]$

for $j \in [\![2,|\Lambda|]\!]$ **do**

Construct $\varphi^\dagger \in \mathcal{A}_I$ defined by

$$\forall\, i \in [\![0,r-1]\!],\ \forall\, \ell \in [\![k^*_i+1, k^*_{i+1}]\!],\ \varphi^\dagger(\ell) = \begin{cases} \ell + s_i & \text{if } \ell \in [\![k^*_i+1, k^*_{i+1}-s_i]\!] \\ \ell + s_i - (k^*_{i+1}-k^*_i) & \text{if } \ell \in [\![k^*_{i+1}-s_i+1, k^*_{i+1}]\!] \end{cases}$$

$\rho(\varphi^\dagger) \leftarrow \lambda_j - \lambda_{j-1}$

$s_i \leftarrow s_i + 1,\ \ \forall\, i \in I_j$

$\sigma_I(\cdot \mid \theta_I) \leftarrow \rho,\ \ \forall\, \theta_I \in \Theta_I$

return σ_I

(i) Each inspection plan selected with positive probability allocates the resources with highest detection rates to the first zone, the ones with the next highest detection rates to the second zone, and so on:

$$\forall\, \varphi \in \text{supp}(\sigma_I),\ \forall\, i \in [\![0,r-1]\!],\ \bigcup_{\ell \in [\![k^*_i+1, k^*_{i+1}]\!]} \varphi(\ell) = [\![k^*_i+1, k^*_{i+1}]\!]. \tag{8}$$

(ii) The detection probability of each location is given by

$$\forall\, i \in [\![0,r-1]\!],\ \forall\, \ell \in [\![k^*_i+1, k^*_{i+1}]\!],\ \zeta_\ell(\sigma_I) = \frac{1}{k^*_{i+1}-k^*_i} \sum_{k=k^*_i+1}^{k^*_{i+1}} \mathbb{E}_{\theta_I \sim \pi_I}[d_k^{\theta_I}] =: \bar{\zeta}_i. \tag{9}$$

(iii) Each inspector type implements an identical probability distribution with support of size at most $n - r + 1$.

From Proposition 2, we find that Algorithm 3 efficiently constructs an inspection strategy σ_I that allocates the best inspection resources to the first zones and is such that the detection probability of every location within zone $[\![k^*_i+1, k^*_{i+1}]\!]$ (for $i \in [\![0, r-1]\!]$) is identical and equal to $\bar{\zeta}_i$. This is achieved by having every inspector type allocating each resource $\ell \in [\![k^*_i+1, k^*_{i+1}]\!]$ uniformly with probability $1/(k^*_{i+1}-k^*_i)$ across locations in zone $[\![k^*_i+1, k^*_{i+1}]\!]$. Notably, the inspection strategy σ_I is such that each inspector type implements the same probability distribution $\rho \in \Delta(\mathcal{A}_I)$, which is itself computed in $O(n^2)$ time, has a linear-size support, and can be easily implemented by the inspector in practice.

EXAMPLE 3. We illustrate Algorithm 3 using the game instance defined in Example 1, for which Algorithm 1 returns the demarcation locations $k^*_0 = 0$, $k^*_1 = 3$, $k^*_2 = 4$, and $k^*_3 = 5$. Figure 3 shows the inspector's mixed strategy produced by Algorithm 3. It cycles the three best resources within the first zone, deterministically allocates the fourth resource in the second zone, and does not inspect the third zone (represented by the fictitious resource). The support of the inspector's strategy is of size 3, which is consistent with the bound established in Proposition 2. △

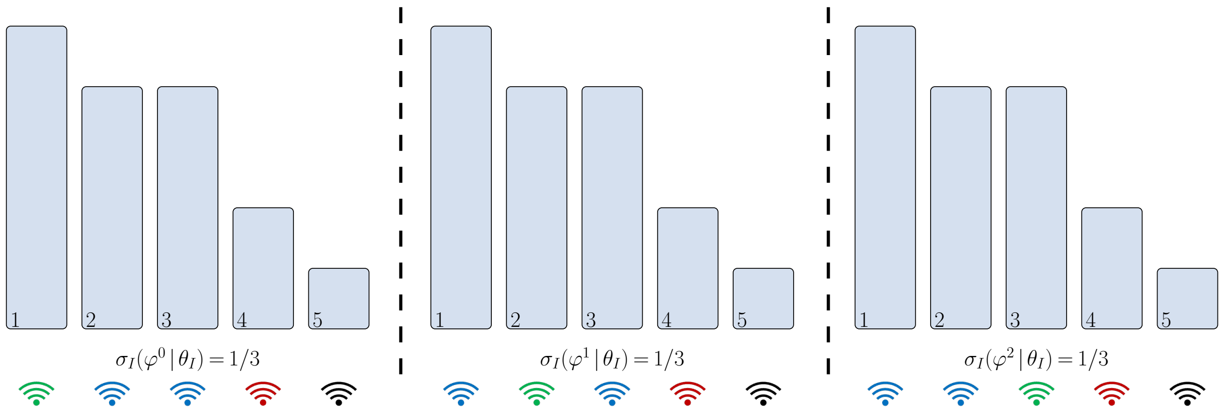

Figure 3 Inspector's mixed strategy as computed by Algorithm 3. Each section illustrates an inspection plan in its support.

Given the properties of the allocation and inspection strategies constructed by Algorithms 2 and 3, respectively, we can now solve the game Γ in the following theorem:

THEOREM 1. *Let $(k_0^*, \ldots, k_r^*)$ be demarcation locations computed from Algorithm 1, $\sigma_A^* \in \Delta_A$ be the allocation strategy computed from Algorithm 2, and $\sigma_I^* \in \Delta_I$ be the inspection strategy computed from Algorithm 3. Then, the strategy profile (σ_I^*, σ_A^*) is a Nash equilibrium of the incomplete-information game Γ. The value of the game is given by*

$$\sum_{i=0}^{r-1} \frac{1}{k_{i+1}^* - k_i^*} \cdot \Big(\sum_{k=k_i^*+1}^{k_{i+1}^*} \mathbb{E}_{\theta_I \sim \pi_I}[d_k^{\theta_I}] \Big) \cdot \Big(\sum_{k=k_i^*+1}^{k_{i+1}^*} \sum_{s \in \psi^0(k)} \mathbb{E}_{\theta_A \sim \pi_A}[v_s^{\theta_A}] \Big) = \sum_{i=0}^{r-1} (k_{i+1}^* - k_i^*) \cdot \bar{\zeta}_i \cdot \bar{\xi}_i, \tag{10}$$

and can be computed in $O(n(|\Theta_I| + m|\Theta_A|))$ time.

In Theorem 1, we show that a Nash equilibrium (σ_I^*, σ_A^*) of the game Γ can be efficiently computed from Algorithms 1-3. From the analytical properties derived in Propositions 1-2, we can characterize the locations' detection probabilities and expected damage values in equilibrium. We recall that these quantities account for the uncertainty players face regarding their opponent's resources. Specifically, given the incomplete information the adversary has regarding the inspector's inspection resources—characterized by the type distribution π_I—the adversary expects that any illegal resource allocated within zone $[\![k_i^* + 1, k_{i+1}^*]\!]$ (for $i \in [\![0, r-1]\!]$) will be detected with probability $\bar{\zeta}_i$. Similarly, given the distribution of adversary types π_A—built for instance from historical data—the inspector expects that the total damage value in each location within zone $[\![k_i^* + 1, k_{i+1}^*]\!]$ is given by $\bar{\xi}_i$.

Overall, we observe that the players' expected payoff in equilibrium primarily depends on the interplay between the distribution of illegal resource damage values, the uncertainty pertaining to the adversary, and the location capacities. It impacts the computation of demarcation locations, which in turn define the zones that dictate where player resources must be coordinated.

Our two-step approach circumvents the combinatorial complexity of building the adversary's equilibrium strategy by first computing the target expected damage values in each location. Then, Algorithm 2 iteratively constructs the adversary's strategy to achieve these target values. Importantly, the properties satisfied by the target expected damage values and demarcation locations computed from Algorithm 1 ensure the feasibility of constructing the adversary's strategy while satisfying game-theoretic conditions, guaranteeing the validity of Algorithm 2. To the best of our knowledge, Algorithm 2 is the first combinatorial algorithm developed in security games for coordinating the allocation of heterogeneous resources in locations with heterogeneous capacities to construct equilibrium strategies.

We end our analysis of the game Γ with properties satisfied by its Nash equilibria constructed with our approach:

COROLLARY 1. *The adversary's equilibrium strategy constructed from Algorithms 1 and 2 does not depend on the information pertaining to the inspector's resources. Similarly, the probability distribution implemented by every inspector type following Algorithms 1 and 3 does not depend on the information available to the adversary pertaining to the inspector's resources.*

From the structure of the Nash equilibrium constructed from Algorithms 1-3, we observe the surprising property that the players' equilibrium strategies solely depend on capacities of the locations and the information pertaining to the adversary (i.e., stochastic types and distribution of illegal resource damage values). The adversary does not need to know the inspection resources available to the inspector, nor their detection rates, to construct an equilibrium allocation strategy. As a result, the incomplete-information game Γ has the interesting property that the inspector has no value in obfuscating the characteristics of the inspection resources they possess.

In the next section, we leverage the equilibrium properties of Γ to determine the optimal acquisition of inspection resources for the inspector.

4. Inspection Resource Acquisition

4.1. Problem Definition

We now consider the tactical-level problem where the inspector aims to acquire inspection resources for best monitoring the system, given the uncertainty they face regarding the adversary who will be utilizing the system to carry out illegal activities. In this setting, the inspector first purchases inspection resources with possibly different characteristics, given a budget and the information they have regarding adversary types. Then, the inspector and adversary interact according to the game Γ.

Formally, we model this problem as a two-stage game of incomplete information. In the first stage, the inspector is given a budget $\hat{b}$ to purchase inspection resources of different types $\mathcal{T}$. Each inspection resource of type $\tau \in \mathcal{T}$ has a monetary cost b_τ and a detection rate d_τ. The inspector's decision can be represented as a feasible acquisition plan in $\mathcal{X} := \{x \in \mathbb{Z}^{\mathcal{T}}_{\geq 0} \mid \sum_{\tau \in \mathcal{T}} b_\tau x_\tau \leq \hat{b}\}$, where x_τ determines the number of acquired inspection resources of type $\tau \in \mathcal{T}$. We assume that the uncertainty pertaining to the adversary types (Θ_A, π_A) is known to the inspector when selecting their inspection resources (e.g., from historical data) and is independent of the inspector's acquisition. We also consider the generic information setting where the adversary's uncertainty regarding the inspector type $(\Theta_I(x), \pi_I(x))$ possibly depends on the inspector's acquisition x. For instance, the adversary may be able to observe the inspector's acquisition, in which case the adversary has full information. Alternatively, they may only have partial information regarding each inspection resource type being purchased, without knowing their exact detection rates. Given an acquisition x, we denote by $\theta_I^0(x) \in \Theta_I(x)$ the true inspector type (possibly unknown to the adversary).

In the second stage, the players interact according to the inspection game $\Gamma(x)$, which depends on the inspector's acquisition x. Given the players' equilibrium strategies, the objective of the inspector is to select an acquisition x that maximizes the expected damage value of detected illegal resources, conditioned on their true type $\theta_I^0(x)$. The overall two-stage game can be formulated as the following mathematical program with equilibrium constraints:

$$\begin{aligned}
\Psi: \ \max \ & \mathbb{E}_{\varphi \sim \sigma_I^*(\cdot \,|\, \theta_I^0(x))}[u(\varphi, \sigma_A^*) \mid \theta_I^0(x), \pi_A)] \\
\text{s.t. } & \sigma_I^*(\cdot \,|\, \theta_I) \in \underset{\sigma_I(\cdot \,|\, \theta_I) \in \Delta(\mathcal{A}_I)}{\arg\max} \ \mathbb{E}_{\varphi \sim \sigma_I(\cdot \,|\, \theta_I)}[u(\varphi, \sigma_A^* \,|\, \theta_I, \pi_A)], \ \forall\, \theta_I \in \Theta_I(x) \\
& \sigma_A^*(\cdot \,|\, \theta_A) \in \underset{\sigma_A(\cdot \,|\, \theta_A) \in \Delta(\mathcal{A}_A)}{\arg\min} \ \mathbb{E}_{\psi \sim \sigma_A(\cdot \,|\, \theta_A)}[u(\sigma_I^*, \psi \,|\, \pi_I(x), \theta_A)], \ \forall\, \theta_A \in \Theta_A \\
& x \in \mathcal{X}.
\end{aligned}$$

In the resource acquisition problem Ψ, the inspector maximizes the payoff for their true type, knowing that the adversary will select an allocation strategy that accounts for the uncertainty regarding inspector types. We also note that this problem is equivalent to a three-stage game, where the inspector first acquires resources, then selects an inspection strategy, and finally the adversary selects an allocation plan based on their type after observing the acquisition and inspection strategy.

4.2. Solution Method

We leverage the equilibrium properties of the inspection game Γ to solve the resource acquisition problem Ψ in the following theorem:

THEOREM 2. *Let $(k_0^*, \dots, k_r^*)$ and $(\bar{\xi}_0, \dots, \bar{\xi}_{r-1})$ be demarcation locations and target expected damage values computed from Algorithm 1. The resource acquisition problem Ψ can be solved using the following Bounded Generalized Multiple-Choice Knapsack Problem:*

$$\max \quad \sum_{i=0}^{r-1} \sum_{\tau \in \mathcal{T}} \bar{\xi}_i d_\tau x_{\tau,i} \tag{11a}$$

$$\text{s.t.} \quad \sum_{i=0}^{r-1} \sum_{\tau \in \mathcal{T}} b_\tau x_{\tau,i} \leq \hat{b}, \tag{11b}$$

$$\sum_{\tau \in \mathcal{T}} x_{\tau,i} \leq k_{i+1}^* - k_i^*, \, \forall\, i \in [\![0, r-1]\!], \tag{11c}$$

$$x_{\tau,i} \in \mathbb{Z}_{\geq 0}, \qquad \forall\, \tau \in \mathcal{T},\, \forall\, i \in [\![0, r-1]\!], \tag{11d}$$

where $x_{\tau,i}$ represents the number of acquired inspection resources of type $\tau \in \mathcal{T}$ to be allocated within zone $[\![k_i^* + 1, k_{i+1}^*]\!]$, for $i \in [\![0, r-1]\!]$.

From Theorem 2, we find that the resource acquisition problem Ψ can be solved using the integer program (11). Despite the complexity of the mathematical program with equilibrium constraints Ψ, such simplification is made possible thanks to the independence of some Nash equilibria of the inspection game Γ with respect to the information pertaining to the inspector's resources (Corollary 1). As a result, the adversary's equilibrium strategy σ_A^* constructed from Algorithms 1 and 2 does not depend on the inspector's first-stage acquisition decision, providing the equilibrium expected damage values in each location. Furthermore, we also show that φ^0 is always a best response to σ_A^*, simplifying the objective function of each feasible acquisition decision. Finally, the equivalence between (11) and Ψ is sealed from the fact that at optimality of (11), the inspection resources with highest detection rates are acquired for the first zones, satisfying the necessary game-theoretic conditions (8).

Therefore, the resource acquisition problem Ψ can be efficiently solved for real-world instances. In fact, if the monetary costs b_τ (for $\tau \in \mathcal{T}$) and the budget $\hat{b}$ are rational numbers—which is typically the case—then (11) can be solved in pseudo-polynomial time using dynamic programming (Kellerer et al. 2004). We note that this analysis is applicable to a broader range of settings, where for instance the adversary also faces uncertainty regarding the inspector's budget and the inspection resource acquisition costs.

5. Case Studies: Strategic Drug Interdiction at U.S. Maritime Ports

We illustrate the practical implications of our modeling framework for a drug interdiction setting at U.S. maritime ports. We investigate the two-stage game Ψ in two case studies involving ports in California and Florida using historical and synthetic data. Our analysis demonstrates how adversarial uncertainty and the geographic distribution of smuggling capacity influence acquisition and deployment decisions in equilibrium.

5.1. Instance Construction

Port Traffic: In each case study, we consider a drug-enforcement agency (the inspector) interested in acquiring and allocating inspection resources across a set L of maritime ports to detect and intercept illegal drugs concealed by a transnational cartel (the adversary) in cargo containers. To estimate the weekly smuggling capacity c_ℓ at each port $\ell \in L$, we first compute in Table 1 the weekly number of containers—measured in Twenty-foot Equivalent Units (TEUs)—transiting through each port using the 2022 import-tonnage data (U.S. Army Corps of Engineers 2022) and the average TEU container weight (Hari Menon 2022). We then assume that 15% of inbound TEU containers originate from cartel-controlled regions, and that the cartel can compromise 30% of those containers, each hiding up to 10kg of drugs. The weekly smuggling capacity c_ℓ (in kg) at each port $\ell \in L$ is then given by $(\text{Weekly TEUs at port } \ell) \times 0.15 \times 0.3 \times 10$.

Table 1 Weekly container traffic

(a) California ports

Port	# Containers (TEUs)
Long Beach, CA	38,451
Los Angeles, CA	27,634
Richmond, CA	7,364
Oakland, CA	5,877
Stockton, CA	2,483
Oxnard Harbor District, CA	1,573
Redwood City, CA	1,037
San Diego, CA	977
San Francisco, CA	920
Sacramento, CA	811

(b) Florida ports

Port	# Containers (TEUs)
Everglades, FL	5,643
Tampa, FL	5,388
Jacksonville, FL	5,010
Miami, FL	3,928
Canaveral, FL	3,175
Manatee County, FL	2,967
Panama City, FL	457
Palm Beach, FL	395

Inspection Resources: The drug-enforcement agency receives an annual budget of $1M to acquire mobile inspection teams, each specialized in a detection technology (e.g., canine units, X-ray scanners, manual inspections). Inspection resources vary in cost and detection accuracy (Table 2). Once acquired, inspection resources will be deployed weekly throughout the upcoming year, with at most one per port. Following regulatory constraints that balance security and trade flow (United States Government Accountability Office 2012, House Subcommittee on Coast Guard and Maritime Transportation 2015), each inspection resource can inspect at most 3.8% of containers in a port weekly, and its detection rate d_τ is then given by $0.038 \times (\text{Accuracy of resource } \tau)$.

Smuggled Drugs: The cartel aims to smuggle five drug types: fentanyl, heroin, cocaine, methamphetamine, and cannabis. Without loss of generality, we assume that drugs are packaged into 1kg bags, with damage values represented by their selling prices (United Nations Office on Drugs and Crime 2024), displayed in Figure 4.

Table 2 Cost and accuracy of inspection resources

Inpection Resource	Accuracy	Cost (USD)
Team of German shepherd dogs and their handler's	86.8% Jezierski et al. (2014)	400,000
Team of English cocker spaniel dog and their handler's	82% Jezierski et al. (2014)	360,000
Team operating X-ray scanners	80% Bendahan (2017)	330,000
Team of labrador retriever dog and their handler's	78.8% Jezierski et al. (2014)	315,000
Team of terrier dog and their handler's	67% Jezierski et al. (2014)	200,000
Team of police trainees	50%	150,000

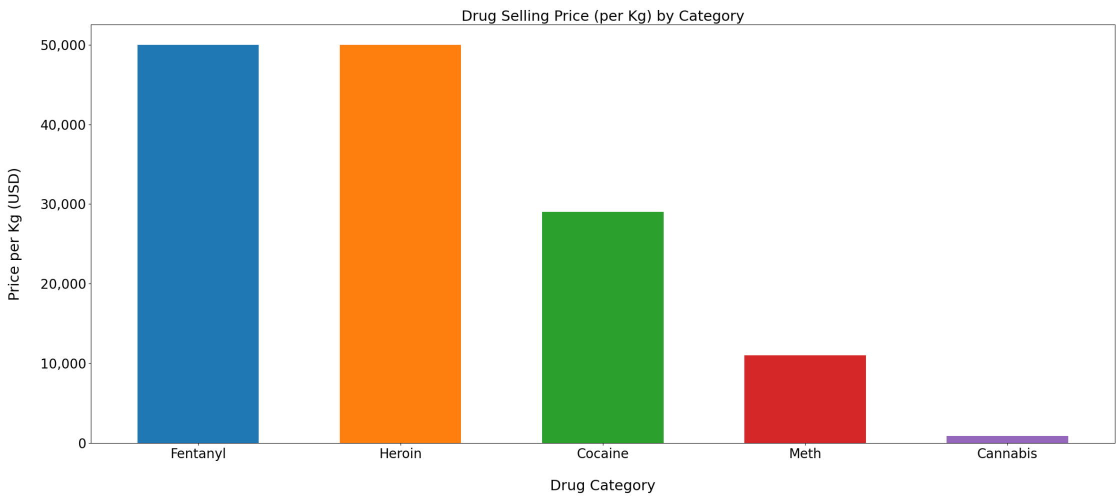

Figure 4 Average drug selling prices by category (USD/kg) in the U.S. between 2017–2022

We then evaluate the volume of each drug type smuggled by the cartel into the U.S. from 2017 until 2022 using drug seizure data (United Nations Office on Drugs and Crime 2024) and the estimated 10% interdiction rate of the U.S. Cast Guard (U.S. Coast Guard 2020); see Figure 5.

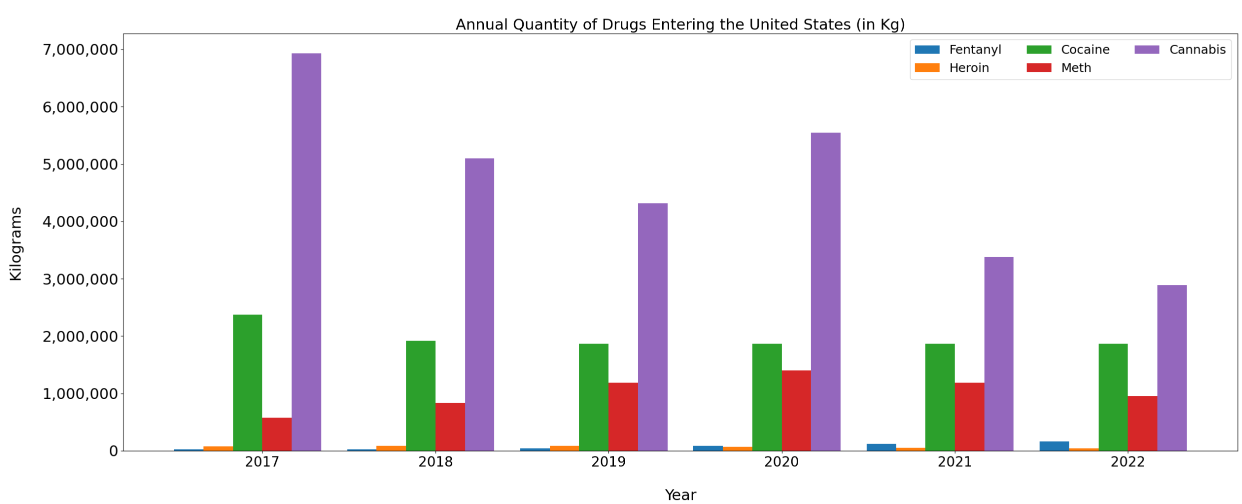

Figure 5 Estimated drug import volumes by category (kg/year) into the U.S. between 2017–2022

For each case study (California and Florida), we assume the agency forecasts that each week of the upcoming year, the cartel smuggles one of six possible combinations of narcotics cargo. Each such "cartel type" θ_A is defined by scaling the annual volumes of drug types smuggled between 2017 and 2022 to weekly quantities and adjusting for the proportion of national container traffic transiting through the corresponding U.S. state (U.S. Army Corps of Engineers 2022). This formulation captures the possibility that the upcoming year's drug cargo will resemble one of the historical profiles from recent years.

Player Uncertainty: From Corollary 1, we assume without loss of generality that the cartel has perfect information regarding the agency's acquired inspection resources. In contrast, we model the uncertainty faced by the agency with a probability distribution π_A over the set Θ_A of six cartel types computed from historical data. We also measure the agency's uncertainty associated with a cartel type distribution π_A using the normalized Shannon entropy, defined by $-\frac{1}{\log_2(|\Theta_A|)} \sum_{\theta_A \in \Theta_A} \pi_A(\theta_A) \cdot \log_2(\pi_A(\theta_A))$, ranging from 0 (perfect information) to 1 (maximum uncertainty).

To generate a diverse set of priors π_A that span all entropy levels, we use the following formula:

$$\forall\, \theta_A \in \Theta_A,\ \pi_A(\theta_A) = \frac{(x_{\theta_A})^\alpha}{\sum_{\theta'_A \in \Theta_A} (x_{\theta'_A})^\alpha}, \tag{12}$$

by drawing for each cartel type θ_A a uniform variable $x_{\theta_A} \sim \mathcal{U}[0,1]$ and varying the exponent $\alpha \geq 1$. For each case study, we generate two million cartel type distributions to reflect a broad spectrum of practical scenarios.

5.2. Computational Analysis

For each case study and cartel type distribution (12), we compute an optimal acquisition plan of the two-stage game Ψ by solving the bounded generalized multiple-choice knapsack problem (11) using Gurobi. Its optimal value provides the expected damage value of intercepted drugs in equilibrium when the cartel and agency allocate their resources at the start of each week according to the strategies computed with Algorithms 2 and 3, respectively.

Value of Intelligence: Figure 6 displays the distributions of equilibrium payoffs stratified by entropy. In both case studies, we observe that the expected value of detected and seized drugs increases monotonically as entropy decreases. However, this increase is not uniform: In the California study, the marginal gain from reducing entropy from 1 to 0.9 is negligible compared to the gain from reducing it from 0.4 to 0.3. This suggests that a minimum threshold of adversarial insight may be necessary for additional intelligence to yield significant operational value.

Overall, the analysis quantifies the impact of uncertainty on inspection performance: In the California (resp. Florida) study, reducing entropy from 1 to 0 yields an average annual increase of \$20.04M (resp. \$7.38M) in the total expected value of seized drugs. These findings underscore the value of intelligence:

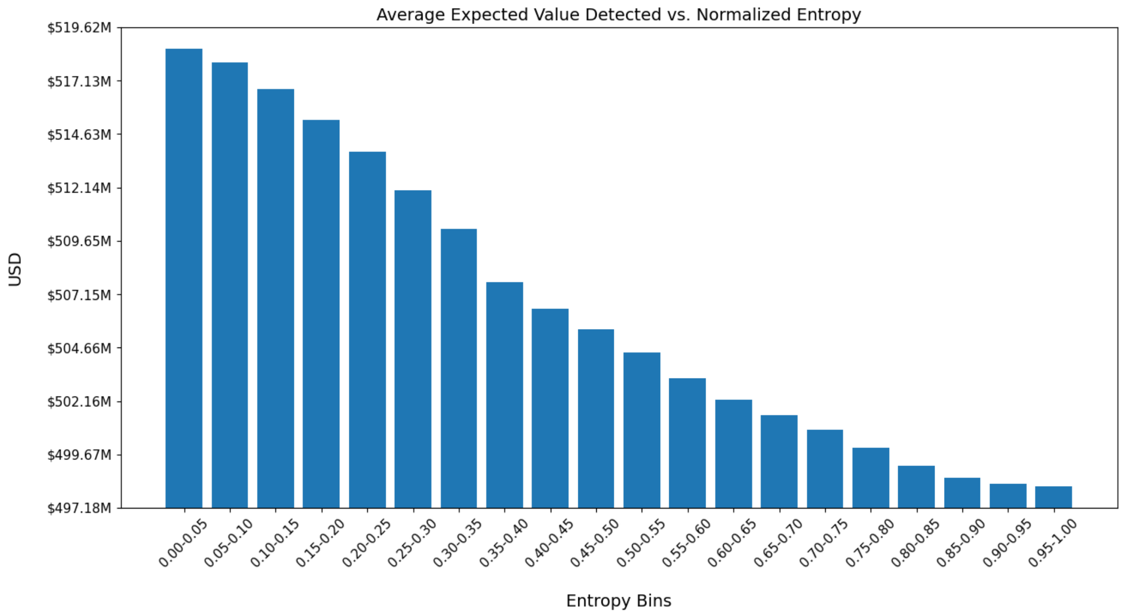

(a) California study

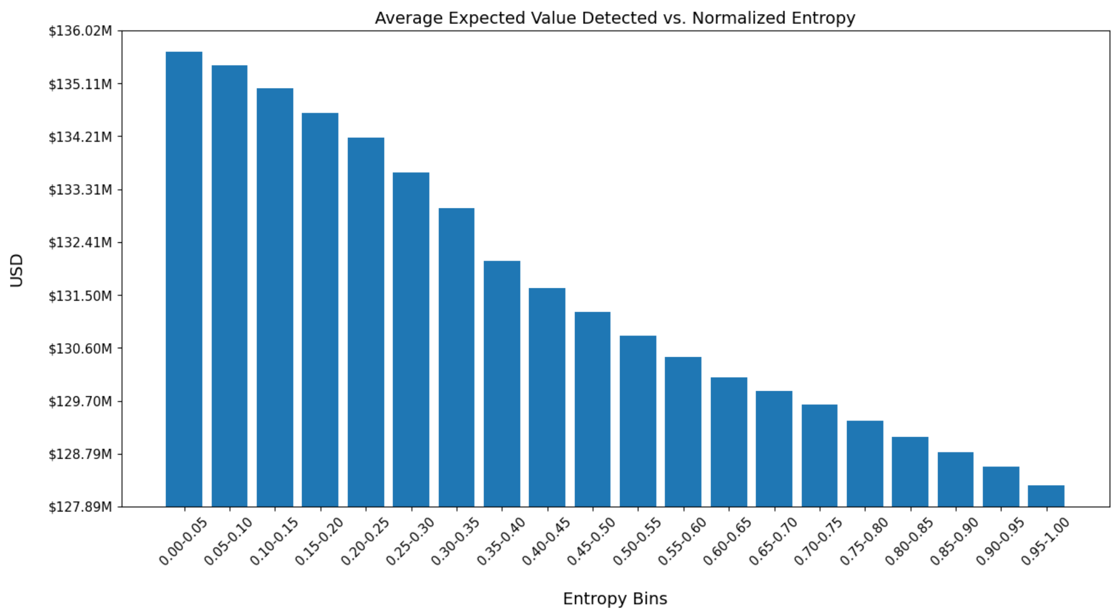

(b) Florida study

Figure 6 Distributions of expected values of detected drugs in equilibrium as functions of normalized entropy. We note that the average *total* expected value of shipped drugs is nearly identical across entropy bins.

as uncertainty regarding the cartel's resources declines, the agency is able to more effectively coordinate inspections and prioritize high-value interdictions.

Sensitivity to Cargo Dispersion: Next, we examine how the dispersion of smuggling capacity across ports influences the agency's optimal inspection resource acquisition. Figure 7 compares the cumulative distributions of port smuggling capacities in the two case studies. The plot shows that cargo traffic in California is significantly more concentrated than in Florida: for example, the two largest ports in California handle 75.8% of the state's total traffic, while in Florida, the top two ports account for only 40.9%.

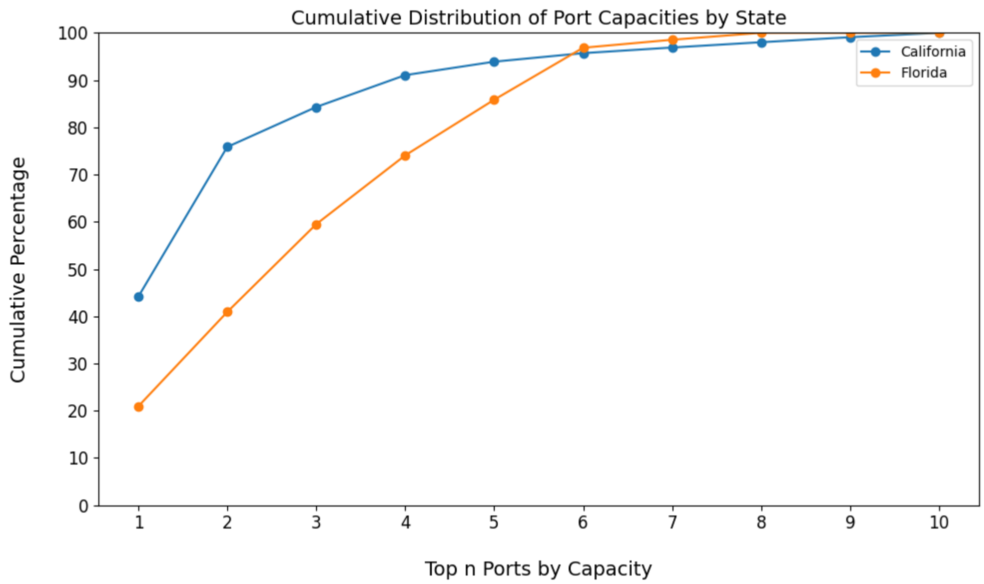

Figure 7 Cumulative smuggling capacities across the ports in California and Florida

To analyze the agency's acquisition behavior, we compute the average cost of the optimally acquired inspection resources for each cartel type distribution. Figure 8 plots these average acquisition costs as a function of entropy for both case studies.

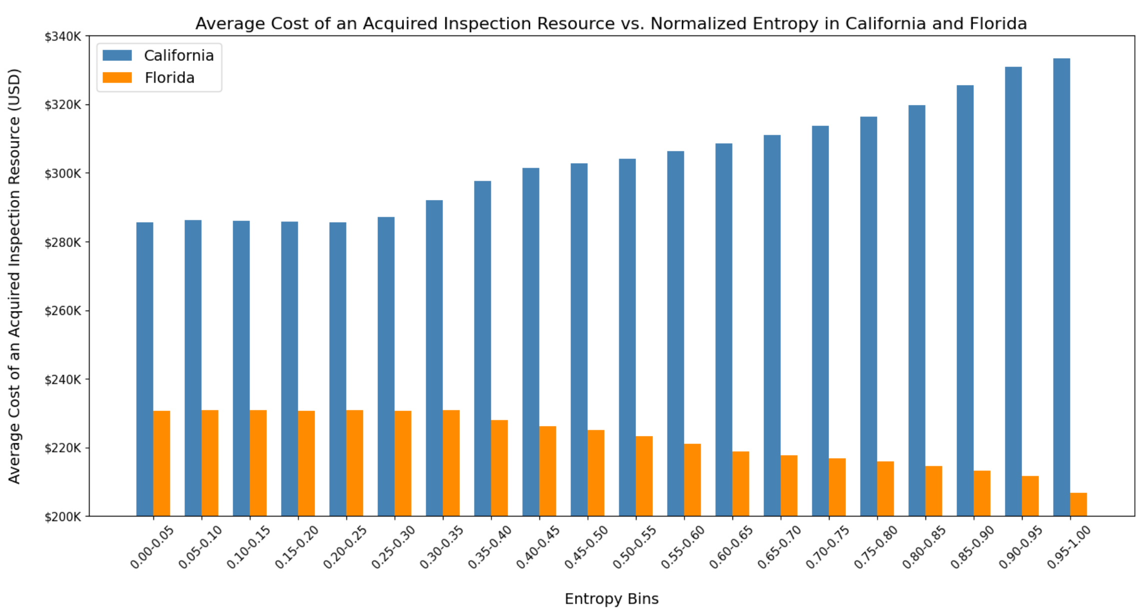

Figure 8 Average cost of acquired inspection resources at optimality as a function of entropy

In the Florida study, we find that under high entropy, the agency tends to acquire more numerous, lower-cost, and less accurate inspection resources. This behavior reflects the dispersed nature of cargo traffic across many mid-sized ports: allocating the budget to just a few high-accuracy resources would risk leaving large volumes of cargo unmonitored under high uncertainty. As the agency gains intelligence about the cartel's operations (i.e., as entropy decreases), it reallocates spending toward fewer, higher-accuracy resources, enabling more targeted deployments to intercept high-value drugs.

Surprisingly, the California study reveals the opposite acquisition pattern. When facing high entropy, the agency opts for fewer but more accurate (and costlier) inspection resources. Because cargo traffic is heavily concentrated in a small number of ports (see Figure 7), the limited available information incentivizes the agency to focus its resources on these ports to maximize the *total* volume of drugs seized, even if those drugs are less valuable. As entropy decreases, the agency diversifies its resource portfolio, acquiring a greater number of less accurate resources to cover more ports when intelligence suggests they may receive more valuable drug shipments.

Managerial Insights: This analysis reveals that both the structure of cargo traffic and the agency's information about adversarial behavior significantly shape optimal acquisition and deployment decisions. While reducing adversarial uncertainty has clear, quantifiable value on interdiction efforts (reaching tens of millions of dollars annually), its impact on the acquisition of inspection resources is more sensitive to the dispersion of cargo traffic. When adversarial uncertainty is high, concentrated traffic favors high-accuracy resources, while dispersed traffic incentivizes broader coverage. In contrast, more adversarial information incentivizes the agency to acquire a mix of low- and high-accuracy resources, achieving a desired trade-off between detection and coverage. These findings suggest that inspection planning should consider uncertainty and system structure as interdependent drivers of effective interdiction policy.

6. Conclusion

In this article, we introduced the first inspection game in which both the inspector and the adversary allocate multiple, heterogeneous resources across capacitated locations while facing uncertainty regarding their opponent's resources. We formulated this problem as a two-player zero-sum game with incomplete information, involving inspection resources that may differ in their detection capabilities and illegal resources that are associated with heterogeneous damage values. The objective of the inspector (resp. adversary) is to maximize (resp. minimize) the expected damage value of detected illegal resources, given the uncertainty they face regarding their opponent's resources, modeled via stochastic player types.

Despite the combinatorial nature of this large-scale game, we leveraged its structure to develop a three-step approach for efficiently computing Nash equilibria: First, we partitioned the locations into zones and characterized the locations' expected total damage value of hidden items in equilibrium. We then designed a combinatorial algorithm that feasibly coordinates the illegal resources of each adversary type to match the equilibrium expected damage values while satisfying best-response conditions. Finally, we constructed an inspection strategy by cycling inspection resources for each inspector type, equalizing the locations' marginal detection probabilities within each zone. Our polynomial-time approach constructs a Nash equilibrium with linear support, which can be deployed practically in real security settings.

From our equilibrium analysis, we could quantify the vulnerability and inspection priority of each location, as a function of the resource damage values of each adversary type and the location capacities. We also showed the surprising result that equilibrium strategies are independent of the inspector's resource capabilities, indicating that concealing them from the adversary yields no strategic advantage. By leveraging these game-theoretic results, we extended our analysis to a two-stage tactical-level problem where the inspector first acquires inspection resources given a budget, and then randomizes their allocation to detect the adversary's illegal resources. We showed that the optimal acquisition plan can be computed in pseudo-polynomial time by solving a bounded generalized multiple-choice knapsack problem using dynamic programming.

By applying our methodology to a case study on drug interdiction at U.S. maritime ports using U.N. drug-seizure records and U.S. port data, we showed that reductions in informational entropy (e.g., through better intelligence) translate into an annual increase of the expected value of detected drugs by approximately \$20.04 million (resp. \$7.38 million) in California (resp. Florida). We also found that the spatial distribution of smuggling capacities crucially shapes resource acquisition: In the more concentrated California setting, the inspector's optimal set of acquired resources tilts towards a few high-accuracy inspection resources under high adversarial uncertainty, whereas the more diffuse Florida setting rewards larger numbers of cheaper and less accurate inspection resources. Our results provide actionable guidance for security agencies charged with deploying heterogeneous inspection resources while operating under uncertainty. Furthermore, our methodology relies on the first combinatorial algorithm for constructing equilibrium strategies by coordinating heterogeneous resources across capacitated locations, with applications beyond the proposed inspection game.

This work can be extended in several directions. One avenue is to allow detection probabilities to depend jointly on the inspection technology and the type of contraband, capturing technology-specific strengths and weaknesses. Another is to model overlapping coverage by permitting multiple inspection resources at a single location and accounting for their combined detection probability, an important feature in critical-infrastructure monitoring. Finally, embedding the stage game within a stochastic, repeated framework—where resource inventories evolve and players update beliefs—would illuminate the dynamic interplay between intelligence gathering and deployment.

Acknowledgments

This work was supported by the Office of Naval Research under Grant No. N00014-24-1-2047.

References

Alpern S, Fokkink R, Kikuta K (2010) On ruckle's conjecture on accumulation games. *SIAM Journal on Control and Optimization* 48:5073–5083.

Avenhaus R, Canty MJ (2009) Inspection games. Meyers RA, ed., *Encyclopedia of Complexity and Systems Science* (New York, NY: Springer), ISBN 978-0-387-75888-6, URL http://dx.doi.org/10.1007/978-0-387-30440-3_287.

Avenhaus R, von Stengel B (2004) Efficient distributions of arms-control inspection effort. *Naval Research Logistics* 51(1):1–27, URL http://dx.doi.org/10.1002/nav.10096.

Avenhaus R, Zamir S (2002) Inspection games. Aumann RJ, Hart S, eds., *Handbook of Game Theory with Economic Applications*, volume 3, chapter 51, 1947–1987 (Elsevier), URL http://dx.doi.org/10.1016/S1574-0005(02)03014-X.

Bahamondes B, Dahan M (2024) Hide-and-seek game with capacitated locations and imperfect detection. *Decision Analysis* 21(2):110–124, URL http://dx.doi.org/10.1287/deca.2023.0012.

Basilico N, Nittis GD, Gatti N (2016) Multi-resource defensive strategies for patrolling games with alarm systems. arXiv:1606.02221.

Bendahan J (2017) Vehicle and cargo scanning for contraband. *Physics Procedia* 90:242–255.

Boone JH, Dahan M (2025) Inspection game with imperfect detection technology. *Naval Research Logistics* URL http://dx.doi.org/10.1002/nav.22241, early View, online ahead of print.

Borrero JS, Prokopyev OA, Sauré D (2019) Sequential interdiction with incomplete information and learning. *Operations Research* 67(1):72–89, URL http://dx.doi.org/10.1287/opre.2018.1773.

Clarkson J, Lin KY, Glazebrook KD (2022) A classical search game in discrete locations. *Mathematics of Operations Research* 48(2):687–707.

De Wulf L, Sokol JB (2005) *Customs Modernization Handbook* (World Bank Publications).

Dresher M (1962) A sampling inspection problem in arms control agreements: A game theoretical analysis. Research Memorandum RM-2972-ARPA, RAND Corporation.

Dziubiński M, Roy J (2018) Hide and seek game with multiple resources. *International Symposium on Algorithmic Game Theory*, 82–86 (Springer).

Freund Y, Schapire R (1999) Adaptive game playing using multiplicative weights. *Games and Economic Behavior* 29(1):79–103.

Gal S, Casas J (2014) Succession of hide–seek and pursuit–evasion at heterogeneous locations. *Journal of the Royal Society Interface* 11(94):20140062.

Gal S, Kohn R (2016) Search games with imperfect detection and prior uncertainty. *European Journal of Operational Research* 249(3):867–879, URL http://dx.doi.org/10.1016/j.ejor.2015.09.042.

Hari Menon (2022) What is teu in shipping - everything you wanted to know. URL https://www.marineinsight.com/maritime-law/teu-in-shipping-everything-you-wanted-to-know/, accessed: 2025-03-17.

Hellerstein L, Lidbetter T, Pirutinsky D (2019) Solving zero-sum games using best-response oracles with applications to search games. *Operations Research* 67(3):731–743.

House Subcommittee on Coast Guard and Maritime Transportation (2015) Prevention of and Response to the Arrival of a Dirty Bomb at a U.S. Port. Hearing before the Subcommittee on Coast Guard and Maritime Transportation of the Committee on Transportation and Infrastructure, House of Representatives, URL https://www.govinfo.gov/content/pkg/CHRG-114hhrg97310/html/CHRG-114hhrg97310.htm, 114th Congress, First Session, October 27, 2015.

Jezierski T, Adamkiewicz E, Walczak M, Sobczyńska M, Gorecka-Bruzda A, Ensminger J, Papet E (2014) Efficacy of drug detection by fully-trained police dogs varies by breed, training level, type of drug and search environment. *Forensic Science International* 237:112–118.

Jorgic D, Gottesdiener L, Cooke K, Eisenhammer S (2024) How fentanyl traffickers are exploiting a u.s. trade law to kill americans. *Reuters Investigates — Fentanyl Express, Part 2* URL https://www.reuters.com/investigates/special-report/drugs-fentanyl-shipping/, accessed 10 Jun 2025.

Kellerer H, Pferschy U, Pisinger D (2004) *The Multiple-Choice Knapsack Problem*, 317–347 (Berlin, Heidelberg: Springer Berlin Heidelberg).

Kiekintveld C, Marecki J, Tambe M (2011) Approximation methods for infinite bayesian stackelberg games: Modeling distributional uncertainty. *Proc. AAMAS*, 1005–1012.

Kikuta K, Ruckle W (1997) Accumulation games, part 1: Noisy search. *Journal of Optimization Theory and Applications* 94:395–408.

Kikuta K, Ruckle W (2002) Continuous accumulation games on discrete locations. *Naval Research Logistics* 49:60–77.

Letchford J, Conitzer V (2013) Solving security games on graphs via marginal probabilities. *Proc. AAAI*, 591–597.

Lidbetter T, Lin KY (2019) Searching for multiple objects in multiple locations. *European Journal of Operational Research* 278(2):709–720.

Lipton R, Markakis E, Mehta A (2003) Playing large games using simple strategies. *Proceedings of the 4th ACM Conference on Electronic Commerce*, 36–41.

McCann B, Dahan M (2022) Network inspection using heterogeneous sensors for detecting strategic attacks. *Proceedings of the 55th Hawaii International Conference on System Sciences* .

McCarthy SM, Tambe M, Kiekintveld C, Gore M, Killion A (2016) Preventing illegal logging: Simultaneous optimization of resource teams and tactics for security. *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 30.

Myerson RB (1991) *Game Theory: Analysis of Conflict* (Harvard University Press).

Nguyen TH, Wang Y, Sinha A, Wellman MP (2019) Deception in finitely repeated security games. *Proc. AAAI*.

Pirani F, Jabbar M, Katzenbeisser S (2021) A bayesian game for cyber deception against remote attacks on automotive systems. *IEEE Transactions on Dependable and Secure Computing* 18(5):2211–2225, URL http://dx.doi.org/10.1109/TDSC.2021.3069876.

Pita J, Jain M, Marecki J, Ordóñez F, Portway C, Tambe M, Western C, Paruchuri P, Kraus S (2008) Deployed ARMOR protection: The application of a game theoretic model for security at the Los Angeles International Airport. *Proc. 7th Internat. Conf. Autonomous Agents Multiagent Sys.* (IFAAMAS, Richland, SC), 125–132.

Port of Barcelona (2020) Non-intrusive scanners: make the invisible visible. URL https://piernext.portdebarcelona.cat/en/technology/non-intrusive-scanners-make-the-invisible-visible/, accessed 10 Jun 2025.

United Nations Conference on Trade and Development (2020) Global actors gather to forge a common front against illicit trade. URL https://unctad.org/news/global-actors-gather-forge-common-front-against-illicit-trade, accessed 10 Jun 2025.

United Nations Office on Drugs and Crime (2024) World drug report 2024. United Nations, URL https://www.unodc.org/unodc/en/data-and-analysis/world-drug-report-2024.html.

United States Government Accountability Office (2012) Supply chain security: Cbp needs to conduct regular assessments of its cargo targeting system. Report to Congressional Requesters GAO-13-9, United States Government Accountability Office, URL https://www.gao.gov/

assets/gao-13-9.pdf.

US Army Corps of Engineers (2022) Waterborne commerce statistics center: Tonnage for selected u.s. ports in 2022. Digital Library of the U.S. Army Corps of Engineers, URL https://usace.contentdm.oclc.org/digital/collection/p16021coll2/id/14580/.

US Coast Guard (2020) Fiscal year 2020 annual performance report. Technical report, U.S. Coast Guard, retrieved https://www.uscg.mil/Portals/0/documents/budget/FY_2020_USCG_APR.pdf.

US Customs and Border Protection (2024) Cbp access — 2024.02 — focused enforcement operations to combat illicit fentanyl. URL https://www.cbp.gov/about/congressional-resources/cbp-access/2024-02, reports Operation Plaza Spike results (Oct 2023–Jun 2024): 418,000 lb drugs, 15,000 lb fentanyl, nearly 3,000 weapons, and over 654,000 rounds of ammunition seized. Accessed 10 Jun 2025.

US Department of Homeland Security, Science & Technology Directorate (2023) Securing our ports of entry, one scan (or thousands) at a time. https://www.dhs.gov/science-and-technology/news/2023/10/26/feature-article-securing-our-ports-entry-one-scan-or-thousands-time, discusses efficacy of large-scale X-ray/gamma scanners; accessed 9 Jun 2025.

Von Neumann J (1953) A certain zero-sum two-person game equivalent to the optimal assignment problem. *Contributions to the Theory of Games* 2(0):5–12.

von Stengel B (2014) Recursive inspection games. *arXiv preprint arXiv:1412.0129* URL https://arxiv.org/abs/1412.0129.

Yang R, Ford B, Tambe M, Lemieux A (2014) Adaptive resource allocation for wildlife protection against illegal poachers. *Proc. AAMAS*, 453–460.

Zoroa N, Fernandez-Saez M, Zoroa P (2004) Search and ambush games with capacities. *Journal of Optimization Theory and Applications* 123:431–450.

Proofs of Statements

Proof of Lemma 1. To simplify the expressions, we denote by $\xi_\ell^0 := \sum_{s\in\psi^0(\ell)} \mathbb{E}_{\theta_A\sim\pi_A}[v_s^{\theta_A}]$ the expected damage value in each location $\ell \in [\![1,n]\!]$ when each adversary type $\theta_A \in \Theta_A$ selects the allocation plan $\psi^0 \in \mathcal{A}_A$. The target expected damage values are then given by $\bar{\xi}_i = \frac{1}{k_{i+1}^*-k_i^*}\sum_{\ell=k_i^*+1}^{k_{i+1}^*} \xi_\ell^0$ for every $i \in [\![0, r-1]\!]$.

Let $i \in [\![1, r-1]\!]$. Using the definition of k_i^*, we obtain the following property:

$$\begin{aligned}
&\frac{1}{k_i^*-k_{i-1}^*}\sum_{\ell=k_{i-1}^*+1}^{k_i^*}\xi_\ell^0 > \frac{1}{k_{i+1}^*-k_{i-1}^*}\sum_{\ell=k_{i-1}^*+1}^{k_{i+1}^*}\xi_\ell^0\\
\iff&\frac{k_{i+1}^*-k_{i-1}^*}{(k_i^*-k_{i-1}^*)(k_{i+1}^*-k_i^*)}\sum_{\ell=k_{i-1}^*+1}^{k_i^*}\xi_\ell^0 > \frac{1}{k_{i+1}^*-k_i^*}\sum_{\ell=k_{i-1}^*+1}^{k_{i+1}^*}\xi_\ell^0\\
\iff&\left(\frac{1}{k_i^*-k_{i-1}^*}+\frac{1}{k_{i+1}^*-k_i^*}\right)\sum_{\ell=k_{i-1}^*+1}^{k_i^*}\xi_\ell^0 > \frac{1}{k_{i+1}^*-k_i^*}\left(\sum_{\ell=k_{i-1}^*+1}^{k_i^*}\xi_\ell^0+\sum_{\ell=k_i^*+1}^{k_{i+1}^*}\xi_\ell^0\right)\\
\iff&\frac{1}{k_i^*-k_{i-1}^*}\sum_{\ell=k_{i-1}^*+1}^{k_i^*}\xi_\ell^0 > \frac{1}{k_{i+1}^*-k_i^*}\sum_{\ell=k_i^*+1}^{k_{i+1}^*}\xi_\ell^0 \iff \bar{\xi}_{i-1} > \bar{\xi}_i,
\end{aligned}$$

thus proving that the target expected damage values are decreasing.

Next, we consider $i \in [\![0, r-1]\!]$. When $k \in [\![k_i^*+1, k_{i+1}^*-1]\!]$, from the definition of k_{i+1}^*, we obtain the following property:

$$\begin{aligned}
&\frac{1}{k-k_i^*}\sum_{\ell=k_i^*+1}^{k}\xi_\ell^0 \le \frac{1}{k_{i+1}^*-k_i^*}\sum_{\ell=k_i^*+1}^{k_{i+1}^*}\xi_\ell^0\\
\iff&\frac{1}{k-k_i^*}\left(\sum_{\ell=k_i^*+1}^{k}\xi_\ell^0-\sum_{\ell=k_i^*+1}^{k_{i+1}^*}\xi_\ell^0\right) \le \left(\frac{1}{k_{i+1}^*-k_i^*}-\frac{1}{k-k_i^*}\right)\sum_{\ell=k_i^*+1}^{k_{i+1}^*}\xi_\ell^0\\
\iff&-\frac{1}{k-k_i^*}\sum_{\ell=k+1}^{k_{i+1}^*}\xi_\ell^0 \le \frac{k-k_{i+1}^*}{(k_{i+1}^*-k_i^*)(k-k_i^*)}\sum_{\ell=k_i^*+1}^{k_{i+1}^*}\xi_\ell^0\\
\iff&\frac{1}{k_{i+1}^*-k}\sum_{\ell=k+1}^{k_{i+1}^*}\xi_\ell^0 \ge \frac{1}{k_{i+1}^*-k_i^*}\sum_{\ell=k_i^*+1}^{k_{i+1}^*}\xi_\ell^0,
\end{aligned}$$

thus proving (3). □

Proof of Proposition 1. Let us consider the start of the first iteration of the outer while loop 3-10. We observe that $\sigma_A \in \Delta_A$ when the algorithm is initialized. Furthermore, since for every $\theta_A \in \Theta_A$, $\text{supp}(\sigma_A(\cdot\,|\,\theta_A)) = \{\psi^0\}$, then (4) and (5) are automatically satisfied, and (6) is a consequence of (3).

Let $[\![k_i^*+1, k_{i+1}^*]\!]$ be the zone considered at that iteration. By definition of the outer while loop, there exists $\ell' \in [\![k_i^*+1, k_{i+1}^*]\!]$ such that $\xi_{\ell'}(\sigma_A) \neq \bar{\xi}_i$. If for every $\ell \in [\![k_i^*+1, k_{i+1}^*]\!]$, $\xi_\ell(\sigma_A) \geq \bar{\xi}_i$, then we obtain the following contradiction:

$$(k_{i+1}^* - k_i^*)\bar{\xi}_i \overset{(4)}{=} \sum_{\ell=k_i^*+1}^{k_{i+1}^*} \xi_\ell(\sigma_A) = \xi_{\ell'}(\sigma_A) + \sum_{\ell \in [\![k_i^*+1, k_{i+1}^*]\!] \setminus \{\ell'\}} \xi_\ell(\sigma_A) > (k_{i+1}^* - k_i^*)\bar{\xi}_i.$$

Thus, $j^- = \max\{\ell \in [\![k_i^*+1, k_{i+1}^*]\!] \mid \xi_\ell(\sigma_A) < \bar{\xi}_i\}$ is well-defined.

Next, if $\{\ell \in [\![j^-+1, k_{i+1}^*]\!] \mid \xi_\ell(\sigma_A) > \bar{\xi}_i\} = \emptyset$, then by definition of j^-, we have $\xi_\ell(\sigma_A) = \bar{\xi}_i$ for every $\ell \in [\![j^-+1, k_{i+1}^*]\!]$ and we obtain the following contradiction:

$$\sum_{\ell=j^-}^{k_{i+1}^*} \xi_\ell(\sigma_A) < \bar{\xi}_i + \sum_{\ell=j^-+1}^{k_{i+1}^*} \xi_\ell(\sigma_A) = (k_{i+1}^* - j^- + 1)\bar{\xi}_i \overset{(6)}{\leq} \sum_{\ell=j^-}^{k_{i+1}^*} \xi_\ell(\sigma_A).$$

Thus, $j^+ = \min\{\ell \in [\![j^-+1, k_{i+1}^*]\!] \mid \xi_\ell(\sigma_A) > \bar{\xi}_i\}$ is well-defined.

Next, we must prove that for each iteration of the inner while loop 6-10,

$$\exists\, \theta'_A \in \Theta_A,\ \exists\, \psi' \in \text{supp}(\sigma_A(\cdot \mid \theta'_A)),\ \exists\, S \subseteq \psi'(j^-) \mid |S| = c_{j^+} \text{ and } \sum_{s \in S} v_s^{\theta'_A} < \sum_{s \in \psi'(j^+)} v_s^{\theta'_A}.$$

To prove this result, we consider the mixed strategy $\sigma_A \in \Delta_A$ at the start of the iteration of the outer while loop 3-10, prior to running the inner while loop. Then, for every $\theta_A \in \Theta_A$ and every $\psi \in \text{supp}(\sigma_A(\cdot \mid \theta_A))$, we construct $\widehat{\psi} \in \mathcal{A}_A$ defined by

$$\widehat{\psi}(\ell) = \begin{cases} \psi(j^-) \setminus S_\psi \cup \psi(j^+) & \text{if } \ell = j^- \\ S_\psi & \text{if } \ell = j^+ \\ \psi(\ell) & \text{otherwise} \end{cases}$$

given arbitrary sets $S_\psi \subseteq \psi(j^-)$ satisfying $|S_\psi| = c_{j^+}$. Next, we construct $\widehat{\sigma}_A \in \Delta_A$, defined by $\widehat{\sigma}_A(\widehat{\psi} \mid \theta_A) = \sigma_A(\psi \mid \theta_A)$ for every $\psi \in \text{supp}(\sigma_A(\cdot \mid \theta_A))$ and every $\theta_A \in \Theta_A$. We then obtain the following:

$$\begin{aligned} \xi_{j^-}(\widehat{\sigma}_A) - \xi_{j^-}(\sigma_A) &= \mathbb{E}_{\theta_A \sim \pi_A}\left[\sum_{\widehat{\psi} \in \mathcal{A}_A} \widehat{\sigma}_A(\widehat{\psi} \mid \theta_A) \sum_{s \in \widehat{\psi}(j^-)} v_s^{\theta_A} - \sum_{\psi \in \mathcal{A}_A} \sigma_A(\psi \mid \theta_A) \sum_{s \in \psi(j^-)} v_s^{\theta_A}\right] \\ &= \mathbb{E}_{\theta_A \sim \pi_A}\left[\sum_{\psi \in \mathcal{A}_A} \sigma_A(\psi \mid \theta_A)\left(\sum_{s \in \psi(j^+)} v_s^{\theta_A} - \sum_{s \in S_\psi} v_s^{\theta_A}\right)\right]. \end{aligned}$$

Furthermore:

$$\xi_{j^-}(\widehat{\sigma}_A) = \mathbb{E}_{\theta_A \sim \pi_A}\left[\mathbb{E}_{\widehat{\psi} \sim \widehat{\sigma}_A(\cdot \mid \theta_A)}\left[\sum_{s \in \widehat{\psi}(j^-)} v_s^{\theta_A}\right]\right] \geq \xi_{j^+}(\sigma_A) > \bar{\xi}_i.$$

From (5), we deduce that for every $\theta_A \in \Theta_A$ and every $\psi \in \operatorname{supp}(\sigma_A(\cdot \,|\, \theta_A))$, $\sum_{s\in S_\psi} v_s^{\theta_A} \le \sum_{s\in\psi(j+)} v_s^{\theta_A}$. We then obtain the following:

$$\bar{\xi}_i < \xi_{j^-}(\widehat{\sigma}_A) = \xi_{j^-}(\sigma_A) + \mathbb{E}_{\theta_A\sim\pi_A}\left[\mathbb{E}_{\psi\sim\sigma_A(\cdot\,|\,\theta_A)}\left[\sum_{s\in\psi(j+)} v_s^{\theta_A} - \sum_{s\in S_\psi} v_s^{\theta_A}\right]\right]$$
$$= \xi_{j^-}(\sigma_A) + \mathbb{E}_{\theta_A\sim\pi_A}\left[\mathbb{E}_{\psi\sim\sigma_A(\cdot\,|\,\theta_A)}\left[\left(\sum_{s\in\psi(j+)} v_s^{\theta_A} - \sum_{s\in S_\psi} v_s^{\theta_A}\right)\cdot \mathbb{1}_{\left\{\sum\limits_{s\in\psi(j+)} v_s^{\theta_A} > \sum\limits_{s\in S_\psi} v_s^{\theta_A}\right\}}\right]\right].$$

A similar derivation provides:

$$\bar{\xi}_i > \xi_{j^+}(\widehat{\sigma}_A) = \xi_{j^+}(\sigma_A) - \mathbb{E}_{\theta_A\sim\pi_A}\left[\mathbb{E}_{\psi\sim\sigma_A(\cdot\,|\,\theta_A)}\left[\left(\sum_{s\in\psi(j+)} v_s^{\theta_A} - \sum_{s\in S_\psi} v_s^{\theta_A}\right)\cdot \mathbb{1}_{\left\{\sum\limits_{s\in\psi(j+)} v_s^{\theta_A} > \sum\limits_{s\in S_\psi} v_s^{\theta_A}\right\}}\right]\right].$$

From these inequalities, we deduce that by reassigning the probabilities from all allocation plans $\psi \in \operatorname{supp}(\sigma_A(\cdot\,|\,\theta_A))$ (for every $\theta_A \in \Theta_A$) for which $\sum_{s\in S_\psi} v_s^{\theta_A} < \sum_{s\in\psi(j+)} v_s^{\theta_A}$ to the corresponding allocation plans $\widehat{\psi}$, then the resulting mixed strategy achieves an expected damage value at location j^- (resp. j^+) that is strictly larger (resp. strictly smaller) than the target expected damage value $\bar{\xi}_i$.

Since the inner while loop 6-10 iteratively reassigns these probabilities, we then deduce that as long as $\xi_{j^-}(\sigma_A) < \bar{\xi}_i$ and $\xi_{j^+}(\sigma_A) > \bar{\xi}_i$, there exist $\theta'_A \in \Theta_A$, $\psi' \in \operatorname{supp}(\sigma_A(\cdot\,|\,\theta'_A))$, and $S \subseteq \psi'(j)$ such that $|S| = c_k$ and $\sum_{s\in S} v_s^{\theta'_A} < \sum_{s\in\psi'(k)} v_s^{\theta'_A}$, ensuring that line 7 is well-defined.

Considering an iteration of the inner while loop, let $\psi^\dagger \in \mathcal{A}_A$ be defined in Line 8, p be defined in Line 9, and σ'_A be defined by:

$$\sigma'_A(\psi\,|\,\theta_A) = \begin{cases} \sigma_A(\psi'\,|\,\theta'_A) - p & \text{if } \psi = \psi' \text{ and } \theta_A = \theta'_A \\ p & \text{if } \psi = \psi^\dagger \text{ and } \theta_A = \theta'_A \\ \sigma_A(\psi\,|\,\theta_A) & \text{otherwise.} \end{cases}$$

We note that $\psi^\dagger \notin \operatorname{supp}(\sigma_A(\cdot\,|\,\theta'_A))$ (i.e., $\sigma_A(\psi^\dagger\,|\,\theta'_A) = 0$), as a consequence of (5). Since, $p \le \sigma_A(\psi'\,|\,\theta'_A) \le 1$, we then conclude that $\sigma'_A \in \Delta_A$.

Furthermore, since $p \le \dfrac{\bar{\xi}_i - \xi_{j^-}(\sigma_A)}{\pi_A(\theta'_A)\cdot(\sum\limits_{s\in\psi'(j+)} v_s^{\theta'_A} - \sum\limits_{s\in S} v_s^{\theta'_A})}$, then we obtain the following:

$$\xi_{j^-}(\sigma'_A) = \xi_{j^-}(\sigma_A) - \pi_A(\theta'_A)\cdot p \cdot \sum_{s\in\psi'(j^-)} v_s^{\theta'_A} + \pi_A(\theta'_A)\cdot p\cdot \sum_{s\in\psi^\dagger(j^-)} v_s^{\theta'_A}$$
$$= \xi_{j^-}(\sigma_A) + \pi_A(\theta'_A)\cdot p\cdot\left(\sum_{s\in\psi'(j+)} v_s^{\theta'_A} - \sum_{s\in S} v_s^{\theta'_A}\right) \le \bar{\xi}_i. \qquad \text{(EC.1)}$$

Similarly, since $p \le \dfrac{\xi_{j^+}(\sigma_A) - \bar{\xi}_i}{\pi_A(\theta'_A)\cdot(\sum\limits_{s\in\psi'(j+)} v_s^{\theta'_A} - \sum\limits_{s\in S} v_s^{\theta'_A})}$, then we obtain the following:

$$\xi_{j^+}(\sigma'_A) = \xi_{j^+}(\sigma_A) - \pi_A(\theta'_A)\cdot p\cdot\left(\sum_{s\in\psi'(j+)} v_s^{\theta'_A} - \sum_{s\in\psi^\dagger(j+)} v_s^{\theta'_A}\right)$$

$$= \xi_{j^+}(\sigma_A) - \pi_A(\theta'_A) \cdot p \cdot \left(\sum_{s \in \psi'(j^+)} v_s^{\theta'_A} - \sum_{s \in S} v_s^{\theta'_A} \right) \geq \bar{\xi}_i. \quad \text{(EC.2)}$$

Importantly, if $p = \dfrac{\bar{\xi}_i - \xi_{j^-}(\sigma_A)}{\pi_A(\theta'_A) \cdot (\sum\limits_{s \in \psi'(j^+)} v_s^{\theta'_A} - \sum\limits_{s \in S} v_s^{\theta'_A})}$, then $\xi_{j^-}(\sigma'_A) = \bar{\xi}_i$, and if $p = \dfrac{\xi_{j^+}(\sigma_A) - \bar{\xi}_i}{\pi_A(\theta'_A) \cdot (\sum\limits_{s \in \psi'(j^+)} v_s^{\theta'_A} - \sum\limits_{s \in S} v_s^{\theta'_A})}$, then $\xi_{j^+}(\sigma'_A) = \bar{\xi}_i$. In both cases, the inner while loop terminates. In conclusion, the first iteration of the outer while loop is well-defined.

Next, let us assume that a given iteration of the outer while loop is well-defined, and properties $(i) - (iv)$ are satisfied for that iteration. We will then prove that they also hold for the subsequent iteration. We let $i \in [\![0, r-1]\!]$ be the zone number selected at the current iteration. Following the same argument as above, (5) implies that at the end of each iteration of the inner while loop, $\sigma_A \in \Delta_A$, and the inner while loop terminates.

(i) During the inner while loop, probabilities are reassigned to allocation plans obtained by swapping illegal resources between j^- and j^+ initially assigned by allocation plans in the support of the current mixed strategy. Since j^- and j^+ are both in the same zone, then by inductive hypothesis, the illegal resources stay within the same zone as initially allocated by ψ^0 for each adversary type. Therefore, (4) holds.

(ii) Let σ_A (resp. σ'_A) be the mixed strategy at the start (resp. end) of the iteration of the outer while loop. By construction, $\xi_\ell(\sigma'_A) = \xi_\ell(\sigma_A)$ for every $\ell \in [\![1, n]\!] \setminus \{j^-, j^+\}$. Furthermore, (5) holds for every $i' \in [\![0, r-1]\!] \setminus \{i\}$ and every $(\ell_1, \ell_2) \in [\![k^*_{i'} + 1, k^*_{i'+1}]\!]^2$ such that $\ell_1 < \ell_2$, $\xi_{\ell_1}(\sigma'_A) \neq \bar{\xi}_{i'}$, and $\xi_{\ell_2}(\sigma'_A) \neq \bar{\xi}_{i'}$ by inductive hypothesis since illegal resources from these locations are not reallocated. Similarly, (5) holds for every $(\ell_1, \ell_2) \in ([\![k^*_i + 1, k^*_{i+1}]\!] \setminus \{j^-, j^+\})^2$ such that $\ell_1 < \ell_2$, $\xi_{\ell_1}(\sigma'_A) \neq \bar{\xi}_i$, and $\xi_{\ell_2}(\sigma'_A) \neq \bar{\xi}_i$. It also holds for every $(\ell_1, \ell_2) \in [\![k^*_i + 1, k^*_{i+1}]\!]^2$ such that $\ell_1 < \ell_2$, ℓ_1 or $\ell_2 \in \{j^-, j^+\}$, $\xi_{\ell_1}(\sigma'_A) \neq \bar{\xi}_i$, $\xi_{\ell_2}(\sigma'_A) \neq \bar{\xi}_i$, $\theta_A \in \Theta_A$, and $\psi \in \text{supp}(\sigma'_A(\cdot \,|\, \theta_A))$ that was not constructed during any iteration of the inner while loop (at that iteration of the outer while loop).

Case 1: $\xi_{j^-}(\sigma'_A) = \bar{\xi}_i$. Consider $\psi^\dagger \in \mathcal{A}_A$ constructed during an iteration of the inner while loop from $\psi' \in \text{supp}(\sigma_A(\cdot \,|\, \theta'_A))$ and $\theta'_A \in \Theta_A$. Consider $\ell_1 < j^+$ such that $\xi_{\ell_1}(\sigma'_A) \neq \bar{\xi}_i$. By definition of j^- and j^+, we obtain that $\ell_1 < j^-$. Then, for every $s_1 \in \psi^\dagger(\ell_1)$ and every $s_2 \in \psi^\dagger(j^+) \subseteq \psi'(j^-)$, $v_{s_1}^{\theta'_A} \leq v_{s_2}^{\theta'_A}$ by inductive hypothesis. Similarly, consider $\ell_2 > j^+$ such that $\xi_{\ell_2}(\sigma'_A) \neq \bar{\xi}_i$. Then, for every $s_1 \in \psi^\dagger(j^+) \subseteq \psi'(j^-)$ and every $s_2 \in \psi^\dagger(\ell_2)$, $v_{s_1}^{\theta'_A} \leq v_{s_2}^{\theta'_A}$ by inductive hypothesis.

Case 2: $\xi_{j^+}(\sigma'_A) = \bar{\xi}_i$. Consider $\psi^\dagger \in \mathcal{A}_A$ constructed during an iteration of the inner while loop from $\psi' \in \text{supp}(\sigma_A(\cdot \,|\, \theta'_A))$ and $\theta'_A \in \Theta_A$. Let $\ell_1 < j^-$ such that $\xi_{\ell_1}(\sigma'_A) \neq \bar{\xi}_i$. Then, for every $s_1 \in \psi^\dagger(\ell_1)$ and every $s_2 \in \psi^\dagger(j^-) \subseteq \psi'(j^-) \cup \psi'(j^+)$, $v_{s_1}^{\theta'_A} \leq v_{s_2}^{\theta'_A}$ by inductive hypothesis. Next, let $\ell_2 > j^-$ such that $\xi_{\ell_2}(\sigma'_A) \neq \bar{\xi}_i$. By definition of j^- and j^+, we obtain that $\ell_2 > j^+$. Then, for every $s_1 \in \psi^\dagger(j^-) \subseteq \psi'(j^-) \cup \psi'(j^+)$ and every $s_2 \in \psi^\dagger(\ell_2)$, $v_{s_1}^{\theta'_A} \leq v_{s_2}^{\theta'_A}$.

In both cases, we conclude that (5) holds for σ'_A.

(iii) Let σ_A (resp. σ'_A) be the mixed strategy at the start (resp. end) of the iteration of the outer while loop. Since $\xi_\ell(\sigma'_A) = \xi_\ell(\sigma_A)$ for every $\ell \in [\![1,n]\!] \setminus \{j^-, j^+\}$, then (6) holds for every $i' \in [\![0, r-1]\!] \setminus \{i\}$ and every $k \in [\![k^*_{i'}+1, k^*_{i'+1}]\!]$.

Next, since $\xi_{j^-}(\sigma'_A) + \xi_{j^+}(\sigma'_A) = \xi_{j^-}(\sigma_A) + \xi_{j^+}(\sigma_A)$, we obtain the following by inductive hypothesis:

$$\forall\, k \in [\![k^*_i+1, j^- - 1]\!] \cup [\![j^+, k^*_{i+1}]\!],\ \frac{1}{k^*_{i+1}-k} \sum_{\ell=k+1}^{k^*_{i+1}} \xi_\ell(\sigma'_A) = \frac{1}{k^*_{i+1}-k} \sum_{\ell=k+1}^{k^*_{i+1}} \xi_\ell(\sigma_A) \geq \bar{\xi}_i.$$

Finally, if $k \in [\![j^-, j^+ - 1]\!]$, then we know that $\bar{\xi}_{j^+}(\sigma'_A) \geq \bar{\xi}_i$ and for every $\ell \in [\![k+1, k^*_{i+1}]\!] \setminus \{j^+\}$, $\xi_\ell(\sigma'_A) = \xi_\ell(\sigma_A) \geq \bar{\xi}_i$ by definition of j^-, providing us with the desired inequality:

$$\frac{1}{k^*_{i+1}-k} \sum_{\ell=k+1}^{k^*_{i+1}} \xi_\ell(\sigma'_A) \geq \bar{\xi}_i.$$

Thus, (6) holds.

(iv) The proof is analogous to the one derived above for the first iteration.

Therefore, we conclude by induction that $(i)-(iv)$ hold at each iteration of the outer while loop, ensuring that Algorithm 2 is well-defined.

We next compute the algorithm's runtime: Constructing the initial mixed strategy σ_A and computing $\xi_\ell(\sigma_A)$ for every $\ell \in [\![1,n]\!]$ can be done in $O(|\Theta_A| m)$. Next, we observe that at each iteration of the outer while loop, the number of locations with expected damage values equal to their target values increases by at least one. Furthermore, for any zone, the last iteration—which results in all target values being met in that zone—ensures that the target expected damage values at both j^- and j^+ are achieved. Then, the number of iterations of the outer loop is at most $k^*_{i+1} - k^*_i - 1$ for every $i \in [\![0, r-1]\!]$, and Algorithm 2 terminates with a maximum of $n-r$ iterations of the outer while loop.

Let us consider the t-th iteration of the outer while loop. Notably, at the start of the iteration, $|\operatorname{supp}(\sigma_A)| \leq |\Theta_A| + t - 1$. Locations j^- and j^+ can be computed in $O(k^*_{i+1} - k^*_i)$ time. Furthermore, the inner while loop can be implemented by simply going once through the support of σ_A and constructing $\psi^\dagger$ whenever the condition in Line 7 is satisfied, without going over the new allocation plans added to the support of σ_A within that same t-th iteration of the outer while loop. Thus, the inner while loop can be implemented in at most $|\operatorname{supp}(\sigma_A)| \leq |\Theta_A| + t - 1$ iterations. At each iteration, the set $S \subseteq \psi'(j^-)$ can be selected as the c_{j^+} resources allocated by ψ' at location j^- with the largest damage values, which can be implemented in $O(c_{j^+})$ by maintaining a nondecreasing order of resources allocated by each allocation plan in $\operatorname{supp}(\sigma_A)$. The corresponding sums $\sum_{s\in S} v_s^{\theta'_A}$ and $\sum_{s\in\psi'(j^+)} v_s^{\theta'_A}$ can also be computed in $O(c_{j^+})$, and Line 7 can be implemented in $O(c_{j^+})$. Then, $\psi^\dagger$ can be constructed in $O(m)$, while maintaining the nondecreasing order of damage values in each locations, by simply swapping S and $\psi'(j^+)$. The remaining lines of the inner while loop can be implemented in $O(1)$. Finally, the expected damage values in each location for the new σ_A can be updated in $O(1)$, since only the values at locations j^- and j^+ must be adjusted using the expression (EC.1)-(EC.2), involving quantities that have already been computed above. Thus, the running time of Algorithm 2 is of the order of:

$$\begin{aligned}&m|\Theta_A|+\sum_{i=0}^{r-1}(k^*_{i+1}-k^*_i)(k^*_{i+1}-k^*_i-1)+\sum_{t=1}^{n-r}(|\Theta_A|+t-1)m\\ \le &m|\Theta_A|+(n-r+1)(n-r)+(n-r)m|\Theta_A|+m\frac{(n-r-1)(n-r)}{2}=O(mn(n+|\Theta_A|)).\end{aligned}$$

The inequality leverages the fact that $\max\{\sum_{i=0}^{r-1}x_i(x_i-1) \mid \sum_{i=0}^{r-1}x_i=n, \text{ and } x_i\ge 1\ \forall\, i\in[\![0,r-1]\!]\}$ reaches its maximum at any extreme point of its feasible region, e.g., $(1,\dots,1,n-r+1)$, with an optimal value of $(n-r+1)(n-r)$.

In conclusion, Algorithm 2 runs in $O(mn(n+|\Theta_A|))$ time and returns a mixed strategy with support of size at most $|\Theta_A|+n-r$. Note that the probability distribution of each adversary type has a support of size at most $n-r+1$. □

Proof of Proposition 2. Lines 1-5 can be implemented in $O(n^2)$ time as follows: We first create for each $i\in[\![0,r-1]\!]$ the ordered list $[(0,i),(\frac{1}{k^*_{i+1}-k^*_i},i)\dots,(1,i)]$ of size $k^*_{i+1}-k^*_i+1$. We then concatenate the r lists, forming a single list of size $n+r$, which we sort in nondecreasing order of the values of the tuples' first element. Then, for the j-th unique value λ_j of the tuples' first element (for every j), we create the set I_j by taking the union of the tuples' second element that are associated with λ_j.

Note that the size of the resulting set Λ is at most $2+\sum_{i=0}^{r-1}(k^*_{i+1}-k^*_i-1)=n-r+2$ (since the values 0 and 1 are shared between all r lists). Then, for every $j\in[\![2,|\Lambda|]\!]$, constructing $\varphi^\dagger$ and updating the values s_i (for $i\in I_j$) require $O(n)$ operations. Thus, the for loop can be implemented in $O(n^2)$, and returns a vector ρ with support of size $|\Lambda|-1\le n-r+1$. Finally, σ_I can be implemented in $O(n|\Theta_I|)$ time using a dictionary that copies for each inspector type the values of ρ assigned to its support. Thus, Algorithm 3 can be implemented in $O(n(n+|\Theta_I|))$ time.

We observe that by definition, each iteration of the for loop constructs an allocation plan $\varphi^\dagger\in\mathcal{A}_I$ that is not already in the support of ρ. Furthermore, $\lambda_1=0$ and $\lambda_{|\Lambda|}=1$, which gives us:

$$\sum_{\varphi\in\mathcal{A}_I}\rho(\varphi)=\sum_{j=2}^{|\Lambda|}(\lambda_j-\lambda_{j-1})=\lambda_{|\Lambda|}-\lambda_1=1.$$

Therefore, $\rho\in\Delta(\mathcal{A}_I)$ is a probability distribution (with support size at most $n-r+1$), and $\sigma_I\in\Delta_I$ is a mixed strategy.

By construction, (8) automatically holds, and the probability distribution ρ satisfies the following:

$$\forall\, i\in[\![0,r-1]\!],\ \forall\,(\ell,k)\in[\![k^*_i+1,k^*_{i+1}]\!]^2,\ \mathbb{P}_{\varphi\sim\rho}(\varphi(\ell)=k)=\frac{1}{k^*_{i+1}-k^*_i}.$$

We can compute the resulting detection probability (9) as follows:

$$\forall\, i \in [\![0, r-1]\!], \forall\, \ell \in [\![k_i^*+1, k_{i+1}^*]\!],\ \zeta_\ell(\sigma_I) = \mathbb{E}_{\theta_I\sim\pi_I}[\mathbb{E}_{\varphi\sim\sigma_I(\cdot\,|\,\theta_I)}[d_{\varphi(\ell)}^{\theta_I}]]$$

$$= \mathbb{E}_{\varphi\sim\rho}[\mathbb{E}_{\theta_I\sim\pi_I}[\sum_{k=k_i^*+1}^{k_{i+1}^*} d_k^{\theta_I}\mathbb{1}_{\{\varphi(\ell)=k\}}]] = \sum_{k=k_i^*+1}^{k_{i+1}^*}\mathbb{E}_{\theta_I\sim\pi_I}[d_k^{\theta_I}]\mathbb{P}_{\varphi\sim\rho}(\varphi(\ell)=k)$$

$$= \frac{1}{k_{i+1}^*-k_i^*}\sum_{k=k_i^*+1}^{k_{i+1}^*}\mathbb{E}_{\theta_I\sim\pi_I}[d_k^{\theta_I}].$$

□

Proof of Theorem 1. Let $(k_0^*,\dots,k_r^*)$ be the demarcation locations computed from Algorithm 1, $\sigma_A^*\in\Delta_A$ be the allocation strategy computed from Algorithm 2, and $\sigma_I^*\in\Delta_I$ be the inspection strategy computed from Algorithm 3. We now consider an inspector type $\theta_I\in\Theta_I$. For any inspection plan $\varphi\in\mathcal{A}_I$ they select, the expected payoff they receive if the adversary selects σ_A^* is given by

$$u(\varphi,\sigma_A^*\,|\,\theta_I,\pi_A) = \mathbb{E}_{\theta_A\sim\pi_A}[\mathbb{E}_{\psi\sim\sigma_A^*(\cdot\,|\,\theta_A)}[\sum_{\ell\in L} d_{\varphi(\ell)}^{\theta_I}\cdot\sum_{j\in\psi(\ell)} v_j^{\theta_A}]] = \sum_{\ell\in L} d_{\varphi(\ell)}^{\theta_I}\cdot\xi_\ell(\sigma_A^*).$$

From Lemma 1 and Proposition 1, we deduce that φ is a best response to σ_A^* if it satisfies

$$\forall\, i\in[\![0,r-1]\!],\quad \bigcup_{\ell\in[\![k_i^*+1,k_{i+1}^*]\!]}\varphi(\ell) = [\![k_i^*+1, k_{i+1}^*]\!]. \tag{EC.3}$$

From (8) in Proposition 2, we conclude that $\sigma_I^*(\cdot\,|\,\theta_I)$ is a best response to σ_A^*.

Next, we consider an adversary type $\theta_A\in\Theta_A$. For any allocation plan $\psi\in\mathcal{A}_A$ they select, the expected payoff they receive if the inspector selects σ_I^* is given by

$$u(\sigma_I^*,\psi\,|\,\pi_I,\theta_A) = \mathbb{E}_{\theta_I\sim\pi_I}[\mathbb{E}_{\varphi\sim\sigma_I^*(\cdot\,|\,\theta_A)}[\sum_{\ell\in L} d_{\varphi(\ell)}^{\theta_I}\cdot\sum_{j\in\psi(\ell)} v_j^{\theta_A}]] = \sum_{\ell\in L}\zeta_\ell(\sigma_I^*)\cdot\sum_{j\in\psi(\ell)} v_j^{\theta_A}.$$

From (9) in Proposition 2, we deduce that ψ is a best response to σ_I^* if it satisfies

$$\forall\, i\in[\![0,r-1]\!],\quad \bigcup_{\ell\in[\![k_i^*+1,k_{i+1}^*]\!]}\psi(\ell) = [\![1+\sum_{j=1}^{k_i^*}c_j, \sum_{j=1}^{k_{i+1}^*}c_j]\!]. \tag{EC.4}$$

From (4) in Proposition 1, we obtain that $\sigma_A^*(\cdot\,|\,\theta_A)$ is a best response to σ_I^*. In conclusion, (σ_I^*,σ_A^*) is an NE of Γ.

The corresponding value of the game Γ is given by

$$u(\sigma_I^*,\sigma_A^*\,|\,\pi_I,\pi_A) = \sum_{\ell\in L}\zeta_\ell(\sigma_I^*)\xi_\ell(\sigma_A^*)$$

$$\overset{(9),\text{Prop. }1}{=} \sum_{i=0}^{r-1}\sum_{\ell=k_i^*+1}^{k_{i+1}^*}\Big(\frac{1}{k_{i+1}^*-k_i^*}\sum_{k=k_i^*+1}^{k_{i+1}^*}\mathbb{E}_{\theta_I\sim\pi_I}[d_k^{\theta_I}]\Big)\cdot\Big(\frac{1}{k_{i+1}^*-k_i^*}\sum_{k=k_i^*+1}^{k_{i+1}^*}\sum_{s\in\psi^0(k)}\mathbb{E}_{\theta_A\sim\pi_A}[v_s^{\theta_A}]\Big)$$

$$= \sum_{i=0}^{r-1}\frac{1}{k_{i+1}^*-k_i^*}\cdot\Big(\sum_{k=k_i^*+1}^{k_{i+1}^*}\mathbb{E}_{\theta_I\sim\pi_I}[d_k^{\theta_I}]\Big)\cdot\Big(\sum_{k=k_i^*+1}^{k_{i+1}^*}\sum_{s\in\psi^0(k)}\mathbb{E}_{\theta_A\sim\pi_A}[v_s^{\theta_A}]\Big) = \sum_{i=0}^{r-1}(k_{i+1}^*-k_i^*)\cdot\bar{\zeta}_i\cdot\bar{\xi}_i.$$

Computing the value of the game first requires determining the demarcation locations. Algorithm 1 has at most n iterations, where each iteration $i \geq 0$ computes up to $n-i$ quantities $\frac{1}{k-k_i^*}\sum_{\ell=k_i^*+1}^{k}\sum_{s\in\psi^0(\ell)}\mathbb{E}_{\theta_A\sim\pi_A}[v_s^{\theta_A}]$. By storing such a quantity for a given k, we can compute the quantity associated with $k+1$ by first multiplying by $k-k_i^*$, then adding c_{k+1} expectation terms—each requiring $O(|\Theta_A|)$ operations—and finally by dividing by $k+1-k_i^*$. The runtime of Algorithm 1 is then of the order of $\sum_{i=0}^{n-1}\sum_{k=i+1}^{n} c_k|\Theta_A| \leq nm|\Theta_A|$. The value of the game can be expressed as $\sum_{i=0}^{r-1}\bar{\xi}_i \cdot \sum_{k=k_i^*+1}^{k_{i+1}^*}\mathbb{E}_{\theta_I\sim\pi_I}[d_k^{\theta_I}]$, which requires $O(n|\Theta_I|)$ additional operations. Therefore, the value of the game Γ can be computed in $O(n(|\Theta_I|+m|\Theta_A|))$ time. □

Proof of Theorem 2. Given the incomplete information regarding adversary types (Θ_A,π_A), let $(k_0^*,\dots,k_r^*)$ and $(\bar{\xi}_0,\dots,\bar{\xi}_{r-1})$ be the demarcation locations and target expected damage values computed from Algorithm 1. For simplicity, we will show that the resource acquisition problem Ψ is equivalent to the following Multiple-Choice Knapsack Problem:

$$\max \quad \sum_{i=0}^{r-1}\sum_{\ell=k_i^*+1}^{k_{i+1}^*}\sum_{\tau\in\mathcal{T}}\bar{\xi}_i d_\tau y_{\tau,\ell} \tag{EC.5a}$$

$$\text{s.t.} \quad \sum_{\ell\in L}\sum_{\tau\in\mathcal{T}} b_\tau y_{\tau,\ell} \leq \hat{b}, \tag{EC.5b}$$

$$\sum_{\tau\in\mathcal{T}} y_{\tau,\ell} \leq 1, \qquad \forall\, \ell\in L, \tag{EC.5c}$$

$$y_{\tau,\ell}\in\{0,1\}, \qquad \forall\,\tau\in\mathcal{T},\ \forall\,\ell\in L. \tag{EC.5d}$$

Problems (11) and (EC.5) are equivalent via the mapping $x_{\tau,i} = \sum_{\ell=k_i^*+1}^{k_{i+1}^*} y_{\tau,\ell}$ for every $\tau\in\mathcal{T}$ and every $i\in[\![0,r-1]\!]$.

Next, we consider a feasible acquisition $x\in\mathcal{X}$ and we let $d_1^{\theta_I^0(x)},\dots,d_n^{\theta_I^0(x)}$ be the associated inspection resources indexed in nonincreasing order of their detection rates, given the true inspector type $\theta_I^0(x)$. From Corollary 1, the allocation strategy $\sigma_A^*\in\Delta_A$ constructed by Algorithm 2 is an equilibrium strategy for the adversary, regardless of the inspector's acquisition. Thus, the expected damage value in each location within zone $[\![k_i^*+1,k_{i+1}^*]\!]$ (for $i\in[\![0,r-1]\!]$) is always given by $\bar{\xi}_i$. Furthermore, from Corollary 1, the probability distribution $\rho\in\Delta(\mathcal{A}_I)$ constructed by Algorithm 3 is a best response to σ_A^* for any inspector type, regardless of the resources available to them. Since $\varphi^0\in\operatorname{supp}(\rho)$, then φ^0 is also a best response to σ_A^*. Therefore, the objective value of acquisition x in Ψ is given by $\sum_{i=0}^{r-1}\sum_{\ell=k_i^*+1}^{k_{i+1}^*}\bar{\xi}_i d_\ell^{\theta_I^0(x)}$. We can then construct a feasible solution y to (EC.5) with identical objective value by selecting $y_{\tau,\ell} = \mathbb{1}_{\{d_\tau = d_\ell^{\theta_I^0(x)}\}}$ for every $\tau\in\mathcal{T}$ and every $\ell\in L$.

We now consider an optimal solution y to (EC.5). Since the target expected damage values are decreasing (Lemma 1), then y allocates inspection resources in nonincreasing order of their detection rates (i.e., $y_{\tau,\ell} = y_{\tau',\ell'} = 1$ for $\ell<\ell'$ implies that $d_\tau \geq d_{\tau'}$). We can then construct the feasible acquisition $x\in\mathcal{X}$, defined by $x_\tau = \sum_{\ell \in L} y_{\tau,\ell}$ for every $\tau \in \mathcal{T}$. Using the same game-theoretic arguments as above, we obtain that the objective value of x in Ψ is identical to the objective value of y in (EC.5).

In conclusion, the optimal values of Ψ, (EC.5), and (11) are identical, and an optimal acquisition plan can be computed by solving (11). $\square$