STRATEGIC RESOURCE COORDINATION FOR DETECTING ILLEGAL ACTIVITY

A Dissertation
Presented to
The Academic Faculty

By

Bastián Bahamondes Pizarro

In Partial Fulfillment
of the Requirements for the Degree
Doctor of Philosophy in Operations Research
School of Industrial and Systems Engineering

Georgia Institute of Technology

December 2024

© Bastián Bahamondes Pizarro 2024

STRATEGIC RESOURCE COORDINATION FOR DETECTING ILLEGAL ACTIVITY

Thesis committee:

Dr. Mathieu Dahan, Advisor
School of Industrial and Systems Engineering
Georgia Institute of Technology

Dr. Santanu Dey
School of Industrial and Systems Engineering
Georgia Institute of Technology

Dr. Alejandro Toriello
School of Industrial and Systems Engineering
Georgia Institute of Technology

Dr. Saurabh Amin
Department of Civil and Environmental Engineering
Massachusetts Institute of Technology

Dr. Juba Ziani
School of Industrial and Systems Engineering
Georgia Institute of Technology

Date approved: August 13, 2024

ACKNOWLEDGMENTS

First and foremost, I would like to express my deepest gratitude to my advisor, Mathieu Dahan. His guidance, support, and encouragement have been essential throughout this journey. I am sincerely grateful for his dedication to my work and the invaluable feedback he provided on both my research and life advise. I am also very thankful to Professors Alejandro Toriello, Juba Ziani, Santanu Dey, and Saurabh Amin for agreeing to serve on my committee and for their kind support of my dissertation.

I would like to extend my heartfelt thanks to my roommates, especially Nauman, Adolfo and Manuel, whose support and friendship have been a constant source of strength throughout this process. Their patience, understanding, humor, and our insightful conversations about life made the challenges of Ph.D. life much more bearable. I will dearly miss our Sunday bike rides along the BeltLine in Atlanta, followed by coffee and chocolate pistachio croissants at Little Tart.

A special thanks goes to the Chilean community in ISyE, who supported and advised me throughout my time in the program: Alfredo, Sebastián, Ramón, Francisco, Daniela, Ignacio, Jaime and Heraldo. Their advice and friendship greatly contributed to my personal and academic growth. I will miss our celebrations, barbecues and soccer matches.

I am also grateful to my fellow Ph.D. students from ISyE: Filipe, Cindy, Sushil, Nidhima, Bobak, Onkar, Haden, Joe, and Pichamon. Each of you contributed in your own way to enriching my journey, and for that, I am deeply appreciative.

I would also like to extend my appreciation to the ISyE staff, especially Amanda, Dawn, Jonathan and May. Their efficiency, dedication, and assistance were vital to my success.

Finally, I want to acknowledge the unconditional support of my family, who have been with me every step of the way, providing love, encouragement, and motivation.

Thank you all for making this achievement possible.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

LIST OF ACRONYMS

CG Column Generation

FG Forward Greedy

LP Linear Program

MIP Mixed-Integer Program

MWU Multiplicative Weights Update

NE Nash Equilibria

RG Reverse Greedy

SUMMARY

In an increasingly complex and interconnected world, ensuring security and resilience requires effective allocation of inspection resources to detect illegal activities. The evolving nature of threats, coupled with resourceful adversaries and limited inspection resources, makes it imperative to develop strategic inspection operations. Challenges include coordinating multiple resources, accounting for imperfect detection capabilities and asymmetric valuations of targets. New opportunities, such as advances in sensing technologies and data analytics, offer potential solutions to enhance the effectiveness of inspection operations.

This thesis leverages game theory for the strategic coordination of inspection resources, focusing on Nash Equilibria (NE) as the main solution concept. It aims to provide valuable insights and efficient algorithms for inspection operations across various security domains.

Chapter 2 examines a variant of the hide-and-seek game, motivated by the challenge of detecting smuggled commodities hidden by criminal organizations. In this game, a seeker inspects multiple hiding locations to find multiple items hidden by a hider. Each hiding location has a maximum hiding capacity and a probability of detecting its hidden items upon inspection. The seeker (resp. hider) aims to minimize (resp. maximize) the expected number of undetected items. We develop a two-step solution approach to compute NE for this zero-sum game. First, we solve a lower-dimensional continuous game to derive closed-form expressions for the equilibrium marginal distributions. Second, we design a combinatorial algorithm to compute mixed strategies that satisfy these marginal distributions. Our approach reveals novel equilibrium behaviors influenced by the complex interplay of game parameters and computes NE in quadratic time with linear support.

Chapter 3 explores a nonzero-sum variant of the hide-and-seek game, driven by the asymmetric valuations that security agencies and criminal organizations place on the outcomes of their interactions. Here, a seeker inspects multiple locations with unit hiding capacities to find items hidden by a hider. Each location is associated with different utility

values for the seeker and hider. The seeker (resp. hider) aims to maximize the utility from inspected (resp. uninspected) locations containing hidden items. We extend the previous two-step approach to obtain NE by deriving closed-form expressions for the equilibrium marginal distributions and computing compatible mixed strategies, resulting in a quadratic-time algorithm for solving this nonzero-sum game. Our analysis not only reveals complex equilibrium behaviors influenced by the players' asymmetric and heterogeneous valuations, but also addresses strategic interactions in various contexts beyond security domains, such as animal behavior and political campaigns. By offering both an intuitive analysis and an efficient solution method, this work bridges a gap in the study of equilibrium behavior in nonzero-sum games of strategic mismatch.

Chapter 4 addresses strategic inspection problems in critical infrastructure resilience through a network inspection game, where a defender positions detectors on a network to detect multiple attacks on its components caused by an attacker. Each detector location has a probability of detecting attacks within its monitored components. The defender (resp. attacker) aims to minimize (resp. maximize) the expected number of undetected attacks. This model extends the hide-and-seek game of Chapter 2 by allowing for detection from multiple locations. To compute NE for this large-scale zero-sum game, we formulate a linear program with a small number of constraints and solve it using Column Generation. We provide an exact mixed-integer program for the pricing problem, which entails computing a defender's pure best response, and leverage its supermodular structure to derive two efficient approaches for obtaining approximate NE with theoretical guarantees: a Column Generation and a Multiplicative Weights Update (MWU) algorithm with approximate best responses. Each iteration of our MWU algorithm requires computing a projection under the unnormalized relative entropy, for which we provide a closed-form solution and a linear-time algorithm. Our computational results in real-world gas distribution networks demonstrate the performance and scalability of our solution approaches.

CHAPTER 1
INTRODUCTION

1.1 Motivation: Strategic Resource Coordination in Security Domains

In an increasingly complex and interconnected world, the challenges of ensuring security and resilience span various domains. Central to addressing these challenges is the effective allocation of inspection resources to optimize security outcomes. Border security requires authorities to allocate X-ray scanners, sniffer dogs, and customs officers for inspecting cargo to detect smuggled goods and contraband. Similarly, cybersecurity demands the distribution of firewalls, intrusion detection systems, and threat monitoring software to monitor and address potential breaches across numerous network points. Airport security involves the allocation of screening resources and security personnel to effectively identify prohibited items in passengers' luggage. Protecting critical infrastructure networks, such as those managed by utility companies, requires prioritizing inspections using drones, surveillance cameras, and maintenance crews to guard against potential sabotage or terrorism.

Each of these scenarios is compounded by the evolving nature of adversarial threats, where resourceful criminal organizations and malicious actors continuously develop sophisticated methods to evade detection. Additionally, inspection resources are often limited, making it impossible to inspect every target at all times. For instance, in port security, there are insufficient personnel and technology to inspect every cargo container [1]. Similarly, airport security cannot screen every piece of luggage with the same intensity due to time constraints and passenger volume [2], and cybersecurity teams cannot monitor all network traffic at all times [3]. In critical infrastructure protection, resources such as drones and ground patrols are limited, making continuous monitoring of every segment of power grids or water distribution networks unfeasible [4]. These limitations compel security agencies to anticipate and counter these evolving tactics effectively, posing the necessity of strategic inspection methods.

Traditional deterministic inspection strategies, which rely on fixed patterns or schedules, often fall short due to their predictability. This predictability enables adversaries to anticipate inspection routines, identify vulnerabilities, and time their illicit activities to avoid detection. In contrast, randomized inspection strategies introduce necessary uncertainty into the inspection process, making it difficult for adversaries to predict where and when inspections will occur. This unpredictability enhances the deterrent effect, as potential offenders are less inclined to engage in illegal activities without knowing their likelihood of detection. Furthermore, randomized strategies allow for a more flexible allocation of limited inspection resources, ensuring critical areas receive appropriate attention while preserving an element of surprise.

A crucial aspect when implementing randomized inspection strategies is the effective coordination in the allocation of inspection resources. In this context, resource coordination refers to assigning each set of inspection targets the appropriate inspection probability, represented through the frequency of inspections, to ensure that over time, resources are distributed to optimize defenders' goals while maintaining essential unpredictability. Without proper coordination, randomized strategies might suffer from coverage gaps or redundancies, undermining their effectiveness. By strategically coordinating inspection resources, decision-makers can enhance overall security, optimize resource allocation, and make complex systems more resilient to adversarial threats. However, the heterogeneity of inspection targets, with varying levels of risk, detection capability, and criticality, necessitates a nuanced approach to resource coordination. Additionally, the asymmetric valuations of potential targets between defenders and attackers complicate strategic planning. Resource constraints and complex network structures further challenge security operations.

Game theory has emerged as a powerful framework for addressing the complexities of security challenges by modeling the interactions between security decision-makers and criminal organizations as strategic games [see, for example, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15, 16, 17]. These interactions can take the form of hiders and seekers in cargo inspections at ports or defenders and attackers in scenarios of critical infrastructure protection. By framing these interactions as strategic games, game theory enables the analysis and optimal coordination of inspection resources, accounting for the rational behavior of the parties involved. This approach allows for the anticipation of potential actions and reactions, leading to more robust and adaptive security measures.

In this thesis, we leverage game theory to explore the intricacies of strategic resource coordination in various security domains. We focus on optimizing the coordinated deployment of inspection resources through game-theoretic models, aiming to provide valuable insights into the incentives driving decision-making processes and offering efficient algorithms to compute effective randomized resource allocations. Our models often assume that agents cannot observe their opponent's decisions before allocating their resources. This assumption applies to scenarios such as cargo inspections at ports, where criminal organizations load containers with illicit goods at the origin port without knowing the likelihood of inspection at the destination port. Similarly, in critical infrastructure protection, inspection strategies are typically kept confidential to prevent adversaries from exploiting known patterns. For example, in a cyber-physical attack, a remote hacker gaining control over an industrial control system can create disruptions without knowing the on-site inspection schedules performed by drones and crew patrols over the critical infrastructure network. This no-observability assumption leads us to consider Nash Equilibria (NE) as the central solution concept of our models, which we aim to characterize analytically when possible or approximate with efficient algorithms for more complex problem structures.

By addressing these challenges, our models contribute to more effective and robust protection against a range of threats across various security domains, including but not limited to the following scenarios.

1.1.1 Detection of Illegal and Smuggled Goods

The detection of illegal and trafficked goods, such as in port and customs environments, is a critical component of global security. Ports and customs are crucial nodes in international trade, handling vast quantities of goods daily, but their high throughput makes them attractive targets for criminal organizations seeking to smuggle illegal commodities such as drugs, weapons, and contraband. The strategic concealment of these goods within legitimate cargo presents a significant challenge for inspection authorities, requiring advanced detection strategies.

The problem of detecting illegal substances in port containers is well-documented [18, 19, 20, 21]. In 2023, Spanish authorities intercepted a shipment of over 9,500 kilograms of cocaine hidden within a container of bananas [22], demonstrating the sophisticated concealment techniques used by traffickers. Recently, European ports have seen a surge in drug seizures, with traffickers employing increasingly sophisticated methods to evade detection [23, 1]. These incidents highlight the evolving nature of smuggling tactics and the need for robust detection frameworks.

Given the vast volume of cargo processed in ports, conducting a comprehensive inspection of every container is impractical. Therefore, port authorities must strategically allocate their limited inspection resources to optimize the detection of illegal goods. The heterogeneity of hiding locations adds to this challenge, as containers vary in size and illegal items can also be concealed in different parts of ships, such as anchor compartments, the hull, or engine rooms (see Figure 1.1). Furthermore, the detection capabilities of inspection resources may vary depending on the specific locations being inspected, complicating the allocation process even further.

In Chapter 2, we address this challenge by modeling the interactions between inspection authorities (seekers) and smugglers (hiders) as a zero-sum game. The seeker's objective is to strategically allocate inspection resources across locations that differ in hiding capacities and detection capabilities to minimize the expected number of undetected illegal goods.

Figure 1.1: Heterogeneous hiding locations used by drug traffickers on cargo ships. Source: InSight Crime [23]

Conversely, the hider's goal is to maximize this number by selecting effective concealment strategies. By analyzing these interactions through game-theoretic models, we can develop efficient inspection strategies and insights to counteract smuggling tactics.

1.1.2 Deterrence of Illegal Activity

In security domains, a critical aspect is the concept of deterrence [24]. In the context of strategic resource coordination for inspection, deterrence refers to the strategic allocation of inspection resources not only to maximize the immediate utility by uncovering illegal activities but also to influence the behavior of potential offenders by creating a credible threat of detection. The goal is to reduce the overall incidence of illegal activity by making it riskier for offenders to engage in such behavior.

Deterrence can be examined through the asymmetric valuations held by security agencies and criminal organizations for the outcomes of their interactions. For instance, in the context of deterring pickpocket activities in a metropolitan subway system (see Figure 1.2), while pickpockets place more value on the immediate success of their thefts, the city's police department may place additional value on the long-term impact of deterrence. By strategically allocating undercover officers throughout the subway system, the police aim to not only create uncertainty for potential pickpockets but also to establish a persistent covert presence that discourages criminal activities, promoting a safer environment for commuters.

Figure 1.2: An undercover officer of the New York City Police Department's Citywide Pickpocket Unit watches for thieves in a subway station. Source: amNewYork Metro [25]

In Chapter 3, we address the concept of deterrence through a hide-and-seek game where both the seeker and the hider hold asymmetric valuations over the hiding locations. This renders their interactions as a nonzero-sum game, with applications extending beyond security domains to include animal behavior, political campaigns, election audits, and general scenarios of strategic mismatch. By understanding these dynamics, security agencies can design inspection protocols that not only target high-risk areas but also deter potential offenders through strategic uncertainty. This approach can significantly enhance the effectiveness of inspections in various security domains, reducing the overall incidence of illegal activities and improving long-term security outcomes.

1.1.3 Critical Infrastructure Resilience

The resilience of critical infrastructure relies on the strategic coordination of inspection resources to detect and mitigate disruptions caused by malicious cyber-phyisical attackers. Power grids, transportation and telecommunication networks, and water, oil, and gas distribution systems form the backbone of modern societies. Yet they are particularly vulnerable to these attacks, which can cause severe economic, social, and environmental consequences.

Several high-profile incidents underscore the urgent need for robust inspection strategies. In 2015, Ukraine's power grid was compromised by a sophisticated cyberattack, resulting in widespread power outages that affected 225,000 customers [26]. Similarly, the 2021 Colonial Pipeline ransomware attack disrupted fuel supply along the East Coast of the United States, causing significant economic losses and revealing vulnerabilities in the protection of essential distribution networks [27]. That same year, a cyberattack on the Oldsmar water treatment plant in Florida nearly resulted in the contamination of the water supply for 15,000 residents, demonstrating the potential for severe public health risks [28].

Inspection systems for critical infrastructure are often affected by location-specific conditions that undermine their detection capabilities. For instance, high-resolution infrared thermography (IRT) can detect irregular temperature patterns indicating gas leaks, but it may be less effective for buried pipelines or in adverse weather conditions [29]. In-pipe inspection robots face obstacles in complex pipeline geometries, limiting their efficacy [30]. Unmanned aerial systems (UAS) can serve as highly effective inspection resources for critical infrastructure, particularly in the context of pipeline inspection. They can detect, quantify, and map natural gas leaks, even in remote areas (see Figure 1.3). However, they performance can be similarly impacted by environmental factors [4].

Another feature of critical infrastructures is that their network structures allow their components to be monitored from multiple locations. While overlapping areas of monitoring can help mitigate undermined detection capabilities, the interactions between inspec-

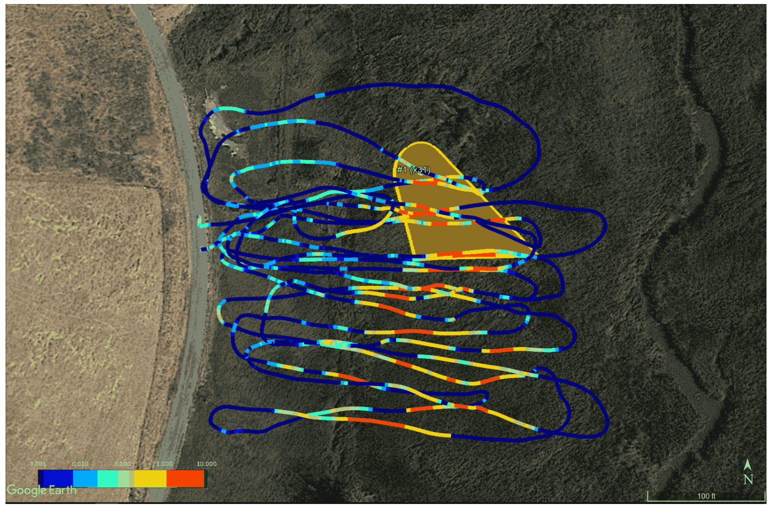

Figure 1.3: Survey map generated by ABB's HoverGuard™ gas leak detection drone after field investigation. Source: ABB Press Release [31]

tion resources at different surveillance points render the probability of an attack remaining undetected a nonadditive function of the set of inspected locations, even under standard independence assumptions. This overlapping feature adds a layer of complexity to the coordination of inspection resources in settings with network structures.

In Chapter 4, we study a zero-sum network inspection game between a defender and attacker to address the optimal allocation of inspection resources in critical infrastructure resilience. This chapter tackles the challenges posed by local conditions undermining detection capabilities and overlapping monitoring of network components. Our model extends the hide-and-seek game from Chapter 2, which can be interpreted as a variant of the network inspection game where each network component can be inspected from exactly one location, i.e. without the overlapping feature. We develop efficient algorithms for computing exact and approximate NE, providing effective methods to enhance the resilience and security of critical infrastructure.

1.2 Thesis Contributions

Hide-and-Seek Game with Capacitated Locations and Imperfect Detection

In Chapter 2, we consider a variant of the hide-and-seek game in which a seeker inspects multiple hiding locations to find multiple items hidden by a hider. Each hiding location has a maximum hiding capacity and a probability of detecting its hidden items upon inspection. The objective of the seeker (resp. hider) is to minimize (resp. maximize) the expected number of undetected items. This model uniquely explores the joint interplay of multiple players' resources with hiding locations that are heterogeneous in both their hiding capacities and detection capabilities, extending previous hide-and-seek games (Gal and Casas [32], Dziubiński and Roy [15], Creasey [33]) to strategic inspection problems, where a security agency is tasked with coordinating multiple inspection resources to detect and seize illegal commodities hidden by a criminal organization.

To solve this large-scale zero-sum game, we leverage its structure and show that its mixed-strategy NE can be characterized using their unidimensional marginal distributions, which are NE of a lower-dimensional continuous zero-sum game. This leads to a two-step approach for efficiently solving our game: First, we analytically solve the continuous game and derive closed-form expressions for the equilibrium marginal distributions. Second, we design a combinatorial algorithm to coordinate the players' resources in heterogeneously capacitated locations—extending a previous algorithm of Dziubiński and Roy [15] for unit hiding capacities—and compute mixed strategies that satisfy the equilibrium marginal distributions. We show that this solution approach computes NE of this hide-and-seek game in quadratic time with linear support. Our analysis reveals novel equilibrium behaviors driven by a complex interplay between the game parameters, captured by our closed-form solutions. We also provide examples of different equilibrium regime patterns and conduct a computational parametric analysis, demonstrating the value of efficiently coordinating multiple inspection resources while accounting for the remaining problem features.

Hide-and-Seek Game with Asymmetric Valuations

In Chapter 3, we examine a nonzero-sum variant of the hide-and-seek game where a seeker inspects multiple locations with unit hiding capacities to find items hidden by a hider. Each location is associated with two (possibly distinct) utility values, representing the respective valuations that the players assign to it. The seeker (resp. hider) aims to maximize the additive utility obtained from the inspected (resp. uninspected) locations containing hidden items. This model extends the zero-sum hide-and-seek game of Dziubiński and Roy [15] and is motivated by the asymmetric valuations that security agencies and criminal organizations place on the outcomes of their interactions. For example, in deterring pickpocket activity, a police department values the long-term deterrence effects of their captures, while criminals value the short-term success of their thefts. Our model is relevant to scenarios of strategic mismatch, where one party benefits from matching (resp. mismatching) their opponent's strategies. Such scenarios arise in various strategic contexts beyond security domains, including animal behavior, political campaigns, and market entry decisions.

Following the two-step approach developed in Chapter 2, we first analytically characterize the NE of this nonzero-sum game by deriving closed-form expressions for the marginal inspection and hiding probabilities in equilibrium at each location. We then compute mixed strategies that satisfy these equilibrium marginal distributions, extending the quadratic-time solution approach used in the zero-sum setting to this nonzero-sum context. Our equilibrium characterization reveals an intricate interaction driven by the players' asymmetric and heterogeneous valuations and their resource constraints.

Our hide-and-seek game also relates to more complex security and inspection games (e.g., Korzhyk *et al.* [8], Clanin and Bhattacharya [34], Deutsch [35]), where players gain additional utilities from locations containing hidden items. While these models have primarily been studied from an algorithmic perspective, with less emphasis on equilibrium behavior, our approach bridges this gap by offering an intuitive analysis of the strategic incentives in equilibrium and presenting an efficient solution approach for solving the game.

Strategic Network Inspection with Location-Specific Detection Capabilities

In Chapter 4, we consider a two-person network inspection game where a defender positions detectors on a network to detect multiple attacks on its components caused by an attacker. Each detector location is associated with a subset of components that can be monitored by the defender upon the placement of a detector. However, detection is imperfect, with each location having a probability of detecting attacks within its monitored components. The objective of the defender (resp. attacker) is to minimize (resp. maximize) the expected number of undetected attacks. This model extends the hide-and-seek game from Chapter 2 by allowing for the possibility of detection from multiple locations, as the sets of monitored components may overlap. This overlapping feature enables us to address strategic inspection problems with network structures, such as safeguarding critical infrastructure against disruptions caused by cyber-physical attackers. It also extends the classic hide-and-seek game by von Neumann [36] (see also Flood [37]) to more complex overlapping structures and multiple players' resources, as well as the seminal network interdiction game by Washburn and Wood [38] to scenarios involving multiple resources for both the interdictor and evader. Furthermore, it extends the network inspection model of Dahan *et al.* [16] by introducing location-specific detection capabilities.

To compute NE for this large-scale zero-sum game, we leverage the structure of the payoff function and represent the attacker's mixed strategies in terms of their marginal probabilities of targeting each component. From this representation, we formulate a linear program with a linear number of constraints, which we solve via Column Generation. We provide an exact mixed-integer program for the pricing problem, which entails computing a defender's pure best response—i.e., a deterministic detector positioning—against a given marginal attack strategy. We show that the pricing problem is NP-hard and cannot be approximated to any constant factor in polynomial time unless P = NP. We therefore leverage its supermodular structure to derive two efficient approaches to obtain approximate NE with instance-dependent approximation guarantees: A Column Generation and a

Multiplicative Weights Update (MWU) algorithms, both relying in fast approximate best responses obtained through the Forward Greedy or Reverse Greedy algorithms.

Our implementation of the MWU algorithm operates on the attacker's marginal probabilities rather than traditional probability distributions. This requires computing, at each iteration, a projection with respect to the unnormalized relative entropy onto the full-dimensional capped simplex polytope. We provide a closed-form solution for this projection problem and derive a linear-time algorithm for its efficient computation, thereby enabling fast iterations of the MWU algorithm. Additionally, we conduct a computational study on real-world gas distribution networks, illustrating the performance and scalability of our solution approaches.

CHAPTER 2
HIDE-AND-SEEK GAME WITH CAPACITATED LOCATIONS AND IMPERFECT DETECTION

2.1 Introduction

In this chapter, we study a variant of the *hide-and-seek* game in which two players, the hider and the seeker, coordinate multiple resources across a set of discrete locations. Specifically, the hider determines where to allocate multiple items within capacitated locations, while the seeker inspects a limited number of locations to detect the hidden items. However, detection is supposed to be imperfect: When inspecting a location, the seeker finds the items with a location-specific probability that captures the local effects undermining the seeker's detection capabilities. The seeker (resp. hider) aims to minimize (resp. maximize) the expected number of items that are undetected. The objective of this work is to efficiently solve this large-scale zero-sum game—that is, to compute a mixed-strategy Nash (or saddle-point) equilibrium and the value of the game—and gather insights on the players' decisions in equilibrium.

Our model is motivated by security applications involving for instance a security agency, such as a coast guard at a port, seeking to conduct inspections on multiple containers and ship compartments used by a criminal organization to conceal illegal commodities such as drugs or weapons [39, 1]. An analogous scenario arises at security screenings of passengers' luggage at airports [40]. Another motivating application involves a utility company tasked with coordinating multiple imperfect sensors to inspect its service network against failures caused by a malicious cyber-physical attacker who is able to target multiple components of the network [41]. A fourth application of interest concerns auditing election results [42, 43]. In such problems, an auditor allocates a limited number of election officials across several polling locations in order to detect electoral fraud by means of recounts. The fraudster may be a malicious organization who is interested in manipulating the results by deploying a malware in the voting machines located in specific polling locations, which are selected based on their respective number of voters [44]. To capture these strategic interactions and gather insights of the decision-making processes involved in a tractable manner, we adopt a zero-sum game framework.

Our work belongs to the broader class of search games. Previous related models have considered either a single hidden item and imperfect detection capabilities [32], or multiple hidden items in locations with unit capacities and perfect detection capabilities [15], which reduces the applicability of the results in security settings. In more realistic environments, the presence of multiple hidden items and imperfect detection occurs *simultaneously*, and their interplay impacts the agents' decisions. On the other hand, the combinatorial nature of the players' strategies prevents us from solving the game using standard linear programming techniques or approximation algorithms [45, 46, 47].

Hence, in this chapter we focus on the following research questions: *(i) How to optimally coordinate multiple imperfect inspection resources to detect multiple hidden commodities in heterogeneous locations? (ii) How are the optimal inspection and hiding strategies jointly impacted by the detection, location, and players' characteristics?*

2.1.1 Contributions

The contributions of this chapter are summarized as follows:

- We propose a novel hide-and-seek game Γ featuring multiple resources for both players, heterogeneous hiding capacities and imperfect detection capabilities, building upon previous models in the literature [32, 15, 33]. Our model uniquely explores the joint interplay of these distinctive features—some of which have been studied separately in prior works—contributing to a richer understanding of search games and offering valuable insights for addressing more realistic scenarios that arise in security domains.

– We show that a strategy profile is an equilibrium of Γ if and only if the corresponding marginal inspection probabilities and expected numbers of hidden items at each location form a pure equilibrium of a lower dimensional continuous game $\widetilde{\Gamma}$ (Proposition 1). From this equivalence, we adopt a two-step approach for solving the hide-and-seek game Γ.

– In the first step, we analytically solve the continuous game $\widetilde{\Gamma}$ (Theorem 1), avoiding the need for linear programming techniques often utilized in previous works that follow this solution approach [see, for example, 48, 7, 49, 50, 51, 52]. Through our analysis, we identify a set of parameters that encompass the relative difference between the players' resources, as well as a set of threshold values that divide the potential number of resources allocated by the hider, resulting in distinct equilibrium regime patterns within each threshold interval. Moreover, we show that our closed-form solutions describe all pure equilibria of $\widetilde{\Gamma}$ except for certain edge cases (Proposition 2).

– In the second step of our solution approach, we compute a mixed strategy of Γ that is consistent with the unidimensional marginal distributions obtained in equilibrium of $\widetilde{\Gamma}$. We design a generic algorithm to construct a probability distribution over the set of resource allocations of *at most* r resources into n *heterogeneously capacitated* locations, satisfying prescribed expected number of resources allocated in each location (Algorithm 1). This extends the algorithm of [15] that coordinates the allocation of *exactly* r resources into n *single-capacity* locations, which cannot be applied in our setting. Our overall approach solves the hide-and-seek game Γ in time $O(n^2)$ and returns equilibrium strategies with linear supports (Theorem 2 and Corollary 1). This contrasts with linear programming approaches, which are not strongly polynomial and may return strategies that cannot be implemented as easily in practice due to their larger supports.

– We provide novel game-theoretic insights of the equilibrium regimes, as well as quantify how the players' decisions and locations' criticalities are jointly impacted by the detection probabilities, the hiding capacities and the number of resources of each player. To complement our findings, we conduct a computational parametric analysis, revealing the value of efficiently coordinating multiple inspection resources while accounting for the other problem features. Our results can be used to inform security decision makers on tactical decisions, such as investing in additional or upgraded inspection resources, in order to balance inspection costs and detection performance.

2.1.2 Literature Review

The hide-and-seek game is a two-person zero-sum game introduced by von Neumann [36]. In its original version, the hider and the seeker interact on a square matrix of nonnegative entries: The hider selects an entry a_{ij} and the seeker selects either a row or a column of the matrix. If the seeker's selection contains the entry chosen by the hider, then the hider pays the seeker a_{ij}; otherwise, the seeker pays the hider a_{ij}. This game has been studied as a general model of strategic mismatch [53] and its equilibrium strategies are well known [36, 37, 54].

More generally, the hide-and-seek game belongs to the category of search games, in which a searcher is concerned with the optimal way of distributing multiple inspection resources to find one or multiple objects hidden by an adversary—the hider—in a search space. This interaction is commonly represented as a zero-sum game. We refer the reader to the books by [55] and [56], as well as the survey of [57] for an introduction to the field.

Search games featuring imperfect detection have been extensively examined in the literature, with considerable emphasis on the setting of finding a *single* object—which can be identified as the hider—hidden in one of several discrete locations, each one characterized by its detection probability. The objective of minimizing the cost (e.g., time) incurred by the searcher has been examined by [58, 59, 60, 61, 62, 63, 64], and [65]. A variant in which the searcher aims to maximize the probability of finding the hider by a deadline, by possibly performing multiple inspections per location, has been studied by [66] and [67]. The objective of maximizing the probability of finding the hider while performing at most one inspection per location has been considered by [68, 32, 69, 70], and [71]. Other objectives have been examined by [72] and [73].

Search games involving *multiple* objects hidden by a hider have been less studied in the literature. Examples of this setting are the accumulation game [74, 75] and the caching game [76, 77, 78], where the hider wins if a certain critical level of commodities remain hidden after the inspections. The objective of minimizing the cost incurred by the searcher to find all the hidden objects has been examined by [79, 80, 81], and [82]. [15] considered the hide-and-seek game in a set of valued locations where the seeker minimizes the sum of the values of the uninspected locations containing hidden objects. Nonetheless, to the best of our knowledge, the existing literature on search games involving multiple hidden objects has predominantly focused on settings where either perfect detection capabilities are assumed or at most a single object is hidden in each location. Accounting for heterogeneity in *both* hiding capacities and detection probabilities will permit developing more realistic inspection strategies and enhance security measures in various real-world scenarios, such as container inspections at ports or baggage screenings at airports.

Among the aforementioned works, our game is most closely related to those by Gal and Casas [32] and Dziubiński and Roy [15]. The former studied a pursuit-evasion model of the interaction between a prey and a predator. The prey chooses a location to hide from the predator, who is able to inspect multiple locations with imperfect detection. The predator (resp. prey) seeks to maximize (resp. minimize) the probability of capture. Our work builds upon this model by allowing multiple preys to coordinately hide in heterogeneously capacitated locations, extending the objective to the expected number of preys uncaught by the predator. Although the subject of animal behavior is beyond the scope of our work, our model extension addresses analogous situations arising in security domains. In contrast to [15], our model does not consider heterogeneously valued locations, but it extends their setting by introducing imperfect detection and allowing the hider to hide multiple items in each box, subject to capacity constraints. Similarly, [33] considers a single hidden item in valued locations, perfect detection, and penalties for the hider in case the seeker finds the hidden object. Our model extends this setting to multiple hidden items in the case of unitary location values, and penalties given by the detection probabilities.

On the other hand, our hide-and-seek game can be used to model a class of instances of strategic sensor placement problems in which the sets of monitored components associated with the sensors are mutually disjoint. In such instances, our results are directly applicable and generalize the equilibrium characterizations from [38, 83], and [16]. Instances with disjoint monitoring arise in situations where it is desirable to reduce sensor interference or the energetic cost of the network [84, 85]. In other security games, disjoint monitoring is naturally satisfied [86, 8, 43, 87].

The two-step approach of characterizing equilibrium strategies in terms of marginal distributions and then computing compatible mixed strategies has been previously proposed in the literature [see, for example, 48, 7, 49, 50, 51, 15, 52]. Our implementation is closely related to that of Dziubiński and Roy [15]. They provide a combinatorial algorithm for constructing a mixed strategy with linear support in quadratic time with respect to the number of locations. Their algorithm iteratively decomposes a given vector representing the marginal probabilities of allocating one resource in each location into a convex combination of a linear number of discrete resource allocations. Our algorithm extends this decomposition to the more general case of marginal distributions representing expected numbers of resources allocated in heterogeneously capacitated locations, and where the budget of resources does not need to be fully utilized, as opposed to Dziubiński and Roy's setting.

Outline. The rest of this chapter is organized as follows. In Section 2.2, we formulate our hide-and-seek game. We then analytically solve a lower dimensional continuous game to characterize and parametrically analyze the marginal inspection probabilities and expected numbers of hidden items across locations in equilibrium in Section 2.3. In Section 2.4, we design a combinatorial algorithm to coordinate the players' resources and compute an equilibrium of the hide-and-seek game. We then provide some concluding remarks in Section 2.5. Finally, examples and proofs of our results are listed in Appendix A.

2.2 Model Description

We consider a hide-and-seek game involving a seeker who is looking for multiple homogeneous items hidden by a hider in a search space consisting of a set of n hiding locations $[\![1, n]\!] := \{1, \dots, n\}$. The locations are *capacitated*; namely, the hider can hide up to $c_i \in \mathbb{Z}_{>0}$ items in each location i. We let $m := \sum_{i=1}^{n} c_i$ be the total hiding capacity. In addition, the seeker has *imperfect* location-specific detection capabilities. Specifically, we assume that when location i is inspected, each item hidden within that location is found by the seeker with probability p_i, independently of other hidden items. We refer to p_i as the *detection probability* of location i.

We assume that both the hider and seeker are strategic, and hence we adopt a game-theoretic framework to study their behaviors. We define a simultaneous-move two-person zero-sum matrix game Γ between a seeker S (she) and a hider H (he). S can select up to $r_{\text{S}} \in \mathbb{Z}_{>0}$ hiding locations to inspect and H can select up to $r_{\text{H}} \in \mathbb{Z}_{>0}$ items to hide. We assume both players have perfect information regarding the hiding capacities, detection probabilities, as well as their own and their opponent's number of resources. To model the players' action sets, we first define a generic set of feasible resource allocations given a capacity vector $b \in \mathbb{Z}^n_{>0}$ and a resource budget $r \in \mathbb{Z}_{>0}$ as follows:

$$\mathcal{A}(b,r) := \left\{ z \in \mathbb{Z}^n : \ 0 \leq z_i \leq b_i, \ \forall \, i \in [\![1,n]\!], \text{ and } \sum_{i=1}^{n} z_i \leq r \right\}. \tag{2.1}$$

The set $\mathcal{A}(b,r)$ contains all the vectors in $\mathbb{Z}^n$ representing allocations of up to r resources within the locations $i \in [\![1,n]\!]$, respecting their capacities given by b_i. Then, the pure action sets for S and H are given by $\mathcal{A}_{\mathrm{S}} := \mathcal{A}(\mathbf{1}_n, r_{\mathrm{S}})$ and $\mathcal{A}_{\mathrm{H}} := \mathcal{A}(c, r_{\mathrm{H}})$ respectively, where $\mathbf{1}_n$ is the vector of ones in $\mathbb{Z}^n$ and $c = (c_i)_{i \in [\![1,n]\!]}$ is the vector of hiding capacities. For every $x \in \mathcal{A}_{\mathrm{S}}$, $x_i = 1$ if S inspects location $i \in [\![1,n]\!]$ and $x_i = 0$ otherwise, and for each $y \in \mathcal{A}_{\mathrm{H}}$, y_i represents the number of items that H hides at location i.

We consider that the the players are interested in the average number of *undetected* items. Given our detection assumptions, the number of items that S detects in location i when S plays $x \in \mathcal{A}_{\mathrm{S}}$ and H plays $y \in \mathcal{A}_{\mathrm{H}}$ follows a binomial distribution with parameters y_i and $p_i x_i$. We then define the *payoff matrix* of the game Γ as $(u(x,y))_{x \in \mathcal{A}_{\mathrm{S}}, \, y \in \mathcal{A}_{\mathrm{H}}}$, where

$$u(x,y) := \sum_{i=1}^{n} (1 - p_i x_i) y_i, \qquad \forall \, (x,y) \in \mathcal{A}_{\mathrm{S}} \times \mathcal{A}_{\mathrm{H}}. \tag{2.2}$$

We allow the players to use *mixed strategies*, defined as probability distributions over their sets of pure actions. To this aim, we first define the set of probability distributions over the set of feasible resource allocations $\mathcal{A}(b,r)$. This set is given by $\Delta(b,r) := \left\{ \sigma \in [0,1]^{\mathcal{A}(b,r)} : \sum_{z \in \mathcal{A}(b,r)} \sigma_z = 1 \right\}$. Then, the sets of mixed strategies for S and H are given by $\Delta_{\mathrm{S}} := \Delta(\mathbf{1}_n, r_{\mathrm{S}})$ and $\Delta_{\mathrm{H}} := \Delta(c, r_{\mathrm{H}})$, respectively. For every mixed strategy $\sigma^{\mathrm{S}} \in \Delta_{\mathrm{S}}$ (resp. $\sigma^{\mathrm{H}} \in \Delta_{\mathrm{H}}$), σ^{S}_x (resp. σ^{H}_y) is the probability that action $x \in \mathcal{A}_{\mathrm{S}}$ (resp. $y \in \mathcal{A}_{\mathrm{H}}$) is executed by S (resp. H). Given a strategy profile $(\sigma^{\mathrm{S}}, \sigma^{\mathrm{H}}) \in \Delta_{\mathrm{S}} \times \Delta_{\mathrm{H}}$, the expected payoff is then defined as $U(\sigma^{\mathrm{S}}, \sigma^{\mathrm{H}}) := \sum_{x \in \mathcal{A}_{\mathrm{S}}} \sum_{y \in \mathcal{A}_{\mathrm{H}}} \sigma^{\mathrm{S}}_x \sigma^{\mathrm{H}}_y u(x,y)$. We assume that S (resp. H) aims to minimize (resp. maximize) U.

Our game Γ is relevant to settings where a security agency, such as a coast guard, is interested in inspecting containers that can be potentially used by a criminal organization to conceal illegal commodities (e.g., drugs, weapons). In such settings, $[\![1, n]\!]$ represents the set of containers, and for each $i \in [\![1, n]\!]$, c_i represents the maximum number of illegal commodities that can be stored in container i. The coast guard can coordinate multiple teams to simultaneously conduct inspections on a maximum of r_{S} containers, while the criminal organization has r_{H} units of illegal commodities to conceal within the containers. The detection probabilities p_i for $i \in [\![1, n]\!]$ capture the local effects that might undermine the detection capabilities of the inspection units, e.g., container characteristics that can impact the efficacy of drug-sniffing dogs [88]. The objective of the coast guard (resp. criminal organization) is to minimize (resp. maximize) the number of illegal commodities that are undetected. Then, a mixed inspection (resp. hiding) strategy is achieved via a randomized schedule of coordinated operations for the coast guard (resp. criminal organization).

We are interested in the *Nash* (or *saddle-point*) *equilibria* of the game Γ, i.e., strategy profiles $(\sigma^{\text{S}^*}, \sigma^{\text{H}^*}) \in \Delta_{\text{S}} \times \Delta_{\text{H}}$ that satisfy

$$U(\sigma^{\text{S}^*}, \sigma^{\text{H}}) \leq U(\sigma^{\text{S}^*}, \sigma^{\text{H}^*}) \leq U(\sigma^{\text{S}}, \sigma^{\text{H}^*}), \quad \forall\, \sigma^{\text{S}} \in \Delta_{\text{S}},\ \forall\, \sigma^{\text{H}} \in \Delta_{\text{H}}.$$

We refer to $U(\sigma^{\text{S}^*}, \sigma^{\text{H}^*})$ as the *value of the game* Γ. Since Γ is a zero-sum matrix game, both the equilibria and value of the game exist, and a strategy profile $(\sigma^{\text{S}^*}, \sigma^{\text{H}^*}) \in \Delta_{\text{S}} \times \Delta_{\text{H}}$ is an equilibrium of Γ if and only if σ^{S^*} is an optimal solution of the *minmax* problem $\min_{\sigma^{\text{S}} \in \Delta_{\text{S}}} \max_{\sigma^{\text{H}} \in \Delta_{\text{H}}} U(\sigma^{\text{S}}, \sigma^{\text{H}})$, and σ^{H^*} is an optimal solution of the *maxmin* problem $\max_{\sigma^{\text{H}} \in \Delta_{\text{H}}} \min_{\sigma^{\text{S}} \in \Delta_{\text{S}}} U(\sigma^{\text{S}}, \sigma^{\text{H}})$ ([89], see also [90]).

From this equivalence, each player's strategy in equilibrium of Γ can be separately computed by reformulating the minmax and maxmin problems as a pair of dual linear programs. Nonetheless, these linear programs become computationally challenging even for small-sized instances, as their numbers of variables and constraints grow combinatorially with r_{S} and r_{H}. Algorithms for computing approximate equilibria [e.g. 45, 46, 47] are also inapplicable for realistic instances of the game Γ.

We adopt a two-step solution approach for solving the game. First in Section 2.3, we reduce the dimensionality of the problem by characterizing its equilibria using the marginal inspection probability and expected number of hidden items in each location. Then in Section 2.4, we derive an algorithm to coordinate the players' resources and recover mixed strategies that are consistent with the characterized marginal inspection probabilities and expected numbers of hidden items in equilibrium.

2.3 Analytical Characterization of Equilibrium Strategies

In this section, we show that the equilibria of the game Γ can be characterized using their marginal inspection probabilities and expected number of hidden items in each location. We prove that these unidimensional quantities are pure equilibrium strategies of a smaller-sized continuous game, which we solve analytically. This analytical characterization permits us to quantify the impact of the problem parameters on the players' decisions in equilibrium.

2.3.1 Continuous Equivalence

To simplify our analysis of the game Γ, we first derive properties of generic randomized resource allocations. Given a capacity vector $b \in \mathbb{Z}^n_{>0}$ and a resource budget $r \in \mathbb{Z}_{>0}$, we denote by $\widetilde{\mathcal{A}}(b, r) := \{\rho \in \mathbb{R}^n : \ 0 \leq \rho_i \leq b_i, \ \forall\, i \in [\![1, n]\!], \text{ and } \sum_{i=1}^n \rho_i \leq r\}$ the linear programming relaxation of the set of feasible resource allocations $\mathcal{A}(b, r)$. Then, for every probability distribution $\sigma \in \Delta(b, r)$ over $\mathcal{A}(b, r)$, we denote as $\rho(\sigma) = (\rho_i(\sigma))_{i \in [\![1,n]\!]}$ the vector of expected numbers of resources allocated at each location, given by $\rho_i(\sigma) := \sum_{z \in \mathcal{A}(b,r)} z_i \sigma_z$, for every $i \in [\![1, n]\!]$. We present the following relation between $\Delta(b, r)$ and $\widetilde{\mathcal{A}}(b, r)$:

Lemma 1. *Consider a capacity vector $b \in \mathbb{Z}^n_{>0}$, a resource budget $r \in \mathbb{Z}_{>0}$, and a vector $\rho' \in \mathbb{R}^n$. Then, $\rho' \in \widetilde{\mathcal{A}}(b, r)$ if and only if there exists a probability distribution $\sigma \in \Delta(b, r)$ that satisfies $\rho_i(\sigma) = \rho'_i$ for all $i \in [\![1, n]\!]$.*

Lemma 1 is a consequence of the integrality of the polytope $\widetilde{\mathcal{A}}(b, r)$. Thus, $\widetilde{\mathcal{A}}(b, r)$ is the set of vectors representing the expected number of resources allocated at each location resulting from a probability distribution in $\Delta(b, r)$. In particular, for every inspection strategy $\sigma^{\mathrm{S}} \in \Delta_{\mathrm{S}}$ and hiding strategy $\sigma^{\mathrm{H}} \in \Delta_{\mathrm{H}}$, $\rho(\sigma^{\mathrm{S}})$ and $\rho(\sigma^{\mathrm{H}})$ respectively represent the vectors of marginal inspection probabilities and expected numbers of hidden items across the locations.

Lemma 1 permits us to relate the game Γ to a continuous zero-sum game $\widetilde{\Gamma}$ between S and H, who respectively select a vector of marginal inspection probabilities $\rho^{\mathrm{S}} \in \widetilde{\mathcal{A}}_{\mathrm{S}} := \widetilde{\mathcal{A}}(\mathbf{1}_n, r_{\mathrm{S}})$ and a vector of expected numbers of hidden items $\rho^{\mathrm{H}} \in \widetilde{\mathcal{A}}_{\mathrm{H}} := \widetilde{\mathcal{A}}(c, r_{\mathrm{H}})$, and their payoff in $\widetilde{\Gamma}$ is given by

$$\tilde{u}(\rho^{\mathrm{S}}, \rho^{\mathrm{H}}) := \sum_{i=1}^{n} (1 - p_i \rho_i^{\mathrm{S}}) \rho_i^{\mathrm{H}}, \qquad \forall\, (\rho^{\mathrm{S}}, \rho^{\mathrm{H}}) \in \widetilde{\mathcal{A}}_{\mathrm{S}} \times \widetilde{\mathcal{A}}_{\mathrm{H}}, \tag{2.3}$$

which S (resp. H) aims to minimize (resp. maximize). A *pure* Nash (or saddle-point) equilibrium of $\widetilde{\Gamma}$ is a pure strategy profile $(\rho^{\mathrm{S}^*}, \rho^{\mathrm{H}^*}) \in \widetilde{\mathcal{A}}_{\mathrm{S}} \times \widetilde{\mathcal{A}}_{\mathrm{H}}$ satisfying $\tilde{u}(\rho^{\mathrm{S}^*}, \rho^{\mathrm{H}}) \leq \tilde{u}(\rho^{\mathrm{S}^*}, \rho^{\mathrm{H}^*}) \leq \tilde{u}(\rho^{\mathrm{S}}, \rho^{\mathrm{H}^*})$ for every $(\rho^{\mathrm{S}}, \rho^{\mathrm{H}}) \in \widetilde{\mathcal{A}}_{\mathrm{S}} \times \widetilde{\mathcal{A}}_{\mathrm{H}}$. Since $\widetilde{\Gamma}$ is a continuous game with biaffine payoff function and compact strategy spaces, the existence of a pure equilibrium is guaranteed and the value of the game $\widetilde{\Gamma}$ is equal to $\tilde{u}(\rho^{\mathrm{S}^*}, \rho^{\mathrm{H}^*})$ [91, Thm 5.2].

Proposition 1. *The games Γ and $\widetilde{\Gamma}$ are related as follows:*

- *For every strategy profile $(\sigma^{\mathrm{S}}, \sigma^{\mathrm{H}}) \in \Delta_{\mathrm{S}} \times \Delta_{\mathrm{H}}$ of Γ, $U(\sigma^{\mathrm{S}}, \sigma^{\mathrm{H}}) = \tilde{u}(\rho(\sigma^{\mathrm{S}}), \rho(\sigma^{\mathrm{H}}))$.*
- *Let $(\sigma^{\mathrm{S}^*}, \sigma^{\mathrm{H}^*}) \in \Delta_{\mathrm{S}} \times \Delta_{\mathrm{H}}$ be an equilibrium of Γ. Then, $(\rho(\sigma^{\mathrm{S}^*}), \rho(\sigma^{\mathrm{H}^*}))$ is a pure equilibrium of $\widetilde{\Gamma}$.*
- *Let $(\rho^{\mathrm{S}^*}, \rho^{\mathrm{H}^*}) \in \widetilde{\mathcal{A}}_{\mathrm{S}} \times \widetilde{\mathcal{A}}_{\mathrm{H}}$ be a pure equilibrium of $\widetilde{\Gamma}$. Then, every strategy profile $(\sigma^{\mathrm{S}^*}, \sigma^{\mathrm{H}^*}) \in \Delta_{\mathrm{S}} \times \Delta_{\mathrm{H}}$ satisfying $\rho(\sigma^{\mathrm{S}^*}) = \rho^{\mathrm{S}^*}$ and $\rho(\sigma^{\mathrm{H}^*}) = \rho^{\mathrm{H}^*}$ is an equilibrium of Γ.*
- *The values of the games Γ and $\widetilde{\Gamma}$ are identical.*

From this proposition, we deduce that equilibria in mixed strategies of the game Γ can be characterized from equilibria in pure strategies of the game $\widetilde{\Gamma}$. In particular, given a pure equilibrium (ρ^{S^*}, ρ^{H^*}) of $\widetilde{\Gamma}$, ρ^{S^*} (resp. ρ^{H^*}) represents the marginal inspection probabilities (resp. expected numbers of hidden items) at every location in an equilibrium of Γ. This provides a significant computational advantage, as these quantities can be represented with vectors of size n, while the players' strategies in Γ require vectors of exponentially large sizes. In addition, marginal inspection probabilities and expected numbers of hidden items more conveniently quantify the criticality of locations for each player. For the remainder of Section 2.3, we focus on analytically solving the game $\widetilde{\Gamma}$.

2.3.2 Preliminary Analysis

We next derive the intuition behind the players' pure equilibrium strategies of the game $\widetilde{\Gamma}$. From (2.3), the payoff $\tilde{u}(\rho^{S}, \rho^{H})$ in $\widetilde{\Gamma}$ is the sum over the hiding locations $i \in [\![1, n]\!]$ of the expected numbers of hidden items that remain undetected, namely, $\left(1 - p_i\rho_i^{S}\right)\rho_i^{H}$. We refer to $1 - p_i\rho_i^{S}$ as the *marginal undetection probability* of location i, that is, the probability that the hidden items at i remain undetected when S plays a marginal inspection strategy $\rho^{S} \in \widetilde{\mathcal{A}}_{S}$. We also refer to $p_i\rho_i^{H}$ as the *expected detection performance* at location $i \in [\![1, n]\!]$, which represents the expected number of hidden items that S is able to detect by inspecting location i when H plays an expected hiding strategy $\rho^{H} \in \widetilde{\mathcal{A}}_{H}$. These quantities will guide our analysis, as H's incentive is to hide items in locations with highest marginal undetection probabilities, and S's incentive is to inspect locations with highest expected detection performances.

Due to the players' incentives, we can easily show that when $r_S = r_H = 1$, the equilibrium strategies (ρ^{S^*}, ρ^{H^*}) satisfy $\rho_i^{S^*} = \rho_i^{H^*} = (1/p_i)/\left(\sum_{j=1}^{n} 1/p_j\right)$ for every $i \in [\![1, n]\!]$, as in [32]. In other words, S inspects each location i with marginal probability proportional to $1/p_i$ so as to equalize the marginal undetection probability of every location. Similarly, H equalizes the expected detection performance of every location.

If we now we consider a general number of player resources, an analogous intuition would suggest that $\rho_i^{\mathrm{S}^*} = (r_{\mathrm{S}}/p_i)/\left(\sum_{j=1}^n 1/p_j\right)$ and $\rho_i^{\mathrm{H}^*} = (r_{\mathrm{H}}/p_i)/\left(\sum_{j=1}^n 1/p_j\right)$ for every $i \in [\![1, n]\!]$. However, due to the multiplicity of players' resources and the heterogeneity of the hiding capacities and detection probabilities, the inequalities $(r_{\mathrm{S}}/p_i)/\left(\sum_{j=1}^n 1/p_j\right) \leq 1$ or $(r_{\mathrm{H}}/p_i)/\left(\sum_{j=1}^n 1/p_j\right) \leq c_i$ may be violated for some locations, preventing ρ^{S^*} and ρ^{H^*} from belonging to $\widetilde{\mathcal{A}}_{\mathrm{S}}$ and $\widetilde{\mathcal{A}}_{\mathrm{H}}$, respectively. In such cases, S (resp. H) cannot ensure the desired level of inspection (resp. hiding) to these locations, rendering them more attractive for H (resp. less attractive for S).

To address the feasibility issues that arise from our model's features, we introduce the following quantities: We define the *detection potential* of each location $i \in [\![1, n]\!]$ as $p_i c_i$, i.e., the maximum expected number of detected items when S inspects location i. Henceforth, we order the locations such that $p_1 c_1 \leq \cdots \leq p_n c_n$. Moreover, for any given $i \in [\![0, n-1]\!]$, we define a bijective mapping $\pi^i : [\![1, n-i]\!] \to [\![i+1, n]\!]$ that satisfies $p_{\pi^i(1)} \leq \cdots \leq p_{\pi^i(n-i)}$, i.e., that orders the set of locations $[\![i+1, n]\!]$ by their detection probabilities. For convenience, we define $p_0 c_0 := 0$ and $p_{\pi^i(0)} := 0$ for every $i \in [\![0, n-1]\!]$. We also denote $S_k^i := \sum_{j=k}^{n-i} 1/p_{\pi^i(j)}$ for every $i \in [\![0, n-1]\!]$ and $k \in [\![1, n-i+1]\!]$.

We remark that in equilibrium when $r_{\mathrm{S}} \geq n$, S fully inspects (i.e., with marginal probability 1) each location and H greedily hides his items in the locations with smallest detection probabilities. Similarly, in equilibrium when $r_{\mathrm{H}} \geq m$, H saturates the capacity of each location and S greedily allocates her resources to fully inspect the locations with the highest detection potentials. Henceforth, we consider that $r_{\mathrm{S}} \in [\![1, n-1]\!]$ and $r_{\mathrm{H}} \in [\![1, m-1]\!]$.

From the discussion above, H may not be able to ensure the desired expected detection performance for the locations with lowest detection potentials. If we denote by $[\![1, i]\!]$ such locations, then S's incentive is to not inspect them, as she will prefer allocating her resources among the remaining locations $[\![i+1, n]\!]$ with higher expected detection performances. In the latter set of locations, S's incentive is to equalize the marginal undetection probabilities. However, this may not be possible if the detection probabilities p_j for

$j \in [\![i+1, n]\!]$ are heterogeneous. Instead, S can fully inspect the k_i most unreliable locations $\{\pi^i(1), \ldots, \pi^i(k_i)\}$, that is, those with lowest detection probabilities, and equalize the marginal undetection probabilities across locations $\{\pi^i(k_i+1), \ldots, \pi^i(n-i)\}$. For every $i \in [\![0, n-1]\!]$, the value of k_i is given by the following expression:

$$k_i := \max\left\{k \in [\![0, n-i]\!] :\ k + p_{\pi^i(k)} S^i_{k+1} < r_{\mathrm{S}}\right\}.$$

As mentioned above, locations for which S cannot achieve the desired level of inspection become more attractive for H. Thus, H's incentive is to saturate the capacities of the ℓ_i most unreliable locations of $[\![i+1, n]\!]$, that is, $\{\pi^i(1), \ldots, \pi^i(\ell_i)\}$, and equalize the expected detection performance across locations $\{\pi^i(\ell_i+1), \ldots, \pi^i(n-i)\}$. In addition, H must ensure that the expected detection performance in the latter set of locations is at least that of the locations in $[\![1, i]\!]$. By doing so, H discourages S from reallocating her resources to $[\![1, i]\!]$ (we recall that $[\![1, i]\!]$ was initially uninspected by S). For every $i \in [\![0, n-1]\!]$, the value of ℓ_i is given by the following expression:

$$\ell_i := \max\left\{\ell \in [\![0, n-i]\!] :\ \sum_{j=1}^{i} c_j + \sum_{j=1}^{\ell} c_{\pi^i(j)} + p_i c_i S^i_{\ell+1} < r_{\mathrm{H}}\right\}.$$

We note that ℓ_i exists when the number of items to hide satisfies $r_{\mathrm{H}} > \sum_{j=1}^{i} c_j + p_i c_i S^i_1$.

The interplay between k_i and ℓ_i will play an important role in solving the game $\widetilde{\Gamma}$. Furthermore, it is crucial to determine the critical number i of locations with smallest detection potentials that will be saturated by H and uninspected by S in equilibrium.

2.3.3 Analytical Characterization

From the discussion above, we observe that the players' decisions in equilibrium depend on the hiding capacities, detection probabilities, detection potentials, numbers of resources, and the parameters k_i and ℓ_i. To derive closed-form expressions that capture this complex interplay in equilibrium of $\widetilde{\Gamma}$, we define the following key quantities:

$$\tau_{-1} \coloneqq 0, \qquad \tau_i \coloneqq \sum_{j=1}^{i} c_j + \sum_{j=1}^{k_i} c_{\pi^i(j)} + p_{i+1} c_{i+1} S^i_{k_i+1}, \qquad \forall\, i \in [\![0, n-1]\!],$$
$$\nu_i \coloneqq \sum_{j=1}^{i} c_j + \sum_{j=1}^{k_i} c_{\pi^i(j)} + p_i c_i S^i_{k_i+1}, \qquad \forall\, i \in [\![0, n-1]\!].$$

First, we show that these quantities are thresholds that partition the continuous interval $[0, m]$ in the following manner.

Lemma 2. *The numbers $\tau_{-1}, \ldots, \tau_{n-1}$, and $\nu_0, \ldots, \nu_{n-1}$, satisfy $\tau_{-1} = 0$, $\tau_{n-1} = m$, and $\tau_{i-1} \leq \nu_i \leq \tau_i$ for all $i \in [\![0, n-1]\!]$.*

Therefore, the interval $[0, m]$ where r_{H} lies is subdivided by the thresholds τ_i and ν_i. We let $i^* \in [\![0, n-1]\!]$ be the unique index satisfying $\tau_{i^*-1} < r_{\mathrm{H}} \leq \tau_{i^*}$. In the following theorem, we show that the subinterval in which r_{H} resides corresponds to a precise configuration of the parameters and determines a specific equilibrium regime.

Theorem 1 (Regime Pattern 1)**.** *If $i^* = 0$ and $\tau_{-1} < r_{\mathrm{H}} \leq \nu_0$, then $\ell_{0^*} < k_{0^*}$ and a strategy profile $(\rho^{\mathrm{S}^*}, \rho^{\mathrm{H}^*}) \in \widetilde{\mathcal{A}}_{\mathrm{S}} \times \widetilde{\mathcal{A}}_{\mathrm{H}}$ is an equilibrium of the game $\widetilde{\Gamma}$ if it satisfies:*

$$\begin{cases} \rho_i^{\mathrm{S}^*} = 1 & \text{if } i \in \mathcal{J} \cup \{\pi^0(\ell_0+1)\}, \\ \dfrac{p_{\pi^0(\ell_0+1)}}{p_i} \leq \rho_i^{\mathrm{S}^*} \leq 1 & \text{if } i \in \mathcal{K} \setminus \{\pi^0(\ell_0+1)\}, \\ \displaystyle\sum_{i=1}^{n} \rho_i^{\mathrm{S}^*} \leq r_{\mathrm{S}}, \end{cases} \tag{2.4}$$

$$\rho_i^{\mathrm{H}^*} = \begin{cases} c_i & \text{if } i \in \mathcal{J}, \\ r_{\mathrm{H}} - \displaystyle\sum_{j=1}^{\ell_0} c_{\pi^0(j)} & \text{if } i = \pi^0(\ell_0+1), \\ 0 & \text{if } i \in \mathcal{K} \setminus \{\pi^0(\ell_0+1)\}, \end{cases} \tag{2.5}$$

where $\mathcal{J} \coloneqq \{\pi^0(1), \ldots, \pi^0(\ell_0)\}$ and $\mathcal{K} \coloneqq \{\pi^0(\ell_0+1), \ldots, \pi^0(n)\}$. The value of the game *is given by*

$$\tilde{u}(\rho^{S^*}, \rho^{H^*}) = r_H - \sum_{j=1}^{\ell_0} p_{\pi^0(j)} c_{\pi^0(j)} - p_{\pi^0(\ell_0+1)} \left(r_H - \sum_{j=1}^{\ell_0} c_{\pi^0(j)} \right).$$

In Regime Pattern 1, H's number of resources is too small relative to S's capability of inspecting hiding locations. As a consequence, H's equilibrium strategy consists in greedily hiding items into the locations with smallest detection probabilities. This results in H saturating the locations in $\mathcal{J} = \{\pi^0(1), \dots, \pi^0(\ell_0)\}$ and allocating his remaining $r_H - \sum_{i \in \mathcal{J}} c_i$ resources in $\pi^0(\ell_0 + 1)$. The feasibility of H's strategy is guaranteed by the definition of ℓ_0.

Since S has enough resources, as guaranteed by $k_0 \geq \ell_0 + 1 > \ell_0$, her strategy fully inspects the locations in $\mathcal{J} \cup \{\pi^0(\ell_0 + 1)\}$, and allocates sufficient resources across the remaining locations to ensure that their marginal undetection probabilities are at most that of location $\pi^0(\ell_0 + 1)$. This prevents H from reallocating some items from $\mathcal{J}$. Interestingly, the inequalities in (2.4) indicate that S's equilibrium strategy can be achieved without necessarily utilizing all her resources.

Theorem 1 (Regime Pattern 2)**.** *If $\nu_{i^*} < r_H \leq \tau_{i^*}$, then $k_{i^*} \leq \ell_{i^*}$ and a strategy profile $(\rho^{S^*}, \rho^{H^*}) \in \widetilde{\mathcal{A}}_S \times \widetilde{\mathcal{A}}_H$ is an equilibrium of the game $\widetilde{\Gamma}$ if it satisfies:*

$$\rho_i^{S^*} = \begin{cases} 0 & \text{if } i \in \mathcal{I}, \\ 1 & \text{if } i \in \mathcal{J}, \\ \dfrac{r_S - k_{i^*}}{p_i S^{i^*}_{k_{i^*}+1}} & \text{if } i \in \mathcal{K}, \end{cases} \tag{2.6}$$

$$\rho_i^{H^*} = \begin{cases} c_i & \text{if } i \in \mathcal{I} \cup \mathcal{J}, \\ \dfrac{r_H - \sum_{j=1}^{i^*} c_j - \sum_{j=1}^{k_{i^*}} c_{\pi^{i^*}(j)}}{p_i S^{i^*}_{k_{i^*}+1}} & \text{if } i \in \mathcal{K}, \end{cases} \tag{2.7}$$

where $\mathcal{I} \coloneqq \{1, \dots, i^*\}$, $\mathcal{J} \coloneqq \{\pi^{i^*}(1), \dots, \pi^{i^*}(k_{i^*})\}$, and $\mathcal{K} \coloneqq \{\pi^{i^*}(k_{i^*} + 1), \dots, \pi^{i^*}(n -$

$i^*)\}$. *The value of the game is given by*

$$\tilde{u}(\rho^{S^*}, \rho^{H^*}) = r_H - \sum_{j=1}^{k_{i^*}} p_{\pi^{i^*}(j)} c_{\pi^{i^*}(j)} - \frac{\left(r_S - k_{i^*}\right)\left(r_H - \sum_{j=1}^{i^*} c_j - \sum_{j=1}^{k_{i^*}} c_{\pi^{i^*}(j)}\right)}{S^{i^*}_{k_{i^*}+1}}.$$

In Regime Pattern 2, S does not inspect the set of locations $\mathcal{I} = \{1, \ldots, i^*\}$, and instead focuses her resources on the remaining locations. Due to the heterogeneity of the detection probabilities in the set $\{i^* + 1, \ldots, n\}$, S fully inspects its k_{i^*} most unreliable locations, namely $\mathcal{J} = \{\pi^{i^*}(1), \ldots, \pi^{i^*}(k_{i^*})\}$. Then, S allocates her $r_S - k_{i^*}$ resources across the locations in $\mathcal{K} = \{\pi^{i^*}(k_{i^*} + 1), \ldots, \pi^{i^*}(n - i^*)\}$, equalizing their marginal undetection probabilities. The feasibility of S's strategy is guaranteed by the definition of k_{i^*}.

Conversely, H's strategy saturates all locations in $\mathcal{I}$ that are not inspected by S, as well as all locations in $\mathcal{J}$, for which S cannot equalize their marginal undetection probabilities. Then, H allocates his remaining $r_H - \sum_{i \in \mathcal{I} \cup \mathcal{J}} c_i$ resources, equalizing the detection performance across locations in $\mathcal{K}$. The feasibility and equilibrium guarantee of H's strategy are a consequence of the subinterval in which r_H belongs. Indeed, since $\nu_{i^*} < r_H$, then $k_{i^*} \leq \ell_{i^*}$, which implies H can saturate the locations in $\mathcal{I} \cup \mathcal{J}$, while still providing a sufficiently high expected detection performance to the locations in $\mathcal{K}$. This prevents S from reallocating inspection resources from $\mathcal{K}$. Furthermore, since $r_H \leq \tau_{i^*}$, H can feasibly equalize the expected detection performances across $\mathcal{K}$.

Theorem 1 (Regime Pattern 3)**.** *If* $i^* \geq 1$ *and* $\tau_{i^*-1} < r_H \leq \nu_{i^*}$*, then* $\ell_{i^*} < k_{i^*}$ *and a* *strategy profile* $(\rho^{S^*}, \rho^{H^*}) \in \widetilde{\mathcal{A}}_S \times \widetilde{\mathcal{A}}_H$ *is an equilibrium of the game* $\widetilde{\Gamma}$ *if it satisfies:*

$$\rho_i^{S^*} = \begin{cases} 0 & \text{if } i \in \mathcal{I} \setminus \{i^*\}, \\ r_S - \ell_{i^*} - p_{\pi^{i^*}(\ell_{i^*}+1)} S^{i^*}_{\ell_{i^*}+1} & \text{if } i = i^*, \\ 1 & \text{if } i \in \mathcal{J} \cup \{\pi^{i^*}(\ell_{i^*}+1)\}, \\ \dfrac{p_{\pi^{i^*}(\ell_{i^*}+1)}}{p_i} & \text{if } i \in \mathcal{K} \setminus \{\pi^{i^*}(\ell_{i^*}+1)\}, \end{cases} \tag{2.8}$$

$$\rho_i^{H^*} = \begin{cases} c_i & \text{if } i \in \mathcal{I} \cup \mathcal{J}, \\ r_H - \sum_{j=1}^{i^*} c_j - \sum_{j=1}^{\ell_{i^*}} c_{\pi^{i^*}(j)} - p_{i^*} c_{i^*} S^{i^*}_{\ell_{i^*}+2} & \text{if } i = \pi^{i^*}(\ell_{i^*}+1), \\ \dfrac{p_{i^*} c_{i^*}}{p_i} & \text{if } i \in \mathcal{K} \setminus \{\pi^{i^*}(\ell_{i^*}+1)\}, \end{cases} \tag{2.9}$$

where $\mathcal{I} \coloneqq \{1, \ldots, i^*\}$*,* $\mathcal{J} \coloneqq \{\pi^{i^*}(1), \ldots, \pi^{i^*}(\ell_{i^*})\}$*, and* $\mathcal{K} \coloneqq \{\pi^{i^*}(\ell_{i^*}+1), \ldots, \pi^{i^*}(n - i^*)\}$*. The value of the game is given by*

$$\begin{aligned} \tilde{u}(\rho^{S^*}, \rho^{H^*}) = r_H - p_{i^*}\left(r_S - \ell_{i^*} - p_{\pi^{i^*}(\ell_{i^*}+1)} S^{i^*}_{\ell_{i^*}+1}\right) c_{i^*} - \sum_{j=1}^{\ell_{i^*}} p_{\pi^{i^*}(j)} c_{\pi^{i^*}(j)} \\ - p_{\pi^{i^*}(\ell_{i^*}+1)} \left(r_H - \sum_{j=1}^{i^*} c_j - \sum_{j=1}^{\ell_{i^*}} c_{\pi^{i^*}(j)} \right). \end{aligned}$$

In Regime Pattern 3, we observe a more complex behavior from the players' equilibrium strategies. First, H cannot equalize the expected detection performance across locations in $\mathcal{I} = \{1, \ldots, i^*\}$ due to their capacity constraints. Thus, locations in $\mathcal{I}$ are saturated by H and initially left uninspected by S. Second, since $\ell_{i^*} + 1 \leq k_{i^*}$, S cannot achieve the desired marginal undetection probability across $\mathcal{J} \cup \{\pi^{i^*}(\ell_{i^*}+1)\} = \{\pi^{i^*}(1), \ldots, \pi^{i^*}(\ell_{i^*}+1)\}$. Therefore, H's incentive is to saturate locations in $\mathcal{J}$, and equalize the expected detection performance across $\mathcal{K} = \{\pi^{i^*}(\ell_{i^*}+1), \ldots, \pi^{i^*}(n-i^*)\}$ to that of i^*. However, H is left with $r_H - \sum_{i \in \mathcal{I} \cup \mathcal{J}} c_i - p_{i^*} c_{i^*} S^{i^*}_{\ell_{i^*}+1} > 0$ resources that he additionally allocates to location $\pi^{i^*}(\ell_{i^*}+1)$, which S fully inspects. The feasibility and equilibrium guarantees of H's strategy follow from the definition of ℓ_{i^*}.

As a result, S fully inspects $\mathcal{J} \cup \{\pi^{i^*}(\ell_{i^*}+1)\}$, and equalizes the marginal undetection probabilities across $\mathcal{K} \setminus \{\pi^{i^*}(\ell_{i^*}+1)\}$ to that of $\pi^{i^*}(\ell_{i^*}+1)$. This is possible since $\ell_{i^*}+1 \leq k_{i^*}$. However, S still has $r_{\mathrm{S}} - \ell_{i^*} - 1 - p_{\pi^{i^*}(\ell_{i^*}+1)} S^{i^*}_{\ell_{i^*}+2} > 0$ resources that she now allocates to the location in $\mathcal{I}$ (previously left uninspected) with highest expected detection performance, namely, i^*. The feasibility and equilibrium guarantees for S's strategy are a consequence of $i^* \geq 1$ and $r_{\mathrm{H}} > \tau_{i^*-1}$. In particular, the resulting marginal undetection probability in i^* is at least that of locations in $\mathcal{K}$, ensuring H will not reallocate items from i^* to $\mathcal{K}$.

From Theorem 1 and thanks to a carefully selected set of thresholds and parameters, we can derive closed-form expressions for the equilibria of the game $\widetilde{\Gamma}$. First, given $i^* \in [\![0, n-1]\!]$ satisfying $\tau_{i^*-1} < r_{\mathrm{H}} \leq \tau_{i^*}$, we generally find that $[\![1, i^*]\!]$ represents the set of locations for which H cannot equalize the expected detection performance due to their small detection potentials, as intuited in Section 2.3.2. When this occurs, S's incentive is to utilize her resources to inspect the locations $[\![i^*+1, n]\!]$ with higher expected detection performance, leaving the locations $[\![1, i^*]\!]$ uninspected while being saturated by H. However, the players' behaviors in the remaining locations differ depending on the subinterval in which r_{H} belongs. Indeed, we find that the threshold ν_{i^*} determines the relation between k_{i^*} and ℓ_{i^*}, which in turn impacts the set of locations $\mathcal{J}$ that S fully inspects and that H saturates. It also dictates how the players should allocate their remaining resources across the locations in equilibrium. We refer the reader to Appendix A.1 for examples illustrating each equilibrium regime pattern.

Our equilibrium regime patterns generalize the results from the pursuit-evasion game of [32], which can be derived from ours by setting $c_i = 1$ for all $i \in [\![1, n]\!]$ and $r_{\mathrm{H}} = 1$. In such a game, we can show that $i^* = 0$ and $\ell_0 = 0$, so the equilibrium regimes observed by the authors are special instances of Regime Pattern 1 when $k_0 \geq 1$ and Regime Pattern 2 when $k_0 = 0$. We also note that the general equilibrium behaviors described in Theorem 1 have not been observed in previously studied models that considered homogeneous detection probabilities [15, 16] or one unit of resources for one or both players [38, 54].

In addition, if we allow the vector of capacities and the players' resources to be continuous in the game $\widetilde{\Gamma}$, and if we consider the set Ψ of parameters for which $\widetilde{\Gamma}$ is nontrivial, i.e.,

$$\Psi := \left\{ (n, p, c, r_{\mathrm{S}}, r_{\mathrm{H}}) \,:\, n \in \mathbb{Z}_{>0},\ p \in (0,1]^n,\ c \in \mathbb{R}^n_{>0},\ r_{\mathrm{S}} \in (0,n),\ r_{\mathrm{H}} \in (0, \sum_{i=1}^{n} c_i) \right\}, \tag{2.10}$$

then we obtain the following stronger result:

Proposition 2. *The set of parameters for which conditions* (2.4)–(2.9) *are necessary and sufficient for* $(\rho^{\mathrm{S}^*}, \rho^{\mathrm{H}^*}) \in \widetilde{\mathcal{A}}_{\mathrm{S}} \times \widetilde{\mathcal{A}}_{\mathrm{H}}$ *to be a pure equilibrium of* $\widetilde{\Gamma}$ *is dense in* Ψ.

Proposition 2 shows that the analytical expressions in Theorem 1 describe *all* pure equilibria of the game $\widetilde{\Gamma}$, except for some edge cases that arise when key quantities of the game are tied or reach boundary values (see Proposition 8 in Appendix A.2).

2.3.4 Parametric Analysis

We continue our analysis by illustrating the impact of the players' resources on the equilibrium regimes of the game $\widetilde{\Gamma}$. To this end, we consider a hide-and-seek instance defined by 6 locations with detection probabilities $p = (1/8, 1/4, 1/3, 1, 4/5, 5/6)$, a total of 18 hiding capacities distributed according to the capacity vector $c = (2, 2, 4, 2, 3, 5)$ (see Example 1 in Appendix A.1 for an illustration). In Figure 2.1, we plot the regions determined by the subintervals $[\tau_{i-1}, \nu_i]$ and $[\nu_i, \tau_i]$ for each $i \in [\![0, n-1]\!]$ as a function of r_{S} and r_{H}.

We first observe that given a specific regime, the values of r_{S} and r_{H} for which that regime holds form a complex region that may even be disconnected. Indeed, the borders that represent the values of the thresholds τ_i and ν_i are defined by step functions of r_{S} through the parameter k_i. For certain values of r_{S}, some thresholds coincide, making certain

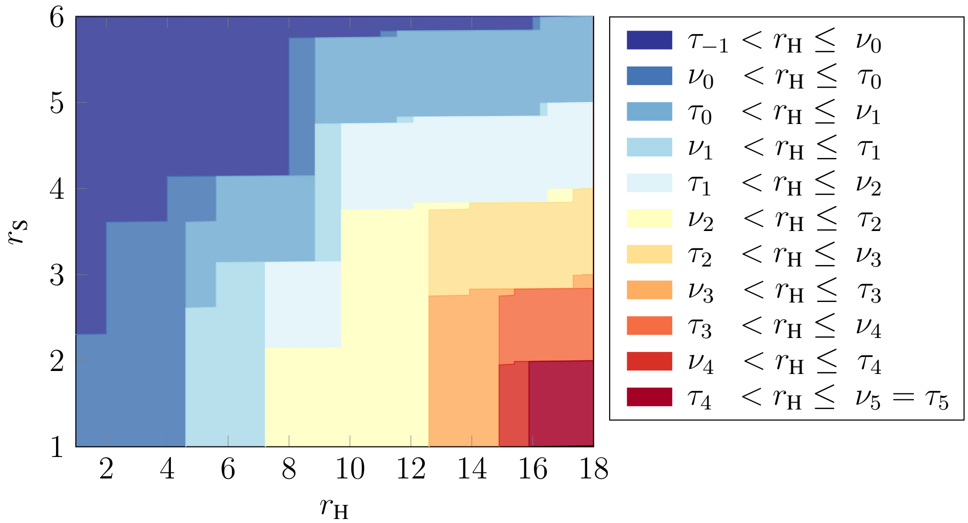

Figure 2.1: Equilibrium regions as a function of the number of resources r_{S} and r_{H} when detection probabilities and hiding capacities are respectively given by $p = (1/8, 1/4, 1/3, 1, 4/5, 5/6)$ and $c = (2, 2, 4, 2, 3, 5)$.

regimes unattainable for any value of r_{H}.

Interestingly, when the number of inspection resources r_{S} is high, very few equilibrium regimes are possible. In such cases, S has enough resources to not leave any location uninspected, resulting in $i^* = 0$. Conversely, when r_{S} is low, S's strategy is highly sensitive to the number of items to hide r_{H}. If r_{H} is small, then H can equalize the expected detection performance across all locations and S should inspect all locations with positive probability. As r_{H} increases, S must carefully determine which locations to inspect and prefers leaving i^* locations with smallest detection potential uninspected to improve the inspection performance in the remaining locations. Furthermore, when r_{S} is low, nearly all the regimes that are achievable follow Regime Pattern 2 (for different values of i^*) since $k_{i^*} \leq \ell_{i^*}$ is most likely to hold for small amounts of inspection resources. Similarly, when r_{H} is high, a single inspection resource incentivizes S to focus her inspection on the last location 6 (i.e., $i^* = 5$). As r_{S} increases, S can allocate more resources to i^* according to (2.8) (Regime Pattern 3) until a new regime emerges with a smaller number of uninspected locations i^*.

By leveraging the closed-form expressions in Theorem 1, we illustrate in Figure 2.2 the impact of the players' resources on the value of the game.

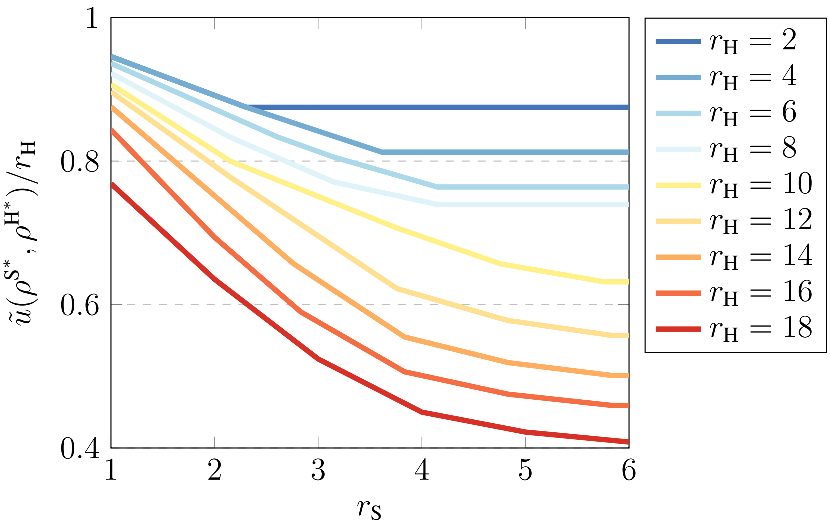

Figure 2.2: Fraction of undetected items in equilibrium as a function of r_S, for different values of r_H.

Specifically, Figure 2.2 compares the fractions of undetected items in equilibrium for different amounts of players' resources. It shows that for every fixed r_H, the gain in performance by having additional inspection resources exhibits diminishing returns. However, as r_S increases, the reduction in undetection at different values of r_H becomes more significant, emphasizing the value of efficiently coordinating the inspection resources. For $r_H \leq 8$, we observe that the undetection reaches a plateau as Regime Pattern 1 emerges when r_S is sufficiently large. In this regime, S fully inspects all locations where H hides resources, and no improvement can be achieved by utilizing more inspection resources. When $r_H \geq 10$, only Regime Pattern 2 and Regime Pattern 3 emerge, for which H hides some items at locations inspected with marginal probabilities strictly smaller than 1. This results in the undetection strictly decreasing in r_S, with changes dictated by the analytical expressions in Theorem 1.

This analysis illustrates the complex interplay between the problem parameters (detection probabilities, hiding capacities, amounts of resources), and quantifies their impact on the players' decisions in equilibrium. More generally, our closed-form expressions can be utilized by security agencies for tactical decisions such as investing in additional inspection resources. As shown by Figure 2.2, adding inspection resources will have a different impact on detection performance based on the number of hidden items, and must be carefully accounted for. Our analytical study can also be used to determine the value of increasing the detection probabilities by upgrading the detection technology. Overall, our analysis can be utilized by decision makers interested in finding a desired tradeoff between cost of inspection operations and detection performance.

2.4 Equilibrium Computation

In Theorem 1, we derived analytical expressions for pure equilibria $(\rho^{\mathrm{S}^*}, \rho^{\mathrm{H}^*}) \in \widetilde{\mathcal{A}}_{\mathrm{S}} \times \widetilde{\mathcal{A}}_{\mathrm{H}}$ of the continuous game $\widetilde{\Gamma}$. From Proposition 1, ρ^{S^*} and ρ^{H^*} respectively represent marginal inspection probabilities and expected numbers of hidden items at each location in equilibrium of the game Γ. To solve Γ, we must now determine the coordination of r_{S} inspection resources and r_{H} items to hide in order to satisfy these quantities. Specifically, we seek to efficiently compute mixed strategies $(\sigma^{\mathrm{S}^*}, \sigma^{\mathrm{H}^*}) \in \Delta_{\mathrm{S}} \times \Delta_{\mathrm{H}}$ satisfying $\rho(\sigma^{\mathrm{S}^*}) = \rho^{\mathrm{S}^*}$ and $\rho(\sigma^{\mathrm{H}^*}) = \rho^{\mathrm{H}^*}$. From Proposition 1, this will ensure that $(\sigma^{\mathrm{S}^*}, \sigma^{\mathrm{H}^*})$ is an equilibrium of Γ.

Then, we aim to solve the following generic resource coordination problem: Given a capacity vector $b \in \mathbb{Z}_{>0}^n$, a resource budget $r \in \mathbb{Z}_{>0}$, and a prescribed expected resource allocation vector $\rho \in \widetilde{\mathcal{A}}(b, r)$, find a feasible solution of $\{\sigma \in \mathbb{R}_{\geq 0}^{\mathcal{A}(b,r)} : \sum_{z \in \mathcal{A}(b,r)} \sigma_z z = \rho, \ \sum_{z \in \mathcal{A}(b,r)} \sigma_z = 1\}$, which is guaranteed to exist by Lemma 1. Although this problem involves an exponential number of variables, Carathéodory's theorem proves the existence of a solution with a support of size at most $n+1$. In fact, this problem can be solved in polynomial time using the ellipsoid method. However, this method is known to be practically inefficient [92]. Hence, we derive another algorithm to efficiently solve the feasibility problem.

Specifically, we extend the algorithm proposed by [15] that computes in time $O(n^2)$

a probability distribution with linear support over the set $\{z \in \{0,1\}^n : \sum_{i=1}^n z_i = r\}$ consistent with prescribed unidimensional marginal probabilities that sum up to *exactly* r resources. The general idea of each iteration of our algorithm is to express a given vector ρ in $\widetilde{\mathcal{A}}(b,r)$ as a convex combination of a vector in $\mathcal{A}(b,r)$ (i.e., with only integer components) and a vector in $\widetilde{\mathcal{A}}(b,r)$ with one more integer component than ρ has. To this end, the algorithm initially allocates $\lfloor \rho_i \rfloor$ resources at each location $i \in [\![1,n]\!]$, and then determines where to allocate some of the remaining $\bar{r} := r - \sum_{i=1}^n \lfloor \rho_i \rfloor$ resources given the fractional part of each component of ρ, defined as $\bar{\rho}_i := \rho_i - \lfloor \rho_i \rfloor$. If $\sum_{i=1}^n \bar{\rho}_i = \bar{r}$, one can use the algorithm of [15] to coordinate the allocation of the $\bar{r}$ remaining resources. However, their algorithm does not handle when $\sum_{i=1}^n \bar{\rho}_i$ is noninteger, as in Regime Pattern 1 (see Example 1 in Appendix A.1). Such instances arise when $\sum_{i=1}^n \bar{\rho}_i < \bar{r}$ and some allocations must necessarily deploy *fewer* than $\bar{r}$ resources.

To handle the general case, our algorithm extension determines the maximum number of locations q to allocate the remaining resources $\bar{r}$. Naturally, q is upper bounded by $\bar{r}$ and the number of positive components of $\bar{\rho}$. Then, the algorithm carefully assigns positive probability to a resource allocation that first assigns $\lfloor \rho_i \rfloor$ resources at each location, and then assigns one additional resource at each of the q locations with highest fractional parts $\bar{\rho}_i$. The algorithm then updates the vector $\bar{\rho}$ so that the new vector contains at least one more integer component than $\bar{\rho}$ has. The algorithm iterates until $\bar{\rho} \in \{0,1\}^n$, at which point it assigns the remaining probability to $\lfloor \rho \rfloor + \bar{\rho} \in \mathcal{A}(b,r)$. We refer the reader to Algorithm 1 for the detailed pseudocode.

We note that at each iteration of the while loop (Lines 3–15), the algorithm selects e and δ so that the current vector $\bar{\rho}$ can be expressed as a convex combination of $e \in \{0,1\}^n$ and the new vector $\bar{\rho}$ (Line 14), which has at least one more integer component than the current $\bar{\rho}$ does. This guarantees that the algorithm terminates. At termination, the algorithm expresses ρ as a convex combination of resource allocations in $\mathcal{A}(b,r)$. This can be translated into the following theorem:

Algorithm 1: Resource Coordination

Input : A capacity vector $b \in \mathbb{Z}^n_{>0}$, a resource budget $r \in \mathbb{Z}_{>0}$, and an expected resource allocation vector $\rho \in \widetilde{\mathcal{A}}(b, r)$

Output: A probability distribution $\sigma \in \Delta(b, r)$ satisfying $\sum_{z \in \mathcal{A}(b,r)} \sigma_z z = \rho$

$\sigma \leftarrow \mathbf{0}_{\mathcal{A}(b,r)}, \quad \bar{\rho} \leftarrow \rho - \lfloor \rho \rfloor, \quad \bar{r} \leftarrow r - \sum_{i=1}^{n} \lfloor \rho_i \rfloor$
$\gamma \leftarrow 1$
while $\bar{\rho} \notin \{0, 1\}^n$ **do**
  $\theta \leftarrow$ Permutation of $[\![1, n]\!]$ such that $\bar{\rho}_{\theta(1)} \geq \cdots \geq \bar{\rho}_{\theta(n)}$
  $q \leftarrow \min\{\bar{r}, |\{i \in [\![1, n]\!] : \bar{\rho}_i > 0\}|\}$
  if $q < n$ **then**
    $\delta \leftarrow \min\{\bar{\rho}_{\theta(q)}, 1 - \bar{\rho}_{\theta(q+1)}\}$
  else
    $\delta \leftarrow \bar{\rho}_{\theta(q)}$
  $e \leftarrow \mathbf{0}_n$
  foreach $j \in \{1, \dots, q\}$ **do**
    $e_{\theta(j)} \leftarrow 1$
  $\sigma_{\lfloor \rho \rfloor + e} \leftarrow \sigma_{\lfloor \rho \rfloor + e} + \gamma\delta$
  $\bar{\rho} \leftarrow \frac{1}{1-\delta}(\bar{\rho} - \delta e)$
  $\gamma \leftarrow \gamma(1 - \delta)$
$\sigma_{\lfloor \rho \rfloor + \bar{\rho}} \leftarrow \sigma_{\lfloor \rho \rfloor + \bar{\rho}} + \gamma$
return σ

Theorem 2. *Given a capacity vector $b \in \mathbb{Z}^n_{>0}$, a resource budget $r \in \mathbb{Z}_{>0}$, and an expected resource allocation vector $\rho \in \widetilde{\mathcal{A}}(b, r)$, Algorithm 1 returns a probability distribution $\sigma \in \Delta(b, r)$ satisfying $\sum_{z \in \mathcal{A}(b,r)} \sigma_z z = \rho$. Furthermore, σ has a support of size at most $n + 1$ and is computed in time $O(n^2)$.*

Thus, Algorithm 1 matches the support size guaranteed by Carathéodory's theorem on the polytope $\widetilde{\mathcal{A}}(b, r)$. The algorithm performs at most n iterations. Furthermore, by reutilizing the sorting of $\bar{\rho}$ of the previous iteration, we can implement each iteration (except for the first one) in time $O(n)$, which guarantees an overall running time of $O(n^2)$.

We can now summarize the overall solution approach for solving the game Γ. First, marginal inspection probabilities and expected numbers of hidden items in each location are computed according to Theorem 1: Sorting the detection potentials $(p_i c_i)_{i \in [\![1,n]\!]}$ and the detection probabilitites $(p_i)_{i \in [\![1,n]\!]}$ in non-decreasing order requires $O(n \log n)$ steps. Then, for every $i \in [\![0, n-1]\!]$, the mapping π^i, parameters k_i and ℓ_i, and thresholds τ_i and ν_i

can be computed in time $O(n)$. Identifying the subinterval $[\tau_{i^*-1}, \nu_{i^*}]$ or $[\nu_{i^*}, \tau_{i^*}]$ in which r_{H} belongs can be performed in $O(\log n)$ steps, and evaluating the expressions from Theorem 1 requires $O(n)$ more steps. Therefore, computing equilibrium marginal inspection probabilities and expected numbers of hidden items can be implemented in time $O(n^2)$. Finally, Algorithm 1 computes the mixed strategies consistent with these unidimensional quantities in time $O(n^2)$. Hence, we obtain the final result:

Corollary 1. *The game* Γ *can be solved in time* $O(n^2)$ *with equilibrium strategies of support size at most* $n + 1$ *each.*

Thus, we obtain an efficient solution approach for solving the large-scale hide-and-seek game Γ with multiple resources and heterogeneous locations. Furthermore, it provides solutions with small supports that can be easily implemented by security decision makers.

2.5 Concluding Remarks

In this chapter, we investigated a hide-and-seek game in which a seeker inspects locations to find items hidden by a hider. We extended previous models in the literature by considering the coordination of multiple resources for both players in locations with heterogeneous hiding capacities and probabilities of detecting hidden items. The objective of the seeker (resp. hider) is to minimize (resp. maximize) the expected number of undetected items. To compute equilibrium strategies and the value of this large-scale zero-sum game, we adopted a solution approach that first derives analytical equilibrium properties and then efficiently coordinates the players' resources. In particular, we showed that the marginal inspection probabilities and expected numbers of hidden items in each location in equilibrium form a pure equilibrium of a continuous zero-sum game. By carefully selecting a set of parameters and thresholds, we analytically solved this continuous game. Then, we derived a quadratic time algorithm that coordinates the players' resources to satisfy the characterized equilibrium marginal distributions and computes an equilibrium of the hide-and-seek game with linear support.

Our analysis revealed new equilibrium behaviors previously absent from the literature on search games, as a result of our game's features. Generally when detection potentials are heterogeneous, as the number of hidden items increases, the seeker should refrain from inspecting locations with smallest detection potentials and instead concentrate the inspection efforts on the remaining locations for increased detection performance. Within the inspected locations, the seeker should deterministically inspect the least reliable ones and randomize her remaining resources across the most reliable locations. In response, the hider will saturate the locations that are the most attractive and randomize his remaining resources among the other locations. These decisions are dictated by a complex interplay between the game parameters, captured by our closed-form solutions. They show the value of efficiently coordinating multiple inspection resources while accounting for the multiple items hidden in locations with heterogeneous capacities and detection probabilities.

More broadly, our insights and solution approach can be used to inform security agencies interested in scheduling multiple inspections to detect and seize illegal commodities hidden by a criminal organization. Our two-step approach can be utilized to coordinate inspection units following a randomized weekly schedule of operations. Furthermore, our parametric analysis can guide tactical decisions to invest in additional or upgraded inspection resources to achieve an optimal tradeoff between inspection costs and detection performance.

This work can be extended in multiple directions by considering additional features that would widen the applicability of our results. One extension is to consider different commodities hidden by the hider with heterogeneous values and detection probabilities, as well as different inspection resources with different detection characteristics. Another extension is to incorporate heterogeneous valuations of the locations to extend the original setting by von Neumann [36]. Finally, an interesting research direction is to consider a repeated version of the hide-and-seek game with players learning the initially unknown characteristics of their opponent while they interact.

CHAPTER 3

HIDE-AND-SEEK GAME WITH ASYMMETRIC VALUATIONS

3.1 Introduction

In this chapter, we examine a variant of the *hide-and-seek* game where two competing players, a seeker and a hider, coordinate multiple resources in a set of discrete locations. Namely, the hider conceals multiple items within the locations, while the seeker allocates inspection resources over a subset of locations to uncover hidden items. Both players are constrained to allocate at most one resource per location, and determine their allocations without observing the other player's decisions. Furthermore, the players hold asymmetric and heterogeneous valuations over the locations. Specifically, each location is associated with two (possibly distinct) positive utility values, one for the seeker and another for the hider. The seeker (resp. hider) aims to maximize the additive utility obtained from the inspected (resp. uninspected) locations containing hidden items. We are interested in efficiently computing mixed-strategy Nash Equilibria (NE) of this large-sacale nonzero-sum game, and gather insights of the players' behavior in equilibirum.

Hide-and-seek games generalize the classical Matching Pennies game and serve as fundamental models of strategic mismatch. Their applications include security domains, where a security agency schedules inspections to detect illegal commodities concealed by a criminal organization in containers [93]; animal behavior, capturing the pursue-evasion dynamics between a prey and a predator [32]; election audits to prevent result manipulation [43]; political campaigns, where a political party benefits from campaigning in districts not matching an incumbent [53], and market entry games, where firms decide whether to enter a market challenge an established competitor [94].

Our model extends the hide-and-seek game of [15], where both players engage in the same interactions, except that they share identical—i.e. symmetric—valuations for each location. This symmetry feature renders their model equivalent to a zero-sum game. While Dziubiński and Roy provide insights into the players' equilibrium behaviors, essentially showing that both players compete for the most valued locations, it remains unclear how they would allocate their resources in scenarios where the most valued locations for one player do not align with those of their opponent. Furthermore, the asymmetry in the valuations in our model results in a nonzero-sum game, wherein minmax and maxmin strategies are not necessarily NE, ruling out solution approaches based on linear programming ([89], see also [90]) or efficient approximation algorithms designed for zero-sum games [45, 46, 47].

Addressing asymmetric valuations enables a better understanding of the players' behaviors in various scenarios involving strategic mismatches. For instance, in security domains, a criminal organization hiding illegal items in port containers may value concealment based on the time invested and the likelihood of successful smuggling, whereas a security agency's inspections are influenced by the potential deterrent effect on future smuggling activities. Similarly, in the interaction between a prey and a predator, the prey values the situation in terms of the inherent risk to its survival, while the predator considers it merely a choice of its next meal. In political campaigns, asymmetric valuations can capture the differing strategic priorities of political parties, where one party might focus on campaigning in districts that are crucial for swinging votes, while the opposing party might prioritize strongholds to consolidate their base.

Therefore, in this chapter we investigate the following research questions: *(i) How do rational agents coordinate multiple homogeneous resources in scenarios of strategic mismatch with additive utilities, considering they hold asymmetric and heterogeneous valuations over the locations where their resources are allocated? (ii) How are the NE jointly impacted by the players' valuations and number of resources?*

3.1.1 Contributions

In this chapter, we present the following contributions:

- We investigate a hide-and-seek game involving multiple resources for both the seeker and the hider, and where both players hold asymmetric and heterogeneous valuations over the hiding locations. This extends previous models in the literature [32, 15, 93] to encompass simultaneous-move nonzero-sum games. Our study advances the understanding of rational agent behavior in scenarios characterized by strategic mismatch, including those arising in security domains, animal behavior, political campaigns, electoral fraud detection, and beyond.

- We provide an analytical characterization of the NE for our hide-and-seek game (Theorem 3). To this end, we first represent the seeker and hider's strategies in terms of the marginal probabilities of inspecting and hiding an item at each location, respectively. This alternative representation allows us to analyze an equivalent lower-dimensional continuous game, which we solve analytically. Our analysis involves carefully partitioning the set of locations and formulating a special class of marginal inspection and hiding strategies that reflect the players' incentives, driven by their asymmetric and heterogeneous valuations and their resource constraints. These strategies enable us to derive two key parameters that are crucial for describing the NE of the game. We then identify a set of threshold values that classify different equilibrium regimes based on parametric intervals for the number of hider's resources and provide closed-form expressions for the marginal inspection and hiding probabilities at each location in equilibrium.

- We provide game-theoretic insights into how the players select their marginal inspection and hiding strategies in each equilibrium regime, elucidating how the aforementioned key game parameters and threshold values determine the feasibility and equilibrium conditions of each regime. Our equilibrium characterization results in an $O(n \log n)$ algorithm for computing NE's marginal inspection and hiding probabilities, extending the qualitative and algorithmic results of [15] to the nonzero-sum setting. Using previous results for computing mixed strategies with linear support size that satisfy prescribed marginal distributions [93], this yields an $O(n^2)$ algorithm for computing mixed-strategy NE of our hide-and-seek game.

3.1.2 Literature Review

The hide-and-seek game was first introduced by von Neumann [36] as a zero-sum game in a square matrix with nonnegative entries. In this model, the hider selects an entry a_{ij}, while the seeker chooses a row or column. If the chosen row or column contains the hider's entry, the hider pays a_{ij} to the seeker; otherwise, the seeker pays a_{ij} to the hider. The equilibria of this game have also been analyzed by [37] and [95].

Hide-and-seek games, which generalize the classical Matching Pennies game, have been extensively applied across various strategic contexts where one party, the seeker (resp. hider), benefits from matching (resp. mismatching) their opponents' strategies. In the field of animal behavior, [32] examined a zero-sum pursuit-evasion game between a prey (hider) and a predator (seeker), where the predator inspects multiple locations to find the prey, who hides in one of these locations. In their model, detection is imperfect; even if the predator inspects the location where the prey is hiding, capture occurs with a probability that depends on that specific location. [93] extended this model to the security domain, where a security agency schedules inspections to detect illegal commodities hidden by a criminal organization in containers. To this aim, they introduced a hider who allocates multiple items across locations with heterogeneous hiding capacities and imperfect detection capabilities.

Another application of hide-and-seek games arises in the field of steganography, which involves concealing information within a message or physical object that is then sent to an intended recipient. [96] analyzed a game where a steganographer (hider) flips a limited number of bits in a file. Subsequently, Nature sends either the original or the modified file to a steganalyst (seeker), who must determine whether the file she received is the original one by querying one of its bits. [97] aimed to analyze the phenomenon of fashion, interpreting it as the result of interactions between rebels (hiders) and conformists (seekers). To this end, they extended the Matching Pennies game to a multiplayer setting within a network, where each node represents a player and the edges represent their social interactions. Election audits, as studied by [43], utilize strategic mismatch principles to design audits that prevent result manipulation, ensuring electoral integrity. Political campaigns also leverage these strategic insights, as noted by [53], as a political party can gain an advantage by campaigning in districts that do not align with an incumbent's strengths.

Within the current literature, the most closely related game to ours is that of [15], who examined a hide-and-seek game where both the hider and seeker allocate multiple resources across a set of discrete locations. However, unlike our model, both players share identical valuations for each location, rendering their strategic interactions equivalent to a zero-sum game. A similar variant, also a zero-sum game but with the hider restricted to having a single resource, was analyzed by [33], who extended a model by [59]. In this variant, the hider can still receive utility when his resource is found by the seeker, although it is smaller than the utility he would receive if the resource remains undetected. Another related model is the nonzero-sum inspection game studied by [35] [see also 98, 99], where an inspector (seeker) examines a set of locations to identify violations incurred by the employers of an inspectee (hider). In this game, it is assumed that the inspector's primary goal is to incentivize compliance at each location rather than to profit from detected violations, while the inspectee profits (resp. loses) from undetected (resp. detected) violations at each location.

Our game, along with several of the aforementioned games, belongs to the richer class of security games, where a defender allocates multiple defending resources to cover a set of targets, aiming to protect them from the actions of a malicious attacker [see, for example, 5, 86, 7, 100, 101, 12, 50, 102, and the references therein]. Drawing an analogy with our game, the defender is the seeker, the attacker is the hider, and the targets represent the hiding locations. Two key features render these games significantly more challenging to analyze than ours. First, the defender's resources can be allocated to *schedules*, meaning a single resource can cover multiple targets simultaneously, allowing for potentially overlapping coverage. Second, each target is associated with *two* utility values for each player. Specifically, for each attacked target, the defender (resp. attacker) receives a certain utility if the target is covered, and a smaller (resp. higher) utility if the target is uncovered. In this context, our game is a special instance of a security game with unit-sized schedules—i.e., each defender resource can cover a single target—where the defender (resp. attacker) gains positive utility from each match (resp. mismatch) and no utility from each mismatch (resp. match). Despite their simpler structure, unit-sized schedules have been effectively utilized in real-world applications of security games [see, e.g., 5, 100]. Nevertheless, a comprehensive understanding of the players' behaviors in equilibrium in security games with unit-sized schedules has not been fully established.

While most of the literature has focused on the Stackelberg version of security games—where the defender commits to a strategy first, and the attacker selects his move next—there have been fewer studies analyzing the simultaneous-move version—as in our setting—where NE is the standard solution concept [8, 103, 104, 34]. However, these works have primarily focused on developing efficient algorithms to compute NE for security games with unit-sized schedules, with less emphasis on the understanding of the incentives that drive the players' behaviors in equilibrium. For instance, in the context of such security games with additive utilities and multiple players' resources, [8] provided a $O(n^2)$ time algorithm to compute the NE's marginal probabilities of attacking and defending each target—where n is the number of targets. Nonetheless, their approach is intricate, and it provides limited intuition regarding the players' incentives or equilibrium behavior. Using different methods, [104] and [34] derived structural formulas for such NE marginal distributions, revealing multiple equilibrium regimes. However, it remains unclear from their analysis how the various game parameters jointly determine the specific equilibrium regime in which the players engage.

On the other hand, [15] derived an analytical characterization of the equilibrium's marginal inspection and hiding probabilities for their zero-sum hide-and-seek game. Their analysis reveals qualitative features of the equilibria, and yields a $O(n \log n)$ time algorithm to compute the equilibrium marginal distributions. However, it is not obvious how their analysis extends to asymmetric players' valuations, or whether the complexity bound remains valid beyond the zero-sum framework. Similarly, [35] provided an $O(n \log n)$ algorithm for the aforementioned inspection game and described its NE in terms of the algorithm's outputs. In contrast, our approach delves into the players' incentives to derive key game parameters, enabling us to analytically characterize the NE's marginal resource allocations. This characterization, in turn, leads to an efficient algorithm for solving our game.

Finally, another important class of games involving strategic mismatch are Colonel-Blotto games [see, e.g., 105, 106, 107, 43, 108, 109]. In these games, players allocate multiple resources across various locations, denoted as battlefields. In contrast to our game, Colonel-Blotto games allow both players to allocate multiple resources to each location. The utility associated with each battlefield is awarded to the player who allocates the highest number of resources there, with predefined tie-breaking rules. The structure of these games requires different techniques than those presented in this chapter for equilibrium computation.

Outline. The rest of this chapter is organized as follows. In Section 3.2, we formulate our hide-and-seek game. In Section 3.3 we review preliminary results that will allow us to address our game. In Section 3.4, we analytically characterize the NE of a lower-dimensional continuous game, providing closed-form expressions and insights for the marginal inspection and hiding probabilities across locations in equilibrium. We then provide some concluding remarks in Section 3.5. Finally, the proofs of our results are listed in Appendix B.1.

3.2 Model Description

We consider the following variant of the hide-and-seek game involving a seeker and hider, who allocate resources over a set of *locations* $\mathcal{T}$. Specifically, the seeker (resp. hider) allocates multiple homogeneous *inspection resources* (resp. *hiding resources* or *items*) within the locations. We let $n := |\mathcal{T}|$ be the number of locations, and for $i, j \in \mathbb{Z}_{\geq 0}$, we denote by $[\![i, j]\!] := \{i, i+1, \ldots, j\}$ the discrete interval between i and j.

We assume that both the seeker and the hider are strategic, leading us to study their behaviors within a game-theoretic framework. To this end, we define a simultaneous-move bimatrix game Γ between the seeker (S, *she*) and the hider (H, *he*). S can select up to $r_{\mathrm{S}} \in \mathbb{Z}_{>0}$ locations to inspect, while H can select up to $r_{\mathrm{H}} \in \mathbb{Z}_{>0}$ resources to hide. A pure strategy for S (resp. H) consists of a resource allocation of at most r_{S} (resp. r_{H}) inspection (resp. hiding) resources to the locations, constrained to at most one resource in each location. We denote by $\mathcal{A}_{\mathrm{S}} := \left\{x \in \{0,1\}^{\mathcal{T}} : \sum_{t \in \mathcal{T}} x_t \leq r_{\mathrm{S}}\right\}$ and $\mathcal{A}_{\mathrm{H}} := \left\{y \in \{0,1\}^{\mathcal{T}} : \sum_{t \in \mathcal{T}} y_t \leq r_{\mathrm{H}}\right\}$ the action sets for S and H, respectively. For each $x \in \mathcal{A}_{\mathrm{S}}$ and $t \in \mathcal{T}$, we say that location t is *inspected* if $x_t = 1$, and *uninspected* if $x_t = 0$. Similarly, for each $y \in \mathcal{A}_{\mathrm{H}}$ and $t \in \mathcal{T}$, we say location t *contains a hidden item* if $y_t = 1$ and is *empty* if $y_t = 0$.

We assume that locations not containing hidden items do not affect either player's utility, and that each player's utility is additive over the locations containing hidden items. Namely, for each location t, S receives a utility $U_{\mathrm{S}}(t) > 0$ if the location contains a hidden item and is inspected. Similarly, when hiding an item in a location t that is uninspected by S, H receives utility $U_{\mathrm{H}}(t) > 0$. Alternatively, we refer to $U_{\mathrm{S}}(t)$ (resp. $U_{\mathrm{H}}(t)$) as S's (resp. H's) *valuation* for location t. Then, with a slight abuse of notation, for each pure strategy profile $(x, y) \in \mathcal{A}_{\mathrm{S}} \times \mathcal{A}_{\mathrm{H}}$, the overall utilities of S and H are respectively $U_{\mathrm{S}}(x, y) := \sum_{t \in \mathcal{T}} U_{\mathrm{S}}(t) x_t y_t$ and $U_{\mathrm{H}}(x, y) := \sum_{t \in \mathcal{T}} U_{\mathrm{H}}(t)(1 - x_t) y_t$. Thus, the payoff matrices of S and H in the game Γ are respectively given by $[U_{\mathrm{S}}(x,y)]_{x\in\mathcal{A}_{\mathrm{S}},\, y\in\mathcal{A}_{\mathrm{H}}}$ and $[U_{\mathrm{H}}(x,y)]_{x\in\mathcal{A}_{\mathrm{S}},\, y\in\mathcal{A}_{\mathrm{H}}}$. We assume both players have perfect information regarding their own and their opponents' valuations and action sets, but determine their allocations without observing their opponents' decisions.

We allow the players to use mixed strategies. A mixed strategy for S (resp. H) is a probability distribution over the set of pure strategies $\mathcal{A}_{\mathrm{S}}$ (resp. $\mathcal{A}_{\mathrm{H}}$). We denote by $\Delta_{\mathrm{S}} := \left\{\sigma^{\mathrm{S}} \in [0,1]^{\mathcal{A}_{\mathrm{S}}} : \sum_{x\in\mathcal{A}_{\mathrm{S}}} \sigma^{\mathrm{S}}_x = 1\right\}$ and $\Delta_{\mathrm{H}} := \left\{\sigma^{\mathrm{H}} \in [0,1]^{\mathcal{A}_{\mathrm{H}}} : \sum_{y\in\mathcal{A}_{\mathrm{H}}} \sigma^{\mathrm{H}}_y = 1\right\}$ the sets of mixed strategies for S and H, respectively. Then, for each mixed strategy profile $(\sigma^{\mathrm{S}}, \sigma^{\mathrm{H}}) \in \Delta_{\mathrm{S}} \times \Delta_{\mathrm{H}}$, the expected utilities of S and H are given by $\mathbf{U}_{\mathrm{S}}(\sigma^{\mathrm{S}}, \sigma^{\mathrm{H}}) := \sum_{x\in\mathcal{A}_{\mathrm{S}}} \sum_{y\in\mathcal{A}_{\mathrm{H}}} \sigma^{\mathrm{S}}_x \, \sigma^{\mathrm{H}}_y \, U_{\mathrm{S}}(x,y)$ and $\mathbf{U}_{\mathrm{H}}(\sigma^{\mathrm{S}}, \sigma^{\mathrm{H}}) := \sum_{x\in\mathcal{A}_{\mathrm{S}}} \sum_{y\in\mathcal{A}_{\mathrm{H}}} \sigma^{\mathrm{S}}_x \, \sigma^{\mathrm{H}}_y \, U_{\mathrm{H}}(x,y)$, respectively.

Our game has practical applications in various scenarios of strategic mismatch, particularly in security settings. For instance, let us consider a metropolitan transit authority responsible for preventing pickpocketing activities by a criminal organization within a city's subway system. Here, the set of locations $\mathcal{T}$ includes stations, platforms, tunnels, and entrances, each with different levels of security and passenger traffic. H, representing the pickpocketing organization, targets locations where the likelihood of successfully stealing from passengers is high, such as less monitored stations, busy entrances, or crowded transfer points. Conversely, S, represented by the security team, prioritizes locations that can effectively deter future pickpocketing attempts, including high-visibility stations and entrances where security measures can be prominently displayed. The security team aims to find a randomized inspection strategy that maximizes the overall deterrent effect, while the criminal organization goal is to devise a randomized allocation of pickpockets to maximize their overall success.

We are interested in the study of strategy profiles $(\sigma^{\mathrm{S}^*}, \sigma^{\mathrm{H}^*}) \in \Delta_{\mathrm{S}} \times \Delta_{\mathrm{H}}$ that represent

Nash equilibria (NE) of the game Γ, i.e., that satisfy

$$\begin{aligned}\mathbf{U}_{\mathrm{S}}(\sigma^{\mathrm{S}^*}, \sigma^{\mathrm{H}^*}) &\geq \mathbf{U}_{\mathrm{S}}(\sigma^{\mathrm{S}}, \sigma^{\mathrm{H}^*}), \quad \forall\, \sigma^{\mathrm{S}} \in \Delta_{\mathrm{S}},\\ \mathbf{U}_{\mathrm{H}}(\sigma^{\mathrm{S}^*}, \sigma^{\mathrm{H}^*}) &\geq \mathbf{U}_{\mathrm{H}}(\sigma^{\mathrm{S}^*}, \sigma^{\mathrm{H}}), \quad \forall\, \sigma^{\mathrm{H}} \in \Delta_{\mathrm{H}}.\end{aligned}$$

Since Γ is a finite bimatrix game, a NE is guaranteed to exist ([110], see also [111]), although it may not be unique. Nonetheless, due to the combinatorial number of players' strategies, standard algorithms for computing NE in bimatrix games [e.g., 112, 113] suffer from the curse of dimensionality, rendering the game Γ challenging to solve. Therefore, we leverage its structure to derive an efficient and intuitive solution approach.

3.3 Preliminaries

The strategies of the game Γ can be equivalently represented in terms of the marginal probabilities that each location is inspected or contains a hidden item. Namely, for every mixed inspection strategy $\sigma^{\mathrm{S}} \in \Delta_{\mathrm{S}}$, we let $\rho(\sigma^{\mathrm{S}})$ be the vector in $\mathbb{R}^{\mathcal{T}}$ whose t-th component is given by $\rho_t(\sigma^{\mathrm{S}}) \coloneqq \sum_{x \in \mathcal{A}_{\mathrm{S}}} \sigma^{\mathrm{S}}_x\, x_t$. Similarly, for every mixed hiding strategy $\sigma^{\mathrm{H}} \in \Delta_{\mathrm{H}}$, we let $\rho(\sigma^{\mathrm{H}})$ be the vector in $\mathbb{R}^{\mathcal{T}}$ whose t-th component is given by $\rho_t(\sigma^{\mathrm{H}}) \coloneqq \sum_{y \in \mathcal{A}_{\mathrm{H}}} \sigma^{\mathrm{H}}_y\, y_t$. Then, $\rho_t(\sigma^{\mathrm{S}})$ (resp. $\rho_t(\sigma^{\mathrm{H}})$) is the marginal probability that location t is inspected (resp. contains a hidden item) when S (resp. H) plays σ^{S} (resp. σ^{H}). Due to the additive structure of the functions $U_{\mathrm{S}}(x, y)$ and $U_{\mathrm{H}}(x, y)$, we can represent the players' expected utilities in terms of these marginal probabilities.

Lemma 3. *For every mixed strategy profile* $(\sigma^{\mathrm{S}}, \sigma^{\mathrm{H}}) \in \Delta_{\mathrm{S}} \times \Delta_{\mathrm{H}}$*, S's and H's expected utilities respectively satisfy* $\mathbf{U}_{\mathrm{S}}(\sigma^{\mathrm{S}}, \sigma^{\mathrm{H}}) = \sum_{t \in \mathcal{T}} U_{\mathrm{S}}(t)\, \rho_t(\sigma^{\mathrm{S}})\, \rho_t(\sigma^{\mathrm{H}})$ *and* $\mathbf{U}_{\mathrm{H}}(\sigma^{\mathrm{S}}, \sigma^{\mathrm{H}}) = \sum_{t \in \mathcal{T}} U_{\mathrm{H}}(t) \left(1 - \rho_t(\sigma^{\mathrm{S}})\right) \rho_t(\sigma^{\mathrm{H}})$.

From Lemma 3, the players' expected utilities $\mathbf{U}_{\mathrm{S}}(\sigma^{\mathrm{S}}, \sigma^{\mathrm{H}})$ and $\mathbf{U}_{\mathrm{H}}(\sigma^{\mathrm{S}}, \sigma^{\mathrm{H}})$ are given by the sums of the players' expected utilities obtained from each location. Specifically, at each location $t \in \mathcal{T}$, S obtains a expected utility of $U_{\mathrm{S}}(t)\,\rho_t(\sigma^{\mathrm{S}})\,\rho_t(\sigma^{\mathrm{H}})$, while H gets a expected utility of $U_{\mathrm{H}}(t)\left(1-\rho_t(\sigma^{\mathrm{S}})\right)\rho_t(\sigma^{\mathrm{H}})$.

Conversely, given a vector ρ^{S} (resp. ρ^{H}) $\in \mathbb{R}^{\mathcal{T}}$, a mixed strategy $\sigma^{\mathrm{S}} \in \Delta_{\mathrm{S}}$ (resp. $\sigma^{\mathrm{H}} \in \Delta_{\mathrm{H}}$) is *compatible* with ρ^{S} (resp. ρ^{H}) if $\rho(\sigma^{\mathrm{S}}) = \rho^{\mathrm{S}}$ (resp. $\rho(\sigma^{\mathrm{H}}) = \rho^{\mathrm{H}}$). In other words, σ^{S} (resp. σ^{H}) is compatible with ρ^{S} (resp. ρ^{H}) if the latter is the vector of marginal inspection (resp. hiding) probabilities resulting from the former. Compatibility enables a bidirectional mapping between the spaces of mixed strategies and marginal probabilities. The following lemma states necessary and sufficient conditions for compatibility.

Lemma 4 (see Lemma 1 and Theorem 2)**.** *Consider a generic resource budget $r \in \mathbb{Z}_{>0}$, and a generic set of actions $\mathcal{A} \coloneqq \left\{z \in \{0,1\}^{\mathcal{T}} : \sum_{t\in\mathcal{T}} z_t \leq r\right\}$. Then, given a and a vector $\rho \in \mathbb{R}^{\mathcal{T}}$, there exists a probability distribution $\sigma \in \Delta \coloneqq \left\{\sigma \in [0,1]^{\mathcal{A}} : \sum_{z\in\mathcal{A}} \sigma_z = 1\right\}$ satisfying $\rho_t(\sigma) = \rho_t$ for all $t \in \mathcal{T}$ if and only if $\rho \in [0,1]^{\mathcal{T}}$ and $\sum_{t\in\mathcal{T}} \rho_t \leq r$. Furthermore, such σ with a support of size at most $n+1$ can be computed in time $O(n^2)$.*

From Lemma 4, we can alternatively represent the mixed strategies of S and H through their vectors of marginal probabilities, given by the feasible solutions of the polytopes $\widetilde{\mathcal{A}}_{\mathrm{S}} \coloneqq \left\{\rho^{\mathrm{S}} \in [0,1]^{\mathcal{T}} : \sum_{t\in\mathcal{T}} \rho^{\mathrm{S}}_t \leq r_{\mathrm{S}}\right\}$ and $\widetilde{\mathcal{A}}_{\mathrm{H}} \coloneqq \left\{\rho^{\mathrm{H}} \in [0,1]^{\mathcal{T}} : \sum_{t\in\mathcal{T}} \rho^{\mathrm{H}}_t \leq r_{\mathrm{H}}\right\}$, respectively. This represents a computational advantage, as S's and H's mixed strategies require vectors of sizes $\sum_{k=0}^{r_{\mathrm{S}}} \binom{r_{\mathrm{S}}}{k}$ and $\sum_{k=0}^{r_{\mathrm{H}}} \binom{r_{\mathrm{H}}}{k}$ respectively, while the vectors in $\widetilde{\mathcal{A}}_{\mathrm{S}}$ and $\widetilde{\mathcal{A}}_{\mathrm{H}}$ are of size n. Additionally, given $\rho^{\mathrm{S}} \in \widetilde{\mathcal{A}}_{\mathrm{S}}$ (resp. $\rho^{\mathrm{H}} \in \widetilde{\mathcal{A}}_{\mathrm{H}}$), a compatible mixed strategy $\sigma^{\mathrm{S}} \in \Delta_{\mathrm{S}}$ (resp. $\sigma^{\mathrm{H}} \in \Delta_{\mathrm{H}}$) with support of size at most $n+1$ can be computed in time $O(n^2)$.

Consequently, for the remainder of this chapter we focus our study on the continuous game $\widetilde{\Gamma}$ between S and H, who respectively select vectors of marginal resource allocations $\rho^{\mathrm{S}} \in \widetilde{\mathcal{A}}_{\mathrm{S}}$ and $\rho^{\mathrm{H}} \in \widetilde{\mathcal{A}}_{\mathrm{H}}$, and where their payoffs are respectively given by

$$\mathbf{U}_{\mathrm{S}}(\rho^{\mathrm{S}},\rho^{\mathrm{H}}) \coloneqq \sum_{t\in\mathcal{T}} U_{\mathrm{S}}(t)\,\rho^{\mathrm{S}}_t\,\rho^{\mathrm{H}}_t, \quad \mathbf{U}_{\mathrm{H}}(\rho^{\mathrm{S}},\rho^{\mathrm{H}}) \coloneqq \sum_{t\in\mathcal{T}} U_{\mathrm{H}}(t)\left(1-\rho^{\mathrm{S}}_t\right)\rho^{\mathrm{H}}_t. \tag{3.1}$$

A *pure* NE of $\widetilde{\Gamma}$ is a marginal strategy profile $(\rho^{\mathrm{S}^*}, \rho^{\mathrm{H}^*}) \in \widetilde{\mathcal{A}}_{\mathrm{S}} \times \widetilde{\mathcal{A}}_{\mathrm{H}}$ satisfying

$$\mathbf{U}_{\mathrm{S}}(\rho^{\mathrm{S}^*}, \rho^{\mathrm{H}^*}) \geq \mathbf{U}_{\mathrm{S}}(\rho^{\mathrm{S}}, \rho^{\mathrm{H}^*}), \quad \forall\, \rho^{\mathrm{S}} \in \widetilde{\mathcal{A}}_{\mathrm{S}},$$
$$\mathbf{U}_{\mathrm{H}}(\rho^{\mathrm{S}^*}, \rho^{\mathrm{H}^*}) \geq \mathbf{U}_{\mathrm{H}}(\rho^{\mathrm{S}^*}, \rho^{\mathrm{H}}), \quad \forall\, \rho^{\mathrm{H}} \in \widetilde{\mathcal{A}}_{\mathrm{H}}.$$

In other words, none of the players can achieve a higher utility by marginally relocating their resources to other locations. The existence of a pure NE of the game $\widetilde{\Gamma}$ is guaranteed by a classical result of Debreu, Glicksberg and Fan [see, e.g., 114, Thm. 1.2]. We are interested in characterizing the pure NE of the continuous nonzero-sum hide-and-seek game $\widetilde{\Gamma}$.

3.4 Analytical Characterization of Nash Equilibrium Strategies

3.4.1 Preliminary Analysis

We first develop some intuition behind the players' equilibrium strategies in $\widetilde{\Gamma}$. From the payoff structure of $\widetilde{\Gamma}$ (3.1), the players' strategy decisions are driven by the expected utilities they can obtain from each location, given their opponent's marginal strategy. Specifically, given H's marginal hiding strategy $\rho^{\mathrm{H}} \in \widetilde{\mathcal{A}}_{\mathrm{H}}$, a best response for S consists in inspecting r_{S} locations with the highest values of $U_{\mathrm{S}}(t)\, \rho^{\mathrm{H}}_t$. Conversely, given S's marginal inspection strategy $\rho^{\mathrm{S}} \in \widetilde{\mathcal{A}}_{\mathrm{S}}$, a best response for H consists in hiding resources in r_{H} locations with the highest values of $U_{\mathrm{H}}(t)\left(1 - \rho^{\mathrm{S}}_t\right)$. We observe that these best responses may not be unique, as ties can occur in the quantities guiding the players' strategies. Additionally, we note that if $r_{\mathrm{S}} = n$, any NE involves S fully inspecting—i.e., with marginal probability 1—every location while H can adopt any hiding strategy. Therefore, for the remainder of this section, we assume $r_{\mathrm{S}} \in [\![1, n-1]\!]$.

To illustrate the intricacies of our model, let us first examine the feasibility of the players' equalizing strategies——that is, strategies that yield their opponent the same expected utility, irrespective of their actions. For S, this entails selecting a constant $\lambda_{\mathrm{S}} \in \mathbb{R}_{\geq 0}$ and setting $\rho_t^{\mathrm{S}^*} = 1 - \lambda_{\mathrm{S}}/U_{\mathrm{H}}(t)$ for every $t \in \mathcal{T}$, ensuring H's expected utility from every location where he hides an item is equal to λ_{S}. Similarly, H would select a constant $\lambda_{\mathrm{H}} \in \mathbb{R}_{\geq 0}$ and set $\rho_t^{\mathrm{H}^*} = \lambda_{\mathrm{H}}/U_{\mathrm{S}}(t)$ for every $t \in \mathcal{T}$, equalizing S's expected utility from each inspected location to λ_{H}. If in addition we let both players to fully utilize their resources—that is, $\sum_{t\in\mathcal{T}} \rho_t^{\mathrm{S}^*} = r_{\mathrm{S}}$ and $\sum_{t\in\mathcal{T}} \rho_t^{\mathrm{H}^*} = r_{\mathrm{H}}$—then the values of λ_{S} and λ_{H} are determined by

$$\lambda_{\mathrm{S}} = \frac{n - r_{\mathrm{S}}}{\sum_{t\in\mathcal{T}} \frac{1}{U_{\mathrm{H}}(t)}}, \qquad \lambda_{\mathrm{H}} = \frac{r_{\mathrm{H}}}{\sum_{t\in\mathcal{T}} \frac{1}{U_{\mathrm{S}}(t)}}.$$

However, due to the multiplicity of players' resources and the heterogeneity of their valuations, it is possible that the inequalities $\rho_t^{\mathrm{S}^*} \geq 0$ or $\rho_t^{\mathrm{H}^*} \leq 1$ may not hold for certain locations. These violations render ρ^{S^*} or ρ^{H^*} infeasible, as they must belong to $\widetilde{\mathcal{A}}_{\mathrm{S}}$ and $\widetilde{\mathcal{A}}_{\mathrm{H}}$, respectively. In such cases, S may prefer not to inspect locations with $\rho_t^{\mathrm{S}^*} < 0$, while H may opt to saturate the locations with $\rho_t^{\mathrm{H}^*} > 1$, setting their marginal hiding probabilities to 1.

Although S naturally prefers inspecting locations with the highest values of $U_{\mathrm{S}}(t)$, she must consider that her utility can only be derived from locations containing items hidden by H. Moreover, due to the asymmetry in the players' valuations, S's most preferred locations for inspection may not align with those most preferred by H for hiding items. Therefore, S must ensure that the locations she selects for inspection become enticing enough for H to hide items. The heterogeneity of H's valuations will play a crucial role in this endeavor, as it will be easier for S to incentivize H to hide items in locations with relatively higher values of $U_{\mathrm{H}}(t)$.

To address the challenges arising from the asymmetric and heterogeneous valuations, we will carefully define multiple partitions of the set of locations $\mathcal{T}$, each one comprising four sets, with certain sets being possibly empty. These partitions will help us characterize the players' behaviors in equilibrium. Along with these partitions, we will also define a special class of inspection and hiding strategies, and two key parameters that will be crucial for solving of our game.

A typical partition of $\mathcal{T}$ begins with a set $\mathcal{I}_i \coloneqq \{t_1, \dots, t_i\}$ containing i carefully selected locations, which we will formally define later in a sequential manner. For now, let us just assume that these locations hold relatively low value for S but relatively high value for H, and that t_i is S's most valued location in $\mathcal{I}_i$.

Locations in $\mathcal{I}_i$ are enticing for H to hide items, but not attractive enough for S to inspect due to their low values of $U_{\mathrm{S}}(t)$. Hence, S's incentive is to avoid inspecting these locations and instead allocate her inspection resources towards locations in $\mathcal{T} \setminus \mathcal{I}_i$. In this regard, the value of $U_{\mathrm{S}}(t_i)$ will serve as a threshold for S: She will only inspect locations in $\mathcal{T} \setminus \mathcal{I}_i$ satisfying $U_{\mathrm{S}}(t) \geq U_{\mathrm{S}}(t_i)$, while leaving the remaining locations uninspected.

On the other hand, S's selected locations to inspect should be among those that are enticing for H; these are H's most valued locations in $\mathcal{T} \setminus \mathcal{I}_i$. Thus, given $\mathcal{I}_i$, we sort the remaining $n - i$ locations in $\mathcal{T} \setminus \mathcal{I}_i$ by nonincreasing H's valuation. To this end, we let π^i be a labeling satisfying $\mathcal{T} \setminus \mathcal{I}_i = \left\{t_{\pi^i(1)}, \dots, t_{\pi^i(n-i)}\right\}$ and $U_{\mathrm{H}}(t_{\pi^i(1)}) \geq \cdots \geq U_{\mathrm{H}}(t_{\pi^i(n-i)})$. To account for edge cases in our analysis, we also let $U_{\mathrm{H}}(t_{\pi^i(n-i+1)}) \coloneqq 0$. The labeling π^i, together with an index $k \in \{1, \dots, n-i\}$, induce a partition of $\mathcal{T} \setminus \mathcal{I}_i$ into three more sets, thus completing the overall partition of $\mathcal{T}$ into the following four sets:

$$
\begin{aligned}
\mathcal{I}_i &\coloneqq \{t_1, \dots, t_i\}\,, \\
\mathcal{J}_i(k) &\coloneqq \left\{t \in \left\{t_{\pi^i(1)}, \dots, t_{\pi^i(k)}\right\} \,:\, U_{\mathrm{S}}(t) \geq U_{\mathrm{S}}(t_i)\right\}, \\
\widetilde{\mathcal{J}}_i(k) &\coloneqq \left\{t \in \left\{t_{\pi^i(1)}, \dots, t_{\pi^i(k)}\right\} \,:\, U_{\mathrm{S}}(t) < U_{\mathrm{S}}(t_i)\right\}, \\
\mathcal{K}_i(k) &\coloneqq \left\{t_{\pi^i(k+1)}, \dots, t_{\pi^i(n-i)}\right\}.
\end{aligned}
$$

In these definitions, the sets $\mathcal{J}_i(k)$ and $\widetilde{\mathcal{J}}_i(k)$ partition the set of H's k most valued locations in $\mathcal{T} \setminus \mathcal{I}_i$. Specifically, $\mathcal{J}_i(k)$ contains the locations within $\left\{t_{\pi^i(1)}, \dots, t_{\pi^i(k)}\right\}$ for which S's valuation is at least that of *any* location in $\mathcal{I}_i$—as they satisfy $U_{\mathrm{S}}(t) \geq U_{\mathrm{S}}(t_i)$—whereas $\widetilde{\mathcal{J}}_i(k)$ contains the remaining locations within $\left\{t_{\pi^i(1)}, \dots, t_{\pi^i(k)}\right\}$. Finally, $\mathcal{K}_i(k)$

represents H's $n-i-k$ least valued locations of $\mathcal{T}\setminus\mathcal{I}_i$, i.e., $\{t_{\pi^i(k+1)},\ldots,t_{\pi^i(n-i)}\}$.

From the above discussion, S's incentive is then to inspect the locations of a set $\mathcal{J}_i(k)$, for some $k\in\{1,\ldots,n-i\}$, while avoiding $\widetilde{\mathcal{J}}_i(k)$ due to their low values of $U_{\mathrm{S}}(t)$, and $\mathcal{K}_i(k)$ due to their low values of $U_{\mathrm{H}}(t)$. To make these inspected locations appealing to H—thus ensuring S derives utility from their inspection—S may equalize H's expected utilities in each location of $\mathcal{J}_i(k)$ to $U_{\mathrm{H}}(t_{\pi^i(k)})$, rendering them more attractive for H than any location in $\mathcal{K}_i(k)$. This can be accomplished by setting $\rho_t^{\mathrm{S}}=1-U_{\mathrm{H}}(t_{\pi^i(k)})/U_{\mathrm{H}}(t)$ for every $t\in\mathcal{J}_i(k)$. From S's resource constraint, the maximum number of locations that S is able to inspect under this strategy is determined by

$$k_i\coloneqq\max\left\{k\in[\![1,n-i]\!]:\sum_{t\in\mathcal{J}_i(k)}\left(1-\frac{U_{\mathrm{H}}(t_{\pi^i(k)})}{U_{\mathrm{H}}(t)}\right)<r_{\mathrm{S}}\right\}.\tag{3.2}$$

We recall that locations in $\mathcal{I}_i$ are valuable for H, but hold relatively low value for S. Therefore, H's hiding strategy involves hiding items in all locations within $\mathcal{I}_i$. In the remaining set of locations $\mathcal{T}\setminus\mathcal{I}_i$, H's incentive is to hide items in his most valued locations, that is, $\mathcal{J}_i(\ell)\cup\widetilde{\mathcal{J}}_i(\ell)=\{t_{\pi^i(1)},\ldots,t_{\pi^i(\ell)}\}$, for some $\ell\in\{1,\ldots,n-i\}$, while avoid hiding items in the less valued locations in $\mathcal{K}_i(\ell)$. Similarly to $\mathcal{I}_i$, locations in $\widetilde{\mathcal{J}}_i(\ell)$ are not appealing for S, but are enticing for H. Thus, H will also hide items in all such locations. At locations within $\mathcal{J}_i(\ell)$, H's incentive is to equalize S's expected utilities to $U_{\mathrm{S}}(t_i)$, rendering them at least as attractive for inspections than the uninspected locations $\mathcal{I}_i\cup\widetilde{\mathcal{J}}_i(\ell)$. This can be achieved by setting $\rho_t^{\mathrm{H}}=U_{\mathrm{S}}(t_i)/U_{\mathrm{S}}(t)$ for every $t\in\mathcal{J}_i(\ell)$. From H's resource constraint, the maximum number of such locations is given by

$$\ell_i\coloneqq\max\left\{\ell\in[\![0,n-i]\!]:\ |\mathcal{I}_i|+|\widetilde{\mathcal{J}}_i(\ell)|+\sum_{t\in\mathcal{J}_i(\ell)}\frac{U_{\mathrm{S}}(t_i)}{U_{\mathrm{S}}(t)}<r_{\mathrm{H}}\right\}.$$

Although these inspection and hiding strategies do not necessarily constitute a NE, the parameters k_i and ℓ_i arising from their description will play a crucial role in solving the game $\widetilde{\Gamma}$.

We now describe the iterative process to construct the sets $\mathcal{I}_i$. We start by defining $\mathcal{I}_0 := \emptyset$, $U_\mathrm{S}(t_0) := 0$ and $U_\mathrm{H}(t_0) := +\infty$. We note that under these definitions, the labeling π^0 sorts *all* the locations in $\mathcal{T}$ by nonincreasing H's valuation. Furthermore, since $U_\mathrm{S}(t) > U_\mathrm{S}(t_0)$ for every $t \in \mathcal{T}$, we obtain $\mathcal{J}_0(k) = \left\{t_{\pi^0(1)}, \ldots, t_{\pi^0(k)}\right\}$, $\widetilde{\mathcal{J}}_0(k) = \emptyset$, and $\mathcal{K}_0(k) = \left\{t_{\pi^0(k+1)}, \ldots, t_{\pi^0(n)}\right\}$ for every $k \in \{1, \ldots, n\}$. Next, for $i \geq 0$, we define the next critical location t_{i+1} as

$$t_{i+1} \in \operatorname*{arg\,min}_{t \in \mathcal{J}_i(k_i)} U_\mathrm{S}(t).$$

If there are multiple locations attaining the minimum, we arbitrarily choose one. In other words, t_{i+1} is S's least valued location among those in $\mathcal{J}_i(k_i)$ where S focuses her resources, according to the aforementioned inspection strategy. Next, we define the following index parameter:

$$m^* := \min\left\{m \in [\![1, n - r_\mathrm{S}]\!] : \, |\mathcal{J}_m(k_m)| \leq r_\mathrm{S}\right\}.$$

We stop this iterative process once the critical location t_{m^*} has been constructed.

Lemma 5. *The following statements hold:*

- *The index m^* is well defined.*

- *For every $i \in [\![0, m^* - 1]\!]$, the set $\mathcal{J}_i(k_i)$ is nonempty. Therefore, up to an arbitrary tie-breaking rule, the locations $t_1, \ldots, t_{m^*}$ are well defined.*

- *The locations $t_1, \ldots, t_{m^*}$ satisfy $0 =: U_\mathrm{S}(t_0) < U_\mathrm{S}(t_1) \leq \cdots \leq U_\mathrm{S}(t_{m^*})$.*

Lemma 5 guarantees that the index parameter m^*, as well as the locations $t_1, \ldots, t_{m^*}$ are well defined. Furthermore, from the inequalities $0 =: U_\mathrm{S}(t_0) < U_\mathrm{S}(t_1) \leq \cdots \leq U_\mathrm{S}(t_i)$, we retrieve one of the properties of $\mathcal{I}_i$ that we had anticipated, namely, that t_i is S's most valued location in $\mathcal{I}_i$. Regarding H's valuations, we provide the following bounds.

Lemma 6. *For every $i \in [\![0, m^*]\!]$,*

$$U_{\mathrm{H}}(t) > \frac{|\mathcal{J}_i(k_i)| - r_{\mathrm{S}}}{\sum_{t \in \mathcal{J}_i(k_i)} 1/U_{\mathrm{H}}(t)} \geq U_{\mathrm{H}}(t') \qquad \forall\, t \in \mathcal{I}_i \cup \mathcal{J}_i(k_i) \cup \widetilde{\mathcal{J}}_i(k_i),\ \forall\, t' \in \mathcal{K}_i(k_i).$$

From Lemma 6, we formally retrieve another property of $\mathcal{I}_i$ that we had anticipated. Specifically, H's valuations over the locations in $\mathcal{I}_i$ are relatively high, as reflected by their lower bound $(|\mathcal{J}_i(k_i)| - r_{\mathrm{S}}) \,/\, \sum_{t \in \mathcal{J}_i(k_i)} 1/U_{\mathrm{H}}(t)$. This bound also applies to H's valuations over the locations in $\mathcal{J}_i(k_i) \cup \widetilde{\mathcal{J}}_i(k_i) = \left\{t_{\pi^i(1)}, \ldots, t_{\pi^i(k_i)}\right\}$, and serves as an upper bound over H's valuations over the locations in $\mathcal{K}_i(k_i)$, reflecting that they are less valuable for H than any other location.

3.4.2 Analytical Characterization

From the previous discussion, we observe that the players' decisions in equilibrium depend on their valuations for each location, their numbers of resources, and the parameters k_i and ℓ_i. Next, it is crucial to determine the critical index i that will establish the specific partition of the locations in equilibrium of the game $\widetilde{\Gamma}$. To this end, we first define the following key quantities:

$$\nu_i \coloneqq |\mathcal{I}_i| + |\widetilde{\mathcal{J}}_i(k_i)| + \sum_{t \in \mathcal{J}_i(k_i)} \frac{U_{\mathrm{S}}(t_i)}{U_{\mathrm{S}}(t)}, \qquad \forall\, i \in [\![0, m^*]\!], \tag{3.3}$$

$$\tau_{-1} \coloneqq 0, \qquad \tau_i \coloneqq |\mathcal{I}_i| + |\widetilde{\mathcal{J}}_i(k_i)| + \sum_{t \in \mathcal{J}_i(k_i)} \frac{U_{\mathrm{S}}(t_{i+1})}{U_{\mathrm{S}}(t)}, \quad \forall\, i \in [\![0, m^*-1]\!], \qquad \tau_{m^*} \coloneqq n. \tag{3.4}$$

These quantities are thresholds that partition the continuous interval $[0, n]$ in the following manner.

Lemma 7. *The numbers $\nu_0, \ldots, \nu_{m^*}$ and $\tau_{-1}, \ldots, \tau_{m^*}$ and satisfy $\tau_{-1} = \nu_0 = 0$, $\tau_{m^*} = n$ and $\tau_{i-1} \leq \nu_i \leq \tau_i$ for every $i \in [\![0, m^*]\!]$.*

Therefore, the interval $[0, n]$ where r_{H} lies is subdivided by the thresholds τ_i and ν_i. We let $i^* \in [\![0, m^*]\!]$ be the unique index satisfying $\tau_{i^*-1} < r_{\mathrm{H}} \leq \tau_{i^*}$. In the following theorem, we show that the subinterval in which r_{H} resides corresponds to a precise configuration of the parameters and determines a specific equilibrium regime of the game $\widetilde{\Gamma}$. Additionally, we provide game-theoretic insights for each equilibrium regime and refer all proof details to Appendix B.1.

Theorem 3 (Regime Pattern 1)**.** *If $i^* \in [\![0, m^*-1]\!]$ and $\tau_{i^*-1} < r_{\mathrm{H}} \leq \nu_{i^*}$, then $\ell_{i^*} < k_{i^*}$ and a strategy profile $(\rho^{\mathrm{S}^*}, \rho^{\mathrm{H}^*}) \in \widetilde{\mathcal{A}}_{\mathrm{S}} \times \widetilde{\mathcal{A}}_{\mathrm{H}}$ is a pure NE of the game $\widetilde{\Gamma}$ if it satisfies:*

$$\rho_t^{\mathrm{S}^*} = \begin{cases} \varepsilon_{\mathrm{S}} & \text{if } t = t_{i^*}, \\ 1 - \dfrac{U_{\mathrm{H}}(t_{\pi^{i^*}(\ell_{i^*}+1)})}{U_{\mathrm{H}}(t)} & \text{if } t \in \mathcal{J}_{i^*}(\ell_{i^*}), \\ 0 & \text{if } t \in \mathcal{I}_{i^*} \setminus \{t_{i^*}\} \cup \widetilde{\mathcal{J}}_{i^*}(\ell_{i^*}) \cup \mathcal{K}_{i^*}(\ell_{i^*}), \end{cases} \tag{3.5}$$

$$\rho_t^{\mathrm{H}^*} = \begin{cases} \dfrac{U_{\mathrm{S}}(t_{i^*})}{U_{\mathrm{S}}(t)} & \text{if } t \in \mathcal{J}_{i^*}(\ell_{i^*}), \\ 1 & \text{if } t \in \mathcal{I}_{i^*} \cup \widetilde{\mathcal{J}}_{i^*}(\ell_{i^*}), \\ \varepsilon_{\mathrm{H}} & \text{if } t = t_{\pi^{i^*}(\ell_{i^*}+1)}, \\ 0 & \text{if } t \in \mathcal{K}_{i^*}(\ell_{i^*}) \setminus \left\{ t_{\pi^{i^*}(\ell_i^*+1)} \right\}, \end{cases} \tag{3.6}$$

where ε_{S} and ε_{H} are given by

$$\varepsilon_{\mathrm{S}} := r_{\mathrm{S}} - \sum_{t \in \mathcal{J}_{i^*}(\ell_{i^*})} \left(1 - \frac{U_{\mathrm{H}}(t_{\pi^{i^*}(\ell_{i^*}+1)})}{U_{\mathrm{H}}(t)} \right), \quad \varepsilon_{\mathrm{H}} := r_{\mathrm{H}} - |\mathcal{I}_{i^*}| - |\widetilde{\mathcal{J}}_{i^*}(\ell_{i^*})| - \sum_{t \in \mathcal{J}_{i^*}(\ell_{i^*})} \frac{U_{\mathrm{S}}(t_{i^*})}{U_{\mathrm{S}}(t)}. \tag{3.7}$$

The equilibrium payoffs are given by

$$\begin{aligned} &\mathbf{U}_{\mathrm{S}}(\rho^{\mathrm{S}^*}, \rho^{\mathrm{H}^*}) = U_{\mathrm{S}}(t_{i^*})\, r_{\mathrm{S}}, \\ &\mathbf{U}_{\mathrm{H}}(\rho^{\mathrm{S}^*}, \rho^{\mathrm{H}^*}) = \left(r_{\mathrm{H}} - |\mathcal{I}_{i^*}| - |\widetilde{\mathcal{J}}_{i^*}(\ell_{i^*})| \right) U_{\mathrm{H}}(t_{\pi^{i^*}(\ell_{i^*}+1)}) - \varepsilon_{\mathrm{S}}\, U_{\mathrm{H}}(t_{i^*}) + \sum_{t \in \mathcal{I}_{i^*} \cup \widetilde{\mathcal{J}}_{i^*}(\ell_{i^*})} U_{\mathrm{H}}(t). \end{aligned}$$

In Regime Pattern 1, H's number of resources satisfies $\tau_{i^*-1} < r_{\mathrm{H}} \leq \nu_{i^*} \leq \tau_{i^*}$. In particular, the inequality $r_{\mathrm{H}} \leq \nu_{i^*}$ implies that $\ell_{i^*} + 1 \leq k_{i^*}$. Due to the magnitude of r_{H}, characterized by the inequality $\tau_{i^*-1} < r_{\mathrm{H}}$, S prioritizes her inspections among locations holding relatively large values of $U_{\mathrm{S}}(t)$. Consequently, S initially refrains from inspecting locations in $\mathcal{I}_{i^*}$, due to their low values of $U_{\mathrm{S}}(t)$, as indicated by the inequalities $U_{\mathrm{S}}(t) \leq U_{\mathrm{S}}(t_{i^*})$ for every $t \in \mathcal{I}_{i^*}$ (Lemma 5). Instead, S focuses her inspection efforts on those locations in $\mathcal{T} \setminus \mathcal{I}_{i^*}$ where her valuations are at least $U_{\mathrm{S}}(t_{i^*})$.

According to the intuition developed in Section 3.4.1, the inequality $\ell_{i^*} + 1 \leq k_{i^*}$ implies that S has sufficient inspection resources so as to equalize H's utilities in his $\ell_{i^*} + 1$ most valued locations within such set—i.e., the locations in $\mathcal{J}_{i^*}(\ell_{i^*} + 1)$—to the value of $U_{\mathrm{H}}(t_{\pi^{i^*}(\ell_{i^*}+1)})$. However, setting H's utility to $U_{\mathrm{H}}(t_{\pi^{i^*}(\ell_{i^*}+1)})$ in location $t_{\pi^{i^*}(\ell_{i^*}+1)}$ requires S to provide no inspection at all to that location. Hence, we equivalently use the sets $\mathcal{J}_{i^*}(\ell_{i^*})$, as well as $\widetilde{\mathcal{J}}_{i^*}(\ell_{i^*})$ and $\mathcal{K}_{i^*}(\ell_{i^*})$—instead of $\mathcal{J}_{i^*}(\ell_{i^*}+1)$, $\widetilde{\mathcal{J}}_{i^*}(\ell_{i^*}+1)$ and $\mathcal{K}_{i^*}(\ell_{i^*}+1)$—to describe S's equilibrium inspection strategy in (3.5).

After equalizing H's utilities in $\mathcal{J}_{i^*}(\ell_{i^*} + 1)$, S has ε_{S} additional resources left, which she allocates to an initially uninspected location that provides her the highest utility, that is, location t_{i^*}. Interestingly, ε_{S} satisfies the following bounds:

$$0 < \varepsilon_{\mathrm{S}} \leq 1 - \frac{U_{\mathrm{H}}(t_{\pi^{i^*}(\ell_{i^*}+1)})}{U_{\mathrm{H}}(t_{i^*})}. \tag{3.8}$$

In particular, unless the upper bound on ε_{S} is tight, (3.8) shows that S is unable to equalize H's utility from t_{i^*} to the value of $U_{\mathrm{H}}(t_{\pi^{i^*}(\ell_{i^*}+1)})$, as she does for the locations in $\mathcal{J}_{i^*}(\ell_{i^*} + 1)$. Instead, H's utility from t_{i^*}, given by $(1 - \varepsilon_{\mathrm{S}})\, U_{\mathrm{H}}(t_{i^*})$, is greater than $U_{\mathrm{H}}(t_{\pi^{i^*}(\ell_{i^*}+1)})$, rendering it more enticing than any location within $\mathcal{J}_{i^*}(\ell_{i^*}+1)$. Furthermore, $U_{\mathrm{H}}(t_{\pi^{i^*}(\ell_{i^*}+1)})$ satisfies the following inequalities:

$$U_{\mathrm{H}}(t) \geq U_{\mathrm{H}}(t_{\pi^{i^*}(\ell_{i^*}+1)}) \geq U_{\mathrm{H}}(t'), \qquad \forall\, t \in \mathcal{I}_{i^*} \setminus \{t_{i^*}\} \cup \widetilde{\mathcal{J}}_{i^*}(\ell_{i^*})\ \forall\, t' \in \mathcal{K}_{i^*}(\ell_{i^*}).$$

These inequalities, together with (3.8), ensure the feasibility of S's strategy and render $\mathcal{I}_{i^*} \cup \widetilde{\mathcal{J}}_{i^*}(\ell_{i^*})$ as the most enticing set of locations for H, followed by those in $\mathcal{J}_{i^*}(\ell_{i^*})$, and finally by those in $\mathcal{K}_{i^*}(\ell_{i^*})$.

On the other hand, H's hiding strategy (3.6) involves saturating the locations within $\mathcal{I}_{i^*} \setminus \{t_{i^*}\} \cup \widetilde{\mathcal{J}}_{i^*}(k_{i^*})$, which are left uninspected by S due to their relatively low values of $U_{\mathrm{S}}(t)$, as well as location t_{i^*}, which as we mentioned, is among the most enticing ones for H. Within the locations in $\mathcal{J}_{i^*}(\ell_{i^*})$, H's equalizes S's utilities to $U_{\mathrm{S}}(t_{i^*})$. By doing so, H desincentivizes S to relocate her inspection resources towards the uninspected locations in $\mathcal{I}_{i^*} \setminus \{t_{i^*}\} \cup \widetilde{\mathcal{J}}_{i^*}(k_{i^*})$. After allocating his resources in this manner, H still has ε_{H} resources left, that he allocates in the initially empty location that provides him the highest utility, that is, $t_{\pi^{i^*}(\ell_{i^*}+1)}$. The quantity ε_{H} satisfies the following bounds:

$$0 < \varepsilon_{\mathrm{H}} \leq \min\left\{\frac{U_{\mathrm{S}}(t_{i^*})}{U_{\mathrm{S}}(t_{\pi^{i^*}(\ell_{i^*}+1)})},\, 1\right\}. \tag{3.9}$$

Thus, unless the upper bound on ε_{H} is tight, (3.9) shows that H can only set S's utility in location $t_{\pi^{i^*}(\ell_{i^*}+1)}$, given by $U_{\mathrm{H}}(t_{\pi^{i^*}(\ell_{i^*}+1)})\,\varepsilon_{\mathrm{H}}$, to a value that is smaller than $U_{\mathrm{S}}(t_{i^*})$, rendering such location less attractive for S to inspect than any location in $\mathcal{J}_{i^*}(\ell_{i^*})$. Moreover, $U_{\mathrm{S}}(t_{i^*})$ satisfies the following inequalities:

$$U_{\mathrm{S}}(t_{i^*}) \geq U_{\mathrm{S}}(t), \qquad \forall\, t \in \mathcal{I}_{i^*} \cup \widetilde{\mathcal{J}}_{i^*}(\ell_{i^*}).$$

These inequalities, together with (3.9), guarantee that H's strategy is feasible and that S's most attractive locations for inspections are those in $\mathcal{J}_{i^*}(\ell_{i^*}) \cup \{t_{i^*}\}$.

Theorem 3 (Regime Pattern 2)**.** *If $i^* \in [\![0, m^*-1]\!]$ and $\nu_{i^*} < r_{\mathrm{H}} \leq \tau_{i^*}$, then $k_{i^*} \leq \ell_{i^*}$ and* *a strategy profile $(\rho^{\mathrm{S}^*}, \rho^{\mathrm{H}^*}) \in \widetilde{\mathcal{A}}_{\mathrm{S}} \times \widetilde{\mathcal{A}}_{\mathrm{H}}$ is a pure NE of the game Γ if it satisfies:*

$$\rho_t^{\mathrm{S}^*} = \begin{cases} 1 - \dfrac{\lambda_{\mathrm{S}}}{U_{\mathrm{H}}(t)} & \text{if } t \in \mathcal{J}_{i^*}(k_{i^*}), \\ 0 & \text{if } t \in \mathcal{I}_{i^*} \cup \widetilde{\mathcal{J}}_{i^*}(k_{i^*}) \cup \mathcal{K}_{i^*}(k_{i^*}), \end{cases} \tag{3.10}$$

$$\rho_t^{\mathrm{H}^*} = \begin{cases} \dfrac{\lambda_{\mathrm{H}}}{U_{\mathrm{S}}(t)} & \text{if } t \in \mathcal{J}_{i^*}(k_{i^*}), \\ 1 & \text{if } t \in \mathcal{I}_{i^*} \cup \widetilde{\mathcal{J}}_{i^*}(k_{i^*}), \\ 0 & \text{if } t \in \mathcal{K}_{i^*}(k_{i^*}), \end{cases} \tag{3.11}$$

where λ_{S} and λ_{H} are given by

$$\lambda_{\mathrm{S}} \coloneqq \frac{|\mathcal{J}_{i^*}(k_{i^*})| - r_{\mathrm{S}}}{\sum_{t \in \mathcal{J}_{i^*}(k_{i^*})} \frac{1}{U_{\mathrm{H}}(t)}}, \qquad \lambda_{\mathrm{H}} \coloneqq \frac{r_{\mathrm{H}} - |\mathcal{I}_{i^*}| - |\widetilde{\mathcal{J}}_{i^*}(k_{i^*})|}{\sum_{t \in \mathcal{J}_{i^*}(k_{i^*})} \frac{1}{U_{\mathrm{S}}(t)}}. \tag{3.12}$$

The equilibrium's expected utilities are given by

$$\begin{aligned} \mathbf{U}_{\mathrm{S}}(\rho^{\mathrm{S}^*}, \rho^{\mathrm{H}^*}) &= \lambda_{\mathrm{H}}\, r_{\mathrm{S}}, \\ \mathbf{U}_{\mathrm{H}}(\rho^{\mathrm{S}^*}, \rho^{\mathrm{H}^*}) &= \lambda_{\mathrm{S}} \left(r_{\mathrm{H}} - |\mathcal{I}_{i^*}| - |\widetilde{\mathcal{J}}_{i^*}(k_{i^*})| \right) + \sum_{t \in \mathcal{I}_{i^*} \cup \widetilde{\mathcal{J}}(k_{i^*})} U_{\mathrm{H}}(t). \end{aligned}$$

In Regime Pattern 2, H's number of resources is greater than that of Regime Pattern 1, as reflected by the inequalities $\tau_{i^*-1} \leq \nu_{i^*} < r_{\mathrm{H}} \leq \tau_{i^*}$. In particular, the condition $\nu_{i^*} < r_{\mathrm{H}}$ implies that $k_{i^*} \leq \ell_{i^*}$. Following the intuition derived from our preliminary analysis in Section 3.4.1, the inequality $k_{i^*} \leq \ell_{i^*}$ indicates that S can only equalize the utilities of k_{i^*} out of H's ℓ_{i^*} most valued locations in $\mathcal{T} \setminus \mathcal{I}_{i^*}$ whose valuations for S are at least $U_{\mathrm{S}}(t_{i^*})$. Consequently, S's equilibrium inspection strategy focuses on the locations within $\mathcal{J}_{i^*}(k_{i^*})$, while refraining from inspecting $\mathcal{I}_{i^*} \cup \widetilde{\mathcal{J}}_{i^*}(k_{i^*})$ due to their relatively low values of $U_{\mathrm{S}}(t)$—which do not exceed $U_{\mathrm{S}}(t_{i^*})$—and $\mathcal{K}_{i^*}(k_{i^*})$ due to the heterogeneity of the values of $U_{\mathrm{H}}(t)$. Within $\mathcal{J}_{i^*}(k_{i^*})$, S utilizes all her inspection resources to equalize H's utility from each location to λ_{S}. From Lemma 6, λ_{S} satisfies the following bounds:

$$U_{\text{H}}(t) > \lambda_{\text{S}} \geq U_{\text{H}}(t') \qquad \forall\, t \in \mathcal{I}_{i^*} \cup \mathcal{J}_{i^*}(k_{i^*}) \cup \widetilde{\mathcal{J}}_{i^*}(k_{i^*}),\ \forall\, t' \in \mathcal{K}_{i^*}(k_{i^*}).$$

These inequalities guarantee that S's inspection strategy is feasible and that H's utilities are such that locations in $\mathcal{I}_{i^*} \cup \widetilde{\mathcal{J}}_{i^*}(k_{i^*})$ become the most enticing for him, followed by the locations in $\mathcal{J}_{i^*}(k_{i^*})$, and finally by those in $\mathcal{K}_{i^*}(k_{i^*})$.

Conversely, H's hiding strategy involves saturating all his most enticing locations, that is, those in $\mathcal{I}_{i^*} \cup \widetilde{\mathcal{J}}_{i^*}(k_{i^*})$, which are left uninspected by S. Then, H allocates his remaining resources to his next most enticing set of locations, $\mathcal{J}_{i^*}(k_{i^*})$, equalizing S's utilities across such locations to λ_{H}. The inequalities $\nu_{i^*} < r_{\text{H}} \leq \tau_{i^*}$ guarantee that H has sufficient resources to implement this strategy, leaving no resources for hiding in the less enticing set of locations $\mathcal{K}_{i^*}(k_{i^*})$. Furthermore, λ_{H} satisfies the following bounds:

$$U_{\text{S}}(t) \geq \lambda_{\text{H}} > U_{\text{S}}(t') \qquad \forall\, t \in \mathcal{J}_{i^*}(k_{i^*}),\ \forall\, t' \in \mathcal{I}_{i^*} \cup \widetilde{\mathcal{J}}_{i^*}(k_{i^*}).$$

These inequalities ensure the feasibility of H's strategy and render S's utilities from locations in $\mathcal{J}_{i^*}(k_{i^*})$ as the most attractive for S to inspect.

Theorem 3 (Regime Pattern 3)**.** *If $i^* = m^*$ and $\nu_{m^*} < r_{\text{H}} \leq \tau_{m^*}$, then $k_{m^*} = n - m^*$ and a strategy profile $(\rho^{\text{S}^*}, \rho^{\text{H}^*}) \in \widetilde{\mathcal{A}}_{\text{S}} \times \widetilde{\mathcal{A}}_{\text{H}}$ is a pure NE of the game $\widetilde{\Gamma}$ if it satisfies:*

$$\rho_t^{\text{S}^*} = \begin{cases} 1 & \text{if } t \in \mathcal{J}_{m^*}(n-m^*), \\ 0 & \text{if } t \in \mathcal{I}_{m^*} \cup \widetilde{\mathcal{J}}_{m^*}(n-m^*), \end{cases} \tag{3.13}$$

$$\begin{cases} \dfrac{U_{\text{S}}(t_{m^*})}{U_{\text{S}}(t)} \leq \rho_t^{\text{H}^*} \leq 1 & \text{if } t \in \mathcal{J}_{m^*}(n-m^*), \\ \rho_t^{\text{H}^*} = 1 & \text{if } t \in \mathcal{I}_{m^*} \cup \widetilde{\mathcal{J}}_{m^*}(n-m^*), \\ \displaystyle\sum_{t \in \mathcal{T}} \rho_t^{\text{H}^*} \leq r_{\text{H}}. & \end{cases} \tag{3.14}$$

The equilibrium's expected utilities satisfy

$$U_{\mathrm{S}}(t_{m^*})\, r_{\mathrm{S}} \leq \mathbf{U}_{\mathrm{S}}(\rho^{\mathrm{S}^*}, \rho^{\mathrm{H}^*}) \leq \sum_{t \in \mathcal{J}_{m^*}(n-m^*)} U_{\mathrm{S}}(t),$$

$$\mathbf{U}_{\mathrm{H}}(\rho^{\mathrm{S}^*}, \rho^{\mathrm{H}^*}) = \sum_{t \in \mathcal{I}_{m^*} \cup \widetilde{\mathcal{J}}_{n-m^*}(n-m^*)} U_{\mathrm{H}}(t).$$

In Regime Pattern 3, H's number of resources satisfies $\nu_{m^*} < r_{\mathrm{H}} \leq \tau_{m^*} = n$, representing the interval where r_{H} takes the largest values in our classification of equilibrium regime patterns. From the definition of m^*, we have $|\mathcal{J}_{m^*}(k_{m^*})| \leq r_{\mathrm{S}}$. On the other hand, the definition of k_{m^*} implies $|\mathcal{J}_{m^*}(k_{m^*})| \geq r_{\mathrm{S}}$. Thus, $|\mathcal{J}_{m^*}(k_{m^*})| = r_{\mathrm{S}}$. Furthermore, this equality also implies $k_{m^*} = n - m^*$. As a result, the set of locations $\mathcal{T}$ is partitioned into the three sets $\mathcal{I}_{m^*}$, $\mathcal{J}_{m^*}(n-m^*)$ and $\widetilde{\mathcal{J}}_{m^*}(n-m^*)$.

On the other hand, the definition of k_{m^*} implies $|\mathcal{J}_{m^*}(k_{m^*})| \geq r_{\mathrm{S}}$ and $k_{m^*} = n - m^*$. Thus, $|\mathcal{J}_{m^*}(k_{m^*})| = |\mathcal{J}_{m^*}(n-m^*)| = r_{\mathrm{S}}$. As a result, the set of locations $\mathcal{T}$ is partitioned into the three sets $\mathcal{I}_{m^*}$, $\mathcal{J}_{m^*}(n-m^*)$ and $\widetilde{\mathcal{J}}_{m^*}(n-m^*)$.

We recall that S's valuations for locations in $\mathcal{J}_{m^*}(n-m^*)$ are no less than $U_{\mathrm{S}}(t_{m^*})$. Moreover, from Lemma 5 and the definition of $\widetilde{\mathcal{J}}_{m^*}(n-m^*)$, S's valuations for the locations in $\mathcal{I}_{m^*} \cup \widetilde{\mathcal{J}}_{m^*}(n-m^*)$ are at most $U_{\mathrm{S}}(t_{m^*})$. Therefore, the set $\mathcal{J}_{m^*}(n-m^*)$ contains S's r_{S} most valued locations. Consequently, S provides full inspection to these locations, as reflected by (3.13).

Conversely, H's hiding strategy (3.14) does not need to fully utilize his resources. In fact, after saturating all the uninspected locations in $\mathcal{I}_{m^*} \cup \widetilde{\mathcal{J}}_{m^*}(n-m^*)$, H only needs to provide a level of hiding to the remaining locations in $\mathcal{J}_{m^*}(n-m^*)$ so as to ensure that S's utility from each such location is at least $U_{\mathrm{S}}(t_{m^*})$. By doing so, the locations in $\mathcal{J}_{m^*}(n-m^*)$ remain the most attractive for S to inspect. The inequality $\nu_{m^*} < r_{\mathrm{H}}$ guarantees the feasibility of H's hiding strategy.

We note that since H's hiding strategy is not unique, S's expected utility depends on the total number of resources H allocates. At the minimum, H equalizes S's utilities to the exact value of $U_{\mathrm{S}}(t_{m^*})$, rendering S a utility of $r_{\mathrm{S}}\, U_{\mathrm{S}}(t_{m^*})$. However, as H's hides more resources within the locations $\mathcal{J}_{m^*}(n-m^*)$, S's overall utility increases, up to the value of $\sum_{t\in\mathcal{J}_{m^*}(n-m^*)} U_{\mathrm{S}}(t)$ when H utilizes $\tau_{m^*} = n$ resources.

From Theorem 3, and due to a careful partition of the set of locations induced by a set of key thresholds and parameters, we can derive closed-form expressions for the NE of the continuous game $\widetilde{\Gamma}$. In general, given $i^* \in [\![0, m^*-1]\!]$ satisfying $\tau_{i^*-1} < r_{\mathrm{H}} \leq \tau_{i^*}$, we generally find that $\mathcal{I}_{i^*} = \{t_1, \ldots, t_{i^*}\}$ represents the set of locations with valuations too small for S to inspect—with the possible exception of t_{i^*}—yet sufficiently high to be enticing for H, who hides items in all of them. This set is constructed iteratively and its number of locations i^* results from the relative difference between the players' number of resources, encoded through the inequality $\tau_{i^*-1} < r_{\mathrm{H}}$. Next, due to constraints arising from the heterogeneity of the players' valuations and their limited resources, both players allocate their remaining resources to H's $\min\{k_{i^*}, \ell_{i^*}\}$ most valued locations within $\mathcal{T}\setminus\mathcal{I}_{i^*}$. Specifically, within this set, the locations in $\widetilde{\mathcal{J}}_{i^*}(\min\{k_{i^*}, \ell_{i^*}\})$—i.e., those for which S's valuations are below the threshold $U_{\mathrm{S}}(t_{i^*})$—are not appealing to S and are therefore left uninspected but saturated by H. The remaining locations in this set, comprising $\mathcal{J}_{i^*}(\min\{k_{i^*}, \ell_{i^*}\})$, are appealing to both players. Hence, they randomize their remaining resources—in S's case, all her resources—over these locations, equalizing their opponent's utilities at each location. Finally, H's $n - i - \min\{k_{i^*}, \ell_{i^*}\}$ less valued locations within $\mathcal{T} \setminus \mathcal{I}_{i^*}$, comprising the set $\mathcal{K}_{i^*}(\min\{k_{i^*}, \ell_{i^*}\})$, are not appealing to either player—with the possible exception of $t_{\pi^{i^*}(\min\{k_{i^*}, \ell_{i^*}\}+1)}$—so they do not allocate resources there.

Nonetheless, the relation between k_{i^*} and ℓ_{i^*}, encoded through the inequalities $r_{\mathrm{H}} \leq \nu_{i^*}$ if $k_{i^*} \leq \ell_{i^*}$ and $\nu_{i^*} < r_{\mathrm{H}}$ if $\ell_{i^*} < k_{i^*}$, dictates additional considerations in determining the NE of the game $\widetilde{\Gamma}$. Specifically, if $k_{i^*} \leq \ell_{i^*}$, both players utilize all their remaining resources for randomization over $\mathcal{J}_{i^*}(\min\{k_{i^*}, \ell_{i^*}\})$. Conversely, if $\ell_{i^*} < k_{i^*}$, both players have a small number of remaining resources after such randomization, which they allocate to their most valued previously unallocated location, t_{i^*} for S and $t_{\pi^{i^*}(\ell_{i^*}+1)}$ for H.

We can now outline the overall approach for solving the game Γ. First, we compute the marginal inspection and hiding probabilities in equilibrium of the game $\widetilde{\Gamma}$, according to Theorem 3. Sorting H's valuations in nonincreasing order requires $O(n \log n)$ steps. For each $i \in [\![0, m^* - 1]\!]$, the mapping π^i, parameters k_i and ℓ_i, and thresholds τ_i and ν_i can be computed in $O(n)$ time. Identifying the subinterval $[\tau_{i^*-1}, \nu_{i^*}]$ or $[\nu_{i^*}, \tau_{i^*}]$ that contains r_{H} can be done in $O(\log n)$ steps. Evaluating the expressions from Theorem 3 requires an additional $O(n)$ steps. Therefore, computing equilibrium marginal inspection and hiding probabilities can be implemented in $O(n \log n)$ time. Finally, from Lemma 4, we can compute mixed strategies compatible with these marginal probabilities and whose support size is at most $n + 1$ in $O(n^2)$ time. Hence, we deduce the following theorem.

Theorem 4. *Therefore, the game Γ can be solved in time $O(n^2)$ with equilibrium strategies of support size at most $n + 1$ each.*

Thus, we have developed an efficient solution method for the large-scale nonzero-sum hide-and-seek game Γ, accounting for multiple players' resources and heterogeneous and asymmetric valuations over the locations. Additionally, our approach yields solutions with small supports, facilitating straightforward implementation by strategic decision makers.

3.5 Concluding Remarks

In this chapter, we investigated a hide-and-seek game where a seeker selects a randomized allocation of inspection resources across a set of locations to find multiple items hidden by a hider. We extended previous models by considering players with multiple resources who hold asymmetric and heterogeneous valuations over the hiding locations. The seeker's (resp. hider's) objective is to maximize the additive value obtained from the inspected (resp. uninspected) locations containing hidden items. The asymmetry in the players' valuations broadens the scope of strategic mismatch models in the literature, extending applications in security domains, animal behavior, political campaigns, etc., into the nonzero-sum setting.

To compute mixed-strategy NE of this large-scale nonzero-sum game, we analyzed the pure NE of an equivalent lower-dimensional continuous game. In this continuous game, players' strategies represent the marginal probabilities of inspection and hiding at each location in the original game. We solved the continuous game analytically, providing closed-form expressions for the players' equilibrium strategies, which depend on key thresholds and parameters derived from the players' valuations, incentives, and resource constraints. Additionally, we offered game-theoretic insights to better understand the players' behavior in equilibrium.

Our equilibrium analysis reveals that, generally, there exists a subset of locations that hold relatively high value for the seeker but relatively low value for the hider. The seeker discards these locations from inspections and instead randomizes the allocation of her inspection resources within the remaining locations. Specifically, she focuses on the hider's most valued locations within this set, provided the seeker's valuation of these locations is at least as high as any location in the initially discarded subset. Conversely, the hider fully targets the seeker's initially discarded locations, as well as his most valued locations within the remaining set, as long as they are not inspected by the seeker. The hider then randomizes his remaining resources within the same set of locations where the seeker randomizes her inspection resources. A final subset of locations, containing the hider's least valued locations, is neither inspected by the seeker nor targeted by the hider. In some equilibrium regimes, the players are left with a small number of remaining resources, which they allocate greedily to the unallocated location providing them the highest utility. Our equilibrium insights and solution approach can be leveraged by security agencies interested in designing inspection strategies to detect and deter the activities of criminal organizations, such as a gang of pickpockets operating in subway stations of a city's public transit system.

This work can be extended in several directions. One such extension is to consider the more general setting found in security games with unit-sized schedules, where each location is associated with two utility values for each player. This can be further extended to include heterogeneously-sized schedules, where closed-form solutions may be possible under the assumption of mutually disjoint inspection schedules. Another interesting extension is to examine the Stackelberg version of the game, where the existing algorithms fall short of providing a comprehensive understanding of the players' incentives driving the equilibrium. Lastly, an intriguing research direction is to explore a version of the hide-and-seek game in which the valuations of one or both players over the locations are private and must be learned through repeated interactions.

CHAPTER 4
STRATEGIC NETWORK INSPECTION WITH LOCATION-SPECIFIC DETECTION CAPABILITIES

4.1 Introduction

Modern societies' welfare significantly relies on critical infrastructure networks, including power grids, transportation and telecommunication systems, and water, oil or gas distribution pipelines. However, these systems have increasingly become targets for malicious cyber-physical attacks, as evidenced by past incidents that have significantly disrupted their operations. In 2015, Ukraine experienced several foreign cyberattacks that compromised the supervisory control and data acquisition (SCADA) systems managing its power grid. These attacks resulted in widespread power outages, affecting 225,000 customers [26]. In 2021, the Colonial Pipeline ransomware attack led to substantial disruptions in fuel supply along the East Coast of the United States, causing losses of $4.4 million USD [27]. In the same year, a cyberattacker gained remote access to a local computer system of the water treatment plant of Oldsmar, Florida, and increased the levels of sodium hydroxide (caustic soda) from 100 ppm to 11,100 ppm, a concentration that can severely harm human tissue. Fortunately, an on-site operator promptly identified the unauthorized change and restored chemical levels, preventing any significant impact to the water quality consumed by the 15,000 residents of the city [28].

Such targeted attacks pose a severe threat to the reliability and functionality of critical infrastructures, highlighting the urgent need for robust inspection systems capable of detecting both random anomalies and adversarial attacks. Developing effective inspection mechanisms has become a paramount requirement to ensure the continuous operation, safety, and resilience of these networks [115, 116, 117].

Numerous efforts have been deployed in the development of inspection technologies that monitor critical parameters such as temperature, pressure, flow rates, and voltage levels to assess the operational status and performance of various networks. These technological advances, often relying on online data acquisition and continual monitoring through the flexible allocation of sensors within the network, enable operators to promptly identify anomalies, faults, or potential threats and perform proactive maintenance tasks and timely mitigation responses [118].

In the context of large-scale networks and costly detection technology, the limited availability of sensing resources can pose significant limitations for the continuous monitoring of all the network components. Moreover, sensing reliability can be locally undermined by several factors, such as network topology, obstacles, and harsh environmental conditions. For instance, detection systems based on infrared thermography (IRT) can identify irregular temperature patterns around gas pipelines, signaling potential leaks. Yet, high-resolution infrared cameras are expensive, and IRT's efficiency might significantly decrease for buried pipelines or within the presence of covering materials such as concrete [29]. Unmanned aerial systems (UAS) are increasingly used in the oil and gas industry for pipeline surveillance, equipped with optical cameras or IRT sensors to capture high-resolution images and videos. Nevertheless, adverse weather conditions such as intense sunlight, rain, and strong winds can impair visibility and inspection efficacy [4]. Similarly, in-pipe inspection robots are devices specially crafted for navigating the interior of water, gas, and oil pipelines. They rely on wheels, tracks, or other mechanisms to generate traction to traverse the pipes. However, their effectiveness can be hindered within complex pipeline geometries, where valves, T-sections, elbows, and varying diameters present obstacles to thorough inspection [30].

Hence, in this chapter we aim to investigate the following research question: *How to effectively coordinate limited inspection resources within a networked system with location-specific detection capabilities, in order to optimize the detection of multiple adversarial* *attacks on its components?*

To address this question, we adopt a game-theoretic approach to the network inspection problem, where two players, a defender and an attacker, strategically coordinate multiple resources within a network. Specifically, the defender coordinates the allocation of inspection resources, which we refer to as *detectors*, while the attacker determines the distribution of attack resources. The network where the players interact comprises *nodes* and *components*. The nodes serve as prospective locations for the placement of detectors by the defender, and the components represent the attacker's potential targets. The allocation of a detector within a node enables the defender to monitor a specific subset of components—which we refer to as *monitoring set*—and potentially detect attacks targeting them. To reflect the intricacies of networked systems arising in critical infrastructure settings, we allow these monitoring sets to overlap, enabling the defender to simultaneously monitor components from multiple locations. Additionally, our network model accounts for imperfect detection, which may result in the defender overlooking targeted components during inspection. Namely, we assume that upon inspection, the defender identifies an attack on a component with a location-specific probability that captures the local effects undermining the detection capabilities. The defender (resp. attacker) aims to minimize (resp. maximize) the expected number of undetected attacks.

The objective of this chapter is to efficiently compute exact and approximate Nash Equilibria (NE) for this large-scale zero-sum game. Although equilibrium strategies can be obtained by solving a pair of dual linear programs (LPs), this approach is impractical for realistic instances of the game, as the number of players' strategies grows combinatorially with the size of the network and the number of available resources.

4.1.1 Contributions

The contributions of this chapter can be summarized as follows:

– We investigate a two-person zero-sum network inspection game involving multiple homogeneous resources available to both players, imperfect detection in overlapping monitoring sets characterized by location-specific detection probabilities, and players' payoff given by the expected number of *undetected* attacks. Our model extends previous works in the literature, which typically involve a single unit of resource for one or both players [36, 38, 54], assume perfect detection capabilities [16], or restrict detection to nonoverlapping monitoring sets [93]. More generally, our contribution expands the literature on combinatorial games and randomized robust optimization with *supermodular* structure, which has been less explored compared to similar settings with submodular structure, enhancing our understanding of strategic interactions in complex networked systems.

– We leverage the structure of our game to compactly represent the attacker's strategies in terms of marginal probabilities of targeting individual network components. Using this representation, we formulate an LP with a small number of constraints that can be solved via Column Generation (CG) to obtain exact NE (Proposition 3). The pricing problem, corresponding to the defender's pure best response problem, consists in finding a detector positioning that minimizes the supermodular expected number of undetected attacks, subject to a cardinality constraint. We show that this problem is NP-hard (Proposition 4) and provide a compact mixed-integer programming (MIP) formulation (Proposition 5). Next, leveraging the supermodular structure of the defender's best response problem, we rely on both the forward and reverse greedy algorithms to compute approximate best responses, whose guarantees depend on the curvature of the defender's payoff function given the attacker's strategy [119, 120]. Nonetheless, we show that the defender's best response problem becomes inapproximable once such curvature parameter reaches its maximum value (Proposition 6).

– We also propose two efficient solution approaches to compute approximate equilibrium strategies for our game. First, we devise a CG algorithm with approximate best responses for the pricing problem, computed by the aforementioned greedy algorithms. This method generates approximate equilibria by inheriting the approximation guarantees for the pricing problem. We extend a result by [121] by allowing early termination of CG at the cost of a small additive error, and by simplifying the convergence criterion (Theorem 6). Second, following the approach of Krause *et al.* [122], we integrate the Multiplicative Weights Update (MWU) algorithm with an approximation algorithm for the defender's best response problem. This integration yields similar approximation guarantees to those of CG, but in polynomial time (Theorem 7). However, the combinatorial number of attacker's strategies in our game renders the update step in Krause *et al.*'s implementation of the MWU algorithm computationally prohibitive. We address this issue by implementing the MWU algorithm on the attacker's marginal probabilities, instead of the original mixed strategies used in the standard implementation. This adaptation requires solving an additional projection problem at each iteration to recover feasibility of the attacker's marginal attack probabilities after the update step.

– We analytically solve the projection problem arising in our implementation of the MWU algorithm by characterizing the projection under the unnormalized relative entropy onto the *full-dimensional* capped simplex polytope (Theorem 8). To achieve this result, we identify a precise condition that determines whether the resource constraint is either tight or not at optimality. To handle the former case, we employ a Lagrangian formulation of the projection problem and circumvent the nonsmoothness of the unnormalized relative entropy by dualizing the equality constraint and then solving independent single-variable optimization problems for each entry of the projection. The resulting closed-form solution involves a key index parameter, which we compute using a linear-time algorithm (Theorem 9). Our algorithm's running time matches that of previous optimal algorithms designed for variants of the *lower-dimensional* capped simplex and related polytopes [123, 124, 125, 126], and is faster than other more general-purpose algorithms intended for computing Bregman projections [127, 128].

- We conduct a computational study to evaluate the performance of our solution approaches in real-world instances from gas distribution networks across the US and Europe. By considering various network configurations, we provide insights into the performance of our exact and approximate solution methods across different problem instances. Furthermore, we analyze the quality of approximations and scalability of our approximate solution methods for a fixed instance with different detector allocation scenarios, providing insights into the practical applicability of each method. Our computational results validate the theoretical underpinnings of our proposed algorithms and offer practical guidelines for efficiently addressing real-world network security challenges.

4.1.2 Literature Review

Game theory has emerged as a powerful tool applied to security related problems [see, for example, 5, 6, 9, 10, 11, 13, 14]. One of such practical applications arises in problems of strategic sensor placement for network inspection, in which a defender positions sensors in a subset of prospective locations of a network to detect attacks caused by a strategic attacker, who can target one or multiple network components. Such models typically account for the detection range of the sensors: Positioning a sensor allows the defender to monitor a subset of components—often referred to as a monitoring set—and potentially detect attacks occurring within it. As a result, attacks may be detected from multiple locations. This overlapping feature renders such games challenging to solve.

Dahan *et al.* [16] investigated a restricted version of our model under the assumption of perfect detection—that is, with sensors that operate without failures within their detection range—and obtained approximate equilibria by means of minimum set covers and maximum set packings. [17] extended the model of Dahan *et al.* by considering network components with heterogeneous criticality values. Conversely, [83, 93] introduced imperfect location-specific detection capabilities to the network inspection model and solved the game restricted to nonoverlapping (i.e., mutually disjoint) monitoring sets. This assumption renders the model akin to a hide-and-seek game with imperfect detection and heterogeneous hiding capacities, allowing the authors to derive closed-form equilibria in terms of marginal inspection probabilities and expected number of items hidden within each hiding location. The more general setting with overlapping monitoring sets and imperfect detection, considered in this chapter, first appeared in [83], who only proposed heuristic solutions based on minimum weighted set covers. In contrast, we derive both exact and approximate solution approaches with proven approximation guarantees.

Our model extends the classic hide-and-seek game of von Neumann [36] (see [54] for a version with imperfect detection), where a hider selects an entry a_{ij} of a square matrix of nonnegative entries, and a seeker selects either a row or a column, aiming to find the hider. If the seeker's choice contains the hider's selection, the hider pays the seeker a_{ij}; otherwise, the seeker pays the hider a_{ij}. Drawing an analogy with our game, nodes represent rows and columns, network components represent matrix entries, and each node's monitoring set contains the entries within its associated row or column. Similarly, our model also extends the seminal zero-sum network interdiction game of Washburn and Wood [38], where an interdictor places an inspection checkpoint, with location-specific detection capabilities, along an arc in a directed graph to interdict an evader traversing a path. The interdictor (resp. evader) aims to maximize (resp. minimize) the interdiction probability. In our model, nodes represent arcs, components represent paths, and each node's monitoring set contains the paths traversing its corresponding arc. Washburn and Wood focused on a single inspector and evader, showing polynomial-time equilibrium computation using network flow techniques. An extension of this model to multiple checkpoints and evader's strategies in the form of network flows was examined by [121]. Our extension handles scenarios with multiple checkpoints and coordinated evaders, making it suitable for security applications involving the interdiction of criminal organizations transporting illegal commodities across networks with a reasonable number of paths, which is typically the case in illicit supply chains.

Another relevant work is the game of [129], an extension of the influential security game of [6], involving multiple attacker and defender resources and nonadditive utilities. Unlike our model, they assume an unconstrained number of defender resources and a more general nonzero-sum utility structure. They showed that computing NE is equivalent to a combinatorial optimization problem over a set system defined by the defender's pure strategy space. A different, but still related sensor positioning game is that of [41], who addressed imperfect detection in a networked control system using a linear filter for processing noisy sensor measurements, and derived equilibrium results through structured systems and graph theory methods.

Our game is also connected to the literature on randomized robust optimization, particularly in settings where an optimizer is concerned with maximizing (resp. minimizing) the worst-case utility (resp. loss) over a collection of submodular (resp. supermodular) functions. However, the existing body of work has primarily emphasized the submodular perspective, leaving the supermodular counterpart arising in our game relatively unexplored. Our work contributes in addressing this gap, broadening the applicability of the techniques in this field.

[122] examined a randomized robust sensor positioning problem akin to a zero-sum game in which the objective was given by a submodular function of the selected set of sensor locations that measured the detection performance against a worst-case intrusion scenario. They extended the analysis of [45] of the MWU algorithm to show that it can be used to efficiently compute approximately optimal mixed strategies for the sensing player, provided that such player can compute approximate best responses—for example, $(1 - 1/e)$-approximations through the use of the greedy algorithm for submodular maximization [130]—the adversary has a polynomial number of strategies, and the payoff

function takes values within the interval $[0, 1]$. In such scenarios, the approximation guarantees associated with the best response extend to the mixed strategy returned by the MWU algorithm, with an additional absolute error that can be made arbitrarily small. [131] utilized this approach to optimize the minmax of potentially non-convex functions, applying it to robust classification with neural networks and robust influence maximization. [132] extended these results for payoff functions taking values within $\mathbb{R}_{\geq 0}$. Using different notions of approximation, [128] employed similar techniques to solve combinatorial zero-sum games with bilinear payoffs, where each player's strategy space corresponds to the vertices of a polytope.

[133] also addresed the submodular setting, but against an adversary with an exponential number of strategies. Using methods from continuous and stochastic optimization, he developed randomized pseudopolynomial-time algorithms with $(1 - 1/e)^2$-approximation guarantees for the maximizer. [134] examined minmax optimization problems where the minimization (resp. maximization) variable in the objective is continuous and convex (resp. discrete and submodular). They showed hardness results and introduced approximation algorithms that integrate nonlinear optimization techniques with the greedy algorithm for submodular maximization.

Our approximate solution approach is related to that of Krause *et al.* However, the supermodularity inherent in our game's payoff prevents the attainment of constant-factor approximations, leading us to rely on instance-dependent approximations contingent upon the curvature of the payoff function [119, 120]. Furthermore, our implementation of the MWU algorithm updates the attacker's unidimensional marginal distributions instead of their original mixed strategies, which requires an additional projection step with respect to the unnormalized relative entropy to recover feasibility of the updated marginals. This projection step generalizes the normalization rule in the standard implementation of the MWU algorithm.

The MWU algorithm [45] is a special case of the Follow the Perturbed Leader algorithm [135], and versions of it introducing projection steps have been proposed by [136, 137, 138] and [126], but primarily in the context of online learning over structured concept classes, where the focus is on minimizing the learner's regret in the adversarial setting. A comprehensive survey on the MWU algorithm and its applications can be found in [139]. An alternative version of the MWU algorithm, which also employs approximate best responses, has been proposed by [47] to obtain approximate equilibria with multiplicative guarantees in zero-sum search games, including hide-and-seek. However, their approach is not directly applicable to our model, as it requires one player to have a linear number of strategies, whereas in our game *both* players face combinatorially many strategies.

Efficient algorithms for computing projections with respect to the unnormalized relative entropy, as well as more general Bregman divergences, onto the simplex and related polytopes—including variants of the capped simplex and the permutahedron—have been studied by [123, 136, 140, 127, 124, 128, 125] and [126]. Unlike these prior works, our game allows players to utilize fewer resources than their respective budgets. Hence, our contribution focuses on the study of projections onto the *full-dimensional* capped simplex. This polytope encompasses vectors within the unit cube, with the constraint that the sum of their entries is *at most* a given resource budget. This contrasts with the *lower-dimensional* capped simplex polytope, which requires the sum of their entries to be *equal* to that budget. We present a proof that yields a closed-form solution for our projection problem, extending the technique of [141] used for analytically solving the projection problem under the Euclidean distance onto the capped simplex. Our solutions lead us to derive a linear-time algorithm that is faster than more general-purpose algorithms for minimizing separable convex functions on base polytopes of polymatroids [127, 128], and that does not require adding perturbations to the relative entropy function to handle its nonsmoothness [124].

Outline. This chapter is organized as follows. In Section 4.2, we introduce the network inspection game. We then review preliminary results that help us investigate our game in

Section 4.3. In Sections 4.4 and 4.5, we derive algorithmic approaches for computing exact and approximate equilibrium strategies, respectively. In Section 4.6, we present our computational results. We then discuss our conclusions and plans for future work in Section 4.7. Finally, the proofs of our results are detailed in Appendix C.1.

4.2 Model Description

We consider a network model comprising a set of *components* $\mathcal{E}$ that can be potential targets for an attacker, and a set of *nodes* $\mathcal{V}$ that serve as detector locations for a defender with the purpose of inspecting the network and detecting the attacks. We let $m := |\mathcal{E}|$ and $n := |\mathcal{V}|$. By placing a detector at a node $v \in \mathcal{V}$, the defender can monitor a subset of components $\mathcal{C}_v \subseteq \mathcal{E}$, which we refer to as the *monitoring set* of node v. Without loss of generality, we assume that each monitoring set is nonempty and that each component can be monitored from at least one node. An example of a network model is shown in Figure 4.1.

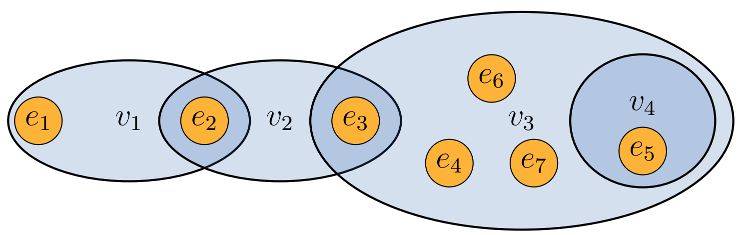

Figure 4.1: Example of a network model. The set of nodes is $\mathcal{V} = \{v_1, v_2, v_3, v_4\}$, and the set of components is $\mathcal{E} = \{e_1, e_2, e_3, e_4, e_5, e_6, e_7\}$. The monitoring sets are $\mathcal{C}_{v_1} = \{e_1, e_2\}$, $\mathcal{C}_{v_2} = \{e_2, e_3\}$, $\mathcal{C}_{v_3} = \{e_3, e_4, e_5, e_6, e_7\}$ and $\mathcal{C}_{v_4} = \{e_5\}$.

Due to inherent characteristics undermining the detection capabilities at each location, we assume imperfect detection of attacks. Specifically, by placing a detector at a node v, the defender identifies an attack on each targeted component in $\mathcal{C}_v$ with a probability p_v, independent of other targeted components and detectors positioned. We refer to p_v as the *detection probability* of node v. Therefore, an attack on a component will only be detected if the defender successfully identifies it through at least one of the detectors positioned at locations monitoring such component.

We assume that both the attacker and the defender are strategic, and hence adopt a game-theoretic framework to study their behaviors. To this aim, we define a simultaneous-move two-person zero-sum game Γ between a defender D and an attacker A. In this game, both players face exogenous resource constraints: D can select up to $r_{\mathrm{D}} \in \mathbb{Z}_{>0}$ nodes from $\mathcal{V}$ to place detectors and monitor a subset of components in $\mathcal{E}$, and A can select up to $r_{\mathrm{A}} \in \mathbb{Z}_{>0}$ components of $\mathcal{E}$ to target. Without loss of generality, we assume that each node can host at most one detector and each component can be the target of at most one attack. Hence, we consider that $r_{\mathrm{D}} \leq n$ and $r_{\mathrm{A}} \leq m$. Formally, let $\mathcal{D} \coloneqq \{S \subseteq \mathcal{V} : |S| \leq r_{\mathrm{D}}\}$ and $\mathcal{A} \coloneqq \{T \subseteq \mathcal{E} : |T| \leq r_{\mathrm{A}}\}$ be the action sets for D and A respectively. We also assume that both players have complete information regarding the network structure (nodes, components, monitoring sets), detection probabilities, and their own and their opponent's action sets.

We consider that in the event of a successful detection of an attack, D can start a response mechanism to mitigate its damage. Thus, in our model, an attack is successful if and only if it remains undetected by D. Let $u : \mathcal{D} \times \mathcal{A} \to \mathbb{R}$ be the *undetection function*, defined as the average number of undetected attacks:

$$u(S,T) \coloneqq \sum_{e \in T} \prod_{v \in S:\, e \in \mathcal{C}_v} (1 - p_v), \quad \forall\, (S,T) \in \mathcal{D} \times \mathcal{A}.$$

We use $u(S,e)$ (resp. $u(v,T)$) to denote the case when $T = \{e\}$ (resp. $S = \{v\}$), for some $e \in \mathcal{E}$ (resp. $v \in \mathcal{V}$). We note that the undetection function satisfies $u(S,T) = \sum_{e \in T} u(S,e)$ for every $(S,T) \in \mathcal{D} \times \mathcal{A}$. The term $u(S,e) = \prod_{v \in S:\, e \in \mathcal{C}_v} (1 - p_v)$ represents the undetection probability of an attack on e when D selects the detector positioning $S \in \mathcal{D}$.

We allow the players to use mixed strategies, defined as probability distributions over their sets of pure actions. To this aim, we let

$$\Delta_{\mathrm{D}} \coloneqq \left\{ \sigma^{\mathrm{D}} \in [0,1]^{\mathcal{D}} : \sum_{S \in \mathcal{D}} \sigma^{\mathrm{D}}_S = 1 \right\}, \quad \Delta_{\mathrm{A}} \coloneqq \left\{ \sigma^{\mathrm{A}} \in [0,1]^{\mathcal{A}} : \sum_{T \in \mathcal{A}} \sigma^{\mathrm{A}}_T = 1 \right\},$$

be the sets of mixed strategies of D and A, respectively. For each $\sigma^{\mathrm{D}} \in \Delta_{\mathrm{D}}$ (resp. $\sigma^{\mathrm{A}} \in \Delta_{\mathrm{A}}$), σ^{D}_S (resp. σ^{A}_T) represents the probability that action $S \in \mathcal{D}$ (resp. $T \in \mathcal{A}$) is executed.

Given a strategy profile $(\sigma^{\mathrm{D}}, \sigma^{\mathrm{A}}) \in \Delta_{\mathrm{D}} \times \Delta_{\mathrm{A}}$, the expected number of undetected attacks is then defined as $U(\sigma^{\mathrm{D}}, \sigma^{\mathrm{A}}) := \sum_{S \in \mathcal{D}} \sum_{T \in \mathcal{A}} \sigma^{\mathrm{D}}_S \, \sigma^{\mathrm{A}}_T \, u(S,T)$. We assume that D (resp. A) aims to minimize (resp. maximize) U. For ease of notation, we use $U(S, \sigma^{\mathrm{A}})$ (resp. $U(\sigma^{\mathrm{D}}, T)$) to denote the case when $\sigma^{\mathrm{D}}_S = 1$ (resp. $\sigma^{\mathrm{A}}_T = 1$) for some $S \in \mathcal{D}$ (resp. $T \in \mathcal{A}$), and omit the brackets when $S = \{v\}$ (resp. $T = \{e\}$).

Our game finds applications in situations involving strategic sensor allocation for the detection of attacks on critical infrastructure networks, particularly those of utilities, such as gas or water distribution companies. In these settings, the utility's network node set $\mathcal{V}$ represents locations (e.g., node junctions, tanks, reservoirs) where the defender is able to position sensors to measure relevant network parameters in real-time, such as pressure and temperature. The set of components $\mathcal{E}$ represents the distribution pipes, which are the potential targets of a malicious cyber-physical attacker, who might manipulate the control systems responsible for managing the pressure in specific sections of the pipeline to an unsustainable level, leading to pipe bursts or leaks. When a sensor is positioned at a node v, it can detect anomalies occurring in a neighboring set of pipes $\mathcal{C}_v$, corresponding to its monitoring set. However, local effects, such as interference from nearby electronic equipment and physical obstructions such as trees or large structures have the potential to disrupt sensor signals. Such disruptions can lead to signal degradation and inaccuracies in sensor readings, encoded in the detection probabilities p_v. The utility (resp. attacker) aims to find a randomized inspection strategy (resp. attack strategy) so as to minimize (resp. maximize) the expected number of undetected attacks.

Our subject of interest is the study of strategy profiles $(\sigma^{\mathrm{D}^*}, \sigma^{\mathrm{A}^*})$ that represent *Nash Equilibria* (NE) of the game Γ, i.e., that satisfy

$$U(\sigma^{\mathrm{D}^*}, \sigma^{\mathrm{A}}) \leq U(\sigma^{\mathrm{D}^*}, \sigma^{\mathrm{A}^*}) \leq U(\sigma^{\mathrm{D}}, \sigma^{\mathrm{A}^*}), \quad \forall\, \sigma^{\mathrm{D}} \in \Delta_{\mathrm{D}},\ \forall\, \sigma^{\mathrm{A}} \in \Delta_{\mathrm{A}}.$$

We refer to $U(\sigma^{\mathrm{D}^*}, \sigma^{\mathrm{A}^*})$ as the *value of the game* Γ. Since Γ is a two-person zero-sum game with a finite number of player actions, a NE is guaranteed to exist and the value of the game is unique. Alternatively, NE can be characterized by strategy profiles $(\sigma^{\mathrm{D}^*}, \sigma^{\mathrm{A}^*})$ where both σ^{D^*} and σ^{A^*} are respectively optimal solutions of the following pair of minmax and maxmin problems:

$$\min_{\sigma^{\mathrm{D}} \in \Delta_{\mathrm{D}}} \max_{\sigma^{\mathrm{A}} \in \Delta_{\mathrm{A}}} U(\sigma^{\mathrm{D}}, \sigma^{\mathrm{A}}), \qquad \max_{\sigma^{\mathrm{A}} \in \Delta_{\mathrm{A}}} \min_{\sigma^{\mathrm{D}} \in \Delta_{\mathrm{D}}} U(\sigma^{\mathrm{D}}, \sigma^{\mathrm{A}}).$$

Furthermore, the optimal values of both these problems are identical and coincide with the value of the game ([89]; see also [90]). From this equivalence, σ^{D^*} and σ^{A^*} are also referred to as *optimal strategies* of the game Γ and can be computed by solving the following Linear Program (LP):

$$\begin{aligned} \min_{\sigma^{\mathrm{D}} \in \mathbb{R}^{\mathcal{D}},\, z \in \mathbb{R}} \quad & z \\ \text{subject to} \quad & \sum_{S \in \mathcal{D}} \sigma^{\mathrm{D}}_S \, u(S,T) \leq z \quad \forall T \in \mathcal{A}, \\ & \sum_{S \in \mathcal{D}} \sigma^{\mathrm{D}}_S = 1 \\ & \sigma^{\mathrm{D}}_S \geq 0 \quad \forall S \in \mathcal{D}. \end{aligned} \tag{LP}$$

(LP) and its dual are reformulations of the minmax and maxmin problems, respectively. Thus, equilibirum inspection (resp. attack) strategies and the value of the game Γ are given by the optimal primal (resp. dual) solutions and optimal value of (LP), respectively. Nonetheless, since the cardinality of $\mathcal{D}$ (resp. $\mathcal{A}$) grows combinatorially with r_{D} (resp. r_{A}), this LP becomes computationally challenging to solve, even for small-sized networks. Hence, we derive solution approaches that leverage the structure of the undetection function, allowing us to compute exact and approximate equilibria of the game Γ.

4.3 Preliminaries

In our analysis, it will be useful to characterize the attacker's strategies in terms of their marginal probabilities of targeting each component. For every mixed attack strategy $\sigma^{\mathrm{A}} \in \Delta_{\mathrm{A}}$, we denote as $\rho(\sigma^{\mathrm{A}}) = (\rho_e(\sigma^{\mathrm{A}}))_{e \in \mathcal{E}}$ the vector of marginal probabilities of targeting each component when A plays σ^{A}, given by

$$\rho_e(\sigma^{\mathrm{A}}) \coloneqq \sum_{T \in \mathcal{A} \colon e \in T} \sigma^{\mathrm{A}}_T, \quad \forall\, e \in \mathcal{E}. \tag{4.1}$$

The additive structure of the undetection function $u(S,T)$ with respect to A's pure actions permits us to represent the expected number of undetected attacks in terms of these marginal attack probabilities and the undetection probabilities resulting from D's mixed inspection strategy.

Lemma 8. *For every* $(\sigma^{\mathrm{D}}, \sigma^{\mathrm{A}}) \in \Delta_{\mathrm{D}} \times \Delta_{\mathrm{A}}$, $U(\sigma^{\mathrm{D}}, \sigma^{\mathrm{A}}) = \sum_{e \in \mathcal{E}} \rho_e(\sigma^{\mathrm{A}})\, U(\sigma^{\mathrm{D}}, e)$.

From Lemma 8, the expected number of undetected attacks is given by the sum, over each component, of the probability that such component is targeted by A and the attack remains undetected by D. While every mixed attack strategy induces a vector of marginal attack probabilities in $\mathbb{R}^{\mathcal{E}}$, we are also interested in the inverse question of identifying the vectors in $\mathbb{R}^{\mathcal{E}}$ that represent marginal attack probabilities resulting from some mixed attack strategy. This bidirectional mapping will facilitate transitioning between the spaces of mixed strategies and marginal attack probabilities. The following lemma provides necessary and sufficient conditions that such vectors must satisfy.

Lemma 9 (see Lemma 1 and Theorem 2)**.** *Consider a resource budget* $r_{\mathrm{A}} \in \mathbb{Z}_{>0}$*, and a vector* $\rho \in \mathbb{R}^{\mathcal{E}}$*. Then, there exists a probability distribution* $\sigma^{\mathrm{A}} \in \Delta_{\mathrm{A}}$ *satisfying* $\rho_e(\sigma^{\mathrm{A}}) = \rho_e$ *for all* $e \in \mathcal{E}$ *if and only if* $\rho \in [0,1]^{\mathcal{E}}$ *and* $\sum_{e \in \mathcal{E}} \rho_e \leq r_{\mathrm{A}}$*. Furthermore, such* σ^{A} *with a support of size at most* $m+1$ *can be computed in time* $O(m^2)$.

As a consequence of Lemma 9, the vectors in $\mathbb{R}^{\mathcal{E}}$ that represent marginal attack probabilities resulting from A's mixed strategies are precisely the feasible solutions of the *full-dimensional capped simplex* polytope:

$$\Theta_{\mathrm{A}} := \left\{ \rho \in [0,1]^{\mathcal{E}} : \sum_{e \in \mathcal{E}} \rho_e \leq r_{\mathrm{A}} \right\}.$$

This provides an alternative representation of A's strategies as vectors in Θ_{A}, which offers a computational advantage, as the elements of Θ_{A} are vectors of size m, while A's mixed strategies require vectors of size $\sum_{k=0}^{r_{\mathrm{A}}} \binom{m}{k}$.

In view of Lemmas Lemma 8 and Lemma 9, for the remainder of this chapter we denote by $U(\sigma^{\mathrm{D}}, \rho^{\mathrm{A}}) := \sum_{e \in \mathcal{E}} \rho_e^{\mathrm{A}} U(\sigma^{\mathrm{D}}, e)$ the expected number of undetected attacks when D plays the *mixed* inspection strategy $\sigma^{\mathrm{D}} \in \Delta_{\mathrm{D}}$ and A plays the *marginal* attack strategy $\rho^{\mathrm{A}} \in \Theta_{\mathrm{A}}$. The reader can verify that equilibria of the game Γ can be equivalently represented by strategy profiles $(\sigma^{\mathrm{D}^*}, \rho^{\mathrm{A}^*}) \in \Delta_{\mathrm{D}} \times \Theta_{\mathrm{A}}$ satsifying $U(\sigma^{\mathrm{D}^*}, \rho^{\mathrm{A}}) \leq U(\sigma^{\mathrm{D}^*}, \rho^{\mathrm{A}^*}) \leq U(\sigma^{\mathrm{D}}, \rho^{\mathrm{A}^*})$ for every $\sigma^{\mathrm{D}} \in \Delta_{\mathrm{D}}$ and $\rho^{\mathrm{A}} \in \Theta_{\mathrm{A}}$, and that the value of the game is $U(\sigma^{\mathrm{D}^*}, \rho^{\mathrm{A}^*})$. Similarly, $(\sigma^{\mathrm{D}^*}, \rho^{\mathrm{A}^*})$ is an equilibrium of Γ if and only if σ^{D^*} is an optimal solution of the minmax problem $\min_{\sigma^{\mathrm{D}} \in \Delta_{\mathrm{D}}} \max_{\rho^{\mathrm{A}} \in \Theta_{\mathrm{A}}} U(\sigma^{\mathrm{D}}, \rho^{\mathrm{A}})$, and ρ^{A^*} is an optimal solution of the maxmin problem $\max_{\rho^{\mathrm{A}} \in \Theta_{\mathrm{A}}} \min_{\sigma^{\mathrm{D}} \in \Delta_{\mathrm{D}}} U(\sigma^{\mathrm{D}}, \rho^{\mathrm{A}})$.

We remark that the undetection function is impacted by the simultaneous positioning of D's detectors. This nonadditive characteristic prevents us from compactly representing inspection strategies using only the marginal probabilities of inspecting each node, and instead requires considering probability distributions over the large set of detector positionings $\mathcal{D}$. Nonetheless, we note that in the special case where the monitoring sets are mutually disjoint, the undetection function becomes additive in D's strategies, rendering the game equivalent to the hide-and-seek game of Chapter 2, where D is the seeker, A is the hider, and $\mathcal{V}$ is the set of hiding locations. In this context, each hiding location v has a hiding capacity given by $|\mathcal{C}_v|$, and a probability of detecting hidden items by the seeker,

p_v. In such instances, as shown in Chapter 2 (see also [93]), equilibrium strategies of the game Γ can be analytically characterized in terms of the marginal probabilities of inspecting each node and the expected number of components targeted in each monitoring set. Furthermore, the game can be solved in time $O(n^2)$.

4.4 Exact Solution Method

In this section, we introduce an exact solution approach for solving the game Γ. We recall that equilibrium strategies of Γ are optimal primal and dual solutions of (LP), which is challenging to solve due to its combinatorial number of variables and constraints. However, by leveraging the compact representation of A's strategies described in Section 4.3, we can formulate an alternative LP with a significantly smaller number of constraints.

Proposition 3. (LP) *is equivalent to:*

$$\begin{aligned} \min_{\sigma^{\mathrm{D}}\in\mathbb{R}^{\mathcal{D}},\, \lambda\in\mathbb{R}^{\mathcal{E}},\, \gamma\in\mathbb{R}} \quad & r_{\mathrm{A}}\gamma + \sum_{e\in\mathcal{E}} \lambda_e \\ \text{subject to} \quad & \gamma + \lambda_e \geq \sum_{S\in\mathcal{D}} \sigma^{\mathrm{D}}_S\, u(S,e) \quad && \forall\, e\in\mathcal{E}, \\ & \sum_{S\in\mathcal{D}} \sigma^{\mathrm{D}}_S = 1 \\ & \sigma^{\mathrm{D}}_S \geq 0 && \forall\, S\in\mathcal{D}, \\ & \lambda_e \geq 0 && \forall\, e\in\mathcal{E}, \\ & \gamma \geq 0. \end{aligned} \tag{LP($\mathcal{D}$)}$$

In particular, let $(\sigma^{\mathrm{D}^*}, \lambda^*, \gamma^*) \in \mathbb{R}^{\mathcal{D}} \times \mathbb{R}^{\mathcal{E}} \times \mathbb{R}$ *be an optimal solution of* (LP($\mathcal{D}$))*, and let* $\rho^{\mathrm{A}^*} \in \mathbb{R}^{\mathcal{E}}$ *be the vector of optimal dual variables associated with its first set of constraints. Then,* $(\sigma^{\mathrm{D}^*}, \rho^{\mathrm{A}^*})$ *is an equilibrium of the game* Γ*, and the value of the game is given by* $r_{\mathrm{A}}\gamma^* + \sum_{e\in\mathcal{E}} \lambda^*_e$.

From Proposition 3, we observe that mixed inspection strategies and marginal attack strategies in equilibrium of Γ are respectively optimal primal and dual solutions of an LP with $|\mathcal{D}| + m + 1$ variables and $m + 1$ constraints. In particular, the size of (LP($\mathcal{D}$)) is independent of the number of attack resources r_{A}. This is a consequence of the alternative representation of the attacker's strategies in terms of their marginal probabilities, which, as feasible solutions of Θ_{A}, only require $m + 1$ constraints for their description.

We can solve (LP($\mathcal{D}$)) using a Column Generation (CG) algorithm, provided that we can derive a relatively small formulation of the pricing problem. To this aim, we consider a subset of D's pure actions $\mathcal{I} \subseteq \mathcal{D}$ and define the associated restricted master problem as (LP($\mathcal{I}$)). Let $(\rho^{\mathrm{A}}, \nu) \in \mathbb{R}^{\mathcal{E}} \times \mathbb{R}$ be optimal dual variables associated with the first set of constraints and second constraint of (LP($\mathcal{I}$)), respectively. One can easily check that $\rho^{\mathrm{A}} \in \Theta_{\mathrm{A}}$, and that it corresponds to the marginal attack probabilities of A's best response to the optimal solution of (LP($\mathcal{I}$)). Consequently, the reduced costs of the variables γ and λ_e are nonnegative. On the other hand, the reduced costs of the variables σ_S^{D} are given by

$$\bar{c}_S = -\nu + \sum_{e \in \mathcal{E}} \rho_e^{\mathrm{A}}\, u(S, e) = -\nu + U(S, \rho^{\mathrm{A}}), \quad \forall\, S \in \mathcal{D}.$$

Therefore, the pricing problem is:

$$\text{Find } S^* \in \arg\min_{S \in \mathcal{D}} U(S, \rho^{\mathrm{A}}). \tag{DBR}$$

In other words, the pricing problem (DBR) refers to finding a deterministic detector positioning that minimizes the expected number of undetected attacks against A playing the marginal attack strategy ρ^{A}. This problem corresponds to D's pure best response problem, which can be solved using the techniques we describe next.

4.4.1 Exact Solutions for Defender's Best Response Problem

We now turn into the computation of pure best responses for D. Generally, the combinatorial nature and nonlinearity of the undetection function U render Problem (DBR) challenging to solve, even for instances where the monitoring sets exhibit minimal overlap.

Proposition 4. *Problem* (DBR) *is* NP*-hard, even if* $r_{\mathrm{A}} = 1$*, the detection probabilities are homogeneous, and every component belongs to at most two monitoring sets.*

The proof of Proposition 4 follows by a reduction from the NP-hard problem VERTEX COVER. Despite this negative result, an exact solution for Problem (DBR) can be computed using a Mixed-Integer Program (MIP). For this purpose, we first set an arbitrary enumeration of the detector locations, denoted as $\mathcal{V} = \{v_1, \ldots, v_n\}$. Then, for every node $v_i \in \mathcal{V}$ we define a binary variable x_i as

$$x_i = \begin{cases} 1 & \text{if a detector is positioned at node } v_i, \\ 0 & \text{otherwise.} \end{cases}$$

Thus, if $S = \{v_i \in \mathcal{V} : \ x_i = 1\}$ is a detector positioning determined by the variables x, we have $S \in \mathcal{D}$ if and only if $\sum_{v_i \in \mathcal{V}} x_i \leq r_{\mathrm{D}}$. Next, we must linearize the term $u(S, e) = \prod_{v_i \in S:\, e \in \mathcal{C}_{v_i}} (1 - p_{v_i})$ in the undetection function. To this end, we define the following decision variables:

$$u_{e,i} = \prod_{v_j \in \mathcal{V}:\, j \leq i} \Big(1 - p_{v_j} \mathbb{1}_{\mathcal{C}_{v_j}}(e) x_j\Big) \quad \forall\, e \in \mathcal{E}, \forall\, v_i \in \mathcal{V},$$

where for every $e \in \mathcal{E}$ and $v_j \in \mathcal{V}$, $\mathbb{1}_{\mathcal{C}_{v_j}}(e) \in \{0, 1\}$ is a parameter defined as $\mathbb{1}_{\mathcal{C}_{v_j}}(e) = 1$ if and only if $e \in \mathcal{C}_{v_j}$. Then, $u(S, e) = u_{e,n}$ for all $e \in \mathcal{E}$. Finally, we can linearize the variables $u_{e,i}$ by leveraging the following relations derived from their definition:

$$\begin{aligned} \big(1 - p_{v_1} \mathbb{1}_{\mathcal{C}_{v_1}}(e) x_1\big) &= u_{e,1} && \forall\, e \in \mathcal{E}, \\ u_{e,i} \Big(1 - p_{v_{i+1}} \mathbb{1}_{\mathcal{C}_{v_{i+1}}}(e) x_{i+1}\Big) &= u_{e,i+1} && \forall\, e \in \mathcal{E}, \forall\, v_i \in \mathcal{V} \setminus \{v_n\}, \\ 0 \leq u_{e,i} &\leq 1 && \forall\, v_i \in \mathcal{V}. \end{aligned}$$

This results in the following MIP formulation for (DBR).

Proposition 5. *Consider an arbitrary enumeration of the set of detector locations as* $\mathcal{V} = \{v_1, \ldots, v_n\}$. *Then, Problem* (DBR) *can be formulated as the following MIP:*

$$
\begin{aligned}
\min_{x\in\mathbb{R}^n,\, u\in\mathbb{R}^{\mathcal{E}\times\{1,\ldots,n\}}} \quad & \sum_{e\in\mathcal{E}} \rho_e^{\mathrm{A}}\, u_{e,n} & \\
\textit{subject to} \quad & \sum_{v_i\in\mathcal{V}} x_i \le r_{\mathrm{D}} & \\
& 1 - p_{v_1} \mathbb{1}_{\mathcal{C}_{v_1}}(e) x_1 \le u_{e,1} & \forall\, e \in \mathcal{E}, \\
& u_{e,i}\left(1 - p_{v_{i+1}} \mathbb{1}_{\mathcal{C}_{v_{i+1}}}(e)\right) \le u_{e,i+1} & \forall\, e \in \mathcal{E},\, \forall\, v_i \in \mathcal{V}\setminus\{v_n\}, \\
& u_{e,i} - x_{i+1} \le u_{e,i+1} & \forall\, e \in \mathcal{E},\, \forall\, v_i \in \mathcal{V}\setminus\{v_n\}, \\
& 0 \le u_{e,i} \le 1 & \forall\, v_i \in \mathcal{V}, \\
& x_i \in \{0,1\} & \forall\, v_i \in \mathcal{V}.
\end{aligned}
\tag{MIP}
$$

In particular, let $(x^*, u^*) \in \mathbb{R}^n \times \mathbb{R}^{\mathcal{E}\times\{1,\ldots,n\}}$ *be an optimal solution of* (MIP). *Then,* $S^* := \{v_i \in \mathcal{V} :\ x_i^* = 1\}$ *is an optimal solution of* (DBR), *and* $U(S^*, \rho^{\mathrm{A}}) = \sum_{e\in\mathcal{E}} \rho_e^{\mathrm{A}}\, u_{e,n}^*$.

(MIP) comprises $O(mn)$ variables and constraints, and it can be used to compute the variable σ_S^{D} of the master problem (LP($\mathcal{D}$)) with lowest reduced cost. Consequently, equilibrium strategies for Γ can be computed by solving (LP($\mathcal{D}$)) using a CG algorithm, which utilizes (MIP) to solve the pricing problem.

4.5 Approximate Solution Methods

We now delve into the computation of approximate equilibrium strategies for the game Γ. We begin by analyzing the structure of Problem (DBR) and explore the utility of approximation algorithms. We then propose two solution approaches that leverage these approximations: one utilizing Column Generation (CG), and the other employing the Multiplicative Weights Update (MWU) algorithm.

4.5.1 Approximate Solutions for Defender's Best Response Problem

We first examine structural properties of Problem (DBR) that will allow us to efficiently compute approximate best responses for a large class of instances.

Lemma 10. *For every fixed $\rho^{\mathrm{A}} \in \Theta_{\mathrm{A}}$, the set function $U(S, \rho^{\mathrm{A}}) := \sum_{e\in\mathcal{E}} \rho_e^{\mathrm{A}}\, u(S, e)$, defined for every $S \subseteq \mathcal{V}$, is nondecreasing and supermodular.*

Therefore, for a fixed A's marginal attack strategy $\rho^{\mathrm{A}} \in \Theta_{\mathrm{A}}$, D's best response problem (DBR) consists in minimizing the nonnegative and nonincreasing supermodular function $U(\cdot, \rho^{\mathrm{A}})$, subject to a cardinality constraint given by $S \in \mathcal{D}$. Let $S^* \in \mathcal{D}$ be an optimal solution of (DBR). For $\alpha \geq 1$, a detector positioning $\widehat{S} \in \mathcal{D}$ is an α-approximate best response for D if it satisfies

$$U(S^*, \rho^{\mathrm{A}}) \leq U(\widehat{S}, \rho^{\mathrm{A}}) \leq \alpha U(S^*, \rho^{\mathrm{A}}). \tag{4.2}$$

To approximate (DBR), [119] and [120] showed that a *reverse* greedy algorithm—also referred to as greedy descent or stingy algorithm—finds a solution whose approximation guarantee depends on the *curvature* of the function $U(\cdot, \rho^{\mathrm{A}})$, defined as

$$c := 1 - \min_{v\in\mathcal{V}:\, U_{\emptyset}(v,\rho^{\mathrm{A}})>0} \frac{U_{\mathcal{V}\setminus\{v\}}(v, \rho^{\mathrm{A}})}{U_{\emptyset}(v, \rho^{\mathrm{A}})},$$

where $U_S(v, \rho^{\mathrm{A}}) := U(S, \rho^{\mathrm{A}}) - U(S \cup \{v\}, \rho^{\mathrm{A}})$ is the marginal decrease in the expected number of undetected attacks when $v \in \mathcal{V}$ is added to the set of detector locations $S \subseteq \mathcal{V}$. We let $c := 0$ if $U_{\emptyset}(v, \rho^{\mathrm{A}}) = 0$ for every $v \in \mathcal{V}$. The curvature parameter c lies in the interval $[0, 1]$ and measures how far $U(\cdot, \rho^{\mathrm{A}})$ is from being an additive set function. In particular, when $c = 0$, the function $U(\cdot, \rho^{\mathrm{A}})$ is additive.

In practical scenarios, it is often the case that each component can be inspected from a small number of detector locations. In such cases, we can provide the following upper bound on the curvature of $U(\cdot, \rho^{\mathrm{A}})$.

Lemma 11. *The curvature parameter c of the function $U(\cdot, \rho^{\mathrm{A}})$ satisfies $c \leq 1 - (1 - \max_{v \in \mathcal{V}} p_v)^d$, where d is the maximum number of locations that can monitor a component.*

The bound in Lemma 11 is independent of the attacker's marginal strategy ρ^{A}, and gives useful information in the case of imperfect detection—that is, $p_v < 1$ for every $v \in \mathcal{V}$. However, it offers no insights in the presence of perfect detection. In fact, even if all detectors are perfect, it is not necessarily the case that $c = 1$. For a counterexample, consider instances with mutually disjoint monitoring sets; in such scenarios, $c = 0$ irrespective of the detection probabilities.

Starting from the full set of nodes $\mathcal{V}$, each iteration of the reverse greedy algorithm selects a node that minimizes the marginal decrease in the expected number of undetected attacks that would result from its removal. Then, it removes such node from the current solution. This process is repeated until r_{D} nodes are left, resulting in a feasible detector positioning. We refer to Algorithm 2 for a pseudocode of the reverse greedy algorithm.

Algorithm 2: Reverse Greedy Algorithm for Approximate D's Best Response

Initialize $S \leftarrow \mathcal{V}$
while $|S| > r_{\mathrm{D}}$ **do**
    $v \in \arg\min_{w \in \mathcal{V}} U_{S \setminus \{w\}}(w, \rho^{\mathrm{A}})$
    $S \leftarrow S \setminus \{v\}$
return S

The reverse greedy algorithm exhibits the following approximation guarantee.

Theorem 5 ([120])**.** *The reverse greedy algorithm (Algorithm 2) returns a detector positioning $\widehat{S} \in \mathcal{D}$ satisfying $U(\widehat{S}, \rho^{\mathrm{A}}) \leq (1/(1-c))\, U(S^*, \rho^{\mathrm{A}})$.*

Therefore, by applying the bound provided in Lemma 11 to Theorem 5, we derive the following instance-dependent approximation result.

Corollary 2. *The reverse greedy algorithm (Algorithm 2) returns a detector positioning $\widehat{S} \in \mathcal{D}$ satisfying $U(\widehat{S}, \rho^{\mathrm{A}}) \leq \left(1/\left(1 - \max_{v \in \mathcal{V}} p_v\right)^d\right) U(S^*, \rho^{\mathrm{A}})$, where d represents the maximum number of locations that can monitor a component.*

From Theorem 5, the reverse greedy algorithm finds a detector positioning for which the expected number of undetected attacks against A's marginal attack strategy ρ^{A} is within a factor of $1/(1-c)$ of that of D's best response. This approximation factor is strictly increasing in the curvature parameter c, and it degrades as c approaches 1. In the limit $c = 1$, we can show that Problem (DBR) becomes notably challenging to approximate.

Proposition 6. *If* $c = 1$*, then, unless* P = NP*, there is no polynomial time approximation algorithm for* (DBR).

Proposition 6 is a consequence of the fact that deciding whether the optimal value of (DBR) is zero is NP-hard. In such cases, inequality (4.2) implies that any α-approximate best response for D must necessarily be an optimal solution of (DBR). Thus, computing such an approximation in polynomial time is infeasible, unless P = NP.

In scenarios with expensive detection resources, the number of available detectors r_{D} may be significantly smaller than the number of nodes n. In such instances, the reverse greedy algorithm needs to remove a large number of nodes from $\mathcal{V}$ before termination. This can potentially lower the performance of solution approaches that frequently utilize reverse greedy as a subroutine. To mitigate this issue, we also consider the *forward* greedy algorithm as an approximation method for (DBR). This algorithm begins with an empty set of nodes, and in each iteration, selects a node that maximizes the marginal decrease in the expected number of undetected attacks, subsequently adding it to the current solution, until r_{D} nodes are selected. We refer to Algorithm 3 for a pseudocode of the forward greedy algorithm.

Algorithm 3: Forward Greedy Algorithm for Approximate D's Best Response

Initialize $S \leftarrow \emptyset$
while $|S| < r_{\mathrm{D}}$ **do**
  $v \in \arg\max_{w \in \mathcal{V}} U_S(w, \rho^{\mathrm{A}})$
  $S \leftarrow S \cup \{v\}$
return S

By leveraging the approximation guarantee of [142] for submodular maximization with cardinality constraints, we can show that the forward greedy algorithm provides the following guarantee.

Proposition 7. *The forward greedy algorithm (Algorithm 3) returns a detector positioning $\widehat{S} \in \mathcal{D}$ satisfying $U(\widehat{S}, \rho^{\mathrm{A}}) \leq \left(\frac{1-\exp(-c)}{c}\right) U(S^*, \rho^{\mathrm{A}}) + \left(1 - \frac{1-\exp(-c)}{c}\right) r_{\mathrm{A}}$.*

The approximation guarantee of Proposition 7 is weaker than that of Theorem 5. In particular, when the optimal value of (DBR) is zero, the reverse greedy algorithm returns an optimal solution, whereas the forward greedy algorithm may yield a solution whose relative approximation error is unbounded.

4.5.2 Column Generation with Approximate Best Response

The exact CG algorithm of Section 4.4 requires solving (MIP) in each iteration to compute optimal solutions for the pricing problem (DBR), which may be computationally expensive. Alternatively, we can utilize approximation algorithms—such as the forward and reverse greedy algorithms—to obtain fast solutions for (DBR). In such cases, we can show that the multiplicative approximation guarantee is inherited by the strategy profile returned by the CG algorithm. Additionally, we can enable early termination once the reduced cost of the variable of (LP($\mathcal{D}$)) associated with the approximate best response is sufficiently close to zero, at the cost of an additional small additive error.

Theorem 6. *Let $\varepsilon \geq 0$. Consider a solution $(\widehat{\sigma}^{\mathrm{D}}, \widehat{\lambda}, \widehat{\gamma}) \in \mathbb{R}^{\mathcal{D}} \times \mathbb{R}^{\mathcal{E}} \times \mathbb{R}$ of* (LP($\mathcal{D}$))*, together with a vector of dual variables $\widehat{\rho}^{\mathrm{A}} \in \mathbb{R}^{\mathcal{E}}$ associated with its first set of constraints, attained through a CG algorithm which in each iteration solves the restricted master problem* (LP($\mathcal{I}$)) *and computes an α-approximation $\widehat{S} \in \mathcal{D}$ for* (DBR)*, and that terminates once the reduced cost of the variable $\sigma^{\mathrm{D}}_{\widehat{S}}$ satisfies $\bar{c}_{\widehat{S}} \geq -\varepsilon$. Then, the strategy profile $(\widehat{\sigma}^{\mathrm{D}}, \widehat{\rho}^{\mathrm{A}})$*

satisfies $U(\widehat{\sigma}^{\mathrm{D}}, \widehat{\rho}^{\mathrm{A}}) = r_{\mathrm{A}}\widehat{\gamma} + \sum_{e\in\mathcal{E}} \widehat{\lambda}_e$ *and*

$$U(\widehat{\sigma}^{\mathrm{D}}, \rho^{\mathrm{A}}) \leq U(\widehat{\sigma}^{\mathrm{D}}, \widehat{\rho}^{\mathrm{A}}) \leq \alpha\, U(\sigma^{\mathrm{D}}, \widehat{\rho}^{\mathrm{A}}) + \varepsilon, \quad \forall\, \sigma^{\mathrm{D}} \in \Delta_{\mathrm{D}},\ \forall\, \rho^{\mathrm{A}} \in \Theta_{\mathrm{A}}.$$

Furthermore, $\widehat{\sigma}^{\mathrm{D}}$ *and* $\widehat{\rho}^{\mathrm{A}}$ *respectively satisfy the following bounds for the worst-case expected number of undetected attacks with respect to the value of the game:*

$$\begin{aligned} U(\sigma^{\mathrm{D}^*}, \rho^{\mathrm{A}^*}) \leq \max_{\rho^{\mathrm{A}}\in\Theta_{\mathrm{A}}} U(\widehat{\sigma}^{\mathrm{D}}, \rho^{\mathrm{A}}) &\leq \alpha\, U(\sigma^{\mathrm{D}^*}, \rho^{\mathrm{A}^*}) + \varepsilon, \\ \frac{1}{\alpha}\left(U(\sigma^{\mathrm{D}^*}, \rho^{\mathrm{A}^*}) - \varepsilon\right) \leq \min_{\sigma^{\mathrm{D}}\in\Delta_{\mathrm{D}}} U(\sigma^{\mathrm{D}}, \widehat{\rho}^{\mathrm{A}}) &\leq U(\sigma^{\mathrm{D}^*}, \rho^{\mathrm{A}^*}). \end{aligned}$$

The strategy profile $(\widehat{\sigma}^{\mathrm{D}}, \widehat{\rho}^{\mathrm{A}})$ returned by the CG algorithm with α-approximate best response for (DBR) described in Theorem 6 is such that D can decrease the expected number of undetected attacks by a factor of at most $1/\alpha - \varepsilon/U(\widehat{\sigma}^{\mathrm{D}}, \widehat{\rho}^{\mathrm{A}})$ through unilateral deviation from $(\widehat{\sigma}^{\mathrm{D}}, \widehat{\rho}^{\mathrm{A}})$, while A's strategy $\widehat{\rho}^{\mathrm{A}}$ is already a best response to $\widehat{\sigma}^{\mathrm{D}}$. We observe that the parameter $\varepsilon \geq 0$ used for the early termination criterion translates into an additive error in the approximation guarantees. In particular, when $\varepsilon = 0$, $\widehat{\sigma}^{\mathrm{D}}$ is an α-approximation of the optimal inspection strategy σ^{D^*}, whereas $\widehat{\rho}^{\mathrm{A}}$ is an $(1/\alpha)$-approximation of the optimal marginal attack strategy ρ^{A^*}. In the case of imperfect detection, Corollary Corollary 2 guarantees an approximation factor of $\alpha \leq 1/\left(1 - \max_{v\in\mathcal{V}} p_v\right)^d$, where d is the maximum number of locations that can monitor a component.

Interestingly, the termination criterion stated in Theorem 6 only requires the reduced cost of the variable in (LP($\mathcal{D}$)) associated with the α-approximate best response $\widehat{S}$ to be at least $-\varepsilon$. In particular, using strong duality arguments, we can show that it is unnecessary to require that $\widehat{S}$ already belong to the set of columns $\mathcal{I}$ upon termination, a condition stipulated by [121]. This observation potentially leads to significantly faster convergence in practice.

4.5.3 Multiplicative Weights Update with Approximate Best Response

While the CG algorithm with approximate best response of Section 4.5.2 provides approximate equilibrium strategies for the game Γ, its running time may remain exponential due to the potential need to add an exponential number of columns before convergence. We now propose a second solution approach for computing approximate equilibria of Γ in polynomial time. Here, we adapt the method introduced by [122], which employs the MWU algorithm in conjunction with an approximate best response oracle to compute randomized sensing strategies for a game with submodular structure.

Although our game shares similarities with that of Krause *et al.*, a key distinction makes ours more intricate to solve. In their setting, it is assumed that the adversary has a polynomial number of pure strategies, allowing each iteration of the MWU algorithm—which involves updating the adversary's mixed strategy vector—to be implementable in polynomial time. In contrast, in our game, the adversary A has $|\mathcal{A}| = \sum_{k=0}^{r_{\mathrm{A}}} \binom{m}{k}$ pure strategies, a number that grows combinatorially with r_{A}, rendering the update step of the adversary's strategies computationally expensive. To address this challenge, we consider an implementation of the MWU algorithm that updates A's marginal attack probabilities, instead of their original mixed strategies.

We initialize our implementation of the MWU algorithm by setting uniform marginal attack probabilities, namely, $\rho_e^{(1)} = r_{\mathrm{A}}/m$ for every $e \in \mathcal{E}$. Each iteration begins with a vector of marginal attack probabilities $\rho^{(t)} \in \Theta$ for A, after which D computes an approximate best response $S^{(t)} \in \mathcal{D}$. Next, A uses $S^{(t)}$ to update $\rho^{(t)}$ to a new marginal attack strategy $\rho^{(t+1)}$. This update is performed in two steps: First, an intermediate updated vector $\widetilde{\rho}^{(t+1)}$ is constructed by applying the multiplicative weights rule, that is, $\widetilde{\rho}_e^{(t+1)} \leftarrow \rho_e^{(t)} \exp\left(\eta u(S^{(t)}, e)\right)$ for every $e \in \mathcal{E}$, where $\eta > 0$ is a parameter of the algorithm. This step aims to increase the marginal attack probabilities for components with a higher undetection probability under $S^{(t)}$. However, the resulting vector $\widetilde{\rho}^{(t+1)}$ may not be a feasible solution of Θ_{A}, and therefore, may not be the marginal distribution resulting from some mixed strategy in Δ_{A} (Lemma 9). To rectify this issue, we proceed to the second part of the update step, wherein we project $\widetilde{\rho}^{(t+1)}$ onto Θ_{A} with respect to the unnormalized relative entropy, which is defined as

$$\mathrm{D}_{\mathrm{RE}}\left(\rho \,\|\, \widetilde{\rho}\right) \coloneqq \sum_{e \in \mathcal{E}} \left(\rho_e \ln \frac{\rho_e}{\widetilde{\rho}_e} + \widetilde{\rho}_e - \rho_e\right), \quad \forall\, \rho \in \mathbb{R}^{\mathcal{E}}_{\geq 0}, \forall\, \widetilde{\rho} \in \mathbb{R}^{\mathcal{E}}_{>0}.$$

We refer to Algorithm 4 for a pseudocode of our implementation of the MWU algorithm.

Algorithm 4: Multiplicative Weights Update with Approximate Best Response

Parameters: $\tau \in \mathbb{Z}_{>0}$: number of iterations, $\eta \in \mathbb{R}_{>0}$: multiplicative weights factor

1 Initialize $\rho_e^{(1)} \leftarrow \frac{r_{\mathrm{A}}}{m}$, $\forall\, e \in \mathcal{E}$
2 **for** $t = 1, \dots, \tau$ **do**
    /* Compute Defender's α-approximate Best Response */
3   Let $S^{(t)} \in \mathcal{D}$ satisfying $U(S^{(t)}, \rho^{(t)}) \leq \alpha \min_{S \in \mathcal{D}} U(S, \rho^{(t)})$
    /* Update Attacker's Marginal Probabilities */
4   **for** $e \in \mathcal{E}$ **do**
5     $\widetilde{\rho}_e^{(t+1)} \leftarrow \rho_e^{(t)} \exp\left(\eta u(S^{(t)}, e)\right)$
    /* Project $\widetilde{\rho}^{(t+1)}$ with respect to unnormalized relative entropy */
6   Let $\rho^{(t+1)} \in \arg\min_{\rho \in \Theta_{\mathrm{A}}} \mathrm{D}_{\mathrm{RE}}\left(\rho \,\middle\|\, \widetilde{\rho}^{(t+1)}\right)$
  /* Return average distributions */
7 **return** $\widehat{\sigma}^{\mathrm{D}} \coloneqq \frac{1}{\tau}\sum_{t=1}^{\tau} \mathbb{1}_{S^{(t)}}$ *and* $\widehat{\rho}^{\mathrm{A}} \coloneqq \frac{1}{\tau}\sum_{t=1}^{\tau} \rho^{(t)}$

The MWU algorithm returns a mixed inspection strategy $\widehat{\sigma}^{\mathrm{D}} \coloneqq \frac{1}{\tau}\sum_{t=1}^{\tau} \mathbb{1}_{S^{(t)}}$, where $\mathbb{1}_{S^{(t)}} \in \{0,1\}^{\mathcal{D}}$ is the characteristic vector of $S^{(t)}$—that is, for every $S \in \mathcal{D}$, $\mathbb{1}_{S^{(t)}}(S) = 1$ if and only if $S = S^{(t)}$. In other words, $\widehat{\sigma}^{\mathrm{D}}_S$ is equal to the relative frequency with which S was selected as D's best response (Line 3) throughout the execution of the algorithm. Therefore, the size of the support of $\widehat{\sigma}^{\mathrm{D}}$ is at most the number of iterations τ. The MWU algorithm also returns a marginal attack strategy $\widehat{\rho}^{\mathrm{A}} \coloneqq \frac{1}{\tau}\sum_{t=1}^{\tau} \rho^{(t)}$, given by the average marginal attack probabilities obtained after completing each update step (Line 6). These strategies satisfy the following approximation guarantees.

Theorem 7. *Let* $\varepsilon > 0$. *Then, after* $\tau \geq 4r_{\mathrm{A}}^2 \max\left\{\ln(m/r_{\mathrm{A}}), 1\right\}/\varepsilon^2$ *iterations, Algorithm 4* *with* $\eta = \sqrt{\max\{\ln(m/r_{\mathrm{A}}), 1\}/\tau}$ *returns a strategy profile* $(\widehat{\sigma}^{\mathrm{D}}, \widehat{\rho}^{\mathrm{A}}) \in \Delta_{\mathrm{D}} \times \Theta_{\mathrm{A}}$ *satisfying*

$$\frac{1}{\alpha}\left(U(\widehat{\sigma}^{\mathrm{D}}, \rho^{\mathrm{A}}) - \varepsilon\right) \leq U(\widehat{\sigma}^{\mathrm{D}}, \widehat{\rho}^{\mathrm{A}}) \leq \alpha U(\sigma^{\mathrm{D}}, \widehat{\rho}^{\mathrm{A}}) + \varepsilon, \quad \forall\, \sigma^{\mathrm{D}} \in \Delta_{\mathrm{D}},\ \forall\, \rho^{\mathrm{A}} \in \Theta_{\mathrm{A}}.$$

Furthermore, $\widehat{\sigma}^{\mathrm{D}}$ *and* $\widehat{\rho}^{\mathrm{A}}$ *respectively satisfy the following bounds for the worst-case expected number of undetected attacks with respect to the value of the game:*

$$\begin{aligned} U(\sigma^{\mathrm{D}^*}, \rho^{\mathrm{A}^*}) &\leq \max_{\rho^{\mathrm{A}} \in \Theta_{\mathrm{A}}} U(\widehat{\sigma}^{\mathrm{D}}, \rho^{\mathrm{A}}) \leq \alpha U(\sigma^{\mathrm{D}^*}, \rho^{\mathrm{A}^*}) + \varepsilon, \\ \frac{1}{\alpha}\left(U(\sigma^{\mathrm{D}^*}, \rho^{\mathrm{A}^*}) - \varepsilon\right) &\leq \min_{\sigma^{\mathrm{D}} \in \Delta_{\mathrm{D}}} U(\sigma^{\mathrm{D}}, \widehat{\rho}^{\mathrm{A}}) \leq U(\sigma^{\mathrm{D}^*}, \rho^{\mathrm{A}^*}). \end{aligned}$$

From Theorem 7, the strategy profile $(\widehat{\sigma}^{\mathrm{D}}, \widehat{\rho}^{\mathrm{A}})$ returned by the MWU algorithm with α-approximate best response (Algorithm 4) offers similar approximation guarantees to those provided by the CG with approximate best response from Theorem 6. However, unlike the latter, in this case, A's marginal attack strategy $\widehat{\rho}^{\mathrm{A}}$ is not necessarily a best response to $\widehat{\sigma}^{\mathrm{D}}$. Instead, A can increase the expected number of undetected attacks by at most a factor of $\alpha + \varepsilon/U(\widehat{\sigma}^{\mathrm{D}}, \widehat{\rho}^{\mathrm{A}})$ through unilateral deviation from $(\widehat{\sigma}^{\mathrm{D}}, \widehat{\rho}^{\mathrm{A}})$. Moreover, the MWU algorithm cannot be executed with $\varepsilon = 0$, as it requires $\Omega(1/\varepsilon^2)$ iterations to achieve convergence.

We next delve into the projection step of Algorithm 4 (Line 6). Given $\widetilde{\rho} \in \mathbb{R}^{\mathcal{E}}_{>0}$, the projection problem consists in finding a vector $\rho \in \Theta_{\mathrm{A}}$ that minimizes the unnormalized relative entropy with respect to $\widetilde{\rho}$, given by $\mathrm{D}_{\mathrm{RE}}\left(\rho \,\|\, \widetilde{\rho}\right)$. This is a strictly convex minimization problem, which renders the projection unique. In the following proposition, we show that this projection problem has a closed-form solution.

Theorem 8. *Let* $\widetilde{\rho} \in \mathbb{R}^{\mathcal{E}}_{>0}$*, and let* $\rho^* \in \mathbb{R}^{\mathcal{E}}$ *denote its projection onto* Θ_{A} *with respect to the unnormalized relative entropy.*

1. *If* $\sum_{e \in \mathcal{E}} \min\{\widetilde{\rho}_e, 1\} \leq r_{\mathrm{A}}$*, then* ρ^* *is given by* $\rho^*_e = \min\{\widetilde{\rho}_e,\, 1\}$ *for every* $e \in \mathcal{E}$*.*

2. *If* $\sum_{e \in \mathcal{E}} \min\{\widetilde{\rho}_e, 1\} > r_{\mathrm{A}}$*, let us sort the elements of* $\mathcal{E}$ *such that* $\widetilde{\rho}_{e_1} \geq \cdots \geq \widetilde{\rho}_{e_m}$*,*

breaking ties arbitrarily, and define the parameter

$$k^* := \max \left\{ k \in \{0, \dots, r_{\text{A}}\} : \, k + \frac{1}{\widetilde{\rho}_{e_k}} \sum_{j=k+1}^{m} \widetilde{\rho}_{e_j} \leq r_{\text{A}} \right\},$$

where we let $\widetilde{\rho}_{e_0} := +\infty$. *Then,* ρ^* *is given by*

$$\rho_e^* = \min \left\{ \frac{r_{\text{A}} - k^*}{\sum_{j=k^*+1}^{m} \widetilde{\rho}_{e_j}} \widetilde{\rho}_e, \, 1 \right\}, \quad \forall \, e \in \mathcal{E}.$$

From Theorem 8, we obtain an analytic characterization for the projection of a vector $\widetilde{\rho} \in \mathbb{R}^{\mathcal{E}}_{>0}$ onto Θ_{A} under the unnormalized relative entropy. In general, this projection is given by a truncated scaling of $\widetilde{\rho}$, namely $\rho_e^* = \min \{\mu \widetilde{\rho}_e, 1\}$ for every $e \in \mathcal{E}$, where the scaling factor μ is determined by two cases. If $\sum_{e \in \mathcal{E}} \min\{\widetilde{\rho}_e, 1\} \leq r_{\text{A}}$, the truncated scaling with $\mu = 1$ is already an optimal solution of the projection problem. Indeed, $\mathrm{D}_{\text{RE}} \left(\cdot \,\|\, \widetilde{\rho} \right)$ is separable, and its e-th summand is minimized in the interval $[0, 1]$ by $\min \{\widetilde{\rho}_e, 1\}$. On the other hand, if $\sum_{e \in \mathcal{E}} \min\{\widetilde{\rho}_e, 1\} > r_{\text{A}}$, the truncated scaling with $\mu = 1$ becomes infeasible for Θ_{A}. Furthermore, KKT conditions are not directly applicable to the projection problem, due to $\mathrm{D}_{\text{RE}} \left(\cdot \,\|\, \widetilde{\rho} \right)$ being nonsmooth. To handle this case, we use a Lagrangian dual approach that helps us determine that the appropriate scaling factor μ is given by the solution of the equation $\sum_{e \in \mathcal{E}} \min \{\mu \widetilde{\rho}_e, 1\} = r_{\text{A}}$. Sorting the elements of $\mathcal{E}$ such that $\widetilde{\rho}_{e_1} \geq \cdots \geq \widetilde{\rho}_{e_m}$ yields $\mu = (r_{\text{A}} - k^*)/ \sum_{j=k^*+1}^{m} \widetilde{\rho}_{e_j}$, where k^* is the integer parameter given in the theorem.

As a direct consequence of Theorem 8, the projection of $\widetilde{\rho}$ onto Θ_{A} can be computed in time $O(m \log m)$, where the running time is dominated by the sorting of the elements of $\mathcal{E}$ required to compute k^*. However, it is possible to compute the projection in linear time. Indeed, the function $s(k) := k + 1/\widetilde{\rho}_{e_k} \sum_{j=k+1}^{m} \widetilde{\rho}_{e_j}$ arising in the definition of the parameter k^* is nondecreasing in k. Furthermore, given $k \in \{0, \dots, m\}$, we can partition the set of components $\mathcal{E}$ into three sets, namely $\mathcal{F}_{\text{high}} := \{e \in \mathcal{E} : \, \widetilde{\rho}_e > \widetilde{\rho}_{e_k}\}$, $\mathcal{F}_{\text{eq}} := \{e \in \mathcal{E} : \, \widetilde{\rho}_e = \widetilde{\rho}_{e_k}\}$ and $\mathcal{F}_{\text{low}} := \{e \in \mathcal{E} : \, \widetilde{\rho}_e < \widetilde{\rho}_{e_k}\}$, so that $s(k) = |\mathcal{F}_{\text{high}}| + |\mathcal{F}_{\text{eq}}| + (1/\widetilde{\rho}_{e_k}) \sum_{e \in \mathcal{F}_{\text{low}}} \widetilde{\rho}_e$. Thus, each evaluation of $s(k)$ is determined by a unique component e_k that induces a splitting of $\mathcal{E}$. This enables us to perform a binary search without sorting $\mathcal{E}$ to find k^* by selecting in each iteration a splitting component e_k—which can be performed in linear time using, for example, the Quickselect algorithm [143]—and testing whether $s(k) \leq r_{\text{A}}$. In the affirmative case, we then continue the search by selecting the splitting component of the next iteration among the elements of $\mathcal{F}_{\text{low}}$; otherwise, we select it among the elements of $\mathcal{F}_{\text{high}}$. We refer to Algorithm 5 for the pseudocode of the projection algorithm.

Algorithm 5: Fast Relative Entropy Projection onto Θ_{A}

Input : A vector $\widetilde{\rho} \in \mathbb{R}^{\mathcal{E}}_{>0}$ and a resource budget $r_{\text{A}} \in \mathbb{Z}_{>0}$.
Output: A vector $\rho^* \in \mathbb{R}^{\mathcal{E}}$ satisfying $\rho^* \in \underset{\rho \in \Theta_{\text{A}}}{\arg\min}\, \mathrm{D}_{\text{RE}}\left(\rho \,\|\, \widetilde{\rho}\right)$.

if $\sum_{e \in \mathcal{E}} \min\{\widetilde{\rho}_e, 1\} \leq r_{\text{A}}$ **then**
  $\mu \leftarrow 1$
else
  Let $\mathcal{F} \leftarrow \mathcal{E},\ k \leftarrow 0,\ s \leftarrow 0$
  while $\mathcal{F} \neq \emptyset$ **do**
    Let $e \leftarrow$ index of $\left\lceil \frac{|\mathcal{F}|}{2} \right\rceil$-th largest entry of $(\widetilde{\rho}_e)_{e \in \mathcal{F}}$
    $\mathcal{F}_{\text{high}} \leftarrow \{e' \in \mathcal{F} :\ \widetilde{\rho}_{e'} > \widetilde{\rho}_e\}$
    $\mathcal{F}_{\text{low}} \leftarrow \{e' \in \mathcal{F} :\ \widetilde{\rho}_{e'} < \widetilde{\rho}_e\}$
    $\mathcal{F}_{\text{eq}} \leftarrow \{e' \in \mathcal{F} :\ \widetilde{\rho}_{e'} = \widetilde{\rho}_e\}$
    if $k + |\mathcal{F}_{\text{high}}| + |\mathcal{F}_{\text{eq}}| + (1/\widetilde{\rho}_e)\left(\sum_{e' \in \mathcal{F}_{\text{low}}} \widetilde{\rho}_{e'} + s\right) \leq r_{\text{A}}$ **then**
      $\mathcal{F} \leftarrow \mathcal{F}_{\text{low}}$
      $k \leftarrow k + |\mathcal{F}_{\text{high}}| + |\mathcal{F}_{\text{eq}}|$
    else
      $\mathcal{F} \leftarrow \mathcal{F}_{\text{high}}$
      $s \leftarrow s + \widetilde{\rho}_e |\mathcal{F}_{\text{eq}}| + \sum_{e' \in \mathcal{F}_{\text{low}}} \widetilde{\rho}_{e'}$
  $\mu \leftarrow (r_{\text{A}} - k)/s$
return $\rho^*_e := \min\{\mu \widetilde{\rho}_e,\ 1\}$ *for all* $e \in \mathcal{E}$

Theorem 9. *Given* $\widetilde{\rho} \in \mathbb{R}^{\mathcal{E}}_{>0}$*, Algorithm 5 computes the projection of* $\widetilde{\rho}$ *onto* Θ_{A} *with respect to the unnormalized relative entropy in time* $O(m)$.

Thanks to Theorem 9, we can compute projections under the unnormalized relative entropy onto Θ_{A} in linear time. Consequently, the computational cost per iteration of Algorithm 4 is determined by the cost of computing a defender's α-approximate best response

(Line 3)—which takes $O(n(n - r_{\mathrm{D}}))$ (resp. $O(nr_{\mathrm{D}})$) evaluations of the undetection function if we use the reverse (resp. forward) greedy algorithm—and the cost of computing the projection of $\widetilde{\rho}^{(t+1)}$ onto Θ_{A} with respect to the unnormalized relative entropy (Line 6), which takes time $O(m)$ using Algorithm 5. Finally, the convergence of Algorithm 4 requires $r_{\mathrm{A}}^2 \max\left\{\ln(m/r_{\mathrm{A}}), 1\right\}/\varepsilon^2$ iterations (Theorem 7), resulting in an overall running time polynomial in m, n, r_{A}, r_{D}, and $1/\varepsilon^2$.

4.6 Computational Study

We test our solution approaches for solving strategic network inspection problems on several natural gas distribution networks across Europe and the United States. The instances from Europe were constructed from the SciGRID_gas database [144], representing the countries of France (EU FR), Germany (EU DE), Italy (EU IT), Poland (EU PL), and Great Britain (EU GB). The instances from United States were obtained from the open datasets of the Homeland Infrastructure Foundation-Level Data [145], and represent the states of Texas (US TX), Oklahoma (US OK), Lousiana (US LA), New York (US NY), and California (US CA).

In this case study, we aim to find inspection strategies for monitoring these networks to detect potential attacks on their pipelines, which can be caused by malicious intervention on their cyber-physical systems or by unauthorized access and trespassing. We consider compressor stations and processing plants from the datasets as the nodes $\mathcal{V}$ for detector locations, and the pipelines as the components $\mathcal{E}$ that the attacker is interested in targeting. To identify disruptions, such as gas leakages indicative of pipe bursts, the defender leverages detectors in the form of unmanned aerial systems equipped with infrared thermography technology to conduct pipeline patrols around each node. For each node, we define its monitoring set as the collection of pipelines intersecting a circular area of radius 50 km around its geographic location, and set its detection probability randomly from a Uniform$[0.5, 1]$ distribution. All LPs and MIPs were solved using Gurobi 9.1.2 [146] on a

11th Gen Intel(R) Core(TM) i7-1195G7 @ 2.90GHz processor.

We begin by evaluating the performance of the exact CG algorithm applied to (LP($\mathcal{D}$)), using (MIP) to compute exact solutions to the pricing problem (DBR), as described in Section 4.4. For each network, we set r_{A} as 2% of the components, and r_{D} as the minimum number of detectors capable of monitoring 80% of the components, except for EU DE and EU FR, which due to their larger sizes, we adjusted r_{D} to cover 70% and 57% of the components, respectively. The results are shown in Table 4.1.

Table 4.1: Computational results for CG on different gas distribution networks

Network	$\lvert\mathcal{V}\rvert$	$\lvert\mathcal{E}\rvert$	r_{D}	r_{A}	$U(\sigma^{\mathrm{D}^*}, \sigma^{\mathrm{A}^*})$	**Time [s]**
EU GB	103	116	13	3	1.68	1.56
US CA	51	616	9	13	9.54	2.77
US NY	47	1,731	7	35	27.07	8.95
US LA	196	2,680	11	54	39.67	209.16
EU PL	256	309	16	7	4.31	465.81
US OK	642	3,237	12	65	46.68	2,759.92
US TX	304	6,762	30	136	106.32	4,826.02
EU IT	302	349	21	7	4.01	9,600.55
EU FR	828	964	25	20	15.16	15,644.50
EU DE	638	866	23	18	12.85	22,568.95

Table 4.1 shows the exact value of the game Γ and the CPU time (in seconds) obtained through CG. Generally, we observe that the running times quickly increase with the size of the network, being the number of nodes more relevant than the number of components. However, additional factors such as the overlapping structure of the monitoring sets and the magnitude of r_{D} relative to the number of nodes may influence performance, as evidenced by the cases of US TX and US OK. Overall, Table 4.1 shows that CG can solve Γ to optimality for real-world networks and a relatively small number of detectors.

Next, we examine the performance of our approximate solution approaches on the same network instances outlined in Table 4.1. Specifically, we evaluate the performance of the CG algorithm to solve (LP($\mathcal{D}$)) using both Forward Greedy (FG) and Reverse Greedy (RG) algorithms to approximate (DBR), as described in Section 4.5.2. We denote these algorithms as (CG-FG) and (CG-RG), respectively. Similarly, we asses the MWU algorithm utilizing both FG and RG algorithms for approximating (DBR), as detailed in Section 4.5.3. We denote these algorithms as (MWU-FG) and (MWU-RG), respectively. We set $\varepsilon = 0.001m$ for all methods. The results are shown in Table 4.2.

Table 4.2: Computational results for approximate solution methods CG-FG, CG-RG, MWU-FG, and MWU-RG for network instances described in Table 4.1

	% Relative Error				Time [s]			
Network	**CG-FG**	**CG-RG**	**MWU-FG**	**MWU-RG**	**CG-FG**	**CG-RG**	**MWU-FG**	**MWU-RG**
EU GB	0.00	1.13	1.07	2.18	0.25	0.22	3.05	8.64
US CA	0.39	1.00	1.23	1.98	0.72	0.97	9.69	15.33
US NY	0.32	0.67	1.07	1.15	3.45	4.31	24.33	38.58
US LA	0.79	1.39	0.85	1.75	32.89	37.28	58.16	331.95
EU PL	0.82	4.14	1.27	5.36	1.25	1.30	12.08	50.97
US OK	0.90	2.31	0.72	3.70	91.27	111.22	137.95	1,833.34
US TX	0.41	0.44	0.84	1.05	508.13	478.30	148.08	358.38
EU IT	1.44	4.91	1.26	6.60	2.58	2.72	18.64	64.39
EU FR	1.02	2.36	1.05	3.82	39.86	34.48	48.56	341.25
EU DE	0.93	3.07	1.16	4.08	23.55	18.81	40.50	265.31

Table 4.2 shows the relative approximation error of the worst-case expected number of undetected attacks with respect to the value of the game Γ and the CPU time (in seconds) of the inspection strategies generated by the four approximation methods. Generally, we observe a substantial reduction in running times for the five instances that took more than $1,000$ seconds to solve under CG—as shown in Table 4.1—with an average reduction of 94.2% achieved by the four approximation methods. Interestingly, although the forward greedy algorithm theoretically provides a weaker approximation guarantee compared to the reverse greedy algorithm, both CG-FG and MWU-FG achieve significantly smaller relative errors, typically less than 1.5%, compared to their counterparts, CG-RG and MWU-RG, which yield solutions within 6.6% of optimality. Furthermore, the running times of CG-FG and MWU-FG are also generally shorter than those of CG-RG and MWU-RG, respectively.

We also compare the performance of both the exact and approximate solution methods on a single network for different values of r_{D} and fixed r_{A}. To this end, we select the gas distribution network of France (EU FR) as our benchmark (see Figure 4.2).

We first analyze the results of our experiments on EU FR for a relatively small number

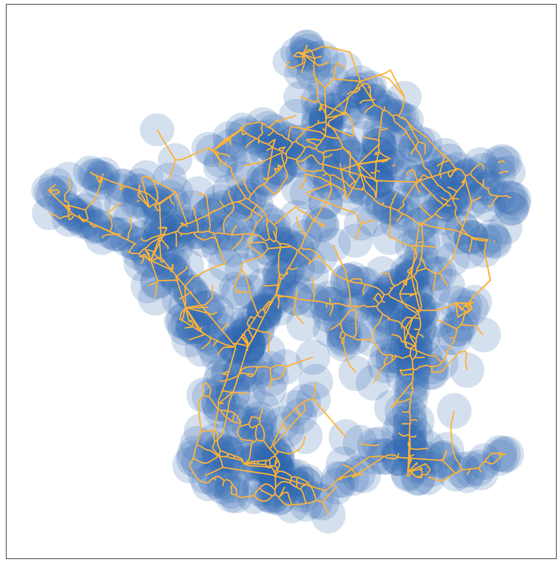

Figure 4.2: Layout of the benchmark natural gas distribution network EU FR, with circular inspection zones of radius 50 km around each detector location.

of detectors, depicted in Figures Figure 4.3a and Figure 4.3b.

Figure 4.3a shows the relative error of the four approximation methods as a function of the number of detectors. The results illustrate that the error of all methods increase with the number of detectors, reflecting the increasing difficulty to effectively coordinating the detectors by the greedy algorithms for approximate best response. In line with the results from Table 4.2, the error of CG-RG and MWU-RG increase at a significantly faster rate with the number of detectors, compared to that of CG-FG and MWU-FG, which consistently remain smaller than 1.1%, being CG-FG the method achieving the smallest error for almost every value of r_{D}.

Figure 4.3b illustrates, using a logarithmic scale, the running times of the exact CG and the four approximation methods. We note that the running time of CG does not scale efficiently with the number of available detectors r_{D}. In fact, CG exhibits running times almost two orders of magnitude larger than those of MWU-FG and MWU-RG, which remain relatively constant. Moreover, MWU-FG outperforms MWU-RG by one order of magnitude. As anticipated in Section 4.5.1, this difference arises due to the small size of r_{D} relative to n, resulting in fewer iterations being required by forward greedy compared

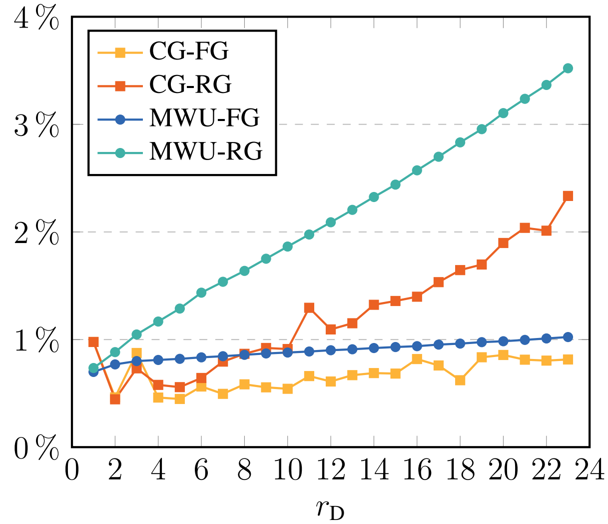

(a) Relative error with respect to value of the game ($r_{\mathrm{A}} = 20$).

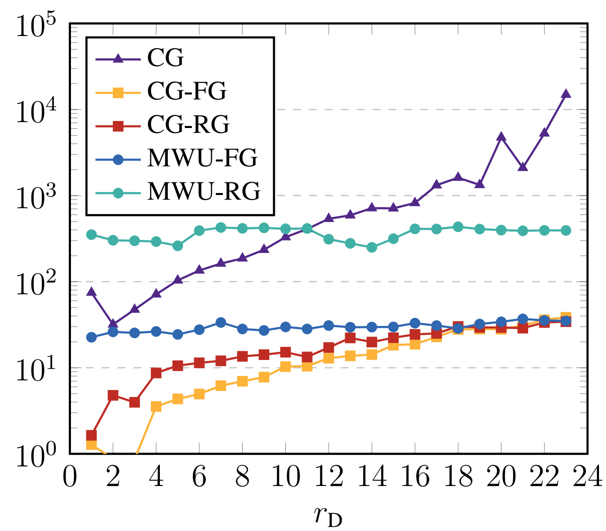

(b) Running time (in seconds, $r_{\mathrm{A}} = 20$).

Figure 4.3: Computational results for small number of detectors on benchmark network EU FR

to reverse greedy. Nonetheless, both CG-FG and CG-RG exhibit faster, albeit increasing, running times, eventually reaching levels comparable to MWU-FG for the largest values of r_{D} in this range. This motivates us to further examine the performance of these methods for larger numbers of detectors, as depicted in Figures Figure 4.4a and Figure 4.4b.

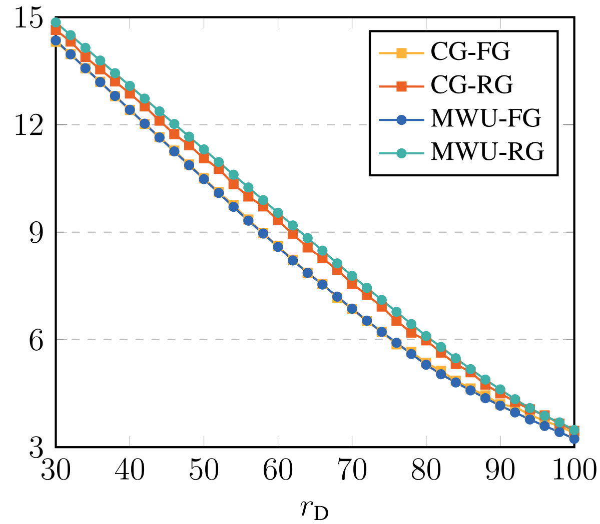

(a) Worst-case expected number of undetected attacks ($r_{\mathrm{A}} = 20$).

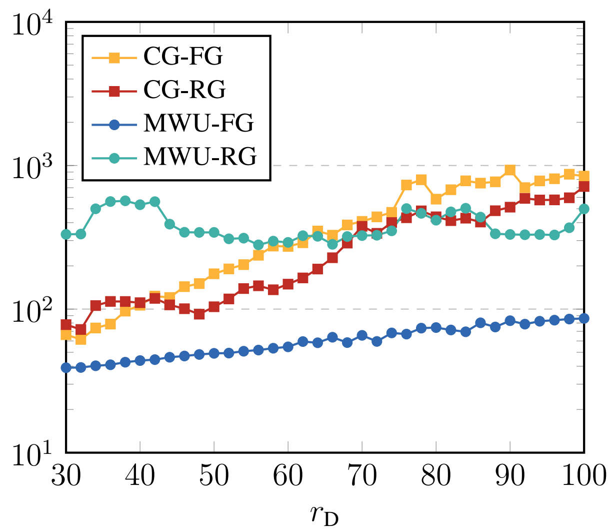

(b) Running time (in seconds, $r_{\mathrm{A}} = 20$).

Figure 4.4: Computational results for large number of detectors on benchmark network EU FR

For larger values of r_{D}, exact solutions of the game Γ become computationally prohibitive, preventing us from plotting the approximation error. Hence, Figure 4.4a illustrates the worst-case expected number of undetected attacks for the inspection strategies generated by all four approximation methods. Consistent with the trends observed in Figure 4.3a, both CG-FG and MWU-FG outperform CG-RG and MWU-RG. Moreover, in Figure 4.4b, while the running times of MWU-FG and MWU-RG remain within similar orders of magnitude as those shown in Figure 4.3b, we observe a notable increase in the running times of both CG-FG and CG-RG with the number of detectors. In this scenario, MWU-FG significantly outperforms the other three methods by up to one order of magnitude for the largest values of r_{D}. Overall, MWU-FG emerges as the preferred method for a large number of detectors.

Finally, we refer to the support sizes of the inspection strategies generated by our solution approaches. In all experiments conducted on EU FR, CG returned mixed inspection strategies that randomized over 91 distinct detector positionings, whereas CG-FG and CG-RG generated inspection strategies that randomized over between 123 and 146 detector positionings. This stands in high contrast to the inspection strategies generated by MWU-FG and MWU-RG, which consistently exhibited support sizes approximately equal to the number of iterations of the MWU algorithms, around 6,000. Smaller support sizes entail less intricate inspection schedules, but they can pose computational challenges for larger networks. Security decision-makers should be mindful of this tradeoff.

4.7 Concluding Remarks

Motivated by real-world applications in safeguarding critical infrastructures, we investigated a large-scale zero-sum network inspection game where a defender strategically allocates multiple detectors in locations with imperfect and heterogeneous detection capabilities, while an attacker targets multiple network components. The defender (resp. attacker) aims to minimize (resp. maximize) the expected number of undetected attacks.

We leveraged the inherent structure of the game, which is supermodular for the defender and additive for the attacker, to develop both exact and approximate solution methods utilizing Column Generation (CG) and Multiplicative Weights Update (MWU) algorithms. These approaches employed efficient algorithms as subroutines to compute either exact or approximate defender's best responses. Specifically, we achieved exact NE using CG with an exact MIP formulation for the defender's best response problem. For approximate NE, we utilized both CG and MWU in conjunction with greedy algorithms for approximate defender's best responses. Across all methods, we represented the attacker's strategies in terms of unidimensional marginal probabilities. In the case of the MWU algorithm, this required addressing an additional relative entropy projection problem within each iteration, for which we derived a closed-form solution that can be computed in linear time.

Our computational study on real-world gas distribution networks revealed that our exact solution method can efficiently compute equilibrium inspection strategies when the defender has access to a relatively small number of detectors. However, for a larger number of detectors, our approximation methods provide the only scalable solution approaches. Among these, the MWU algorithm exhibits better performance than a CG algorithm when both utilize the forward greedy algorithm for computing approximate defender's best responses. This enhanced performance comes at the expense of generating inspection strategies with larger support sizes.

Our solution approaches can be applied to other combinatorial zero-sum games featuring submodular or supermodular structure for one player and additive structure for the adversary. Examples of such settings include the network inspection model with heterogeneous component criticalities proposed by [17], or the flow interdiction model introduced by [121].

We propose the following directions for future work. First, exploring approximate NE for structured instances, such as monitoring sets forming minimum set covers or maximum set packings, could enhance our comprehension of strategic interacions in scenarios with complex network structures, and potentially lead to the development of more efficient algorithms. Second, we advocate for investigating nonzero-sum variants of our game. Unfortunately, the minmax formulation underlying our CG algorithm, and the MWU algorithm do not necessarily converge to NE in nonzero-sum games. Hence, alternative techniques capable of providing satisfactory approximations in these scenarios need to be developed. Finally, another interesting research direction involves analyzing a variant of the game under incomplete information, where players face uncertainty regarding their adversary's number of resources or the detection probabilities, which can be learned through repeated interactions.

Appendices

APPENDIX A

SUPPLEMENT TO CHAPTER 2

A.1 Examples of Equilibrium Regime Patterns

Example 1. *Consider the hide-and-seek model represented in Figure A.1 and assume that* $r_{\mathrm{S}} = 5$ *and* $r_{\mathrm{H}} = 3$*. In this case,* $0 = \tau_{-1} < r_{\mathrm{H}} \leq \nu_0 = 8$ *and* $i^* = 0$*. Furthermore,*

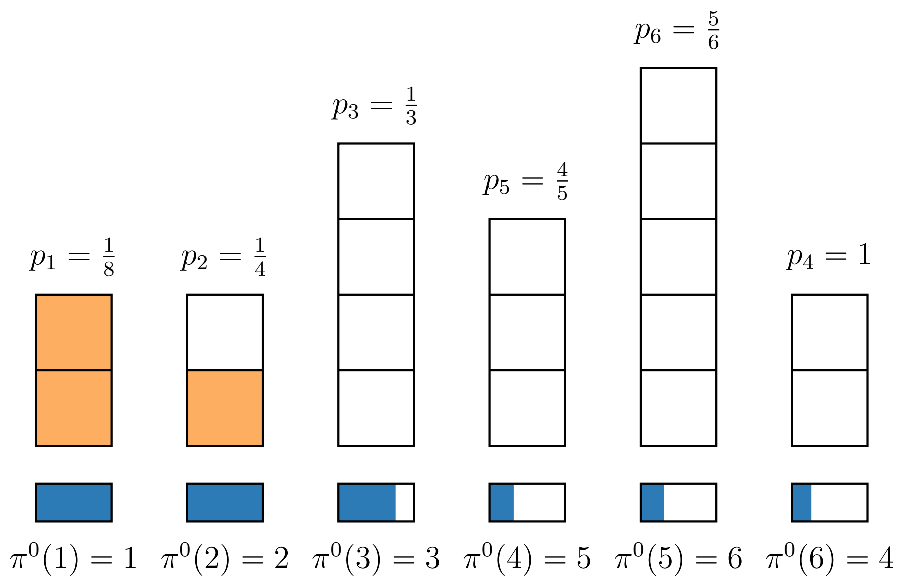

Figure A.1: Equilibrium Regime Pattern 1 when $r_{\mathrm{S}} = 5$ and $r_{\mathrm{H}} = 3$. The hiding capacity of each location is represented by the corresponding number of squares. Marginal inspection probabilities (resp. expected numbers of hidden items) in equilibrium are represented by the blue (resp. orange) colors.

$1 = \ell_0 < k_0 = 3$ *and* $\pi^0(1) = 1$*,* $\pi^0(2) = 2$*,* $\pi^0(3) = 3$*,* $\pi^0(4) = 5$*,* $\pi^0(5) = 6$*,* $\pi^0(6) = 4$*. From Regime Pattern 1, we have* $\mathcal{J} = \{1\}$*,* $\mathcal{K} = \{2, 3, 4, 5, 6\}$*, and an equilibrium strategy* *profile* $(\rho^{\mathrm{S}^*}, \rho^{\mathrm{H}^*}) \in \widetilde{\mathcal{A}}_{\mathrm{S}} \times \widetilde{\mathcal{A}}_{\mathrm{H}}$ *of the game* $\widetilde{\Gamma}$ *is given by*

$$\rho_i^{\mathrm{S}^*} = \begin{cases} 1 & \textit{if } i = 1, \\ 1 & \textit{if } i = 2, \\ 3/4 & \textit{if } i = 3, \\ 1/4 & \textit{if } i = 4, \\ 5/16 & \textit{if } i = 5, \\ 3/10 & \textit{if } i = 6, \end{cases} \qquad \rho_i^{\mathrm{H}^*} = \begin{cases} 2 & \textit{if } i = 1, \\ 1 & \textit{if } i = 2, \\ 0 & \textit{if } i = 3, \\ 0 & \textit{if } i = 4, \\ 0 & \textit{if } i = 5, \\ 0 & \textit{if } i = 6. \end{cases}$$

An equilibrium inspection strategy of the game Γ *is given by* $\sigma^{\mathrm{S}^*} \in \Delta_{\mathrm{S}}$ *such that* $\sigma^{\mathrm{S}^*}_{(1,1,0,0,0,0)} = 11/80$, $\sigma^{\mathrm{S}^*}_{(1,1,0,0,0,1)} = 9/80$, $\sigma^{\mathrm{S}^*}_{(1,1,1,1,0,0)} = 1/4$, $\sigma^{\mathrm{S}^*}_{(1,1,1,0,1,0)} = 5/16$, $\sigma^{\mathrm{S}^*}_{(1,1,1,0,0,1)} = 3/16$ *and* $\sigma^{\mathrm{S}^*}_x = 0$ *otherwise. An equilibrium hiding strategy of the game* Γ *is given by* $\sigma^{\mathrm{H}^*} \in \Delta_{\mathrm{H}}$ *such that* $\sigma^{\mathrm{H}^*}_{(2,1,0,0,0,0)} = 1$ *and* $\sigma^{\mathrm{H}^*}_y = 0$ *otherwise. We note that S can implement her equilibrium strategy by randomizing over resource allocations of up to* $4 < r_{\mathrm{S}}$ *inspection resources. Furthermore, H's equilibrium strategy is pure. The value of the games* Γ *and* $\widetilde{\Gamma}$ *is* $2 + 1/2$.
$\triangle$

Example 2. *Consider the hide-and-seek model represented in Figure A.2 and assume that* $r_{\mathrm{S}} = 3$ *and* $r_{\mathrm{H}} = 7$. *In this case,* $449/80 = \nu_0 < r_{\mathrm{H}} \leq \tau_1 = 289/40$ *and* $i^* = 1$. *Furthermore,* $k_1 = \ell_1 = 1$ *and* $\pi^1(1) = 2$, $\pi^1(2) = 3$, $\pi^1(3) = 5$, $\pi^1(4) = 6$, $\pi^1(5) = 4$. *From Regime Pattern 2, we have* $\mathcal{I} = \{1\}$, $\mathcal{J} = \{2\}$, $\mathcal{K} = \{3, 4, 5, 6\}$, *and an equilibrium*

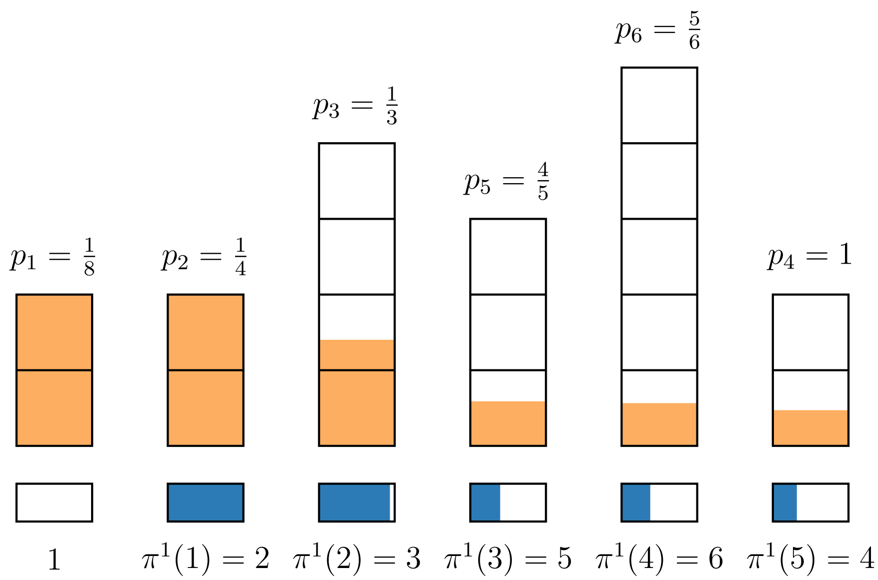

Figure A.2: Equilibrium Regime Pattern 2 when $r_{\mathrm{S}} = 3$ and $r_{\mathrm{H}} = 6$. The hiding capacity of each location is represented by the corresponding number of squares. Marginal inspection probabilities (resp. expected numbers of hidden items) in equilibrium are represented by the blue (resp. orange) colors.

strategy profile $(\rho^{\mathrm{S}^*}, \rho^{\mathrm{H}^*}) \in \widetilde{\mathcal{A}}_{\mathrm{S}} \times \widetilde{\mathcal{A}}_{\mathrm{H}}$ *of the game* $\widetilde{\Gamma}$ *is given by*

$$\rho_i^{\mathrm{S}^*} = \begin{cases} 0 & \textit{if } i = 1, \\ 1 & \textit{if } i = 2, \\ 40/43 & \textit{if } i = 3, \\ 40/129 & \textit{if } i = 4, \\ 50/129 & \textit{if } i = 5, \\ 16/43 & \textit{if } i = 6, \end{cases} \qquad \rho_i^{\mathrm{H}^*} = \begin{cases} 2 & \textit{if } i = 1, \\ 2 & \textit{if } i = 2, \\ 60/43 & \textit{if } i = 3, \\ 20/43 & \textit{if } i = 4, \\ 25/43 & \textit{if } i = 5, \\ 24/43 & \textit{if } i = 6. \end{cases}$$

An equilibrium inspection strategy of the game Γ *is given by* $\sigma^{\mathrm{S}^*} \in \Delta_{\mathrm{S}}$ *such that* $\sigma^{\mathrm{S}^*}_{(0,1,1,1,0,0)} = 31/129$*,* $\sigma^{\mathrm{S}^*}_{(0,1,1,0,1,0)} = 50/129$*,* $\sigma^{\mathrm{S}^*}_{(0,1,1,0,0,1)} = 13/43$*,* $\sigma^{\mathrm{S}^*}_{(0,1,0,1,0,1)} = 3/43$*, and* $\sigma^{\mathrm{S}^*}_x = 0$ *otherwise. An equilibrium hiding strategy of the game* Γ *is given by* $\sigma^{\mathrm{H}^*} \in \Delta_{\mathrm{H}}$ *such that* $\sigma^{\mathrm{H}^*}_{(2,2,0,2,1,0)} = 1/43$*,* $\sigma^{\mathrm{H}^*}_{(2,2,2,1,0,0)} = 18/43$*,* $\sigma^{\mathrm{H}^*}_{(2,2,1,0,1,1)} = 24/43$ *and* $\sigma^{\mathrm{H}^*}_y = 0$ *otherwise. The value of the games* Γ *and* $\widetilde{\Gamma}$ *is* $5 + 49/86$*.* $\triangle$

Example 3. *Consider the hide-and-seek model represented in Figure A.3 and assume that $r_{\mathrm{S}} = 4$ and $r_{\mathrm{H}} = 10$. In this case, $389/40 = \tau_1 < r_{\mathrm{H}} \leq \nu_2 = 33/2$ and $i^* = 2$.*

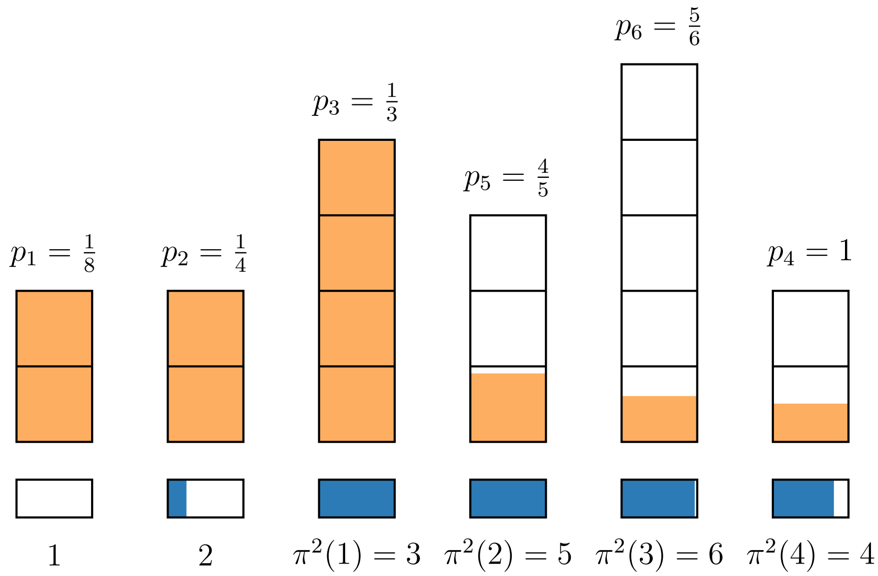

Figure A.3: Equilibrium Regime Pattern 3 when $r_{\mathrm{S}} = 4$ and $r_{\mathrm{H}} = 10$. The hiding capacity of each location is represented by the corresponding number of squares. Marginal inspection probabilities (resp. expected numbers of hidden items) in equilibrium are represented by the blue (resp. orange) colors.

Furthermore, $k_2 = 3 > 1 = \ell_2$ and $\pi^2(1) = 3$, $\pi^2(2) = 5$, $\pi^2(3) = 6$, $\pi^2(4) = 4$. From Regime Pattern 3, we have $\mathcal{I} = \{1,2\}$, $\mathcal{J} = \{3\}$, $\mathcal{K} = \{4,5,6\}$, and an equilibrium strategy profile $(\rho^{\mathrm{S}^*}, \rho^{\mathrm{H}^*}) \in \widetilde{\mathcal{A}}_{\mathrm{S}} \times \widetilde{\mathcal{A}}_{\mathrm{H}}$ of the game $\widetilde{\Gamma}$ is given by

$$\rho_i^{\mathrm{S}^*} = \begin{cases} 0 & \text{if } i = 1, \\ 6/25 & \text{if } i = 2, \\ 1 & \text{if } i = 3, \\ 4/5 & \text{if } i = 4, \\ 1 & \text{if } i = 5, \\ 24/25 & \text{if } i = 6, \end{cases} \qquad \rho_i^{\mathrm{H}^*} = \begin{cases} 2 & \text{if } i = 1, \\ 2 & \text{if } i = 2, \\ 4 & \text{if } i = 3, \\ 1/2 & \text{if } i = 4, \\ 9/10 & \text{if } i = 5, \\ 3/5 & \text{if } i = 6. \end{cases}$$

An equilibrium inspection strategy of the game Γ is given by $\sigma^{\mathrm{S}^*} \in \Delta_{\mathrm{S}}$ such that $\sigma^{\mathrm{S}^*}_{(0,1,1,1,1,0)} =$

$1/25$, $\sigma^{\mathrm{S}^*}_{(0,1,1,0,1,1)} = 1/5$, $\sigma^{\mathrm{S}^*}_{(0,1,1,0,0,1)} = 13/43$, $\sigma^{\mathrm{S}^*}_{(0,0,1,1,1,1)} = 19/25$, *and* $\sigma^{\mathrm{S}^*}_x = 0$ *otherwise. An equilibrium hiding strategy of the game* Γ *is given by* $\sigma^{\mathrm{H}^*} \in \Delta_{\mathrm{H}}$ *such that* $\sigma^{\mathrm{H}^*}_{(2,2,4,1,0,1)} = 1/2$, $\sigma^{\mathrm{H}^*}_{(2,2,4,0,1,1)} = 1/10$, $\sigma^{\mathrm{H}^*}_{(2,2,4,0,2,0)} = 2/5$ *and* $\sigma^{\mathrm{H}^*}_y = 0$ *otherwise. The value of the games* Γ *and* $\widetilde{\Gamma}$ *is* $6 + 71/75$. $\triangle$

A.2 Proofs of Statements

A.2.1 Proofs of Section 2.3

Proof of Lemma 1. Let $b \in \mathbb{Z}^n_{>0}$ be a capacity vector and $r \in \mathbb{Z}_{>0}$ be a resource budget. First, we observe that $\widetilde{\mathcal{A}}(b, r)$ is the convex hull of $\mathcal{A}(b, r)$. Indeed, since the row vector $\mathbf{1}_n^\top$ is totally unimodular, then (2.1) is an ideal formulation of $\mathcal{A}(b, r)$ and the polytope $\widetilde{\mathcal{A}}(b, r)$ has integral extreme points, which we denote $z^1, \ldots, z^I$. Consider a vector $\rho' \in \mathbb{R}^n$. If $\rho' \in \widetilde{\mathcal{A}}(b, r)$, then there exist $\lambda_1, \ldots, \lambda_I \in [0, 1]$ such that $\rho' = \sum_{i=1}^I \lambda_i z^i$ and $\sum_{i=1}^I \lambda_i = 1$. Then, $\sigma \in [0, 1]^{\mathcal{A}(b,r)}$ defined by $\sigma_{z^i} = \lambda_i$ for $i \in [\![1, I]\!]$ and $\sigma_z = 0$ if $z \notin \{z^1, \ldots, z^I\}$ is a probability distribution in $\Delta(b, r)$ and satisfies $\rho_i(\sigma) = \rho'_i$ for all $i \in [\![1, n]\!]$. Conversely, if there exists $\sigma \in \Delta(b, r)$ that satisfies $\rho'_i = \rho_i(\sigma) = \sum_{z \in \mathcal{A}(b,r)} z_i \sigma_z$ for all $i \in [\![1, n]\!]$, then ρ' is a convex combination of elements on $\mathcal{A}(b, r)$, and belongs to its convex hull, $\widetilde{\mathcal{A}}(b, r)$. $\square$

Proof of Proposition 1. We prove each statement.

– Consider a strategy profile $(\sigma^{\mathrm{S}}, \sigma^{\mathrm{H}}) \in \Delta_{\mathrm{S}} \times \Delta_{\mathrm{H}}$. Then,

$$U(\sigma^{\mathrm{S}}, \sigma^{\mathrm{H}}) \overset{(2.2)}{=} \sum_{i=1}^{n} \left(1 - p_i \sum_{x \in \mathcal{A}_{\mathrm{S}}} x_i \sigma^{\mathrm{S}}_x\right) \sum_{y \in \mathcal{A}_{\mathrm{H}}} y_i \sigma^{\mathrm{H}}_y \overset{(2.3)}{=} \tilde{u}(\rho(\sigma^{\mathrm{S}}), \rho(\sigma^{\mathrm{H}})). \tag{A.1}$$

– Let $(\sigma^{\mathrm{S}^*}, \sigma^{\mathrm{H}^*}) \in \Delta_{\mathrm{S}} \times \Delta_{\mathrm{H}}$ be an equilibrium of Γ. We must show that $(\rho(\sigma^{\mathrm{S}^*}), \rho(\sigma^{\mathrm{H}^*}))$ is a pure equilibrium of $\widetilde{\Gamma}$. First, from Lemma 1, we know that $(\rho(\sigma^{\mathrm{S}^*}), \rho(\sigma^{\mathrm{H}^*})) \in \widetilde{\mathcal{A}}_{\mathrm{S}} \times \widetilde{\mathcal{A}}_{\mathrm{H}}$. Then, given $(\rho^{\mathrm{S}}, \rho^{\mathrm{H}}) \in \widetilde{\mathcal{A}}_{\mathrm{S}} \times \widetilde{\mathcal{A}}_{\mathrm{H}}$, let $(\sigma^{\mathrm{S}}, \sigma^{\mathrm{H}}) \in \Delta_{\mathrm{S}} \times \Delta_{\mathrm{H}}$ satisfying $\rho(\sigma^{\mathrm{S}}) = \rho^{\mathrm{S}}$ and $\rho(\sigma^{\mathrm{H}}) = \rho^{\mathrm{H}}$ (Lemma 1). Using equilibrium conditions in Γ, we obtain:

$$\tilde{u}(\rho(\sigma^{S^*}), \rho(\sigma^{H^*})) \overset{(A.1)}{=} U(\sigma^{S^*}, \sigma^{H^*}) \leq U(\sigma^{S}, \sigma^{H^*}) \overset{(A.1)}{=} \tilde{u}(\rho^{S}, \rho(\sigma^{H^*})),$$
$$\text{and } \tilde{u}(\rho(\sigma^{S^*}), \rho(\sigma^{H^*})) \overset{(A.1)}{=} U(\sigma^{S^*}, \sigma^{H^*}) \geq U(\sigma^{S^*}, \sigma^{H}) \overset{(A.1)}{=} \tilde{u}(\rho(\sigma^{S^*}), \rho^{H}).$$

Thus, $(\rho(\sigma^{S^*}), \rho(\sigma^{H^*}))$ is a pure equilibrium of $\widetilde{\Gamma}$.

- Let $(\rho^{S^*}, \rho^{H^*}) \in \widetilde{\mathcal{A}}_S \times \widetilde{\mathcal{A}}_H$ be a pure equilibrium of $\widetilde{\Gamma}$ and let $(\sigma^{S^*}, \sigma^{H^*}) \in \Delta_S \times \Delta_H$ be a strategy profile that satisfies $\rho(\sigma^{S^*}) = \rho^{S^*}$ and $\rho(\sigma^{H^*}) = \rho^{H^*}$ (Lemma 1). Using equilibrium conditions in $\widetilde{\Gamma}$, we obtain:

$$\forall\, \sigma^{S} \in \Delta_S,\ U(\sigma^{S^*}, \sigma^{H^*}) \overset{(A.1)}{=} \tilde{u}(\rho^{S^*}, \rho^{H^*}) \leq \tilde{u}(\rho(\sigma^{S}), \rho^{H^*}) \overset{(A.1)}{=} U(\sigma^{S}, \sigma^{H^*}),$$
$$\forall\, \sigma^{H} \in \Delta_H,\ U(\sigma^{S^*}, \sigma^{H^*}) \overset{(A.1)}{=} \tilde{u}(\rho^{S^*}, \rho^{H^*}) \geq \tilde{u}(\rho^{S^*}, \rho(\sigma^{H})) \overset{(A.1)}{=} U(\sigma^{S^*}, \sigma^{H}).$$

Thus, $(\sigma^{S^*}, \sigma^{H^*})$ is an equilibrium of Γ.

- Given an equilibrium $(\sigma^{S^*}, \sigma^{H^*}) \in \Delta_S \times \Delta_H$ of Γ, we know that $(\rho(\sigma^{S^*}), \rho(\sigma^{H^*}))$ is a pure equilibrium of $\widetilde{\Gamma}$ and $U(\sigma^{S^*}, \sigma^{H^*}) \overset{(A.1)}{=} \tilde{u}(\rho(\sigma^{S^*}), \rho(\sigma^{H^*}))$. Therefore, the values of the games Γ and $\widetilde{\Gamma}$ are identical.

□

We note that when some detection probabilities are identical, permutations π^i ordering $[\![i+1, n]\!]$ by their detection probabilities may not be unique. To simplify our proofs, we assume without loss of generality that π^i maintains the order between identical detection probabilities, i.e., $\pi^i(j) < \pi^i(k)$ when $1 \leq j < k \leq n - i$ and $p_{\pi^i(j)} = p_{\pi^i(k)}$, thus rendering π^i unique for every $i \in [\![0, n-1]\!]$.

Before proving Lemma 2 we need the following auxiliary lemmas.

Lemma 12. *Let $i \in \llbracket 1, n-1 \rrbracket$ and $j^* \in \llbracket 1, n-i+1 \rrbracket$ be such that $\pi^{i-1}(j^*) = i$. Then:*

$$\pi^{i-1}(j) = \begin{cases} \pi^i(j) & \text{if } j \in \llbracket 1, j^*-1 \rrbracket, \\ i & \text{if } j = j^*, \\ \pi^i(j-1) & \text{if } j \in \llbracket j^*+1, n-i+1 \rrbracket. \end{cases} \tag{A.2}$$

Moreover, for every $k \in \llbracket 0, n-i+1 \rrbracket$,

$$\sum_{j=1}^{k} c_{\pi^{i-1}(j)} = \begin{cases} \sum_{j=1}^{k} c_{\pi^i(j)} & \text{if } k \in \llbracket 0, j^*-1 \rrbracket, \\ \sum_{j=1}^{k-1} c_{\pi^i(j)} + c_i & \text{if } k \in \llbracket j^*, n-i+1 \rrbracket, \end{cases} \tag{A.3}$$

and

$$S_{k+1}^{i-1} = \begin{cases} S_{k+1}^{i} + \frac{1}{p_i} & \text{if } k \in \llbracket 0, j^*-1 \rrbracket, \\ S_k^i & \text{if } k \in \llbracket j^*, n-i+1 \rrbracket. \end{cases} \tag{A.4}$$

Proof of Lemma 12. Let $i \in \llbracket 1, n-1 \rrbracket$ and $j^* \in \llbracket 1, n-i+1 \rrbracket$ be such that $\pi^{i-1}(j^*) = i$. The permutation π^{i-1} sorts locations $\llbracket i, n \rrbracket$ in order of non-decreasing detection probabilities. After removing $p_i = p_{\pi^{i-1}(j^*)}$ from the chain of inequalities, we obtain $p_{\pi^{i-1}(1)} \leq \cdots \leq p_{\pi^{i-1}(j^*-1)} \leq p_{\pi^{i-1}(j^*+1)} \leq \cdots \leq p_{\pi^{i-1}(n-i+1)}$, which sorts $\llbracket i+1, n \rrbracket$ by detection probabilities, thus providing (A.2).

As a consequence, for every $k \in [\![0, j^* - 1]\!]$,

$$\begin{aligned}
\sum_{j=1}^{k} c_{\pi^{i-1}(j)} &= \sum_{j=1}^{k} c_{\pi^{i}(j)}, \\
S_{k+1}^{i-1} &= \sum_{j=k+1}^{j^*-1} \frac{1}{p_{\pi^{i-1}(j)}} + \frac{1}{p_i} + \sum_{j=j^*+1}^{n-i+1} \frac{1}{p_{\pi^{i-1}(j)}} \\
&= \frac{1}{p_i} + \sum_{j=k+1}^{j^*-1} \frac{1}{p_{\pi^{i}(j)}} + \sum_{j=j^*+1}^{n-i+1} \frac{1}{p_{\pi^{i}(j-1)}} = \frac{1}{p_i} + S_{k+1}^{i}.
\end{aligned}$$

Similarly, for every $k \in [\![j^*, n - i + 1]\!]$,

$$\begin{aligned}
\sum_{j=1}^{k} c_{\pi^{i-1}(j)} &= \sum_{j=1}^{j^*-1} c_{\pi^{i}(j)} + c_i + \sum_{j=j^*+1}^{k} c_{\pi^{i}(j-1)} = \sum_{j=1}^{k-1} c_{\pi^{i}(j)} + c_i, \\
S_{k+1}^{i-1} &= \sum_{j=k+1}^{n-i+1} \frac{1}{p_{\pi^{i-1}(j)}} = \sum_{j=k+1}^{n-i+1} \frac{1}{p_{\pi^{i}(j-1)}} = S_k^i.
\end{aligned}$$

□

Lemma 13. *For every $i \in [\![0, n-1]\!]$, k_i exists and*

$$K_i \coloneqq \left\{ k \in [\![0, n-i]\!] : \, k + p_{\pi^i(k)} S_{k+1}^i < r_{\mathrm{S}} \right\} = [\![0, k_i]\!].$$

Furthermore, for every $i \in [\![1, n-1]\!]$ and $j^* \in [\![1, n-i+1]\!]$ such that $\pi^{i-1}(j^*) = i$,

$$k_{i-1} \leq \begin{cases} k_i & \text{if } k_{i-1} \in [\![0, j^* - 1]\!], \\ k_i + 1 & \text{if } k_{i-1} \in [\![j^*, n-i+1]\!]. \end{cases} \tag{A.5}$$

Proof of Lemma 13. Let $i \in [\![0, n-1]\!]$. Since $p_{\pi^i(0)} S_1^i = 0 < r_{\mathrm{S}}$, then $0 \in K_i$ and k_i exists. Next, we consider $k \in [\![1, n-i]\!]$. If $k \in K_i$, then $k - 1 \in K_i$, as shown below:

$$r_{\mathrm{S}} > k + p_{\pi^i(k)} S_{k+1}^i = k - 1 + p_{\pi^i(k)} \left(S_{k+1}^i + \frac{1}{p_{\pi^i(k)}} \right) \geq k - 1 + p_{\pi^i(k-1)} S_k^i,$$

where we used the fact that $k \in K_i$ and π^i orders locations in $[\![i+1, n]\!]$ by their detection probabilities. Therefore, $K_i = [\![0, k_i]\!]$.

We next analyze k_i as a function of i. Let $i \in [\![1, n-1]\!]$ and $j^* \in [\![1, n-i+1]\!]$ be such that $\pi^{i-1}(j^*) = i$. If $k_{i-1} \in [\![0, j^*-1]\!]$, we obtain:

$$r_{\mathrm{S}} > k_{i-1} + p_{\pi^{i-1}(k_{i-1})} S^{i-1}_{k_{i-1}+1} \overset{\text{(A.2),(A.4)}}{\geq} k_{i-1} + p_{\pi^{i}(k_{i-1})} S^{i}_{k_{i-1}+1}.$$

Since $k_{i-1} \leq j^* - 1 \leq n - i$, we deduce that $k_{i-1} \in K_i$, and $k_{i-1} \leq k_i$.

If $k_{i-1} \in [\![j^*+1, n-i+1]\!]$, we obtain:

$$r_{\mathrm{S}} \overset{\text{(A.2),(A.4)}}{>} k_{i-1} + p_{\pi^{i}(k_{i-1}-1)} S^{i}_{k_{i-1}} > k_{i-1} - 1 + p_{\pi^{i}(k_{i-1}-1)} S^{i}_{k_{i-1}-1+1}.$$

Since $k_{i-1} - 1 \leq n - i$, then $k_{i-1} - 1 \in K_i$ and $k_{i-1} - 1 \leq k_i$.

Finally, if $k_{i-1} = j^*$, we obtain:

$$r_{\mathrm{S}} \overset{\text{(A.4)}}{>} k_{i-1} + p_{\pi^{i-1}(k_{i-1}-1)} S^{i}_{k_{i-1}} \overset{\text{(A.2)}}{>} k_{i-1} - 1 + p_{\pi^{i}(k_{i-1}-1)} S^{i}_{k_{i-1}-1+1}.$$

Thus, $k_{i-1} - 1 \in K_i$ and $k_{i-1} - 1 \leq k_i$. Note that we used throughout that $p_{\pi^i(0)} = 0$. $\square$

Lemma 14. *For every $i \in [\![0, n-1]\!]$, if $r_{\mathrm{H}} > \sum_{j=1}^{i} c_j + p_i c_i S^i_1$, then ℓ_i exists and*

$$L_i \coloneqq \left\{ \ell \in [\![0, n-i]\!] : \sum_{j=1}^{i} c_j + \sum_{j=1}^{\ell} c_{\pi^i(j)} + p_i c_i S^i_{\ell+1} < r_{\mathrm{H}} \right\} = [\![0, \ell_i]\!].$$

Proof of Lemma 14. Let $i \in [\![0, n-1]\!]$ and suppose that $r_{\mathrm{H}} > \sum_{j=1}^{i} c_j + p_i c_i S^i_1$. Then, $0 \in L_i$ and ℓ_i exists.

Next, consider $\ell \in [\![1, n-i]\!]$. If $\ell \in L_i$, then $\ell - 1 \in L_i$, as shown below:

$$r_{\mathrm{H}} - \sum_{j=1}^{i} c_j > \sum_{j=1}^{\ell-1} c_{\pi^i(j)} + \frac{p_{\pi^i(\ell)}}{p_{\pi^i(\ell)}} c_{\pi^i(\ell)} + p_i c_i S^i_{\ell+1} \geq \sum_{j=1}^{\ell-1} c_{\pi^i(j)} + p_i c_i S^i_{\ell},$$

where we used the fact that $\ell \in L_i$ and locations in $[\![1, n]\!]$ are ordered by their detection potentials. Therefore, $L_i = [\![0, \ell_i]\!]$. $\square$

We are now ready to prove Lemma 2.

Proof of Lemma 2. First, $\tau_{-1} = 0$ by definition. Moreover, we observe that $k_{n-1} \in \{0, 1\}$, which implies that $\tau_{n-1} = m$: If $k_{n-1} = 0$, then $\tau_{n-1} = \sum_{j=1}^{n-1} c_j + p_n c_n / p_n = m$. If $k_{n-1} = 1$, then $\tau_{n-1} = \sum_{j=1}^{n-1} c_j + c_n = m$. We next show that $\tau_{i-1} \leq \nu_i \leq \tau_i$ for all $i \in [\![0, n-1]\!]$.

We note that the inequality $\nu_i \leq \tau_i$ follows directly from the fact that $p_i c_i \leq p_{i+1} c_{i+1}$. Thus, it only remains to show that $\tau_{i-1} \leq \nu_i$. This is trivial for $i = 0$, so we may assume that $i \in [\![1, n-1]\!]$. We note that

$$\nu_i - \tau_{i-1} = c_i + \sum_{j=1}^{k_i} c_{\pi^i(j)} - \sum_{j=1}^{k_{i-1}} c_{\pi^{i-1}(j)} - p_i c_i \left(S^{i-1}_{k_{i-1}+1} - S^i_{k_i+1} \right). \tag{A.6}$$

Let $j^* \in [\![1, n-i+1]\!]$ be such that $\pi^{i-1}(j^*) = i$. If $k_{i-1} \in [\![0, j^*-1]\!]$, we obtain:

$$\begin{aligned}
\nu_i - \tau_{i-1} &\overset{(A.3),(A.4)}{=} c_i + \sum_{j=1}^{k_i} c_{\pi^i(j)} - \sum_{j=1}^{k_{i-1}} c_{\pi^i(j)} - p_i c_i \left(S^i_{k_{i-1}+1} - S^i_{k_i+1} + \frac{1}{p_i} \right) \\
&\overset{(A.5)}{=} \sum_{j=k_{i-1}+1}^{k_i} \frac{p_{\pi^i(j)} c_{\pi^i(j)} - p_i c_i}{p_{\pi^i(j)}} \geq 0.
\end{aligned}$$

If $k_{i-1} \in [\![j^*, n-i+1]\!]$, we obtain:

$$\begin{aligned}
\nu_i - \tau_{i-1} &\overset{(A.3),(A.4)}{=} \sum_{j=1}^{k_i} c_{\pi^i(j)} - \sum_{j=1}^{k_{i-1}-1} c_{\pi^i(j)} - p_i c_i \left(S^i_{k_{i-1}} - S^i_{k_i+1} \right) \\
&\overset{(A.5)}{=} \sum_{j=k_{i-1}}^{k_i} \frac{p_{\pi^i(j)} c_{\pi^i(j)} - p_i c_i}{p_{\pi^i(j)}} \geq 0.
\end{aligned}$$

$\square$

The following lemma derives properties satisfied by our auxiliary parameters:

Lemma 15. *Let $i \in [\![0, n-1]\!]$ be such that $r_{\mathrm{H}} > \tau_{i-1}$. Then, the following statements hold:*

- *ℓ_i exists. Furthermore, when $i \geq 1$, let $j^* \in [\![1, n-i+1]\!]$ satisfying $\pi^{i-1}(j^*) = i$. Then,*

$$k_{i-1} \leq \begin{cases} \ell_i & \text{if } k_{i-1} \in [\![0, j^*-1]\!] \\ \ell_i + 1 & \text{if } k_{i-1} \in [\![j^*, n-i+1]\!]. \end{cases} \tag{A.7}$$

- *If $\nu_i < r_{\mathrm{H}}$, then $k_i \leq \ell_i$. If $r_{\mathrm{H}} \leq \nu_i$, then $k_i > \ell_i$.*

Proof of Lemma 15. We prove each statement.

- Let $i \in [\![0, n-1]\!]$. If $i = 0$ and $r_{\mathrm{H}} > \tau_{-1} = 0$, then $0 \in L_0$ and ℓ_0 exists. We now assume that $i \in [\![1, n-1]\!]$ and $r_{\mathrm{H}} > \tau_{i-1}$. Let $j^* \in [\![1, n-i+1]\!]$ be such that $\pi^{i-1}(j^*) = i$. If $k_{i-1} \in [\![0, j^*-1]\!]$, then:

$$r_{\mathrm{H}} > \tau_{i-1} \overset{\text{(A.3),(A.4)}}{=} \sum_{j=1}^{i} c_j + \sum_{j=1}^{k_{i-1}} c_{\pi^i(j)} + p_i c_i S^i_{k_{i-1}+1}.$$

  Since $k_{i-1} \leq j^* - 1 \leq n - i$, then $k_{i-1} \in L_i$, ℓ_i exists, and $k_{i-1} \leq \ell_i$.

  If $k_{i-1} \in [\![j^*, n-i+1]\!]$, then:

$$r_{\mathrm{H}} \overset{\text{(A.3),(A.4)}}{>} \sum_{j=1}^{i} c_j + \sum_{j=1}^{k_{i-1}-1} c_{\pi^i(j)} + p_i c_i S^i_{k_{i-1}-1+1}.$$

  Since $k_{i-1} - 1 \leq n - i$, then $k_{i-1} - 1 \in L_i$, ℓ_i exists, and $k_{i-1} - 1 \leq \ell_i$.

- Let $i \in [\![0, n-1]\!]$ be such that $r_{\mathrm{H}} > \tau_{i-1}$. If $r_{\mathrm{H}} > \nu_i$, then by definition of ν_i, $k_i \in L_i$ and $k_i \leq \ell_i$. On the other hand, if $r_{\mathrm{H}} \leq \nu_i$, then $k_i \notin L_i$. Since ℓ_i exists and $L_i = [\![0, \ell_i]\!]$ (Lemma 14), then $k_i > \ell_i$.

$\square$

We can now prove Theorem 1. Let $r_{\mathrm{S}} \in [\![1, n-1]\!]$, $r_{\mathrm{H}} \in [\![1, m-1]\!]$, and let $i^* \in [\![0, n-1]\!]$ be the unique index satisfying $\tau_{i^*-1} < r_{\mathrm{H}} \leq \tau_{i^*}$.

Proof of Theorem 1 (Regime Pattern 1). Suppose that $i^* = 0$ and $\tau_{-1} < r_{\mathrm{H}} \leq \nu_0$. From Lemma 15, we know that $k_0 > \ell_0$. Let $\rho^{\mathrm{S}^*} \in \mathbb{R}^n$ and $\rho^{\mathrm{H}^*} \in \mathbb{R}^n$ satisfying (2.4) and (2.5), respectively. We will show that $(\rho^{\mathrm{S}^*}, \rho^{\mathrm{H}^*}) \in \widetilde{\mathcal{A}}_{\mathrm{S}} \times \widetilde{\mathcal{A}}_{\mathrm{H}}$ and is an equilibrium of $\widetilde{\Gamma}$.

First, we note that $\ell_0 < k_0 \leq n$, which implies that $\ell_0 + 1 \leq n$. Thus, $S^0_{\ell_0+2}$ is well defined and $\mathcal{K} \neq \emptyset$. For every $i \in \mathcal{K} \setminus \{\pi^0(\ell_0+1)\}$, $p_{\pi^0(\ell_0+1)} \leq p_i$. Furthermore, since $\ell_0 + 1 \in K_0$, we obtain:

$$\sum_{i=1}^{n} \rho_i^{\mathrm{S}^*} = \ell_0 + 1 + p_{\pi^0(\ell_0+1)} S^0_{\ell_0+2} < r_{\mathrm{S}}. \tag{A.8}$$

Thus, $\rho^{\mathrm{S}^*} \in \widetilde{\mathcal{A}}_{\mathrm{S}}$. Next, we show that $\rho^{\mathrm{H}^*} \in \widetilde{\mathcal{A}}_{\mathrm{H}}$. Since $\ell_0 \in L_0$, $\ell_0 + 1 \notin L_0$, and $\ell_0 + 1 \leq n$, then:

$$0 < r_{\mathrm{H}} - \sum_{j=1}^{\ell_0} c_{\pi^0(j)} = \rho^{\mathrm{H}^*}_{\pi^0(\ell_0+1)} = r_{\mathrm{H}} - \sum_{j=1}^{\ell_0+1} c_{\pi^0(j)} + c_{\pi^0(\ell_0+1)} \leq c_{\pi^0(\ell_0+1)}. \tag{A.9}$$

Since $\sum_{i=1}^{n} \rho_i^{\mathrm{H}^*} = r_{\mathrm{H}}$, we then conclude that $\rho^{\mathrm{H}^*} \in \widetilde{\mathcal{A}}_{\mathrm{H}}$.

Next, we show that $(\rho^{\mathrm{S}^*}, \rho^{\mathrm{H}^*})$ is an equilibrium of $\widetilde{\Gamma}$. To this end, we first note the following:

$$\begin{aligned}
\forall\, \rho^{\mathrm{H}} \in \widetilde{\mathcal{A}}_{\mathrm{H}},\ \min_{\rho^{\mathrm{S}} \in \widetilde{\mathcal{A}}_{\mathrm{S}}} \tilde{u}(\rho^{\mathrm{S}}, \rho^{\mathrm{H}}) \overset{(2.3)}{=} \sum_{i=1}^{n} \rho_i^{\mathrm{H}} \ - \ \max_{\rho^{\mathrm{S}}} \ & \sum_{i=1}^{n} p_i \rho_i^{\mathrm{H}} \rho_i^{\mathrm{S}} \\
\text{s.t.} \ & \sum_{i=1}^{n} \rho_i^{\mathrm{S}} \leq r_{\mathrm{S}} \\
& 0 \leq \rho_i^{\mathrm{S}} \leq 1, \quad \forall\, i \in [\![1, n]\!].
\end{aligned} \tag{A.10}$$

Thus, a best response to $\rho^{\mathrm{H}} \in \widetilde{\mathcal{A}}_{\mathrm{H}}$ is an optimal solution to a continuous knapsack problem with n different (fractional) objects of unitary weights and a knapsack capacity equal to r_{S}. Each object $i \in [\![1, n]\!]$ has a profit equal to $p_i \rho_i^{\mathrm{H}}$. An optimal solution consists in filling the capacity of the knapsack with the objects with highest profits.

Then, given ρ^{H^*} satisfying (2.5), the "profit" of each object in the knapsack problem (A.10) is given by:

$$p_i \rho_i^{\mathrm{H}^*} = \begin{cases} p_i c_i & \text{if } i \in \mathcal{J}, \\ p_{\pi^0(\ell_0+1)} \left(r_{\mathrm{H}} - \displaystyle\sum_{j=1}^{\ell_0} c_{\pi^0(j)} \right) & \text{if } i = \pi^0(\ell_0+1), \\ 0 & \text{if } i \in \mathcal{K} \setminus \{\pi^0(\ell_0+1)\}. \end{cases} \tag{A.11}$$

Since $r_{\mathrm{S}} \overset{\text{(A.8)}}{>} \ell_0 + 1 = |\mathcal{J}| + 1$, then a best response to ρ^{H^*} will select all the objects in $\mathcal{J} \cup \{\pi^0(\ell_0+1)\}$ (and might not entirely fill the knapsack). Hence, ρ^{S^*} defined in (2.4) is a best response to ρ^{H^*}.

To show that ρ^{H^*} is a best response to ρ^{S^*}, we similarly observe the following:

$$\begin{aligned} \forall\, \rho^{\mathrm{S}} \in \widetilde{\mathcal{A}}_{\mathrm{S}},\ \max_{\rho^{\mathrm{H}} \in \widetilde{\mathcal{A}}_{\mathrm{H}}} \tilde{u}(\rho^{\mathrm{S}}, \rho^{\mathrm{H}}) \overset{(2.3)}{=} \max_{\rho^{\mathrm{H}}}\ & \sum_{i=1}^{n} (1 - p_i \rho_i^{\mathrm{S}}) \rho_i^{\mathrm{H}} \\ \text{s.t. } & \sum_{i=1}^{n} \rho_i^{\mathrm{H}} \leq r_{\mathrm{H}} \\ & 0 \leq \rho_i^{\mathrm{H}} \leq c_i, \qquad \forall\, i \in [\![1, n]\!]. \end{aligned} \tag{A.12}$$

Thus, a best response to $\rho^{\mathrm{S}} \in \widetilde{\mathcal{A}}_{\mathrm{S}}$ is an optimal solution to another continuous knapsack problem with n different (fractional) objects of unitary weights and a knapsack capacity equal to r_{H}. Each object $i \in [\![1, n]\!]$ is available c_i times and has a profit equal to $(1 - p_i \rho_i^{\mathrm{S}})$. An optimal solution consists in selecting as many copies as possible of the objects with highest profits until filling the capacity of the knapsack.

Then, given ρ^{S^*} satisfying (2.4), the profit of each object in the knapsack problem

(A.12) is given by:

$$1 - p_i\rho_i^{\mathrm{S}^*} = \begin{cases} 1 - p_i & \text{if } i \in \mathcal{J} \cup \{\pi^0(\ell_0+1)\}, \\ 1 - p_{\pi^0(\ell_0+1)} & \text{if } i \in \mathcal{K} \setminus \{\pi^0(\ell_0+1)\}. \end{cases} \tag{A.13}$$

By definition of π^0, we have the following inequalities: $1-p_{\pi^0(1)} \geq \cdots \geq 1-p_{\pi^0(\ell_0+1)}$. Since (A.9) implies that $r_{\mathrm{H}} > \sum_{j=1}^{\ell_0} c_{\pi^0(j)}$, then a best response to ρ^{S^*} consist in selecting all copies of the objects in $\mathcal{J}$ and any fraction of the objects in $\mathcal{K}$ until the knapsack is full. Hence, ρ^{H^*} defined in (2.5) is a best response to ρ^{S^*}.

Thus, $(\rho^{\mathrm{S}^*}, \rho^{\mathrm{H}^*})$ is an equilibrium of $\widetilde{\Gamma}$. Furthermore, the value of the game is given by:

$$\tilde{u}(\rho^{\mathrm{S}^*}, \rho^{\mathrm{H}^*}) = r_{\mathrm{H}} - \sum_{j=1}^{\ell_0} p_{\pi^0(j)}c_{\pi^0(j)} - p_{\pi^0(\ell_0+1)}\left(r_{\mathrm{H}} - \sum_{j=1}^{\ell_0} c_{\pi^0(j)}\right).$$

□

Proof of Theorem 1 (Regime Pattern 2). Suppose that $\nu_{i^*} < r_{\mathrm{H}} \leq \tau_{i^*}$. From Lemma 15, we know that $k_{i^*} \leq \ell_{i^*}$. Let $\rho^{\mathrm{S}^*} \in \mathbb{R}^n$ and $\rho^{\mathrm{H}^*} \in \mathbb{R}^n$ satisfying (2.6) and (2.7), respectively. We will analogously show that $(\rho^{\mathrm{S}^*}, \rho^{\mathrm{H}^*}) \in \widetilde{\mathcal{A}}_{\mathrm{S}} \times \widetilde{\mathcal{A}}_{\mathrm{H}}$ and is an equilibrium of $\widetilde{\Gamma}$.

First, we note that $k_{i^*} < n - i^*$. Indeed, if $k_{i^*} = n - i^*$, then $r_{\mathrm{H}} > \nu_{i^*} = \sum_{j=1}^{n} c_j = m$, which is a contradiction. Therefore, $\mathcal{K} \neq \emptyset$ and $S_{k_{i^*}+1}^{i^*} > 0$. By definition of k_{i^*}, we obtain

$$r_{\mathrm{S}} > k_{i^*} + p_{\pi^{i^*}(k_{i^*})} S_{k_{i^*}+1}^{i^*} \geq k_{i^*}, \tag{A.14}$$

which implies that $\rho_i^{\mathrm{S}^*} \geq 0$ for every $i \in \mathcal{K}$. Furthermore, since $k_{i^*} + 1 \leq n - i^*$ and $k_{i^*} + 1 \notin K_{i^*}$, then:

$$r_{\mathrm{S}} \leq k_{i^*} + 1 + p_{\pi^{i^*}(k_{i^*}+1)}\left(S_{k_{i^*}+1}^{i^*} - \frac{1}{p_{\pi^{i^*}(k_{i^*}+1)}}\right) = k_{i^*} + p_{\pi^{i^*}(k_{i^*}+1)} S_{k_{i^*}+1}^{i^*}. \tag{A.15}$$

Thus, for every $i \in \mathcal{K}$, $\rho_i^{\mathrm{S}^*} \leq 1$. Finally,

$$\sum_{i=1}^{n} \rho_i^{\mathrm{S}^*} = |\mathcal{J}| + \frac{r_{\mathrm{S}} - k_{i^*}}{S_{k_{i^*}+1}^{i^*}} S_{k_{i^*}+1}^{i^*} = r_{\mathrm{S}}.$$

Therefore, $\rho^{\mathrm{S}^*} \in \widetilde{\mathcal{A}}_{\mathrm{S}}$. Next, we show that $\rho^{\mathrm{H}^*} \in \widetilde{\mathcal{A}}_{\mathrm{H}}$. Since $k_{i^*} \leq \ell_{i^*}$, then $k_{i^*} \in L_{i^*}$ and

$$r_{\mathrm{H}} > \sum_{j=1}^{i^*} c_j + \sum_{j=1}^{k_{i^*}} c_{\pi^{i^*}(j)} + p_{i^*} c_{i^*} S_{k_{i^*}+1}^{i^*} \geq \sum_{j=1}^{i^*} c_j + \sum_{j=1}^{k_{i^*}} c_{\pi^{i^*}(j)}. \tag{A.16}$$

Thus, $\rho_i^{\mathrm{H}^*} \geq 0$ for every $i \in \mathcal{K}$. Furthermore, since $r_{\mathrm{H}} \leq \tau_{i^*}$ and $p_{i^*+1} c_{i^*+1} \leq p_i c_i$ for every $i \in \mathcal{J} \cup \mathcal{K}$, we obtain:

$$\forall\, i \in \mathcal{K},\ \rho_i^{\mathrm{H}^*} = \frac{r_{\mathrm{H}} - \sum_{j=1}^{i^*} c_j - \sum_{j=1}^{k_{i^*}} c_{\pi^{i^*}(j)}}{p_i S_{k_{i^*}+1}^{i^*}} \leq \frac{p_{i^*+1} c_{i^*+1}}{p_i} \leq c_i. \tag{A.17}$$

Finally,

$$\sum_{i=1}^{n} \rho_i^{\mathrm{H}^*} = \sum_{i=1}^{i^*} c_i + \sum_{j=1}^{k_{i^*}} c_{\pi^{i^*}(j)} + \left(r_{\mathrm{H}} - \sum_{j=1}^{i^*} c_j - \sum_{j=1}^{k_{i^*}} c_{\pi^{i^*}(j)} \right) \frac{S_{k_{i^*}+1}^{i^*}}{S_{k_{i^*}+1}^{i^*}} = r_{\mathrm{H}}.$$

Therefore, $\rho^{\mathrm{H}^*} \in \widetilde{\mathcal{A}}_{\mathrm{H}}$.

Next, we show that $(\rho^{\mathrm{S}^*}, \rho^{\mathrm{H}^*})$ is an equilibrium of $\widetilde{\Gamma}$. Given ρ^{H^*} satisfying (2.7), the profit of each object in the knapsack problem (A.10) is given by:

$$p_i \rho_i^{\mathrm{H}^*} = \begin{cases} p_i c_i & \text{if } i \in \mathcal{I} \cup \mathcal{J}, \\ \dfrac{r_{\mathrm{H}} - \sum_{j=1}^{i^*} c_j - \sum_{j=1}^{k_{i^*}} c_{\pi^{i^*}(j)}}{S_{k_{i^*}+1}^{i^*}} & \text{if } i \in \mathcal{K}. \end{cases} \tag{A.18}$$

We know that for every $i \in \mathcal{I}$ and every $l \in \mathcal{J}$:

$$p_i c_i \leq p_{i^*} c_{i^*} \overset{\text{(A.16)}}{<} \frac{r_{\mathrm{H}} - \sum_{j=1}^{i^*} c_j - \sum_{j=1}^{k_{i^*}} c_{\pi^{i^*}(j)}}{S_{k_{i^*}+1}^{i^*}} \overset{\text{(A.17)}}{\leq} p_{i^*+1} c_{i^*+1} \leq p_l c_l. \tag{A.19}$$

Therefore, the objects in $\mathcal{J}$ are the most profitable, followed by the objects in $\mathcal{K}$ that have equal profit, followed by the objects in $\mathcal{I}$. We now must determine bounds on r_{S}. We showed in (A.14) that $r_{\mathrm{S}} > k_{i^*} = |\mathcal{J}|$. An upper bound is given as follows:

$$r_{\mathrm{S}} \overset{\text{(A.15)}}{\leq} k_{i^*} + 1 + \sum_{j=k_{i^*}+2}^{n-i^*} \frac{p_{\pi^{i^*}(k_{i^*}+1)}}{p_{\pi^{i^*}(j)}} \leq n - i^* = |\mathcal{J}| + |\mathcal{K}|.$$

Thus, a best response to ρ^{H^*} selects all the objects in $\mathcal{J}$ and any fraction of the objects in $\mathcal{K}$ until the knapsack is full. Hence, ρ^{S^*} defined in (2.6) is a best response to ρ^{H^*}. Then, given ρ^{S^*} satisfying (2.6), the profit of each object in the knapsack problem (A.12) is given by:

$$1 - p_i \rho_i^{\mathrm{S}^*} = \begin{cases} 1 & \text{if } i \in \mathcal{I}, \\ 1 - p_i & \text{if } i \in \mathcal{J}, \\ 1 - \dfrac{r_{\mathrm{S}} - k_{i^*}}{S^{i^*}_{k_{i^*}+1}} & \text{if } i \in \mathcal{K}. \end{cases} \tag{A.20}$$

We have the following inequalities:

$$\forall\, i \in \mathcal{J},\ 1 - \frac{r_{\mathrm{S}} - k_{i^*}}{S^{i^*}_{k_{i^*}+1}} \overset{\text{(A.14)}}{<} 1 - p_{\pi^{i^*}(k_{i^*})} \leq 1 - p_i < 1. \tag{A.21}$$

Therefore, the objects in $\mathcal{I}$ are the most profitable, followed by the objects in $\mathcal{J}$, followed by the objects in $\mathcal{K}$ that have equal profit. To determine which objects will be selected given r_{H}, we recall that $r_{\mathrm{H}} < m = \sum_{i=1}^{n} c_i$. Furthermore, (A.16) implies that $r_{\mathrm{H}} > \sum_{i \in \mathcal{I} \cup \mathcal{J}} c_i$.

Thus, a best response to ρ^{S^*} consists in selecting all copies of the objects in $\mathcal{I}$ and $\mathcal{J}$, and in selecting any fraction of the objects in $\mathcal{K}$ until the knapsack is full. Hence, ρ^{H^*} defined in (2.7) is a best response to ρ^{S^*}.

As a consequence, $(\rho^{\mathrm{S}^*}, \rho^{\mathrm{H}^*})$ is an equilibrium of $\widetilde{\Gamma}$. Furthermore, the value of the game is given by:

$$\tilde{u}(\rho^{S^*}, \rho^{H^*}) = r_H - \sum_{j=1}^{k_{i^*}} p_{\pi^{i^*}(j)} c_{\pi^{i^*}(j)} - \frac{\left(r_S - k_{i^*}\right)\left(r_H - \sum_{j=1}^{i^*} c_j - \sum_{j=1}^{k_{i^*}} c_{\pi^{i^*}(j)}\right)}{S^{i^*}_{k_{i^*}+1}}.$$

$\square$

Proof of Theorem 1 (Regime Pattern 3). Finally, we consider the case when $i^* \geq 1$ and $\tau_{i^*-1} < r_H \leq \nu_{i^*}$. From Lemma 15, we know that $k_{i^*} > \ell_{i^*}$. Let $\rho^{S^*} \in \mathbb{R}^n$ and $\rho^{H^*} \in \mathbb{R}^n$ satisfying (2.8) and (2.9), respectively. We will analogously show that $(\rho^{S^*}, \rho^{H^*}) \in \widetilde{\mathcal{A}}_S \times \widetilde{\mathcal{A}}_H$ and is an equilibrium of $\widetilde{\Gamma}$.

Similarly, we note that $\ell_{i^*} < k_{i^*} \leq n - i^*$, which implies that $\ell_{i^*} + 1 \leq n - i^*$. Thus, $S^{i^*}_{\ell_{i^*}+2}$ is well defined and $\mathcal{K} \neq \emptyset$. Since $\ell_{i^*} + 1 \in K_{i^*}$, then

$$r_S > p_{\pi^{i^*}(\ell_{i^*}+1)} S^{i^*}_{\ell_{i^*}+2} + \ell_{i^*} + 1 = p_{\pi^{i^*}(\ell_{i^*}+1)} S^{i^*}_{\ell_{i^*}+1} + \ell_{i^*}. \tag{A.22}$$

Thus, $\rho^{S^*}_{i^*} > 0$. Next, we will show by contradiction the following upper bound:

$$r_S - \ell_{i^*} - p_{\pi^{i^*}(\ell_{i^*}+1)} S^{i^*}_{\ell_{i^*}+1} \leq \min\left\{\frac{p_{\pi^{i^*}(\ell_{i^*}+1)}}{p_{i^*}}, 1\right\}. \tag{A.23}$$

Let us assume that (A.23) does not hold, and let $j^* \in [\![1, n - i^* + 1]\!]$ satisfying $\pi^{i^*-1}(j^*) = i^*$. If $\ell_{i^*} + 1 \leq j^* - 1$, then:

$$\begin{aligned}\frac{p_{\pi^{i^*}(\ell_{i^*}+1)}}{p_{i^*}} = \min\left\{\frac{p_{\pi^{i^*}(\ell_{i^*}+1)}}{p_{i^*}}, 1\right\} &< r_S - \ell_{i^*} - 1 - p_{\pi^{i^*}(\ell_{i^*}+1)} S^{i^*}_{\ell_{i^*}+2} \\ &\overset{(A.2),(A.4)}{=} r_S - \ell_{i^*} - 1 - p_{\pi^{i^*-1}(\ell_{i^*}+1)} S^{i^*-1}_{\ell_{i^*}+2} + \frac{p_{\pi^{i^*}(\ell_{i^*}+1)}}{p_{i^*}},\end{aligned}$$

which implies that $\ell_{i^*} + 1 \leq k_{i^*-1}$. However, by Lemma 15, this can only occur when $k_{i^*-1} \geq j^*$, for which we obtain the following contradiction $j^* \leq k_{i^*-1} \overset{(A.7)}{\leq} \ell_{i^*} + 1 \leq j^* - 1$.

If on the other hand $\ell_{i^*}+1 \geq j^*$, then $j^* < \ell_{i^*}+2 \leq n-i^*+1$ and

$$1 = \min\left\{\frac{p_{\pi^{i^*-1}(\ell_{i^*}+2)}}{p_{i^*}},\, 1\right\} = \min\left\{\frac{p_{\pi^{i^*}(\ell_{i^*}+1)}}{p_{i^*}},\, 1\right\} \overset{\text{(A.2),(A.4)}}{<} r_{\mathrm{S}} - \ell_{i^*} - 1 - p_{\pi^{i^*-1}(\ell_{i^*}+2)} S^{i^*-1}_{\ell_{i^*}+3},$$

which implies that $\ell_{i^*}+2 \leq k_{i^*-1}$, thus contradicting (A.7). Therefore, (A.23) holds. Finally,

$$\sum_{i=1}^{n} \rho_i^{\mathrm{S}^*} = r_{\mathrm{S}} - \ell_{i^*} - p_{\pi^*(\ell_{i^*}+1)} S^{i^*}_{\ell_{i^*}+1} + |\mathcal{J}| + 1 + p_{\pi^*(\ell_{i^*}+1)} S^{i^*}_{\ell_{i^*}+2} = r_{\mathrm{S}},$$

which implies that $\rho^{\mathrm{S}^*} \in \widetilde{\mathcal{A}}_{\mathrm{S}}$.

Next, we show that $\rho^{\mathrm{H}^*} \in \widetilde{\mathcal{A}}_{\mathrm{H}}$. Since locations in $[\![1, n]\!]$ are ordered by their detection potentials, then for every $i \in \mathcal{K},\ p_i c_i \geq p_{i^*} c_{i^*}$. Since $\ell_{i^*} \in L_{i^*}$, then:

$$\rho^{\mathrm{H}^*}_{\pi^{i^*}(\ell_{i^*}+1)} = r_{\mathrm{H}} - \sum_{j=1}^{i^*} c_j - \sum_{j=1}^{\ell_{i^*}} c_{\pi^{i^*}(j)} - p_{i^*} c_{i^*} S^{i^*}_{\ell_{i^*}+1} + \frac{p_{i^*} c_{i^*}}{p_{\pi^{i^*}(\ell_{i^*}+1)}} > \frac{p_{i^*} c_{i^*}}{p_{\pi^{i^*}(\ell_{i^*}+1)}} \geq 0. \tag{A.24}$$

Furthermore, since $\ell_{i^*}+1 \notin L_{i^*}$ and $\ell_{i^*}+1 \leq n-i^*$, then:

$$\rho^{\mathrm{H}^*}_{\pi^{i^*}(\ell_{i^*}+1)} = r_{\mathrm{H}} - \sum_{j=1}^{i^*} c_j - \sum_{j=1}^{\ell_{i^*}+1} c_{\pi^{i^*}(j)} - p_{i^*} c_{i^*} S^{i^*}_{\ell_{i^*}+2} + c_{\pi^{i^*}(\ell_{i^*}+1)} \leq c_{\pi^{i^*}(\ell_{i^*}+1)}. \tag{A.25}$$

Finally,

$$\sum_{i=1}^{n} \rho_i^{\mathrm{H}^*} = \sum_{i \in \mathcal{I} \cup \mathcal{J}} c_i + r_{\mathrm{H}} - \sum_{i=1}^{i^*} c_j - \sum_{j=1}^{\ell_{i^*}} c_{\pi^*(j)} - p_{i^*} c_{i^*} S^{i^*}_{\ell_{i^*}+2} + p_{i^*} c_{i^*} S^{i^*}_{\ell_{i^*}+2} = r_{\mathrm{H}}.$$

Therefore, $\rho^{\mathrm{H}^*} \in \widetilde{\mathcal{A}}_{\mathrm{H}}$.

Next, we show that $(\rho^{\mathrm{S}^*}, \rho^{\mathrm{H}^*})$ is an equilibrium of $\widetilde{\Gamma}$. Given ρ^{H^*} satisfying (2.9), the profit of each object in the knapsack problem (A.10) is given by:

$$p_i\rho_i^{\mathrm{H}^*} = \begin{cases} p_i c_i & \text{if } i \in \mathcal{I} \cup \mathcal{J}, \\ p_{\pi^{i^*}(\ell_{i^*}+1)} \left(r_{\mathrm{H}} - \sum_{j=1}^{i^*} c_j - \sum_{j=1}^{\ell_{i^*}} c_{\pi^{i^*}(j)} - p_{i^*} c_{i^*} S^{i^*}_{\ell_{i^*}+2} \right) & \text{if } i = \pi^{i^*}(\ell_{i^*}+1), \\ p_{i^*} c_{i^*} & \text{if } i \in \mathcal{K} \setminus \{\pi^{i^*}(\ell_{i^*}+1)\}. \end{cases} \tag{A.26}$$

Since the locations in $[\![1, n]\!]$ are ordered by their detection potentials, then we know that $p_1 c_1 \leq \cdots \leq p_{i^*} c_{i^*} \leq p_j c_j$ for every $j \in \mathcal{J}$. Furthermore, we have the following inequality:

$$p_{i^*} c_{i^*} \overset{\text{(A.24)}}{<} p_{\pi^{i^*}(\ell_{i^*}+1)} \left(r_{\mathrm{H}} - \sum_{j=1}^{i^*} c_j - \sum_{j=1}^{\ell_{i^*}} c_{\pi^{i^*}(j)} - p_{i^*} c_{i^*} S^{i^*}_{\ell_{i^*}+2} \right). \tag{A.27}$$

Thus, the objects in $\mathcal{J} \cup \{\pi^{i^*}(\ell_{i^*}+1)\}$ are the most profitable, followed by the objects in $\{i^*\} \cup \mathcal{K} \setminus \{\pi^{i^*}(\ell_{i^*}+1)\}$ that have equal profit, followed by the objects in $\mathcal{I} \setminus \{i^*\}$.

From (A.22), we know that $r_{\mathrm{S}} > \ell_{i^*} + 1 = |\mathcal{J} \cup \{\pi^{i^*}(\ell_{i^*}+1)\}|$. Furthermore, an upper bound is given as follows:

$$r_{\mathrm{S}} \overset{\text{(A.23)}}{\leq} \ell_{i^*} + 1 + \sum_{j=\ell_{i^*}+1}^{n-i^*} \frac{p_{\pi^{i^*}(\ell_{i^*}+1)}}{p_{\pi^{i^*}(j)}} \leq n - i^* + 1 = |\{i^*\} \cup \mathcal{J} \cup \mathcal{K}|.$$

Therefore, one best response to ρ^{H^*} will select all the objects in $\mathcal{J} \cup \{\pi^{i^*}(\ell_{i^*}+1)\}$ and will select any fraction of the objects in $\{i^*\} \cup \mathcal{K} \setminus \{\pi^{i^*}(\ell_{i^*}+1)\}$ until the knapsack is full. Hence, ρ^{S^*} defined in (2.8) is a best response to ρ^{H^*}.

Then, given ρ^{S^*} satisfying (2.8), the profit of each object in the knapsack problem

(A.12) is given by:

$$1 - p_i\rho_i^{\mathrm{S}^*} = \begin{cases} 1 & \text{if } i \in \mathcal{I} \setminus \{i^*\} \\ 1 - p_{i^*}(r_{\mathrm{S}} - \ell_{i^*} - p_{\pi^*(\ell_{i^*}+1)}S^{i^*}_{\ell_{i^*}+1}) & \text{if } i = i^* \\ 1 - p_i & \text{if } i \in \mathcal{J} \cup \{\pi^{i^*}(\ell_{i^*}+1)\}, \\ 1 - p_{\pi^{i^*}(\ell_{i^*}+1)} & \text{if } i \in \mathcal{K} \setminus \{\pi^{i^*}(\ell_{i^*}+1)\}. \end{cases} \tag{A.28}$$

By definition of π^{i^*}, we have the following inequalities: $1 - p_{\pi^{i^*}(1)} \geq \cdots \geq 1 - p_{\pi^{i^*}(\ell_{i^*}+1)}$. Furthermore, (A.23) implies that $1 - p_{i^*}(r_{\mathrm{S}} - \ell_{i^*} - p_{\pi^*(\ell_{i^*}+1)}S^{i^*}_{\ell_{i^*}+1}) \geq 1 - p_{\pi^{i^*}(\ell_{i^*}+1)}$.

Since (A.24) implies that $r_{\mathrm{H}} > \sum_{j=1}^{i^*} c_j + \sum_{j=1}^{\ell_{i^*}} c_{\pi^{i^*}(j)}$, then one best response to ρ^{S^*} selects all copies of the objects in $\mathcal{I} \cup \mathcal{J}$ and selects any fraction of the objects in $\mathcal{K}$ until the knapsack is full. Hence, ρ^{H^*} defined in (2.9) is a best response to ρ^{S^*}.

Thus, $(\rho^{\mathrm{S}^*}, \rho^{\mathrm{H}^*})$ is an equilibrium of $\widetilde{\Gamma}$. Furthermore, the value of the game is given by:

$$\begin{aligned} \tilde{u}(\rho^{\mathrm{S}^*}, \rho^{\mathrm{H}^*}) = & r_{\mathrm{H}} - p_{i^*}\left(r_{\mathrm{S}} - \ell_{i^*} - p_{\pi^{i^*}(\ell_{i^*}+1)}S^{i^*}_{\ell_{i^*}+1}\right)c_{i^*} - \sum_{j=1}^{\ell_{i^*}} p_{\pi^{i^*}(j)}c_{\pi^{i^*}(j)} \\ & - p_{\pi^{i^*}(\ell_{i^*}+1)}\left(r_{\mathrm{H}} - \sum_{j=1}^{i^*} c_j - \sum_{j=1}^{\ell_{i^*}} c_{\pi^{i^*}(j)}\right). \end{aligned}$$

□

Proposition 8. *Let* $r_{\mathrm{S}} \in [\![1, n-1]\!]$ *and* $r_{\mathrm{H}} \in [\![1, m-1]\!]$. *Let* $i^* \in [\![0, n-1]\!]$ *satisfying* $\tau_{i^*-1} < r_{\mathrm{H}} \leq \tau_{i^*}$.

- Regime Pattern 1: $i^* = 0$ and $\tau_{-1} < r_{\mathrm{H}} \leq \nu_0$. *Suppose also that* $p_{\pi^0(\ell_0)} < p_{\pi^0(\ell_0+1)} < 1$ *and* $p_{\pi^0(\ell_0+1)} < p_{\pi^0(\ell_0+2)}$ *if* $\ell_0 \leq n-2$. *Then,* (2.4) *and* (2.5) *are necessary and sufficient conditions for a strategy profile* $(\rho^{\mathrm{S}^*}, \rho^{\mathrm{H}^*}) \in \widetilde{\mathcal{A}}_{\mathrm{S}} \times \widetilde{\mathcal{A}}_{\mathrm{H}}$ *to be an equilibrium of* $\widetilde{\Gamma}$.

- Regime Pattern 2: $\nu_{i^*} < r_{\mathrm{H}} \leq \tau_{i^*}$. *Suppose also that* $r_{\mathrm{H}} < \tau_{i^*}$ *and* $r_{\mathrm{S}} < k_{i^*} + 1 +$

$p_{\pi^{i^*}(k_{i^*}+1)} S^{i^*}_{k_{i^*}+2}$. *Then,* (2.6) *and* (2.7) *are necessary and sufficient conditions for a strategy profile* $(\rho^{\mathrm{S}^*}, \rho^{\mathrm{H}^*}) \in \widetilde{\mathcal{A}}_{\mathrm{S}} \times \widetilde{\mathcal{A}}_{\mathrm{H}}$ *to be an equilibrium of* $\widetilde{\Gamma}$.

– Regime Pattern 3: $i^* \geq 1$ and $\tau_{i^*-1} < r_{\mathrm{H}} \leq \nu_{i^*}$. *Suppose also that* $p_{i^*-1}c_{i^*-1} < p_{i^*}c_{i^*} < p_{i^*+1}c_{i^*+1}$, $r_{\mathrm{H}} < \sum_{j=1}^{i^*} c_j + \sum_{j=1}^{\ell_{i^*}+1} c_{\pi^{i^*}(j)} + p_{i^*}c_{i^*}S^{i^*}_{\ell_{i^*}+2}$, $r_{\mathrm{S}} < k_{i^*-1} + 1 + p_{\pi^{i^*-1}(k_{i^*-1}+1)} S^{i^*-1}_{k_{i^*-1}+2}$, $p_{\pi^{i^*}(\ell_{i^*})} < p_{\pi^{i^*}(\ell_{i^*}+1)} < 1$, *and* $p_{\pi^{i^*}(\ell_{i^*}+1)} < p_{\pi^{i^*}(\ell_{i^*}+2)}$ *if* $\ell_{i^*} \leq n - i^* - 2$. *Then,* (2.8) *and* (2.9) *are necessary and sufficient conditions for a strategy profile* $(\rho^{\mathrm{S}^*}, \rho^{\mathrm{H}^*}) \in \widetilde{\mathcal{A}}_{\mathrm{S}} \times \widetilde{\mathcal{A}}_{\mathrm{H}}$ *to be an equilibrium of* $\widetilde{\Gamma}$.

Proof of Proposition 8. We prove each statement.

– *Regime Pattern 1:* $i^* = 0$ *and* $\tau_{-1} < r_{\mathrm{H}} \leq \nu_0$. We additionally consider the following non-edge case assumptions: $p_{\pi^0(\ell_0)} < p_{\pi^0(\ell_0+1)} < 1$ and $p_{\pi^0(\ell_0+1)} < p_{\pi^0(\ell_0+2)}$ if $\ell_0 \leq n-2$.

  Let $\rho^{\mathrm{H}^*} \in \widetilde{\mathcal{A}}_{\mathrm{H}}$ satisfying (2.5) and let $\rho^{\mathrm{S}'}$ be an equilibrium strategy for S in $\widetilde{\Gamma}$. Since $\rho^{\mathrm{S}'}$ is a best response to ρ^{H^*}, it is an optimal solution to the continuous knapsack problem (A.10). The profits of each object are given by (A.11). Furthermore, (A.9) and the inequality $r_{\mathrm{S}} > |\mathcal{J}| + 1$ imply that any best response to ρ^{H^*} must select all the objects in $\mathcal{J} \cup \{\pi^0(\ell_0+1)\}$. Therefore, $\rho^{\mathrm{S}'}_i = 1$ for every $i \in \mathcal{J} \cup \{\pi^0(\ell_0+1)\}$. Next, since ρ^{H^*} is a best response to $\rho^{\mathrm{S}'}$, then it is an optimal solution to the continuous knapsack problem (A.12). Next, we write the dual of (A.12) associated with $\rho^{\mathrm{S}'}$:

$$\begin{aligned} \min_{\alpha,\beta} \quad & r_{\mathrm{H}}\,\alpha + \sum_{i=1}^{n} c_i\beta_i \\ \text{subject to} \quad & \alpha + \beta_i \geq 1 - p_i\rho^{\mathrm{S}'}_i \qquad \forall\, i \in [\![1,n]\!], \\ & \alpha \geq 0 \\ & \beta_i \geq 0 \qquad\qquad\qquad \forall\, i \in [\![1,n]\!]. \end{aligned} \tag{A.29}$$

  Let (α^*, β^*) be an optimal solution of the dual problem (A.29). Since ρ^{H^*} is a best response to $\rho^{\mathrm{S}'}$, then it is an optimal solution to the continuous knapsack problem (A.12). Since $0 \overset{(A.9)}{<} \rho^{\mathrm{H}^*}_{\pi^0(\ell_0+1)}$, then at optimality of the dual (A.29), $\alpha^* = 1 - p_{\pi^0(\ell_0+1)}\rho^{\mathrm{S}'}_{\pi^0(\ell_0+1)} - \beta^*_{\pi^0(\ell_0+1)} = 1 - p_{\pi^0(\ell_0+1)} - \beta^*_{\pi^0(\ell_0+1)} \leq 1 - p_{\pi^0(\ell_0+1)}$.

Finally, for every $i \in \mathcal{K} \setminus \{\pi^0(\ell_0+1)\}$, $\rho^{\mathrm{H}^*}_{\pi^0(\ell_0+1)} = 0 < c_{\pi^0(\ell_0+1)}$. Thus, by complementary slackness, $\beta^*_i = 0$ and $\rho^{\mathrm{S}'}_i \geq (1-\alpha^*)/p_i \geq p_{\pi^0(\ell_0+1)}/p_i$ for every $i \in \mathcal{K} \setminus \{\pi^0(\ell_0+1)\}$. In conclusion, $\rho^{\mathrm{S}'}$ satisfies (2.4).

Let $\rho^{\mathrm{H}'}$ be an equilibrium strategy for H in $\widetilde{\Gamma}$, and consider $\rho^{\mathrm{S}^*} \in \mathcal{A}_{\mathrm{S}}$ satisfying

$$\begin{aligned} \rho^{\mathrm{S}^*}_i &= 1 \qquad \text{if } i \in \mathcal{J} \cup \{\pi^0(\ell_0+1)\}, \\ \frac{p_{\pi^0(\ell_0+1)}}{p_i} < \rho^{\mathrm{S}^*}_i &< 1 \qquad \text{if } i \in \mathcal{K} \setminus \{\pi^0(\ell_0+1)\}, \\ \sum_{i=1}^{n} \rho^{\mathrm{S}^*}_i &< r_{\mathrm{S}}. \end{aligned}$$

Such a vector exists as a consequence of (A.8) and since $p_{\pi^0(\ell_0+1)} < p_i$ for every $i \in \mathcal{K} \setminus \{\pi^0(\ell_0+1)\}$ under the non-edge case assumptions. Then, $\rho^{\mathrm{H}'}$ is a best response to ρ^{S^*} and is an optimal solution to the continuous knapsack problem (A.12). The profits of each object are given by:

$$\begin{aligned} &\forall\, i \in \mathcal{J} \cup \{\pi^0(\ell_0+1)\},\ 1 - p_i\rho^{\mathrm{S}^*}_i = 1 - p_i \\ &\forall\, i \in \mathcal{K} \setminus \{\pi^0(\ell_0+1)\},\ 1 - p_i\rho^{\mathrm{S}^*}_i < 1 - p_{\pi^0(\ell_0+1)}. \end{aligned}$$

Under the non-edge case assumptions, $1-p_{\pi^0(1)} \geq \cdots \geq 1-p_{\pi^0(\ell_0)} > 1-p_{\pi^0(\ell_0+1)} > 0$. Since $\sum_{j=1}^{\ell_0} c_{\pi^0(j)} < r_{\mathrm{H}} < \sum_{j=1}^{\ell_0+1} c_{\pi^0(j)}$, then any best response to ρ^{S^*} selects all copies of the objects in $\mathcal{J}$ and fills the remaining of the knapsack with objects in $\pi^0(\ell_0+1)$. Therefore $\rho^{\mathrm{H}'}$ satisfies (2.5).

– *Regime Pattern 2:* $\nu_{i^*} < r_{\mathrm{H}} \leq \tau_{i^*}$. We additionally consider the following non-edge case assumptions: $r_{\mathrm{H}} < \tau_{i^*}$ and $r_{\mathrm{S}} < k_{i^*} + 1 + p_{\pi^{i^*}(k_{i^*}+1)} S^{i^*}_{k_{i^*}+2}$.

Let $(\rho^{\mathrm{S}^*}, \rho^{\mathrm{H}^*}) \in \widetilde{\mathcal{A}}_{\mathrm{S}} \times \widetilde{\mathcal{A}}_{\mathrm{H}}$ satisfying (2.6) and (2.7), and let $(\rho^{\mathrm{S}'}, \rho^{\mathrm{H}'}) \in \widetilde{\mathcal{A}}_{\mathrm{S}} \times \widetilde{\mathcal{A}}_{\mathrm{H}}$ be any equilibrium of the game $\widetilde{\Gamma}$. Since $\widetilde{\Gamma}$ is a zero-sum game, then $(\rho^{\mathrm{S}'}, \rho^{\mathrm{H}^*})$ is also an equilibrium of $\widetilde{\Gamma}$.

Since $\rho^{\mathrm{S}'}$ is a best response to ρ^{H^*}, it is an optimal solution to the continuous knapsack problem (A.10). The profits of each object are given by (A.18) and satisfy inequalities (A.19). Furthermore, $r_{\mathrm{H}} < \tau_{i^*}$ implies that:

$$\frac{r_{\mathrm{H}} - \sum_{j=1}^{i^*} c_j - \sum_{j=1}^{k_{i^*}} c_{\pi^{i^*}(j)}}{S^{i^*}_{k_{i^*}+1}} < p_{i^*+1} c_{i^*+1}.$$

Since $|\mathcal{J}| < r_{\mathrm{S}} \leq |\mathcal{J}| + |\mathcal{K}|$, then any best response to ρ^{H^*} must select all the objects in $\mathcal{J}$, must not select any object in $\mathcal{I}$, and must entirely fill the knapsack. Thus, $\rho^{\mathrm{S}'}_i = 0$ for every $i \in \mathcal{I}$, $\rho^{\mathrm{S}'}_i = 1$ for every $i \in \mathcal{J}$, and $\sum_{i=1}^n \rho^{\mathrm{S}'}_i = r_{\mathrm{S}}$.

Next, since ρ^{H^*} is a best response to $\rho^{\mathrm{S}'}$, then it is an optimal solution to the continuous knapsack problem (A.12). Let (α^*, β^*) be an optimal solution of the dual problem (A.29). Since $0 < \rho^{\mathrm{H}^*}_i < c_i$ for every $i \in \mathcal{K}$, then by complementary slackness, $\beta^*_i = 0$ and $\rho^{\mathrm{S}'}_i = (1 - \alpha^*)/p_i$ for every $i \in \mathcal{K}$. Since $\rho^{\mathrm{S}'}$ must fill the knapsack (A.10) entirely, then:

$$r_{\mathrm{S}} = \sum_{i \in \mathcal{I} \cup \mathcal{J} \cup \mathcal{K}} \rho^{\mathrm{S}'}_i = k_{i^*} + (1 - \alpha^*) S^{i^*}_{k_{i^*}+1}.$$

Thus, for every $i \in \mathcal{K}$, $\rho^{\mathrm{S}'}_i = (r_{\mathrm{S}} - k_{i^*})/(p_i S^{i^*}_{k_{i^*}+1})$. In conclusion, $\rho^{\mathrm{S}'}$ satisfies (2.6).

Similarly, $(\rho^{\mathrm{S}^*}, \rho^{\mathrm{H}'})$ is an equilibrium of $\widetilde{\Gamma}$. Then, $\rho^{\mathrm{H}'}$ is a best response to ρ^{S^*} and is an optimal solution to the continuous knapsack problem (A.12). The profits of each object are given by (A.20) and satisfy inequalities (A.21). Since $\sum_{i \in \mathcal{I} \cup \mathcal{J}} c_i < r_{\mathrm{H}} < m$, then any best response to ρ^{S^*} must select all copies of the objects in $\mathcal{I}$ and $\mathcal{J}$. Therefore, $\rho^{\mathrm{H}'}_i = c_i$ for every $i \in \mathcal{I} \cup \mathcal{J}$. Furthermore, $k_{i^*} + 1 \notin K_{i^*}$ and the non-edge case assumptions imply that:

$$1 - \frac{r_{\mathrm{S}} - k_{i^*}}{S^{i^*}_{k_{i^*}+1}} \overset{\text{(A.15)}}{>} 1 - p_{\pi^{i^*}(k_{i^*}+1)} \geq 0. \tag{A.30}$$

Therefore, $\rho^{\mathrm{H}'}$ must entirely fill the knapsack and $\sum_{i=1}^{n} \rho_i^{\mathrm{H}'} = r_{\mathrm{H}}$.

Next, since ρ^{S^*} is a best response to $\rho^{\mathrm{H}'}$, then it is an optimal solution to the continuous knapsack problem in (A.10). The dual of (A.10) associated with $\rho^{\mathrm{H}'}$ is given by:

$$\begin{aligned} \min_{\eta,\xi} \quad & r_{\mathrm{S}}\,\eta + \sum_{i=1}^{n} \xi_i \\ \text{subject to} \quad & \eta + \xi_i \geq p_i \rho_i^{\mathrm{H}'} \qquad \forall\, i \in [\![1, n]\!], \\ & \eta \geq 0 \\ & \xi_i \geq 0 \qquad\qquad\quad\; \forall\, i \in [\![1, n]\!]. \end{aligned} \tag{A.31}$$

Let (η^*, ξ^*) be an optimal solution of the dual problem (A.31). Since (A.14) and (A.30) imply that $0 < \rho_i^{\mathrm{S}^*} < 1$ for every $i \in \mathcal{K}$, then by complementary slackness, $\xi_i^* = 0$ and $\rho_i^{\mathrm{H}'} = \eta^*/p_i$ for every $i \in \mathcal{K}$. Since $\rho^{\mathrm{H}'}$ must fill the knapsack (A.12) entirely, then:

$$r_{\mathrm{H}} = \sum_{i \in \mathcal{I} \cup \mathcal{J} \cup \mathcal{K}} \rho_i^{\mathrm{H}'} = \sum_{j=1}^{i^*} c_j + \sum_{j=1}^{k_{i^*}} c_{\pi^{i^*}(j)} + \eta S^{i^*}_{k_{i^*}+1}.$$

Therefore, $\rho^{\mathrm{H}'}$ satisfies (2.7).

– *Regime Pattern 3: $i^* \geq 1$ and $\tau_{i^*-1} < r_{\mathrm{H}} \leq \nu_{i^*}$.* We additionally consider the following non-edge case assumptions: $p_{i^*-1}c_{i^*-1} < p_{i^*}c_{i^*} < p_{i^*+1}c_{i^*+1}$, $r_{\mathrm{H}} < \sum_{j=1}^{i^*} c_j + \sum_{j=1}^{\ell_{i^*}+1} c_{\pi^{i^*}(j)} + p_{i^*}c_{i^*}S^{i^*}_{\ell_{i^*}+2}$, $r_{\mathrm{S}} < k_{i^*-1} + 1 + p_{\pi^{i^*-1}(k_{i^*-1}+1)} S^{i^*-1}_{k_{i^*-1}+2}$, $p_{\pi^{i^*}(\ell_{i^*})} < p_{\pi^{i^*}(\ell_{i^*}+1)} < 1$, and $p_{\pi^{i^*}(\ell_{i^*}+1)} < p_{\pi^{i^*}(\ell_{i^*}+2)}$ if $\ell_{i^*} \leq n - i^* - 2$.

Let $(\rho^{\mathrm{S}^*}, \rho^{\mathrm{H}^*}) \in \widetilde{\mathcal{A}}_{\mathrm{S}} \times \widetilde{\mathcal{A}}_{\mathrm{H}}$ satisfying (2.8) and (2.9), and let $(\rho^{\mathrm{S}'}, \rho^{\mathrm{H}'}) \in \widetilde{\mathcal{A}}_{\mathrm{S}} \times \widetilde{\mathcal{A}}_{\mathrm{H}}$ be any equilibrium of the game $\widetilde{\Gamma}$. Since $\rho^{\mathrm{S}'}$ is a best response to ρ^{H^*}, it is an optimal solution to the continuous knapsack problem (A.10). The profits of each object are given by (A.26). Under the non-edge case assumptions, $p_1 c_1 \leq \cdots \leq p_{i^*-1} c_{i^*-1} < p_{i^*} c_{i^*} < p_{i^*+1} c_{i^*+1} \leq \cdots \leq p_n c_n$. From (A.27) and the fact that $|\mathcal{J}| + 1 < r_{\mathrm{S}} \leq |\mathcal{J}| + |\mathcal{K}| + 1$, we deduce that any best response to ρ^{H^*} selects all the objects in $\mathcal{J} \cup \{\pi^{i^*}(\ell_{i^*} + 1)\}$, does not select any object in $\mathcal{I} \setminus \{i^*\}$, and fills the knapsack (A.10) entirely. Therefore, $\rho_i^{\mathrm{S}'} = 0$ for every $i \in \mathcal{I} \setminus \{i^*\}$, $\rho_i^{\mathrm{S}'} = 1$ for every $i \in \mathcal{J} \cup \{\pi^{i^*}(\ell_{i^*} + 1)\}$. and $\sum_{i=1}^n \rho_i^{\mathrm{S}'} = r_{\mathrm{S}}$.

Since ρ^{H^*} is a best response to $\rho^{\mathrm{S}'}$, then it is an optimal solution to the continuous knapsack problem (A.12). Under the non-edge case assumptions, $0 < \rho^{\mathrm{H}^*}_{\pi^{i^*}(\ell_{i^*}+1)} < c_{\pi^{i^*}(\ell_{i^*}+1)}$. Thus, at optimality of the dual (A.29), $\alpha^* = 1 - p_{\pi^{i^*}(\ell_{i^*}+1)} \rho^{\mathrm{S}'}_{\pi^{i^*}(\ell_{i^*}+1)} = 1 - p_{\pi^{i^*}(\ell_{i^*}+1)}$. Furthermore, since $0 < \rho_i^{\mathrm{H}^*} < c_i$ for every $i \in \mathcal{K} \setminus \{\pi^{i^*}(\ell_{i^*} + 1)\}$ under the non-edge case assumptions, then $\rho_i^{\mathrm{S}'} = (1 - \alpha^*)/p_i = p_{\pi^{i^*}(\ell_{i^*}+1)}/p_i$. Since $\rho^{\mathrm{H}'}$ must fill the knapsack (A.10) entirely, then:

$$\rho_{i^*}^{\mathrm{S}'} = r_{\mathrm{S}} - \ell_{i^*} - 1 - p_{\pi^{i^*}(\ell_{i^*}+1)} S^{i^*}_{\ell_{i^*}+2} = r_{\mathrm{S}} - \ell_{i^*} - p_{\pi^{i^*}(\ell_{i^*}+1)} S^{i^*}_{\ell_{i^*}+1}.$$

Therefore, $\rho^{\mathrm{S}'}$ satisfies (2.8).

Similarly, $(\rho^{\mathrm{S}^*}, \rho^{\mathrm{H}'})$ is an equilibrium of $\widetilde{\Gamma}$. Then, $\rho^{\mathrm{H}'}$ is a best response to ρ^{S^*} and is an optimal solution to the continuous knapsack problem (A.12). The profits of each object are given by (A.28). Under the non-edge case assumptions, $1 - p_{\pi^{i^*}(1)} \geq \cdots \geq 1 - p_{\pi^{i^*}(\ell_{i^*})} > 1 - p_{\pi^{i^*}(\ell_{i^*}+1)} > 0$. We next show that:

$$r_{\mathrm{S}} - \ell_{i^*} - p_{\pi^{i^*}(\ell_{i^*}+1)} S^{i^*}_{\ell_{i^*}+1} < \min \left\{ \frac{p_{\pi^{i^*}(\ell_{i^*}+1)}}{p_{i^*}}, 1 \right\}. \tag{A.32}$$

Let us assume that (A.32) does not hold and let $j^* \in [\![1, n - i^* + 1]\!]$ satisfying $\pi^{i^*-1}(j^*) = i^*$. If $\ell_{i^*} + 1 \leq j^* - 1$, then (A.23) implies that $r_{\mathrm{S}} = \ell_{i^*} + 1 +$

$p_{\pi^{i^*-1}(\ell_{i^*}+1)} S^{i^*-1}_{\ell_{i^*}+2}$, which contradicts the non-edge case assumption. If on the other hand $\ell_{i^*}+1 \geq j^*$, then $j^* < \ell_{i^*}+2 \leq n-i^*+1$ and $r_{\mathrm{S}} = \ell_{i^*}+2+p_{\pi^{i^*-1}(\ell_{i^*}+2)} S^{i^*-1}_{\ell_{i^*}+3}$, which also contradicts the non-edge case assumptions. Therefore, (A.32) holds.

Since $\sum_{i \in \mathcal{I} \cup \mathcal{J}} c_i < r_{\mathrm{H}}$, then any best response to ρ^{S^*} must select all copies of the objects in $\mathcal{I}$ and $\mathcal{J}$, and must fill the knapsack (A.12) entirely. Therefore, $\rho^{\mathrm{H}'}_i = c_i$ for every $i \in \mathcal{I} \cup \mathcal{J}$ and $\sum_{i=1}^n \rho^{\mathrm{H}'}_i = r_{\mathrm{H}}$. Since ρ^{S^*} is a best response to $\rho^{\mathrm{H}'}$, then it is an optimal solution to the continuous knapsack problem (A.10). Under the non-edge case assumptions, $0 < \rho^{\mathrm{S}^*}_i < 1$ for every $i \in \{i^*\} \cup \mathcal{K} \setminus \{\pi^{i^*}(\ell_{i^*}+1)\}$. Therefore, at optimality of the dual (A.31), $\eta^* = p_{i^*} \rho^{\mathrm{H}'}_{i^*} = p_{i^*} c_{i^*}$, and $\rho^{\mathrm{H}'}_{i^*} = \eta^*/p_i = p_{i^*} c_{i^*}/p_i$ for every $i \in \mathcal{K} \setminus \{\pi^{i^*}(\ell_{i^*}+1)\}$. Finally, since $\rho^{\mathrm{H}'}$ fills the knapsack (A.12) entirely, then:

$$\rho^{\mathrm{H}'}_{\pi^{i^*}(\ell_{i^*}+1)} = r_{\mathrm{H}} - \sum_{j=1}^{i^*} c_j - \sum_{j=1}^{\ell_{i^*}} c_{\pi^{i^*}(j)} - p_{i^*} c_{i^*} S^{i^*}_{\ell_{i^*}+2}.$$

In conclusion, $\rho^{\mathrm{H}'}$ satisfies (2.9).

$\square$

Proof of Proposition 2. In this proof, we allow the capacity vector c and the players' resources r_{S} and r_{H} to be continuous in the game $\widetilde{\Gamma}$. Let Ψ be the set of parameters given by (2.10) for which $\widetilde{\Gamma}$ is nontrivial. First, we note that

$$\Psi' := \{(n, p, c, r_{\mathrm{S}}, r_{\mathrm{H}}) \in \Psi \,:\, p_i < 1 \,\forall\, i \in [\![1, n]\!],\ p_i \neq p_j \text{ and } p_i c_i \neq p_j c_j \,\forall\, i \neq j \in [\![1, n]\!]\}$$

is a dense subset of Ψ. Next, we consider an instantiation of the game parameters $(n, p, c, r_{\mathrm{S}}, r_{\mathrm{H}}) \in \Psi'$. We order the indices such that $p_i c_i < p_{i+1} c_{i+1}$ for every $i \in [\![1, n-1]\!]$. Let $i^* \in [\![0, n-1]\!]$ such that $\tau_{i^*-1} < r_{\mathrm{H}} \leq \tau_{i^*}$.

We first consider the case of Regime Pattern 1, i.e., $i^* = 0$ and $r_{\mathrm{H}} \leq \nu_0$. Proposition 8 implies that all pure equilibria of the game $\widetilde{\Gamma}$ with the parameters $(n, p, c, \hat{r}_{\mathrm{S}}, \hat{r}_{\mathrm{H}}) \in \Psi'$

satisfy the corresponding equilibrium conditions (2.4) and (2.5).

We next consider the case of Regime Pattern 2, i.e., $\nu_{i^*} < r_{\mathrm{H}} \leq \tau_{i^*}$. We then consider new player resources $\hat{r}_{\mathrm{S}} = r_{\mathrm{S}} - \varepsilon$ and $\hat{r}_{\mathrm{H}} = r_{\mathrm{H}} - \varepsilon$ for $\varepsilon > 0$ arbitrarily small. To avoid confusion, we denote the corresponding parameters that depend on $\hat{r}_{\mathrm{S}}$ and $\hat{r}_{\mathrm{H}}$ as $\hat{k}_i$, $\hat{\ell}_i$, $\hat{\tau}_i$, $\hat{\nu}_i$, and $\hat{i}^*$.

By definition of k_{i^*}, and for arbitrarily small ε, we obtain:

$$k_{i^*} + p_{\pi^{i^*}(k_{i^*})} S^{i^*}_{k_{i^*}+1} < \hat{r}_{\mathrm{S}} < r_{\mathrm{S}} \leq k_{i^*} + 1 + p_{\pi^{i^*}(k_{i^*}+1)} S^{i^*}_{k_{i^*}+2}.$$

Thus, $\hat{k}_{i^*} = k_{i^*}$. This implies that $\hat{\tau}_{i^*} = \tau_{i^*}$ and $\hat{\nu}_{i^*} = \nu_{i^*}$. We then deduce the following inequalities for arbitrarily small ε: $\hat{\nu}_{i^*} = \nu_{i^*} < \hat{r}_{\mathrm{H}} < r_{\mathrm{H}} \leq \tau_{i^*} = \hat{\tau}_{i^*}$. Thus, $\hat{i}^* = i^*$.

Since $\nu_{i^*} < \hat{r}_{\mathrm{H}} < \tau_{i^*}$ and $\hat{r}_{\mathrm{S}} < k_{i^*} + 1 + p_{\pi^{i^*}(k_{i^*}+1)} S^{i^*}_{k_{i^*}+2}$, then Proposition 8 implies that all pure equilibria of the game $\widetilde{\Gamma}$ with the parameters $(n, p, c, \hat{r}_{\mathrm{S}}, \hat{r}_{\mathrm{H}})$ for arbitrarily small $\varepsilon > 0$ satisfy the corresponding equilibrium conditions (2.6) and (2.7). Furthermore $(n, p, c, \hat{r}_{\mathrm{S}}, \hat{r}_{\mathrm{H}})$ is arbitrarily close to $(n, p, c, r_{\mathrm{S}}, r_{\mathrm{H}})$.

Finally, we consider the case of Regime Pattern 3, i.e., $i^* \geq 1$ and $\tau_{i^*-1} < r_{\mathrm{H}} \leq \nu_{i^*}$. We then consider new player resources $\hat{r}_{\mathrm{S}} = r_{\mathrm{S}} - \varepsilon$ and $\hat{r}_{\mathrm{H}} = r_{\mathrm{H}} - \varepsilon$ for $\varepsilon > 0$ arbitrarily small. Similarly, we denote the corresponding auxiliary parameters as $\hat{k}_i$, $\hat{\ell}_i$, $\hat{\tau}_i$, $\hat{\nu}_i$, and $\hat{i}^*$.

Using a similar derivation as above, we deduce that for arbitrarily small ε, $\hat{k}_{i^*} = k_{i^*}$ and $\hat{\nu}_{i^*} = \nu_{i^*}$. Then, by definition of k_{i^*-1}, we obtain:

$$k_{i^*-1} + p_{\pi^{i^*-1}(k_{i^*-1})} S^{i^*-1}_{k_{i^*-1}+1} < \hat{r}_{\mathrm{S}} < r_{\mathrm{S}} \leq k_{i^*-1} + 1 + p_{\pi^{i^*-1}(k_{i^*-1}+1)} S^{i^*-1}_{k_{i^*-1}+2}.$$

Thus, $\hat{k}_{i^*-1} = k_{i^*-1}$ and $\hat{\tau}_{i^*-1} = \tau_{i^*-1}$. Then, we obtain that $\hat{i}^* = i^*$ since $\hat{\tau}_{i^*-1} = \tau_{i^*-1} < \hat{r}_{\mathrm{H}} < r_{\mathrm{H}} \leq \nu_{i^*} = \hat{\nu}_{i^*}$. Finally, by definition of ℓ_{i^*}, we obtain:

$$\sum_{j=1}^{i^*} c_j + \sum_{j=1}^{\ell_{i^*}} c_{\pi^{i^*}(j)} + p_{i^*} c_{i^*} S^{i^*}_{\ell_{i^*}+1} < \hat{r}_{\mathrm{H}} < r_{\mathrm{H}} \leq \sum_{j=1}^{i^*} c_j + \sum_{j=1}^{\ell_{i^*}+1} c_{\pi^{i^*}(j)} + p_{i^*} c_{i^*} S^{i^*}_{\ell_{i^*}+2}.$$

Thus, $\hat{\ell}_{i^*} = \ell_{i^*}$. Since $r_{\mathrm{H}} < \sum_{j=1}^{i^*} c_j + \sum_{j=1}^{\ell_{i^*}+1} c_{\pi^{i^*}(j)} + p_{i^*} c_{i^*} S^{i^*}_{\ell_{i^*}+2}$ and $r_{\mathrm{S}} < k_{i^*-1} + 1 + p_{\pi^{i^*-1}(k_{i^*-1}+1)} S^{i^*-1}_{k_{i^*-1}+2}$, then Proposition 8 implies that all pure equilibria of the game $\widetilde{\Gamma}$ with the parameters $(n, p, c, \hat{r}_{\mathrm{S}}, \hat{r}_{\mathrm{H}})$ for arbitrarily small $\varepsilon > 0$ satisfy the corresponding equilibrium conditions (2.8) and (2.9). Furthermore $(n, p, c, \hat{r}_{\mathrm{S}}, \hat{r}_{\mathrm{H}})$ is arbitrarily close to $(n, p, c, r_{\mathrm{S}}, r_{\mathrm{H}})$. $\square$

A.2.2 Proofs of Section 2.4

Before proving Theorem 2, we show that Algorithm 1 is well defined and terminates. We denote as $\kappa^* \in \mathbb{Z}_{\geq 0} \cup \{+\infty\}$ the number of iterations of the while loop (Lines 3–15) in Algorithm 1. In the remainder of this section, we denote by $(\gamma^k)_{k \in [\![1, \kappa^*+1]\!]}$, $(\bar{\rho}^k)_{k \in [\![1, \kappa^*+1]\!]}$, $(\theta^k)_{k \in [\![1, \kappa^*]\!]}$, $(q^k)_{k \in [\![1, \kappa^*]\!]}$, $(\delta^k)_{k \in [\![1, \kappa^*]\!]}$ and $(e^k)_{k \in [\![1, \kappa^*]\!]}$ the iterates of γ, $\bar{\rho}$, θ, q, δ and e respectively generated by Algorithm 1.

Proposition 9. *Each iteration of Algorithm 1 is well defined. In particular,*

$$\forall\, k \in [\![1, \kappa^* + 1]\!],\ \bar{\rho}^k \in [0,1]^n \ \textit{and} \ \sum_{i=1}^{n} \bar{\rho}_i^k \leq \bar{r}, \tag{A.33}$$

$$\forall\, k \in [\![1, \kappa^*]\!],\ q^k \in [\![1, n]\!] \ \textit{and} \ \delta^k \in [0,1). \tag{A.34}$$

Proof of Proposition 9. We show (A.33) and (A.34) by induction. We first consider $k = 1$. By construction, $\bar{\rho}^1 = \rho - \lfloor \rho \rfloor \in [0,1]^n$. Furthermore, by definition of $\widetilde{\mathcal{A}}(b, r)$, we obtain:

$$\bar{r} = r - \sum_{i=1}^{n} \lfloor \rho_i \rfloor = r - \sum_{i=1}^{n} \rho_i + \sum_{i=1}^{n} \bar{\rho}_i^1 \geq \sum_{i=1}^{n} \bar{\rho}_i^1. \tag{A.35}$$

Next, q^1 is constructed when the algorithm initiates the while loop (Lines 3–15), that is, when $\bar{\rho}^1 \notin \{0,1\}^n$. Since $\bar{\rho}^1 \geq \mathbf{0}_n$ and $\bar{r} \in \mathbb{Z}$, then $1 \leq |\{i \in [\![1, n]\!] :\ \bar{\rho}_i^1 > 0\}| \leq n$ and $1 \leq \lceil \sum_{i=1}^{n} \bar{\rho}_i^1 \rceil \overset{\text{(A.35)}}{\leq} \bar{r}$. Therefore, $1 \leq q^1 \leq n$. Finally, δ^1 is well defined since $q^1 \in [\![1, n]\!]$, and $\delta^1 \in [0,1]$ as a consequence of $\bar{\rho}^1 \in [0,1]^n$. We next show by contradiction that $\delta^1 < 1$. Indeed, if $\delta^1 = 1$, then we first deduce that for every $j \in [\![1, q^1]\!]$, $1 \geq \bar{\rho}^1_{\theta^1(j)} \geq \bar{\rho}^1_{\theta^1(q^1)} \geq$

$\delta^1 = 1$. If $q^1 = n$, then this contradicts $\bar{\rho}^1 \notin \{0,1\}^n$. If $q^1 < n$, then, we derive the following inequalities:

$$q^1 = \sum_{j=1}^{q^1} \bar{\rho}^1_{\theta^1(j)} \leq \sum_{i=1}^{n} \bar{\rho}^1_i \overset{\text{(A.35)}}{\leq} \bar{r}, \tag{A.36}$$

$$q^1 \leq \left|\left\{i \in [\![1,n]\!] : \bar{\rho}^1_i = 1\right\}\right| \leq \left|\left\{i \in [\![1,n]\!] : \bar{\rho}^1_i > 0\right\}\right|. \tag{A.37}$$

If $q^1 = \bar{r}$, (resp. $q^1 = |\{i \in [\![1,n]\!] : \bar{\rho}^1_i > 0\}|$) then (A.36) (resp. (A.37)) implies that $\bar{\rho}^1_{\theta^1(j)} = 0$ for every $j \in [\![q^1+1,n]\!]$. This also contradicts $\bar{\rho}^1 \notin \{0,1\}^n$. Thus, $\delta^1 < 1$.

Next, we assume that (A.33) and (A.34) hold for $k \in [\![1,\kappa^*]\!]$. Since $\delta^k < 1$, then we obtain:

$$\forall\, j \in [\![1,q^k]\!],\ 0 \leq \frac{\bar{\rho}^k_{\theta^k(j)} - \bar{\rho}^k_{\theta^k(q^k)}}{1-\delta^k} \leq \frac{\bar{\rho}^k_{\theta^k(j)} - \delta^k}{1-\delta^k} = \bar{\rho}^{k+1}_{\theta^k(j)} \leq \frac{1-\delta^k}{1-\delta^k} = 1$$

and if $q^k < n$, then :

$$\forall\, j \in [\![q^k+1,n]\!],\ 0 \leq \frac{\bar{\rho}^k_{\theta^k(j)}}{1-\delta^k} = \bar{\rho}^{k+1}_{\theta^k(j)} \leq \frac{\bar{\rho}^k_{\theta^k(j)}}{\bar{\rho}^k_{\theta^k(q^k+1)}} \leq 1.$$

Therefore, for every $i \in [\![1,n]\!]$, $\bar{\rho}^{k+1}_i \in [0,1]$. Next, we show that $\sum_{i=1}^n \bar{\rho}^{k+1}_i \leq \bar{r}$:

$$\sum_{i=1}^{n} \bar{\rho}^{k+1}_i = \sum_{j=1}^{q^k} \frac{\bar{\rho}^k_{\theta^k(j)} - \delta^k}{1-\delta^k} + \sum_{j=q^k+1}^{n} \frac{\bar{\rho}^k_{\theta^k(j)}}{1-\delta^k} = \frac{1}{1-\delta^k}\left(\sum_{i=1}^{n} \bar{\rho}^k_i - q^k\delta^k\right).$$

If $q^k = \bar{r}$, then:

$$\sum_{i=1}^{n} \bar{\rho}^{k+1}_i \leq \frac{1}{1-\delta^k}\left(\bar{r} - \bar{r}\delta^k\right) = \bar{r}.$$

If on the other hand $q^k = \left|\left\{i \in [\![1,n]\!] :\ \bar{\rho}_i^k > 0\right\}\right|$, then $\bar{\rho}^k \in [0,1]^n$ implies that:

$$\begin{aligned}\sum_{i=1}^{n} \bar{\rho}_i^{k+1} &= \frac{1}{1-\delta^k}\left(\sum_{i=1}^{n} \bar{\rho}_i^k - \left|\left\{i \in [\![1,n]\!] :\ \bar{\rho}_i^k > 0\right\}\right| \delta^k\right) \\ &\leq \frac{1}{1-\delta^k}\left(\sum_{i=1}^{n} \bar{\rho}_i^k - \delta^k \sum_{\{i \in [\![1,n]\!] :\, \bar{\rho}_i^k > 0\}} \bar{\rho}_i^k\right) = \sum_{i=1}^{n} \bar{\rho}_i^k \leq \bar{r}.\end{aligned}$$

Therefore, $\sum_{i=1}^{n} \bar{\rho}_i^{k+1} \leq \bar{r}$. Since $\bar{\rho}^{k+1} \in [0,1]^n$, then the same argument as the one derived for $k = 1$ can be applied to conclude that if $k < \kappa^*$ and $\bar{\rho}^{k+1} \notin \{0,1\}^n$, then $q^{k+1} \in [\![1,n]\!]$ and $\delta^{k+1} \in [0,1)$. In conclusion, (A.33) and (A.34) hold by induction. $\square$

Proposition 10. *Algorithm 1 terminates after $\kappa^* \leq n$ iterations of the while loop (Lines 3–15). In particular, for every $k \in [\![1,\kappa^*]\!]$, $\delta^k > 0$, and*

$$\left|\left\{i \in [\![1,n]\!] :\ \bar{\rho}_i^{k+1} \in \{0,1\}\right\}\right| > \left|\left\{i \in [\![1,n]\!] :\ \bar{\rho}_i^{k} \in \{0,1\}\right\}\right|.$$

Proof of Proposition 10. Let $k \in [\![1,\kappa^*]\!]$. First, we show that $\delta^k > 0$. From the inequality $q^k \leq \left|\left\{i \in [\![1,n]\!] :\ \bar{\rho}_i^k > 0\right\}\right|$, it follows that $\bar{\rho}^k_{\theta^k(q^k)} > 0$. Next, we show by contradiction that if $q^k < n$, then $\bar{\rho}^k_{\theta^k(q^k+1)} < 1$. Indeed, i f instead $q^k < n$ and $\bar{\rho}^k_{\theta^k(q^k+1)} = 1$, then we first deduce that $q^k < \left|\left\{i \in [\![1,n]\!] :\ \bar{\rho}_i^k > 0\right\}\right|$. Furthermore,

$$\bar{r} \overset{\text{(A.33)}}{\geq} \sum_{i=1}^{n} \bar{\rho}_i^k \geq \sum_{j=1}^{q^k+1} \bar{\rho}^k_{\theta^k(j)} = q^k + 1 > q^k.$$

This contradicts the definition of q^k. Therefore if $q^k < n$, then $\bar{\rho}^k_{\theta^k(q^k+1)} < 1$, which in turn implies that $\delta^k > 0$.

We now show that $\left|\left\{i \in [\![1,n]\!] :\ \bar{\rho}_i^{k+1} \in \{0,1\}\right\}\right| > \left|\left\{i \in [\![1,n]\!] :\ \bar{\rho}_i^{k} \in \{0,1\}\right\}\right|$. Let $j' \in [\![1,n]\!]$ be such that $\bar{\rho}^k_{\theta^k(j')} = 0$. Necessarily, $j' > \left|\left\{i \in [\![1,n]\!] :\ \bar{\rho}_i^k > 0\right\}\right| \geq q^k$, which implies that $\bar{\rho}^{k+1}_{\theta^k(j')} = \bar{\rho}^k_{\theta^k(j')}/(1-\delta^k) = 0$.

Next, we consider $j' \in [\![1,n]\!]$ such that $\bar{\rho}^k_{\theta^k(j')} = 1$. Then, $j' \leq q^k$, as implied by the following inequalities:

$$
\begin{aligned}
j' &= \sum_{j=1}^{j'} \bar{\rho}^k_{\theta^k(j)} \leq \sum_{i=1}^{n} \bar{\rho}^k_i \overset{\text{(A.33)}}{\leq} \bar{r}, \\
j' &\leq \left|\left\{i \in [\![1,n]\!] : \bar{\rho}^k_i = 1\right\}\right| \leq \left|\left\{i \in [\![1,n]\!] : \bar{\rho}^k_i > 0\right\}\right|.
\end{aligned}
$$

Thus, $j' \leq q^k$ and $\bar{\rho}^{k+1}_{\theta^k(j')} = (\bar{\rho}^k_{\theta^k(j')} - \delta^k)/(1-\delta^k) = 1$. This shows that the inequality $\left|\left\{i \in [\![1,n]\!] : \bar{\rho}^{k+1}_i \in \{0,1\}\right\}\right| \geq \left|\left\{i \in [\![1,n]\!] : \bar{\rho}^k_i \in \{0,1\}\right\}\right|$ holds. To ensure a strict inequality, we must show that one fractional component of $\bar{\rho}^k$ becomes 0 or 1 in $\bar{\rho}^{k+1}$.

We know that $0 < \delta^k < 1$. If $\delta^k = \bar{\rho}^k_{\theta^k(q^k)}$, then $\bar{\rho}^{k+1}_{\theta^k(q^k)} = 0$. If $q^k < n$ and $\delta^k = 1 - \bar{\rho}^k_{\theta^k(q^k+1)}$, then $\bar{\rho}^{k+1}_{\theta^k(q^k+1)} = 1$. In both cases, a fractional component of $\bar{\rho}^k$ becomes 0 or 1 in $\bar{\rho}^{k+1}$. In conclusion $\left|\left\{i \in [\![1,n]\!] : \bar{\rho}^{k+1}_i \in \{0,1\}\right\}\right| > \left|\left\{i \in [\![1,n]\!] : \bar{\rho}^k_i \in \{0,1\}\right\}\right|$.

Since for every $k \in [\![1,\kappa^*+1]\!]$, $\left|\left\{i \in [\![1,n]\!] : \bar{\rho}^k_i \in \{0,1\}\right\}\right| \leq n$, then the algorithm must terminate after at most n iterations of the while loop (Lines 3–15). Therefore $\kappa^* \leq n$. □

Now that we proved that Algorithm 1 is well defined and terminates, we can show Theorem 2.

Proof of Theorem 2. Consider a capacity vector $b \in \mathbb{Z}^n_{>0}$, a resource budget $r \in \mathbb{Z}_{>0}$, and an expected resource allocation vector $\rho \in \widetilde{\mathcal{A}}(b,r)$. For convenience, we denote $e^{\kappa^*+1} := \bar{\rho}^{\kappa^*+1} \in \{0,1\}^n$. We first show that for every $k \in [\![1,\kappa^*+1]\!]$, $\lfloor \rho \rfloor + e^k \in \mathcal{A}(b,r)$.

In Proposition 10, we showed that for every $k \in [\![1,\kappa^*]\!]$, if a component $i \in [\![1,n]\!]$ satisfies $\bar{\rho}^k_i = 0$, then $\bar{\rho}^{k+1}_i = 0$. Thus, for every $k \in [\![1,\kappa^*+1]\!]$, $\bar{\rho}^k_i > 0$ only if $\bar{\rho}^1_i > 0$. By definition of $\widetilde{\mathcal{A}}(b,r)$ and since $b_i \in \mathbb{Z}$, we deduce that if $\bar{\rho}^1_i > 0$, then $b_i \geq \lceil \rho_i \rceil =$

$\lfloor \rho_i \rfloor + \lceil \bar{\rho}_i^1 \rceil = \lfloor \rho_i \rfloor + 1$. Furthermore,

$$\forall\, k \in [\![1, \kappa^*]\!],\ \sum_{i=1}^{n} (\lfloor \rho_i \rfloor + e_i^k) = q^k + \sum_{i=1}^{n} \lfloor \rho_i \rfloor \leq \bar{r} + \sum_{i=1}^{n} \lfloor \rho_i \rfloor = r,$$
$$\text{and } \sum_{i=1}^{n} (\lfloor \rho_i \rfloor + e_i^{\kappa^*+1}) = \sum_{i=1}^{n} \bar{\rho}_i^{\kappa^*+1} + \sum_{i=1}^{n} \lfloor \rho_i \rfloor \overset{\text{(A.33)}}{\leq} \bar{r} + \sum_{i=1}^{n} \lfloor \rho_i \rfloor = r.$$

Therefore, for every $k \in [\![1, \kappa^* + 1]\!]$, $\lfloor \rho \rfloor + e^k \in \mathcal{A}(b, r)$.

Next, we show that σ returned by the algorithm is a probability distribution. We first note that for every $k \in [\![1, \kappa^* + 1]\!]$, $\gamma^k \geq 0$. Furthermore,

$$\sum_{z \in \mathcal{A}(b,r)} \sigma_z = \gamma^{\kappa^*+1} + \sum_{k=1}^{\kappa^*} \gamma^k \delta^k = \gamma^{\kappa^*+1} + \sum_{k=1}^{\kappa^*} (\gamma^k - \gamma^{k+1}) = \gamma^{\kappa^*+1} + \gamma^1 - \gamma^{\kappa^*+1} = 1.$$

Therefore, $\sigma \in \Delta(b, r)$. We now show that σ returned by the algorithm is consistent with the vector ρ. To this end, we note the following equality:

$$\forall\, k \in [\![1, \kappa^*]\!],\ \forall\, i \in [\![1, n]\!],\ \gamma^k \bar{\rho}_i^k - \gamma^{k+1} \bar{\rho}_i^{k+1} = \gamma^k (\bar{\rho}_i^k - (1 - \delta^k) \bar{\rho}_i^{k+1}) = \gamma^k \delta^k e_i^k.$$

Then, we obtain:

$$\begin{aligned}\forall\, i \in [\![1, n]\!],\ \sum_{z \in \mathcal{A}(b,r)} z_i \sigma_z &= \gamma^{\kappa^*+1} (\lfloor \rho_i \rfloor + \bar{\rho}_i^{\kappa^*+1}) + \sum_{k=1}^{\kappa^*} \gamma^k \delta^k (\lfloor \rho_i \rfloor + e_i^k) \\ &= \lfloor \rho_i \rfloor \sum_{z \in \mathcal{A}(b,r)} \sigma_z + \gamma^{\kappa^*+1} \bar{\rho}_i^{\kappa^*+1} + \gamma^1 \bar{\rho}_i^1 - \gamma^{\kappa^*+1} \bar{\rho}_i^{\kappa^*+1} \\ &= \lfloor \rho_i \rfloor + \bar{\rho}_i^1 = \rho_i.\end{aligned}$$

Thus, σ returned by Algorithm 1 is consistent with the vector ρ.

Since $\kappa^* \leq n$, then the support of σ is of size at most $n + 1$. Finally, we argue that Algorithm 1 runs in time $O(n^2)$. Indeed, the first iteration of the while loop (Lines 3–15) can be implemented in time $O(n \log n)$ by using an efficient sorting algorithm (e.g. merge sort) to sort $\bar{\rho}^1$ and create the permutation θ^1. Fortunately, the subsequent iterations can be implemented in time $O(n)$. Indeed, we note that for every $k \in [\![1, \kappa^* - 1]\!]$, $\bar{\rho}^{k+1}_{\theta^k(1)} \geq \cdots \geq \bar{\rho}^{k+1}_{\theta^k(q^k)}$ and $\bar{\rho}^{k+1}_{\theta^k(q^k+1)} \geq \cdots \geq \bar{\rho}^{k+1}_{\theta^k(n)}$. Therefore, we can sort $\bar{\rho}^{k+1}$ and create the permutation θ^{k+1} by merging and sorting the lists $(\bar{\rho}^{k+1}_{\theta^k(1)}, \ldots, \bar{\rho}^{k+1}_{\theta^k(q^k)})$ and $(\bar{\rho}^{k+1}_{\theta^k(q^k+1)}, \ldots, \bar{\rho}^{k+1}_{\theta^k(n)})$ that are already sorted. This operation can be carried out in time $O(n)$. Since the number of iterations of the while loop (Lines 3–15) is upper bounded by n, then the overall running time of Algorithm 1 is $O(n^2)$. $\square$

APPENDIX B

SUPPLEMENT TO CHAPTER 3

B.1 Proofs of Statements

Proof of Lemma 5. We prove each statement.

- We first show that the index m^* is well defined, i.e., that there exists $m \in [\![1, n - r_\mathrm{S}]\!]$ satisfying $|\mathcal{J}_m(k_m)| \leq r_\mathrm{S}$. To this end, we note that the definition of $\mathcal{J}_{n-r_\mathrm{S}}(k)$ implies $\mathcal{J}_{n-r_\mathrm{S}}(k) \subseteq \mathcal{T} \setminus \mathcal{I}_{n-r_\mathrm{S}}$. Then, it follows that $|\mathcal{J}_{n-r_\mathrm{S}}(k)| \leq |\mathcal{T} \setminus \mathcal{I}_{n-r_\mathrm{S}}| = r_\mathrm{S}$ for every $k \in [\![1, r_\mathrm{S}]\!]$. Since the definition of k_{n-r_S} states that $k_{n-r_\mathrm{S}} \in [\![1, r_\mathrm{S}]\!]$, we deduce that $|\mathcal{J}_{n-r_\mathrm{S}}(k_{n-r_\mathrm{S}})| \leq r_\mathrm{S}$.

- We next show that $\mathcal{J}_i(k_i) \neq \emptyset$ for every $i \in [\![0, m^* - 1]\!]$. We begin by showing that $\mathcal{J}_0(k_0) \neq \emptyset$. By definition, $k_0 \in [\![1, n]\!]$. In particular, $k_0 \geq 1$. Then, $\mathcal{J}_0(k_0) = \left\{t_{\pi^0(1)}, \ldots, t_{\pi^0(k_0)}\right\} \neq \emptyset$. Next, for $i \in [\![1, m^* - 1]\!]$, from the definition of m^*, it follows that $|\mathcal{J}_i(k_i)| > r_\mathrm{S}$. Thus, $\mathcal{J}_i(k_i) \neq \emptyset$.

- Finally, we argue that $0 =: U_\mathrm{S}(t_0) < U_\mathrm{S}(t_1) \leq \cdots \leq U_\mathrm{S}(t_{m^*})$. By definition, we have $U_\mathrm{S}(t_0) := 0$. Also by definition, for every $i \in [\![0, m^* - 1]\!]$, we have $t_{i+1} \arg\min_{t \in \mathcal{J}_i(k_i)} U_\mathrm{S}(t)$. In particular, $t_{i+1} \in \mathcal{J}_i(k_i)$. From the condition in the definition of $\mathcal{J}_i(k_i)$, it follows that $U_\mathrm{S}(t_{i+1}) \geq U_\mathrm{S}(t_i) > 0$, from which we deduce the sequence of inequalities.

$\square$

We note that when some utilities are identical, labelings π^i ordering $\mathcal{T} \setminus \mathcal{I}_i$ by their utilities for H may not be unique. To simplify our proofs, we assume without loss of generality that π^i maintains the order between identical valuations, i.e., $\pi^i(j) < \pi^i(k)$

when $1 \leq j < k \leq n - i$ and $U_{\mathrm{H}}(t_{\pi^i(j)}) = U_{\mathrm{H}}(t_{\pi^i(k)})$, thus rendering π^i unique for every $i \in [\![0, m^*]\!]$.

Before proving Lemma 7, we need the following auxiliary lemmas.

Lemma 16. *Let $i \in [\![1, m^*]\!]$ and $j^* \in [\![1, n-i+1]\!]$ be such that $t_{\pi^{i-1}(j^*)} = t_i$. Then:*

$$t_{\pi^{i-1}(j)} = \begin{cases} t_{\pi^i(j)} & \text{if } j \in [\![1, j^*-1]\!], \\ t_i & \text{if } j = j^*, \\ t_{\pi^i(j-1)} & \text{if } j \in [\![j^*+1, n-i+1]\!]. \end{cases} \tag{B.1}$$

Moreover, for every $k \in [\![0, n-i+1]\!]$,

$$\left\{t_{\pi^{i-1}(1)}, \ldots, t_{\pi^{i-1}(k)}\right\} = \begin{cases} \left\{t_{\pi^i(1)}, \ldots, t_{\pi^i(k)}\right\} & \text{if } k \in [\![0, j^*-1]\!], \\ \left\{t_{\pi^i(1)}, \ldots, t_{\pi^i(k-1)}\right\} \cup \{t_i\} & \text{if } k \in [\![j^*, n-i+1]\!]. \end{cases} \tag{B.2}$$

Proof of Lemma 16. Let $i \in [\![1, m^*]\!]$ and $j^* \in [\![1, n-i+1]\!]$ be such that $t_{\pi^{i-1}(j^*)} = t_i$. The permutation π^{i-1} sorts the locations in $\mathcal{T} \setminus \mathcal{I}_{i-1}$ in order of nonincreasing H's utilities. After removing $t_i = t_{\pi^{i-1}(j^*)}$ from the chain of inequalities, we obtain $U_{\mathrm{H}}(t_{\pi^{i-1}(1)}) \geq \cdots \geq U_{\mathrm{H}}(t_{\pi^{i-1}(j^*-1)}) \geq U_{\mathrm{H}}(t_{\pi^{i-1}(j^*+1)}) \geq \cdots \geq U_{\mathrm{H}}(t_{\pi^{i-1}(n-i+1)})$, which sorts $\mathcal{T} \setminus \mathcal{I}_i$ by nonincreasing H's utilities, thus providing (B.1). Consequently, for every $k \in [\![0, j^*-1]\!]$, we have

$$\left\{t_{\pi^{i-1}(1)}, \ldots, t_{\pi^{i-1}(k)}\right\} = \left\{t_{\pi^i(1)}, \ldots, t_{\pi^i(k)}\right\}.$$

Similarly, for every $k \in [\![j^*, n-i+1]\!]$,

$$\left\{t_{\pi^{i-1}(1)}, \ldots, t_{\pi^{i-1}(j^*-1)}, t_i, t_{\pi^{i-1}(j^*+1)}, \ldots, t_{\pi^{i-1}(k)}\right\} = \left\{t_{\pi^i(1)}, \ldots, t_{\pi^i(k-1)}\right\} \cup \{t_i\}.$$

$\square$

Lemma 17. *For every $i \in [\![0, m^*]\!]$, the set*

$$K_i := \left\{ k \in [\![1, n-i]\!] : \sum_{t \in \mathcal{J}_i(k)} \left(1 - \frac{U_{\mathrm{H}}(t_{\pi^i(k)})}{U_{\mathrm{H}}(t)}\right) < \ r_{\mathrm{S}} \right\}$$

is nonempty, and therefore the parameter k_i is well defined. Furthermore, the function $f(k) := \sum_{t \in \mathcal{J}_i(k)} \left(1 - U_{\mathrm{H}}(t_{\pi^i(k)})/U_{\mathrm{H}}(t)\right)$ is nondecreasing over $[\![1, n-i]\!]$.

Proof of Lemma 17. First, we note that the equality $\sum_{t \in \mathcal{J}_i(1)} \left(1 - U_{\mathrm{H}}(t_{\pi^i(1)})/U_{\mathrm{H}}(t)\right) = 0$ holds as either $\mathcal{J}_i(1) = \left\{t_{\pi^i(1)}\right\}$ or $\mathcal{J}_i(1) = \emptyset$. Therefore, $1 \in K_i$ and $k_i = \max K_i$ is well defined.

Next, let $k \in [\![1, n-i-1]\!]$. Since either $\mathcal{J}_i(k+1) = \mathcal{J}_i(k) \cup \left\{t_{\pi^i(k+1)}\right\}$ or $\mathcal{J}_i(k+1) = \mathcal{J}_i(k)$, we have $f(k+1) = \sum_{t \in \mathcal{J}_i(k)} \left(1 - U_{\mathrm{H}}(t_{\pi^i(k+1)})/U_{\mathrm{H}}(t)\right)$. Then, using that $U_{\mathrm{H}}(t_{\pi^i(k)}) \geq U_{\mathrm{H}}(t_{\pi^i(k+1)})$, we have

$$\begin{aligned} f(k+1) = \sum_{t \in \mathcal{J}_i(k+1)} \left(1 - \frac{U_{\mathrm{H}}(t_{\pi^i(k+1)})}{U_{\mathrm{H}}(t)}\right) &= \sum_{t \in \mathcal{J}_i(k)} \left(1 - \frac{U_{\mathrm{H}}(t_{\pi^i(k+1)})}{U_{\mathrm{H}}(t)}\right) \\ &\geq \sum_{t \in \mathcal{J}_i(k)} \left(1 - \frac{U_{\mathrm{H}}(t_{\pi^i(k)})}{U_{\mathrm{H}}(t)}\right) = f(k). \end{aligned}$$

Hence, f is nondecreasing over $[\![1, n-i]\!]$. $\square$

Lemma 18. *The following statements hold for every $i \in [\![1, m^*]\!]$:*

- $\mathcal{J}_{i-1}(k_{i-1}) = \mathcal{J}_i(k_{i-1} - 1) \cup \{t_i\}$
- $\widetilde{\mathcal{J}}_{i-1}(k_{i-1}) = \widetilde{\mathcal{J}}_i(k_{i-1} - 1)$.
- $k_i \geq k_{i-1} - 1$.

Proof of Lemma 18. We prove each statement.

- We first show that $\mathcal{J}_{i-1}(k_{i-1}) = \mathcal{J}_i(k_{i-1} - 1) \cup \{t_i\}$. Let $j^* \in [\![1, n-i+1]\!]$ be such that $t_{\pi^{i-1}(j^*)} = t_i$. Since $t_i \in \mathcal{J}_{i-1}(k_{i-1}) \subseteq \left\{t_{\pi^{i-1}(1)}, \ldots, t_{\pi^{i-1}(k_{i-1})}\right\}$, we must have $j^* \in [\![1, k_{i-1}]\!]$, or equivalently, $k_{i-1} \in [\![j^*, n-i+1]\!]$.

Let $t \in \mathcal{J}_{i-1}(k_{i-1})$. Since $k_{i-1} \in [\![j^*, n-i+1]\!]$, from (B.2) we have $t \in \mathcal{J}_{i-1}(k_{i-1}) \subseteq \left\{t_{\pi^{i-1}(1)}, \ldots, t_{\pi^{i-1}(k_{i-1})}\right\} = \left\{t_{\pi^{i}(1)}, \ldots, t_{\pi^{i}(k_{i-1}-1)}\right\} \cup \{t_i\}$. Furthermore, from the definition of t_i, we have $U_{\mathrm{S}}(t) \geq U_{\mathrm{S}}(t_i)$ for every $t \in \mathcal{J}_{i-1}(k_{i-1})$. Therefore, $t \in \mathcal{J}_i(k_{i-1}-1) \cup \{t_i\}$.

Conversely, let $t \in \mathcal{J}_i(k_{i-1}-1) \cup \{t_i\}$. Using (B.2) again, we get $t \in \mathcal{J}_i(k_{i-1}-1) \cup \{t_i\} \subseteq \left\{t_{\pi^{i}(1)}, \ldots, t_{\pi^{i}(k_{i-1}-1)}\right\} \cup \{t_i\} = \left\{t_{\pi^{i-1}(1)}, \ldots, t_{\pi^{i-1}(k_{i-1})}\right\}$. Moreover, from the definition of $\mathcal{J}_i(k_{i-1}-1)$, we have $U_{\mathrm{S}}(t) \geq U_{\mathrm{S}}(t_i)$, and by Lemma 5, $U_{\mathrm{S}}(t_i) \geq U_{\mathrm{S}}(t_{i-1})$ as well. Hence, $U_{\mathrm{S}}(t) \geq U_{\mathrm{S}}(t_{i-1})$, and thus $t \in \mathcal{J}_{i-1}(k_{i-1})$.

– We next show that $\widetilde{\mathcal{J}}_{i-1}(k_{i-1}) = \widetilde{\mathcal{J}}_i(k_{i-1}-1)$. Using that $\mathcal{J}_{i-1}(k_{i-1}) = \mathcal{J}_i(k_{i-1}-1) \cup \{t_i\}$ and the definitions of $\widetilde{\mathcal{J}}_{i-1}(k_{i-1})$ and $\widetilde{\mathcal{J}}_i(k_{i-1}-1)$, it follows that

$$\begin{aligned}
\widetilde{\mathcal{J}}_{i-1}(k_{i-1}) &= \left\{t_{\pi^{i-1}(1)}, \ldots, t_{\pi^{i-1}(k_{i-1})}\right\} \setminus \mathcal{J}_{i-1}(k_{i-1}) \\
&\overset{\text{(B.2)}}{=} \left(\left\{t_{\pi^{i}(1)}, \ldots, t_{\pi^{i}(k_{i-1}-1)}\right\} \cup \{t_i\}\right) \setminus \mathcal{J}_{i-1}(k_{i-1}) \\
&= \left(\left\{t_{\pi^{i}(1)}, \ldots, t_{\pi^{i}(k_{i-1}-1)}\right\} \cup \{t_i\}\right) \setminus (\mathcal{J}_i(k_{i-1}-1) \cup \{t_i\}) \\
&= \left\{t_{\pi^{i}(1)}, \ldots, t_{\pi^{i}(k_{i-1}-1)}\right\} \setminus \mathcal{J}_i(k_{i-1}-1) \\
&= \widetilde{\mathcal{J}}_i(k_{i-1}-1).
\end{aligned}$$

– Finally, we show that $k_i \geq k_{i-1}-1$. If $k_{i-1} \leq 2$, there is nothing to show, as $k_i \geq 1$ by definition. Thus, we assume that $k_{i-1} \in [\![3, n-i+1]\!]$. Let $j^* \in [\![1, n-i+1]\!]$ be such that $t_{\pi^{i-1}(j^*)} = t_i$. We already argued that $k_{i-1} \in [\![j^*, n-i+1]\!]$. Then, we consider two cases. First, if $k_{i-1} \in [\![j^*+1, n-i+1]\!]$, from the definition of k_{i-1}, it follows that

$$
\begin{aligned}
r_{\mathrm{S}} > \sum_{t \in \mathcal{J}_{i-1}(k_{i-1})} \left(1 - \frac{U_{\mathrm{H}}(t_{\pi^{i-1}(k_{i-1})})}{U_{\mathrm{H}}(t)}\right) &= \sum_{t \in \mathcal{J}_i(k_{i-1}-1) \cup \{t_i\}} \left(1 - \frac{U_{\mathrm{H}}(t_{\pi^{i-1}(k_{i-1})})}{U_{\mathrm{H}}(t)}\right) \\
&\overset{(\mathrm{B.1})}{=} \sum_{t \in \mathcal{J}_i(k_{i-1}-1) \cup \{t_i\}} \left(1 - \frac{U_{\mathrm{H}}(t_{\pi^{i}(k_{i-1}-1)})}{U_{\mathrm{H}}(t)}\right) \\
&\geq \sum_{t \in \mathcal{J}_i(k_{i-1}-1)} \left(1 - \frac{U_{\mathrm{H}}(t_{\pi^{i}(k_{i-1}-1)})}{U_{\mathrm{H}}(t)}\right).
\end{aligned}
$$

Since $1 \leq k_{i-1} - 1 \leq n - i$, then $k_{i-1} - 1 \in K_i$ and $k_{i-1} - 1 \leq k_i$ follows by definition of k_i.

Second, we consider the case $k_{i-1} = j^*$. Using a similar approach, we have

$$
\begin{aligned}
r_{\mathrm{S}} > \sum_{t \in \mathcal{J}_{i-1}(k_{i-1})} \left(1 - \frac{U_{\mathrm{H}}(t_{\pi^{i-1}(k_{i-1})})}{U_{\mathrm{H}}(t)}\right) &= \sum_{t \in \mathcal{J}_i(k_{i-1}-1) \cup \{t_i\}} \left(1 - \frac{U_{\mathrm{H}}(t_i)}{U_{\mathrm{H}}(t)}\right) \\
&= \sum_{t \in \mathcal{J}_i(k_{i-1}-1)} \left(1 - \frac{U_{\mathrm{H}}(t_i)}{U_{\mathrm{H}}(t)}\right) \\
&\geq \sum_{t \in \mathcal{J}_i(k_{i-1}-1)} \left(1 - \frac{U_{\mathrm{H}}(t_{\pi^{i}(k_{i-1}-1)})}{U_{\mathrm{H}}(t)}\right),
\end{aligned}
$$

where the last inequality follows from

$$
U_{\mathrm{H}}(t_{\pi^{i}(k_{i-1}-1)}) \overset{(\mathrm{B.1})}{=} U_{\mathrm{H}}(t_{\pi^{i-1}(k_{i-1}-1)}) \geq U_{\mathrm{H}}(t_{\pi^{i-1}(k_{i-1})}) = U_{\mathrm{H}}(t_{\pi^{i-1}(j^*)}) = U_{\mathrm{H}}(t_i).
$$

Thus, $k_{i-1} - 1 \in K_i$ again, and therefore $k_{i-1} - 1 \leq k_i$ by definition of k_i.

□

Lemma 19. *For every $i \in [\![0, m^*]\!]$, if $r_{\mathrm{H}} > \tau_{i-1}$, then the set*

$$
L_i := \left\{ \ell \in [\![0, n-i]\!] : |\mathcal{I}_i| + |\widetilde{\mathcal{J}}_i(\ell)| + \sum_{t \in \mathcal{J}_i(\ell)} \frac{U_{\mathrm{S}}(t_i)}{U_{\mathrm{S}}(t)} < r_{\mathrm{H}} \right\}
$$

is nonempty, and therefore the parameter ℓ_i is well defined. Furthermore, the function $g(\ell) \coloneqq |\mathcal{I}_i| + |\widetilde{\mathcal{J}}_i(\ell)| + \sum_{t \in \mathcal{J}_i(\ell)} U_{\mathrm{S}}(t_i)/U_{\mathrm{S}}(t)$ is nondecreasing over $[\![0, n-i]\!]$.

Proof of Lemma 19. If $r_{\mathrm{H}} > \tau_{i-1}$, we have

$$\begin{aligned} r_{\mathrm{H}} > \tau_{i-1} &= |\mathcal{I}_{i-1}| + |\widetilde{\mathcal{J}}_{i-1}(k_{i-1})| + \sum_{t \in \mathcal{J}_{i-1}(k_{i-1})} \frac{U_{\mathrm{S}}(t_i)}{U_{\mathrm{S}}(t)} \\ &\overset{\text{Lemma 18}}{=} |\mathcal{I}_{i-1}| + |\widetilde{\mathcal{J}}_i(k_{i-1}-1)| + \sum_{t \in \mathcal{J}_i(k_{i-1}-1) \cup \{t_i\}} \frac{U_{\mathrm{S}}(t_i)}{U_{\mathrm{S}}(t)} \\ &= |\mathcal{I}_i| + |\widetilde{\mathcal{J}}_i(k_{i-1}-1)| + \sum_{t \in \mathcal{J}_i(k_{i-1}-1)} \frac{U_{\mathrm{S}}(t_i)}{U_{\mathrm{S}}(t)}. \end{aligned}$$

Moreover, $1 \leq k_{i-1} \leq n-i$ implies $0 \leq k_{i-1}-1 \leq n-i-1 \leq n-i$. Thus, $k_{i-1}-1 \in L_i$. Therefore, $\ell_i = \max\ L_i$ is well defined.

Next, let $\ell \in [\![0, n-i-1]\!]$. Then, by considering the cases $t_{\pi^i(\ell+1)} \in \mathcal{J}_i(\ell+1)$ and $t_{\pi^i(\ell+1)} \in \widetilde{\mathcal{J}}_i(\ell+1)$, it follows that

$$\begin{aligned} g(\ell+1) &= |\mathcal{I}_i| + |\widetilde{\mathcal{J}}_i(\ell+1)| + \sum_{t \in \mathcal{J}_i(\ell+1)} \frac{U_{\mathrm{S}}(t_i)}{U_{\mathrm{S}}(t)} \\ &= |\mathcal{I}_i| + |\widetilde{\mathcal{J}}_i(\ell)| + \sum_{t \in \mathcal{J}_i(\ell)} \frac{U_{\mathrm{S}}(t_i)}{U_{\mathrm{S}}(t)} + \min\left\{ \frac{U_{\mathrm{S}}(t_i)}{U_{\mathrm{S}}(t_{\pi^i(\ell+1)})}, 1 \right\} \geq g(\ell). \end{aligned}$$

Hence, g is nondecreasing over $[\![0, n-i]\!]$. □

Lemma 20. *Let $i \in [\![1, m^*]\!]$, and suppose that $r_{\mathrm{H}} > \tau_{i-1}$. Then, the following statements hold:*

- $\ell_i \geq k_{i-1} - 1$.
- $r_{\mathrm{H}} \leq \nu_i$ *if and only if* $k_i > \ell_i$.

Proof of Lemma 20. Suppose that $r_{\mathrm{H}} > \tau_{i-1}$. We prove each statement.

- In proof of Lemma 19, we deduced that $r_{\mathrm{H}} > \tau_{i-1}$ implies that $k_{i-1} - 1 \in L_i$. Therefore, from the definition of ℓ_i, we conclude that $k_{i-1} - 1 \leq \ell_i$.

– From the definition of ν_i, the inequality $\nu_i < r_{\mathrm{H}}$ is equivalent to

$$|\mathcal{I}_i| + |\widetilde{\mathcal{J}}_i(k_i)| + \sum_{t \in \mathcal{J}_i(k_i)} \frac{U_{\mathrm{S}}(t_i)}{U_{\mathrm{S}}(t)} < r_{\mathrm{H}}. \tag{B.3}$$

Then, $k_i \in L_i$, and thus $k_i \leq \ell_i$ follows from the definition of L_i. Conversely, suppose that $r_{\mathrm{H}} \leq \nu_i$, and let us assume by contradiction that $k_i \leq \ell_i$. By reversing inequality (B.3), we deduce $k_i \notin L_i$. Furthermore, from Lemma 19, the function $g(\ell) \coloneqq |\mathcal{I}_i| + |\widetilde{\mathcal{J}}_i(\ell)| + \sum_{t \in \mathcal{J}_i(\ell)} U_{\mathrm{S}}(t_i)/U_{\mathrm{S}}(t)$ is nondecreasing over $[\![0, n-i]\!]$. Then, $\nu_i = g(k_i) \leq g(\ell_i) < r_{\mathrm{H}}$, which is a contradiction. Hence, $k_i > \ell_i$.

□

We are now ready to prove Lemma 7.

Proof of Lemma 7. First, we note that since $\mathcal{I}_0 = \emptyset$, $\widetilde{\mathcal{J}}_0(k_0) = \emptyset$ and $U_{\mathrm{S}}(t_0) = 0$, we have $\nu_0 = |\mathcal{I}_0| + |\widetilde{\mathcal{J}}_0(k_0)| + \sum_{t \in \mathcal{J}_0(k_0)} U_{\mathrm{S}}(t_0)/U_{\mathrm{S}}(t) = 0$. Thus, $\tau_{-1} = \nu_0 = 0$. Also, $\tau_{m^*} = n$ holds by definition of τ_{m^*}.

We next show that $\tau_{i-1} \leq \nu_i \leq \tau_i$ for every $i \in [\![0, m^*]\!]$. From the definitions of ν_i and τ_i, as well as the inequality $U_{\mathrm{S}}(t_i) \leq U_{\mathrm{S}}(t_{i+1})$ (Lemma 5), we directly deduce that $\nu_i \leq \tau_i$ for every $i \in [\![0, m^* - 1]\!]$. Furthermore, for $i = m^*$, we have

$$\begin{aligned}
\nu_{m^*} = |\mathcal{I}_{m^*}| + |\widetilde{\mathcal{J}}_{m^*}(k_{m^*})| + \sum_{t \in \mathcal{J}_{m^*}(k_{m^*})} \frac{U_{\mathrm{S}}(t_{m^*})}{U_{\mathrm{S}}(t)} &\leq |\mathcal{I}_{m^*}| + |\widetilde{\mathcal{J}}_{m^*}(k_{m^*})| + |\mathcal{J}_{m^*}(k_{m^*})| \\
&= m^* + k_{m^*} \leq m^* + n - m^* = \tau_{m^*},
\end{aligned}$$

where the first inequality holds since $U_{\mathrm{S}}(t) \geq U_{\mathrm{S}}(t_m)$ for every $t \in \mathcal{J}_{m^*}(k_{m^*})$, and the second inequality follows from the definition of k_{m^*}, which implies that $k_{m^*} \leq n - m^*$. Therefore, $\nu_{m^*} \leq \tau_{m^*}$.

Thus, it remains to show that $\tau_{i-1} \leq \nu_i$ for every $i \in [\![1, m^*]\!]$. To this aim, we consider the differences $\nu_i - \tau_{i-1}$. It follows that

$$\begin{aligned}
\nu_i - \tau_{i-1} &= |\mathcal{I}_i| + |\widetilde{\mathcal{J}}_i(k_i)| + \sum_{t\in\mathcal{J}_i(k_i)} \frac{U_{\mathrm{S}}(t_i)}{U_{\mathrm{S}}(t)} - \left(|\mathcal{I}_{i-1}| + |\widetilde{\mathcal{J}}_{i-1}(k_{i-1})| + \sum_{t\in\mathcal{J}_{i-1}(k_{i-1})} \frac{U_{\mathrm{S}}(t_i)}{U_{\mathrm{S}}(t)}\right) \\
&= 1 + \left(|\widetilde{\mathcal{J}}_i(k_i)| - |\widetilde{\mathcal{J}}_{i-1}(k_{i-1})|\right) + \left(\sum_{t\in\mathcal{J}_i(k_i)} \frac{U_{\mathrm{S}}(t_i)}{U_{\mathrm{S}}(t)} - \sum_{t\in\mathcal{J}_{i-1}(k_{i-1})} \frac{U_{\mathrm{S}}(t_i)}{U_{\mathrm{S}}(t)}\right) \\
&\overset{\text{Lemma 18}}{=} 1 + \left(|\widetilde{\mathcal{J}}_i(k_i)| - |\widetilde{\mathcal{J}}_i(k_{i-1}-1)|\right) + \left(\sum_{t\in\mathcal{J}_i(k_i)} \frac{U_{\mathrm{S}}(t_i)}{U_{\mathrm{S}}(t)} - \sum_{t\in\mathcal{J}_i(k_{i-1}-1)\cup\{t_i\}} \frac{U_{\mathrm{S}}(t_i)}{U_{\mathrm{S}}(t)}\right) \\
&= \left(|\widetilde{\mathcal{J}}_i(k_i)| - |\widetilde{\mathcal{J}}_i(k_{i-1}-1)|\right) + \left(\sum_{t\in\mathcal{J}_i(k_i)} \frac{U_{\mathrm{S}}(t_i)}{U_{\mathrm{S}}(t)} - \sum_{t\in\mathcal{J}_i(k_{i-1}-1)} \frac{U_{\mathrm{S}}(t_i)}{U_{\mathrm{S}}(t)}\right) \\
&\overset{\text{Lemma 18}}{\geq} 0.
\end{aligned}$$

In the last inequality we used the fact that $k_i \geq k_{i-1} - 1$ (Lemma 18), so $\widetilde{\mathcal{J}}_i(k_{i-1}-1) \subseteq \widetilde{\mathcal{J}}_i(k_i)$ and $\mathcal{J}_i(k_{i-1}-1) \subseteq \mathcal{J}_i(k_i)$. $\square$

Before proving Theorem 3, we need the following auxiliary lemmas.

Lemma 21. *For every $i \in [\![0, m^*]\!]$,*

$$U_{\mathrm{H}}(t_{\pi^i(k_i+1)}) \leq \frac{|\mathcal{J}_i(k_i)| - r_{\mathrm{S}}}{\sum_{t\in\mathcal{J}_i(k_i)} \frac{1}{U_{\mathrm{H}}(t)}} < U_{\mathrm{H}}(t_{\pi^i(k_i)}).$$

Proof of Lemma 21. From Lemma 17, the function $f(k) := \sum_{t\in\mathcal{J}_i(k)} \left(1 - U_{\mathrm{H}}(t_{\pi^i(k)})/U_{\mathrm{H}}(t)\right)$ is nondecreasing in $[\![1, n-i]\!]$. Using this fact and the definition of k_i, it follows that

$$\sum_{t\in\mathcal{J}_i(k_i)} \left(1 - \frac{U_{\mathrm{H}}(t_{\pi^i(k_i)})}{U_{\mathrm{H}}(t)}\right) < r_{\mathrm{S}} \leq \sum_{t\in\mathcal{J}_i(k_i+1)} \left(1 - \frac{U_{\mathrm{H}}(t_{\pi^i(k_i+1)})}{U_{\mathrm{H}}(t)}\right). \tag{B.4}$$

From the definition of $\mathcal{J}_i(k_i+1)$, we have either $\mathcal{J}_i(k_i+1) = \mathcal{J}_i(k_i) \cup \{t_{\pi^i(k_i+1)}\}$ or $\mathcal{J}_i(k_i+1) = \mathcal{J}_i(k_i)$. In the former case, the term corresponding to $t_{\pi^i(k_i+1)}$ in the upper bound in (B.4) is equal to zero. In the latter case, such term is not part of the sum. Therefore, the sum in the upper bound of (B.4) can be equivalently written over $t \in \mathcal{J}_i(k_i)$, instead of over $t \in \mathcal{J}_i(k_i + 1)$. Using this fact, we obtain the following equivalent version of (B.4):

$$\sum_{t \in \mathcal{J}_i(k_i)} \left(1 - \frac{U_{\mathrm{H}}(t_{\pi^i(k_i)})}{U_{\mathrm{H}}(t)}\right) < r_{\mathrm{S}} \leq \sum_{t \in \mathcal{J}_i(k_i)} \left(1 - \frac{U_{\mathrm{H}}(t_{\pi^i(k_i+1)})}{U_{\mathrm{H}}(t)}\right).$$

Finally, rearranging terms provides the lemma. □

Lemma 22. *For every $i \in [\![0, m^*]\!]$,*

$$\frac{|\mathcal{J}_i(k_i)| - r_{\mathrm{S}}}{\sum_{t \in \mathcal{J}_i(k_i)} \frac{1}{U_{\mathrm{H}}(t)}} < U_{\mathrm{H}}(t), \qquad \forall\, t \in \mathcal{I}_i.$$

Proof of Lemma 22. We divide the proof into three steps. In the first step, we show that

$$\frac{|\mathcal{J}_i(k_i)| - r_{\mathrm{S}}}{\sum_{t \in \mathcal{J}_i(k_i)} \frac{1}{U_{\mathrm{H}}(t)}} < U_{\mathrm{H}}(t_i), \qquad \forall\, i \in [\![0, m^*]\!]. \tag{B.5}$$

To this end, we first observe that if $i = 0$, there is nothing to show, as $U_{\mathrm{H}}(t_0) \coloneqq +\infty$.

For $i \in [\![1, m^*]\!]$, we consider two cases. First, if $U_{\mathrm{H}}(t_{\pi^i(k_i)}) \leq U_{\mathrm{H}}(t_i)$, then (B.5) follows directly by Lemma 21. Second, if $U_{\mathrm{H}}(t_i) < U_{\mathrm{H}}(t_{\pi^i(k_i)})$, let $j^* \in [\![1, k_{i-1}]\!]$ be the index such that $t_i = t_{\pi^{i-1}(j^*)}$. We next argue that the inequality $U_{\mathrm{H}}(t_i) < U_{\mathrm{H}}(t_{\pi^i(k_i)})$ implies $j^* = k_{i-1}$ and $k_i = k_{i-1} - 1$. Indeed, let us first suppose that $j^* \in [\![1, k_{i-1} - 1]\!]$. Then,

$$U_{\mathrm{H}}(t_{\pi^{i-1}(j^*)}) = U_{\mathrm{H}}(t_i) < U_{\mathrm{H}}(t_{\pi^i(k_i)}) \overset{\text{Lemma 18}}{\leq} U_{\mathrm{H}}(t_{\pi^i(k_{i-1}-1)}) \overset{\text{(B.1)}}{=} U_{\mathrm{H}}(t_{\pi^{i-1}(k_{i-1})}).$$

Therefore, $U_{\mathrm{H}}(t_{\pi^{i-1}(j^*)}) < U_{\mathrm{H}}(t_{\pi^{i-1}(k_{i-1})})$. This implies that $j^* > k_{i-1}$, which is a contradiction. Thus, $j^* = k_{i-1}$. Next, we show that $k_i = k_{i-1} - 1$. We recall that $k_i \geq k_{i-1} - 1$ (Lemma 18). Let us suppose, by contradiction, that $k_i \geq k_{i-1}$. Then,

$$U_{\mathrm{H}}(t_{\pi^{i-1}(j^*)}) = U_{\mathrm{H}}(t_i) < U_{\mathrm{H}}(t_{\pi^i(k_i)}) \leq U_{\mathrm{H}}(t_{\pi^i(k_{i-1})}) \overset{\text{(B.1)}}{=} U_{\mathrm{H}}(t_{\pi^{i-1}(k_{i-1}+1)}).$$

Hence, $U_{\mathrm{H}}(t_{\pi^{i-1}(j^*)}) < U_{\mathrm{H}}(t_{\pi^{i-1}(k_{i-1}+1)})$. This implies $j^* > k_{i-1} + 1$, which is also a contradiction. Therefore, $k_i = k_{i-1} - 1$ holds. Then, using these facts, it follows that

$$\begin{aligned}\sum_{t\in\mathcal{J}_i(k_i)}\left(1-\frac{U_{\mathrm{H}}(t_i)}{U_{\mathrm{H}}(t)}\right) = \sum_{t\in\mathcal{J}_i(k_{i-1}-1)}\left(1-\frac{U_{\mathrm{H}}(t_i)}{U_{\mathrm{H}}(t)}\right) &= \sum_{t\in\mathcal{J}_i(k_{i-1}-1)\cup\{t_i\}}\left(1-\frac{U_{\mathrm{H}}(t_i)}{U_{\mathrm{H}}(t)}\right)\\ &\overset{\text{Lemma 18}}{=} \sum_{t\in\mathcal{J}_{i-1}(k_{i-1})}\left(1-\frac{U_{\mathrm{H}}(t_i)}{U_{\mathrm{H}}(t)}\right)\\ &= \sum_{t\in\mathcal{J}_{i-1}(k_{i-1})}\left(1-\frac{U_{\mathrm{H}}(t_{\pi^{i-1}(k_{i-1})})}{U_{\mathrm{H}}(t)}\right) < r_{\mathrm{S}},\end{aligned}$$

where the last inequality is by definition of k_{i-1}. Therefore,

$$\sum_{t\in\mathcal{J}_i(k_i)}\left(1-\frac{U_{\mathrm{H}}(t_i)}{U_{\mathrm{H}}(t)}\right) < r_{\mathrm{S}}.$$

Rearranging terms, we obtain (B.5).

In the second step of the proof, we show that

$$\frac{|\mathcal{J}_i(k_i)| - r_{\mathrm{S}}}{\sum_{t\in\mathcal{J}_i(k_i)}\frac{1}{U_{\mathrm{H}}(t)}} > \frac{|\mathcal{J}_{i+1}(k_{i+1})| - r_{\mathrm{S}}}{\sum_{t\in\mathcal{J}_{i+1}(k_{i+1})}\frac{1}{U_{\mathrm{H}}(t)}}, \qquad \forall\, i \in [\![0, m^* - 1]\!]. \tag{B.6}$$

To this end, we recall from Lemma 18 that $k_{i+1} \geq k_i - 1$. Then, we consider the following two cases. First if $k_{i+1} \geq k_i$, it follows that

$$\begin{aligned}\frac{|\mathcal{J}_i(k_i)| - r_{\mathrm{S}}}{\sum_{t\in\mathcal{J}_i(k_i)}\frac{1}{U_{\mathrm{H}}(t)}} \overset{\text{Lemma 21}}{\geq} U_{\mathrm{H}}(t_{\pi^i(k_i+1)}) \overset{\text{(B.1)}}{=} U_{\mathrm{H}}(t_{\pi^{i+1}(k_i)}) &\geq U_{\mathrm{H}}(t_{\pi^{i+1}(k_{i+1})})\\ &\overset{\text{Lemma 21}}{>} \frac{|\mathcal{J}_{i+1}(k_{i+1})| - r_{\mathrm{S}}}{\sum_{t\in\mathcal{J}_{i+1}(k_{i+1})}\frac{1}{U_{\mathrm{H}}(t)}}.\end{aligned}$$

Now, let us assume that $k_{i+1} = k_i - 1$. We recall that the following property of real numbers, called the *mediant inequality*: Let $a, c \geq 0$ and $b, d > 0$. Then, $a/b > c/d$ implies $a/b > (a+c)/(b+d) > c/d$. In this case, we let $a \coloneqq 1$, $b \coloneqq 1/U_{\mathrm{H}}(t_{i+1})$,

$c \coloneqq |\mathcal{J}_{i+1}(k_{i+1})| - r_{\mathrm{S}}$ and $d \coloneqq \sum_{t\in\mathcal{J}_{i+1}(k_{i+1})} \frac{1}{U_{\mathrm{H}}(t)}$. Then, $a/b > c/d$ is equivalent to

$$U_{\mathrm{H}}(t_{i+1}) > \frac{|\mathcal{J}_{i+1}(k_{i+1})| - r_{\mathrm{S}}}{\sum_{t\in\mathcal{J}_{i+1}(k_{i+1})} \frac{1}{U_{\mathrm{H}}(t)}},$$

which holds by (B.5). From the mediant inequality, we deduce that $(a+c)/(b+d) > c/d$, which translates to

$$\frac{|\mathcal{J}_{i+1}(k_{i+1})| - r_{\mathrm{S}} + 1}{\sum_{t\in\mathcal{J}_{i+1}(k_i-1)} \frac{1}{U_{\mathrm{H}}(t)} + \frac{1}{U_{\mathrm{H}}(t_{i+1})}} > \frac{|\mathcal{J}_{i+1}(k_{i+1})| - r_{\mathrm{S}}}{\sum_{t\in\mathcal{J}_{i+1}(k_{i+1})} \frac{1}{U_{\mathrm{H}}(t)}}. \tag{B.7}$$

Therefore,

$$\begin{aligned}\frac{|\mathcal{J}_i(k_i)| - r_{\mathrm{S}}}{\sum_{t\in\mathcal{J}_i(k_i)} \frac{1}{U_{\mathrm{H}}(t)}} \overset{\text{Lemma 18}}{=} \frac{|\mathcal{J}_{i+1}(k_i-1)\cup\{t_{i+1}\}| - r_{\mathrm{S}}}{\sum_{t\in\mathcal{J}_{i+1}(k_i-1)\cup\{t_{i+1}\}} \frac{1}{U_{\mathrm{H}}(t)}} &= \frac{|\mathcal{J}_{i+1}(k_{i+1})| - r_{\mathrm{S}} + 1}{\sum_{t\in\mathcal{J}_{i+1}(k_i-1)} \frac{1}{U_{\mathrm{H}}(t)} + \frac{1}{U_{\mathrm{H}}(t_{i+1})}} \\ &\overset{\text{(B.7)}}{>} \frac{|\mathcal{J}_{i+1}(k_{i+1})| - r_{\mathrm{S}}}{\sum_{t\in\mathcal{J}_{i+1}(k_{i+1})} \frac{1}{U_{\mathrm{H}}(t)}}.\end{aligned}$$

from which we deduce (B.6).

Finally, in the third step of the proof, we let $t_j \in \mathcal{I}_i \setminus \{t_i\}$, so $j < i$. Then, by repeatedly applying (B.6), we obtain

$$U_{\mathrm{H}}(t_j) \overset{\text{(B.5)}}{>} \frac{|\mathcal{J}_j(k_j)| - r_{\mathrm{S}}}{\sum_{t\in\mathcal{J}_j(k_j)} \frac{1}{U_{\mathrm{H}}(t)}} \overset{\text{(B.6)}}{>} \cdots \overset{\text{(B.6)}}{>} \frac{|\mathcal{J}_i(k_i)| - r_{\mathrm{S}}}{\sum_{t\in\mathcal{J}_i(k_i)} \frac{1}{U_{\mathrm{H}}(t)}}, \qquad \forall\, t_j \in \mathcal{I}_i \setminus \{t_i\}\,. \tag{B.8}$$

Therefore, from (B.5) and (B.8), we conclude the proof. □

Building on Lemmas 21 and 22, we can now proceed to prove Lemma 6.

Proof of Lemma 6. Let $i \in [\![0, m^*]\!]$. From Lemma 21, we have

$$U_{\mathrm{H}}(t_{\pi^i(k_i)}) > \frac{|\mathcal{J}_i(k_i)| - r_{\mathrm{S}}}{\sum_{t\in\mathcal{J}_i(k_i)} \frac{1}{U_{\mathrm{H}}(t)}} \geq U_{\mathrm{H}}(t_{\pi^i(k_i+1)}).$$

We note that for every $t \in \mathcal{J}_i(k_i) \cup \widetilde{\mathcal{J}}_i(k_i) = \left\{t_{\pi^i(1)}, \ldots, t_{\pi^i(k_i)}\right\}$, it holds that $U_{\mathrm{H}}(t) \geq U_{\mathrm{H}}(t_{\pi^i(k_i)})$. Similarly, for every $t' \in \mathcal{K}_i(k_i)$, we have $U_{\mathrm{H}}(t_{\pi^i(k_i+1)}) \geq U_{\mathrm{H}}(t')$. Furthermore, from Lemma 22, it follows that

$$U_{\mathrm{H}}(t) > \frac{|\mathcal{J}_i(k_i)| - r_{\mathrm{S}}}{\sum_{t\in\mathcal{J}_i(k_i)} \frac{1}{U_{\mathrm{H}}(t)}}, \qquad \forall\, t \in \mathcal{I}_i.$$

Combining these inequalities, we deduce the lemma. □

Lemma 23. *For every $i \in [\![1, m^*]\!]$, if $\tau_{i-1} < r_{\mathrm{H}} \leq \nu_i$, then*

$$0 < r_{\mathrm{S}} - \sum_{t\in\mathcal{J}_i(\ell_i)} \left(1 - \frac{U_{\mathrm{H}}(t_{\pi^i(\ell_i+1)})}{U_{\mathrm{H}}(t)}\right) \leq 1 - \frac{U_{\mathrm{H}}(t_{\pi^i(\ell_i+1)})}{U_{\mathrm{H}}(t_i)}.$$

Proof of Lemma 23. From Lemma 20, we have $\ell_i < k_i$, or equivalently, $\ell_i + 1 \leq k_i$. Furthermore, from Lemma 17, the function $f(k) \coloneqq \sum_{t\in\mathcal{J}_i(k)} \big(1 - U_{\mathrm{H}}(t_{\pi^i(k)})/U_{\mathrm{H}}(t)\big)$ is nondecreasing over $[\![1, n-i]\!]$. Using these facts, it follows that

$$\begin{aligned}\sum_{t\in\mathcal{J}_i(\ell_i)} \left(1 - \frac{U_{\mathrm{H}}(t_{\pi^i(\ell_i+1)})}{U_{\mathrm{H}}(t)}\right) = \sum_{t\in\mathcal{J}_i(\ell_i+1)} \left(1 - \frac{U_{\mathrm{H}}(t_{\pi^i(\ell_i+1)})}{U_{\mathrm{H}}(t)}\right) &\leq \sum_{t\in\mathcal{J}_i(k_i)} \left(1 - \frac{U_{\mathrm{H}}(t_{\pi^i(k_i)})}{U_{\mathrm{H}}(t)}\right) \\ &< r_{\mathrm{S}},\end{aligned}$$

where the last inequality holds by definition of k_i. Hence,

$$0 < r_{\mathrm{S}} - \sum_{t\in\mathcal{J}_i(\ell_i)} \left(1 - \frac{U_{\mathrm{H}}(t_{\pi^i(\ell_i+1)})}{U_{\mathrm{H}}(t)}\right).$$

Next, using the relations between π^{i-1} and π^{i}, as well as k_{i-1} and ℓ_i, it follows that

$$
\begin{aligned}
r_{\mathrm{S}} &\overset{\text{Lemma 21}}{\leq} \sum_{t \in \mathcal{J}_{i-1}(k_{i-1})} \left(1 - \frac{U_{\mathrm{H}}(t_{\pi^{i-1}(k_{i-1}+1)})}{U_{\mathrm{H}}(t)}\right) \\
&\overset{\text{Lemma 18}}{=} \sum_{t \in \mathcal{J}_{i}(k_{i-1}-1) \cup \{t_i\}} \left(1 - \frac{U_{\mathrm{H}}(t_{\pi^{i-1}(k_{i-1}+1)})}{U_{\mathrm{H}}(t)}\right) \\
&\overset{\text{(B.2)}}{=} \sum_{t \in \mathcal{J}_{i}(k_{i-1}-1)} \left(1 - \frac{U_{\mathrm{H}}(t_{\pi^{i}(k_{i-1})})}{U_{\mathrm{H}}(t)}\right) + \left(1 - \frac{U_{\mathrm{H}}(t_{\pi^{i}(k_{i-1})})}{U_{\mathrm{H}}(t_i)}\right) \\
&= \sum_{t \in \mathcal{J}_{i}(k_{i-1})} \left(1 - \frac{U_{\mathrm{H}}(t_{\pi^{i}(k_{i-1})})}{U_{\mathrm{H}}(t)}\right) + \left(1 - \frac{U_{\mathrm{H}}(t_{\pi^{i}(k_{i-1})})}{U_{\mathrm{H}}(t_i)}\right) \\
&\overset{\text{Lemma 20}}{\leq} \sum_{t \in \mathcal{J}_{i}(\ell_i+1)} \left(1 - \frac{U_{\mathrm{H}}(t_{\pi^{i}(\ell_i+1)})}{U_{\mathrm{H}}(t)}\right) + \left(1 - \frac{U_{\mathrm{H}}(t_{\pi^{i}(\ell_i+1)})}{U_{\mathrm{H}}(t_i)}\right) \\
&= \sum_{t \in \mathcal{J}_{i}(\ell_i)} \left(1 - \frac{U_{\mathrm{H}}(t_{\pi^{i}(\ell_i+1)})}{U_{\mathrm{H}}(t)}\right) + \left(1 - \frac{U_{\mathrm{H}}(t_{\pi^{i}(\ell_i+1)})}{U_{\mathrm{H}}(t_i)}\right).
\end{aligned}
$$

Therefore,

$$
r_{\mathrm{S}} \leq \sum_{t \in \mathcal{J}_{i}(\ell_i)} \left(1 - \frac{U_{\mathrm{H}}(t_{\pi^{i}(\ell_i+1)})}{U_{\mathrm{H}}(t)}\right) + \left(1 - \frac{U_{\mathrm{H}}(t_{\pi^{i}(\ell_i+1)})}{U_{\mathrm{H}}(t_i)}\right),
$$

from which we conclude the proof. □

Lemma 24. *For every $i \in [\![1, m^*]\!]$, if $\tau_{i-1} < r_{\mathrm{H}}$, then*

$$
0 < r_{\mathrm{H}} - |\mathcal{I}_i| - |\widetilde{\mathcal{J}}_i(\ell_i)| - \sum_{t \in \mathcal{J}_i(\ell_i)} \frac{U_{\mathrm{S}}(t_i)}{U_{\mathrm{S}}(t)} \leq \min\left\{ \frac{U_{\mathrm{S}}(t_i)}{U_{\mathrm{S}}(t_{\pi^{i}(\ell_i+1)})},\, 1 \right\}.
$$

Proof of Lemma 24. From Lemma 19, it follows that the function $g(\ell) \coloneqq |\mathcal{I}_i| + |\widetilde{\mathcal{J}}_i(\ell)| + \sum_{t \in \mathcal{J}_i(\ell)} U_{\mathrm{S}}(t_i)/U_{\mathrm{S}}(t)$ is nondecreasing over $[\![0, n-i]\!]$. Then, from the definitions of ℓ_i and the sets $\mathcal{J}_i(\ell_i+1)$ and $\widetilde{\mathcal{J}}_i(\ell_i+1)$, we have

$$
\begin{aligned}
|\mathcal{I}_i| + |\widetilde{\mathcal{J}}_i(\ell_i)| + \sum_{t \in \mathcal{J}_i(\ell_i)} \frac{U_{\mathrm{S}}(t_i)}{U_{\mathrm{S}}(t)} < r_{\mathrm{H}} &\le |\mathcal{I}_i| + |\widetilde{\mathcal{J}}_i(\ell_i+1)| + \sum_{t \in \mathcal{J}_i(\ell_i+1)} \frac{U_{\mathrm{S}}(t_i)}{U_{\mathrm{S}}(t)} \\
&= |\mathcal{I}_i| + |\widetilde{\mathcal{J}}_i(\ell_i)| + \sum_{t \in \mathcal{J}_i(\ell_i)} \frac{U_{\mathrm{S}}(t_i)}{U_{\mathrm{S}}(t)} \\
&\quad + \min\left\{ \frac{U_{\mathrm{S}}(t_i)}{U_{\mathrm{S}}(t_{\pi^i(\ell_i+1)})}, 1 \right\},
\end{aligned}
$$

from which we deduce the lemma. □

Now, we are ready to prove Theorem 3. Let $r_{\mathrm{S}}, r_{\mathrm{H}} \in [\![1, n-1]\!]$, and let $i^* \in [\![0, m^*]\!]$ be the unique index satisfying $\tau_{i^*-1} < r_{\mathrm{H}} \le \tau_{i^*}$.

Proof of Theorem 3 (Regime Pattern 1). Suppose that $i^* \in [\![0, m^*-1]\!]$ and $\tau_{i^*-1} < r_{\mathrm{H}} \le \nu_{i^*}$. We note that since $\tau_{-1} = \nu_0 = 0$ (Lemma 7), we must have $i^* \ge 1$. From Lemma 20, we know that $\ell_{i^*} < k_{i^*}$. Let $\rho^{\mathrm{S}^*} \in \mathbb{R}^{\mathcal{T}}$ and $\rho^{\mathrm{H}^*} \in \mathbb{R}^{\mathcal{T}}$ satisfying (3.5) and (3.6), respectively. We will show that $(\rho^{\mathrm{S}^*}, \rho^{\mathrm{H}^*}) \in \widetilde{\mathcal{A}}_{\mathrm{S}} \times \widetilde{\mathcal{A}}_{\mathrm{H}}$ and is an equilibrium of $\widetilde{\Gamma}$.

First, we show that $\rho^{\mathrm{S}^*} \in \widetilde{\mathcal{A}}_{\mathrm{S}}$. From Lemma 23, it follows that the quantity $\varepsilon_{\mathrm{S}} := r_{\mathrm{S}} - \sum_{t \in \mathcal{J}_{i^*}(\ell_{i^*})} \big(1 - U_{\mathrm{H}}(t_{\pi^{i^*}(\ell_{i^*}+1)})/U_{\mathrm{H}}(t)\big)$ in (3.7) satisfies the inequalities

$$
0 < \varepsilon_{\mathrm{S}} \le 1 - \frac{U_{\mathrm{H}}(t_{\pi^{i^*}(\ell_{i^*}+1)})}{U_{\mathrm{H}}(t_{i^*})}. \tag{B.9}
$$

This implies $0 \le \rho^{\mathrm{S}^*}_{t_{i^*}} \le 1$. Furthermore, since $U_{\mathrm{H}}(t) \ge U_{\mathrm{H}}(t_{\pi^{i^*}(\ell_{i^*}+1)})$ for every $t \in \mathcal{J}_{i^*}(\ell_{i^*})$ and $U_{\mathrm{H}}(t) > 0$ for every $t \in \mathcal{T}$, we also have $0 \le \rho^{\mathrm{S}^*}_t \le 1$ for every $t \in \mathcal{J}_{i^*}(\ell_{i^*})$. Additionally, from the definition of ε_{S}, we obtain $\sum_{t \in \mathcal{T}} \rho^{\mathrm{S}^*}_t = r_{\mathrm{S}}$. Thus, $\rho^{\mathrm{S}^*} \in \widetilde{\mathcal{A}}_{\mathrm{S}}$.

We now show that $\rho^{\mathrm{H}^*} \in \widetilde{\mathcal{A}}_{\mathrm{H}}$. From Lemma 24, the quantity $\varepsilon_{\mathrm{H}} := r_{\mathrm{H}} - |\mathcal{I}_{i^*}| - |\widetilde{\mathcal{J}}_{i^*}(\ell_{i^*})| - \sum_{t \in \mathcal{J}_{i^*}(\ell_{i^*})} U_{\mathrm{S}}(t_{i^*})/U_{\mathrm{S}}(t)$ in (3.7) satisfies the following bounds:

$$
0 < \varepsilon_{\mathrm{H}} \le \min\left\{ \frac{U_{\mathrm{S}}(t_{i^*})}{U_{\mathrm{S}}(t_{\pi^{i^*}(\ell_{i^*}+1)})}, 1 \right\}. \tag{B.10}
$$

This ensures that $0 \leq \rho^{\mathrm{H}^*}_{t_{\pi^{i^*}(\ell_{i^*}+1)}} \leq 1$. Additionally, from the definition of $\mathcal{J}_{i^*}(\ell_{i^*})$, we know that $U_{\mathrm{S}}(t_i) \leq U_{\mathrm{S}}(t)$ for every $t \in \mathcal{J}_{i^*}(\ell_{i^*})$. Thus, $0 \leq \rho^{\mathrm{H}^*}_t \leq 1$ for every $t \in \mathcal{J}_{i^*}(\ell_{i^*})$. Furthermore, the definition of ε_{H} implies $\sum_{t\in\mathcal{T}} \rho^{\mathrm{H}^*}_t = r_{\mathrm{H}}$. Therefore, $\rho^{\mathrm{H}^*} \in \widetilde{\mathcal{A}}_{\mathrm{H}}$.

Next, we show that $(\rho^{\mathrm{S}^*}, \rho^{\mathrm{H}^*})$ is an equilibrium of $\widetilde{\Gamma}$. To this end, we first note that for a fixed $\rho^{\mathrm{H}} \in \widetilde{\mathcal{A}}_{\mathrm{H}}$, a best response for S is given a by an optimal solution of

$$
\begin{aligned}
\max_{\rho^{\mathrm{S}}\in\widetilde{\mathcal{A}}_{\mathrm{S}}} \mathbf{U}_{\mathrm{S}}(\rho^{\mathrm{S}}, \rho^{\mathrm{H}}) \;=\; \max_{\rho^{\mathrm{S}}\in\mathbb{R}^{\mathcal{T}}} \quad & \sum_{t\in\mathcal{T}} U_{\mathrm{S}}(t)\, \rho^{\mathrm{H}}_t\, \rho^{\mathrm{S}}_t \\
\text{subject to} \quad & \sum_{t\in\mathcal{T}} \rho^{\mathrm{S}}_t \leq r_{\mathrm{S}} \\
& 0 \leq \rho^{\mathrm{S}}_t \leq 1 \qquad \forall\, t \in \mathcal{T}.
\end{aligned}
\tag{B.11}
$$

In other words, a best response to $\rho^{\mathrm{H}} \in \widetilde{\mathcal{A}}_{\mathrm{H}}$ is an optimal solution to a continuous knapsack problem with n different (fractional) objects of unitary weights and a knapsack capacity equal to r_{S}. Each object $t \in \mathcal{T}$ has a profit equal to $U_{\mathrm{S}}(t)\, \rho^{\mathrm{H}}_t$. An optimal solution for (B.11) consists in filling the capacity of the knapsack with the objects with highest profits.

Then, given ρ^{H^*} satisfying (3.6), the profit of each object in the knapsack problem (B.11) is given by:

$$
U_{\mathrm{S}}(t)\rho^{\mathrm{H}^*}_t = \begin{cases}
U_{\mathrm{S}}(t_{i^*}) & \text{if } t \in \mathcal{J}_{i^*}(\ell_{i^*}), \\
U_{\mathrm{S}}(t) & \text{if } t \in \mathcal{I}_{i^*} \cup \widetilde{\mathcal{J}}_{i^*}(\ell_{i^*}), \\
\varepsilon_{\mathrm{H}}\, U_{\mathrm{S}}(t_{\pi^{i^*}(\ell_{i^*}+1)}) & \text{if } t = t_{\pi^{i^*}(\ell_{i^*}+1)}, \\
0 & \text{if } t \in \mathcal{K}_{i^*}(\ell_{i^*}) \setminus \left\{ t_{\pi^{i^*}(\ell^*_i+1)} \right\}.
\end{cases}
$$

From Lemma 5, the location t_{i^*} satisfies $U_{\mathrm{S}}(t) \leq U_{\mathrm{S}}(t_{i^*})$ for every $t \in \mathcal{I}_{i^*}$. Additionally, from the definition of $\widetilde{\mathcal{J}}_{i^*}(\ell_{i^*})$, we have $U_{\mathrm{S}}(t) < U_{\mathrm{S}}(t_{i^*})$ for every $t \in \widetilde{\mathcal{J}}_{i^*}(\ell_{i^*})$. Furthermore, from (B.10), it follows that $\varepsilon_{\mathrm{H}}\, U_{\mathrm{S}}(t_{\pi^{i^*}(\ell_{i^*}+1)}) \leq U_{\mathrm{S}}(t_{i^*})$. Therefore, combining these inequalities, the profits in the in the knapsack problem (B.11) satisfy

$$\begin{aligned} U_{\mathrm{S}}(t_{i^*}) &\geq U_{\mathrm{S}}(t), && \forall\, t \in \mathcal{I}_{i^*} \cup \widetilde{\mathcal{J}}_{i^*}(\ell_{i^*}), \\ U_{\mathrm{S}}(t_{i^*}) &\geq \varepsilon_{\mathrm{H}}\, U_{\mathrm{S}}(t_{\pi^{i^*}(\ell_{i^*}+1)}). \end{aligned} \tag{B.12}$$

Thus, from inequalities (B.12), the most "profitable" items in the continuous knapsack problem (B.11) are those in $\mathcal{J}_{i^*}(\ell_{i^*}) \cup \{t_{i^*}\}$. Then, a best response for S consists in assigning $\rho^{\mathrm{S}}_t = 1$ to r_{S} items in $\mathcal{J}_{i^*}(\ell_{i^*}) \cup \{t_{i^*}\}$. This is possible if $r_{\mathrm{S}} \leq |\mathcal{J}_{i^*}(\ell_{i^*})| + 1$. To see this inequality holds, we note that

$$\begin{aligned} r_{\mathrm{S}} &\overset{\text{Lemma 23}}{\leq} \sum_{t \in \mathcal{J}_i(\ell_i) \cup \{t_{i^*}\}} \left(1 - \frac{U_{\mathrm{H}}(t_{\pi^i(\ell_i+1)})}{U_{\mathrm{H}}(t)}\right) \\ &= |\mathcal{J}_i(\ell_i)| + 1 - \sum_{t \in \mathcal{J}_i(\ell_i) \cup \{t_{i^*}\}} \frac{U_{\mathrm{H}}(t_{\pi^i(\ell_i+1)})}{U_{\mathrm{H}}(t)} \leq |\mathcal{J}_i(\ell_i)| + 1. \end{aligned}$$

This best response yields S a utility of $U_{\mathrm{S}}(t_{i^*})\, r_{\mathrm{S}} = \mathbf{U}_{\mathrm{S}}(\rho^{\mathrm{S}^*}, \rho^{\mathrm{H}^*})$.

On the other hand, for a fixed $\rho^{\mathrm{S}} \in \widetilde{\mathcal{A}}_{\mathrm{S}}$, a best response for H is given a by an optimal solution of

$$\begin{aligned} \max_{\rho^{\mathrm{H}} \in \widetilde{\mathcal{A}}_{\mathrm{H}}} \mathbf{U}_{\mathrm{H}}(\rho^{\mathrm{S}}, \rho^{\mathrm{H}}) \;=\; \max_{\rho^{\mathrm{H}} \in \mathbb{R}^{\mathcal{T}}} \quad & \sum_{t \in \mathcal{T}} U_{\mathrm{H}}(t)\, \left(1 - \rho^{\mathrm{S}}_t\right)\, \rho^{\mathrm{H}}_t \\ \text{subject to} \quad & \sum_{t \in \mathcal{T}} \rho^{\mathrm{H}}_t \leq r_{\mathrm{H}} \\ & 0 \leq \rho^{\mathrm{H}}_t \leq 1 \qquad \forall\, t \in \mathcal{T}. \end{aligned} \tag{B.13}$$

Similarly, a best response to $\rho^{\mathrm{S}} \in \widetilde{\mathcal{A}}_{\mathrm{S}}$ is an optimal solution to a continuous knapsack problem with n different (fractional) objects of unitary weights and a knapsack capacity equal to r_{H}, and where each object $t \in \mathcal{T}$ has a profit equal to $U_{\mathrm{H}}(t)\, \left(1 - \rho^{\mathrm{S}}_t\right)$. An optimal solution for (B.13) consists in filling the capacity of the knapsack with the objects with highest profits.

Thus, given ρ^{S^*} satisfying (3.5), the profit of each object in the knapsack problem (B.13) is given by:

$$U_{\mathrm{H}}(t)\left(1-\rho_t^{\mathrm{S}^*}\right)=\begin{cases}(1-\varepsilon_{\mathrm{S}})\,U_{\mathrm{H}}(t_{i^*}) & \text{if } t=t_{i^*},\\ U_{\mathrm{H}}(t_{\pi^{i^*}(\ell_{i^*}+1)}) & \text{if } \mathcal{J}_{i^*}(\ell_{i^*}),\\ U_{\mathrm{H}}(t) & \text{if } t\in\mathcal{I}_{i^*}\setminus\{t_{i^*}\}\cup\widetilde{\mathcal{J}}_{i^*}(\ell_{i^*})\cup\mathcal{K}_{i^*}(\ell_{i^*}).\end{cases}$$

From (B.9), we have $(1-\varepsilon_{\mathrm{S}})\,U_{\mathrm{H}}(t_{i^*})\geq U_{\mathrm{H}}(t_{\pi^{i^*}(\ell_{i^*}+1)})$. Moreover, from the definition of $\widetilde{\mathcal{J}}_{i^*}(\ell_{i^*})$, we also know that $U_{\mathrm{H}}(t)\geq U_{\mathrm{H}}(t_{\pi^{i^*}(\ell_{i^*}+1)})$ for every $t\in\widetilde{\mathcal{J}}_{i^*}(\ell_{i^*})$. Similarly, the definition of $\mathcal{K}_{i^*}(\ell_{i^*})$ implies $U_{\mathrm{H}}(t_{\pi^{i^*}(\ell_{i^*}+1)})\geq U_{\mathrm{H}}(t)$ for every $t\in\mathcal{K}_{i^*}(\ell_{i^*})$. Furthermore, the locations in $\mathcal{I}_{i^*}\setminus\{t_{i^*}\}=\mathcal{I}_{i^*-1}$ satisfy the following inequalities:

$$\begin{aligned}\min_{t\in\mathcal{I}_{i^*-1}}U_{\mathrm{H}}(t)\overset{\text{Lemma 22}}{>}\frac{|\mathcal{J}_{i^*-1}(k_{i^*-1})|-r_{\mathrm{S}}}{\sum_{t\in\mathcal{J}_{i^*-1}(k_{i^*-1})}\frac{1}{U_{\mathrm{H}}(t)}}&\overset{\text{Lemma 21}}{\geq}U_{\mathrm{H}}(t_{\pi^{i^*-1}(k_{i^*-1}+1)})\\ &\overset{\text{(B.2)}}{=}U_{\mathrm{H}}(t_{\pi^{i^*}(k_{i^*-1})})\\ &\overset{\text{Lemma 20}}{\geq}U_{\mathrm{H}}(t_{\pi^{i^*}(\ell_{i^*}+1)}).\end{aligned}$$

Therefore, from these inequalities, the profits in the knapsack problem (B.13) satisfy

$$\begin{aligned}&U_{\mathrm{H}}(t)\geq U_{\mathrm{H}}(t_{\pi^{i^*}(\ell_{i^*}+1)})\geq U_{\mathrm{H}}(t'), \qquad \forall\, t\in\mathcal{I}_{i^*}\setminus\{t_{i^*}\}\cup\widetilde{\mathcal{J}}_{i^*}(\ell_{i^*})\ \forall\, t'\in\mathcal{K}_{i^*}(\ell_{i^*}),\\ &\qquad U_{\mathrm{H}}(t_{\pi^{i^*}(\ell_{i^*}+1)})\geq(1-\varepsilon_{\mathrm{S}})\,U_{\mathrm{H}}(t_{i^*}).\end{aligned} \tag{B.14}$$

Hence, from inequalities (B.14), the most "profitable" items in the continuous knapsack problem (B.13) are those in $\mathcal{I}_{i^*}\cup\widetilde{\mathcal{J}}_{i^*}(\ell_{i^*})$, followed by those in $\mathcal{J}_{i^*}(\ell_{i^*})\cup\left\{t_{\pi^{i^*}(\ell_{i^*}+1)}\right\}$. Consequently, a best response for H involves assigning $\rho_t^{\mathrm{H}}=1$ to all items in $\mathcal{I}_{i^*}\cup\widetilde{\mathcal{J}}_{i^*}(\ell_{i^*})$, and to $r_{\mathrm{H}}-|\mathcal{I}_{i^*}|-|\widetilde{\mathcal{J}}_{i^*}(\ell_{i^*})|$ items in $\mathcal{J}_{i^*}(\ell_{i^*})\cup\left\{t_{\pi^{i^*}(\ell_{i^*}+1)}\right\}$. This is feasible if $|\mathcal{I}_{i^*}|+$

$|\widetilde{\mathcal{J}}_{i^*}(\ell_{i^*})| < r_{\mathrm{H}}$ and $r_{\mathrm{H}} - |\mathcal{I}_{i^*}| - |\widetilde{\mathcal{J}}_{i^*}(\ell_{i^*})| \leq |\mathcal{J}_{i^*}(\ell_{i^*})| + 1$. These inequalities hold because

$$\begin{aligned} 0 < \sum_{t \in \mathcal{J}_{i^*}(\ell_{i^*})} \frac{U_{\mathrm{S}}(t_{i^*})}{U_{\mathrm{S}}(t)} &\overset{\text{(B.13)}}{<} r_{\mathrm{H}} - |\mathcal{I}_{i^*}| - |\widetilde{\mathcal{J}}_{i^*}(\ell_{i^*})| \\ &\overset{\text{(B.13)}}{\leq} \sum_{t \in \mathcal{J}_{i^*}(\ell_{i^*})} \frac{U_{\mathrm{S}}(t_{i^*})}{U_{\mathrm{S}}(t)} + \min\left\{ \frac{U_{\mathrm{S}}(t_{i^*})}{U_{\mathrm{S}}(t_{\pi^{i^*}(\ell_{i^*}+1)})}, 1 \right\} \\ &\leq |\mathcal{J}_{i^*}(\ell_{i^*})| + 1. \end{aligned}$$

Finally, we note that this best response yields H a utility of

$$\begin{aligned} &\sum_{t \in \mathcal{I}_{i^*} \setminus \{t_{i^*}\} \cup \widetilde{\mathcal{J}}_{i^*}(\ell_{i^*})} U_{\mathrm{H}}(t) + (1 - \varepsilon_{\mathrm{S}})\, U_{\mathrm{H}}(t_{i^*}) + U_{\mathrm{H}}(t_{\pi^{i^*}(\ell_{i^*}+1)}) \left(r_{\mathrm{H}} - |\mathcal{I}_{i^*}| - |\widetilde{\mathcal{J}}_{i^*}(\ell_{i^*})| \right) = \\ &\quad \mathbf{U}_{\mathrm{H}}(\rho^{\mathrm{S}^*}, \rho^{\mathrm{H}^*}). \end{aligned}$$

□

Proof of Theorem 3 (Regime Pattern 2). Suppose that $i^* \in [\![0, m^* - 1]\!]$ and $\nu_{i^*} < r_{\mathrm{H}} \leq \tau_{i^*}$. From Lemma 20, we know that $k_{i^*} \leq \ell_{i^*}$. Let $\rho^{\mathrm{S}^*} \in \mathbb{R}^{\mathcal{T}}$ and $\rho^{\mathrm{H}^*} \in \mathbb{R}^{\mathcal{T}}$ satisfying (3.10) and (3.11), respectively. We will analogously show that $(\rho^{\mathrm{S}^*}, \rho^{\mathrm{H}^*}) \in \widetilde{\mathcal{A}}_{\mathrm{S}} \times \widetilde{\mathcal{A}}_{\mathrm{H}}$ and is an equilibrium of $\widetilde{\Gamma}$.

We first show that $\rho^{\mathrm{S}^*} \in \widetilde{\mathcal{A}}_{\mathrm{S}}$. From Lemmas 21 and 22, the equalization constant λ_{S} in (3.12) satisfies the following inequalities:

$$U_{\mathrm{H}}(t_{\pi^{i^*}(k_{i^*}+1)}) \leq \lambda_{\mathrm{S}} = \frac{|\mathcal{J}_{i^*}(k_{i^*})| - r_{\mathrm{S}}}{\sum_{t \in \mathcal{J}_{i^*}(k_{i^*})} \frac{1}{U_{\mathrm{H}}(t)}} < \min\left\{ U_{\mathrm{H}}(t_{\pi^{i^*}(k_{i^*})}),\, \min_{t \in \mathcal{I}_{i^*}} U_{\mathrm{H}}(t) \right\}. \quad \text{(B.15)}$$

From the definition of $\mathcal{K}_{i^*}(k_{i^*})$, we have $U_{\mathrm{H}}(t) \leq U_{\mathrm{H}}(t_{\pi^{i^*}(k_{i^*}+1)})$ for every $t \in \mathcal{K}_{i^*}(k_{i^*})$. Therefore, the first inequality in (B.15) implies $U_{\mathrm{H}}(t) \leq \lambda_{\mathrm{S}}$ for every $t \in \mathcal{K}_{i^*}(k_{i^*})$. Additionally, from the definitions of $\mathcal{J}_{i^*}(k_{i^*})$ and $\widetilde{\mathcal{J}}_{i^*}(k_{i^*})$, it follows that $U_{\mathrm{H}}(t_{\pi^{i^*}(k_{i^*})}) \leq U_{\mathrm{H}}(t)$ for every $t \in \mathcal{J}_{i^*}(k_{i^*}) \cup \widetilde{\mathcal{J}}_{i^*}(k_{i^*})$. Consequently, the second inequality in (B.15) implies $\lambda_{\mathrm{S}} < U_{\mathrm{H}}(t)$ for every $t \in \mathcal{J}_{i^*}(k_{i^*}) \cup \widetilde{\mathcal{J}}_{i^*}(k_{i^*})$. Thus, (B.15) implies λ_{S} satisfies the following inequalities:

$$U_{\mathrm{H}}(t) \leq \lambda_{\mathrm{S}} < U_{\mathrm{H}}(t'), \qquad \forall\, t \in \mathcal{K}_{i^*}(k_{i^*}),\ \forall\, t' \in \mathcal{I}_{i^*} \cup \mathcal{J}_{i^*}(k_{i^*}) \cup \widetilde{\mathcal{J}}_{i^*}(k_{i^*}). \tag{B.16}$$

In particular, from the inequality $\lambda_{\mathrm{S}} < U_{\mathrm{H}}(t')$ for every $t' \in \mathcal{J}_{i^*}(k_{i^*})$, we deduce that $\rho_t^{\mathrm{S}^*} \geq 0$ for every $t \in \mathcal{J}_{i^*}(k_{i^*})$. Moreover, since $\lambda_{\mathrm{S}} \geq U_{\mathrm{H}}(t_{\pi^{i^*}(k_{i^*}+1)}) > 0$ (B.15), it follows that $\rho_t^{\mathrm{S}^*} \leq 1$ for every $t \in \mathcal{J}_{i^*}(k_{i^*})$. Finally, from the definition of λ_{S}, we also have $\sum_{t\in\mathcal{T}} \rho_t^{\mathrm{S}^*} = r_{\mathrm{S}}$. Thus, $\rho^{\mathrm{S}^*} \in \widetilde{\mathcal{A}}_{\mathrm{S}}$.

Next, we show that $\rho^{\mathrm{H}^*} \in \widetilde{\mathcal{A}}_{\mathrm{H}}$. The inequalities $\nu_{i^*} < r_{\mathrm{H}} \leq \tau_{i^*}$ translate into the following inequalities for the equalization constant λ_{H} in (3.12):

$$U_{\mathrm{S}}(t_{i^*}) < \lambda_{\mathrm{H}} = \frac{r_{\mathrm{H}} - |\mathcal{I}_{i^*}| - |\widetilde{\mathcal{J}}_{i^*}(k_{i^*})|}{\sum_{t\in\mathcal{J}_{i^*}(k_{i^*})} \frac{1}{U_{\mathrm{S}}(t)}} \leq U_{\mathrm{S}}(t_{i^*+1}). \tag{B.17}$$

We recall that the locations within $\mathcal{I}_{i^*}$ satisfy $0 =: U_{\mathrm{S}}(t_0) \leq \cdots \leq U_{\mathrm{S}}(t_{i^*})$ (Lemma 5). Moreover, from the definition of $\widetilde{\mathcal{J}}_{i^*}(k_{i^*})$, we have $U_{\mathrm{S}}(t) < U_{\mathrm{S}}(t_{i^*})$ for every $t \in \widetilde{\mathcal{J}}_{i^*}(k_{i^*})$. On the other hand, from the definition of t_{i^*+1}, the locations $t \in \mathcal{J}_{i^*}(k_{i^*})$ satisfy $U_{\mathrm{S}}(t) \geq U_{\mathrm{S}}(t_{i^*+1})$. Therefore, (B.17) implies λ_{H} satisfies

$$U_{\mathrm{S}}(t) < \lambda_{\mathrm{H}} \leq U_{\mathrm{S}}(t'), \qquad \forall\, t \in \mathcal{I}_{i^*} \cup \widetilde{\mathcal{J}}_{i^*}(k_{i^*}),\ \forall\, t' \in \mathcal{J}_{i^*}(k_{i^*}). \tag{B.18}$$

These inequalities imply that $0 \leq \rho_t^{\mathrm{H}^*} \leq 1$ for every $\mathcal{J}_{i^*}(k_{i^*})$. Furthermore, from the definition of λ_{H}, it holds that $\sum_{t\in\mathcal{T}} \rho_t^{\mathrm{H}^*} = r_{\mathrm{H}}$. Thus, $\rho^{\mathrm{H}^*} \in \widetilde{\mathcal{A}}_{\mathrm{H}}$.

Next, we show that $(\rho^{\mathrm{S}^*}, \rho^{\mathrm{H}^*})$ is a NE of the game $\widetilde{\Gamma}$. Given ρ^{H^*} satisfying (3.11), the "profit" of each object in the knapsack problem (B.11) is given by:

$$U_{\mathrm{S}}(t)\, \rho_t^{\mathrm{H}^*} = \begin{cases} \lambda_{\mathrm{H}} & \text{if } t \in \mathcal{J}_{i^*}(k_{i^*}) \\ U_{\mathrm{S}}(t) & \text{if } t \in \mathcal{I}_{i^*} \cup \widetilde{\mathcal{J}}_{i^*}(k_{i^*}), \\ 0 & \text{if } t \in \mathcal{K}_{i^*}(k_{i^*}). \end{cases} \tag{B.19}$$

From the inequalities (B.18), the most "profitable" items in the continuous knapsack problem (B.11) are those in $\mathcal{J}_{i^*}(k_{i^*})$. Then, a best response for S consists in assigning $\rho_t^{\mathrm{S}} = 1$ to r_{S} items in $\mathcal{J}_{i^*}(k_{i^*})$. This is possible if $r_{\mathrm{S}} \leq |\mathcal{J}_{i^*}(k_{i^*})|$. To see this inequality holds, we note that

$$r_{\mathrm{S}} \overset{\text{(B.15)}}{\leq} \sum_{t \in \mathcal{J}_i(k_i)} \left(1 - \frac{U_{\mathrm{H}}(t_{\pi^i(k_i+1)})}{U_{\mathrm{H}}(t)}\right) = |\mathcal{J}_i(k_i)| - \sum_{t \in \mathcal{J}_i(k_i)} \frac{U_{\mathrm{H}}(t_{\pi^i(k_i+1)})}{U_{\mathrm{H}}(t)} \leq |\mathcal{J}_i(k_i)|.$$

This best response yields S a utility of $\lambda_{\mathrm{H}}\, r_{\mathrm{S}} = \mathbf{U}_{\mathrm{S}}(\rho^{\mathrm{S}^*}, \rho^{\mathrm{H}^*})$.

On the other hand, given ρ^{S^*} satisfying (3.10), the "profit" of each object in the knapsack problem (B.13) is given by:

$$U_{\mathrm{H}}(t)\left(1 - \rho_t^{\mathrm{S}^*}\right) = \begin{cases} \lambda_{\mathrm{S}} & \text{if } t \in \mathcal{J}_{i^*}(k_{i^*}), \\ U_{\mathrm{H}}(t) & \text{if } t \in \mathcal{I}_{i^*} \cup \widetilde{\mathcal{J}}_{i^*}(k_{i^*}) \cup \mathcal{K}_{i^*}(k_{i^*}). \end{cases} \tag{B.20}$$

From the inequalities (B.16), the most "profitable" items in the continuous knapsack problem (B.13) are those in $\mathcal{I}_{i^*} \cup \widetilde{\mathcal{J}}_{i^*}(k_{i^*})$, followed by those in $\mathcal{J}_{i^*}(k_{i^*})$. Then, a best response for S consists in assigning $\rho_t^{\mathrm{S}} = 1$ to all items in $\mathcal{I}_{i^*} \cup \widetilde{\mathcal{J}}_{i^*}(k_{i^*})$, as well as to $r_{\mathrm{H}} - |\mathcal{I}_{i^*}| - |\widetilde{\mathcal{J}}_{i^*}(k_{i^*})|$ items in $\mathcal{J}_{i^*}(k_{i^*})$. This is possible if $|\mathcal{I}_{i^*}| + |\widetilde{\mathcal{J}}_{i^*}(k_{i^*})| < r_{\mathrm{H}}$ and $r_{\mathrm{H}} - |\mathcal{I}_{i^*}| - |\widetilde{\mathcal{J}}_{i^*}(k_{i^*})| \leq |\mathcal{J}_{i^*}(k_{i^*})|$. These inequalities are a consequence of

$$0 < \sum_{t \in \mathcal{J}_{i^*}(k_{i^*})} \frac{U_{\mathrm{S}}(t_{i^*})}{U_{\mathrm{S}}(t)} \overset{\text{(B.17)}}{<} r_{\mathrm{H}} - |\mathcal{I}_{i^*}| - |\widetilde{\mathcal{J}}_{i^*}(k_{i^*})| \overset{\text{(B.17)}}{\leq} \sum_{t \in \mathcal{J}_{i^*}(k_{i^*})} \frac{U_{\mathrm{S}}(t_{i^*+1})}{U_{\mathrm{S}}(t)} \leq |\mathcal{J}_{i^*}(k_{i^*})|,$$

where in the last inequality we used that $U_{\mathrm{S}}(t_{i^*+1}) \leq U_{\mathrm{S}}(t)$ for every $t \in \mathcal{J}_{i^*}(k_{i^*})$, which follows from the definition of t_{i^*+1}. This best response yields H a utility of

$$\lambda_{\mathrm{S}}\left(r_{\mathrm{H}} - |\mathcal{I}_{i^*}| - |\widetilde{\mathcal{J}}_{i^*}(k_{i^*})|\right) + \sum_{t \in \mathcal{I}_{i^*} \cup \widetilde{\mathcal{J}}_{i^*}(k_{i^*})} U_{\mathrm{H}}(t) = \mathbf{U}_{\mathrm{H}}(\rho^{\mathrm{S}^*}, \rho^{\mathrm{H}^*}).$$

$\square$

Proof of Theorem 3 (Regime Pattern 3). Suppose that $i^* = m^*$ and $\nu_{m^*} < r_{\mathrm{H}} \leq \tau_{m^*}$. We first show that $k_{m^*} = n - m^*$. From the definition of k_{m^*}, we know that $k_{m^*} \in \{1, \ldots, n - m^*\}$. Let us assume, by contradiction, that $k_{m^*} < n - m^*$. Then, $k_{m^*} + 1 \leq n - m^*$. This implies $U_{\mathrm{H}}(t_{\pi^{m^*}(k_{m^*}+1)}) > 0$—we remark that this is not the case if $k_{m^*} = n - m^*$, as $U_{\mathrm{H}}(t_{\pi^{m^*}(n-m^*+1)}) = 0$ by definition. Then, from Lemma 21, we have

$$U_{\mathrm{H}}(t_{\pi^{m^*}(k_{m^*}+1)}) \leq \frac{|\mathcal{J}_{m^*}(k_{m^*})| - r_{\mathrm{S}}}{\sum_{t \in \mathcal{J}_{m^*}(k_{m^*})} \frac{1}{U_{\mathrm{H}}(t)}} < U_{\mathrm{H}}(t_{\pi^{m^*}(k_{m^*})}). \tag{B.21}$$

In particular, from the first inequality in (B.21) and the fact that $U_{\mathrm{H}}(t_{\pi^{m^*}(k_{m^*}+1)}) > 0$, we deduce that $r_{\mathrm{S}} < |\mathcal{J}_{m^*}(k_{m^*})|$. This is a contradiction, since by definition, the index m^* satisfies $|\mathcal{J}_{m^*}(k_{m^*})| \leq r_{\mathrm{S}}$. Thus, $k_{m^*} = n - m^*$.

We note that the equality $k_{m^*} = n - m^*$ implies $\mathcal{K}_{m^*}(n - m^*) = \emptyset$. Therefore, the set of locations $\mathcal{T}$ is partitioned into the sets $\mathcal{I}_{m^*}$, $\mathcal{J}_{m^*}(n - m^*)$ and $\widetilde{\mathcal{J}}_{m^*}(n - m^*)$.

As we mentioned, the definition of m^* implies $|\mathcal{J}_{m^*}(k_{m^*})| \leq r_{\mathrm{S}}$. Furthermore, since $k_{m^*} = n - m^*$, the first inequality in (B.21) implies $r_{\mathrm{S}} \leq |\mathcal{J}_{m^*}(k_{m^*})|$. Therefore, we have $|\mathcal{J}_{m^*}(k_{m^*})| = |\mathcal{J}_{m^*}(n - m^*)| = r_{\mathrm{S}}$.

Let $\rho^{\mathrm{S}^*} \in \mathbb{R}^{\mathcal{T}}$ and $\rho^{\mathrm{H}^*} \in \mathbb{R}^{\mathcal{T}}$ satisfying (3.13) and (3.14), respectively. We will analogously show that $(\rho^{\mathrm{S}^*}, \rho^{\mathrm{H}^*}) \in \widetilde{\mathcal{A}}_{\mathrm{S}} \times \widetilde{\mathcal{A}}_{\mathrm{H}}$ and is an equilibrium of $\widetilde{\Gamma}$.

We note that the equality $|\mathcal{J}_{m^*}(k_{m^*})| = r_{\mathrm{S}}$ guarantees that $\rho^{\mathrm{S}^*} \in \widetilde{\mathcal{A}}_{\mathrm{S}}$. On the other hand, from the definition of $\mathcal{J}_{m^*}(n - m^*)$, we have $U_{\mathrm{S}}(t_{m^*}) \leq U_{\mathrm{S}}(t)$ for every $t \in \mathcal{J}_{m^*}(n - m^*)$. This implies that the inequalities (3.14) defining ρ^{H^*} are feasible, and that $\rho^{\mathrm{H}^*} \in \widetilde{\mathcal{A}}_{\mathrm{H}}$.

We next show that $(\rho^{\mathrm{S}^*}, \rho^{\mathrm{H}^*}) \in \widetilde{\mathcal{A}}_{\mathrm{S}} \times \widetilde{\mathcal{A}}_{\mathrm{H}}$ and is an equilibrium of $\widetilde{\Gamma}$. Given ρ^{H^*} satisfying (3.14), the profit of each object in the knapsack problem (B.11) is given by:

$$\begin{cases} U_{\mathrm{S}}(t_{m^*}) \leq U_{\mathrm{S}}(t)\rho_t^{\mathrm{H}^*} \leq U_{\mathrm{S}}(t) & \text{if } t \in \mathcal{J}_{m^*}(n - m^*), \\ \qquad\qquad U_{\mathrm{S}}(t)\rho_t^{\mathrm{H}^*} = U_{\mathrm{S}}(t) & \text{if } t \in \mathcal{I}_{m^*} \cup \widetilde{\mathcal{J}}_{m^*}(n - m^*). \end{cases} \tag{B.22}$$

From Lemma 5, we have $U_{\mathrm{S}}(t) \leq U_{\mathrm{S}}(t_{m^*})$ for every $t \in \mathcal{I}_{m^*}$. Moreover, from the definition of $\widetilde{\mathcal{J}}_{m^*}(n-m^*)$ it follows that $U_{\mathrm{S}}(t) < U_{\mathrm{S}}(t_{m^*})$. Thus, $U_{\mathrm{S}}(t) \leq U_{\mathrm{S}}(t_{m^*})$ for every $t \in \mathcal{I}_{m^*} \cup \widetilde{\mathcal{J}}_{m^*}(n-m^*)$. This renders the objects in $\mathcal{J}_{m^*}(n-m^*)$ the most profitable for S. Then, a best response for S consists in assigning $\rho_t^{\mathrm{S}} = 1$ to all the r_{S} items in $\mathcal{J}_{m^*}(n-m^*)$. This yields S a utility of $\sum_{t\in\mathcal{J}_{m^*}(n-m^*)} U_{\mathrm{S}}(t)\rho_t^{\mathrm{H}^*}$, which from (B.22) lies in the continuous interval $[U_{\mathrm{S}}(t_{m^*})r_{\mathrm{S}}, \sum_{t\in\mathcal{J}_{m^*}(n-m^*)} U_{\mathrm{S}}(t)]$.

On the other hand, given ρ^{S^*} satisfying (3.13), the "profit" of each object in the knapsack problem (B.13) is given by:

$$U_{\mathrm{H}}(t)\left(1-\rho_t^{\mathrm{S}^*}\right) = \begin{cases} 0 & \text{if } t \in \mathcal{J}_{m^*}(n-m^*), \\ U_{\mathrm{H}}(t) & \text{if } t \in \mathcal{I}_{m^*} \cup \widetilde{\mathcal{J}}_{m^*}(n-m^*). \end{cases} \tag{B.23}$$

Thus, the objects in $\mathcal{I}_{m^*} \cup \widetilde{\mathcal{J}}_{m^*}(n-m^*)$ are the most profitable for H. Furthermore, the inequality $\nu_{m^*} < r_{\mathrm{H}}$ translates into

$$|\mathcal{I}_{m^*}| + |\widetilde{\mathcal{J}}_{m^*}(n-m^*)| + \sum_{t\in\mathcal{J}_{m^*}(n-m^*)} \frac{U_{\mathrm{S}}(t_{m^*})}{U_{\mathrm{S}}(t)} < r_{\mathrm{H}}.$$

In other words, H has sufficiently many resources to hide items in all locations within $\mathcal{I}_{m^*} \cup \widetilde{\mathcal{J}}_{m^*}(n-m^*)$. Therefore, a best response for H involves setting $\rho_t^{\mathrm{H}^*} = 1$ for every $t \in \mathcal{I}_{m^*} \cup \widetilde{\mathcal{J}}_{m^*}(n-m^*)$. This yields H a utility of $\sum_{t\in\mathcal{I}_{m^*}\cup\widetilde{\mathcal{J}}_{m^*}(n-m^*)} U_{\mathrm{H}}(t) = \mathbf{U}_{\mathrm{H}}(\rho^{\mathrm{S}^*}, \rho^{\mathrm{H}^*})$. $\square$

APPENDIX C

SUPPLEMENT TO CHAPTER 4

C.1 Proofs of Statements

Proof of Lemma 8. Let $(\sigma^{\mathrm{D}}, \sigma^{\mathrm{A}}) \in \Delta_{\mathrm{D}} \times \Delta_{\mathrm{A}}$. Then,

$$\begin{aligned} U(\sigma^{\mathrm{D}}, \sigma^{\mathrm{A}}) = \sum_{S \in \mathcal{D}} \sum_{T \in \mathcal{A}} \sigma^{\mathrm{D}}_S \, \sigma^{\mathrm{A}}_T \, u(S,T) &= \sum_{S \in \mathcal{D}} \sum_{T \in \mathcal{A}} \sigma^{\mathrm{D}}_S \, \sigma^{\mathrm{A}}_T \sum_{e \in T} u(S,e) \\ &= \sum_{S \in \mathcal{D}} \sigma^{\mathrm{D}}_S \sum_{e \in \mathcal{E}} \sum_{T \in \mathcal{A} \colon e \in T} \sigma^{\mathrm{A}}_T \, u(S,e) \\ &= \sum_{S \in \mathcal{D}} \sigma^{\mathrm{D}}_S \sum_{e \in \mathcal{E}} \rho_e(\sigma^{\mathrm{A}}) \, u(S,e) \\ &= \sum_{e \in \mathcal{E}} \rho_e(\sigma^{\mathrm{A}}) \, U(\sigma^{\mathrm{D}}, e). \end{aligned}$$

□

Proof of Proposition 3. Let $z^* \in \mathbb{R}$ be the optimal value of the variable z in (LP). Then,

$$z^* = \min_{\sigma^{\mathrm{D}} \in \Delta_{\mathrm{D}}} \max_{T \in \mathcal{A}} U(\sigma^{\mathrm{D}}, T) = \min_{\sigma^{\mathrm{D}} \in \Delta_{\mathrm{D}}} \max_{\sigma^{\mathrm{A}} \in \Delta_{\mathrm{A}}} U(\sigma^{\mathrm{D}}, \sigma^{\mathrm{A}}) = \min_{\sigma^{\mathrm{D}} \in \Delta_{\mathrm{D}}} \max_{\sigma^{\mathrm{A}} \in \Delta_{\mathrm{A}}} \sum_{e \in \mathcal{E}} \rho_e(\sigma^{\mathrm{A}}) \, U(\sigma^{\mathrm{D}}, e).$$

The first equality follows by the constraints of (LP). The second equality holds since for every fixed $\sigma^{\mathrm{D}} \in \Delta_{\mathrm{D}}$, the function $U(\sigma^{\mathrm{D}}, \sigma^{\mathrm{A}}) = \sum_{T \in \mathcal{A}} \sigma^{\mathrm{A}}_T \, U(\sigma^{\mathrm{D}}, T)$ is linear in σ^{A}, and therefore it attains its maximum over the set of extreme points of Δ_{A}, given by the characteristic vectors of the sets in $\mathcal{A}$. The third equality follows from Lemma 8. Thus, (LP) is equivalent to

$$\min_{\sigma^{\mathrm{D}} \in \Delta_{\mathrm{D}}} \max_{\sigma^{\mathrm{A}} \in \Delta_{\mathrm{A}}} \sum_{e \in \mathcal{E}} \rho_e(\sigma^{\mathrm{A}}) \, U(\sigma^{\mathrm{D}}, e). \tag{C.1}$$

Next, for fixed $\sigma^{\mathrm{D}} \in \Delta_{\mathrm{D}}$, the inner maximization in (C.1) is equivalent to the following LP:

$$\begin{aligned} \max_{\rho^{\mathrm{A}} \in \Theta_{\mathrm{A}}} \sum_{e \in \mathcal{E}} \rho_e^{\mathrm{A}} \, U(\sigma^{\mathrm{D}}, e) = \max_{\rho^{\mathrm{A}} \in \mathbb{R}^{\mathcal{E}}} \quad & \sum_{e \in \mathcal{E}} \rho_e^{\mathrm{A}} \sum_{S \in \mathcal{D}} \sigma_S^{\mathrm{D}} \, u(S, e) \\ \text{subject to} \quad & \sum_{e \in \mathcal{E}} \rho_e^{\mathrm{A}} \leq r_{\mathrm{A}} \\ & 0 \leq \rho_e^{\mathrm{A}} \leq 1 \qquad \forall \, e \in \mathcal{E}. \end{aligned} \tag{C.2}$$

Indeed, by Lemma 9, for any feasible solution of the inner maximization of (C.1), there exists a feasible solution of (C.2) with the same expected number of undetected attacks and vice versa. Now, we consider the dual of (C.2):

$$\begin{aligned} \min_{\lambda \in \mathbb{R}^{\mathcal{E}}, \, \gamma \in \mathbb{R}} \quad & r_{\mathrm{A}} \gamma + \sum_{e \in \mathcal{E}} \lambda_e \\ \text{subject to} \quad & \gamma + \lambda_e \geq \sum_{S \in \mathcal{D}} \sigma_S^{\mathrm{D}} \, u(S, e) \quad \forall \, e \in \mathcal{E}, \\ & \lambda_e \geq 0 \qquad \forall \, e \in \mathcal{E}, \\ & \gamma \geq 0. \end{aligned} \tag{C.3}$$

By strong duality, the optimal values of (C.2) and (C.3) coincide. Thus, the reformulation (LP($\mathcal{D}$)) is obtained by substituting the inner maximization in (C.1) by (C.3). In particular, given an optimal solution $(\sigma^{\mathrm{D}^*}, \lambda^*, \gamma^*) \in \mathbb{R}^{\mathcal{D}} \times \mathbb{R}^{\mathcal{E}} \times \mathbb{R}$ of (LP($\mathcal{D}$)), the value of the game is given by the optimal value $r_{\mathrm{A}} \gamma^* + \sum_{e \in \mathcal{E}} \lambda_e^*$.

Next, let $(\rho^{\mathrm{A}^*}, \nu^*) \in \mathbb{R}^{\mathcal{E}} \times \mathbb{R}$ be an optimal solution of the dual of (LP($\mathcal{D}$)), which is given by

$$
\begin{aligned}
\max_{\rho^{\mathrm{A}} \in \mathbb{R}^{\mathcal{E}},\, \nu \in \mathbb{R}} \quad & \nu \\
\text{subject to} \quad & \nu \le \sum_{e \in \mathcal{E}} \rho^{\mathrm{A}}_e \, u(S,e) \qquad \forall\, S \in \mathcal{D}, \\
& \sum_{e \in \mathcal{E}} \rho^{\mathrm{A}}_e \le r_{\mathrm{A}} \\
& 0 \le \rho^{\mathrm{A}}_e \le 1 \qquad\qquad \forall\, e \in \mathcal{E}.
\end{aligned}
\tag{C.4}
$$

From the equivalence between (LP) and (LP($\mathcal{D}$)), it follows that σ^{D^*} in an optimal solution of $\min_{\sigma^{\mathrm{D}} \in \Delta_{\mathrm{D}}} \max_{\rho^{\mathrm{A}} \in \Theta_{\mathrm{A}}} U(\sigma^{\mathrm{D}}, \rho^{\mathrm{A}})$. On the other hand, we can similarly show that the dual of (LP) is equivalent to (C.4), which in turn is equivalent to $\max_{\rho^{\mathrm{A}} \in \Theta_{\mathrm{A}}} \min_{\sigma^{\mathrm{D}} \in \Delta_{\mathrm{D}}} U(\sigma^{\mathrm{D}}, \rho^{\mathrm{A}})$, for which ρ^{A^*} is an optimal solution. Therefore, $(\sigma^{\mathrm{D}^*}, \rho^{\mathrm{A}^*})$ is an equilibrium of the game Γ. $\square$

Proof of Proposition 4. We show that Problem (DBR) is NP-hard even for instances where $r_{\mathrm{A}} = 1$, every component belongs to at most two monitoring sets, and the detection probabilities are homogeneous. To this aim, we reduce the NP-hard problem VERTEX COVER to (DBR). Given a graph $G = (V, E)$ and a parameter $k \in \mathbb{Z}_{>0}$, the VERTEX COVER problem consists in finding a subset of nodes $S \subseteq V$ of size at most k, such that every edge of G is incident to at least one node of S.

Given an instance of VERTEX COVER, we construct an instance of (DBR) by setting $\mathcal{V} = V$ and $\mathcal{E} = E$. Then, we define the monitoring sets and detection probabilities respectively as $\mathcal{C}_v = \{e \in E :\ v \in e\}$ and $p_v = 1 - 1/(m+1)$ for all $v \in \mathcal{V}$ (we recall that $m := |\mathcal{E}|$). We note that every component belongs to exactly two monitoring sets, and the detection probabilities are homogeneous. Next, we set $r_{\mathrm{A}} = 1$ and let ρ^{A} be the uniform distribution in $\mathcal{E}$, that is $\rho^{\mathrm{A}}_e = 1/m$ for every $e \in \mathcal{E}$. Finally, we let $r_{\mathrm{D}} = k$, so $\mathcal{D} = \{S \subseteq \mathcal{V} :\ |S| \le k\}$. From the selection of the detection probabilities, for every detector positioning $S \in \mathcal{D}$ and component $e \in \mathcal{E}$, we have $u(S,e) \in \{1, 1/(m+1), 1/(m+1)^2\}$.

Specifically, $u(S,e) = 1$ if $e \notin \cup_{v \in S} \mathcal{C}_v$, $u(S,e) = 1/(m+1)$ if there is exactly one $v \in S$ such that $e \in \mathcal{C}_v$, and $u(S,e) = 1/(m+1)^2$ if there are exactly two $v, w \in S$ such that $e \in \mathcal{C}_v \cap \mathcal{C}_w$. Next, we argue that G has a vertex cover of size at most k if and only if there exists a detector positioning $S \in \mathcal{D}$ such that $U(S, \rho^{\mathrm{A}}) \leq 1/(m+1)$.

Suppose G has a vertex cover S of size at most k. Then, $S \in \mathcal{D}$, and since all edges are covered by S, we must have $u(S,e) \leq 1/(m+1)$ for every $e \in \mathcal{E}$. It follows that $U(S, \rho^{\mathrm{A}}) = (1/m) \sum_{e \in \mathcal{E}} u(S,e) \leq 1/(m+1)$. Conversely, suppose there is a detector positioning $S \in \mathcal{D}$ such that $U(S, \rho^{\mathrm{A}}) \leq 1/(m+1)$. Then, $|S| \leq k$ and S is a vertex cover in G. Indeed, suppose there is an edge $e' \in E$ which is not incident to any node in S. Then, we must have $u(S, e') = 1$. If $m = 1$, then e' is the only edge of E and $S = \emptyset$, so $U(S, \rho^{\mathrm{A}}) = u(S, e') = 1$, which is a contradiction. If $m > 1$, then it follows that

$$\begin{aligned}
\frac{1}{m+1} \geq U(S, \rho^{\mathrm{A}}) = \frac{1}{m} \sum_{e \in \mathcal{E}} u(S,e) &= \frac{1}{m} u(S, e') + \frac{1}{m} \sum_{\substack{e \in \mathcal{E}: \\ e \neq e'}} u(S,e) \\
&\geq \frac{1}{m} + \frac{1}{m} \cdot \frac{m-1}{(m+1)^2} > \frac{1}{m},
\end{aligned}$$

which is again a contradiction. □

Proof of Proposition 5. For a fixed enumeration of $\mathcal{V} = \{v_1, \ldots, v_n\}$, let $S \in \mathcal{D}$ be a detector positioning and $(x,u) \in \mathbb{R}^n \times \mathbb{R}^{\mathcal{E} \times \{1,\ldots,n\}}$ be defined as

$$x_i = \mathbb{1}_S(v_i) \qquad \forall\, v_i \in \mathcal{V}, \tag{C.5}$$

$$u_{e,i} = \prod_{v_j \in \mathcal{V}:\, j \leq i} \Big(1 - p_{v_j} \mathbb{1}_{\mathcal{C}_{v_j}}(e) x_j\Big) \qquad \forall\, e \in \mathcal{E}, \forall\, v_i \in \mathcal{V}, \tag{C.6}$$

where $\mathbb{1}_S(v_i) = 1$ (resp. $\mathbb{1}_{\mathcal{C}_{v_j}}(e) = 1$) if $v_i \in S$ (resp. $e \in \mathcal{C}_{v_j}$) and $\mathbb{1}_S(v_i) = 0$ (resp. $\mathbb{1}_{\mathcal{C}_{v_j}}(e) = 0$) otherwise. We first show that (x,u) is a feasible solution of (MIP). To this end, let us consider its constraints:

$$
\begin{aligned}
\sum_{v_i \in \mathcal{V}} x_i &\leq r_{\mathrm{D}}, && && \text{(C.7)}\\
1 - p_{v_1} \mathbb{1}_{\mathcal{C}_{v_1}}(e) x_1 &\leq u_{e,1} && \forall\, e \in \mathcal{E}, && \text{(C.8)}\\
u_{e,i} \left(1 - p_{v_{i+1}} \mathbb{1}_{\mathcal{C}_{v_{i+1}}}(e)\right) &\leq u_{e,i+1} && \forall\, e \in \mathcal{E},\, \forall\, v_i \in \mathcal{V} \setminus \{v_n\}, && \text{(C.9)}\\
u_{e,i} - x_{i+1} &\leq u_{e,i+1} && \forall\, e \in \mathcal{E},\, \forall\, v_i \in \mathcal{V} \setminus \{v_n\}. && \text{(C.10)}\\
0 \leq u_{e,i} &\leq 1 && \forall\, v_i \in \mathcal{V}, && \text{(C.11)}\\
x_i &\in \{0,1\} && \forall\, v_i \in \mathcal{V}. && \text{(C.12)}
\end{aligned}
$$

By definition, (x,u) clearly satisfies (C.7), (C.8), (C.11) and (C.12). Then, it remains to show that (x,u) also satisfies (C.9) and (C.10). From (C.6), it holds that (x,u) satisfies the following (nonlinear) equalities:

$$
u_{e,i} \left(1 - p_{v_{i+1}} \mathbb{1}_{\mathcal{C}_{v_{i+1}}}(e) x_{i+1}\right) = u_{e,i+1} \qquad \forall\, e \in \mathcal{E},\, \forall\, v_i \in \mathcal{V} \setminus \{v_n\}. \tag{C.13}
$$

We note that the left hand side of (C.13) is equivalent to

$$
\max \left\{ u_{e,i} \left(1 - p_{v_{i+1}} \mathbb{1}_{\mathcal{C}_{v_{i+1}}}(e)\right),\, u_{e,i} - x_{i+1} \right\}.
$$

Indeed, if $x_{i+1} = 0$, then from the inequalities $u_{e,i} \geq 0$ and $p_{v_{i+1}} \mathbb{1}_{\mathcal{C}_{v_{i+1}}}(e) \leq 1$ it follows that $u_{e,i} \left(1 - p_{v_{i+1}} \mathbb{1}_{\mathcal{C}_{v_{i+1}}}(e)\right) \leq u_{e,i} = u_{e,i} - x_{i+1}$. On the other hand, if $x_{i+1} = 1$, then the inequalities $u_{e,i} \leq 1$ and $p_{v_{i+1}} \mathbb{1}_{\mathcal{C}_{v_{i+1}}}(e) \leq 1$ imply that $u_{e,i} - x_{i+1} = u_{e,i} - 1 \leq u_{e,i} \left(1 - p_{v_{i+1}} \mathbb{1}_{\mathcal{C}_{v_{i+1}}}(e)\right)$. Thus, (C.13) translates to

$$
\max \left\{ u_{e,i} \left(1 - p_{v_{i+1}} \mathbb{1}_{\mathcal{C}_{v_{i+1}}}(e)\right),\, u_{e,i} - x_{i+1} \right\} = u_{e,i+1} \qquad \forall\, e \in \mathcal{E},\, \forall\, v_i \in \mathcal{V} \setminus \{v_n\},
$$

from which we deduce that (x,u) satisfies (C.9) and (C.10). Therefore, (x,u) is feasible for (MIP). We also note that $u_{e,n} = u(S,e)$ for every $e \in \mathcal{E}$.

Next, let $(x^*, u^*) \in \mathbb{R}^n \times \mathbb{R}^{\mathcal{E}\times\{1,\ldots,n\}}$ be an optimal solution of (MIP), and let $S^* := \{v_i \in \mathcal{V} : x_i^* = 1\}$ be its associated detector positioning. By similar arguments as above, we can show that (C.8)-(C.10) imply that (x^*, u^*) satisfies

$$1 - p_{v_1}\mathbb{1}_{\mathcal{C}_{v_1}}(e)x_1^* \le u_{e,1}^* \qquad \forall\, e \in \mathcal{E}, \tag{C.14}$$

$$u_{e,i}^*\left(1 - p_{v_{i+1}}\mathbb{1}_{\mathcal{C}_{v_{i+1}}}(e)x_{i+1}^*\right) \le u_{e,i+1}^* \qquad \forall\, e \in \mathcal{E},\, \forall\, v_i \in \mathcal{V}\setminus\{v_n\}\,. \tag{C.15}$$

By repeatedly applying these inequalities, it follows that (x^*, u^*) satisfies

$$\begin{aligned} u_{e,n}^* &\ge u_{e,n-1}^*\left(1 - p_{v_n}\mathbb{1}_{\mathcal{C}_{v_n}}(e)x_n^*\right) \\ &\ge u_{e,n-2}^*\left(1 - p_{v_{n-1}}\mathbb{1}_{\mathcal{C}_{v_{n-1}}}(e)x_{n-1}^*\right)\left(1 - p_{v_n}\mathbb{1}_{\mathcal{C}_{v_n}}(e)x_n^*\right) \\ &\ge \cdots \\ &\ge \prod_{v_i\in\mathcal{V}}\left(1 - p_{v_i}\mathbb{1}_{\mathcal{C}_{v_i}}(e)x_i^*\right) = u(S^*, e). \end{aligned}$$

Furthermore, since (MIP) is a minimization problem, for every $e \in \mathcal{E}$ with $\rho_e^{\mathrm{A}} > 0$, $u_{e,n}^*$ attains its lower bound $u(S^*, e)$ by tightening the inequalities (C.14) and (C.15). Therefore,

$$U(S^*, \rho^{\mathrm{A}}) = \sum_{e\in\mathcal{E}} \rho_e^{\mathrm{A}}\, u(S^*, e) = \sum_{e\in\mathcal{E}} \rho_e^{\mathrm{A}}\, u_{e,n}^* \le \sum_{e\in\mathcal{E}} \rho^{\mathrm{A}}\, u_{e,n} = \sum_{e\in\mathcal{E}} \rho^{\mathrm{A}}\, u(S, e) = U(S, \rho^{\mathrm{A}}).$$

Since $S \in \mathcal{D}$ is arbitrary, we conclude that S^* is an optimal solution for (DBR). $\square$

Proof of Lemma 10. We first show that for every fixed $e \in \mathcal{E}$, the function $u(S,e) = \prod_{v\in S\colon e\in\mathcal{C}_v}(1-p_v)$, defined for every $S \subseteq \mathcal{V}$, is nonincreasing and supermodular.

– Let $S, S' \subseteq \mathcal{V}$ be such that $S \subseteq S'$. Then,

$$u(S,e) = \prod_{v\in S\colon e\in\mathcal{C}_v}(1-p_v) \ge \prod_{v\in S'\colon e\in\mathcal{C}_v}(1-p_v) = u(S', e).$$

Thus, $u(\cdot, e)$ is nonincreasing.

– Let $S, S' \subseteq \mathcal{V}$ be such that $S \subseteq S'$, and let $v \notin S'$. Then,

$$\begin{aligned}
u(S,e) - u(S \cup \{v\}, e) &= \prod_{w \in S \colon e \in \mathcal{C}_w} (1 - p_w) - \prod_{w \in S \cup \{v\} \colon e \in \mathcal{C}_w} (1 - p_w) \\
&= p_v \mathbb{1}_{\mathcal{C}_v}(e) \prod_{w \in S \colon e \in \mathcal{C}_w} (1 - p_w) \\
&\geq p_v \mathbb{1}_{\mathcal{C}_v}(e) \prod_{w \in S' \colon e \in \mathcal{C}_w} (1 - p_w) \\
&= u(S', e) - u(S' \cup \{v\}, e).
\end{aligned}$$

Therefore, $u(\cdot, e)$ is supermodular.

Finally, we recall that for every fixed $\rho^{\mathrm{A}} \in \Theta_{\mathrm{A}}$, $U(S, \rho^{\mathrm{A}}) = \sum_{e \in \mathcal{E}} \rho_e^{\mathrm{A}} \, u(S, e)$. In particular, $U(\cdot, \rho^{\mathrm{A}})$ is a nonnegative linear combination of nonincreasing and supermodular set functions. This implies that $U(\cdot, \rho^{\mathrm{A}})$ is nonincreasing and supermodular as well. □

Proof of Lemma 11. For every $\rho^{\mathrm{A}} \in \Theta_{\mathrm{A}}$, we have

$$\begin{aligned}
U_{\mathcal{V} \setminus \{v\}}(v, \rho^{\mathrm{A}}) &= U(\mathcal{V} \setminus \{v\}, \rho^{\mathrm{A}}) - U(\mathcal{V}, \rho^{\mathrm{A}}) \\
&= \sum_{e \in \mathcal{E}} \rho_e^{\mathrm{A}} \, u(\mathcal{V} \setminus \{v\}, e) - \sum_{e \in \mathcal{E}} \rho_e^{\mathrm{A}} \, u(\mathcal{V} \setminus \{v\}, e) \, (1 - p_v \mathbb{1}_{\mathcal{C}_v}(e)) \\
&= p_v \sum_{e \in \mathcal{C}_v} \rho_e^{\mathrm{A}} \, u(\mathcal{V} \setminus \{v\}, e) \\
&= p_v \sum_{e \in \mathcal{C}_v} \rho_e^{\mathrm{A}} \prod_{\substack{w \in \mathcal{V} \setminus \{v\} \colon \\ e \in \mathcal{C}_w}} (1 - p_w) \geq \left(1 - \max_{w \in \mathcal{V}} p_w\right)^d p_v \sum_{e \in \mathcal{C}_v} \rho_e^{\mathrm{A}},
\end{aligned}$$

where the inequality is due to the assumption that every component belongs to at most d monitoring sets. On the other hand,

$$U_{\emptyset}(v, \rho^{\mathrm{A}}) = U(\emptyset, \rho^{\mathrm{A}}) - U(v, \rho^{\mathrm{A}}) = \sum_{e \in \mathcal{E}} \rho_e^{\mathrm{A}} - \sum_{e \in \mathcal{E}} \rho_e^{\mathrm{A}} \, (1 - p_v \mathbb{1}_{\mathcal{C}_v}(e)) = p_v \sum_{e \in \mathcal{C}_v} \rho_e^{\mathrm{A}}.$$

Let us assume that $\rho^{\mathrm{A}} \neq \mathbf{0}$. Then, since $p_v > 0$ for every $v \in \mathcal{V}$, the set of detector locations $\{v \in \mathcal{V} : U_{\emptyset}(v, \rho^{\mathrm{A}}) > 0\}$ is nonempty. It follows that

$$c = 1 - \min_{v \in \mathcal{V} :\, U_{\emptyset}(v,\rho^{\mathrm{A}})>0} \frac{U_{\mathcal{V}\setminus\{v\}}(v, \rho^{\mathrm{A}})}{U_{\emptyset}(v, \rho^{\mathrm{A}})} \leq 1 - \left(1 - \max_{w \in \mathcal{V}} p_w\right)^d.$$

Finally, if $\rho^{\mathrm{A}} = \mathbf{0}$, $U(\cdot, \rho^{\mathrm{A}})$ is identical to zero. Therefore, $c = 0$ by definition, ensuring that the bound also holds. $\square$

Proof of Proposition 6. We first show that if $c = 1$, then deciding whether the optimal value of (DBR) is zero is NP-hard. To this aim, we consider the same reduction from VERTEX COVER of Proposition 4, but we modify the values of the detection probabilities to $p_v = 1$ for every $v \in \mathcal{V}$. We note that for the constructed instance, every $v \in \mathcal{V}$ satisfies

$$\begin{aligned} U_{\mathcal{V}\setminus\{v\}}(v, \rho^{\mathrm{A}}) &= U(\mathcal{V} \setminus \{v\}, \rho^{\mathrm{A}}) - U(\mathcal{V}, \rho^{\mathrm{A}}) = 0, \\ U_{\emptyset}(v, \rho^{\mathrm{A}}) &= U(\emptyset, \rho^{\mathrm{A}}) - U(\{v\}, \rho^{\mathrm{A}}) = \frac{|\mathcal{C}_v|}{m+1}. \end{aligned}$$

Therefore, $c = 1$. Clearly, G has a vertex cover of size at most k if and only if there exists a detector positioning $S \in \mathcal{D}$ such that $U(S, \rho^{\mathrm{A}}) = 0$, which concludes the reduction.

Next, let $S^* \in \mathcal{D}$ be a best response of an instance of (DBR) with $c = 1$. Given any $\alpha \geq 1$, suppose there exists an algorithm that runs in polynomial time and computes a detector positioning $\widehat{S}$ satisfying

$$U(S^*, \rho^{\mathrm{A}}) \leq U(\widehat{S}, \rho^{\mathrm{A}}) \leq \alpha\, U(S^*, \rho^{\mathrm{A}}).$$

These inequalities imply that $U(S^*, \rho^{\mathrm{A}}) = 0$ if an only if $U(\widehat{S}, \rho^{\mathrm{A}}) = 0$. In other words, we can use such algorithm to decide whether the optimal value of (DBR) is zero in polynomial time, which from the above reduction, is not possible unless P=NP. $\square$

Proof of Proposition 7. Let $\rho^{\mathrm{A}} \in \Theta_{\mathrm{A}}$ be fixed. First, we observe that $S^* \in \mathcal{D}$ minimizes the expected number of undetected attacks $U(S, \rho^{\mathrm{A}})$ over $\mathcal{D}$ if and only if S^* maximizes the expected number of *detected* attacks $D(S, \rho^{\mathrm{A}}) := \sum_{e \in \mathcal{E}} \rho_e^{\mathrm{A}} - U(S, \rho^{\mathrm{A}})$ over $\mathcal{D}$.

Since $U(\cdot, \rho^{\mathrm{A}})$ is upper bounded by $\sum_{e \in \mathcal{E}} \rho_e^{\mathrm{A}}$, as well as nonincreasing and supermodular (Lemma 10), it follows that $D(\cdot, \rho^{\mathrm{A}})$ is nonnegative, nondecreasing and submodular. The curvature parameter for the latter function is defined as

$$c := 1 - \min_{v \in \mathcal{V} \colon D_{\emptyset}(v, \rho^{\mathrm{A}}) > 0} \frac{D_{\mathcal{V} \setminus \{v\}}(v, \rho^{\mathrm{A}})}{D_{\emptyset}(v, \rho^{\mathrm{A}})},$$

where $D_S(v, \rho^{\mathrm{A}}) := D(S \cup \{v\}, \rho^{\mathrm{A}}) - D(S, \rho^{\mathrm{A}})$ is the marginal *increase* in the expected number of detected attacks when $v \in \mathcal{V}$ is added to the set of detector locations $S \subseteq \mathcal{V}$. We note that $U_S(v, \rho^{\mathrm{A}}) = D_S(v, \rho^{\mathrm{A}})$ for every $S \subseteq \mathcal{V}$ and $v \in \mathcal{V}$, so both $U(\cdot, \rho^{\mathrm{A}})$ and $D(\cdot, \rho^{\mathrm{A}})$ share the same curvature parameter c.

Let S be the set of nodes selected by the forward greedy algorithm in a given iteration. If $|S| < r_{\mathrm{D}}$, in the next iteration the algorithm selects a node $v \in \arg\max_{w \in \mathcal{V}} U_S(w, \rho^{\mathrm{A}}) = \arg\max_{w \in \mathcal{V}} D_S(w, \rho^{\mathrm{A}})$, and updates S to $S \cup \{v\}$. In other words, the algorithm selects a node maximizing the marginal increase in the expected number of detected attacks against ρ^{A}. Therefore, running the forward greedy algorithm on $U(\cdot, \rho^{\mathrm{A}})$ is equivalent to running it on $D(\cdot, \rho^{\mathrm{A}})$. By [142], the forward greedy algorithm returns a detector positioning $\widehat{S} \in \mathcal{D}$ satisfying

$$D(\widehat{S}, \rho^{\mathrm{A}}) \geq \left(\frac{1 - e^{-c}}{c} \right) D(S^*, \rho^{\mathrm{A}}).$$

Rearranging terms, we obtain

$$\begin{aligned} U(\widehat{S}, \rho^{\mathrm{A}}) &\leq \left(\frac{1 - e^{-c}}{c} \right) U(S^*, \rho^{\mathrm{A}}) + \left(1 - \frac{1 - e^{-c}}{c} \right) \sum_{e \in \mathcal{E}} \rho_e^{\mathrm{A}} \\ &\leq \left(\frac{1 - e^{-c}}{c} \right) U(S^*, \rho^{\mathrm{A}}) + \left(1 - \frac{1 - e^{-c}}{c} \right) r_{\mathrm{A}}. \end{aligned}$$

□

Proof of Theorem 6. Let $(\widehat{\sigma}^{\mathrm{D}}, \widehat{\lambda}, \widehat{\gamma}) \in \mathbb{R}^{\mathcal{D}} \times \mathbb{R}^{\mathcal{E}} \times \mathbb{R}$ and $(\widehat{\rho}^{\mathrm{A}}, \widehat{\nu}) \in \mathbb{R}^{\mathcal{E}} \times \mathbb{R}$ be optimal primal and dual solutions of the restricted master problem (LP($\mathcal{I}$)) upon termination of the CG algorithm with α-approximate best response, that is, when the α-approximate best response $\widehat{S}$ satisfies $\bar{c}_{\widehat{S}} \geq -\varepsilon$.

First, we argue that $\widehat{\nu} = U(\widehat{\sigma}^{\mathrm{D}}, \widehat{\rho}^{\mathrm{A}})$. To this aim, we recall that (LP($\mathcal{I}$)) is equivalent to (LP($\mathcal{D}$)) with the constraint that the support of the variable σ^{D} is contained within the set of columns $\mathcal{I}$. We let $\mathrm{Supp}(\sigma^{\mathrm{D}}) \coloneqq \{S \in \mathcal{D} :\ \sigma^{\mathrm{D}}_S > 0\}$ denote the support of σ^{D}. Then, from similar arguments to those given in the proof of Proposition 3, we can show that (LP($\mathcal{I}$)) and its dual are respectively equivalent to

$$\min_{\substack{\sigma^{\mathrm{D}} \in \Delta_{\mathrm{D}}: \\ \mathrm{Supp}(\sigma^{\mathrm{D}}) \subseteq \mathcal{I}}} \max_{\rho^{\mathrm{A}} \in \Theta_{\mathrm{A}}} U(\sigma^{\mathrm{D}}, \rho^{\mathrm{A}}) \quad \text{and} \quad \max_{\rho^{\mathrm{A}} \in \Theta_{\mathrm{A}}} \min_{\substack{\sigma^{\mathrm{D}} \in \Delta_{\mathrm{D}}: \\ \mathrm{Supp}(\sigma^{\mathrm{D}}) \subseteq \mathcal{I}}} U(\sigma^{\mathrm{D}}, \rho^{\mathrm{A}}).$$

Furthermore, $\widehat{\sigma}^{\mathrm{D}}$ and $\widehat{\rho}^{\mathrm{A}}$ are respectively optimal solutions of the above minmax and maxmin problems. Therefore,

$$\widehat{\nu} = \min_{\substack{\sigma^{\mathrm{D}} \in \Delta_{\mathrm{D}}: \\ \mathrm{Supp}(\sigma^{\mathrm{D}}) \subseteq \mathcal{I}}} U(\sigma^{\mathrm{D}}, \widehat{\rho}^{\mathrm{A}}) \leq U(\widehat{\sigma}^{\mathrm{D}}, \widehat{\rho}^{\mathrm{A}}) \leq \max_{\rho^{\mathrm{A}} \in \Theta_{\mathrm{A}}} U(\widehat{\sigma}^{\mathrm{D}}, \rho^{\mathrm{A}}) = \widehat{\gamma} + \sum_{e \in \mathcal{E}} \widehat{\lambda}_e, \tag{C.16}$$

where the equalities hold since $\widehat{\gamma} + \sum_{e \in \mathcal{E}} \widehat{\lambda}_e$ and $\widehat{\nu}$ are the optimal values of (LP($\mathcal{I}$)) and its dual, respectively. By strong duality, we have $\widehat{\nu} = \widehat{\gamma} + \sum_{e \in \mathcal{E}} \widehat{\lambda}_e$, and therefore all the inequalities in (C.16) are in fact *equalities*. Hence, $\widehat{\nu} = U(\widehat{\sigma}^{\mathrm{D}}, \widehat{\rho}^{\mathrm{A}})$.

Next, let $S^* \in \mathcal{D}$ be a pure best response against $\widehat{\rho}^{\mathrm{A}}$, and let $\widehat{S} \in \mathcal{D}$ be the approximate best response returned by the α-approximation algorithm for (DBR). Then,

$$U(\widehat{\sigma}^{\mathrm{D}}, \widehat{\rho}^{\mathrm{A}}) = \widehat{\nu} \leq U(\widehat{S}, \widehat{\rho}^{\mathrm{A}}) + \varepsilon \leq \alpha\, U(S^*, \widehat{\rho}^{\mathrm{A}}) + \varepsilon. \tag{C.17}$$

The first inequality holds by the algorithm's termination criterion, which is $\bar{c}_{\widehat{S}} = -\widehat{\nu} + U(\widehat{S}, \widehat{\rho}^{\mathrm{A}}) \geq -\varepsilon$, whereas the second inequality is due to $\widehat{S}$ being an α-approximation of

(DBR). On the other hand, from the optimality of S^*, it follows that

$$U(S^*, \widehat{\rho}^{\mathrm{A}}) \leq \sum_{S \in \mathcal{D}} \sigma^{\mathrm{D}}_S U(S, \widehat{\rho}^{\mathrm{A}}) = U(\sigma^{\mathrm{D}}, \widehat{\rho}^{\mathrm{A}}), \quad \forall\, \sigma^{\mathrm{D}} \in \Delta_{\mathrm{D}}. \tag{C.18}$$

Therefore, we have the following bounds for unilateral deviation from $(\widehat{\sigma}^{\mathrm{D}}, \widehat{\rho}^{\mathrm{A}})$:

$$U(\widehat{\sigma}^{\mathrm{D}}, \widehat{\rho}^{\mathrm{A}}) \overset{\text{(C.17)}}{\leq} \alpha\, U(S^*, \widehat{\rho}^{\mathrm{A}}) + \varepsilon \overset{\text{(C.18)}}{\leq} \alpha\, U(\sigma^{\mathrm{D}}, \widehat{\rho}^{\mathrm{A}}) + \varepsilon, \qquad \forall\, \sigma^{\mathrm{D}} \in \Delta_{\mathrm{D}}, \tag{C.19}$$

$$U(\widehat{\sigma}^{\mathrm{D}}, \widehat{\rho}^{\mathrm{A}}) \overset{\text{(C.16)}}{=} \max_{\rho^{\mathrm{A}} \in \Theta_{\mathrm{A}}} U(\widehat{\sigma}^{\mathrm{D}}, \rho^{\mathrm{A}}) \geq U(\widehat{\sigma}^{\mathrm{D}}, \rho^{\mathrm{A}}), \qquad \forall\, \rho^{\mathrm{A}} \in \Theta_{\mathrm{A}}. \tag{C.20}$$

Now, let $(\sigma^{\mathrm{D}^*}, \rho^{\mathrm{A}^*})$ be an equilibrium of the game Γ. Using the definition of equilibrium, we have

$$U(\sigma^{\mathrm{D}^*}, \rho^{\mathrm{A}^*}) \leq U(\widehat{\sigma}^{\mathrm{D}}, \rho^{\mathrm{A}^*}) \overset{\text{(C.20)}}{\leq} U(\widehat{\sigma}^{\mathrm{D}}, \widehat{\rho}^{\mathrm{A}}) \overset{\text{(C.19)}}{\leq} \alpha\, U(\sigma^{\mathrm{D}^*}, \widehat{\rho}^{\mathrm{A}}) + \varepsilon \leq \alpha\, U(\sigma^{\mathrm{D}^*}, \rho^{\mathrm{A}^*}) + \varepsilon. \tag{C.21}$$

Thus, $\widehat{\sigma}^{\mathrm{D}}$ and $\widehat{\rho}^{\mathrm{A}}$ respectively satisfy the following bounds for the worst-case expected number of undetected attacks with respect to the value of the game $U(\sigma^{\mathrm{D}^*}, \rho^{\mathrm{A}^*})$:

$$U(\sigma^{\mathrm{D}^*}, \rho^{\mathrm{A}^*}) \overset{\text{(C.21)}}{\leq} U(\widehat{\sigma}^{\mathrm{D}}, \widehat{\rho}^{\mathrm{A}}) \overset{\text{(C.16)}}{=} \max_{\rho^{\mathrm{A}} \in \Theta_{\mathrm{A}}} U(\widehat{\sigma}^{\mathrm{D}}, \rho^{\mathrm{A}}) \overset{\text{(C.21)}}{\leq} \alpha\, U(\sigma^{\mathrm{D}^*}, \rho^{\mathrm{A}^*}) + \varepsilon,$$

$$\frac{1}{\alpha}\left(U(\sigma^{\mathrm{D}^*}, \rho^{\mathrm{A}^*}) - \varepsilon\right) \overset{\text{(C.21)}}{\leq} \frac{1}{\alpha}\left(U(\widehat{\sigma}^{\mathrm{D}}, \widehat{\rho}^{\mathrm{A}}) - \varepsilon\right) \overset{\text{(C.19)}}{\leq} \min_{\sigma^{\mathrm{D}} \in \Delta_{\mathrm{D}}} U(\sigma^{\mathrm{D}}, \widehat{\rho}^{\mathrm{A}}) \leq U(\sigma^{\mathrm{D}^*}, \widehat{\rho}^{\mathrm{A}}) \leq U(\sigma^{\mathrm{D}^*}, \rho^{\mathrm{A}^*}).$$

□

Before delving into the proof of Theorem 7, we establish some key lemmas. First, we present a standard bound essential for analyzing MWU algorithms. Our proof mirrors that given in [139] for the MWU algorithm applied to restricted distributions.

Lemma 25. *Algorithm 4 satisfies the following bound for the sum of the expected number* *of undetected attacks across τ iterations:*

$$\sum_{t=1}^{\tau} U(S^{(t)}, \rho^{(t)}) \geq \max_{\rho \in \Theta_{\mathrm{A}}} \left\{ \sum_{t=1}^{\tau} U(S^{(t)}, \rho) - \frac{1}{\eta} \mathrm{D}_{\mathrm{RE}} \left(\rho \,\middle\|\, \rho^{(1)} \right) \right\} - \eta \sum_{t=1}^{\tau} \sum_{e \in \mathcal{E}} \rho_e^{(t)} \left(u(S^{(t)}, e) \right)^2 .$$

Proof of Lemma 25. Let $\rho \in \Theta_{\mathrm{A}}$. First, we consider the following Generalized Pythagorean inequality [see e.g., 147, 123]:

$$\mathrm{D}_{\mathrm{RE}} \left(\rho \,\middle\|\, \rho^{(t+1)} \right) + \mathrm{D}_{\mathrm{RE}} \left(\rho^{(t+1)} \,\middle\|\, \widetilde{\rho}^{(t+1)} \right) \leq \mathrm{D}_{\mathrm{RE}} \left(\rho \,\middle\|\, \widetilde{\rho}^{(t+1)} \right) . \tag{C.22}$$

Since $\mathrm{D}_{\mathrm{RE}} \left(\rho^{(t+1)} \,\middle\|\, \widetilde{\rho}^{(t+1)} \right) \geq 0$, it follows that

$$\begin{aligned}
\mathrm{D}_{\mathrm{RE}} \left(\rho \,\middle\|\, \rho^{(t+1)} \right) - \mathrm{D}_{\mathrm{RE}} \left(\rho \,\middle\|\, \rho^{(t)} \right) &= \mathrm{D}_{\mathrm{RE}} \left(\rho \,\middle\|\, \rho^{(t+1)} \right) + \mathrm{D}_{\mathrm{RE}} \left(\rho^{(t+1)} \,\middle\|\, \widetilde{\rho}^{(t+1)} \right) \\
&\quad - \mathrm{D}_{\mathrm{RE}} \left(\rho^{(t+1)} \,\middle\|\, \widetilde{\rho}^{(t+1)} \right) - \mathrm{D}_{\mathrm{RE}} \left(\rho \,\middle\|\, \rho^{(t)} \right) \\
&\overset{\text{(C.22)}}{\leq} \mathrm{D}_{\mathrm{RE}} \left(\rho \,\middle\|\, \widetilde{\rho}^{(t+1)} \right) - \mathrm{D}_{\mathrm{RE}} \left(\rho^{(t+1)} \,\middle\|\, \widetilde{\rho}^{(t+1)} \right) - \mathrm{D}_{\mathrm{RE}} \left(\rho \,\middle\|\, \rho^{(t)} \right) \\
&\leq \mathrm{D}_{\mathrm{RE}} \left(\rho \,\middle\|\, \widetilde{\rho}^{(t+1)} \right) - \mathrm{D}_{\mathrm{RE}} \left(\rho \,\middle\|\, \rho^{(t)} \right) .
\end{aligned}$$

Thus,

$$\mathrm{D}_{\mathrm{RE}} \left(\rho \,\middle\|\, \rho^{(t+1)} \right) - \mathrm{D}_{\mathrm{RE}} \left(\rho \,\middle\|\, \rho^{(t)} \right) \leq \mathrm{D}_{\mathrm{RE}} \left(\rho \,\middle\|\, \widetilde{\rho}^{(t+1)} \right) - \mathrm{D}_{\mathrm{RE}} \left(\rho \,\middle\|\, \rho^{(t)} \right) . \tag{C.23}$$

Next, we derive an upper bound for the right hand side $\mathrm{D}_{\mathrm{RE}} \left(\rho \,\middle\|\, \widetilde{\rho}^{(t+1)} \right) - \mathrm{D}_{\mathrm{RE}} \left(\rho \,\middle\|\, \rho^{(t)} \right)$. From the update step of the algorithm (Lines 4–5), we have $\widetilde{\rho}_e^{(t+1)} = \rho_e^{(t)} \exp \left(\eta u(S^{(t)}, e) \right)$

for every $e \in \mathcal{E}$. Then,

$$
\begin{aligned}
\mathrm{D}_{\mathrm{RE}}\left(\rho \,\middle\|\, \widetilde{\rho}^{(t+1)}\right) - \mathrm{D}_{\mathrm{RE}}\left(\rho \,\middle\|\, \rho^{(t)}\right) &= \sum_{e\in\mathcal{E}} \rho_e \ln \frac{\rho_e^{(t)}}{\widetilde{\rho}_e^{(t+1)}} + \widetilde{\rho}_e^{(t+1)} - \rho_e^{(t)} \\
&= \sum_{e\in\mathcal{E}} \rho_e \ln \exp\left(-\eta u(S^{(t)}, e)\right) + \rho_e^{(t)} \left(\exp\left(\eta u(S^{(t)}, e)\right) - 1\right) \\
&\leq -\eta \sum_{e\in\mathcal{E}} \rho_e u(S^{(t)}, e) \\
&\quad + \sum_{e\in\mathcal{E}} \rho_e^{(t)} \left(1 + \eta u(S^{(t)}, e) + \eta^2 \left(u(S^{(t)}, e)\right)^2 - 1\right) \\
&= -\eta U(S^{(t)}, \rho) + \eta U(S^{(t)}, \rho^{(t)}) + \eta^2 \sum_{e\in\mathcal{E}} \rho_e^{(t)} \left(u(S^{(t)}, e)\right)^2,
\end{aligned}
$$

where we used the inequality $e^x \leq 1 + x + x^2$ for $x \in [-1, 1]$. Hence,

$$
\mathrm{D}_{\mathrm{RE}}\left(\rho \,\middle\|\, \widetilde{\rho}^{(t+1)}\right) - \mathrm{D}_{\mathrm{RE}}\left(\rho \,\middle\|\, \rho^{(t)}\right) \leq -\eta U(S^{(t)}, \rho) + \eta U(S^{(t)}, \rho^{(t)}) + \eta^2 \sum_{e\in\mathcal{E}} \rho_e^{(t)} \left(u(S^{(t)}, e)\right)^2. \tag{C.24}
$$

Combining inequalities (C.23) and (C.24), we obtain

$$
\mathrm{D}_{\mathrm{RE}}\left(\rho \,\middle\|\, \rho^{(t+1)}\right) - \mathrm{D}_{\mathrm{RE}}\left(\rho \,\middle\|\, \rho^{(t)}\right) \leq -\eta U(S^{(t)}, \rho) + \eta U(S^{(t)}, \rho^{(t)}) + \eta^2 \sum_{e\in\mathcal{E}} \rho_e^{(t)} \left(u(S^{(t)}, e)\right)^2.
$$

Summing over t, rearranging terms and using that $\mathrm{D}_{\mathrm{RE}}\left(\rho \,\middle\|\, \rho^{(\tau+1)}\right) \geq 0$, we get

$$
\eta \sum_{t=1}^{\tau} U(S^{(t)}, \rho^{(t)}) \geq \eta \sum_{t=1}^{\tau} U(S^{(t)}, \rho) - \mathrm{D}_{\mathrm{RE}}\left(\rho \,\middle\|\, \rho^{(1)}\right) - \eta^2 \sum_{t=1}^{\tau} \sum_{e\in\mathcal{E}} \rho_e^{(t)} \left(u(S^{(t)}, e)\right)^2.
$$

Finally, since $\rho \in \Theta_{\mathrm{A}}$ is arbitrary, we divide by $\eta > 0$ to conclude the result. □

The next lemma provides an upper bound for the unnormalized relative entropy between an arbitrary vector in Θ_{A} and the initial marginal attack strategy of Algorithm 4.

Lemma 26. *Let* $\rho^{(1)} \in \Theta_{\mathrm{A}}$ *given by* $\rho_e^{(1)} := r_{\mathrm{A}}/m$ *for every* $e \in \mathcal{E}$. *Then,*

$$\mathrm{D}_{\mathrm{RE}}\left(\rho \,\middle\|\, \rho^{(1)}\right) \leq r_{\mathrm{A}} \max\left\{\ln \frac{m}{r_{\mathrm{A}}}, 1\right\}, \quad \forall\, \rho \in \Theta_{\mathrm{A}}.$$

Proof of Lemma 26. The convexity of the function $x \to x\ln(x)$ in $\mathbb{R}_{\geq 0}$—where we let $0\ln(0) := 0$— implies that the function $\mathrm{D}_{\mathrm{RE}}\left(\cdot \,\middle\|\, \rho^{(1)}\right) = \sum_{e\in\mathcal{E}}\left(\rho_e \ln\left(\frac{m}{r_{\mathrm{A}}}\rho_e\right) + \frac{r_{\mathrm{A}}}{m} - \rho_e\right)$ is convex in Θ_{A}, which is a convex set. Furthermore, Θ_{A} is the convex hull of the characteristic vectors of the sets in $\mathcal{A}$ [see, for example, 83]. Let $T \in \mathcal{A}$, and $\mathbb{1}_T \in \{0,1\}^{\mathcal{E}}$ be its characteristic vector, such that for every $e \in \mathcal{E}$, $(\mathbb{1}_T)_e = 1$ if $e \in T$ and $(\mathbb{1}_T)_e = 0$ otherwise. Then,

$$\mathrm{D}_{\mathrm{RE}}\left(\mathbb{1}_T \,\middle\|\, \rho^{(1)}\right) = \sum_{e\in T}\left(\ln\left(\frac{m}{r_{\mathrm{A}}}\right) + \frac{r_{\mathrm{A}}}{m} - 1\right) + \sum_{e\notin T}\frac{r_{\mathrm{A}}}{m} = |T|\ln\left(\frac{m}{r_{\mathrm{A}}}\right) - |T| + r_{\mathrm{A}}.$$

The last term is upper bounded by r_{A} if $\ln(m/r_{\mathrm{A}}) < 1$, and by $r_{\mathrm{A}}\ln(m/r_{\mathrm{A}})$ otherwise. Since T is arbitrary, we obtain

$$\forall\, T \in \mathcal{A}, \quad \mathrm{D}_{\mathrm{RE}}\left(\mathbb{1}_T \,\middle\|\, \rho^{(1)}\right) \leq r_{\mathrm{A}} \max\left\{\ln \frac{m}{r_{\mathrm{A}}}, 1\right\}. \tag{C.25}$$

Now, let $\rho \in \Theta_{\mathrm{A}}$. Then, $\rho = \sum_{i=1}^k \lambda_i \mathbb{1}_{T_i}$ for some sets $T_1, \ldots, T_k \in \mathcal{A}$ and scalars $\lambda_1, \cdots, \lambda_k \in [0,1]$ such that $\sum_{i=1}^k \lambda_i = 1$. From the convexity of $\mathrm{D}_{\mathrm{RE}}\left(\cdot \,\middle\|\, \rho^{(1)}\right)$, it follows that

$$\mathrm{D}_{\mathrm{RE}}\left(\rho \,\middle\|\, \rho^{(1)}\right) = \mathrm{D}_{\mathrm{RE}}\left(\sum_{i=1}^k \lambda_i \mathbb{1}_{T_i} \,\middle\|\, \rho^{(1)}\right) \leq \sum_{i=1}^k \lambda_i \mathrm{D}_{\mathrm{RE}}\left(\mathbb{1}_{T_i} \,\middle\|\, \rho^{(1)}\right) \overset{\text{(C.25)}}{\leq} r_{\mathrm{A}} \max\left\{\ln \frac{m}{r_{\mathrm{A}}}, 1\right\},$$

which concludes the proof. □

Now we are ready to prove Theorem 7.

Proof of Theorem 7. From Lemma 25, after τ iterations of Algorithm 4, we have

$$\sum_{t=1}^{\tau} U(S^{(t)}, \rho^{(t)}) \geq \max_{\rho \in \Theta_{\mathrm{A}}} \left\{ \sum_{t=1}^{\tau} U(S^{(t)}, \rho) - \frac{1}{\eta} \mathrm{D}_{\mathrm{RE}} \left(\rho \,\middle\|\, \rho^{(1)} \right) \right\} - \eta \sum_{t=1}^{\tau} \sum_{e \in \mathcal{E}} \rho_e^{(t)} \left(u(S^{(t)}, e) \right)^2 .$$

Using Lemma 26 to bound the term $\mathrm{D}_{\mathrm{RE}} \left(\rho \,\middle\|\, \rho^{(1)} \right)$, and the facts that $u(S^{(t)}, e) \in [0,1]$ and $\rho^{(t)} \in \Theta_{\mathrm{A}}$, we obtain

$$\frac{1}{\tau} \sum_{t=1}^{\tau} U(S^{(t)}, \rho^{(t)}) \geq \max_{\rho \in \Theta_{\mathrm{A}}} \frac{1}{\tau} \sum_{t=1}^{\tau} U(S^{(t)}, \rho) - \frac{1}{\eta\tau} r_{\mathrm{A}} \max \left\{ \ln \frac{m}{r_{\mathrm{A}}}, 1 \right\} - \eta r_{\mathrm{A}}.$$

Therefore,

$$\max_{\rho \in \Theta_{\mathrm{A}}} \frac{1}{\tau} \sum_{t=1}^{\tau} U(S^{(t)}, \rho) - \frac{1}{\tau} \sum_{t=1}^{\tau} U(S^{(t)}, \rho^{(t)}) \leq \frac{1}{\eta\tau} r_{\mathrm{A}} \max \left\{ \ln \frac{m}{r_{\mathrm{A}}}, 1 \right\} + \eta r_{\mathrm{A}}. \tag{C.26}$$

The right hand side in (C.26) is minimized when $\eta = \sqrt{\max \left\{ \ln(m/r_{\mathrm{A}}), 1 \right\} / \tau}$, which yields the bound

$$\max_{\rho \in \Theta_{\mathrm{A}}} \frac{1}{\tau} \sum_{t=1}^{\tau} U(S^{(t)}, \rho) - \frac{1}{\tau} \sum_{t=1}^{\tau} U(S^{(t)}, \rho^{(t)}) \leq 2 r_{\mathrm{A}} \sqrt{\frac{\max \left\{ \ln \frac{m}{r_{\mathrm{A}}}, 1 \right\}}{\tau}}.$$

Thus, the right hand side in (C.26) is at most ε if

$$\tau \geq \frac{4 r_{\mathrm{A}}^2 \max \left\{ \ln \frac{m}{r_{\mathrm{A}}}, 1 \right\}}{\varepsilon^2}.$$

Thus, after $\tau \geq 4 r_{\mathrm{A}}^2 \max \left\{ \ln(m/r_{\mathrm{A}}), 1 \right\} / \varepsilon^2$ iterations, we obtain the following lower bound:

$$\max_{\rho \in \Theta_{\mathrm{A}}} \frac{1}{\tau} \sum_{t=1}^{\tau} U(S^{(t)}, \rho) - \varepsilon \leq \frac{1}{\tau} \sum_{t=1}^{\tau} U(S^{(t)}, \rho^{(t)}). \tag{C.27}$$

On the other hand, since $S^{(t)}$ is an α-approximation of D's best response against $\rho^{(t)}$, it follows that

$$\frac{1}{\tau}\sum_{t=1}^{\tau} U(S^{(t)},\rho^{(t)}) \leq \frac{1}{\tau}\sum_{t=1}^{\tau} \alpha\, U(S,\rho^{(t)}), \qquad \forall\, S \in \mathcal{D}.$$

Then, by linearity of the expectation, we deduce that

$$\frac{1}{\tau}\sum_{t=1}^{\tau} U(S^{(t)},\rho^{(t)}) \leq \alpha \min_{\sigma\in\Delta_{\mathrm{D}}} \frac{1}{\tau}\sum_{t=1}^{\tau} U(\sigma,\rho^{(t)}). \tag{C.28}$$

Thus, combining (C.27) and (C.28), after $\tau \geq 4r_{\mathrm{A}}^2 \max\{\ln(m/r_{\mathrm{A}}),1\}/\varepsilon^2$ iterations we obtain

$$\max_{\rho\in\Theta_{\mathrm{A}}} \frac{1}{\tau}\sum_{t=1}^{\tau} U(S^{(t)},\rho) - \varepsilon \leq \frac{1}{\tau}\sum_{t=1}^{\tau} U(S^{(t)},\rho^{(t)}) \leq \alpha \min_{\sigma\in\Delta_{\mathrm{D}}} \frac{1}{\tau}\sum_{t=1}^{\tau} U(\sigma,\rho^{(t)}). \tag{C.29}$$

Next, using that $\widehat{\sigma}^{\mathrm{D}} \coloneqq \frac{1}{\tau}\sum_{t=1}^{\tau} \mathbb{1}_{S^{(t)}}$ and $\widehat{\rho}^{\mathrm{A}} \coloneqq \frac{1}{\tau}\sum_{t=1}^{\tau} \rho^{(t)}$, we have the following identities:

$$\frac{1}{\tau}\sum_{t=1}^{\tau} U(S^{(t)},\rho) = \sum_{S\in\mathcal{D}} \frac{1}{\tau}\sum_{t=1}^{\tau} \mathbb{1}_{S^{(t)}}(S) U(S,\rho) = \sum_{S\in\mathcal{D}} \widehat{\sigma}^{\mathrm{D}}_S U(S,\rho) = U(\widehat{\sigma}^{\mathrm{D}},\rho), \quad \forall\, \rho \in \Theta_{\mathrm{A}}. \tag{C.30}$$

$$\frac{1}{\tau}\sum_{t=1}^{\tau} U(\sigma,\rho^{(t)}) = \sum_{e\in\mathcal{E}} \frac{1}{\tau}\sum_{t=1}^{\tau} \rho_e^{(t)} U(\sigma,e) = \sum_{e\in\mathcal{E}} \widehat{\rho}^{\mathrm{A}}_e U(\sigma,e) = U(\sigma,\widehat{\rho}^{\mathrm{A}}), \qquad \forall\, \sigma \in \Delta_{\mathrm{D}}. \tag{C.31}$$

Let $(\sigma^{\mathrm{D}^*},\rho^{\mathrm{A}^*})$ be an equilibrium of the game Γ. Then,

$$\max_{\rho\in\Theta_{\mathrm{A}}} U(\widehat{\sigma}^{\mathrm{D}},\rho) \geq U(\widehat{\sigma}^{\mathrm{D}},\rho^{\mathrm{A}^*}) \geq U(\sigma^{\mathrm{D}^*},\rho^{\mathrm{A}^*}), \tag{C.32}$$

$$\min_{\sigma\in\Delta_{\mathrm{D}}} U(\sigma,\widehat{\rho}^{\mathrm{A}}) \leq U(\sigma^{\mathrm{D}^*},\widehat{\rho}^{\mathrm{A}}) \leq U(\sigma^{\mathrm{D}^*},\rho^{\mathrm{A}^*}). \tag{C.33}$$

Therefore,

$$U(\sigma^{\mathrm{D}^*},\rho^{\mathrm{A}^*})-\varepsilon \overset{(C.32)}{\leq} \max_{\rho\in\Theta_{\mathrm{A}}} U(\widehat{\sigma}^{\mathrm{D}},\rho)-\varepsilon \overset{(C.29),(C.30),(C.31)}{\leq} \alpha \min_{\sigma^{\mathrm{D}}\in\Delta_{\mathrm{D}}} U(\sigma^{\mathrm{D}},\widehat{\rho}^{\mathrm{A}}) \overset{(C.33)}{\leq} \alpha U(\sigma^{\mathrm{D}^*},\rho^{\mathrm{A}^*}), \tag{C.34}$$

from which we deduce the following worst-case bounds for the expected number of undetected attacks for $\widehat{\sigma}^{\mathrm{D}}$ and $\widehat{\rho}^{\mathrm{A}}$ respectively:

$$U(\sigma^{\mathrm{D}^*},\rho^{\mathrm{A}^*}) \leq \max_{\rho^{\mathrm{A}}\in\Theta_{\mathrm{A}}} U(\widehat{\sigma}^{\mathrm{D}},\rho^{\mathrm{A}}) \leq \alpha U(\sigma^{\mathrm{D}^*},\rho^{\mathrm{A}^*})+\varepsilon,$$

$$\frac{1}{\alpha}\big(U(\sigma^{\mathrm{D}^*},\rho^{\mathrm{A}^*})-\varepsilon\big) \leq \min_{\sigma^{\mathrm{D}}\in\Delta_{\mathrm{D}}} U(\sigma^{\mathrm{D}},\widehat{\rho}^{\mathrm{A}}) \leq U(\sigma^{\mathrm{D}^*},\rho^{\mathrm{A}^*}).$$

Finally, the bounds for unilateral deviation from $(\widehat{\sigma}^{\mathrm{D}},\widehat{\rho}^{\mathrm{A}})$ follow from

$$U(\widehat{\sigma}^{\mathrm{D}},\widehat{\rho}^{\mathrm{A}}) \leq \max_{\rho\in\Theta_{\mathrm{A}}} U(\widehat{\sigma}^{\mathrm{D}},\rho) \overset{(C.34)}{\leq} \alpha \min_{\sigma^{\mathrm{D}}\in\Delta_{\mathrm{D}}} U(\sigma^{\mathrm{D}},\widehat{\rho}^{\mathrm{A}})+\varepsilon \leq \alpha U(\widehat{\sigma}^{\mathrm{D}},\widehat{\rho}^{\mathrm{A}})+\varepsilon,$$

$$U(\widehat{\sigma}^{\mathrm{D}},\widehat{\rho}^{\mathrm{A}}) \geq \min_{\sigma^{\mathrm{D}}\in\Delta_{\mathrm{D}}} U(\sigma^{\mathrm{D}},\widehat{\rho}^{\mathrm{A}}) \overset{(C.34)}{\geq} \frac{1}{\alpha}\left(\max_{\rho\in\Theta_{\mathrm{A}}} U(\widehat{\sigma}^{\mathrm{D}},\rho)-\varepsilon\right) \geq \frac{1}{\alpha}\big(U(\widehat{\sigma}^{\mathrm{D}},\widehat{\rho}^{\mathrm{A}})-\varepsilon\big).$$

□

We now delve into the proof of Theorem 8. To this aim, we first need the following auxiliary lemmas.

Lemma 27. *For every $a>0$ and $b\in\mathbb{R}$, the real function*

$$f(x) \coloneqq \begin{cases} x\ln\frac{x}{a}+a-(1-b)x & \textit{if } x>0, \\ a & \textit{if } x=0, \end{cases}$$

is strictly convex in $\mathbb{R}_{\geq 0}$ and has a unique minimizer, given by $y^*=\exp(-b)a$. Consequently, f has a unique minimizer in the interval $[0,1]$, given by $x^*=\min\{\exp(-b)a,1\}$.

Proof of Lemma 27. The first and second derivatives of f in $\big(0,+\infty\big)$ are given by $f'(x)=$

$\ln(x/a) - b$ and $f''(x) = 1/x$, respectively. Therefore, f is strictly convex and its global minimizer can be obtained by setting $f'(y^*) = 0$, which yields $y^* = \exp(-b)a$. $\square$

Lemma 28. *Let $\widetilde{\rho} \in \mathbb{R}^{\mathcal{E}}_{>0}$, and let us sort $\mathcal{E}$ such that $\widetilde{\rho}_{e_1} \geq \cdots \geq \widetilde{\rho}_{e_m}$, breaking ties arbitrarily. Then:*

- *The function $s : \{0, \ldots, m\} \to \mathbb{R}_{\geq 0}$ defined by $s(k) := k + (1/\widetilde{\rho}_{e_k}) \sum_{j=k+1}^{m} \widetilde{\rho}_{e_j}$ (where we let $\widetilde{\rho}_{e_0} := +\infty$, so $1/\widetilde{\rho}_{e_0} := 0$) is nondecreasing. Moreover, for every $k \in \{0, \ldots, m-1\}$, $s(k) < s(k+1)$ if and only if $\widetilde{\rho}_{e_k} > \widetilde{\rho}_{e_{k+1}}$.*
- *Let $k^* := \max\{k \in \{0, \ldots, r_{\mathrm{A}}\} : s(k) \leq r_{\mathrm{A}}\}$. If $k^* \in \{0, \ldots, m-1\}$, then $\widetilde{\rho}_{e_{k^*}} > \widetilde{\rho}_{e_{k^*+1}}$.*

Proof of Lemma 28. For every $k \in \{0, \ldots, m-1\}$, it holds that

$$s(k+1) = k + 1 + \frac{1}{\widetilde{\rho}_{e_{k+1}}} \sum_{j=k+2}^{m} \widetilde{\rho}_{e_j} = k + \frac{1}{\widetilde{\rho}_{e_{k+1}}} \sum_{j=k+1}^{m} \widetilde{\rho}_{e_j} \geq k + \frac{1}{\widetilde{\rho}_{e_k}} \sum_{j=k+1}^{m} \widetilde{\rho}_{e_j} = s(k).$$

Therefore, $s(k) \leq s(k+1)$. Furthermore, the inequality above is strict if an only if $\widetilde{\rho}_{e_k} > \widetilde{\rho}_{e_{k+1}}$. Next, let us assume that $k^* \in \{0, \ldots, m-1\}$. By definition, k^* satisfies $s(k^*) \leq r_{\mathrm{A}} < s(k^*+1)$; hence $\widetilde{\rho}_{e_{k^*}} > \widetilde{\rho}_{e_{k^*+1}}$. $\square$

The following lemma is standard for results involving Bregman projections [see, for example 140, 127], and shows that such projections preserve the relative ordering of the entries within the projected vector.

Lemma 29. *Let $\widetilde{\rho} \in \mathbb{R}^{\mathcal{E}}_{>0}$, and let $\rho^* \in \mathbb{R}^{\mathcal{E}}$ be its projection with respect to the unnormalized relative entropy. For every $e, e' \in \mathcal{E}$, if $\widetilde{\rho}_e \geq \widetilde{\rho}_{e'}$, then $\rho^*_e \geq \rho^*_{e'}$.*

Proof of Lemma 29. Suppose there exist $e, e' \in \mathcal{E}$ such that $\widetilde{\rho}_e \geq \widetilde{\rho}_{e'}$ and $\rho^*_e < \rho^*_{e'}$. Let

$\rho' \in \mathbb{R}^{\mathcal{E}}$ be defined as

$$\rho'_f \coloneqq \begin{cases} \rho^*_f & \text{if } f \neq e, f \neq e', \\ \rho^*_{e'} & \text{if } f = e, \\ \rho^*_e & \text{if } f = e', \end{cases} \qquad \forall\, f \in \mathcal{E}.$$

In other words, ρ' is obtained from ρ^* by exchanging the coordinates e and e'. Then,

$$\mathrm{D}_{\mathrm{RE}}\left(\rho^* \,\|\, \widetilde{\rho}\,\right) - \mathrm{D}_{\mathrm{RE}}\left(\rho' \,\|\, \widetilde{\rho}\,\right) = \rho^*_{e'} \ln \frac{\widetilde{\rho}_e}{\widetilde{\rho}_{e'}} + \rho^*_e \ln \frac{\widetilde{\rho}_{e'}}{\widetilde{\rho}_e} = (\rho^*_{e'} - \rho_e) \ln \frac{\widetilde{\rho}_e}{\widetilde{\rho}_{e'}} \geq 0.$$

Hence, $\rho' \in \Theta_{\mathrm{A}}$ and $\mathrm{D}_{\mathrm{RE}}\left(\rho^* \,\|\, \widetilde{\rho}\,\right) \geq \mathrm{D}_{\mathrm{RE}}\left(\rho' \,\|\, \widetilde{\rho}\,\right)$. Since ρ^* minimizes $\mathrm{D}_{\mathrm{RE}}\left(\cdot \,\|\, \widetilde{\rho}\,\right)$ over Θ_{A}, we must have $\mathrm{D}_{\mathrm{RE}}\left(\rho^* \,\|\, \widetilde{\rho}\,\right) = \mathrm{D}_{\mathrm{RE}}\left(\rho' \,\|\, \widetilde{\rho}\,\right)$, and thus $\rho' \neq \rho^*$ is another minimizer, contradicting the uniqueness of the projection. $\square$

Lemma 30. *Let $\widetilde{\rho} \in \mathbb{R}^{\mathcal{E}}_{>0}$ satisfying $\sum_{e \in \mathcal{E}} \min\{\widetilde{\rho}_e, 1\} > r_{\mathrm{A}}$, and let $\rho^* \in \mathbb{R}^{\mathcal{E}}$ be its projection into Θ_{A} with respect to the unnormalized relative entropy. Then, ρ^* satisfies $\sum_{e \in \mathcal{E}} \rho^*_e = r_{\mathrm{A}}$.*

Proof of Lemma 30. By definition of the projection, we have $\rho^* \in \Theta_{\mathrm{A}}$. In particular, ρ^* satisfies $\sum_{e \in \mathcal{E}} \rho^*_e \leq r_{\mathrm{A}}$. Let us uppose that $\sum_{e \in \mathcal{E}} \rho^*_e < r_{\mathrm{A}}$. Then, there exists $e' \in \mathcal{E}$ such that $\rho^*_{e'} < \min\{\widetilde{\rho}_{e'}, 1\}$; otherwise we would have $\sum_{e \in \mathcal{E}} \rho^*_e \geq \sum_{e \in \mathcal{E}} \min\{\widetilde{\rho}_e, 1\} > r_{\mathrm{A}}$, which contradicts that $\rho^* \in \Theta_{\mathrm{A}}$. Consequently, there exists $\varepsilon > 0$ such that $\rho^*_{e'} + \varepsilon < \min\{\widetilde{\rho}_{e'}, 1\}$ and the vector $\rho' \in \mathbb{R}^{\mathcal{E}}$ defined as

$$\rho'_e \coloneqq \begin{cases} \rho^*_e & \text{if } e \neq e', \\ \rho^*_e + \varepsilon & \text{if } e = e', \end{cases} \qquad \forall\, e \in \mathcal{E},$$

belongs to Θ_{A}. Then, it follows that

$$\begin{aligned}&\mathrm{D}_{\mathrm{RE}}\left(\rho^* \,\|\, \widetilde{\rho}\,\right) - \mathrm{D}_{\mathrm{RE}}\left(\rho' \,\|\, \widetilde{\rho}\,\right) = \\ &\quad \left(\rho^*_{e'} \ln \frac{\rho^*_{e'}}{\widetilde{\rho}_{e'}} + \widetilde{\rho}_{e'} - \rho^*_{e'}\right) - \left((\rho^*_{e'} + \varepsilon) \ln \frac{\rho^*_{e'} + \varepsilon}{\widetilde{\rho}_{e'}} + \widetilde{\rho}_{e'} - (\rho^*_{e'} + \varepsilon)\right) > 0,\end{aligned}$$

where we used the fact that for every $a > 0$, the function

$$f(x) := \begin{cases} x \ln \frac{x}{a} + a - x & \text{if } x > 0, \\ a & \text{if } x = 0 \end{cases}$$

is strictly decreasing in the interval $[0, \min\{a, 1\}]$ (Lemma 27). Thus, $\mathrm{D}_{\mathrm{RE}}\left(\rho^* \,\|\, \widetilde{\rho}\,\right) > \mathrm{D}_{\mathrm{RE}}\left(\rho' \,\|\, \widetilde{\rho}\,\right)$ and $\rho' \in \Theta_{\mathrm{A}}$, contradicting the optimality of the projection ρ^*. $\square$

Now we are ready to prove Theorem 8.

Proof of Theorem 8. We split the proof of Theorem 8 into the two cases given in its statement.

1. We first consider the case $\sum_{e \in \mathcal{E}} \min\{\widetilde{\rho}_e, 1\} \leq r_{\mathrm{A}}$. The vector $\rho^* \in \mathbb{R}^{\mathcal{E}}$ given by $\rho^*_e = \min\{\widetilde{\rho}_e,\, 1\}$ for every $e \in \mathcal{E}$ satisfies $\rho^* \in \Theta_{\mathrm{A}}$, since $\rho^* \in [0,1]^{\mathcal{E}}$ and $\sum_{e \in \mathcal{E}} \rho^*_e = \sum_{e \in \mathcal{E}} \min\{\widetilde{\rho}_e, 1\} \leq r_{\mathrm{A}}$. Then,

$$\begin{aligned}\mathrm{D}_{\mathrm{RE}}\left(\rho \,\|\, \widetilde{\rho}\,\right) &= \sum_{e \in \mathcal{E}} \left(\rho_e \ln \frac{\rho_e}{\widetilde{\rho}_e} + \widetilde{\rho}_e - \rho_e\right) \\ &\geq \sum_{e \in \mathcal{E}} \left(\min\{\widetilde{\rho}_e, 1\} \ln \frac{\min\{\widetilde{\rho}_e, 1\}}{\widetilde{\rho}_e} + \widetilde{\rho}_e - \min\{\widetilde{\rho}_e, 1\}\right) = \mathrm{D}_{\mathrm{RE}}\left(\rho^* \,\|\, \widetilde{\rho}\,\right),\end{aligned}$$

   where the inequality follows by Lemma 27. Thus, ρ^* minimizes $\mathrm{D}_{\mathrm{RE}}\left(\cdot \,\|\, \widetilde{\rho}\,\right)$ over Θ_{A} and therefore it is the projection of $\widetilde{\rho}$.

2. We next consider the case $\sum_{e \in \mathcal{E}} \min\{\widetilde{\rho}_e, 1\} > r_{\mathrm{A}}$. First, we note that in this case it holds that $r_{\mathrm{A}} < m$; otherwise, we would have $m = r_{\mathrm{A}} < \sum_{e \in \mathcal{E}} \min\{\widetilde{\rho}_e, 1\} \leq$

m, which is a contradiction. Moreover, from Lemma 30, the projection ρ^* must satisfy $\sum_{e\in\mathcal{E}} \rho_e^* = r_{\mathrm{A}}$. Thus, ρ^* is the optimal solution of the following optimization problem:

$$\begin{aligned} \min_{\rho\in[0,1]^{\mathcal{E}}} \quad & \sum_{e\in\mathcal{E}} \left(\rho_e \ln \frac{\rho_e}{\widetilde{\rho}_e} + \widetilde{\rho}_e - \rho_e\right) \\ \text{subject to} \quad & \sum_{e\in\mathcal{E}} \rho_e = r_{\mathrm{A}}, \end{aligned} \tag{P}$$

where we let $0 \log 0 \coloneqq 0$ to ensure the objective is continuous. We note that the optimal solution of (P) is unique as the objective is strictly convex. The Lagrangian function of (P) is given by

$$\begin{aligned} \mathcal{L}(\rho,\nu) &= \sum_{e\in\mathcal{E}} \left(\rho_e \ln \frac{\rho_e}{\widetilde{\rho}_e} + \widetilde{\rho}_e - \rho_e\right) + \nu \left(\sum_{e\in\mathcal{E}} \rho_e - r_{\mathrm{A}}\right) \\ &= \sum_{e\in\mathcal{E}} \left(\rho_e \ln \frac{\rho_e}{\widetilde{\rho}_e} + \widetilde{\rho}_e - (1-\nu)\rho_e\right), \qquad \forall\, (\rho,\nu) \in [0,1]^{\mathcal{E}} \times \mathbb{R}. \end{aligned}$$

The Lagrangian dual of (P) is the problem

$$\max_{\nu\in\mathbb{R}} \min_{\rho\in[0,1]^{\mathcal{E}}} \mathcal{L}(\rho,\nu). \tag{D}$$

Problem (P) is convex and satisfies the Slater condition (consider, for example, $\rho \in \mathbb{R}^{\mathcal{E}}$ given by $\rho_e \coloneqq r_{\mathrm{A}}/m$ for every $e \in \mathcal{E}$) . Therefore, strong duality holds, and the optimal values of (P) and (D) are identical. Furthermore, if ν^* is an optimal solution of (D), then ρ^* minimizes $\mathcal{L}(\cdot,\nu^*)$ over $[0,1]^{\mathcal{E}}$.

Given the value of ν^*, we can compute the minimizer of $\mathcal{L}(\cdot,\nu^*)$ over $[0,1]^{\mathcal{E}}$ (i.e., ρ^*) analytically. Indeed, since $\mathcal{L}(\rho,\nu^*)$ is separable over the variables ρ_e, the minimum is obtained by minimizing each real function $f_e(\rho_e) \coloneqq \rho_e \ln(\rho_e/\widetilde{\rho}_e) + \widetilde{\rho}_e - (1-\nu^*)\rho_e$, for every $e \in \mathcal{E}$, over the interval $[0,1]$. This yields $\rho_e^* = \min\{\exp(-\nu^*)\widetilde{\rho}_e, 1\}$ for every $e \in \mathcal{E}$ (Lemma 27).

To determine the value of ν^*, we use primal feasibility. From the constraint $\sum_{e \in \mathcal{E}} \rho_e^* = r_{\mathrm{A}}$, it follows that ν^* must satisfy

$$\sum_{e \in \mathcal{E}} \min\{\exp(-\nu^*)\widetilde{\rho}_e, 1\} = r_{\mathrm{A}}. \tag{C.35}$$

To solve equation (C.35), let us sort the components in $\mathcal{E}$ such that $\widetilde{\rho}_{e_1} \geq \cdots \geq \widetilde{\rho}_{e_m}$. Then, by Lemma 29, the projection ρ^* satisfies $\rho^*_{e_1} \geq \cdots \geq \rho^*_{e_m}$. Let k^* be the largest index in $\{1, \ldots, m\}$ such that $\rho^*_{e_{k^*}} = 1$, if such index exists, and $k^* := 0$ otherwise. We note that $k^* \in \{0, \ldots, r_{\mathrm{A}}\}$, as the constraints $\rho^* \in [0,1]^{\mathcal{E}}$ and $\sum_{e \in \mathcal{E}} \rho_e^* = r_{\mathrm{A}}$ prevent ρ^* from having more than r_{A} entries being equal to 1. Then, equation (C.35) translates into

$$k^* + \exp(-\nu^*) \sum_{j=k^*+1}^{m} \widetilde{\rho}_{e_j} = r_{\mathrm{A}},$$

which yields

$$\exp(-\nu^*) = \frac{r_{\mathrm{A}} - k^*}{\sum_{j=k^*+1}^{m} \widetilde{\rho}_{e_j}}.$$

Thus, we can write

$$\rho_e^* = \begin{cases} 1 & \text{if } e = e_i,\ i = 1, \ldots, k^*, \\ \dfrac{r_{\mathrm{A}} - k^*}{\sum_{j=k^*+1}^{m} \widetilde{\rho}_{e_j}} \widetilde{\rho}_e & \text{if } e = e_i,\ i = k^*+1, \ldots, m. \end{cases} \tag{C.36}$$

Next, it remains to determine the value of k^*. We observe that the following inequalities hold:

$$\frac{r_{\mathrm{A}} - k^*}{\sum_{j=k^*+1}^{m} \widetilde{\rho}_{e_j}} \widetilde{\rho}_{e_{k^*+1}} < 1 \leq \frac{r_{\mathrm{A}} - k^*}{\sum_{j=k^*+1}^{m} \widetilde{\rho}_{e_j}} \widetilde{\rho}_{e_{k^*}}. \tag{C.37}$$

Indeed, the first inequality in (C.37) is equivalent to $\rho^*_{e_{k^*+1}} < 1$, which follows by definition of k^*. On the other hand, since $\rho^*_{e_{k^*}} = \min\{\exp(-\nu^*)\widetilde{\rho}_{e_{k^*}}, 1\}$ and $\rho^*_{e_{k^*}} = 1$, we must have $1 \leq \exp(-\nu^*)\widetilde{\rho}_{e_{k^*}}$, which yields the second inequality in (C.37). Rearranging terms in (C.37), we obtain the following inequalities for r_{A}:

$$k^* + \frac{1}{\widetilde{\rho}_{e_{k^*}}} \sum_{j=k^*+1}^{m} \widetilde{\rho}_{e_j} \leq r_{\mathrm{A}} < k^* + 1 + \frac{1}{\widetilde{\rho}_{e_{k^*+1}}} \sum_{j=k^*+2}^{m} \widetilde{\rho}_{e_j}, \tag{C.38}$$

where we let $\widetilde{\rho}_{e_0} := +\infty$. Let $s : \{0, \ldots, m\} \to \mathbb{R}_{\geq 0}$ be the function defined by $s(k) := k + (1/\widetilde{\rho}_{e_k}) \sum_{j=k+1}^{m} \widetilde{\rho}_{e_j}$ for $k \in \{0, \ldots, m\}$. Therefore, (C.38) is equivalent to the inequalities $s(k^*) \leq r_{\mathrm{A}} < s(k^* + 1)$. Since s is nondecreasing (Lemma 28), we conclude that k^* is the maximum integer in $\{0, \ldots, r_{\mathrm{A}}\}$ whose image under s is at most r_{A}, that is,

$$k^* = \max\left\{k \in \{0, \ldots, r_{\mathrm{A}}\} : \, k + \frac{1}{\widetilde{\rho}_{e_k}} \sum_{j=k+1}^{m} \widetilde{\rho}_{e_j} \leq r_{\mathrm{A}}\right\}.$$

Finally, we note that k^*, as well as the constant $(r_{\mathrm{A}} - k^*)/\sum_{j=1}^{m} \widetilde{\rho}_{e_j}$, do not depend on the tie-breaking rule selected for sorting the entries of $\widetilde{\rho}$. Indeed, if π is any other permutation of $\{1, \ldots, m\}$ satisfying $\widetilde{\rho}_{e_{\pi(1)}} \geq \cdots \geq \widetilde{\rho}_{e_{\pi(m)}}$, then $\widetilde{\rho}_{e_{\pi(j)}} = \widetilde{\rho}_{e_j}$ for every $j \in \{1, \ldots, m\}$.

□

Before proceeding with the proof of Theorem 9, we remark that when $\sum_{e \in \mathcal{E}} \min\{\widetilde{\rho}_e, 1\} \leq r_{\mathrm{A}}$, the fact that Algorithm 5 returns the projection of $\widetilde{\rho}$ follows directly from the first case of Theorem 8. Therefore, we focus our analysis on the case $\sum_{e \in \mathcal{E}} \min\{\widetilde{\rho}_e, 1\} > r_{\mathrm{A}}$.

To account for edge cases, throughout this section we assume that both e_0 and e_{m+1} are two components *not* belonging to $\mathcal{E}$ satisfying $\widetilde{\rho}_{e_0} := +\infty$ and $\widetilde{\rho}_{e_{m+1}} := 0$. We recall that the value of k^*, as well as $\sum_{j=k^*+1}^{m} \widetilde{\rho}_{e_j}$ involved in the closed form of the projection (C.36) do not depend on the tie-breaking rule, as noted in the proof of Theorem 8.

We denote by τ the number of iterations of the while loop of Algorithm 5 (Lines 5–15). We use the superscript (t) to denote the iterates generated by the algorithm, which we define as follows.

At the initialization of Algorithm 5, it sets $\mathcal{F}^{(1)} := \mathcal{E}$, $k^{(1)} := 0$ and $s^{(1)} := 0$. For $t \in \{1, \ldots, \tau\}$, the t-th iteration of the while loop of Algorithm 5 starts with $\mathcal{F}^{(t)} \neq \emptyset$. Then, the algorithm selects $e^{(t)} \in \mathcal{F}^{(t)}$ corresponding to the $\lceil |\mathcal{F}^{(t)}|/2 \rceil$-th largest entry of $(\widetilde{\rho}_e)_{e \in \mathcal{F}^{(t)}}$, breaking ties arbitrarily. Then, the algorithm defines the sets $\mathcal{F}^{(t)}_{\text{high}} := \left\{e \in \mathcal{F}^{(t)} : \widetilde{\rho}_e > \widetilde{\rho}_{e^{(t)}}\right\}$, $\mathcal{F}^{(t)}_{\text{eq}} := \left\{e \in \mathcal{F}^{(t)} : \widetilde{\rho}_e = \widetilde{\rho}_{e^{(t)}}\right\}$ and $\mathcal{F}^{(t)}_{\text{low}} := \left\{e \in \mathcal{F}^{(t)} : \widetilde{\rho}_e < \widetilde{\rho}_{e^{(t)}}\right\}$. Next, it computes the quantity $\gamma^{(t)} := k^{(t)} + |\mathcal{F}^{(t)}_{\text{high}}| + |\mathcal{F}^{(t)}_{\text{eq}}| + (1/\widetilde{\rho}_{e^{(t)}}) \left(\sum_{e \in \mathcal{F}^{(t)}_{\text{low}}} \widetilde{\rho}_e + s^{(t)}\right)$, and generates the next iterates $\mathcal{F}^{(t+1)}$, $k^{(t+1)}$ and $s^{(t+1)}$ by considering two cases:

1. If $\gamma^{(t)} \leq r_{\text{A}}$, then $\mathcal{F}^{(t+1)} := \mathcal{F}^{(t)}_{\text{low}}$, $k^{(t+1)} := k^{(t)} + |\mathcal{F}^{(t)}_{\text{high}}| + |\mathcal{F}^{(t)}_{\text{eq}}|$, and $s^{(t+1)} := s^{(t)}$.

2. If $\gamma^{(t)} > r_{\text{A}}$, then $\mathcal{F}^{(t+1)} := \mathcal{F}^{(t)}_{\text{high}}$, $k^{(t+1)} := k^{(t)}$, and $s^{(t+1)} := s^{(t)} + \widetilde{\rho}_{e^{(t)}} |\mathcal{F}^{(t)}_{\text{eq}}| + \sum_{e \in \mathcal{F}^{(t)}_{\text{low}}} \widetilde{\rho}_e$.

We note that as a consequence of the final iteration, we have $\mathcal{F}^{(\tau+1)} = \emptyset$, and the algorithm returns the vector $\rho^* \in \mathbb{R}^{\mathcal{E}}$ defined as $\rho^*_e := \min\{\mu\widetilde{\rho}_e, 1\}$ for every $e \in \mathcal{E}$, with $\mu := (r_{\text{A}} - k^{(\tau+1)})/s^{(\tau+1)}$.

First, we show that Algorithm 5 runs in linear time.

Lemma 31. *Algorithm 5 runs in time $O(m)$.*

Proof of Lemma 31. Since $e^{(t)}$ is the index of the $\lceil |\mathcal{F}^{(t)}|/2 \rceil$-th largest component of $(\widetilde{\rho}_e)_{e \in \mathcal{F}^{(t)}}$, from the update rule of Algorithm 5, we have $|\mathcal{F}^{(t+1)}| \leq \max\left\{|\mathcal{F}^{(t)}_{\text{low}}|, |\mathcal{F}^{(t)}_{\text{high}}|\right\} \leq |\mathcal{F}^{(t)}|/2$. We also note that $e^{(t)}$ can be computed in time $O(m)$ using the Quickselect algorithm [143].

Furthermore, since the τ-th (i.e., the last) iteration of the while loop generates $\mathcal{F}^{(\tau+1)} = \emptyset$, it follows that the total number of iterations of the while loop is then $\tau \leq \lfloor \log(|\mathcal{F}^{(1)}|) + 1 \rfloor$. By iteratively applying the inequality $|\mathcal{F}^{(t+1)}| \leq |\mathcal{F}^{(t)}|/2$, we obtain

$$|\mathcal{F}^{(t)}| \leq \frac{|\mathcal{F}^{(t-1)}|}{2} \leq \cdots \leq \frac{|\mathcal{F}^{(1)}|}{2^{t-1}}, \quad \forall\, t \in \{2, \ldots, \tau\}.$$

Therefore, the total running time of the while loop in Algorithm 5 is given by

$$\sum_{t=1}^{\tau} O(|\mathcal{F}^{(t)}|) = \sum_{t=1}^{\lfloor \log(|\mathcal{F}^{(1)}|)+1 \rfloor} O\left(\frac{|\mathcal{F}^{(1)}|}{2^{t-1}}\right) = O(|\mathcal{F}^{(1)}|) = O(m).$$

Finally, the initialization and return steps of the algorithm take $O(m)$ as well. Therefore, the overall running time of Algorithm 5 is $O(m)$. $\square$

Next, we show that when $\sum_{e\in\mathcal{E}} \min\{\widetilde{\rho}_e, 1\} > r_{\mathrm{A}}$, Algorithm 5 satisfies the following invariants throughout the execution of the while loop:

I$^{(t)}$. The elements of $\mathcal{F}^{(t)}$ are consecutive, that is, for every $i < j < k \in \{1, \ldots, m\}$, $e_i, e_k \in \mathcal{F}^{(t)}$ implies $e_j \in \mathcal{F}^{(t)}$.

II$^{(t)}$. $k^{(t)} = \ell^{(t)} - 1$, where $\ell^{(t)} := \min\left\{i \in \{1, \ldots, m\} : \, e_i \in \mathcal{F}^{(t)}\right\}$.

III$^{(t)}$. $s^{(t)} = \sum_{j=k^{(t)}+1}^{m} \widetilde{\rho}_{e_j} - \sum_{e\in\mathcal{F}^{(t)}} \widetilde{\rho}_e$.

IV$^{(t)}$. $\mathcal{F}^{(t)}$ contains e_{k^*} or e_{k^*+1}.

The following lemma will serve as a useful tool prior to demonstrating the maintenance of these invariants throughout the algorithm.

Lemma 32. *Assume* I$^{(t)}$, II$^{(t)}$ *and* III$^{(t)}$ *hold, and let $i_t \in \{1, \ldots, m\}$ be such that $e^{(t)} = e_{i_t}$. Then, $\gamma^{(t)} = i_t + (1/\widetilde{\rho}_{e_{i_t}}) \sum_{j=i_t+1}^{m} \widetilde{\rho}_{e_j}$. Furthermore, $\gamma^{(t)} \leq r_{\mathrm{A}}$ if and only if $i_t \in \{1, \ldots, k^*\}$.*

Proof of Lemma 32. Let $r^{(t)}$ be defined as $r^{(t)} := \min\left\{j \in \{1, \ldots, m\} : \, e_j \in \mathcal{F}^{(t)}_{\text{low}}\right\}$ if $\mathcal{F}^{(t)}_{\text{low}} \neq \emptyset$, and $r^{(t)} := m + 1$ if $\mathcal{F}^{(t)}_{\text{low}} = \emptyset$. Since I$^{(t)}$ holds and $\ell^{(t)}$ is the index of the first element of $\mathcal{F}^{(t)}$, we can write $\mathcal{F}^{(t)}_{\text{high}} \cup \mathcal{F}^{(t)}_{\text{eq}} = \{e_{\ell^{(t)}}, e_{\ell^{(t)}+1}, \ldots, e_{r^{(t)}-1}\}$. Moreover, since

$e_{i_t} \in \mathcal{F}_{\text{eq}}^{(t)}$, we have $\{e_{i_t+1}, \ldots, e_{r^{(t)}-1}\} \subseteq \mathcal{F}_{\text{eq}}^{(t)}$. Then, it follows that

$$\begin{aligned}
\gamma^{(t)} &= k^{(t)} + |\mathcal{F}_{\text{high}}^{(t)}| + |\mathcal{F}_{\text{eq}}^{(t)}| + \frac{1}{\widetilde{\rho}_{e^{(t)}}} \left(\sum_{e \in \mathcal{F}_{\text{low}}^{(t)}} \widetilde{\rho}_e + s^{(t)} \right) \\
&\overset{(\text{II}^{(t)},\text{III}^{(t)})}{=} \ell^{(t)} - 1 + |\mathcal{F}_{\text{high}}^{(t)}| + |\mathcal{F}_{\text{eq}}^{(t)}| + \frac{1}{\widetilde{\rho}_{e^{(t)}}} \left(\sum_{e \in \mathcal{F}_{\text{low}}^{(t)}} \widetilde{\rho}_e + \sum_{j=\ell^{(t)}}^{m} \widetilde{\rho}_{e_j} - \sum_{e \in \mathcal{F}^{(t)}} \widetilde{\rho}_e \right) \\
&= \ell^{(t)} - 1 + |\mathcal{F}_{\text{high}}^{(t)}| + |\mathcal{F}_{\text{eq}}^{(t)}| + \frac{1}{\widetilde{\rho}_{e^{(t)}}} \left(\sum_{j=\ell^{(t)}}^{m} \widetilde{\rho}_{e_j} - \sum_{e \in \mathcal{F}_{\text{high}}^{(t)} \cup \mathcal{F}_{\text{eq}}^{(t)}} \widetilde{\rho}_e \right) \\
&= |\{e_1, \ldots, e_{\ell^{(t)}-1}\}| + |\{e_{\ell^{(t)}}, \ldots, e_{r^{(t)}-1}\}| + \frac{1}{\widetilde{\rho}_{e^{(t)}}} \left(\sum_{j=\ell^{(t)}}^{m} \widetilde{\rho}_{e_j} - \sum_{j=\ell^{(t)}}^{r^{(t)}-1} \widetilde{\rho}_{e_j} \right) \\
&= |\{e_1, \ldots, e_{\ell^{(t)}-1}\}| + |\{e_{\ell^{(t)}}, \ldots, e_{i_t}\}| + \frac{1}{\widetilde{\rho}_{e_{i_t}}} \left(\sum_{j=i_t+1}^{r^{(t)}-1} \widetilde{\rho}_{e_j} + \sum_{j=r^{(t)}}^{m} \widetilde{\rho}_{e_j} \right) \\
&= i_t + \frac{1}{\widetilde{\rho}_{e_{i_t}}} \sum_{j=i_t+1}^{m} \widetilde{\rho}_{e_j}.
\end{aligned}$$

Finally, the fact that $\gamma^{(t)} = i_t + (1/\widetilde{\rho}_{e_{i_t}}) \sum_{j=i_t+1}^{m} \widetilde{\rho}_{e_j} \leq r_{\text{A}}$ if and only if $i_t \in \{1, \ldots, k^*\}$ follows directly from the definition of k^* and the fact that the function $s(k) \coloneqq k + (1/\widetilde{\rho}_{e_k}) \sum_{j=k+1}^{m} \widetilde{\rho}_{e_j}$ (with $\widetilde{\rho}_{e_0} \coloneqq +\infty$) is nondecreasing (Lemma 28). $\square$

Lemma 33. *Algorithm 5 satisfies invariants* $\text{I}^{(t)}$*,* $\text{II}^{(t)}$*,* $\text{III}^{(t)}$ *and* $\text{IV}^{(t)}$ *for every* $t \in \{1, \ldots, \tau\}$.

Proof of Lemma 33. We prove the lemma by induction on t. At the initialization of Algorithm 5, we have $\mathcal{F}^{(1)} = \mathcal{E}$, $k^{(1)} = 0$, and $s^{(1)} = 0$. In particular, invariants $\text{I}^{(1)}$, $\text{II}^{(1)}$, and $\text{III}^{(1)}$ hold. Moreover, if $k^* > 0$, then $\mathcal{F}^{(1)}$ contains e_{k^*}, and if $k^* = 0$, $\mathcal{F}^{(1)}$ contains $e_1 = e_{k^*+1}$, so $\text{IV}^{(1)}$ holds as well.

Now, for $t \in \{1, \ldots, \tau - 1\}$, let us suppose invariants $\text{I}^{(t)}$, $\text{II}^{(t)}$, $\text{III}^{(t)}$ and $\text{IV}^{(t)}$ hold. At iteration t of the while loop (Lines 5–15), the algorithm selects a component $e^{(t)} \in \mathcal{F}^{(t)}$, and defines $\mathcal{F}_{\text{high}}^{(t)} = \{e \in \mathcal{F}^{(t)} : \widetilde{\rho}_e > \widetilde{\rho}_{e^{(t)}}\}$, $\mathcal{F}_{\text{low}}^{(t)} = \{e \in \mathcal{F}^{(t)} : \widetilde{\rho}_e < \widetilde{\rho}_{e^{(t)}}\}$ and $\mathcal{F}_{\text{eq}}^{(t)} = \{e \in \mathcal{F}^{(t)} : \widetilde{\rho}_e = \widetilde{\rho}_{e^{(t)}}\}$. Let $i_t \in \{1, \ldots, m\}$ be such that $e^{(t)} = e_{i_t}$. Then, we have two cases:

1. $\gamma^{(t)} \leq r_{\text{A}}$. Then, the algorithm sets $\mathcal{F}^{(t+1)} = \mathcal{F}^{(t)}_{\text{low}}$, $k^{(t+1)} = k^{(t)} + |\mathcal{F}^{(t)}_{\text{high}}| + |\mathcal{F}^{(t)}_{\text{eq}}|$, and $s^{(t+1)} = s^{(t)}$. From I$^{(t)}$, the elements of $\mathcal{F}^{(t)}$ are consecutive, so the elements of $\mathcal{F}^{(t+1)}$ are consecutive as well, hence I$^{(t+1)}$ holds. On the other hand, by definition, $\ell^{(t)}$ is the index of the largest element of $\mathcal{F}^{(t)}$, and $\ell^{(t+1)}$ is the index of the largest element of $\mathcal{F}^{(t+1)} = \mathcal{F}^{(t)}_{\text{low}}$. Since the elements of $\mathcal{F}^{(t)}$ are consecutive, we can write $\mathcal{F}^{(t)}_{\text{high}} \cup \mathcal{F}^{(t)}_{\text{eq}} = \{e_{\ell^{(t)}}, e_{\ell^{(t)}+1}, \ldots, e_{\ell^{(t+1)}-1}\}$, and thus $\ell^{(t+1)} = \ell^{(t)} + |\mathcal{F}^{(t)}_{\text{high}}| + |\mathcal{F}^{(t)}_{\text{eq}}|$. Then, it follows that

$$k^{(t+1)} = k^{(t)} + |\mathcal{F}^{(t)}_{\text{high}}| + |\mathcal{F}^{(t)}_{\text{eq}}| \overset{(\text{II}^{(t)})}{=} \ell^{(t)} - 1 + |\mathcal{F}^{(t)}_{\text{high}}| + |\mathcal{F}^{(t)}_{\text{eq}}| = \ell^{(t+1)} - 1. \quad \text{(C.39)}$$

Therefore, II$^{(t+1)}$ holds. Moreover,

$$\begin{aligned}
s^{(t+1)} = s^{(t)} &\overset{(\text{III}^{(t)})}{=} \sum_{j=k^{(t)}+1}^{m} \widetilde{\rho}_{e_j} - \sum_{e \in \mathcal{F}^{(t)}} \widetilde{\rho}_e \\
&\overset{(\text{II}^{(t)})}{=} \sum_{j=\ell^{(t)}}^{m} \widetilde{\rho}_{e_j} - \sum_{e \in \mathcal{F}^{(t)}} \widetilde{\rho}_e \\
&= \sum_{j=\ell^{(t)}}^{\ell^{(t+1)}-1} \widetilde{\rho}_{e_j} + \sum_{j=\ell^{(t+1)}}^{m} \widetilde{\rho}_{e_j} - \sum_{e \in \mathcal{F}^{(t)}} \widetilde{\rho}_e \\
&= \sum_{e \in \mathcal{F}^{(t)}_{\text{high}} \cup \mathcal{F}^{(t)}_{\text{eq}}} \widetilde{\rho}_e + \sum_{j=\ell^{(t+1)}}^{m} \widetilde{\rho}_{e_j} - \sum_{e \in \mathcal{F}^{(t)}} \widetilde{\rho}_e \\
&= \sum_{j=\ell^{(t+1)}}^{m} \widetilde{\rho}_{e_j} - \sum_{e \in \mathcal{F}^{(t)}_{\text{low}}} \widetilde{\rho}_e \\
&\overset{(\text{C.39})}{=} \sum_{j=k^{(t+1)}+1}^{m} \widetilde{\rho}_{e_j} - \sum_{e \in \mathcal{F}^{(t+1)}} \widetilde{\rho}_e.
\end{aligned}$$

Thus, III$^{(t+1)}$ holds. Next, we show that IV$^{(t+1)}$ holds. From I$^{(t)}$, II$^{(t)}$, III$^{(t)}$ and Lemma 32, we have $\gamma^{(t)} = i_t + (1/\widetilde{\rho}_{e_{i_t}}) \sum_{j=i_t+1}^{m} \widetilde{\rho}_{e_j} \leq r_{\text{A}}$, and therefore, $i_t \in \{1, \ldots, k^*\}$. Moreover, from Lemma 28, it holds that $\widetilde{\rho}_{e_{k^*}} > \widetilde{\rho}_{e_{k^*+1}}$. Hence, we have $\widetilde{\rho}_{e_{i_t}} \geq \widetilde{\rho}_{e_{k^*}} > \widetilde{\rho}_{e_{k^*+1}}$. From IV$^{(t)}$, we have $e_{k^*} \in \mathcal{F}^{(t)}$ or $e_{k^*+1} \in \mathcal{F}^{(t)}$. We then consider the following three cases. First, if $e_{k^*+1} \in \mathcal{F}^{(t)}$, then the inequality $\widetilde{\rho}_{e_{i_t}} > \widetilde{\rho}_{e_{k^*+1}}$ implies $e_{k^*+1} \in \mathcal{F}^{(t)}_{\text{low}} = \mathcal{F}^{(t+1)}$. Second, if $e_{k^*} \in \mathcal{F}^{(t)}$ and $\widetilde{\rho}_{e_{i_t}} > \widetilde{\rho}_{e_{k^*}}$, then $e_{k^*} \in \mathcal{F}^{(t)}_{\text{low}} = \mathcal{F}^{(t+1)}$. Third, if $e_{k^*} \in \mathcal{F}^{(t)}$ and $\widetilde{\rho}_{e_{i_t}} = \widetilde{\rho}_{e_{k^*}}$, we have $e_{k^*} \in \mathcal{F}^{(t)}_{\text{eq}}$. Moreover, $\mathcal{F}^{(t)}_{\text{low}} \neq \emptyset$, as $t < \tau$ and $\mathcal{F}^{(t)}_{\text{low}} = \mathcal{F}^{(t+1)} \neq \emptyset$. From I$^{(t)}$, the elements of $\mathcal{F}^{(t)}$ are consecutive, and $\widetilde{\rho}_{e_{k^*}} > \widetilde{\rho}_{e_{k^*+1}}$ necessarily implies that $e_{k^*+1} \in \mathcal{F}^{(t)}_{\text{low}} = \mathcal{F}^{(t+1)}$. Hence, IV$^{(t+1)}$ holds.

2. $\gamma^{(t)} > r_{\text{A}}$. Then, $\mathcal{F}^{(t+1)} = \mathcal{F}^{(t)}_{\text{high}}$, $k^{(t+1)} = k^{(t)}$, and $s^{(t+1)} = s^{(t)} + \widetilde{\rho}_{e^{(t)}}|\mathcal{F}^{(t)}_{\text{eq}}| + \sum_{e\in\mathcal{F}^{(t)}_{\text{low}}} \widetilde{\rho}_e$. From I$^{(t)}$, the elements of $\mathcal{F}^{(t)}$ are consecutive, so the elements of $\mathcal{F}^{(t)}_{\text{high}}$ are consecutive as well, so I$^{(t+1)}$ holds. On the other hand, by definition, $\ell^{(t)}$ is the index of the largest element of $\mathcal{F}^{(t)}$, and $\ell^{(t+1)}$ is the index of the largest element of $\mathcal{F}^{(t+1)} = \mathcal{F}^{(t)}_{\text{high}}$. Therefore, $\ell^{(t+1)} = \ell^{(t)}$. From, II$^{(t)}$, it follows that

$$k^{(t+1)} = k^{(t)} \overset{(\text{II}^{(t)})}{=} \ell^{(t)} - 1 = \ell^{(t+1)} - 1.$$

Therefore, II$^{(t+1)}$ holds. Moreover,

$$\begin{aligned}
s^{(t+1)} &= s^{(t)} + \widetilde{\rho}_{e^{(t)}}|\mathcal{F}^{(t)}_{\text{eq}}| + \sum_{e\in\mathcal{F}^{(t)}_{\text{low}}} \widetilde{\rho}_e \\
&\overset{(\text{III}^{(t)})}{=} \sum_{j=k^{(t)}+1}^{m} \widetilde{\rho}_{e_j} - \sum_{e\in\mathcal{F}^{(t)}} \widetilde{\rho}_e + \widetilde{\rho}_{e^{(t)}}|\mathcal{F}^{(t)}_{\text{eq}}| + \sum_{e\in\mathcal{F}^{(t)}_{\text{low}}} \widetilde{\rho}_e \\
&= \sum_{j=k^{(t)}+1}^{m} \widetilde{\rho}_{e_j} - \sum_{e\in\mathcal{F}^{(t)}} \widetilde{\rho}_e + \sum_{e\in\mathcal{F}^{(t)}_{\text{eq}}} \widetilde{\rho}_e + \sum_{e\in\mathcal{F}^{(t)}_{\text{low}}} \widetilde{\rho}_e \\
&= \sum_{j=k^{(t)}+1}^{m} \widetilde{\rho}_{e_j} - \sum_{e\in\mathcal{F}^{(t)}_{\text{high}}} \widetilde{\rho}_e \\
&= \sum_{j=k^{(t+1)}+1}^{m} \widetilde{\rho}_{e_j} - \sum_{e\in\mathcal{F}^{(t+1)}} \widetilde{\rho}_e.
\end{aligned}$$

Thus, III$^{(t)}$ holds. Next, we show that IV$^{(t)}$ holds. From I$^{(t)}$, II$^{(t)}$, III$^{(t)}$ and Lemma 32, we have $\gamma^{(t)} = i_t + (1/\widetilde{\rho}_{e_{i_t}}) \sum_{j=i_t+1}^{m} \widetilde{\rho}_{e_j} > r_{\text{A}}$, and therefore, $i_t \in \{k^*+1, \ldots, m\}$. This implies $\widetilde{\rho}_{e_{k^*}} > \widetilde{\rho}_{e_{k^*+1}} \geq \widetilde{\rho}_{e_{i_t}}$. From IV$^{(t)}$, we have $e_{k^*} \in \mathcal{F}^{(t)}$ or $e_{k^*+1} \in \mathcal{F}^{(t)}$. We then consider the following three cases. First, if $e_{k^*} \in \mathcal{F}^{(t)}$, then the inequality $\widetilde{\rho}_{e_{k^*}} > \widetilde{\rho}_{e_{i_t}}$ implies $e_{k^*} \in \mathcal{F}^{(t)}_{\text{high}} = \mathcal{F}^{(t+1)}$. Second, if $e_{k^*+1} \in \mathcal{F}^{(t)}$ and $\widetilde{\rho}_{e_{k^*+1}} > \widetilde{\rho}_{e_{i_t}}$, then $e_{k^*+1} \in \mathcal{F}^{(t)}_{\text{high}} = \mathcal{F}^{(t+1)}$. Third, if $e_{k^*+1} \in \mathcal{F}^{(t)}$ and $\widetilde{\rho}_{e_{k^*+1}} = \widetilde{\rho}_{e_{i_t}}$, we have $e_{k^*+1} \in \mathcal{F}^{(t)}_{\text{eq}}$. Moreover, $\mathcal{F}^{(t)}_{\text{high}} \neq \emptyset$, as $t < \tau$ and $\mathcal{F}^{(t)}_{\text{high}} = \mathcal{F}^{(t+1)} \neq \emptyset$. From I$^{(t)}$, the elements of $\mathcal{F}^{(t)}$ are consecutive, and $\widetilde{\rho}_{e_{k^*}} > \widetilde{\rho}_{e_{k^*+1}}$ necessarily implies that $e_{k^*} \in \mathcal{F}^{(t)}_{\text{high}} = \mathcal{F}^{(t+1)}$. Hence, IV$^{(t+1)}$ holds.

$\square$

Finally, we are ready to prove Theorem 9.

Proof of Theorem 9. We note that if $\sum_{e \in \mathcal{E}} \min\{\widetilde{\rho}_e, 1\} \leq r_{\text{A}}$, the projection of $\widetilde{\rho}$ returned by Algorithm 5 directly follows from the first case of Theorem 8. Therefore, for the rest of the proof, we assume that $\sum_{e \in \mathcal{E}} \min\{\widetilde{\rho}_e, 1\} > r_{\text{A}}$.

Let us consider the τ-th (i.e., the last) iteration of the while loop (Lines 5–15). From Lemma 33, Algorithm 5 satisfies invariants I$^{(\tau)}$, II$^{(\tau)}$, III$^{(\tau)}$, and IV$^{(\tau)}$. At iteration τ, the algorithm selects $e^{(\tau)} \in \mathcal{F}^{(\tau)}$, and defines $\mathcal{F}^{(\tau)}_{\text{high}} = \{e \in \mathcal{F}^{(\tau)} : \widetilde{\rho}_e > \widetilde{\rho}_{e^{(\tau)}}\}$, $\mathcal{F}^{(\tau)}_{\text{low}} = \{e \in \mathcal{F}^{(\tau)} : \widetilde{\rho}_e < \widetilde{\rho}_{e^{(\tau)}}\}$ and $\mathcal{F}^{(\tau)}_{\text{eq}} = \{e \in \mathcal{F}^{(\tau)} : \widetilde{\rho}_e = \widetilde{\rho}_{e^{(\tau)}}\}$. Let $i_\tau \in \{1, \ldots, m\}$ be such that $e^{(\tau)} = e_{i_\tau}$. Then, we have the following two cases:

1. $\gamma^{(\tau)} \leq r_{\text{A}}$. Then, the algorithm sets $\mathcal{F}^{(\tau+1)} = \mathcal{F}^{(\tau)}_{\text{low}}$, $k^{(\tau+1)} = k^{(\tau)} + |\mathcal{F}^{(\tau)}_{\text{high}}| + |\mathcal{F}^{(\tau)}_{\text{eq}}|$, and $s^{(\tau+1)} = s^{(\tau)}$. Moreover, from Lemma 32, $\gamma^{(\tau)} = i_\tau + (1/\widetilde{\rho}_{e_{i_\tau}}) \sum_{j=i_\tau+1}^{m} \widetilde{\rho}_{e_j} \leq r_{\text{A}}$, so $i_\tau \in \{1, \ldots, k^*\}$, which implies $\widetilde{\rho}_{i_\tau} \geq \widetilde{\rho}_{e_{k^*}}$. We also recall that the definition of k^* implies $\widetilde{\rho}_{e_{k^*}} > \widetilde{\rho}_{e_{k^*+1}}$ (Lemma 28). On the other hand, since the τ-th iteration is the last one of the while loop, we must have $\mathcal{F}^{(\tau)}_{\text{low}} = \mathcal{F}^{(\tau+1)} = \emptyset$.

   We next argue that $e_{k^*} \in \mathcal{F}^{(\tau)}$, but $e_{k^*+1} \notin \mathcal{F}^{(\tau)}$. From IV$^{(\tau)}$, we know that $\mathcal{F}^{(\tau)}$ contains e_{k^*} or e_{k^*+1}. Suppose $e_{k^*+1} \in \mathcal{F}^{(\tau)}$. Then, $\mathcal{F}^{(\tau)}_{\text{low}} = \emptyset$ implies $\widetilde{\rho}_{i_\tau} \leq \widetilde{\rho}_{e_{k^*+1}} <$

$\widetilde{\rho}_{e_{k^*}}$, which is a contradiction. Thus, $e_{k^*} \in \mathcal{F}^{(\tau)}$. Additionally, $\mathcal{F}^{(\tau)}_{\text{low}} = \emptyset$ implies $\widetilde{\rho}_{i_\tau} \leq \widetilde{\rho}_{e_{k^*}}$, so $\widetilde{\rho}_{i_\tau} = \widetilde{\rho}_{e_{k^*}}$. In particular, $e_{k^*} \in \mathcal{F}^{(\tau)}_{\text{eq}}$. Moreover, due to the facts that $\widetilde{\rho}_{e_{k^*}} > \widetilde{\rho}_{e_{k^*+1}}$, the elements of $\mathcal{F}^{(\tau)}$ are consecutive by $\mathrm{I}^{(\tau)}$, and $\mathcal{F}^{(\tau)}_{\text{low}} = \emptyset$, it follows that e_{k^*} is the smallest element of $\mathcal{F}^{(\tau)}$. Therefore, we can write $\mathcal{F}^{(\tau)} = \{e_{\ell^{(\tau)}}, e_{\ell^{(\tau)}+1}, \ldots, e_{k^*}\}$ and then $\ell^{(\tau)} = k^* - |\mathcal{F}^{(\tau)}| + 1$. Then,

$$k^{(\tau+1)} = k^{(\tau)} + |\mathcal{F}^{(\tau)}_{\text{high}}| + |\mathcal{F}^{(\tau)}_{\text{eq}}| = \ell^{(\tau)} - 1 + |\mathcal{F}^{(\tau)}| = k^*.$$

In addition,

$$\begin{aligned} s^{(\tau+1)} = s^{(\tau)} &\overset{(\mathrm{III}^{(\tau)})}{=} \sum_{j=k^{(\tau)}+1}^{m} \widetilde{\rho}_{e_j} - \sum_{e \in \mathcal{F}^{(\tau)}} \widetilde{\rho}_e \\ &\overset{(\mathrm{II}^{(\tau)})}{=} \sum_{j=\ell^{(\tau)}}^{m} \widetilde{\rho}_{e_j} - \sum_{e \in \mathcal{F}^{(\tau)}} \widetilde{\rho}_e \\ &= \sum_{j=\ell^{(\tau)}}^{k^*} \widetilde{\rho}_{e_j} + \sum_{j=k^*+1}^{m} \widetilde{\rho}_{e_j} - \sum_{e \in \mathcal{F}^{(\tau)}} \widetilde{\rho}_e \\ &= \sum_{j=k^*+1}^{m} \widetilde{\rho}_{e_j}. \end{aligned}$$

2. $\gamma^{(\tau)} > r_{\text{A}}$. Then, the algorithm sets $\mathcal{F}^{(\tau+1)} = \mathcal{F}^{(\tau)}_{\text{high}}$, $k^{(\tau+1)} = k^{(\tau)}$, and $s^{(\tau+1)} = s^{(\tau)} + \widetilde{\rho}_{e^{(\tau)}} |\mathcal{F}^{(\tau)}_{\text{eq}}| + \sum_{e \in \mathcal{F}^{(\tau)}_{\text{low}}} \widetilde{\rho}_e$. Moreover, from Lemma 32, it follows that $\gamma^{(\tau)} = i_\tau + (1/\widetilde{\rho}_{e_{i_\tau}}) \sum_{j=i_\tau+1}^{m} \widetilde{\rho}_{e_j} > r_{\text{A}}$, so $i_\tau \in \{k^*+1, \ldots, m\}$, which implies $\widetilde{\rho}_{e_{k^*+1}} \geq \widetilde{\rho}_{i_\tau}$. On the other hand, since the τ-th iteration is the last one of the while loop, we must have $\mathcal{F}^{(\tau)}_{\text{high}} = \mathcal{F}^{(\tau+1)} = \emptyset$.

   We next argue that $e_{k^*+1} \in \mathcal{F}^{(\tau)}$, but $e_{k^*} \notin \mathcal{F}^{(\tau)}$. From $\mathrm{IV}^{(\tau)}$, we know that $\mathcal{F}^{(\tau)}$ contains e_{k^*} or e_{k^*+1}. Suppose $e_{k^*} \in \mathcal{F}^{(\tau)}$. Then, $\mathcal{F}^{(\tau)}_{\text{high}} = \emptyset$ implies $\widetilde{\rho}_{i_\tau} \geq \widetilde{\rho}_{e_{k^*}} > \widetilde{\rho}_{e_{k^*+1}}$, which is a contradiction. Thus, $e_{k^*+1} \in \mathcal{F}^{(\tau)}$. Additionally, $\mathcal{F}^{(\tau)}_{\text{high}} = \emptyset$ implies $\widetilde{\rho}_{i_\tau} \leq \widetilde{\rho}_{e_{k^*+1}}$, so $\widetilde{\rho}_{i_\tau} = \widetilde{\rho}_{e_{k^*+1}}$. In particular, $e_{k^*+1} \in \mathcal{F}^{(\tau)}_{\text{eq}}$. Moreover, due to the facts that $\widetilde{\rho}_{e_{k^*}} > \widetilde{\rho}_{e_{k^*+1}}$, the elements of $\mathcal{F}^{(\tau)}$ are consecutive by $\mathrm{I}^{(\tau)}$, and $\mathcal{F}^{(\tau)}_{\text{high}} = \emptyset$, it follows that e_{k^*+1} is the largest element of $\mathcal{F}^{(\tau)}$, so $\ell^{(\tau)} = k^* + 1$. Then,

$$k^{(\tau+1)} = k^{(\tau)} = \ell^{(\tau)} - 1 = k^*. \tag{C.40}$$

In addition,

$$\begin{aligned}
s^{(\tau+1)} &= s^{(\tau)} + \widetilde{\rho}_{e^{(\tau)}} |\mathcal{F}_{\text{eq}}^{(\tau)}| + \sum_{e \in \mathcal{F}_{\text{low}}^{(\tau)}} \widetilde{\rho}_e \\
&\overset{(\text{III}^{(\tau)})}{=} \sum_{j=k^{(\tau)}+1}^{m} \widetilde{\rho}_{e_j} - \sum_{e \in \mathcal{F}^{(\tau)}} \widetilde{\rho}_e + \sum_{e \in \mathcal{F}_{\text{eq}}^{(\tau)}} \widetilde{\rho}_e + \sum_{e \in \mathcal{F}_{\text{low}}^{(\tau)}} \widetilde{\rho}_e \\
&= \sum_{j=k^{(\tau)}+1}^{m} \widetilde{\rho}_{e_j} \\
&\overset{(\text{C.40})}{=} \sum_{j=k^*+1}^{m} \widetilde{\rho}_{e_j}.
\end{aligned}$$

Therefore, after the the execution of the while loop, we have $k^{(\tau+1)} = k^*$ and $s^{(\tau+1)} = \sum_{j=k^*+1}^{m} \widetilde{\rho}_{e_j}$. This ensures that $\mu = (r_{\text{A}} - k^*)/\sum_{j=k^*+1}^{m} \widetilde{\rho}_{e_j}$, and by the second case of Theorem 8, the vector returned by Algorithm 5 is the projection of $\widetilde{\rho}$.

Finally, the fact that Algorithm 5 runs in time $O(m)$ has been shown in Lemma 31. $\square$

REFERENCES

[1] R. Staring, L. Bisschop, R. Roks, E. Brein, and H. van de Bunt, "Drug crime and the port of Rotterdam: About the phenomenon and its approach," in *Organized Crime in the 21st Century: Motivations, Opportunities, and Constraints* (Springer, Cham, Switzerland), H. Nelen and D. Siegel, Eds., 2023, pp. 43–61.

[2] L. A. McLay, A. J. Lee, and S. H. Jacobson, "Risk-based policies for airport security checkpoint screening," *Transportation Science*, vol. 44, no. 3, pp. 333–349, 2010.

[3] A. Sinha, T. H. Nguyen, D. Kar, M. Brown, M. Tambe, and A. X. Jiang, "From physical security to cybersecurity," *Journal of Cybersecurity*, vol. 1, no. 1, pp. 19–35, Nov. 2015.

[4] T. R. Wanasinghe, R. G. Gosine, O. De Silva, G. K. I. Mann, L. A. James, and P. Warrian, "Unmanned aerial systems for the oil and gas industry: Overview, applications, and challenges," *IEEE Access*, vol. 8, pp. 166 980–166 997, 2020.

[5] J. Pita *et al.*, "Deployed ARMOR protection: The application of a game theoretic model for security at the Los Angeles International Airport," in *Proceedings of the 7th International Joint Conference on Autonomous Agents and Multiagent Systems: Industrial Track*, ser. AAMAS '08, Estoril, Portugal: International Foundation for Autonomous Agents and Multiagent Systems, 2008, pp. 125–132.

[6] C. Kiekintveld *et al.*, "Computing optimal randomized resource allocations for massive security games," in *Security and Game Theory: Algorithms, Deployed Systems, Lessons Learned*. Cambridge University Press, 2011, pp. 156–176.

[7] D. Korzhyk, V. Conitzer, and R. Parr, "Complexity of computing optimal stackelberg strategies in security resource allocation games," in *Proceedings of the Twenty-Fourth AAAI Conference on Artificial Intelligence (AAAI)*, AAAI Press, Palo Alto, CA, 2010, pp. 805–810.

[8] D. Korzhyk, V. Conitzer, and R. Parr, "Security games with multiple attacker resources," in *Proceedings of the Twenty-Second International Joint Conference on Artificial Intelligence (IJCAI)*, T. Walsh, Ed., AAAI Press/IJCAI, Menlo Park, CA, 2011, pp. 273–279.

[9] Q. Zhu and T. Basar, "Game-theoretic methods for robustness, security, and resilience of cyberphysical control systems: Games-in-games principle for optimal cross-layer resilient control systems," *IEEE Control Systems Magazine*, vol. 35, no. 1, pp. 46–65, 2015.

[10] A. Gupta, C. Langbort, and T. Başar, "Dynamic games with asymmetric information and resource constrained players with applications to security of cyberphysical systems," *IEEE Transactions on Control of Network Systems*, vol. 4, no. 1, pp. 71–81, 2016.

[11] A. R. Hota, A. A. Clements, S. Sundaram, and S. Bagchi, "Optimal and game-theoretic deployment of security investments in interdependent assets," in *International Conference on Decision and Game Theory for Security*, Springer, 2016, pp. 101–113.

[12] O. Lerma, V. Kreinovich, and C. Kiekintveld, "Linear-time resource allocation in security games with identical fully protective resources," in *Proceedings of the 6th AAAI Conference on Applied Adversarial Reasoning and Risk Modeling (AAAIWS)*, AAAI Press, 2011, pp. 42–48.

[13] D. Bertsimas, E. Nasrabadi, and J. B. Orlin, "On the power of randomization in network interdiction," *Operations Research Letters*, vol. 44, no. 1, pp. 114–120, 2016.

[14] F. Miao, Q. Zhu, M. Pajic, and G. J. Pappas, "A hybrid stochastic game for secure control of cyber-physical systems," *Automatica*, vol. 93, pp. 55–63, 2018.

[15] M. Dziubiński and J. Roy, "Hide and seek game with multiple resources," in *Proceedings of the 11th International Symposium on Algorithmic Game Theory (SAGT)*, X. Deng, Ed., Springer, Cham, Switzerland, 2018, pp. 82–86.

[16] M. Dahan, L. Sela, and S. Amin, "Network inspection for detecting strategic attacks," *Operations Research*, vol. 70, no. 2, pp. 1008–1024, 2022.

[17] J. Milošević, M. Dahan, S. Amin, and H. Sandberg, "Strategic monitoring of networked systems with heterogeneous security levels," *IEEE Transactions on Control of Network Systems*, pp. 1–12, 2023.

[18] M. Pourakbar and R. Zuidwijk, "The role of customs in securing containerized global supply chains," *European Journal of Operational Research*, vol. 271, no. 1, pp. 331–340, 2018.

[19] C. Bruwer, "Smuggling and trafficking of illicit goods by sea," in *Global Challenges in Maritime Security: An Introduction*, L. Otto, Ed. Springer, Cham, 2020, pp. 49–73.

[20] M. Gallien and F. Weigand, "Channeling contraband: How states shape international smuggling routes," *Security Studies*, vol. 30, no. 1, pp. 79–106, 2021.

[21] D. Kim and Y. Tajima, “Smuggling and border enforcement,” *International Organization*, vol. 76, no. 4, pp. 830–867, 2022.

[22] E. Pinedo, I. Landauro, and D. Syamnath, *Spanish police make record cocaine bust in Ecuadorean banana shipment*, Reuters, Ed., https://www.reuters.com/world/spanish-police-make-record-cocaine-bust-ecuadorean-banana-shipment-2023-08-25/, [Online; accessed 24-July-2024], 2023.

[23] K. Jones, *Hide and seek: How drug traffickers get creative at sea*, InSight Crime, Ed., https://insightcrime.org/news/hide-and-seek-drug-traffickers-creative-at-sea/, [Online; accessed 24-July-2024], 2021.

[24] A. Chalfin and J. McCrary, “Criminal deterrence: A review of the literature,” *Journal of Economic Literature*, vol. 55, no. 1, pp. 5–48, Mar. 2017.

[25] D. Moses, *To catch a thief: An exclusive look inside the NYPD’s pickpocket task force*, amNewYork Metro, Ed., https://www.amny.com/news/nypd-pickpocket-task-force-exclusive/, [Online; accessed 24-July-2024], 2023.

[26] R. M. Lee, M. J. Assante, and T. Conway, “Analysis of the cyber attack on the Ukrainian power grid,” Electricity Information Sharing and Analysis Center & SANS Industrial Control Systems, Tech. Rep., 2016.

[27] J. Beerman, D. Berent, Z. Falter, and S. Bhunia, “A review of Colonial Pipeline ransomware attack,” in *2023 IEEE/ACM 23rd International Symposium on Cluster, Cloud and Internet Computing Workshops (CCGridW)*, 2023, pp. 8–15.

[28] B. Jeffries, S. Saravia, C. Carter, and Z. Ankuda, “Cyber risk to mission case study: Oldsmar,” The MITRE Corporation, Tech. Rep., 2022.

[29] M. A. Adegboye, W.-K. Fung, and A. Karnik, “Recent advances in pipeline monitoring and oil leakage detection technologies: Principles and approaches,” *Sensors*, vol. 19, no. 11, 2019.

[30] G. H. Mills, A. E. Jackson, and R. C. Richardson, “Advances in the inspection of unpiggable pipelines,” *Robotics*, vol. 6, no. 4, 2017.

[31] ABB Press Release, *ABB launches world’s fastest, most sensitive drone-based gas leak detection and greenhouse gas measuring system*, https://new.abb.com/news/detail/75959/abb-launches-worlds-fastest-most-sensitive-drone-based-gas-leak-detection-and-greenhouse-gas-measuring-system/, [Online; accessed 24-July-2024], 2021.

[32] S. Gal and J. Casas, “Succession of hide–seek and pursuit–evasion at heterogeneous locations,” *Journal of The Royal Society Interface*, vol. 11, no. 20140062, 94 2014.

[33] P. E. Creasey, “Two hide-search games with rapid strategies for multiple parallel searches,” *Open Computer Science*, vol. 12, no. 1, pp. 171–180, 2022.

[34] J. Clanin and S. Bhattacharya, “Additive security games: Structure and optimization,” *arXiv preprint arXiv:2208.02298*, 2022.

[35] Y. Deutsch, “A polynomial-time method to compute all Nash equilibria solutions of a general two-person inspection game,” *European Journal of Operational Research*, vol. 288, no. 3, pp. 1036–1052, 2021.

[36] J. von Neumann, “A certain zero-sum two-person game equivalent to the optimal assignment problem,” in *Contributions to the Theory of Games (AM-28), Volume II*, H. W. Kuhn and A. W. Tucker, Eds., Princeton University Press, 1953, pp. 5–12.

[37] M. M. Flood, “The hide and seek game of von Neumann,” *Management Science*, vol. 18, no. 5-part-2, pp. 107–109, 1972.

[38] A. Washburn and K. Wood, “Two-person zero-sum games for network interdiction,” *Operations Research*, vol. 43, no. 2, pp. 243–251, 1995.

[39] T. J. Leonard, P. Gallo, and S. Véronneau, “Security challenges in United States sea ports: An overview,” *Journal of Transportation Security*, vol. 8, pp. 41–49, 2015.

[40] R. F. I. Meuter and P. F. Lacherez, “When and why threats go undetected: Impacts of event rate and shift length on threat detection accuracy during airport baggage screening,” *Human Factors*, vol. 58, no. 2, pp. 218–228, 2016.

[41] M. Pirani, J. A. Taylor, and B. Sinopoli, “Strategic sensor placement on graphs,” *Systems & Control Letters*, vol. 148, p. 104 855, 2021.

[42] J. Blocki, N. Christin, A. Datta, A. Procaccia, and A. Sinha, “Audit games with multiple defender resources,” in *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence (AAAI)*, AAAI Press, Palo Alto, CA, 2015.

[43] S. Behnezhad *et al.*, “From battlefields to elections: Winning strategies of Blotto and auditing games,” in *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, SIAM Philadelphia, PA, 2018, pp. 2291–2310.

[44] A. L. Abba, M. Awad, Z. Al-Qudah, and A. H. Jallad, “Security analysis of current voting systems,” in *2017 International Conference on Electrical and Computing Technologies and Applications (ICECTA)*, IEEE, 2017, pp. 1–6.

[45] Y. Freund and R. E. Schapire, “Adaptive game playing using multiplicative weights,” *Games and Economic Behavior*, vol. 29, no. 1, pp. 79–103, 1999.

[46] R. J. Lipton, E. Markakis, and A. Mehta, "Playing large games using simple strategies," in *Proceedings of the 4th ACM Conference on Electronic Commerce (EC)*, ACM, New York, 2003, pp. 36–41.

[47] L. Hellerstein, T. Lidbetter, and D. Pirutinsky, "Solving zero-sum games using best-response oracles with applications to search games," *Operations Research*, vol. 67, no. 3, pp. 731–743, 2019.

[48] D. Koller, N. Megiddo, and B. von Stengel, "Fast algorithms for finding randomized strategies in game trees," in *Proceedings of the Twenty-Sixth Annual ACM Symposium on Theory of Computing (STOC)*, ACM, New York, 1994, pp. 750–759.

[49] N. Immorlica, A. T. Kalai, B. Lucier, A. Moitra, A. Postlewaite, and M. Tennenholtz, "Dueling algorithms," in *Proceedings of the Forty-Third Annual ACM Symposium on Theory of Computing (STOC)*, ACM, New York, 2011, pp. 215–224.

[50] J. Letchford and V. Conitzer, "Solving security games on graphs via marginal probabilities," in *Proceedings of the Twenty-Seventh AAAI Conference on Artificial Intelligence (AAAI)*, AAAI Press, Palo Alto, CA, 2013, pp. 591–597.

[51] H. Chan, A. X. Jiang, K. Leyton-Brown, and R. Mehta, "Multilinear games," in *Proceedings of the 12th International Conference on Web and Internet Economics (WINE)*, Y. Cai and A. Vetta, Eds., Springer, Berlin, Heidelberg, 2016, pp. 44–58.

[52] A. Ahmadinejad, S. Dehghani, M. Hajiaghayi, B. Lucier, H. Mahini, and S. Seddighin, "From duels to battlefields: Computing equilibria of Blotto and other games," *Mathematics of Operations Research*, vol. 44, no. 4, pp. 1304–1325, 2019.

[53] V. P. Crawford and N. Iriberri, "Fatal attraction: Salience, naïveté, and sophistication in experimental "hide-and-seek" games," *American Economic Review*, vol. 97, no. 5, pp. 1731–1750, 2007.

[54] A. R. Karlin and Y. Peres, *Game Theory, Alive*. American Mathematical Society, Providence, RI, 2017.

[55] S. Alpern and S. Gal, *The Theory of Search Games and Rendezvous*. Springer New York, NY, 2003.

[56] S. Alpern, R. Fokkink, L. Gąsieniec, R. Lindelauf, and V. S. Subrahmanian, *Search Theory: A Game Theoretic Perspective*. Springer New York, NY, 2013.

[57] R. Hohzaki, "Search games: Literature and survey," *Journal of the Operations Research Society of Japan*, vol. 59, no. 1, pp. 1–34, 2016.

[58] J. Bram, “A 2-player n-region search game,” Center for Naval Analyses, Operations Evaluation Group, Alexandria, VA, Tech. Rep., 1963.

[59] M. Sakaguchi, “Two-sided search games,” *Journal of the Operations Research Society of Japan*, vol. 16, no. 4, pp. 207–225, 1973.

[60] D. M. Roberts and J. C. Gittins, “The search for an intelligent evader: Strategies for searcher and evader in the two-region problem,” *Naval Research Logistics Quarterly*, vol. 25, no. 1, pp. 95–106, 1978.

[61] J. Gittins and D. Roberts, “The search for an intelligent evader concealed in one of an arbitrary number of regions,” *Naval Research Logistics Quarterly*, vol. 26, no. 4, pp. 651–666, 1979.

[62] W. H. Ruckle, “A discrete search game,” in *Stochastic Games And Related Topics: In Honor of Professor L. S. Shapley* (Springer, Dordrecht, Netherlands), T. E. S. Raghavan, T. S. Ferguson, T. Parthasarathy, and O. J. Vrieze, Eds., 1991, pp. 29–43.

[63] J. Clarkson, K. Y. Lin, and K. D. Glazebrook, “A classical search game in discrete locations,” *Mathematics of Operations Research*, vol. 48, no. 2, pp. 687–707, 2023.

[64] J. Clarkson and K. Y. Lin, “Computing optimal strategies for a search game in discrete locations,” *Preprint, submitted May 17, https://arxiv.org/abs/2305.10342*, 2023.

[65] T. Bui, T. Lidbetter, and K. Y. Lin, “Optimal pure strategies for a discrete search game,” *Preprint, submitted June 19, https://arxiv.org/abs/2306.10908*, 2023.

[66] E. J. Subelman, “A hide–search game,” *Journal of Applied Probability*, vol. 18, no. 3, pp. 628–640, 1981.

[67] K. Y. Lin and D. I. Singham, “Finding a hider by an unknown deadline,” *Operations Research Letters*, vol. 44, no. 1, pp. 25–32, 2016.

[68] R. Hohzaki, K. Iida, and T. Komiya, “Discrete search allocation game with energy constraints,” *Journal of the Operations Research Society of Japan*, vol. 45, no. 1, pp. 93–108, 2002.

[69] S. Gal, S. Alpern, and J. Casas, “Prey should hide more randomly when a predator attacks more persistently,” *Journal of The Royal Society Interface*, vol. 12, no. 20150861, 113 2015.

[70] R. Hohzaki, "A search game with incomplete information on detective capability of searcher," *Contributions to Game Theory and Management*, vol. 10, pp. 129–142, 2017.

[71] S. Alpern, S. Gal, V. Lee, and J. Casas, "A stochastic game model of searching predators and hiding prey," *Journal of The Royal Society Interface*, vol. 16, no. 20190087, 153 2019.

[72] T. Nakai, "Search models with continuous effort under various criteria," *Journal of the Operations Research Society of Japan*, vol. 31, no. 3, pp. 335–352, 1988.

[73] K. Iida, R. Hohzaki, and K. Sato, "Hide-and-search game with the risk criterion," *Journal of the Operations Research Society of Japan*, vol. 37, no. 4, pp. 287–296, 1994.

[74] K. Kikuta and W. H. Ruckle, "Accumulation games, part 1: Noisy search," *Journal of Optimization Theory and Applications*, vol. 94, pp. 395–408, 1997.

[75] K. Kikuta and W. H. Ruckle, "Continuous accumulation games on discrete locations," *Naval Research Logistics*, vol. 49, no. 1, pp. 60–77, 2002.

[76] S. Alpern, R. Fokkink, J. op den Kelder, and T. Lidbetter, "Disperse or unite? A mathematical model of coordinated attack," in *Decision and Game Theory for Security (GameSec 2010)*, T. Alpcan, L. Buttyán, and J. S. Baras, Eds., Springer, Berlin, Heidelberg, 2010, pp. 220–233.

[77] S. Alpern, R. Fokkink, T. Lidbetter, and N. S. Clayton, "A search game model of the scatter hoarder's problem," *Journal of The Royal Society Interface*, vol. 9, pp. 869–879, 70 2012.

[78] E. Csóka and T. Lidbetter, "The solution to an open problem for a caching game," *Naval Research Logistics*, vol. 63, no. 1, pp. 23–31, 2016.

[79] T. Lidbetter, "Search games with multiple hidden objects," *SIAM Journal on Control and Optimization*, vol. 51, no. 4, pp. 3056–3074, 2013.

[80] R. Fokkink, K. Kikuta, and D. Ramsey, "The search value of a set," *Annals of Operations Research*, vol. 256, no. 1, pp. 63–73, 2017.

[81] R. Fokkink, T. Lidbetter, and L. A. Végh, "On submodular search and machine scheduling," *Mathematics of Operations Research*, vol. 44, no. 4, pp. 1431–1449, 2019.

[82] T. Lidbetter and K. Y. Lin, "Searching for multiple objects in multiple locations," *European Journal of Operational Research*, vol. 278, no. 2, pp. 709–720, 2019.

[83] B. Bahamondes and M. Dahan, “Network inspection from locations with imperfect detection capabilities,” in *2022 American Control Conference (ACC)*, IEEE, 2022, pp. 613–620.

[84] M. Cardei and D.-Z. Du, “Improving wireless sensor network lifetime through power aware organization,” *Wireless Networks*, vol. 11, no. 3, pp. 333–340, 2005.

[85] W.-Q. Wang and H. Shao, “Radar-to-radar interference suppression for distributed radar sensor networks,” *Remote Sensing*, vol. 6, no. 1, pp. 740–755, 2014.

[86] C. Kiekintveld, M. Jain, J. Tsai, J. Pita, F. Ordóñez, and M. Tambe, “Computing optimal randomized resource allocations for massive security games,” in *Proceedings of the 8th International Conference on Autonomous Agents and Multiagent Systems (AAMAS) - Volume 1*, IFAAMAS, Richland, SC, 2009, pp. 689–696.

[87] M. Musegaas, L. Schlicher, and H. Blok, “Stackelberg production-protection games: Defending crop production against intentional attacks,” *European Journal of Operational Research*, vol. 297, no. 1, pp. 102–119, 2022.

[88] T. Jezierski *et al.*, “Efficacy of drug detection by fully-trained police dogs varies by breed, training level, type of drug and search environment,” *Forensic Science International*, vol. 237, pp. 112–118, 2014.

[89] J. von Neumann, “Zur theorie der gesellschaftsspiele,” *Mathematische Annalen*, vol. 100, no. 1, pp. 295–320, 1928.

[90] J. von Neumann and O. Morgenstern, *Theory of Games and Economic Behavior (60th Anniversary Commemorative Edition)* (Princeton Classic Editions). Princeton University Press, Princeton, NJ, 2004.

[91] A. Washburn, *Two-Person Zero-Sum Games*. Springer, New York, 2014.

[92] S. Behnezhad, S. Dehghani, M. Derakhshan, M. Hajiaghayi, and S. Seddighin, “Faster and simpler algorithm for optimal strategies of Blotto game,” in *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence (AAAI)*, AAAI Press, Palo Alto, CA, 2017.

[93] B. Bahamondes and M. Dahan, “Hide-and-seek game with capacitated locations and imperfect detection,” *Decision Analysis*, vol. 21, no. 2, pp. 110–124, 2024.

[94] R. Selten and W. Güth, “Equilibrium point selection in a class of market entry games,” in *Games, Economic Dynamics, and Time Series Analysis*, M. Deistler, E. Fürst, and G. Schwödiauer, Eds., Heidelberg: Physica-Verlag HD, 1982, pp. 101–116.

[95] A. Karlin and Y. Peres, *Game Theory, Alive* (Miscellaneous Book Series). American Mathematical Society, 2017, ISBN: 9781470419820.

[96] A. Laszka and D. Szeszlér, "Hide and seek in digital communication: The steganography game," in *Proceedings of the 9th Hungarian-Japanese Symposium on Discrete Mathematics and Its Applications*, S. Iwata, N. Kamiyama, S. Kijima, H. Ono, and Y. Yamauchi, Eds., June 2-5, 2015, Fukuoka, Japan, 2015, pp. 126–136.

[97] Z. Cao and X. Yang, "The fashion game: Network extension of matching pennies," *Theoretical Computer Science*, vol. 540-541, pp. 169–181, 2014, Combinatorial Optimization: Theory of algorithms and Complexity.

[98] Y. Deutsch, B. Golany, and U. G. Rothblum, "Determining all Nash equilibria in a (bi-linear) inspection game," *European Journal of Operational Research*, vol. 215, no. 2, pp. 422–430, 2011.

[99] Y. Deutsch, B. Golany, N. Goldberg, and U. G. Rothblum, "Inspection games with local and global allocation bounds," *Naval Research Logistics (NRL)*, vol. 60, no. 2, pp. 125–140, 2013.

[100] M. Jain *et al.*, "Software assistants for randomized patrol planning for the lax airport police and the federal air marshal service," *Interfaces*, vol. 40, no. 4, pp. 267–290, 2010.

[101] M. Jain, E. Kardeş, C. Kiekintveld, M. Tambe, and F. Ordóñez, "Security games with arbitrary schedules: A branch-and-price approach," in *Security and Game Theory: Algorithms, Deployed Systems, Lessons Learned*. Cambridge University Press, 2011, pp. 177–190.

[102] A. Sinha, F. Fang, B. An, C. Kiekintveld, and M. Tambe, "Stackelberg security games: Looking beyond a decade of success," in *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence (IJCAI-18)*, Stockholm, Sweden: IJCAI, 2018, pp. 5494–5501.

[103] D. Korzhyk, Z. Yin, C. Kiekintveld, V. Conitzer, and M. Tambe, "Stackelberg vs. Nash in security games: An extended investigation of interchangeability, equivalence, and uniqueness," *Journal of Artificial Intelligence Research*, vol. 41, pp. 297–327, 2011.

[104] H. Emadi and S. Bhattacharya, "On security games with additive utility$**$this work was supported in part by nsf cps grant eccs 1739969.," *IFAC-PapersOnLine*, vol. 52, no. 20, pp. 351–356, 2019, 8th IFAC Workshop on Distributed Estimation and Control in Networked Systems NECSYS 2019.

[105] B. Roberson, “The Colonel Blotto game,” *Economic Theory*, vol. 29, pp. 1–24, 2006.

[106] S. Hart, “Discrete Colonel Blotto and General Lotto games,” *International Journal of Game Theory*, vol. 36, no. 3, pp. 441–460, 2008.

[107] R. Golman and S. E. Page, “General blotto: Games of allocative strategic mismatch,” *Public Choice*, vol. 138, pp. 279–299, 2009.

[108] A. Ahmadinejad, S. Dehghani, M. Hajiaghayi, B. Lucier, H. Mahini, and S. Seddighin, “From duels to battlefields: Computing equilibria of blotto and other games,” *Mathematics of Operations Research*, vol. 44, no. 4, pp. 1304–1325, 2019.

[109] S. Behnezhad, S. Dehghani, M. Derakhshan, M. Hajiaghayi, and S. Seddighin, “Fast and simple solutions of Blotto games,” *Operations Research*, vol. 71, no. 2, pp. 506–516, 2023.

[110] J. F. Nash, “Equilibrium points in n-person games,” *Proceedings of the National Academy of Sciences*, vol. 36, no. 1, pp. 48–49, 1950. eprint: https://www.pnas.org/doi/pdf/10.1073/pnas.36.1.48.

[111] N. Nisan, T. Roughgarden, E. Tardos, and V. V. Vazirani, *Algorithmic Game Theory*. Cambridge University Press, 2007.

[112] C. E. Lemke and J. T. Howson, “Equilibrium points of bimatrix games,” *Journal of the Society for Industrial and Applied Mathematics*, vol. 12, no. 2, pp. 413–423, 1964.

[113] D. Avis, G. D. Rosenberg, R. Savani, and B. von Stengel, “Enumeration of Nash equilibria for two-player games,” *Economic Theory*, vol. 42, pp. 9–37, 2010.

[114] D. Fudenberg and J. Tirole, *Game Theory*. MIT press, 1991.

[115] Y. Mo *et al.*, “Cyber–physical security of a smart grid infrastructure,” *Proceedings of the IEEE*, vol. 100, no. 1, pp. 195–209, 2011.

[116] H. Sandberg, S. Amin, and K. H. Johansson, “Cyberphysical security in networked control systems: An introduction to the issue,” *IEEE Control Systems Magazine*, vol. 35, no. 1, pp. 20–23, 2015.

[117] S. Weerakkody and B. Sinopoli, “Challenges and opportunities: Cyber-physical security in the smart grid,” in *Smart Grid Control*, Springer, 2019, pp. 257–273.

[118] A. Vidács and R. Vida, “Wireless sensor network based technologies for critical infrastructure systems,” in *Intelligent Monitoring, Control, and Security of Critical Infrastructure Systems*. Springer, Berlin, Heidelberg, 2015, pp. 301–316.

[119] V. P. Il’ev, “An approximation guarantee of the greedy descent algorithm for minimizing a supermodular set function,” *Discrete Applied Mathematics*, vol. 114, no. 1, pp. 131–146, 2001.

[120] M. Sviridenko, J. Vondrák, and J. Ward, “Optimal approximation for submodular and supermodular optimization with bounded curvature,” *Mathematics of Operations Research*, vol. 42, no. 4, pp. 1197–1218, 2017.

[121] Q. Guo, B. An, Y. Zick, and C. Miao, “Optimal interdiction of illegal network flow,” in *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence*, ser. IJCAI’16, New York, New York, USA: AAAI Press, 2016, pp. 2507–2513, ISBN: 9781577357704.

[122] A. Krause, A. Roper, and D. Golovin, “Randomized sensing in adversarial environments,” in *Proceedings of the Twenty-Second International Joint Conference on Artificial Intelligence - Volume Volume Three*, ser. IJCAI’11, Barcelona, Catalonia, Spain: AAAI Press, 2011, pp. 2133–2139.

[123] M. Herbster and M. K. Warmuth, “Tracking the best linear predictor,” *Journal of Machine Learning Research*, vol. 1, no. 281-309, pp. 10–1162, 2001.

[124] W. Krichene, S. Krichene, and A. Bayen, “Efficient Bregman projections onto the simplex,” in *2015 54th IEEE Conference on Decision and Control (CDC)*, 2015, pp. 3291–3298.

[125] C. H. Lim and S. J. Wright, “Efficient Bregman projections onto the permutahedron and related polytopes,” in *Proceedings of the 19th International Conference on Artificial Intelligence and Statistics*, A. Gretton and C. C. Robert, Eds., ser. Proceedings of Machine Learning Research, vol. 51, Cadiz, Spain: PMLR, May 2016, pp. 1205–1213.

[126] J. Robinson and M. Herbster, “Improved regret bounds for tracking experts with memory,” in *Advances in Neural Information Processing Systems*, M. Ranzato, A. Beygelzimer, Y. Dauphin, P. Liang, and J. W. Vaughan, Eds., vol. 34, Curran Associates, Inc., 2021, pp. 7625–7636.

[127] D. Suehiro, K. Hatano, S. Kijima, E. Takimoto, and K. Nagano, “Online prediction under submodular constraints,” in *Algorithmic Learning Theory*, N. H. Bshouty, G. Stoltz, N. Vayatis, and T. Zeugmann, Eds., Berlin, Heidelberg: Springer Berlin Heidelberg, 2012, pp. 260–274.

[128] S. Gupta, M. Goemans, and P. Jaillet, "Solving combinatorial games using products, projections and lexicographically optimal bases," *arXiv preprint arXiv:1603.00522*, 2016.

[129] S. Wang and N. Shroff, "Security game with non-additive utilities and multiple attacker resources," *Proc. ACM Meas. Anal. Comput. Syst.*, vol. 1, no. 1, Jun. 2017.

[130] G. L. Nemhauser, L. A. Wolsey, and M. L. Fisher, "An analysis of approximations for maximizing submodular set functions—I," *Mathematical programming*, vol. 14, pp. 265–294, 1978.

[131] R. Chen, B. Lucier, Y. Singer, and V. Syrgkanis, "Robust optimization for non-convex objectives," in *Proceedings of the 31st International Conference on Neural Information Processing Systems*, ser. NIPS'17, Long Beach, California, USA: Curran Associates Inc., 2017, pp. 4708–4717, ISBN: 9781510860964.

[132] Y. Kawase and H. Sumita, "Randomized strategies for robust combinatorial optimization," *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 33, no. 01, pp. 7876–7883, 2019.

[133] B. Wilder, "Equilibrium computation and robust optimization in zero sum games with submodular structure," *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 32, no. 1, Apr. 2018.

[134] A. Adibi, A. Mokhtari, and H. Hassani, "Minimax optimization: The case of convex-submodular," in *Proceedings of The 25th International Conference on Artificial Intelligence and Statistics*, G. Camps-Valls, F. J. R. Ruiz, and I. Valera, Eds., ser. Proceedings of Machine Learning Research, vol. 151, PMLR, Mar. 2022, pp. 3556–3580.

[135] A. Kalai and S. Vempala, "Efficient algorithms for online decision problems," *Journal of Computer and System Sciences*, vol. 71, no. 3, pp. 291–307, 2005, Learning Theory 2003.

[136] M. K. Warmuth and D. Kuzmin, "Randomized online PCA algorithms with regret bounds that are logarithmic in the dimension," *Journal of Machine Learning Research*, vol. 9, no. 75, pp. 2287–2320, 2008.

[137] D. P. Helmbold and M. K. Warmuth, "Learning permutations with exponential weights," *Journal of Machine Learning Research*, vol. 10, no. 58, pp. 1705–1736, 2009.

[138] W. M. Koolen, M. K. Warmuth, J. Kivinen, *et al.*, "Hedging structured concepts.," in *COLT*, Citeseer, 2010, pp. 93–105.

[139] S. Arora, E. Hazan, and S. Kale, “The multiplicative weights update method: A meta-algorithm and applications,” *Theory of Computing*, vol. 8, no. 6, pp. 121–164, 2012.

[140] S. Yasutake, K. Hatano, S. Kijima, E. Takimoto, and M. Takeda, “Online linear optimization over permutations,” in *Algorithms and Computation*, T. Asano, S.-i. Nakano, Y. Okamoto, and O. Watanabe, Eds., Berlin, Heidelberg: Springer Berlin Heidelberg, 2011, pp. 534–543, ISBN: 978-3-642-25591-5.

[141] M. S. Ang, J. Ma, N. Liu, K. Huang, and Y. Wang, “Fast projection onto the capped simplex with applications to sparse regression in bioinformatics,” in *Advances in Neural Information Processing Systems*, M. Ranzato, A. Beygelzimer, Y. Dauphin, P. Liang, and J. W. Vaughan, Eds., vol. 34, Curran Associates, Inc., 2021, pp. 9990–9999.

[142] M. Conforti and G. Cornuéjols, “Submodular set functions, matroids and the greedy algorithm: Tight worst-case bounds and some generalizations of the Rado-Edmonds theorem,” *Discrete Applied Mathematics*, vol. 7, no. 3, pp. 251–274, 1984.

[143] C. A. R. Hoare, “Algorithm 65: Find,” *Commun. ACM*, vol. 4, no. 7, pp. 321–322, Jul. 1961.

[144] A. Pluta *et al.*, “SciGRID_gas - Data model of the European gas transport network,” in *2022 Open Source Modelling and Simulation of Energy Systems (OSMSES)*, 2022, pp. 1–7.

[145] HIFLD, *Homeland Infrastructure Foundation-Level Data, Geospatial Management Office. U.S. Department of Homeland Security*, Last accessed: 2024-01-19, 2024.

[146] Gurobi Optimization, LLC, *Gurobi Optimizer Reference Manual*, 2024.

[147] I. Csiszár, “I-divergence geometry of probability distributions and minimization problems,” *The Annals of Probability*, vol. 3, no. 1, pp. 146–158, 1975.