



MÁSTER UNIVERSITARIO EN INGENIERÍA DE TELECOMUNICACIÓN

TRABAJO FIN DE MASTER

GAME THEORY AND COUNTERTERRORISM: DEFENSE STRATEGIES FOR TRAIN STATIONS

Author: Jules Winstel

Director: Francisco Alberto Campos Fernández

Co-Director: Luis Jesús Fernández Palomino

Madrid

June of 2026

Declaration of originality

I declare under my responsibility that the Project presented with the title **Game Theory and Counterterrorism: Defense Strategies for train stations** at the ICAI School of Engineering of the Comillas Pontifical University in the academic year 2025/2026 is of my authorship and has not been presented previously for other purposes. The Project is not plagiarised from any other, either totally or partially, and the information that has been taken from other documents is duly referenced.

Use of Artificial Intelligence¹

I declare under my responsibility that (indicate the correct option):

☐ I have not used Artificial Intelligence in the preparation of this document.

☒ I have used Artificial Intelligence in the preparation of this document and/or Annex B under the conditions allowed by Comillas Pontifical University, i.e. applying Level 2 of the Perkins et al. (2024) Assessment Scale: "AI can be used for pre-task activities such as brainstorming, description and initial research. This level focuses on the use of AI for planning, synthesising and generating ideas, but assessments should emphasise the ability to develop and refine these ideas independently". Specifically, Artificial Intelligence has been used to:

(indicate here the concrete use that has been made of Artificial Intelligence)

- Look for articles on the Internet to constitute the state of the art
- Summarize the content of some of these articles
- Assist in coding the optimization problem in GAMS
- Check for errors and typo in some texts written

¹ This declaration refers to the use of generative Artificial Intelligence to carry out the Project documents (Annex B and Memory). It does not apply to Projects where, by their nature, artificial intelligence must be used as part of them (application of machine learning techniques, neural networks, data analysis...).

Signature (student): Jules Winstel



Date: 10th of June 2026

Authorisation for Project delivery

Thesis supervisor	Thesis deputy supervisor (if any)
Francisco Alberto Campos Fernández	Luis Jesús Fernández
Signature: Fecha: 2026.06.17 09:31:53 +02'00'	Signature: <i>Luis Fernández</i>
Date: 10th of June 2026	Date: 10th of June 2026



MÁSTER UNIVERSITARIO EN INGENIERÍA DE TELECOMUNICACIÓN

TRABAJO FIN DE MASTER

GAME THEORY AND COUNTERTERRORISM: DEFENSE STRATEGIES FOR TRAIN STATIONS

Author: Jules Winstel

Director: Francisco Alberto Campos Fernández

Co-Director: Luis Jesús Fernández Palomino

Madrid

June of 2026

TEORÍA DE JUEGOS Y CONTRATERRORISMO: ESTRATEGIAS DE DEFENSA PARA ESTACIONES DE TREN

Autor: Winstel, Jules.

Director: Campos Fernández, Francisco Alberto.

Co-Director: Fernández Palomino, Luis Jesús.

Entidad Colaboradora: ICAI – Universidad Pontificia Comillas.

RESUMEN DEL PROYECTO

El presente trabajo aborda el problema de asignar de forma óptima unos recursos de seguridad limitados para proteger una red ferroviaria frente a un adversario estratégico. Motivado por los atentados del 11-M de 2004 en Madrid, se desarrolla un marco basado en la teoría de juegos que determina cómo un defensor debe distribuir un presupuesto restringido entre las estaciones de una red con el fin de minimizar el daño esperado causado por un atacante racional.

La principal aportación consiste en un modelo teórico que integra tres elementos pocas veces combinados en la literatura existente: la dependencia temporal de los flujos de pasajeros, un conjunto de estrategias defensivas discretas de coste y eficacia crecientes, y la topología de la red. La interacción entre defensor y atacante se formula como un problema de optimización binivel, cuya función objetivo de tipo fraccional se resuelve mediante un procedimiento de bisección sobre el cociente de utilidades entre atacante y defensor.

El modelo se aplica a una representación simplificada en diez zonas de la red de Cercanías de Madrid, empleando datos de pasajeros de acceso público, y se complementa con un análisis de sensibilidad sobre la racionalidad del atacante, el presupuesto disponible y la granularidad temporal del despliegue. Los resultados muestran que la protección se concentra en las estaciones de mayor tamaño durante las horas punta, y que, al aumentar el presupuesto, el defensor primero amplía la cobertura y solo después mejora las estaciones más críticas con estrategias más costosas. La racionalidad del atacante resulta determinante: frente a un adversario suficientemente racional, la concentración supera a la dispersión, y la política óptima puede incluso dejar parte del presupuesto sin gastar. Por último, coordinar el presupuesto de manera conjunta entre franjas horarias reduce de forma significativa la utilidad del atacante en el peor caso, lo que demuestra que

cuándo proteger importa tanto como dónde hacerlo. En conjunto, estos resultados respaldan la planificación de la seguridad basada en la priorización y ofrecen una herramienta escalable para apoyar la protección de sistemas de transporte reales.

Keywords: teoría de juegos, contraterrorismo, protección de infraestructuras críticas, seguridad ferroviaria, optimización binivel, asignación de recursos.

GAME THEORY AND COUNTERTERRORISM: DEFENSE STRATEGIES FOR TRAIN STATIONS

Author: Winstel, Jules.

Director: Campos Fernández, Francisco Alberto.

Co-Director: Fernández Palomino, Luis Jesús.

Collaborating Entity: ICAI – Universidad Pontificia Comillas.

ABSTRACT

This work addresses the problem of optimally allocating limited security resources to protect a railway network against a strategic adversary. Motivated by the 11-M Madrid train bombings of 2004, a game-theoretic framework is developed to determine how a defender should distribute a constrained budget across the stations of a network in order to minimize the expected harm caused by a rational attacker.

The main contribution is a theoretical model that integrates three elements rarely combined in the existing literature: the time-dependence of passenger flows, a set of discrete defensive strategies of increasing cost and effectiveness, and the topology of the network. The interaction between defender and attacker is formulated as a bilevel optimization problem, whose fractional objective is solved through a bisection procedure on the attacker-to-defender utility ratio.

The model is applied to a simplified, ten-zone representation of the Madrid Cercanías network, using publicly available passenger data, and is complemented by a sensitivity analysis on the attacker's rationality, the available budget, and the temporal granularity of the deployment. The results show that protection concentrates on the largest stations during peak hours, and that as the budget grows the defender first broadens coverage and only later upgrades the most critical stations to costlier strategies. The attacker's rationality proves decisive: against a sufficiently rational adversary, concentration outperforms dispersion, and the optimal policy may even leave part of the budget unspent. Finally, coordinating the budget jointly across time slots significantly reduces the attacker's worst-case utility, showing that *when* to protect matters as much as *where*. Overall, these results support prioritization-based security planning and provide a scalable tool to inform the protection of real transit systems.

Keywords: game theory, counterterrorism, critical infrastructure protection, railway security, bilevel optimization, resource allocation.

Contents

1. Introduction.....	1
2. State of the art.....	3
2.1 Evolutionary security games	3
2.2 Information completeness	4
2.3 Type of rationality	6
2.4 Time-dependence	6
2.5 Resolution methods	7
3. Description of the problem	10
3.1 Assumptions, strategies and objective functions.....	10
3.2 Sequence of actions	12
4. Mathematical Formalization.....	13
4.1 Parameters, Variables & Functions	13
4.2 The Attacker's objective function (under Bounded Rationality).....	17
4.3 The Defender's optimization problem	18
4.4 The bilevel mathematical programming problem	19
5. Resolution method	21
5.1 Binary Search Algorithm	21
5.2 Finding feasibility to an achievable r : a new MIQCP problem.....	23
5.3 Piecewise Linear Approximation with the Multiple-Choice Model	24
5.4 Constraints	26
5.5 The objective function.....	26
5.6 Complete MIQCP Formulation.....	27
6. Case study.....	28
6.1 Context of the Case Study.....	29
6.2 Estimation of the parameters.....	35
6.3 Pure Strategies.....	53
6.4 Results: Influence of the Attacker's Rationality	59
6.5 Results: Influence of the Defender's Budget	76
6.6 Results: Influence of the Time Period Chosen.....	94
6.7 Computational Analysis	106

7. Contributions, Conclusions, and Future Lines of Research	112
7.1 Contributions.....	112
7.2 Conclusions.....	113
7.3 Future lines of research.....	118
8. Appendix: Social Impact.....	121

List of figures

Figure 1: Piecewise Linear Approximation of $f(h) = he^h$	25
Figure 2: Cercanías Network of Madrid	29
Figure 3: Model of the Simplified Network	32
Figure 4: Symbolic Importance of each Station	36
Figure 5: Performance Score and Defense Cost of the Strategies	57
Figure 6: Ratio of Defense Cost to Performance Score	58
Figure 7: Evolution of strategies with λ for stations 3 and 5	61
Figure 8: Evolution of strategies with λ for stations 4 and 9	61
Figure 9: Evolution of strategies with λ for stations 1 and 2	62
Figure 10: Evolution of strategies with λ for station 6	63
Figure 11: Evolution of strategies with λ for station 7	64
Figure 12: Evolution of strategies with λ for station 8	65
Figure 13: Evolution of strategies with λ for station 10	66
Figure 14: Evolution of the Use of Strategy $d1$ with λ	68
Figure 15: Evolution of the Use of Strategy $d2$ with λ	69
Figure 16: Evolution of the Use of Strategy $d3$ with λ	70
Figure 17: Evolution of the Use of Strategy $d4$ with λ	71
Figure 18: Evolution of the Use of Strategy $d5$ with λ	72
Figure 19: Evolution of the Use of Strategy $d6$ with λ	73
Figure 20: Evolution of the Share of the Budget Spent per Strategy with λ	74
Figure 21: Mixed strategy for $s1$ when $\lambda = 50, 125$ and 200	78
Figure 22: Mixed strategy for $s2$ when $\lambda = 50, 125$ and 200	78
Figure 23: Strategy mix for $s5$ when $\lambda = 50, 125$ and 200	81
Figure 24: Strategy mix for $s3$ when $\lambda = 50, 125$ and 200	81
Figure 25: Strategy mix for $s4, s6, s7, s8, s9$ and $s10$ when $\lambda = 50$	83
Figure 26: Strategy mix for $s4, s6, s7, s8, s9$ and $s10$ when $\lambda = 125$	85
Figure 27: Strategy mix for $s4, s6, s7, s8, s9$ and $s10$ when $\lambda = 200$	86
Figure 28: Use of $d1$ when $\lambda = 50, 125$ and 200	87
Figure 29: Use of $d2$ when $\lambda = 50, 125$ and 200	88

Figure 30: Use of $d3$ when $\lambda = 50, 125$ and 200	88
Figure 31: Use of $d4$ when $\lambda = 50, 125$ and 200	89
Figure 32: Use of $d5$ when $\lambda = 50, 125$ and 200	89
Figure 33: Use of $d6$ when $\lambda = 50, 125$ and 200	90
Figure 34: Evolution of the Share of the Budget Spent per Strategy with K_t^D for $\lambda = 50$	91
Figure 35: Evolution of the Share of the Budget Spent per Strategy with K_t^D for $\lambda = 125$	92
Figure 36: Evolution of the Share of the Budget Spent per Strategy with K_t^D for $\lambda = 200$	93
Figure 38: Maximum Attacker's Utility Comparison: Joint vs. Independent	97
Figure 39: Absolute (€) and Relative (%) Budget Spent per Time Slot	99
Figure 40: Heatmap of the Budget Spent per Station and Time Slot	100
Figure 41: Aggregated Strategy Use Over the Day	101
Figure 42: Heatmap of the Attacker's Utility Before and After Allocation	102
Figure 43: Correlation Matrix of Station Parameters	104
Figure 44: Evolution of the Computational Time (t) with the Number of Time Slots (n)	110
Figure 45: Alignment with SDGs	121

List of tables

Table 1: State of the Art.....	9
Table 2: Zone Description of the Simplified Model.....	33
Table 3: Frequency and Occupancy per Time Slot.....	39
Table 4: βh per Time Slot [%].....	39
Table 5: Normalized Weight Matrix of W_{ij} values [%].....	40
Table 6: Example of $\Delta h \times W_{ij}$ values in the time slot 6:00-7:00 [%]	42
Table 7: Number of Passengers Affected per Station and Time Slot	44
Table 8: Betweenness Centrality per Station and Time Slot	45
Table 9: Benchmark of Investments in Stations	46
Table 10: Strategy Comparison for s_7 and s_8	84
Table 11: Hardware and Operating System	106
Table 12: Modeling Environment.....	107
Table 13: Solver Configuration	107
Table 14: Algorithmic Parameters	108
Table 15: Instance Parameters	108

List of acronyms

Acronym	Meaning
SSG	Stackelberg Security Game
QR	Quantal Response
SUQR	Subjective Utility Quantal Response
MILP	Mixed Integer Linear Program
MC	Multiple Choice Model
MIQCP	Mixed-Integer Quadratically Constrained Program
SDG	Sustainable Development Goal

Nomenclature

Sets

Symbol	Description
$\mathcal{S}$	Set of stations
$\mathcal{D}$	Set of defensive strategies
$\mathcal{T}$	Set of discrete time slots

Indices

Symbol	Description
$s \in \mathcal{S}$	Index for stations
$d \in \mathcal{D}$	Index for defensive strategies
$t \in \mathcal{T}$	Index for time slots
i	Index for linearization iteration
k	Index for segmentation interval

Parameters (fixed data, capital letters)

Symbol	Description	Units
$ \mathcal{S} $	Total number of stations	num
$ \mathcal{D} $	Total number of defensive strategies	num

$ \mathcal{T} $	Total number of time slots	num
PS_d	Performance score of the pure defensive strategy d	dimensionless
C_d	Cost of the pure defensive strategy d	€/time slot
S_s	Symbolic impact of attacking the station s	dimensionless
$N_{s,t}$	Number of people in the station s at time t	Persons
$B_{s,t}$	Betweenness centrality of the station s at time t	dimensionless
$G_{s',s''}$	Number of shortest paths in a graph between stations s' and s''	num
$G_{s',s''}(s)$	Number of shortest paths in a graph between stations s' and s'' that pass through station s	num
A_s	Cost of assets in the station s	€
C^A	Cost of the attack for the attacker	€
K_T^D	Budget of the defender over the time period T	€
HC	Cost of Human Casualty	€/person
ND	Cost of Network Delay	€
Aff	Number of passengers affected	Persons
T_{delay}	Average additional travel time	Hours
VoT	Monetary value of time per passenger	€/hour/pers
λ	Rationality parameter of the attacker	dimensionless

$H_{s,t}^{min}$	Lower bound of the interval segmentation	€
$H_{s,t}^{max}$	Upper bound of the interval segmentation	€
$P_{s,t,k}$	k-th point of the segmented interval	€
$A_{s,t,k}^{(i)}$	Slope in the k-th interval	dimensionless
$B_{s,t,k}^{(i)}$	y-intercept in the k-th interval	€
ε	Convergence threshold of the binary search	dimensionless
U_0	Upper bound of the binary search	dimensionless
L_0	Lower bound of the binary search	dimensionless
N_h	Number of passengers present in the network at hour h	Persons
N_{tot}	Number of passengers present in the network during the day	Persons
F_h	Train frequency at hour h	Trains/hour
Θ_h	Hourly occupancy coefficient	dimensionless
W_{ij}	Weight of trips from zone i to zone j	dimensionless
P_i	Population of zone i	Persons
A_j	Attractiveness of zone j	dimensionless
D_{ij}	Average travel time from zone i to zone j	Minutes
$T_{ij}(h)$	Hourly passenger demand between zones i and j	Persons
Δ	Directional split	dimensionless

Λ_{board}	Rate of passengers starting their trip	dimensionless
Λ_{alight}	Rate of passengers finishing their trip	dimensionless
W_{board}	Average time in a station when starting a trip	Minutes
W_{alight}	Average time in a station when finishing a trip	Minutes
E	Exposure to an explosion in a station	dimensionless
P_d	Probability of dying if exposed to the attack	%
P_i	Probability of getting injured if exposed to the attack	%
C_{death}	Cost of death for the society	€
C_{injury}	Cost of injury for the society	€

Decision Variables (lowercase letters)

Symbol	Description	Unit
$x_{s,d,t}$	Probability of applying strategy d at station s and time t	%
$x_{s,t}^A$	Probability of the attacker attacking station s at time t	%
$ps_{s,t}$	Performance score of the mixed defensive strategy at station s and time t	%
$c_{s,t}$	Cost of the mixed defensive strategy at station s and time t	€/time slot
$h_{s,t}$	Harm done at station s and time t	€

$\mathbf{z}_{s,t,k}$	Binary switch variable	$\{0,1\}$
$\mathbf{h}_{s,t,k}$	Value of the harm in the k-th interval	€
$\mathbf{r}$	Current candidate value of the binary search iteration	dimensionless

Functions (lowercase letters)

Symbol	Description	Unit
$\mathbf{n}(x)$	Numerator of the fractional objective function	€
$\mathbf{d}(x)$	Denominator of the fractional objective function	dimensionless
$\mathbf{u}^A$	Utility of the attacker	€
$\mathbf{u}^D$	Utility of the defender	€
$\mathbf{r}_t$	Worst-case attacker utility for time slot t	€
$\mathbf{r}_{joint}$	Worst-case attacker utility in the joint case	€
$\mathbf{r}_{ind}$	Worst-case attacker utility in the independent case	€
$\mathbf{f}_{s,t}^{(i)}$	Function to linearize	€
$\mathbf{l}_{s,t}^{(i)}$	Function linearized	€

1. INTRODUCTION

On March 11, 2004, a series of coordinated, nearly simultaneous bombings targeted the Cercanías commuter train system in Madrid, Spain. At the time, it was the deadliest jihadist attack in European history, killing 193 people and injuring over 1,800 [1]. These attacks, their impact on the population, and the political turmoil following deeply marked the country, and highlighted the urgent need for a restructuring of the Spanish counterterrorism system. This led to the creation of the National Counter-Terrorism Coordination Centre, and strengthened intelligence sharing with the EU, the U.S., and North African countries. Despite these improvements, illegal migration and radicalization have kept Spain, although safer than in 2004, at threat level 4 (high risk) for terrorism since 2015 .

Train networks, as well as other public transit systems, are particularly vulnerable to terrorist attacks. They represent a target that is, on the one hand, easy to hit because they are often crowded and open, and, on the other hand, symbolic of daily life and the presence of the state through its infrastructure [1], [2]. For these reasons, attacks on surface transportation have not shown any signs of decline since 9/11. Moreover, even when not lethal, terrorist threats, such as hoaxes or bomb scares, can severely disrupt transportation, leading to panic and significant economic losses.

In this context, recent acts of sabotage, including cable thefts that affected more than 10,000 travelers, have highlighted the significant impact of exploiting network vulnerabilities [3]. These safety concerns also have a deterrent effect on tourism, eroding public confidence in the system and harming a major pillar of the economy of the world’s second most-visited country [4].

To address these security issues, different approaches based on game theory have been developed over the past decades. Formalized in the 1920s by von Neumann and extended in the 1950s by Nash, this mathematical framework, which studies strategic interactions among decision-makers, has been increasingly applied to infrastructure and network defense since 9/11 [5]. From resource allocation among potential targets to more complex models of

infrastructure protection, game theory allows for the optimization of defensive responses when the range of possibilities is too vast for traditional human supervision. However, this approach requires certain assumptions, such as the amount of information available to the attacker or their capacity for rational decision-making factors that are difficult to assess when dealing with terrorism.

The main objective of this thesis is to model the Madrid Cercanías train network and to develop a dynamic defensive model that can adapt to temporal events while providing optimal coverage and minimizing overall costs. More specifically, this thesis aims to advance the theoretical understanding of defender-attacker resource allocation in public transportation networks by developing a unified framework that overcomes a key limitation of the existing literature: the inability to jointly account for the time-dependent nature of passenger flows, the diversity of available defensive strategies, and the topology of the network. In addition, it seeks to characterize how the optimal defensive policy responds to the fundamental factors that shape real-world security planning, namely the rationality of the attacker, the resources available to the defender, and the temporal horizon over which these resources must be deployed.

2. STATE OF THE ART

Over the last decade, game theory has become a major framework for modeling the interaction between attackers and defenders in critical infrastructures such as power grids, communication networks, and transportation systems. We will review in this section the modeling approaches developed in the past to analyze this type of infrastructures, as well as the tools used to solve the resulting problem.

A first approach to a similar problem has been depicted in [2], in the context of defending a pipeline against terrorists, where each segment of the pipeline has a likelihood of getting attacked depending on various factors. This article introduces key concepts that are relevant to our case, including environmental and social considerations in defense planning, and the defender's problem of allocating limited resources across multiple vulnerable sites.

2.1 *EVOLUTIONARY SECURITY GAMES*

Game theory provides a robust framework to model strategic interactions between defenders and attackers in security contexts. In particular, *evolutionary security games* [6] study how these strategies evolve and adapt over time, as players learn from experience or react to changing environments. Such models are especially relevant for critical infrastructures, where defenders and attackers interact repeatedly and their behavior cannot be considered perfectly rational or static.

Within this family, one of the most studied formulations is the *Stackelberg Security Game* (SSG), introduced by Tambe in 2008 [7]. In this model, the defender acts as a *leader* who allocates limited resources across multiple targets, while the attacker, acting as a *follower*, observes this allocation and then decides which target to attack. They are particularly useful for modeling strategic resource allocation and generating patrol schedules, like in the case of the Los Angeles metro [5], an urban rail network [8], or to protect wildlife from poachers [9].

This concept can also be applied to virtual networks, where attackers' payoffs are harder to understand and must get estimated [10].

However, SSGs represent only a specific case within the broader class of evolutionary security games, since they assume a sequential leader-follower interaction. In contrast, evolutionary models relax these assumptions: they allow both players (or populations of strategies) to adapt dynamically and often use replicator dynamics or similar mechanisms to describe how strategies evolve over time. In all these formulations, several key assumptions are defined regarding:

- The completeness of the players' information.
- The rationality of the players.
- The time-dependence of the game.

The resolution methods and use cases will broadly depend on these three assumptions.

2.2 INFORMATION COMPLETENESS

Regarding the completeness of the information, it is important to distinguish between information available to the *players* and that available to the *modeler*. The distinction concerns who lacks the information. Players may face incomplete information about each other (for example, the defender does not know the attacker's exact preferences). The modeler, on the other hand, may lack information about the players themselves: when solving the game with an algorithm such as EXP3 [3], the modeler does not assume knowledge of the players' utility functions or strategy distributions in advance (these are learned through repeated interaction). In other words, EXP3 reflects the modeler's uncertainty about which strategies the players will actually choose, regardless of how much the players themselves know. In this section, we will discuss the information available to the players. Later, when discussing resolution methods, we will focus more on the information available to the modeler.

One type of game, similar to the pipeline defense game [2], assumes complete information for both players. This occurs when the defender's information is publicly available [2], or when the attacker is assumed to behave as if he had complete information, for statistical reasons. For instance, in [8], where the attacker is the set of fraudsters that is assumed to behave as if it knew all of the defender's patrols, since metro stations are open environments, and patrols are visible.

However, in most practical settings, one of the players does not have complete information, either on the behavior of its adversary, or on the payoff matrix. This is the case in [3], where both players have to estimate reactance in a power grid to base their decisions, or in [9], where neither the poachers nor the rangers know the animal distribution, and are getting information on the payoff while patrolling.

In sequential games, one player plays before the other. The *follower* (generally the attacker) has therefore more information than *leader* (generally the defender). This is the case, most of the time, in SSGs like [5], [9], [10], where the attacker is expected to have full information of the defender's resource allocation. We can also observe this sequential attribute in [11], where both players adapt their moves with respect to the other's last move, converging iteratively towards an equilibrium, using Replicators Dynamics.

Finally, some games also consider the defender as the one that has complete knowledge, as opposed to the attacker. Cognitive hierarchy, depicted in [4], is a concept modeling the attacker with a probability of belonging to a certain level, each level having access to a different amount of information. In [4], the defender has a grid to protect, and knows perfectly all its characteristics, while the attacker can have a certain probability to know only the most loaded link, or even not know anything and attack randomly.

2.3 *TYPE OF RATIONALITY*

Another key modeling assumption concerns the rationality of the players. While information is about what the player *knows*, rationality is about how the player *reasons*. We speak of *perfect rationality* when a player makes choices that will maximize their payoff with respect to the information they have. Conversely, when a player has a positive probability to make sub-optimal actions, we speak of *bounded rationality*. In the reviewed literature, players are assumed to act rationally most of the time, but we can find mentions of bounded rationality in [12] and [9]. In those articles, the attacker's actions are modeled using QR (Quantal Response) and SUQR (Subjective Utility Quantal Response), two behavioral models that give the attacker a higher probability of attacking a target with a higher payoff, while still allowing for the possibility of suboptimal choices. The bounded rationality assumption is justified in [12] as a psychological model of human decision-making, while, as stated in [9], SUQR “was shown to perform best in human subject experiments” when compared to other models.

Bounded rationality thus provides a more realistic framework for modeling human behavior, but it also increases the analytical and computational complexity of the resulting game models.

2.4 *TIME-DEPENDENCE*

Security is not only a matter of allocating sufficient resources to protect targets, but also of doing it at the right time. The defensive measures required to protect a train station on a regular weekday are not the same as those needed on the first day of a major holiday. Decisions and optimal strategies can even differ depending on the hours of the day. This is illustrated in [8], where patrol schedules in metro stations must adapt to the varying behavior of fare evaders throughout the day, and in [9], where the spatial distribution of animals, and therefore the poaching risk, changes over time.

Adding time-dependency to the model also allows to take uncertainties into consideration, like in [5], where unexpected events can occur, disturbing the patrol's schedule.

However, incorporating temporal dynamics substantially increases the mathematical and computational complexity of the problem, often requiring advanced dynamic or multi-stage optimization methods.

2.5 RESOLUTION METHODS

The primary goal of game-theoretic resolution methods is to identify a *Nash equilibrium*, defined as a pair of strategies (A, D) for the attacker and defender such that neither player has an incentive to unilaterally deviate without reducing their own payoff. Multiple Nash equilibria may exist within a game, and these can be refined into *Lyapunov-stable* or *asymptotically stable* equilibria, more robust versions that account for populations of players evolving over time.

Depending on the modeling assumptions, different resolution methods will be considered. As stated earlier, the more rational the players, and the smaller the strategy space, the simpler the algorithm to reach the exact solution. In models with perfectly rational players and a small, fully known payoff matrix, the Nash equilibrium can be computed directly using the Lemke–Howson algorithm [13]. It is the case, this approach is suitable for scenarios such as defending a limited number of pipeline segments, where the defender has around a dozen strategies and the attacker only a few.

When the strategy space becomes too large, such as in the generation of patrol schedules between metro stations, a column generation algorithm is preferred [8]. Column generation is an iterative optimization method that starts with a small subset of variables and progressively adds new ones (columns) that most improve the objective, instead of enumerating all possibilities from the start. When the subset is still too large, a greedy version of the column generation algorithm can be applied, with satisfying results [8].

Adding uncertainties about the possibility to execute patrol’s schedule normally, as in [5], further increases the number of possible strategies. In this case, the problem can be

reformulated as a compact Markov Decision Process and solved using linear programming techniques to determine the equilibrium.

In a more complex and general way, when the defender only knows that the attacker's response function belongs to a parametric family, and must estimate the parameters, a combination of projected gradient descent and hysteresis switching guarantees a finite-time convergence to near-optimal strategy, while maintaining robustness under model mismatch or observation noise [10].

When the modeler has perfect knowledge of the game, but players don't, Replicators Dynamics can be used. Each player is represented by a population of strategies whose frequencies evolve over time: strategies performing better than the population average increase in proportion, while weaker ones decline. Therefore, the equilibrium can be approximated, as in [14], or analytically derived by finding the zeros of the strategy derivatives, as in [11]. Although Replicator Dynamics offers intuitive visualization of population evolution, its convergence point is not necessarily a Nash equilibrium and must therefore be verified afterward.

Finally, in the cases of bounded rationality, like QR or SUQR, a combination of regret minimization (ARROW algorithm), column generation and cutting-plane approach (BLADE algorithm) lead to robust results with reasonable computational complexity [9], [12].

Below, in Table 1, a summary of the different articles of the State of the Art of Evolutionary Games when applied to critical infrastructures and their relevant characteristics is presented.

Reference	Information Completeness	Rationality Type	Time dependence	Move Sequence	Decision Variable	Resolution Method
[2]	Complete information	Perfect rationality	No	Simultaneous	Which segment to attack, which defense strategy to choose	Lemke-Howson algorithm [13]
[5]	Defender doesn't know the type of the attacker, and the events that can occur	Perfect rationality	Yes	Sequential	Patrol schedule	Markov Decision Process
[8]	Complete information	Perfect rationality	Yes	Simultaneous	Patrol schedule	Column Generation with greedy approach
[9]	Both players don't know exact payoff	Bounded rationality (SUQR)	Yes	Sequential	Patrol schedule	ARROW BLADE
[10]	Defender doesn't know the attacker's preferences	Perfect rationality	No	Sequential	Where to allocate defense	Adaptive learning algorithm (gradient-based estimation + optimization)
[3]	Incomplete information about grid values for attacker	Perfect rationality	No	Simultaneous	Which line to attack, where to inject false information	EXP3 algorithm
[11]	Strategies are iteratively revised and improved over time	Replicator Dynamics	No	Sequential	Which strategy to choose in a predefined list	Exact Calculation of fixed points of Replicator Dynamics
[4]	Attackers have different levels of knowledge	Perfect rationality	No	Simultaneous	Choice of a link to attack/defend	Lemke-Howson algorithm [13]
[12]	Defenders can only know the probabilities of attacking for each site	Bounded rationality (Quantal Response)	No	Simultaneous	Choice of a target to attack/defend	• Non-linear approximated algorithm • Adapted Branch-and-Price • BLADE
[14]	Strategies are iteratively revised and improved over time	Replicator Dynamics	Yes	Simultaneous	Both choose strategies in a predefined list	Regret Minimization algorithm (improved Replicator Dynamics)

Table 1: State of the Art

3. DESCRIPTION OF THE PROBLEM

The game we will consider is a type of Stackelberg Security Game, where the attacker is a terrorist looking to maximize the damages, and the defender is the society, looking to minimize the damages. The defender is at the upper level, and it is the one actually taking action assuming the reaction of the attacker to the defender's strategies. In this SSG, the targets will represent elements of a network and are interconnected.

3.1 ASSUMPTIONS, STRATEGIES AND OBJECTIVE FUNCTIONS

- The game is modeled as a Stackelberg Security Game in which the defender acts as the leader and the attacker as the follower. The defender commits to a defensive allocation before the attacker selects a target, while the attacker observes the resulting protection levels and reacts accordingly.
- The game assumes complete information regarding the consequences of defensive allocations. Both players are assumed to understand how defensive measures affect the expected harm associated with each target.
- The defender's strategies are the allocation of defense's resources (which are defensive strategies) to different stations in the network, at a time t in $\mathcal{T}$. It is important to note that one of the defenders' strategies will actually be "do nothing", with a cost of 0 and a performance score of 0. This allows the defender to keep some stations unprotected. This choice has been made because in some cases, it can be more optimal not to defend a station, as shown in [15].
- When decisions are taken over multiple time slots, the defender commits to the entire allocation plan at the beginning of each time slot and cannot modify it afterwards. In particular, defensive allocations cannot be adapted based on events observed in previous time slots.
- The defender will have several "pure strategies" available (i.e. patrols, cameras...) and will be able to form "mixed strategies" based on them (i.e. use the pure strategy X 40% of the time and the pure strategy Y 60% of the time). This means that at any moment, the station

has 40% chance of being protected with pure strategy X, and 60% chance to be protected with pure strategy Y. The pure strategies X and Y, however, cannot be used in the station at the same time.

- The defender is assumed to know the attacker's behavioral response model. Therefore, defensive decisions are not taken against a fixed attack pattern, but rather against the attacker's anticipated reaction. More specifically, the defender recognizes that any modification of the defensive allocation changes the expected harm associated with each station-time pair, which in turn affects the attack probabilities generated by the attacker's behavioral model. Consequently, the defender internalizes the attacker's future response when allocating defensive resources.
- The attacker chooses a station-time pair to attack throughout $\mathcal{T}$. It is assumed that only one attack occurs during the planning horizon, which is a reasonable approximation for terrorist attacks, where the objective is typically to maximize the impact of a single operation. The attacker evaluates the expected harm associated with each station-time pair after observing the defender's allocation and bases his decision on the resulting attractiveness of the targets rather than on the specific defensive measures deployed.
- We do not assume that the attacker is perfectly rational. Instead, attacker behavior is modeled using a Quantal Response (QR) model. Unlike the classical assumption of a perfectly rational attacker who always selects the target with the highest utility, the QR model assumes a probabilistic response: targets associated with higher expected harm are more likely to be attacked, but suboptimal targets retain a positive probability of being selected. This formulation captures realistic limitations in human decision-making, including cognitive biases, perception errors and imperfect judgement. The choice of the QR model is also supported by empirical evidence showing that it predicts attacker behavior more accurately than perfect-rationality models in security games [16]. Furthermore, QR-based approaches have been successfully deployed in operational systems such as PROTECT and IRIS [16]. Adopting the same formulation therefore ensures consistency with the state of the art while providing a tractable framework for optimization.

3.2 *SEQUENCE OF ACTIONS*

The sequence of actions during the Stackelberg Security Game is as follows:

1. The defender allocates defensive resources across stations and time slots while anticipating how the attacker will react to the resulting protection levels.
2. Once the defensive allocation is fixed, each station-time pair is associated with a resulting expected harm.
3. The attacker observes these expected harm values and evaluates the attractiveness of each potential target.
4. Following the QR model, the attacker selects a station-time pair with a probability that increases with its expected harm.
5. The outcome of the game is represented by the attack probabilities associated with each station-time pair and the corresponding expected harm. If a simulation is performed, an attack can then be sampled according to the QR probabilities and the resulting harm can be evaluated.

4. MATHEMATICAL FORMALIZATION

4.1 PARAMETERS, VARIABLES & FUNCTIONS

The parameters (inputs), variables and functions depend on the stations $s \in \mathcal{S}$, the defensive strategies $d \in \mathcal{D}$, and the time $t \in \mathcal{T}$, and they are detailed in the next subsections.

4.1.1 PARAMETERS

Parameters are denoted in capital letters. Their estimation procedure will be detailed in the next section.

- The **Symbolic Importance** of the station $S_s \in [0,1]$, will represent how symbolic a station is for society. For example, Madrid Puerta de Atocha would be assigned a high symbolic importance value due to its historical significance, national relevance, and strong presence in Spain's collective memory.
- The **Number of People** $N_{s,t}$ in a station s at a time t .
- The **Betweenness Centrality** $B_{s,t} \in [0,1]$ calculates how many of the most travelled/high-capacity routes pass through a specific station. If a station is attacked, it would sever the high-volume routes that most of the population. Therefore, it represents the importance of the station in the whole network, in terms of its “centrality”, which can be computed as follows:

$$B_{s,t} = \sum_{s' \neq s''} \frac{G_{s',s'',t}(s)}{G_{s',s'',t}} \quad (1)$$

with:

- $G_{s',s'',t}$ the number of shortest paths between two stations s' and s'' at time t
- $G_{s',s'',t}(s)$ the number of those paths that go through station s at time t

The shortest paths can be found by applying Dijkstra's algorithm with the distance between two stations defined as the inverse of the passenger flow between those stations.

- The **Value of Assets Destroyed** A_s in the station s in €, will represent all the economic damage that is not related to human loss (building, equipment...)
- The **Cost of the Attack** for the attacker C^A , in €, represents how much the attacker spends to prepare and execute its attack.
- The **Cost of Human Casualty** HC , in €/person, represents the cost associated with the people killed or injured during the attack. It will be estimated using official data describing the amount of money the society is eager to spend to prevent a death or an injury (which may vary between countries). It will also take into account the exposure factor, since an attack may not reach everyone in the station.
- The **Cost of Network Delays** ND , in €, represents how the rest of the network is economically affected. It will be estimated based on the next equation:

$$ND = Aff \cdot T_{delay} \cdot VoT \quad (2)$$

with N_{total} the total number of passengers that use the network (in person), T_{delay} the average additional time caused by the station not being available (in hours), and VoT the value of time per passenger (in €/hour/passenger).

- The **Defender's Budget** K^D , in €. It represents the amount of money available to the defender to implement his defensive strategies.
- The **Performance Score** $PS_d \in [0,1]$ of the pure strategy d represents the strategy's efficiency. A performance score close to 1 will be given to strategies making it extremely difficult to attack a station (because it has a scanner check of the luggage, like in the airport, for example).
- The **Cost** C_d , in €, of the pure strategy d , representing the cost of implementing this defend strategy for one time slot.

4.1.2 VARIABLES

Variable $x = (x_t)_{t \in T}$, being $x_t \in [0,1]^{|S| \times |D|}$, represents the matrix of probabilities $x_{s,d,t} \in [0,1]$ of the defender assigning strategy d to station s at time t (being T the transpose operator):

$$x_t = \begin{pmatrix} x_{s_1,t}^T \\ \vdots \\ x_{s_{|\mathcal{S}|},t}^T \end{pmatrix} = \begin{pmatrix} x_{s_1,d_1,t} & \cdots & x_{s_1,d_{|\mathcal{D}|},t} \\ \vdots & \ddots & \vdots \\ x_{s_{|\mathcal{S}|},d_1,t} & \cdots & x_{s_{|\mathcal{S}|},d_{|\mathcal{D}|},t} \end{pmatrix}$$

Note that for each station and time slot, the sum of the probabilities associated with each defensive strategies is 1 (i.e., the sum of each row in the above matrix must be equal to 1):

$$\sum_{d \in \mathcal{D}} x_{s,d,t} = 1, \quad \forall s, \forall t \quad (3)$$

4.1.3 FUNCTIONS

- The **Performance Score** $ps_{s,t}(x) \in [0,1]$ represents the efficiency of the mixed strategy x at the station s . It is defined as the weighted average of the mixed strategies over the pure defensive strategy d , considering the performance score $PS = (PS_{d_1}, \dots, PS_{d_{|\mathcal{D}|}})$:

$$ps_{s,t}(x) = PS \cdot x_{s,t} = \sum_{d \in \mathcal{D}} PS_d \cdot x_{s,d,t} \quad (4)$$

- The **Cost of Defense** $c_{s,t}(x)$, in €, represents the average cost of the mixed strategies. Considering $C = (C_{d_1}, \dots, C_{d_{|\mathcal{D}|}})$, we have:

$$c_{s,t}(x) = C \cdot x_{s,t} = \sum_{d \in \mathcal{D}} C_d \cdot x_{s,d,t} \quad (5)$$

- The **Harm** $h_{s,t}(x)$, in €, of an attack at a station s and a time t , when the defensive allocation is x , will be derived from the previous functions as:

$$h_{s,t}(x) = \frac{1}{1 + ps_{s,t}(x)} \cdot (1 + S_s) \cdot (N_{s,t} \cdot HC + B_{s,t} \cdot ND + A_s - C^A + c_{s,t}(x)) \quad (6)$$

We can note that there is no dependency in d . Indeed, since the defensive allocations have already been made, the attacker's payoff only depends on the station s , the time t , and the defensive allocation x . By separating the terms and using matrix notation, we get:

$$h_{s,t}(x_{s,t}) = \frac{1}{1 + PS \cdot x_{s,t}} \cdot (1 + S_s) \cdot (N_{s,t} \cdot HC + B_{s,t} \cdot ND + A - C^A + C \cdot x_{s,t}) \quad (7)$$

Note that:

- the term $\frac{1}{1 + PS \cdot x_{s,t}}$ represents the probability of the attacker successfully conducting an attack taking into account the impact of the defender's strategies in reducing the harm of the attack. For a defense performance score of 0 (i.e. no defense, $PS \cdot x_{s,t}=0$), the probability of attack is 100%=1/(1+0). When the defense performance score is 1 (the maximum value it can take, i.e., $PS \cdot x_{s,t}=1$), we make the conservative assumption that the probability of success can, at maximum, be reduced by 50%=1/(1+1). This could be adapted with the knowledge of security experts. Note that 1 is added at the denominator to ensure that the fraction never goes above 1 (since $PS \cdot x_{s,t} \in [0,1]$).
- the term $1 + S_s$ represents the symbolic impact. Since $S_s \in [0,1]$, we add 1 to emphasize the multiplicative effect of attacking a symbolic place. We don't consider any case where the station is not symbolic to the point that its harm can be reduced to 0.
- the term $N_{s,t} \cdot HC + B_{s,t} \cdot ND + A_s - C^A + C \cdot x_{s,t}$ represents the economic impact of the attack to station s at a time t . $N_{s,t} \cdot HC$ quantifies the cost of direct human casualty (people in the station at the time of the attack), $B_{s,t} \cdot ND$ represents the cost associated with the network disruption (how will the network be affected if this station is attacked), A_s quantifies the damage on the infrastructure and assets of the station, and $-C^A + C \cdot x_{s,t}$ the costs of the attack and the cost of the defense. The defensive cost $C \cdot x_{s,t}$ is positive, because it corresponds to money spent by society, so we count it as a loss, or a positive harm. On the other hand, the attacking cost $-C^A$ is spent by the terrorist, and is counted as a negative harm, as the terrorist is losing money.

4.2 THE ATTACKER'S OBJECTIVE FUNCTION (UNDER BOUNDED RATIONALITY)

Under QR, the probability of the attacker choosing to attack station s is calculated with an exponential weighting, originally introduced in [17] and widely adopted in the security games literature to capture the bounded rationality of human adversaries. Formally, under the exponential weighting QR model, the probability that the attacker selects target s at time t takes the logit form:

$$x_{s,t}^A(x) = \frac{e^{\lambda \cdot h_{s,t}(x)}}{\sum_{s',t'} e^{\lambda \cdot h_{s',t'}(x)}}, \lambda \in \mathbb{R}_+^* \quad (8)$$

The only information available to the attacker is the harm $h_{s,t}$ for each pair station-time (s, t) . We can see that the greater the harm $h_{s,t}$, the greater the probability of attacking, with λ representing the level of rationality of the attacker:

- $\lambda = 0$ implies that the attacker attacks randomly.
- $\lambda \rightarrow +\infty$ implies that the attacker always attacks the station with the highest payoff.

Intermediate values of λ allow the model to interpolate smoothly between these two extremes, which makes the formulation particularly well-suited to representing realistic human adversaries. Moreover, by construction, for any value of λ :

$$\sum_{s,t} x_{s,t}^A(x) = 1 \quad (9)$$

Therefore, the expected harm, which serves as the attacker's objective function $u^A(x)$ to maximize, is defined as the sum of the attacker's harm at each station and each time, weighted by the probability that the attacker actually targets that station and time:

$$u^A(x) = \sum_{s,t} x_{s,t}^A(x) \cdot h_{s,t}(x) \quad (10)$$

4.3 THE DEFENDER'S OPTIMIZATION PROBLEM

4.3.1 THE DEFENDER'S OBJECTIVE FUNCTION

The defender's objective $u^D(x)$ to maximize is the opposite of the attacker objective function defined above, i.e:

$$u^D(x) = -u^A(x) \quad (11)$$

4.3.2 THE DEFENDER'S CONSTRAINTS

Budget constraint

The first constraint of the defender is the budget constraint, which imposes that the mixed strategy allocating the defensive strategies shouldn't go above the defender's budget K_T^D over T :

$$\sum_{s,d,t} x_{s,d,t} \cdot C_d \leq K_T^D \quad (12)$$

This is a constraint that is transversal in time: since the summation runs over all time slots $t \in \mathcal{T}$, the budget is not allocated independently at each time slot but shared across the whole period. As a consequence, the deployment decisions in different time slots are coupled: spending more in one time slot reduces the budget available for the others.

Mix strategy constraint

The second constraint is the probability constraint described in equation (3).

4.4 THE BILEVEL MATHEMATICAL PROGRAMMING PROBLEM

Originally, if the attacker might be completely rational, the SSG would be formulated as the following bilevel model:

$$\max_x u_D(x, x^A) \quad (13)$$

s.t. (3), (12) and also to:

$$x^A = \arg \max u_A(y, x) \quad (14)$$

Being y the mixed strategies of the attacker and $u_D(x, y)$ and $u_A(y, x)$ the utility function of the defender and the attacker, respectively. However, it is worth emphasizing that, although this is a Stackelberg problem in which the defender (leader) commits to a strategy before the attacker (follower) responds, it has not resulted in a genuine bilevel optimization problem. Under the assumption that the attacker observes and knows the defender's strategy (an assumption embedded in the QR equation), the follower's behavior can be expressed in closed form as a function of the defender's strategy $y = x^A(x)$. This allows us to substitute the attacker's response directly into the defender's problem $u^D(x) = u_D(x^A(x), x)$, collapsing what would normally be a two-level structure into the following simpler single-level formulation.

$$\max \left\{ u^D(x) = u_D(x^A(x), x) = - \sum_{s,t} x_{s,t}^A(x) \cdot h_{s,t}(x) \right\} \quad (15)$$

s.t. (3), (12)

Which is equivalent to:

$$\min \left\{ u^A(x) = \sum_{s,t} x_{s,t}^A(x) \cdot h_{s,t}(x) \right\} \quad (16)$$

s.t. (3), (12)

By substituting the QR probability $x_{s,t}^A(x)$, we get:

$$\min \sum_{s,t} \frac{e^{\lambda \cdot h_{s,t}(x)}}{\sum_{s',t'} e^{\lambda \cdot h_{s',t'}(x)}} \cdot h_{s,t}(x) \quad (17)$$

s.t. (3), (12)

Notice that the denominator is independent of the outer summation indexes s and t . Therefore, we can factor it out to express the defender's objective function as a single fraction as follows:

$$\min_x \frac{\sum_{s,t} (h_{s,t}(x) \cdot e^{\lambda \cdot h_{s,t}(x)})}{\sum_{s',t'} e^{\lambda \cdot h_{s',t'}(x)}} \quad (18)$$

s.t. (3), (12)

By expanding $h_{s,t}(x)$ into its explicit components (as defined in (6)), the full optimization problem of the defender for a specific time t is formulated as follows:

$$\min_x \frac{\sum_{s,t} \left(\alpha_s \cdot \frac{\kappa_{s,t} + c_{s,t}(x)}{1 + p_{s,t}(x)} \cdot e^{\lambda \cdot \alpha_s \cdot \frac{\kappa_{s,t} + c_{s,t}(x)}{1 + p_{s,t}(x)}} \right)}{\sum_{s',t'} e^{\lambda \cdot \alpha_{s'} \cdot \frac{\kappa_{s',t'} + c_{s',t'}(x)}{1 + p_{s',t'}(x)}}} \quad (19)$$

s.t. (3), (12)

being:

$$\kappa_{s,t} = N_{s,t} \cdot HC + A_s + B_{s,t} \cdot ND - C^A \quad (20)$$

$$\alpha_s = 1 + S_s \quad (21)$$

5. RESOLUTION METHOD

The inclusion of the QR model results in a non-linear and non-convex fractional objective function for the defender. Finding a global optimum for such a problem directly might not be computationally possible, especially as the network size and the number of available defensive strategies increase. To guarantee a global optimum within a reasonable computation time, we transform this problem into a sequence of Mixed-Integer Linear Programs (MILPs). In particular, our approach relies on two main techniques: a Binary Search method to eliminate the fractional nature of the objective, and a Multiple Choice (MC) piecewise linear approximation to handle the non-linear exponential components.

5.1 BINARY SEARCH ALGORITHM

We consider separately the numerator $n(x)$ and the denominator $d(x)$ of the function we want to minimize (i.e. of (19)):

$$n(x) = \sum_{s,t} \alpha_s \cdot \frac{\kappa_{s,t} + c_{s,t}(x)}{1 + p_{s,t}(x)} \cdot e^{\lambda \cdot \alpha_s \frac{\kappa_{s,t} + c_{s,t}(x)}{1 + p_{s,t}(x)}} \quad (22)$$

$$d(x) = \sum_{s',t'} e^{\lambda \cdot \alpha_{s'} \frac{\kappa_{s',t'} + c_{s',t'}(x)}{1 + p_{s',t'}(x)}} \quad (23)$$

We notice here that $\forall x, d(x) > 0$.

Our goal is to find the minimum global value p^* of the ratio $\frac{n(x)}{d(x)}$. Instead of minimizing this complex fraction directly, we treat it as a feasibility problem. We test whether a target minimum value r is achievable or not, i.e., when:

$$\frac{n(x)}{d(x)} \leq r \leftrightarrow n(x) \leq r \cdot d(x) \quad (24)$$

Therefore, mathematically $r \geq p^* = \frac{n(x^*)}{d(x^*)}$ if and only if there exists a feasible defensive allocation x such that:

$$r \cdot d(x) - n(x) \geq 0 \quad (25)$$

With the smallest value for r for any value of x . Therefore, the minimization of the fraction for the Defender's point of view is transformed into a Feasibility Check Optimization (CF-OPT). Then, we apply the Binary Search algorithm, developed in [16], by selecting different values of r . If for a specific value of r , there is a feasible solution x , we reduce the value of r . Otherwise, we increase it. Both movements are controlled by an upper U and lower L values for r until we converge towards the optimal solution, i.e, until both bounds U and L coincide in a bisection algorithm. This can be represented through the following pseudocode:

```

1 Input: select  $\varepsilon$ ,  $U$ ,  $L$ ;
2 while  $U - L \geq \varepsilon$  do
3    $r \leftarrow (U+L)/2$ 
4   CheckFeasibility, ie., if thereis  $x^r$  such that  $r \cdot d(x^r) - n(x^r) \geq 0$ ;
6   if feasible then
7      $U \leftarrow r$ 
8   else
9      $L \leftarrow r$ 
10 return  $U$ ,  $L$ ,  $x^r$ ;

```

The number of iterations of this algorithm is bounded by $\mathcal{O}\left(\log\left(\frac{U_0-L_0}{\varepsilon}\right)\right)$, being U_0 and L_0 the initial values of U and L chosen in the first iteration of the algorithm.

5.2 FINDING FEASIBILITY TO AN ACHIEVABLE R : A NEW MIQCP PROBLEM

To evaluate the feasibility of $r \cdot d(x) - n(x) \geq 0$ in step 4 of the above algorithm, we must process the non-linear exponential function $e^{\lambda \cdot h_{s,t}(x)}$. However, the expected harm $h_{s,t}(x)$ depends on the mixed strategy vector x , which is a multi-dimensional probability distribution (see equation (6)). Attempting to apply a piecewise linear approximation directly over a multi-dimensional vector creates a combinatorial explosion. For example, discretizing a 10-strategy space with 5 points per strategy would require 5^{10} binary variables per station, rendering the MILP completely unsolvable.

But when looking closely, we notice that the non-linear QR formula does not strictly care about which specific strategies the defender chose. It only operates on the final calculated scalar value of the attacker's expected utility at that station. Therefore, we bypass the multi-dimensional space by using the continuous intermediate scalar variable $h_{s,t} \in \mathbb{R}$ to represent the attacker's expected damage at station s and time t :

$$h_{s,t} = \alpha_s \cdot \frac{\kappa_{s,t} + c_{s,t}(x)}{1 + ps_{s,t}(x)} \quad (26)$$

Variable $h_{s,t}$ allows us to reduce a multi-dimensional optimization problem into a simple, one-dimensional curve. Thus, we can define the following non-linear functions, which are part of $d(x)$ and $n(x)$, as functions of the variable $h_{s,t}$:

$$f^{(1)}(h_{s,t}) = e^{\lambda \cdot h_{s,t}} = f_{s,t}^{(1)} \quad (27)$$

$$f^{(2)}(h_{s,t}) = h_{s,t} \cdot e^{\lambda \cdot h_{s,t}} = f_{s,t}^{(2)} \quad (28)$$

5.3 *PIECEWISE LINEAR APPROXIMATION WITH THE MULTIPLE-CHOICE MODEL*

With our functions reduced to 1D scalars, we utilize the MC model, first introduced in [18] and shown as equivalent to other approximation methods in [19]. The idea is to divide the domain of the decision variable into distinct segments and to use binary variables to force the optimization solver to select exactly one active segment. As demonstrated in [20], this explicit geometric formulation is highly effective for modern MILP solvers when the binary variables aren't too numerous, as it is the case for us.

To implement the MC model, we must define strict upper and lower bounds $[H_{s,t}^{min}, H_{s,t}^{max}]$ of $h_{s,t}$. Since harm function $h_{s,t}$ is a linear-fractional function over the probability simplex $\{x_{s,d,t} : d \in D, \sum_{d \in D} x_{s,d,t} = 1\}$ the global minimum and maximum must occur at the pure strategies, as shown in [21]. Therefore, we can pre-calculate the mentioned bounds as:

$$H_{s,t}^{min} = \min_{d \in D} \left(\alpha_s \cdot \frac{\kappa_{s,t} + C_d}{1 + PS_d} \right) \quad (29)$$

$$H_{s,t}^{max} = \max_{d \in D} \left(\alpha_s \cdot \frac{\kappa_{s,t} + C_d}{1 + PS_d} \right) \quad (30)$$

Then, we divide this domain into K segments by defining a set of $K + 1$ grid points: $P_{s,t,0}, P_{s,t,1}, \dots, P_{s,t,K}$ where $P_{s,t,0} = H_{s,t}^{min}$ and $P_{s,t,K} = H_{s,t}^{max}$. The k -th segment operates within the boundaries $[P_{s,t,k-1}, P_{s,t,k}]$.

Then we will introduce two new sets of variables for each segment $k \in \{1, \dots, K\}$:

- The **Binary Switch Variable** ($z_{s,t,k} \in \{0, 1\}$) that dictates whether the k -th segment is active ($z_{s,t,k} = 1$) or inactive ($z_{s,t,k} = 0$)
- The **Segmented Utility Variable** ($h_{s,t,k} \geq 0$) that tracks the exact utility value, but *only* if its corresponding segment is active.

For each segment, we calculate two constants:

- The slope of the segment $A_{s,t,k}^{(i)}$ for the function $f^{(i)}$, as follows:

$$A_{s,t,k}^{(i)} = \frac{f^{(i)}(P_{s,t,k}) - f^{(i)}(P_{s,t,k-1})}{P_{s,t,k} - P_{s,t,k-1}} \quad (31)$$

- The y-intercept $B^{(i)}$ of the function $f_s^{(i)}$:

$$B_{s,t,k}^{(i)} = f^{(i)}(P_{s,t,k}) - P_{s,t,k} \cdot \frac{f^{(i)}(P_{s,t,k}) - f^{(i)}(P_{s,t,k-1})}{P_{s,t,k} - P_{s,t,k-1}} \quad (32)$$

Using these constants, we can replace the non-linear curves by a linear sum of intercepts and slopes:

$$l_{s,t}^{(i)} = \sum_{k=1}^K (B_{s,t,k}^{(i)} \cdot z_{s,t,k} + A_{s,t,k}^{(i)} \cdot h_{s,t,k}) \quad (33)$$

In Figure 1, we can see an example of piecewise linear approximation for $f^{(2)}$.




Figure 1: Piecewise Linear Approximation of $f(h) = he^h$

5.4 CONSTRAINTS

We also impose the following linear constraints:

- The exclusivity constraint: exactly one segment must be active at a given time

$$\sum_{k=1}^K z_{s,t,k} = 1, \quad \forall s, \forall t \quad (34)$$

- The bounding constraint: the value of $h_{s,t,k}$ is constrained by the grid points of its segment.

$$z_{s,t,k} \cdot P_{s,t,k-1} \leq h_{s,t,k} \leq z_{s,t,k} \cdot P_{s,t,k}, \quad \forall s, \forall t, \forall k \quad (35)$$

- The utility aggregation: Linking $h_{s,t,k}$ to h_s :

$$\sum_{k=1}^K h_{s,t,k} = h_{s,t}, \quad \forall s, \forall t \quad (36)$$

- The utility bridging constraint: To link the intermediate variable $h_{s,t}$ back to the original mixed strategy probabilities x , the relationship is defined by:

$$h_{s,t} \cdot (1 + p_{s,t}(x)) = \alpha_s \cdot (\kappa_{s,t} + c_{s,t}(x)), \quad \forall s, \forall t \quad (37)$$

This last constraint introduces a bilinear term $h_{s,t} \cdot p_{s,t}(x)$, which transforms the problem into a Mixed-Integer Quadratically Constrained Program (MIQCP). This kind of problem is, however, solvable by modern optimization solvers.

5.5 THE OBJECTIVE FUNCTION

To verify if a target expected utility r is achievable, the following objective function must be maximized:

$$\max_x \sum_{s,t} (r \cdot l_{s,t}^{(1)} - l_{s,t}^{(2)}) \geq 0 \quad (38)$$

If the optimal value of the objective function is positive, then feasibility of $r \cdot d(x) - n(x) \geq 0$ in step 4 of the algorithm is achieved.

5.6 COMPLETE MIQCP FORMULATION

By expanding $l_{s,t}^{(1)}$ and $l_{s,t}^{(2)}$, the final MIQCP can be rewritten as follows:

$$\max_{x,h,z} \sum_{s,t} \sum_{k=1}^K \left[(r \cdot B_{s,t,k}^{(1)} - B_{s,t,k}^{(2)}) \cdot z_{s,t,k} + (r \cdot A_{s,t,k}^{(1)} - A_{s,t,k}^{(2)}) \cdot h_{s,t,k} \right]$$

s.t. (3), (12), (34), (35), (36), (37) and (38).

6. CASE STUDY

This case study examines how the defender's protection strategy across the network responds to variations in the model's key parameters, based on passenger data from the 1st of January 2023. Section 6.1 introduces the geographical context, justifying this choice, and extracting the network to be analyzed. Section 6.2 then presents the methodology used to estimate each parameter involved in the model. Regarding the strategies, section 6.3 will detail the estimate of the costs and the performance scores of each pure strategy. Finally, the results are presented in the last sections. In particular:

- In sections 6.4 and 6.5, we analyze how the defender allocates resources over a single time slot (from 6:00 to 7:00) with respect to two key parameters of the model: the rationality of the attacker λ and the budget of the defender K_t^D .
- In section 6.6, we extend the analysis to several time slots (six slots, from 6:00 to 0:00), examining the impact of distributing the budget jointly versus separately, and studying the defensive allocations throughout the day.
- Finally, section 6.7, evaluates the computational performance of the model and the resolution method over 25 time slots with parameters generated randomly, which correspond to 4 days, assessing the extent to which it can be scaled up.

6.1 CONTEXT OF THE CASE STUDY

6.1.1 THE CERCANÍAS NETWORK OF MADRID

This case study will be conducted on the Cercanías network of Madrid, the suburban railway system operated by Renfe that connects the city center with its metropolitan area. Below can be found the complete map of this network:



Figure 2: Cercanías Network of Madrid

The choice of this particular network as the application case for our model has been motivated by the following reasons:

Historical relevance and exposure to terrorist threats.

As mentioned in the introduction, the Cercanías network of Madrid was the target of the deadliest terrorist attack ever perpetrated on European soil, the 11-M attacks of March 11, 2004, in which ten coordinated explosions on four commuter trains killed 193 people and

injured more than 2,000. This tragic precedent demonstrates that the network is not a hypothetical target but a system with a documented history of being attacked, which makes it both a meaningful and a sensitive case for the application of a security resource allocation model. Moreover, the attack itself revealed structural vulnerabilities (high passenger density, multiple unmonitored access points, and the difficulty of inspecting luggage in a high-throughput environment) that are precisely the kind of features our model aims to capture through the accessibility and permeability parameters.

Symbolic, economic, and strategic value of the network as a potential target.

Beyond its historical exposure, the Cercanías network of Madrid concentrates a set of characteristics that make it an especially attractive target from an attacker's perspective, and therefore a particularly relevant case for a SSG. The network transports several hundreds of thousands of passengers per day, with strong peaks during morning and evening rush hours, when stations such as Atocha or Chamartín gather tens of thousands of commuters in confined spaces within short time windows. This extreme passenger density translates directly into a high expected number of casualties in case of an attack, which is precisely the quantity that the attacker seeks to maximize in our model. In addition, the network serves the capital of Spain and connects key political, economic, and transport hubs (including ministries, business districts, and the country's main railway and airport interchanges) which gives any attack on the system a strong symbolic and media impact that extends well beyond its immediate human cost. From a game-theoretic standpoint, this combination of high physical damage potential and high symbolic value means that the attacker's utility function is sharply differentiated across targets, which is exactly the kind of setting in which optimizing the defender's allocation of limited security resources yields the largest benefit compared to uniform or intuition-based strategies. In other words, Cercanías is not only a plausible target but also a discriminating one, where the added value of a formal security model is most clearly visible.

Manageable structural complexity for first implementation.

The Cercanías network is composed of around 90 stations distributed over 10 lines, but it can be naturally aggregated into a small number of large operational zones. This structure allows the model to be implemented at a smaller scale in its first version, without facing the combinatorial explosion that would arise from modeling a large metro or high-speed rail network at the station level. Working on an aggregated representation of Cercanías therefore offers a good compromise between realism (the zones correspond to genuine operational and security units) and computational complexity, which is essential when validating the PASAQ-based solution method developed in the previous chapter.

Availability of public data on passenger flows.

A key input of the model is the number of passengers $N_{s,t}$ present at each station s during each time period t . Renfe publishes on its official website [22] aggregated statistics on the daily ridership of the Cercanías network. This makes it possible to calibrate the passenger-flow parameters of the model with publicly available, official data, rather than relying on synthetic or proprietary datasets. The availability of this information is a decisive practical argument: it ensures that the case study is reproducible, that its results can be discussed transparently, and that the model can later be updated as new data are released or more precise ones are given.

6.1.2 SIMPLIFYING THE NETWORK

To simplify the complete network, while keeping the main passenger flows and important zones, we will tackle in this case study the following simplified network, composed of 10 Macro-Zones (hereinafter called by zones):




Figure 3: Model of the Simplified Network

These zones can be thought of as clusters of stations, where the passenger flow through each zone approximates the average flow across the geographical area it represents. This aggregation reduces the complexity of the model significantly, allowing the entire Madrid Cercanías network to be captured within only 10 representative zones.

Below can be found a summary of the macro-zones, examples of stations in this zone, and their role in the model:

Macro-Zone ID	Name	Example of stations	Role in Model
S_1	Central Core (Sur)	Atocha, Méndez Álvaro, Delicias, Embajadores	The Hub. Massive alighting (jobs) and transfer point. High capacity
S_2	Chamartín Hub	Chamartín, Fuente de la Mora	North Connector. Major interchange and entry point for northern lines.
S_3	Corredor Norte	Cantoblanco, Alcobendas-S.S. Reyes, Colmenar Viejo, Aeropuerto	Mixed North. Airport and University.
S_4	Corredor Guadalajara	Alcalá de Henares, Torrejón, Coslada, Vicálvaro	Industrial/Residential East. Massive population + Industrial jobs. High flow to Z1/Z2.
S_5	Corredor Sur-Este	Aranjuez, Ciempozuelos, Valdemoro	Outer South. Longer distance commuters, lower frequency.
S_6	Corredor Sur	Parla, Getafe, Villaverde	Residential South. Extremely high population density ("dormitory cities"). High morning flow to Z1/Z2.
S_7	Corredor Suroeste	Leganés, Fuenlabrada, Humanes	The "Heavyweight". Highest frequency (4 min) and highest passenger volume.
S_8	Corredor Móstoles	Móstoles, Móstoles-El Soto, Las Retamas	High-Density Residential Terminus. Represents the western end of the highly congested C-5 line.
S_9	Corredor Oeste	Pozuelo, Las Rozas, Príncipe Pío, Villalba	Mixed West. Wealthier residential areas + universities. Lower density but steady flow.
S_{10}	Corredor Bifurcado Norte	Colmenar Viejo, Tres Cantos, Alcobendas-S.S. Reyes, Valde las fuentes	Northern Dual-Branch Feeder. Captures the bifurcated endpoints of lines C-4a and C-4b.

Table 2: Zone Description of the Simplified Model

To calculate the parameters for each zone, we will assume that a zone is an average of all the stations included in it. Thus, the value of assets A_{s_i} for a zone s_i will be $A_{s_i} = \frac{1}{|s_i|} \sum_{s \in s_i} A_s$.

There are approximately 90 stations in Madrid, and we make the assumption that each macro-zone contains 9 stations. This will be used later when calculating the cost of the defensive strategies.

6.2 *ESTIMATION OF THE PARAMETERS*

6.2.1 SYMBOLIC IMPORTANCE (S_s)

The symbolism level captures the intangible but operationally significant dimension of a station's attractiveness as a terrorist target, beyond its physical footprint or passenger volume. As documented in the aftermath of the 2004 Madrid bombings and consistent with the DHS "Soft Target" methodology, terrorist actors do not select targets solely on the basis of expected casualties, they also seek to maximize psychological impact, media exposure, and the symbolic message conveyed by the attack location. A strike on a historically or politically charged landmark inflicts a disproportionate blow to national morale and institutional confidence relative to its material damage alone. Incorporating into the threat model therefore ensures that the optimizer does not systematically under-protect high-visibility stations simply because they are not the busiest ones.

Following the DHS framework, is estimated along two axes: **Iconic Value**, reflecting the station's historical, architectural, and cultural significance in the national identity, and **Political Significance**, reflecting proximity to governmental, economic, or international institutions whose disruption would carry a strategic message beyond the immediate attack site. Rather than discretizing zones into a few rigid categories, each macro-zone is positioned continuously on these two axes, and its symbolic importance score is derived from its joint position. The resulting distribution is illustrated in the figure below, where a clear separation emerges between a small group of high-symbolism hubs (top-right quadrant) and the bulk of utilitarian commuter corridors (bottom-left quadrant).




Figure 4: Symbolic Importance of each Station

The following considerations have been applied to estimate each symbolic value.

- Central Core: $S_{s_1} = 1$. Atocha is the most symbolically loaded station in the entire Spanish rail network: it was the primary site of the 2004 bombings and sits at the heart of Madrid's historical and institutional core, scoring maximum on both axes.
- Chamartín Hub: $S_{s_2} = 0.8$. Chamartín is Madrid's northern gateway and a key node of the city's Central Business District, giving it high political and economic significance, though its iconic value remains slightly below that of Atocha.
- Corredor Norte: $S_{s_3} = 0.8$. This zone includes Aeropuerto T4 and Nuevos Ministerios, granting it maximum political significance through its role as Spain's international gateway and its proximity to government ministries, even if its iconic value is more moderate.

-
- Corredor Guadalajara: $S_{s_4} = 0.4$. The presence of Alcalá de Henares — a UNESCO World Heritage city and birthplace of Cervantes, gives this zone a non-negligible cultural footprint, but its overall political weight remains low.
 - Corredor Sur-Este: $S_{s_5} = 0.2$. A peripheral commuter corridor with long-distance, low-frequency service and no significant institutional or cultural landmarks, placing it near the origin on both axes.
 - Corredor Sur: $S_{s_6} = 0.2$. Despite its very high passenger volumes, this dormitory-city corridor (Parla, Getafe, Villaverde) carries minimal symbolic weight, as its function is almost purely utilitarian.
 - Corredor Suroeste: $S_{s_7} = 0.3$. The "heavyweight" of the network in passenger terms, but its symbolic profile remains modest; the slight uplift reflects the visibility that a successful attack on the network's busiest corridor would generate.
 - Corredor Móstoles: $S_{s_8} = 0.3$. The western terminus of the highly congested C-5 line; its symbolic value comes primarily from its operational prominence rather than from any iconic or political dimension.
 - Corredor Oeste $S_{s_9} = 0.4$. Includes Príncipe Pío (a recognizable Madrid landmark with historical resonance) and several wealthier residential and university areas, giving it a slightly elevated profile compared to other peripheral corridors.
 - Corredor Norte Bifurcado $S_{s_{10}} = 0.3$. A peripheral feeder zone whose moderate score reflects the partial overlap with northern growth areas, but lacking any major iconic or political anchor of its own.

6.2.2 NUMBER OF PEOPLE ($N_{s,t}$)

The casualty potential of an attack on station at time is directly proportional to the number of passengers $N_{s,t}$ physically present in the station at that moment. Estimating this quantity for the Cercanías Madrid network is non-trivial: Renfe publishes aggregated daily ridership figures, but provides no direct hourly or station-level breakdown. To bridge this gap, we develop a four-step estimation pipeline that combines (i) a temporal disaggregation of daily demand, (ii) a gravity-based spatial distribution between origin and destination zones, (iii) a directional split capturing the asymmetry of commuter flows, and (iv) a queuing-theory application of Little's Law to convert passenger flows into station occupancy.

6.2.2.1 Hourly Network Occupancy (N_h)

Daily ridership cannot be uniformly distributed across the 24 hours of operation, since commuter rail demand is heavily concentrated during morning and evening peaks. To capture this, the number of passengers present in the network during hour h is estimated as:

$$N_h = N_{tot} \cdot \frac{F_h \cdot \Theta_h}{\sum_{t=1}^{24} (F_t \cdot \Theta_t)} \quad (39)$$

Where N_{tot} is the total daily ridership reported by Renfe, F_h is the train frequency at hour h (trains/hour), and $\Theta_h \in [0.1]$ is the hourly occupancy coefficient representing how full the trains are. The product $F_h \cdot \Theta_h$ acts as a proxy for the effective transport supply consumed during hour h , and the normalization ensures consistency with the total daily passenger count. The values used for F_h and Θ_h , reported in the corresponding table, were chosen to reflect the well-known bimodal commuting pattern of Madrid, with peak hours in the 07:00–09:00 and 18:00–20:00 intervals. For notational simplicity, we define β_h such that $N_h = N_{tot} \cdot \beta_h$, with β_h summarized in the second table.

Time slots (h)	Frequency (f_h)	Occupancy Coeff (Θ_h)
06:00 - 07:00	6 trains/h	0.5
07:00 - 09:00	20 trains/h	1
09:00 - 13:00	12 trains/h	0.4
13:00 - 17:00	15 trains/h	0.7
17:00 - 21:00	20 trains/h	0.9
21:00 - 00:00	4 trains/h	0.2

Table 3: Frequency and Occupancy per Time Slot

Hours (h)	β_h
06:00 - 07:00	5
07:00 - 09:00	35
09:00 - 13:00	8
13:00 - 17:00	18
17:00 - 21:00	32
21:00 - 00:00	1

Table 4: β_h per Time Slot [%]

6.2.2.2 Spatial Distribution: The Gravity Model

Once the total hourly volume N_h is known, it must be distributed across all origin–destination (O–D) pairs of macro-zones. To this end, we use a classical gravity model, widely used in transport planning, where the unnormalized weight W_{ij} of trips from zone i to zone j is given by:

$$W_{ij} = \frac{P_i \cdot A_j}{D_{ij}}$$

Here P_i is the population of the origin zone (the pool of potential travelers) and A_j is the attractiveness of the destination zone: capturing jobs, university seats, and special hub workers (e.g. airport, government ministries). D_{ij} is the average travel time in minutes between the geographic centroids of the two zones. The values of P_i and A_j were obtained from the Madrid Town Hall's statistical portal [23], while d_{ij} was extracted from Google Maps queries for a 1:00 p.m. trip between the centroids of each pair of zones (chosen to avoid peak-hour congestion bias). The resulting normalized weight matrix W_{ij} is provided in the corresponding table.

	Dest: s_1	s_2	s_3	s_4	s_5	s_6	s_7	s_8	s_9	s_{10}
Orig: s_1		5.3	0.4	0.6	0.2	0.9	1.0	2.8	0.4	0.7
s_2	8.7		1.3	0.5	0.2	0.6	0.6	1.7	0.6	2.3
s_3	0.4	0.8		0.0	0.0	0.0	0.0	0.1	0.0	0.1
s_4	3.7	2.1	0.3		0.1	0.4	0.4	1.0	0.2	0.4
s_5	1.7	0.9	0.1	0.1		0.2	0.2	0.4	0.1	0.2
s_6	6.5	2.7	0.3	0.4	0.1		0.5	1.3	0.2	0.5
s_7	8.9	3.4	0.4	0.5	0.2	0.6		1.7	0.3	0.6
s_8	7.7	2.9	0.3	0.4	0.1	0.5	0.5		0.2	0.5
s_9	2.0	2.0	0.2	0.2	0.0	0.2	0.2	0.5		0.3
s_{10}	2.6	5.0	0.4	0.2	0.1	0.2	0.2	0.6	0.2	

Table 5: Normalized Weight Matrix of W_{ij} values [%]

6.2.2.3 Directional Split: Capturing Commuting Asymmetry

A purely symmetric gravity formulation would fail to capture a fundamental feature of suburban rail demand: in the morning, flows are overwhelmingly directed from the periphery toward the central business district, while in the evening the direction reverses. To incorporate this asymmetry, the macro-zones are partitioned into two subsets:

$$V_{core} = \{s_1, s_2\}$$

$$V_{perip} = \{s_3, \dots, s_{10}\}$$

A time-dependent directional split $\Delta_{in}(h)$ then specifies the share of inbound (periphery $\rightarrow$ core) flow during hour h , with $\Delta_{out}(h) = 1 - \Delta_{in}(h)$ being the share of outbound (core $\rightarrow$ periphery) flow. The values of $\Delta_{in}(h)$, reported in the corresponding table, are based on the *Encuesta Domiciliaria de Movilidad 2018* [24] and follow the expected pattern: $\Delta_{in} = 0.85$ during the morning peak (07:00–09:00) and $\Delta_{in} = 0.30$ during the evening peak (18:00–20:00). The hourly passenger demand between zones and is then given by:

$$T_{ij}(h) = \begin{cases} \frac{\Delta_{in}(h) \cdot N_h \cdot W_{ij}}{\sum_{p:s_p \in V_{perip}} \sum_{q:s_q \in V_{core}} W_{pq}} & \text{if } s_j \in V_{core} \\ \frac{\Delta_{out}(h) \cdot N_h \cdot W_{ji}}{\sum_{p:s_p \in V_{perip}} \sum_{q:s_q \in V_{core}} W_{qp}} & \text{if } s_i \in V_{core} \end{cases}$$

In this formulation, only the destination matters when classifying a trip as inbound or outbound: what we ultimately track is *where the passenger ends up*, which determines where they pose an exposure risk. An illustrative example of the resulting $\Delta(h) \times W_{ij}$ distribution for the 06:00–07:00 interval is provided in the corresponding table.

	Dest: s_1	s_2	s_3	s_4	s_5	s_6	s_7	s_8	s_9	s_{10}
Orig: s_1		7.0	0.1	0.2	0.1	0.3	0.3	0.9	0.1	0.2
s_2	11.4		0.4	0.2	0.1	0.2	0.2	0.6	0.2	0.8
s_3	0.6	1.0		0.0	0.0	0.0	0.0	0.0	0.0	0.0
s_4	4.8	2.7	0.1		0.0	0.1	0.1	0.3	0.1	0.1
s_5	2.3	1.2	0.0	0.0		0.1	0.1	0.1	0.0	0.1
s_6	8.5	3.5	0.1	0.1	0.0		0.2	0.4	0.1	0.2
s_7	11.5	4.5	0.1	0.2	0.1	0.2		0.5	0.1	0.2
s_8	10.0	3.8	0.1	0.1	0.0	0.2	0.2		0.3	0.2
s_9	2.7	2.6	0.1	0.1	0.0	0.1	0.1	0.6		0.1
s_{10}	3.4	6.5	0.1	0.1	0.0	0.1	0.1	0.8	0.2	

Table 6: Example of $\Delta(h) \times W_{ij}$ values in the time slot 6:00-7:00 [%]

6.2.2.4 From Flow to Occupancy: Little's Law

The matrix $T_{ij}(h)$ provides flows (passengers per hour), but the casualty function requires occupancy (passengers physically present in the station at a given moment). The bridge between the two is provided by Little's Law, a foundational result in queuing theory: in a stationary system, the long-term average number of people present $N_{s,t}$ equals the arrival rate λ multiplied by the average time W each person spends in the system. Applied to a station, this gives:

$$N_{s,t} = (\Lambda_{board,s,t} \cdot W_{board}) + (\Lambda_{aligh,s,t} \cdot W_{aligh})$$

where $\Lambda_{board,s,t} = \sum_{j \in \mathcal{S}} T_{sj}(t)$ is the rate of passengers starting their trip at s at time t , and $\Lambda_{aligh,s,t} = \sum_{i \in \mathcal{S}} T_{is}(t)$ is the rate of passengers ending their trip at s at time t . The average dwell times are set to $W_{board} = 6 \text{ min} = 0.1 \text{ h}$ and $W_{aligh} = 2 \text{ min} \approx 0.033 \text{ h}$, reflecting the fact that boarding passengers typically wait on the platform for the next train (consistent with the average headway during off-peak hours), whereas alighting passengers traverse the station relatively quickly on their way out.

This four-step pipeline produces a station-level, time-dependent occupancy estimate $N_{s,t}$ that is both grounded in publicly available data and theoretically consistent with established transport-planning practice.

In the table below can be seen the values of passenger flows for each time slot of the 01/01/2023:

The number of passengers physically affected by the attack $N_{s,t}$:

	t1	t2	t3	t4	t5	t6
	(6:00-7:00)	(7:00-9:00)	(9:00-13:00)	(13:00-17:00)	(17:00-21:00)	(21:00-00:00)
s1	242	865	74	266	349	13
s2	217	773	70	259	345	14
s3	18	62	8	34	48	2
s4	76	269	26	103	139	6
s5	35	122	12	46	62	3
s6	118	419	38	144	192	8
s7	155	554	49	182	241	10
s8	143	486	46	177	238	10

s9	58	186	18	68	91	4
s10	105	346	32	121	161	7

Table 7: Number of Passengers Affected per Station and Time Slot

6.2.3 BETWEENNESS CENTRALITY ($B_{s,t}$)

To capture the structural role of each macro-zone within the Cercanías network, we compute the weighted betweenness centrality of every node for each time slot of a representative day.

The key modelling choice lies in the definition of the weight of the edge. Since shortest-path algorithms minimize cumulative weight, but we want paths to prefer high-flow links (those carrying the bulk of passenger demand), we define the cost of an edge between two zones as the **inverse of the passenger flow** circulating on it.

In this way, heavily used links translate into short distances and are naturally favored by Dijkstra's algorithm, while low-flow links act as long detours. Edges with zero flow are excluded from the graph altogether, as they are effectively non-existent from a passenger-circulation standpoint.

In practice, for each time slot of a representative day (the flow pattern being assumed to repeat identically across days), an undirected graph is constructed from the O–D flow matrix $T_{ij}(h)$ derived in the previous section, and the betweenness centrality is calculated in Python using the networkx library.

In the table below can be seen the betweenness centrality for each time slot of the 01/01/2023:

The betweenness centrality of the network $B_{s,t}$

	t1	t2	t3	t4	t5	t6
	(6:00-7:00)	(7:00-9:00)	(9:00-13:00)	(13:00-17:00)	(17:00-21:00)	(21:00-00:00)
s1	0.833333	0.833333	0.833333	0.805556	0.805556	0.333333
s2	0.416667	0.416667	0.416667	0.361111	0.194444	0
s3	0	0	0	0	0	0
s4	0	0	0	0	0	0
s5	0	0	0	0	0	0
s6	0	0	0	0	0	0
s7	0	0	0	0	0.138889	0.055556
s8	0	0	0	0	0	0.361111
s9	0	0	0	0	0	0
s10	0	0	0	0	0	0.027778

Table 8: Betweenness Centrality per Station and Time Slot

6.2.4 VALUE OF ASSETS DESTROYED (A_s)

Given the absence of a public inventory of Cercanías physical assets, A_s is estimated at the order-of-magnitude level using a three-tier classification based on passenger volume, which correlates strongly with station size and infrastructure complexity. The values (€8M / €3M / €0.75M) reflect typical ranges observed in Adif's public investment programs for stations of comparable scale, that are displayed in the following benchmark:

Station type	Investment Program	Cost (€)	Source
Major Hub	Improvements in Sol station	7.5 M	[23]
Mid-size urban	Improvements in Parla and Embajadores stations	2.72 M	[24]
Small/peri-urban halt	Improvements in San Cristobal de Los Angeles	556 k	[27]

Table 9: Benchmark of Investments in Stations

Even though it is only the cost of improvements and not the total cost of assets in the station, we assume that this is a pretty good approximation of the value that might get destroyed during a terrorist attack.

Based on this, and with the idea that the value of the assets in the zone must represent the average value of assets in a station:

- For s_1, s_2 (major hubs), $A_s = 8 \text{ M€}$
- For $s_3, s_4, s_6, s_7, s_8, s_9, s_{10}$ (mid-size urban), $A_s = 3 \text{ M€}$
- For s_5 (small halt), $A_s = 0.75 \text{ M€}$

6.2.5 COST OF THE ATTACK (C^A)

Although the attacker's cost does not play a structural role in the Stackelberg equilibrium, it enters the attacker's harm function $h_{s,t}$ as a constant offset that is identical across all targets and therefore does not affect the attacker's relative preference between zones. It must still be assigned a numerical value to keep the payoffs consistent and interpretable. To produce a realistic estimate, we ground our choice in empirical evidence drawn from documented terrorist attacks in Spain and from European-level intelligence reports.

Two historical incidents serve as primary references. The 11-M Madrid train bombings of 2004, which directly motivate the present case study, were carried out at an estimated

operational cost of approximately €105,000 [28]. The 2017 Barcelona attack on La Rambla was executed with a budget of roughly €15,000, according to the Spanish Intelligence Center against Terrorism and Organized Crime. These two figures bracket a wide range of operational scales, from a coordinated, multi-station bombing campaign to a vehicle-ramming attack with rudimentary means.

Europol's 2024 Terrorism Situation and Trend Report (TE-SAT) further highlights a clear and persistent trend toward "low-cost terrorism" in Western Europe: over 75% of terrorist plots are executed with budgets below €30,000 [29]. This statistic provides a robust upper bound for the typical attacker profile our model is designed to represent.

Combining these three sources, we adopt:

$$C^A = 30,000 \text{ €}$$

This value sits at the upper end of the Europol-reported distribution, comfortably above the Barcelona benchmark while remaining well below the 11-M figure, a conservative yet realistic order of magnitude for the operational cost of a plausible attack on a Cercanías zone.

6.2.6 COST OF HUMAN CASUALTY (HC)

A central component of the payoff function is the monetary cost associated with the human casualties caused by an attack. To produce a realistic estimate, we decompose this quantity into three factors: the probability of being physically affected by the explosion (the exposure E), the conditional probabilities of death and injury given exposure (P_d and P_i respectively), and the monetary cost of a death and of an injury (C_{death} and C_{injury}). Formally:

$$HC = E \cdot (P_d \cdot C_{death} + P_i \cdot C_{injury})$$

6.2.6.1 Cost of a death

Following the Dirección General de Tráfico, we adopt the official Spanish value of a statistical life (the amount society is willing to spend to prevent a fatality) set at $C_{dea} = € 2,000,000$ [30].

6.2.6.2 Cost of an injury

According to the Dirección General de Seguros y Fondos de Pensiones, statutory compensation for bodily harm ranges from approximately €10,000 to €400,000, depending on the victim's age and the severity of the injury [31]. We adopt the midpoint of this range, $C_{injury} = € 200,000$, as a representative average across the heterogeneous mix of injuries typically produced by a terrorist explosion.

6.2.6.3 Death-to-injury ratio

Historical analysis of bombing attacks in public spaces consistently shows that fatalities account for roughly one out of every eleven physically affected victims, i.e. a ratio of approximately 1 death for 10 injuries. Combining this ratio with the unit costs above yields an average cost per physically affected person of:

$$P_d \cdot C_{death} + P_i \cdot C_{injury} = \frac{1 \times 2,000,000 + 10 \times 200,000}{11} \approx €363,000 / \text{person}$$

6.2.6.4 Exposure factor

Not every person present in a “Sala de Embarque” at the moment of the attack will be physically harmed: only those within the lethal/injury radius of the blast. Considering an average “Sala de Embarque” surface of approximately 1,250 m² and an explosion impact area of approximately 350 m², we obtain an exposure factor of:

$$E = \frac{350}{1250} \approx 28\%$$

6.2.6.5 Final value

Combining the average cost per affected person with the exposure factor yields:

$$HC = 0.28 \times 363,000 \approx \text{€}100,000 / \text{person}$$

6.2.7 COST OF NETWORK DELAYS (*ND*)

Beyond the direct human toll, a successful attack also disrupts the rail network as a whole: when a station becomes unavailable, passengers are forced to rely on slower or less convenient alternatives, generating a collective loss of time that translates into a tangible economic cost. We estimate this cost using the standard transport-economics formula:

$$ND = Aff \cdot T_{delay} \cdot VoT$$

where *Aff* is the number of passengers affected, T_{delay} is the average additional travel time induced by the disruption, and *VoT* is the monetary value of time per passenger.

6.2.7.1 Number of passengers affected (*Aff*)

Madrid's Estación de Atocha, the main hub of our case study, handles approximately 32 million long-distance and high-speed passengers per year. Dividing by 365 yields a daily flow of about 88,000 passengers for Renfe alone. Accounting for the additional traffic generated by the private operators Ouigo and Iryo, which have entered the Spanish high-speed market in recent years, we round this figure up to $Aff \approx 100,000$ passengers per day

6.2.7.2 Average delay per passenger (T_{delay})

To calibrate T_{delay} , we rely on Renfe's own service standards: tickets are fully reimbursed after a 90-minute delay, which implicitly defines the threshold beyond which passengers are expected to switch to an alternative mode of transport. As a benchmark, a Madrid–Sevilla journey takes approximately 2.5 hours by train versus 6.5 hours by bus, giving a 4-hour penalty

for the modal substitution. We therefore adopt a conservative estimate of $T_{delay} = 4h$ per affected passenger.

6.2.7.3 Value of time (VoT)

According to the European Commission's official transport appraisal guidelines [32], the value of time in Spain is approximately €15/h for business travel and €6/h for leisure travel. Assuming a representative 60/40 business-to-leisure split for long-distance rail passengers, the weighted average value of time is:

$$VoT = 0.6 \times 15 + 0.4 \times 6 = € 11.4 /h$$

6.2.7.4 Final value (ND)

Combining the three components yields:

$$ND = 100,000 \times 4 \times (0.6 \times 15 + 0.4 \times 6) \approx €456,000$$

Rounding up to account for second-order effects not explicitly modelled, such as reputational damage to the operator, missed connections, freight disruption, and the cascading congestion on substitute modes, we adopt:

$$ND = € 500,000$$

This value is then weighted by the betweenness centrality $B_{s,t}$ of each zone in the payoff function, so that the network-delay cost effectively borne after an attack on zone s reflects the structural importance of that zone within the Cercanías network.

6.2.8 DEFENDER'S BUDGET (K_t^D)

The defender's resource constraint K_t^D represents the monetary budget available per hour for security deployment across the Madrid Cercanías network. Since no single public figure directly reports this quantity, we reconstruct it in a top-down way by aggregating the publicly available security expenditures of the two main institutional actors involved: Renfe (the train operator) and Adif (the infrastructure manager). We will then progressively narrow the scope down to the temporal granularity of our model.

According to their respective financial reports:

- Renfe spent € 73.4 M on surveillance and security in 2024 [33].
- Adif will spent € 150 M on security between 2026 and 2028, leading to an annual budget of € 50 M/year.

This yields a combined annual expenditure of € 123.4 M dedicated to the protection of the Spanish rail network. Additional funding flowing directly from the central government for national counter-terrorism purposes is deliberately excluded from this estimate, as its allocation across transport assets is opaque and difficult to attribute with sufficient confidence.

6.2.8.1 Restriction to the Madrid Cercanías network

Since our case study focuses exclusively on Madrid's commuter network, we must isolate the share of this national budget assigned to it. Madrid Cercanías accounts for approximately 40% of total train traffic in Spain (205 M / 507 M in 2023 [34]), and we will assume that the proportion that holds for security spending is lower (~30%), due to the higher cost of protecting high-speed trains. This yields an annual Madrid-Cercanías security budget of:

$$0.30 \cdot € 123.4 M \approx € 37.02 M/year$$

Which accounts for

$$€ 37.02 M/365 = € 101k/day$$

Assuming the same budget for each time slot, and considering that there are 6 time slots during each day, the budget per time slot is :

$$K_t^D = \frac{\text{€ } 101k}{6} \approx \text{€ } 16,900 /time\ slot$$

Even though the time slots have different durations, we will allocate the same budget to each of them. This could be modified in a real-life application, with different budget allocations per time slot depending on their importance. This question will be analyzed in section 6.6.

6.3 *PURE STRATEGIES*

6.3.1 DEFENSE PURE STRATEGIES ($d \in \mathcal{D}$)

At each station and during each time slot, the defender selects one security posture from a discrete set of six strategies, denoted $d_1, d_2, d_3, d_4, d_5, d_6$. The strategies are organized as a nested hierarchy: each posture builds on the previous one by adding a new layer of protection. Each strategy corresponds to a specific combination of human, canine, and technological resources, and is characterized by two quantities:

- a cost per time slot C_d , expressed in euros and entering the budget constraint K_t^D . The 6 time slots being distributed over 18 hours, we will take an average of 3 hours per time slot to calculate the price. Although in practice time slots differ in length, this allows us to give the same importance to each of them.
- a performance score $PS_d \in [0,1]$, representing the probability that the strategy successfully detects or deters an attack attempt at the station where it is deployed. As can be seen, the performance score is treated here as an exogenous parameter. In a real-world deployment of the model, these scores would be determined by domain experts (typically security officers from Renfe and Adif or counter-terrorism specialists) on the basis of historical incident data, red-team exercises, and field experience. The values used in this thesis are therefore illustrative, but the framework itself is agnostic to their precise calibration: any expert-validated set of scores can be plugged into the model without altering its structure.

The six strategies are defined in the following subsections.

6.3.1.1 Strategy d_1 : *Do nothing*

The station receives no specific protection: no action is taken during the time slot. This strategy allows the model to leave a station undefended, if the potential damage isn't worth the cost.

- Cost: € 0/time slot
- Performance Score: 0

6.3.1.2 Strategy d_2 : Active CCTV monitoring

Station CCTV feeds are actively monitored in real time by a trained operator in a control room. This enables the early detection of suspicious behavior (abandoned luggage, loitering, unusual flows) and triggers a coordinated response from external units if needed. This is the cheapest active posture, as it leverages already-installed infrastructure and only requires the marginal cost of a dedicated operator. Assuming one operator per station at €21/h, the cost per time slot for a station is $\text{€ } 20.83 \times 3 = \text{€ } 62.5$. Considering that there are approximately 9 stations in a zone, we have a total cost of € 562.5 per time slot. We choose a performance score of 0.1, because this protection is very basic and unlikely to prevent a terrorist attack.

- Cost: € 562.5/*time slot*
- Performance Score: 0.1

6.3.1.3 Strategy d_3 : Active CCTV Monitoring + Foot Patrol

The CCTV layer is complemented by a team of two uniformed security agents patrolling the concourse and platforms on foot. The patrol provides visible deterrence, a basic intervention capability, and a physical response channel for the alerts raised by the CCTV operator. The four additional agents add $2 \times \text{€ } 20.83 \times 3 \times 9$ to the previous posture, resulting in a cost of € 1687.5.

We improved the performance score to 0.3 in this case.

- Cost: € 1687.5/*time slot*
- Performance Score: 0.3

6.3.1.4 Strategy d_4 : Active CCTV Monitoring + Foot Patrol + Random Screening

On top of the previous posture, random access-point screening is introduced: a subset of passengers is selected at random for bag inspection or handheld metal-detector checks. This option captures most of the deterrent value of systematic screening while remaining compatible with the high throughput of a commuter network, with full screening being operationally infeasible at major Cercanías stations such as Atocha or Sol, which handle tens of thousands of passengers per hour. Random screening typically requires two additional agents at a

dedicated checkpoint, adding $2 \times \text{€ } 20.83 \times 3 \times 9$ to the previous posture, resulting in a cost of € 2812.5.

With this strategy that gives us a way to directly check suspicious behavior, we estimate a performance score of 0.6

- Cost: € 2812.5/*time slot*
- Performance Score: 0.6

6.3.1.5 Strategy d_5 : Active CCTV Monitoring + Foot Patrol + Random Screening + Canine Unit

A K-9 team (one handler and one explosive-detection dog) is added to the posture. Canine units substantially increase the probability of detecting concealed explosive devices, which is the dominant attack modality considered in this study, and complement random screening by allowing targeted secondary inspection of suspicious passengers or luggage. A K-9 unit is more expensive than a standard patrol agent because of the specialized training of both handler and dog (typically 6–12 months), the veterinary and kennel costs, and the limited operational lifespan of the animal. We estimate the loaded cost of a K-9 unit at € 37/h/*station*, adding to the previous posture, for a total of per slot per zone, resulting in a cost of € 3825.

This coverage is close to optimal, and we assume the performance score to be 0.8.

- Cost: € 3825/*time slot*
- Performance Score: 0.8

6.3.1.6 Strategy d_6 : Active CCTV Monitoring + Foot Patrol + Random Screening + Canine Unit + Scanner

The most resource-intensive option: a portable X-ray bag scanner is deployed at the main access point, allowing systematic inspection of luggage in addition to the random screening of passengers themselves. This posture approaches the security level enforced at AVE stations. Unlike a fixed installation, a portable scanner offers the operational flexibility of being

redeployed across stations according to the threat level, but it incurs additional logistical costs related to transport, setup, and secure storage between uses.

The marginal cost of this layer is built up as follows. A professional portable X-ray scanner costs around € 50,000 – € 70,000 and has a useful life of approximately 7 years; amortised over an operational use of ~5,000 hours per year, this yields a capital cost of roughly € 2/h. To this we add one dedicated operator at € 20.83/h, a provision of € 18/h for maintenance, calibration, and consumables (films, batteries, periodic radiation safety checks), an installation and redeployment cost of € 30/h (covering the daily setup/teardown of the unit by a two-person technical team and the associated transport between stations), and a secure storage cost of € 25/h (amortised rental of a certified radiation-safe storage room with restricted access, available 24/7 across the network). The total marginal cost of the scanner layer thus amounts to € 100/h, adding € 2245 and resulting in a total cost for the zone and the time slot of € 3830.

This strategy will have a performance score of 1.

- Cost: € 6075/*time slot*
- Performance Score: 1

6.3.1.7 Summary of defense strategies

Below can be seen a summary of the relation between performance score and defense cost. We can notably see that the strategy d_6 seems very expensive compared to its performance score.




Figure 5: Performance Score and Defense Cost of the Strategies

This can be seen even better on the graph below, representing the ratios $\frac{C_d}{PS_d}$ for $PS_d \neq 0$.




Figure 6: Ratio of Defense Cost to Performance Score

Looking at this graph, it appears that the strategies d_4 and d_5 may have an advantage, although this ratio never appears in the harm function $h_{s,t}$ without other terms in C_d and PS_d appearing. The strategies d_2, d_3 may however be used to protect small stations, and d_6 when the traffic is important, like in Atocha for example.

6.4 *RESULTS: INFLUENCE OF THE ATTACKER'S RATIONALITY*

In this section, we will analyze the influence of the rationality parameter λ . This parameter plays a central role in the model, since it controls how rationally the attacker selects his target: when $\lambda \rightarrow 0$, the attacker behaves as if choosing uniformly at random among the available (segment, slot) pairs, irrespective of their actual payoff, whereas when $\lambda \rightarrow \infty$, the attacker becomes perfectly rational and deterministically selects the option yielding the highest expected utility. Intermediate values of λ correspond to a boundedly rational attacker, who is more likely to choose high-payoff targets but still assigns non-zero probability to less attractive ones. Because λ is, by nature, difficult to estimate from data, it is essential to understand how sensitive the defender's optimal strategy is to its value.

In particular we conduct a sensitivity analysis with respect to λ . We will consider a linear distribution $\{\lambda_k = 10 \cdot k, k \in \{1, \dots, 20\}\}$ to have λ varying from 0 to 200. To isolate its effect from the temporal dimension of the problem, the analysis is restricted to a single time slot, which allows us to focus exclusively on how the defender reallocates resources across stations and strategies as the attacker's rationality varies. The study is organized around three complementary perspectives:

- a. **Mixed Strategy per station:** For each station, we examine how the combination of defensive mixed strategies assigned by the optimizer evolves as λ increases. This reveals whether the defender favors low-cost measures on a large number of station, or concentrated, high-performance ones, when the attacker's degree of rationality changes.
- b. **Aggregate use of each strategy:** We then track the total deployment of each pure strategy across the network as λ varies, which provides a global view of how the type of defensive measures employed evolves with the attacker's behavior. The difference with previous analysis described in a. is that, in a., we look at each station individually and analyze how it is defended using mixed strategies (it includes therefore several pure strategies) over the rationality interval, while in this analysis we look at each pure

strategy individually, and also analyze how it is used over the rationality interval (it includes therefore several stations).

- c. **Budget share per station:** Finally, we look at the share of the total budget allocated to each station as a function of λ . This perspective highlights where the defender chooses to spend and shows how resources shift from a uniform spread (under a near-random attacker) toward the most exposed stations (under a near-rational attacker).

Together, these three views provide a complete picture of how the defender's optimal response adapts to the attacker's rationality, and allow us to identify the values of λ at which qualitative changes in the optimal allocation occur.

6.4.1 DEFENSE MIXED STRATEGY

We now examine, station by station, how the defensive mixed strategy evolves with the attacker's rationality parameter λ . To make the comparison easier to follow, the stations are presented in increasing order of defensive intensity, i.e. starting with those that rely almost exclusively on the cheapest option (d_1 , doing nothing) and ending with those that concentrate the budget on the most comprehensive measure (d_6). Four groups of stations emerge naturally from this order:

- Stations 3 (Aeropuerto) and 5 (Aranjuez) rely exclusively on d_1 (do nothing), regardless of the value of λ . These stations are never selected by the optimizer for any active defense, which reflects the fact that their attacker-utility profile is too low to justify spending budget on them, even against a non-rational attacker that selects targets randomly.




Figure 7: Evolution of strategies with λ for stations 3 and 5

- Stations 4 (Guadalajara) and 9 (Sta Maria Cercedilla Cotos) also rely predominantly on d_1 , but introduce a small share of d_4 (CCTV + Foot Patrol + Random Screening) when the attacker's rationality is low. In other words, when the attacker is close to random and may strike any station with non-negligible probability, the defender finds it worthwhile to deploy a moderate level of protection at these mid-risk stations. As λ increases and the attacker concentrates on the most attractive targets, this protection is withdrawn and reallocated elsewhere.




Figure 8: Evolution of strategies with λ for stations 4 and 9

- Stations 6 (Parla), 7 (Humanes), 8 (Mostoles) and 10 (Colmenar Alcobendas) display a richer behavior, with mixed strategy that vary substantially with λ . These stations sit in the intermediate range of attacker utility, where the optimal level of defense is most sensitive to the attacker's degree of rationality. Each of them is therefore discussed individually in the following paragraphs.
- Stations 1 (Atocha) and 2 (Chamartin), finally, are protected almost exclusively with d_6 (CCTV + Foot Patrol + Random Screening + Canine Unit + Scanner), the most complete and most expensive strategy. A small share of d_5 (CCTV + Foot Patrol + Random Screening + Canine Unit) appears when λ is low, but as soon as the attacker becomes even moderately rational, the defender shifts entirely to d_6 . This is consistent with the fact that these two stations exhibit by far the highest attacker utility in the network, and therefore concentrate the bulk of the defensive budget.




Figure 9: Evolution of strategies with λ for stations 1 and 2

This grouping already reveals a clear pattern: the stations at the two extremes of the risk spectrum (very low or very high attacker utility) display stable allocations across values of λ . The sensitivity to the attacker's rationality is concentrated in the intermediate-risk stations (stations 6, 7, 8 and 10), which we now analyze in detail:

- Station 6 follows a smooth, three-phase pattern. For low rationality ($\lambda \leq 40$), the defender uses a mix of d_1 and d_4 , with d_4 starting at about 44% and decreasing as λ grows. Between $\lambda = 40$ and $\lambda = 100$, the allocation stabilizes on a plateau with

roughly 82% of d_1 and 18% of d_4 . From $\lambda \approx 110$, d_4 progressively disappears, with a small bump of d_5 ($\approx 6\%$) around $\lambda = 120$, and the station is fully abandoned ($d_1 = 1$) from $\lambda = 130$ onwards. The transition is short and almost monotonic.




Figure 10: Evolution of strategies with λ for station 6

- Station 7 is mostly protected by d_4 (between 65% and 85%) for $\lambda \leq 120$. The transition is much sharper than for Station 6: at $\lambda = 130$, d_4 disappears almost completely and is replaced by a strong spike of d_3 ($\approx 95\%$), which decreases quickly over the next two values of λ . After this main spike, a second smaller spike of d_3 ($\approx 19\%$) appears around $\lambda = 180$, together with a small amount of d_6 between $\lambda = 150$ and $\lambda = 180$. The defender therefore does not abandon Station 7 monotonically: he switches from d_4 to d_3 , briefly disengages, comes back with a smaller d_3 effort, and only then leaves the station almost undefended.




Figure 11: Evolution of strategies with λ for station 7

- Station 8 follows the same low- pattern as the others: a stable mix of d_1 ($\approx 45\%$) and d_4 ($\approx 55\%$) for $\lambda \leq 110$. The transition zone, however, is the most active of the four stations and uses almost every available strategy. At $\lambda = 120$, is replaced by a spike of d_5 ($\approx 42\%$); at $\lambda \in [130, 140]$, d_6 appears ($\approx 25\%$ then 18%); at $\lambda = 150$, the defender switches to d_3 ($\approx 37\%$), with d_1 dropping at the same time; and finally d_6 reappears around $\lambda \in [160, 170]$ before the station is fully abandoned at $\lambda = 180$. An important point is that when d_3 is used, d_1 also decreases, so the total probability of having some protection in place remains comparable to when d_6 is used. The defender is essentially choosing between "rarely apply a strong strategy" (d_6 with high d_1) and "more often apply a moderate strategy" (d_3 with lower d_1), and these two options give a similar overall level of deterrence.




Figure 12: Evolution of strategies with λ for station 8

- Station 10 looks like a slightly more complex version of Station 6. For $\lambda \leq 110$, the defender uses a smooth mix of d_1 and d_4 , with d_4 starting at about 52% and decreasing to about 25%. The transition zone is concentrated in a very narrow range: at $\lambda = 120$, d_4 is replaced by a spike of d_5 ($\approx 22\%$), and at $\lambda = 130$, by a spike of d_3 ($\approx 18\%$). From $\lambda = 140$ onwards, the station is fully abandoned. The transition is therefore short but clearly non-monotonic: the defender first tries a more targeted strategy (d_5), then a cheaper one (d_3), before disengaging.




Figure 13: Evolution of strategies with λ for station 10

As it has been seen, the four stations (stations 6, 7, 8 and 10) share the same overall structure: a low-rationality regime where they are partially protected by a mix of d_1 and d_4 , an intermediate plateau where the allocation is stable, and a high-rationality regime where they are fully abandoned. What changes from one station to another is the shape of the transition between the plateau and the abandonment.

Moreover, two main observations stand out. First, the transitions almost always involve short spikes of strategies that are not used elsewhere (d_3 , d_5 , d_6), and these strategies often appear in a non-monotonic order. This shows that near the transition point, several allocations give very similar values, and the optimizer can switch between them with small changes in λ . The example of station 8 is particularly clear: alternating between d_6 with high d_1 , and d_3 with lower d_1 , leads to a comparable level of protection, which explains why the defender can move from one to the other without losing much.

Second, the strategy d_2 (cameras only) is never used in any of the four stations, and, more generally, never used in the network. This suggests that d_2 is dominated: adding cameras without patrols does not bring enough deterrence to justify its cost, while the next step up (d_3 , which adds patrols) is efficient enough to be chosen. In practical terms, cameras only become useful when they are combined with active surveillance.

Overall, the sensitivity of the allocation to λ leads to the defender allocating resources only to s_1 and s_2 when the attacker is highly rational, and protecting 6 of the 10 stations with d_4 when the attacker has low rationality. 4 stations (6 (Parla), 7 (Humanes), 8 (Mostoles) and 10 (Colmenar Alcobendas)) are particularly sensitive to λ during the transition period.

6.4.2 AGGREGATE USE OF EACH STRATEGY

Looking at the use of each strategy across all stations gives a complementary view of the allocation patterns and confirms several of the observations made before.




Figure 14: Evolution of the Use of Strategy d_1 with λ

Strategy 1 (d_1 , no defense) is used in every station except s_1 and s_2 , which are the two highest-value stations and are always strongly protected. For all the other stations, the share of d_1 increases with λ , reaching 100% for high rationality values. This clearly illustrates the progressive abandonment of the lower-value stations as the attacker becomes more rational: when the attacker is almost sure to pick the most valuable target, there is no reason to spend resources on stations he will not attack.

Strategy 2 (d_2 , cameras only) is never used, in any station and for any value of λ . This confirms that this strategy is dominated by combinations of d_1 and d_3 : cameras alone do not bring enough deterrence to justify their cost, while adding patrols (i.e. moving to d_3) makes the strategy efficient enough to be selected.




Figure 15: Evolution of the Use of Strategy d_2 with λ

Strategy 3 (d_3 , cameras + patrols) is only used in short spikes, for λ between 120 and 180, in stations s_7 , s_8 and s_{10} . It never appears outside the transition zone. This shows that d_3 plays the specific role of a transition strategy: it is chosen when the defender is hesitating between maintaining moderate protection and fully disengaging, because it offers a reasonable level of deterrence at a moderate cost.




Figure 16: Evolution of the Use of Strategy d_3 with λ

Strategy 4 (d_4 , cameras + patrols + random screening) is used only for $\lambda < 130$, and across most of the medium-value stations ($s_4, s_6, s_7, s_8, s_9, s_{10}$). It is the dominant strategy when the attacker has low rationality: since the attacker is then likely to spread his attacks across many stations, d_4 is the most efficient way to provide an average level of protection on each of them. Once the attacker becomes more rational, d_4 disappears in favor of either full abandonment (d_1) or stronger strategies (d_5 , cameras+patrols+random screening+dogs, d_6 , full protection) at the most valuable stations.




Figure 17: Evolution of the Use of Strategy d_4 with λ

Strategy 5 (d_5 , cameras + patrols + random screening + dogs) is used in two very different situations. First, it is used heavily at s_1 and s_2 for very low rationality ($\lambda \leq 40$), where it represents the main defense before d_6 takes over. Second, it appears as a single spike around $\lambda = 120$ in three other stations (s_6, s_8, s_{10}), in the transition zone. Outside of these two cases, d_5 is not used.




Figure 18: Evolution of the Use of Strategy d_5 with λ

Strategy 6 (d_6 , full protection) is used almost all the time at s_1 and s_2 , reaching 100% as soon as $\lambda \geq 25$ for s_1 , and as $\lambda \geq 25$ for s_2 . For the other stations, it only appears in the transition zone, mostly at s_7 and s_8 , when the attacker becomes highly rational and the defender briefly invests heavily in a medium-value station before abandoning it. d_6 is therefore the standard choice for the most valuable stations, and an occasional transition strategy for the medium-value ones.




Figure 19: Evolution of the Use of Strategy d_6 with λ

The aggregate view confirms a clear structure. The two top-value stations (s_1 and s_2) are protected by the strongest strategies (d_5 and especially d_6) almost regardless of λ . The medium-value stations rely mainly on d_4 as long as the attacker has low to moderate rationality, and are then progressively abandoned, with d_3 , d_5 and d_6 appearing only as short spikes during the transition. The lowest-value stations stay at d_1 throughout. Finally, d_2 is never used, which is a meaningful result on its own: it indicates that, in this setting, cameras only become a worthwhile investment when they are paired with active surveillance.

6.4.3 BUDGET SHARE PER STATION

Below is presented a graph showing the share of the total budget spent at each station, and giving a clear synthesis of the defender's behavior as the attacker's rationality grows.




Figure 20: Evolution of the Share of the Budget Spent per Strategy with λ

The first observation concerns the two highest-value stations, s_1 and s_2 . Both of them quickly converge to the maximum budget that can be spent on a single station: s_1 already reaches this cap at $\lambda = 20$, and s_2 joins it at $\lambda = 40$. From that point on, they stay at the maximum for every higher value of λ . This is consistent with the strategy analysis: s_1 and s_2 are essentially always defended with d_6 , which is the most expensive option, and the defender invests as much as possible in them as soon as the attacker becomes even moderately rational.

For all the other stations, the share of the total budget decreases monotonically with λ . The medium-value stations ($s_4, s_6, s_7, s_8, s_{10}$) start with non-negligible shares (between 4% and 14% depending on the station) and progressively lose budget as λ increases, until they reach 0%. The lowest-value stations (s_3, s_5, s_9) are almost never funded. This confirms the progressive abandonment already observed in the strategy graphs: as the attacker becomes

more rational, the defender concentrates her resources on the few stations the attacker is actually likely to target.

The most interesting feature of the graph is the behavior of the total budget used. For $\lambda \leq 120$, the defender always spends 100% of the available budget. But from $\lambda=120$ onwards, the total budget used starts to decrease, and it stabilizes around 74% for high values of λ . This threshold has a clear interpretation: when the attacker is rational enough, the probability that he attacks a low-value station, multiplied by the harm such an attack would cause, becomes lower than the cost of the cheapest useful strategy. In other words, defending these stations is no longer worth it: the expected damage avoided is smaller than the price paid to avoid it. The defender then prefers to leave part of the budget unused rather than spending it inefficiently.

This threshold at $\lambda = 120$ also matches the transition zone identified in the strategy graphs, where d_3, d_5 and d_6 briefly appear at medium-value stations before disappearing. It marks the point where the optimization problem changes nature: below this threshold, the defender tries to cover as many stations as possible with the available budget; above it, the defender focuses exclusively on the few stations that truly matter and accepts that part of the budget remains unspent.

6.5 RESULTS: INFLUENCE OF THE DEFENDER'S BUDGET

After analyzing the impact of the attacker's rationality, we now turn to a second key parameter of the model: the total budget K_t^D available to the defender for one time slot. Until now, this budget has been fixed at its baseline value ($K_t^D = € 16,900$). We will now let it vary linearly from € 5,000 to € 30,000 and observe how the optimal allocation reacts.

To make the analysis meaningful across the different regimes identified in the previous section, we will repeat it for three representative values of the attacker's rationality:

- $\lambda = 50$, which corresponds to the low-rationality regime, where the attacker spreads his probability of attack across many stations and the defender protects most of the network with d_4 .
- $\lambda = 125$, which lies inside the transition zone, where the allocation is the most sensitive and several strategies (d_3, d_5, d_6) appear in short spikes.
- $\lambda = 200$, which corresponds to the high-rationality regime, where the attacker concentrates almost surely on the highest-value stations and the defender abandons most of the network because the probability of the attacker targeting these stations is too low to justify this spending.

For each of these three values of λ , we will follow the same structure as in the previous section. We will first look at how the strategy assigned to each station evolves with the budget (subsection 6.4.1), then at the share of the total budget spent at each station and at how much of the global budget is actually used (subsection 6.4.2), and finally at the evolution of the usage of each individual strategy across the network (subsection 6.4.3). This will allow us to see whether increasing the budget mainly leads to better protection of the already-defended stations, to the protection of new stations, or to a different combination depending on the rationality regime.

6.5.1 MIXED STRATEGY PER STATION

In this section, we are going to study, for each station and each value of λ , how the mixed strategy evolves when the budget increases.

6.5.1.1 Stations s_1 and s_2

Stations s_1 and s_2 are the two most valuable stations of the network, and as expected they are the first ones the defender wants to protect. **Error! Reference source not found.** and **Error! Reference source not found.** presents three graphs that show how the allocation at s_1 and s_2 evolves with the available budget for $\lambda = 50$, $\lambda = 125$ and $\lambda = 200$. As can be seen, the three cases follow the same qualitative pattern, which can be described as a progressive climb up the strategy ladder as the budget increases.




Figure 21: Mixed strategy for s_1 when $\lambda = 50, 125$ and 200




Figure 22: Mixed strategy for s_2 when $\lambda = 50, 125$ and 200

Since behavior across the budget range is very similar in both stations s_1 and s_2 , we take s_2 as an illustrative example:

For very low budgets (around $K_t^D = \text{€ } 5,000$), the defender cannot afford to protect s_2 properly. She falls back on d_1 (no defense) for a fraction of the allocation, and uses d_5 for the rest. The share of d_1 at this point ranges from roughly 30% for $\lambda = 50$ to nearly 100% for $\lambda = 200$: when the attacker is rational, s_1 is targeted first, so s_2 is sacrificed when the budget is too tight.

As the budget grows, d_1 progressively disappears and the allocation shifts to d_5 , which becomes the dominant strategy for budgets between approximately $\text{€ } 7000$ and $\text{€ } 11000$.

Beyond a second threshold (around $\text{€ } 11000\text{--}\text{€ } 13000$), d_5 is in turn replaced by d_6 , the full-protection strategy, which then takes 100% of the allocation for every higher budget.

The three regimes of λ differ mainly in the speed and the order of this transition. For $\lambda = 50$, the attacker spreads his attacks across many stations, so s_2 is treated almost symmetrically to s_1 : the defender invests in it as soon as possible, and the switch to d_6 happens around a budget of $\text{€ } 12000$. For $\lambda = 125$ and especially $\lambda = 200$, the attacker concentrates on s_1 first, so s_2 is "served second": the defender keeps s_2 at d_1 for a wider range of low budgets and only switches to d_5 and then d_6 once s_1 is fully protected. This is why the transitions at s_2 happen at slightly higher budgets when λ is high: the defender first saturates s_1 before redirecting additional budget towards s_2 .

This confirms that, for the highest-value stations (s_1 and s_2), the budget acts as a simple capacity constraint: the defender always wants to use the strongest possible strategy, and the only question is whether she can afford it. The intermediate strategies d_2 , d_3 and d_4 are essentially never used at s_1 and s_2 . They are dominated by the $d_1 \rightarrow d_5 \rightarrow d_6$ progression.

6.5.1.2 Stations s_3 and s_5

Stations s_3 and s_5 , in Figure 23 and Figure 24, are at the opposite end of the value spectrum: they are among the least valuable stations of the network.

The graphs below make this very clear: regardless of the rationality λ of the attacker and regardless of the available budget, the allocation stays at 100% at d_1 over the entire budget range.

This means that even when the defender has a very large budget at her disposal (up to € 30000), she still chooses to leave s_3 and s_5 completely undefended. The reason is that the expected loss caused by an attack on these stations is small enough that any euro spent on defending them would yield a better marginal return if reallocated to a more valuable station. In other words, s_3 and s_5 are dominated stations: they are never worth protecting, no matter how rich the defender is or how rational the attacker is.

This is also consistent with what we observed in the previous section: across the full range of λ , s_3 and s_5 were the only two stations using exclusively d_1 . The budget analysis confirms that this is a structural property of these stations, not an artifact of a particular budget level: they are simply too low-value to ever enter the defender's optimal allocation.




Figure 23: Strategy mix for s_5 when $\lambda = 50, 125$ and 200




Figure 24: Strategy mix for s_3 when $\lambda = 50, 125$ and 200

6.5.1.3 Stations s_4, s_6, s_7, s_8, s_9 and s_{10}

Next, we examine the strategic configurations of the remaining medium-value stations (s_4, s_6, s_7, s_8, s_9 and s_{10}). These six remaining stations occupy the middle and lower part of the value ranking, and their behavior is much more sensitive to both the budget and the rationality λ of the attacker than the extreme stations analyzed before. Although their overall behavior is similar, distinct variations emerge when evaluated across different rationality frameworks. To provide an exhaustive analysis, we compare these stations under each specific rationality model. In particular, three clear regimes emerge, one per value of λ .

Low rationality ($\lambda = 50$)

The results for low rationality ($\lambda = 50$) are displayed in Figure 25. When the attacker is almost uniform in his choice of target, the defender wants to spread her protection across many stations. As a result, once s_1 and s_2 are taken care of, the additional budget is progressively redirected to the intermediate stations through the strategy d_4 , which provides a balanced, moderately costly protection. The pattern is the same for all six stations: a long plateau at 100% d_1 at low budgets, then a smooth $d_1 \rightarrow d_4$ transition starting around €12000–€15000, and d_4 becoming dominant for the highest budgets. The most valuable stations among this group (s_7, s_8) even reach a second transition $d_4 \rightarrow d_5$ near a budget of €25000–€27000, showing that with enough money the defender further upgrades their protection. The least valuable (s_9, s_{10}) only reach the d_4 stage by the end of the budget range. This regime is the one where the largest variety of strategies is used.




Figure 25: Strategy mix for s_4 , s_6 , s_7 , s_8 , s_9 and s_{10} when $\lambda = 50$

Intermediate rationality ($\lambda = 125$)

This is the transition regime, displayed in Figure 26, and it is also the most interesting one because we can see a very clear transition for some stations, with often surprising evolutions of the mixed strategies. For the low-value stations s_4 and s_9 , the rationality level is already too high for the defender to protect it, because of the unlikeliness of an attack. It is similar for s_6 and s_{10} for low budget value, but after crossing a threshold (€ 17,000 for s_6 and € 14,000 for

s_{10}), they start to be defended as well, indicating that they are worth defending, but only when everything else is already defended.

The last two stations, s_7 and s_8 , have similar behaviors: they start undefended, while all the budget is allocated to s_1 and s_2 , and then quickly receive partial protection with d_4 , that will grow until we reach a budget of € 16,000. From this point, s_8 will follow the logical path $d_4 \rightarrow d_5 \rightarrow d_6$, with a spike of 35% d_5 around € 17,000 and 35% d_6 after. Station s_7 , however, follows a more unusual path. After d_4 , the station is protected by a mix of d_3 and d_6 (and even a small amount of d_2 for $K_t^D = € 17,000$, the only use of d_2 we will witness). Looking at the exact data, the mixed strategy stabilizes around (74.3% d_3 , 25.7% d_5), while for s_8 , it stabilizes at (2/3 d_1 , 1/3 d_6). To explain the difference, we can look at the performance scores and the costs of the two mixed strategies:

	Cost (€)	Performance Score
(74.3% d_3 , 25.7% d_5)	2,237	0.429
(2/3 d_1 , 1/3 d_6)	2,025	0.333

Table 10: Strategy Comparison for s_7 and s_8

We can see here that s_7 is defended with a better and more costly strategy, which makes sense since it's considered a higher-value station. However, reaching the performance score of 1/3 using the mixed strategy over (d_3, d_5) for s_8 would require the mixed strategy (93.3% d_3 , 6.7% d_5), that costs € 1,830. We could therefore reach the same result for a cheaper cost: the mixed strategy over (d_3, d_5) is structurally better than over (d_1, d_6). The reason why (d_1, d_6) is chosen will be explained with more details in Figure 35, but it is basically because at this point, the defender has "too much money", and the way he chooses to defend the stations doesn't really matter, as long as he reaches the desired performance score. The mixed strategy over (d_1, d_6) is widely used even with budget constraints, but it is because it allows to reach performance scores >0.8 , which is something that you can't do with (d_3, d_5).




Figure 26: Strategy mix for s_4, s_6, s_7, s_8, s_9 and s_{10} when $\lambda = 125$

High rationality ($\lambda = 200$)

The results for high rationality are displayed in Figure 27. When the attacker is sharp, he concentrates almost all of his attention on s_1 and s_2 , and the marginal value of protecting any other station collapses. As a result, the allocation at s_4, s_6, s_7, s_8, s_9 and s_{10} stays at 100% d_1 over essentially the whole budget range, with at most a very small share of d_5 or d_6 appearing at the most valuable of them (s_7) once s_1 and s_2 are fully saturated. This confirms a key insight:

a rational attacker effectively reduces the defender's optimization to the top of the value ranking, and the rest of the network is left undefended even when the budget is large.




Figure 27: Strategy mix for s_4, s_6, s_7, s_8, s_9 and s_{10} when $\lambda = 200$

6.5.2 AGGREGATE USE OF EACH STRATEGY

Once again, looking at the individual behavior of the strategies confirms the insights discovered before.

Strategy 1 (no defense).

We can see in Figure 28:

- Low rationality ($\lambda = 50$). All stations except s_3 and s_5 , progressively leave d_1 as the budget grows, and the order in which they do so follows exactly the value ranking already identified station by station (s_1, s_2 first, then s_7, s_8 then s_4, s_6, s_{10} , and finally s_9). This confirms that with a weak attacker the defender spreads her budget across many stations, protecting them one by one from the most to the least valuable.
- Intermediate rationality ($\lambda = 125$). Only s_1, s_2 and s_7 fully leave d_1 , while s_4, s_8 and s_{10} drop to an intermediate plateau and the rest stays at 100%. This is the aggregate signature of the short d_3 spikes seen earlier on the medium-value stations: they get partial protection but the defender never fully commits to defending them.
- High rationality ($\lambda = 200$). Only s_1 and s_2 leave 100% d_1 , and every other station stays flat at the top regardless of the budget. This matches the previous observation that a rational attacker collapses the problem to the two most valuable stations, and any extra budget is poured into them rather than spread over the network.




Figure 28: Use of d_1 when $\lambda = 50, 125$ and 200

Strategy 2 (cameras only)

Once again, as we can see in Figure 29, d_2 is almost never used, in any station and for any value of λ . We notice a small usage at s_7 for $\lambda = 125$, but this strategy is definitely too weak.




Figure 29: Use of d_2 when $\lambda = 50, 125$ and 200

Strategy 3 (cameras + patrols).

The strategy 3, displayed in Figure 30, is only used for $\lambda = 125$, and only for the station s_7 . Its low usage rate confirms a position of transition strategy, very sensitive to the parameters of the model. Interestingly, when used in s_7 , we notice a plateau from $K_t^D \approx \text{€ } 16,000$, meaning that even with more budget, the defender doesn't want to move to a better strategy, indicating an equilibrium point between the characteristics of s_7 , λ , and the cost of defense: a small change in any of these parameters is likely to trigger a change of strategy.




Figure 30: Use of d_3 when $\lambda = 50, 125$ and 200

Strategy 4 (cameras + patrols + random screening).

We can see in Figure 31 that the use of d_4 decreases with λ , which we already knew, but also that it has a very specific behavior: All the curves of d_4 with respect to the budget are concave,

and reach a maximum point that depends on the station and λ . This makes sense, because we have seen earlier that d_4 is the preferred strategy to defend moderately stations that are worth defending. With the budget increasing, we can first defend more stations that are worth defending, as we can see for $\lambda = 50$, until we have a sufficient budget to defend them with better strategies, and abandon d_4 to replace it with d_5 or d_6 .




Figure 31: Use of d_4 when $\lambda = 50, 125$ and 200

Strategy 5 (cameras + patrols + random screening + dogs).

The behavior of d_5 is very similar to the one of d_4 , as shown in Figure 32, except it appears later for medium-value stations, and disappears quicker for them with λ . We can also see that it lasts a bit longer for high-value stations. A notable difference is that for some stations (s_7, s_8) and rationality level ($\lambda = 50, 125$), it seems to plateau, leading to think that d_5 could be a viable alternative in some case, even with unlimited budget.




Figure 32: Use of d_5 when $\lambda = 50, 125$ and 200

Strategy 6 (full protection).

Strategy 6 is displayed in Figure 33. We can see that for the 3 rationality levels, d_6 is quickly adopted by the two high-value stations, as soon as the available budget allows it. For $\lambda = 50$, no other station uses d_6 , since it's more important to protect all the stations, and only cheaper strategies allow that. For $\lambda = 125$, we can see that a few strategies will reach an equilibrium state using fractions of d_6 , coupled with d_1 . The same goes for $\lambda = 200$ and s_7 .




Figure 33: Use of d_6 when $\lambda = 50, 125$ and 200

Globally, this shows that increasing the budget always favors better strategies, since we can see d_1 decreasing, and the cycle $d_4 \rightarrow d_5 \rightarrow d_6$ in many configurations. It doesn't show, however, which part of the budget is actually used, and if, with more budget available, it is still worth not defending a station.

6.5.3 BUDGET SHARE PER STATION

To understand how the budget distribution changes across the stations, when the total budget increases, we will analyze, for the three levels of rationality, the budget share per station, and the % of the total budget used overall.

For $\lambda = 50$, displayed in Figure 34, we notice that 100% of the budget is always used, because at this level of rationality, many stations are likely to get attacked, and consequently, any new budget available is used to increase their defense.

We can see that for a low budget, s_1 is prioritized, reached by s_2 for $K_t^D \approx \text{€ } 8,000$ as the only two stations defended at 50/50 for this budget level. Then, additional budget goes to s_1 and the remaining medium-value stations ($s_4, s_6, s_7, s_8, s_9, s_{10}$).




Figure 34: Evolution of the Share of the Budget Spent per Strategy with K_t^D for $\lambda = 50$

For $\lambda = 125$, we can see in Figure 35 that there is a threshold, around $K_t^D \approx \text{€ } 17,000$, where the defender no longer uses the whole budget available. From this point, the budget is no longer a constraint: The allocation stays still for the stations in terms of budget (although it might vary in terms of the combination of strategies used).

When all the budget is used, we witness a development very similar to the one in the case $\lambda = 50$, with s_1 prioritized until we have sufficient budget to protect s_1 and s_2 with d_6 .




Figure 35: Evolution of the Share of the Budget Spent per Strategy with K_t^D for $\lambda = 125$

For $\lambda = 200$, displayed in Figure 36, the analysis is very similar. This time, the threshold occurs earlier, for $K_t^D \approx \text{€ } 12,000$, i.e. when s_1 and s_2 can be fully protected with d_6 . From this point, only a small budget is allocated to s_7 (for a small percentage of d_6), and nothing more: For a rational attacker, only three stations are worth defending, and only two fully.




Figure 36: Evolution of the Share of the Budget Spent per Strategy with K_t^D for $\lambda = 200$

In a nutshell, the defender's budget analysis leads to several interesting insights:

- The budget K_t^D is mostly a constraint for low rationality attackers. In high rationality, the probability of the attacker attacking high-value stations is so high (and therefore the probability of him attacking low-value stations is so low) that only a few stations are worth defending
- This hierarchy is also seen when the budget is a constraint: s_1 is protected first, followed by s_2 , and then the other stations.
- When looking in detail at the strategies and the stations, we notice that increasing the budget leads to an increase in the quality of the mixed strategy, as expected, but this quality increase isn't always linear (from d_1 to d_6). We have some cases of transition $(42\% d_1, 58\% d_4) \rightarrow (70\% d_3, 30\% d_5)$, with d_4 disappearing completely and d_3 appearing at 70%.

6.6 RESULTS: INFLUENCE OF THE TIME PERIOD CHOSEN

This analysis will be conducted on a full day (6 time slots), and will study the behavior of the attacker when it must split resources across time slots.

Since we want the budget to act as a constraint, we will use $\lambda = 50$ as the rationality parameter. The analysis will, once again, take place on the 1st of January 2023, and the data estimated by time slot are the one depicted in Table 7 and Table 8, in section 6.2.

In this analysis, instead of optimizing for a single time slot, we will do it for $\mathcal{T} = \{t_1 (6:00 - 7:00), t_2 (7:00 - 9:00), t_3 (9:00 - 13:00), t_4 (13:00 - 17:00), t_5 (17:00 - 21:00), t_6 (21:00 - 00:00)\}$. This will have an impact on the number of passengers affected (physically and for the delays), and on the budget constraint K_T^D . The daily budget will be higher ($K_T^D = € 118,000$), and can be distributed at any time slot.

In section 6.6.1, we will compare two cases:

- A joint formulation where the defender can allocate the full budget of € 118,000 the way he wants across the whole day.
- An independent formulation, where the defender has a fixed budget of € 16,900 for each time slot.

Then analyze, for the joint formulation, the budget allocation per pair (station, time slot) (section 6.6.2), the mixed strategies used throughout the day (section 6.6.3), the effect on the attacker's utility (section 6.6.4), and the parameters that have the biggest influence on the budget spent (section 6.6.5).

6.6.1 COMPARISON OF THE JOINT AND SLOT-INDEPENDENT FORMULATIONS

The two formulations produce objective values that are not directly comparable in their raw form, since they correspond to different levels of temporal aggregation:

- In the joint formulation, the optimization is carried out simultaneously over all six time slots of the day under a single daily budget constraint K_T^D . The reported objective r_{joint}^* (the minimum value of the objective function of the attacker defined in section 5.1, that we obtain with the binary search, for the case of the joint formulation), therefore represents the worst-case attacker utility over the entire day, evaluated under an optimally coordinated allocation of defensive resources across slots.
- In the slot-independent formulation, by contrast, the problem is decomposed into six separate subproblems, each solved under a per-slot budget $K_{slot}^D = \frac{K_T^D}{6}$ (since there are 6 time slots in a day, we allocate 1/6 of the daily budget K_T^D per time slot). This yields six slot-level worst-case utilities r_t^* , one per time slot t . To obtain a meaningful comparison, we exploit the fact that a rational attacker will select the single most favorable (segment, slot) pair over the day. Consequently, the day-level worst-case attacker utility implied by the slot-independent solution is:

$$r_{ind}^* = \max_{t \in \mathcal{T}} r_t^* \quad (40)$$

r_{ind}^* can be compared directly to r_{joint}^* . Indeed, because the joint formulation has strictly more flexibility, because it can shift budget toward the slots with the highest attacker exposure, whereas the independent formulation imposes a rigid equal split, we expect:

$$r_{joint}^* \leq r_{ind}^* \quad (41)$$

The difference between the two quantifies the value of temporal coordination of the defensive budget.

In our experiments, we obtain $r_{joint}^* = 0.056125$, while the slot-independent formulation produces a maximum slot-level utility of $r_{ind}^* = \max_t r_t^* = 0.07664$, attained at the most exposed slot (in this case t_2). The relative improvement obtained by coordinating the budget across slots is therefore

$$\frac{r_{ind}^* - r_{joint}^*}{r_{ind}^*} = \frac{0.07664 - 0.056125}{0.07664} \approx 26.8\%$$

In other words, allowing the defender to allocate the daily budget across time slots in a coordinated manner reduces the worst-case attacker utility by roughly 27% relative to a uniform per-slot allocation. This result confirms the intuition that risk is not evenly distributed throughout the day: some slots are intrinsically more vulnerable, and the joint formulation correctly reallocates resources toward those slots, thereby lowering the maximum exploitable utility available to the attacker.

The contrast between the joint and slot-independent formulations is further illustrated by looking at the resulting attacker utility at the station level, both on average across the day (Figure 37) and at its peak value across slots (Figure 38). These two perspectives are complementary and reveal the underlying mechanism through which the joint formulation achieves its overall advantage.




Figure 37: Average Attacker's Utility Comparison: Joint vs. Independent




Figure 38: Maximum Attacker's Utility Comparison: Joint vs. Independent

Figure 37 **Error! Reference source not found.** shows the average attacker utility per station after defensive resources have been allocated. For most stations, the average is slightly lower under the independent formulation than under the joint one. This may appear counter-intuitive at first, but it is a direct consequence of the way the joint model reallocates resources: by concentrating budget on the slots where the attacker's potential damage is highest, it deliberately under-protects stations and slots in which passenger flows, and therefore the attacker's payoff, are low. These low-risk situations contribute disproportionately to the average, pushing it upward in the joint case. The independent formulation, on the contrary, is forced to spend a fixed budget in every slot, which mechanically reduces the attacker's utility uniformly, including in slots where this protection has little strategic value.

Figure 38 shows the maximum attacker utility per station across the six time slots. Here the ordering is reversed: for every station, the maximum is at least as low under the joint formulation as under the independent one, and strictly lower for the stations where low harm level can be done ($s_4, s_6, s_7, s_8, s_{10}$). This is precisely the quantity that matters from a security standpoint, since the defender's ultimate objective is not to minimize the average exposure but

to prevent the worst-case scenario in which an attacker inflicts the maximum possible damage. By coordinating the budget across slots, the joint formulation is able to channel resources toward the highest-risk (station, slot) pairs and thereby flatten the peak of the attacker's utility distribution.

Taken together, these two figures clarify the trade-off between the two formulations. The independent model achieves a more uniform, and on average slightly lower, level of attacker utility, but at the cost of higher peaks. The joint model accepts a marginally higher average in exchange for substantially lower maxima, which is the relevant criterion in a worst-case-oriented security setting such as the one considered in this thesis.

6.6.2 BUDGET DISTRIBUTION AMONG THE STATIONS

Looking now only at the joint distribution, it is interesting to look at the use of the budget per time slot, since the defender is no longer constrained to use it on a single time slot. The result analysis is displayed on Figure 39.




Figure 39: Absolute (€) and Relative (%) Budget Spent per Time Slot

The total budget spent across the day shows a strongly non-uniform allocation, and reveals that the defender does not use her entire daily budget: the six bars sum to roughly 50% of the €118,000 available, meaning that a significant share of the budget is left unspent because the marginal protection it would provide is not worth its cost.

Among the time slots that do receive funding, the morning peak t_2 clearly dominates with 25.0% of the daily budget (well above the 16.7% that a uniform allocation across the six slots would give). This confirms that the defender concentrates her resources on the slot where passenger flows, and thus potential harm, are the highest. t_5 and t_4 receive moderate shares (10.3% and 8.6%), and the early morning t_1 gets a smaller but still significant 6.4%. In contrast, the midday slot t_3 and the evening slot t_6 receive no budget at all. In these low-flow periods, the expected harm from an attack is small enough that no defensive strategy passes the cost-

benefit threshold, and the defender prefers to leave the network entirely unprotected rather than spend on protection that would not pay off.

Below, in Figure 40, is presented a heatmap of the budget allocation per pair (station, time slot), in €.

On the heatmap, the darker the blue, the higher the budget spent utility (with the maximum budget for an (s, t) pair being the cost of the pure strategy d_6 : € 6,075).




Figure 40: Heatmap of the Budget Spent per Station and Time Slot

We can see that the repartition of the budget follows mainly two conclusions:

- The protection of s_1 and s_2 , which is consistent with the results obtained previously.
- The protection during the time slot t_2 , with 3 stations protected at their maximum (s_1 , s_2 and s_7) during this time slot.

Therefore, when considering the time slots, s_1 and s_2 are no longer the “must-protect” stations. It is for example more important to protect s_4 at t_2 than s_1 at t_3 .

It also shows that, even though the rationality level is low, the most efficient defense is highly concentrated: only 13 of the 60 possible pairs (station, time slot) are protected.

6.6.3 STRATEGIES USED

Regarding the strategy used, Figure 41 shows the aggregated strategy use over the day (i.e. the average strategy probability for each time slot). On this chart, d_1 having an average probability of 0.82 for t_1 means that, across all stations, the strategy d_1 has been used 82% of the time at t_1 .




Figure 41: Aggregated Strategy Use Over the Day

We can gather several insights from this figure:

- As expected with the highly concentrated allocation, d_1 is the strategy the most used for every time slot, except for t_2 , when it is tied with d_6 .
- t_2 is the only time slot where strategies different than d_1 and d_6 are used, with a small share of d_3 and d_4 . d_5 is absent, which is surprising, since it was used often in the previous studies, in sections 6.4 and 6.5.
- Figure 41 reveals a shift in the defender's strategies allocation. It seems that there are no longer medium-value pairs (s, t) to protect, but only pairs to fully protect with d_6 , or not protect at all with d_1 .

6.6.4 EFFECT ON ATTACKER'S UTILITY

Figure 42 displays the heatmap of the attacker's utility before and after the allocation. This allows us to understand better why the defender left some station unprotected, and focused only on a few. On the heatmap, the darker the blue, the higher the attacker utility (which is the damage that could be inflicted to the station).



Figure 42: Heatmap of the Attacker's Utility Before and After Allocation

In the initial heatmap (left), we see some areas where the attacker utility is much higher (in red), with a maximum utility for s_1 at t_2 , of 0.1898. These areas are the zones that the defender should protect in priority, and we can see that they correspond with the heatmap of the budget spent in Figure 40. Thanks to these allocations, the utility of these zones has been divided by two, which is what happens when the average performance score of the mixed strategy is 1.

On the other hand, the defender made the choice of not defending some pairs like (s_7, t_5) , where the utility is higher than (s_6, t_2) for example, even though the defender had budget available. This can be seen as a way to incentivize the attacker to attack stations other than, s_1 and s_2 , so that the harm can be reduced.

6.6.5 INFLUENCE OF OTHER PARAMETERS

In Figure 43 can be seen a correlation matrix of three station parameters (the number of passengers affected $N_{s,t}$, the betweenness centrality $B_{s,t}$, the value of assets A_s , and the symbolism level S_s), across all pairs (s, t) , that will help us understand how they influence the budget spent.




Figure 43: Correlation Matrix of Station Parameters

There are a few insights that we can gather from this matrix:

- All the parameters are positively correlated with the budget spend, with the less correlated being S_s with 0.45. This makes sense, since the harm function is strictly growing with the parameters (see equation (6)).
- The number of passengers $N_{s,t}$ is the parameter with the biggest influence on the budget spent, with a correlation of 0.82: the defender chooses in priority to defend stations where the passenger flow is high.
- Then comes the betweenness centrality, with 0.62, the asset value with 0.60 and finally the symbolic importance, with 0.45.

Consequently, in this model, the priorities of the defender are the following:

- First priority: minimize the number of people directly affected by the attack: the cost of death and injuries is too high with respect to the other costs.
- Second priority: Minimize more ‘economical’ costs, which are the asset values and the network delays.
- Third priority: Minimize the symbolic impact. In this model, the symbolic impact isn’t that much taken into consideration by the defender. This raises the question of whether we should add a coefficient to S_s in the harm formulation to increase its importance, depending on at which correlation level we would consider the model the most realistic.

6.7 COMPUTATIONAL ANALYSIS

6.7.1 EXPERIMENTAL SETUP

Before presenting the results of the computational study, this section details the hardware, software, and solver configuration under which all experiments were carried out. Reporting these parameters is essential to ensure the reproducibility of the results and to allow a fair interpretation of the reported solution times, since runtime measurements are inherently dependent on the underlying computational environment.

6.7.1.1 Hardware and Operating System

Component	Specification
Machine	HONOR MagicBook (personal laptop)
Operating System	Microsoft Windows (x86 64-bit)
CPU	11th Gen Intel(R) Core(TM) i7-11390H @ 3.40GHz (3.42 GHz)
RAM	16.0 GB

Table 11: Hardware and Operating System

6.7.1.2 Modeling Environment

Parameter	Value
Modeling language	GAMS
GAMS version	47.3.0

Platform	WEX-WEI (Windows 64-bit)
License	Small MUD – 5 User License (Universidad Pontificia Comillas, teaching and research)

Table 12: Modeling Environment

6.7.1.3 Solver Configuration

Parameter	Value
Solver	BARON 24.5.8 (Built May 8, 2024, WIN-64)
Problem type	MIQCP (Mixed-Integer Quadratically Constrained Program)
Relative optimality gap (optcr)	1e-4
Absolute optimality gap (optca)	0
Resource limit (reslim)	300
Iteration limit (iterlim)	200000
BARON print level (prlevel)	0 (silent log)
Threads	-1

Table 13: Solver Configuration

BARON was used with its default subsolver configuration as provided by GAMS 47.3.0, without a custom option file forcing a specific LP or NLP subsolver.

6.7.1.4 Algorithmic Parameters (Binary Search)

Parameter	Value
Convergence tolerance ε	1e-4
Maximum number of bisection iteration	50

Table 14: Algorithmic Parameters

6.7.1.5 Instance Parameters

Parameter	Value
Number of stations $ \mathcal{S} $	10
Number of strategies $ \mathcal{D} $	6
Number of linearization segments K	10
Attacker rationality λ	50
Scaling factor	1e9 ($\text{€} \rightarrow B\text{€}$)
Reproducibility seed	1234

Table 15: Instance Parameters

6.7.1.6 Study Design

The computational study consists of solving the same model for increasing numbers of active time slots, $NT \in \{3, 4, 5, \dots, 25\}$, with the daily budget scaling linearly as $K_T^D = 16\,900 \cdot$

NT €, in order to keep the relative defensive capability roughly constant across instances and avoid the problem becoming trivially infeasible or trivially easy as NT grows.

Since the purpose of this study is to characterize the computational behavior of the model (i.e., how solving time scales with problem size) and not to interpret the optimal protection strategy, the two parameters evolving with the time slot are generated randomly:

- The passenger flow $N_{s,t}$ drawn from a uniform distribution $U(0, 1000)$.
- The betweenness centrality $B_{s,t}$ drawn from a uniform distribution $U(0, 1)$.

A fixed random seed (OPTION seed = 1234) is used so that the generated values are identical across runs, ensuring full reproducibility. This randomization approach is justified because:

- The runtime of an MIQCP solver depends primarily on the structure and size of the problem (number of variables, constraints, binaries, density of the constraint matrix), and only secondarily on the specific numerical values of the coefficients.
- Using random inputs avoids any bias that could arise from using real Madrid metro data tuned to particular operational patterns, which might artificially favor or penalize certain instance sizes.
- The objective value “obj” reported in the results table is therefore not meaningful in absolute terms. It is reported only as a sanity check to confirm that the solver did converge to a feasible optimum in each instance.

The other parameters (asset values A_s , symbolic levels S_s , strategy costs C_d , performance scores PS_d) are kept fixed at their nominal values from the original Madrid metro dataset presented in 6.2, since they define the structural dimensions of the problem (number of stations, number of strategies) rather than instance-specific data.

6.7.2 COMPUTATIONAL RESULTS

The results of the analysis are displayed below in Figure 44:




Figure 44: Evolution of the Computational Time (t) with the Number of Time Slots (n)

The results show that the computational time grows approximately as a power law of the number of time slots, with a fitted exponent of 2.17 ($t \approx 2.41 \cdot n^{2.17}$). This polynomial (and roughly quadratic) growth confirms that the proposed model scales well with the temporal granularity of the planning horizon and remains tractable on standard hardware over the full range studied ($3 \leq n \leq 25$), with solve times below one hour even in the worst case.

The fact that the curve oscillates so much around the trend line is not surprising: each instance is built from a different random draw of betweenness and passenger values, and some of these draws simply happen to be harder to solve than others. This is a common behavior for this type of optimization problem, and it does not call into question the overall trend we observe. From a practical point of view, these results are reassuring: in a real-world security planning context, the number of time slots would typically correspond to a handful of operational periods (for instance, morning peak, off-peak, evening peak, and night service) which falls well within the range where the model solves comfortably.

It is also notable that the optimal objective value remains stable across instance sizes, confirming that the variability in solving time is purely computational and not driven by qualitatively different solutions.

7. CONTRIBUTIONS, CONCLUSIONS, AND FUTURE LINES OF RESEARCH

7.1 *CONTRIBUTIONS*

7.1.1 THEORETICAL CONTRIBUTIONS

- The first contribution of this thesis is the development of a framework that integrates **time-dependence, defensive strategies, and network topology** within a single formulation. Most Stackelberg Security Games in the literature treat targets as independent and static entities, focusing on a single dimension of the problem at a time. In contrast, the model proposed here jointly accounts for the temporal evolution of passenger flows across operational periods, a portfolio of qualitatively different defensive measures (patrols, K-9 units, CCTV, access control, and explosive detection scanners) with distinct costs and effectiveness scores, and the topological role of each station captured through its betweenness centrality. To the best of our knowledge, no prior work combines these three dimensions in a single solvable optimization model.
- The thesis also contributes a **bounded-rationality attacker model** based on Quantal Response that remains computationally manageable. The attacker's behavior is described by a logit-type response function, capturing realistic deviations from perfect rationality. The resulting fractional optimization problem is reformulated through a binary search bisection on the fractional ratio, and the remaining non-linear terms are handled by a piecewise linear approximation. The combination of these two techniques yields a model that can be solved by standard MIQCP solvers, without resorting to specialized algorithms.

7.1.2 PRACTICAL CONTRIBUTIONS

- From a practical standpoint, the thesis delivers a reproducible and fully documented **case study on the Madrid Cercanías network**. Every input of the model (the hourly defender

budget reconstructed top-down from public Renfe and Adif security expenditures, the bimodal passenger flow profile calibrated on open Renfe data, and the clustering of stations into operational zones) is derived from publicly available sources and explicitly justified. This transparency ensures that the case study can be reproduced, discussed, and updated as new data become available.

- The thesis also contributes to a **computational study** demonstrating that the model scales polynomially, with solve times growing approximately as the square of the number of time slots and remaining **below one hour on standard hardware** for all instance sizes considered. This confirms the practical applicability of the framework for realistic operational horizons. The complete GAMS implementation is modular and can be readily extended to other transport networks or critical infrastructures, providing a reusable basis for future work.

7.2 CONCLUSIONS

7.2.1 RESULTS OF THE CASE STUDIES

7.2.1.1 Strong Sensitivity to Rationality

- The case studies conducted on the Madrid Cercanías network led to several interesting observations. The most striking is the **strong sensitivity** of the optimal defensive allocation to the attacker's **rationality parameter λ** . When the attacker is modeled as nearly uniform in its choice of target, the defender spreads the budget across many stations using a broad mix of low-cost strategies. As λ increases and the attacker becomes more rational, the allocation gradually concentrates on the highest-value stations, until a clear threshold is reached beyond which only the two central hubs are protected and part of the budget is deliberately left unused. **This non-monotonic and threshold-like behavior** is one of the central qualitative findings of the thesis.

- The analysis identifies a **rationality threshold**, located around $\lambda \approx 120$ in the Cercanías case study, beyond which the nature of the optimal defensive strategy changes qualitatively. Below this threshold, the defender spreads resources across many stations through a smooth mix of low-cost strategies; above it, the allocation concentrates abruptly on a small number of high-value targets, and part of the budget is deliberately left unused. This characterization provides a novel qualitative insight into how attacker sophistication reshapes optimal defensive allocation.

7.2.1.2 Most Protected Stations and Critical Time Slots

The results consistently show that **defensive effort concentrates on high-passenger-flow stations during peak hours**. Because the attacker's expected harm is weighted by passenger flow, the stations belonging to the major interchange hubs receive the strongest and most sophisticated protection (strategies d_5 : Active CCTV Monitoring + Foot Patrol + Random Screening + Canine Unit and d_6 : Active CCTV Monitoring + Foot Patrol + Random Screening + Canine Unit + Scanner), while smaller feeder stations are protected only moderately, or are deliberately left unprotected when the budget is tight.

The bimodal passenger profile directly maps onto the criticality of time slots: the morning and evening commuting peaks emerge as the most critical time slots, absorbing the largest share of the budget. This confirms that **the model correctly identifies the moments of highest exposure rather than spreading resources uniformly across the day**.

7.2.1.3 Expected Behaviors Confirming Model Coherence

Several observations were expected a priori and confirm that the model behaves consistently:

- **Monotonic budget effect:** the attacker's expected utility decreases as the defender's budget grows, confirming the model rewards additional resources coherently.

- Strategy substitution with budget: **at low rationality, as budget increases, the defender first extends coverage** (defending more worthwhile stations with intermediate strategies such as d_4 : Active CCTV Monitoring + Foot Patrol + Random Screening), and then upgrades quality (replacing d_4 by d_5 and d_6 on the stations that matter most). This two-phase "coverage-then-quality" pattern is intuitive and consistent.
- Effect of rationality λ : **higher λ (a more rational attacker) pushes the defender to concentrate protection on the few highest-value targets**, since a rational attacker will always strike the weakest valuable point.

7.2.1.4 Surprising and Novel Findings

Some results were less obvious and constitute the more original insights of this work:

- Deliberately unspent budget: **beyond a certain rationality threshold, the defender stops using the full budget and focuses exclusively on a small number of high-value targets**, accepting that part of the budget remains unused. This is a counter-intuitive but rational result: protecting low-value stations would not meaningfully reduce a sophisticated attacker's expected harm, so the marginal resources are simply not worth deploying.
- Non-monotonic use of intermediate strategies: the use of strategy d_4 is concave in the budget and reaches a maximum that depends on both the station and λ . This shows that **intermediate-cost strategies act as transitional defenses, appearing and then disappearing as the budget grows**.

7.2.1.5 Value of Temporal Coordination

- The comparison between joint and independent budget formulations shows that the **joint formulation consistently outperforms the independent one in worst-case terms**. By reallocating resources toward peak-flow periods, the joint model lowers the maximum

attacker utility across station-slot pairs at the cost of a slightly higher average. This confirms that the temporal coordination of the budget is a genuine source of security gain rather than a mere modeling convenience, and that it matters most precisely in the kind of worst-case-oriented setting that motivates this work.

- Overall, the results highlight the **importance of explicitly modeling the temporal dimension**. Ignoring time, as most existing models do, would systematically lead to **under-protection during peak hours and over-protection during off-peak ones**. The bimodal commuting pattern of Madrid, with morning and evening peaks between 07:00–09:00 and 18:00–20:00, drives most of the allocation and is fully consistent with operational intuition. In a similar spirit, the two central stations of the Cercanías network emerge as the most valuable targets across virtually all scenarios explored, reflecting the combined effect of high passenger throughput and high betweenness centrality. This result is in line with the historical fact that the 2004 Madrid attacks targeted trains converging on Atocha, which provides additional empirical support for the model.

7.2.2 BROADER CONCLUSIONS

- Beyond the specific numerical results, the work carried out in this thesis suggests a few broader lessons about how the protection of critical infrastructures could be approached in practice. The most counter-intuitive of these is probably that **trying to defend everything is not a good idea**. When the attacker is sufficiently rational, the best strategy is not to spread the budget thinly across the whole network, but rather to concentrate it on a small number of high-value targets, **even if this means leaving part of the budget unused**. This goes against the common political and managerial reflex of maximizing visible security spending, and it pleads in favor of a **more targeted, risk-based logic of allocation**.
- A second important lesson is that **the way the attacker is modeled changes the nature of the problem in a fundamental way**. Assuming a fully rational attacker, or on the contrary a completely random one, leads to very different allocations. The Quantal Response

formulation used here lies between these two extremes and is, in our view, the most realistic of the three. The fact that a clear rationality threshold appears, beyond which the optimal allocation suddenly shifts, also shows that small changes in the assumed level of sophistication of the attacker can completely change the recommended strategy. This has a direct practical implication: **the way intelligence services communicate their threat assessments to security planners matters, because even a moderate revision of the perceived attacker rationality can call for a very different defensive posture.**

- The case studies also make clear that **network topology and passenger flows have to be considered together**. A station may be valuable because many passengers go through it, because it occupies a strategic position in the network, or for both reasons at the same time. The model shows that these two dimensions sometimes reinforce each other and sometimes compensate each other, and that allocations based on passenger counts alone, which is still the most common practice in operational settings, are sub-optimal in many realistic configurations.
- It is also worth keeping in mind that a game-theoretic model of this kind is a **decision-support tool, not an oracle**. The numerical values it produces depend on inputs such as the effectiveness of each strategy, its cost, or the assumed rationality of the attacker, all of which are partly subjective and inherently uncertain. **The main contribution of this thesis therefore does not lie so much in the specific allocations obtained, but rather in the structural insights it provides: which trade-offs matter, how they evolve with the parameters, and what kind of strategies remain robust across scenarios.**
- Finally, this work also illustrates that doing serious quantitative research on critical infrastructure protection in an open and reproducible way is both possible and desirable. By relying exclusively on public data and documenting every modeling choice, the thesis tries to show that rigorous analysis of sensitive topics can be carried out transparently, without compromising operational security and while still producing results that are useful to practitioners.

7.3 *FUTURE LINES OF RESEARCH*

Like any modeling work, this thesis had to rely on a number of simplifying assumptions, and several of them could be relaxed in future research. Some of these improvements concern the structure of the model and its mathematical formulation, while others are more applied and relate to the way the case study was built and analyzed. The two subsections below describe these perspectives in turn.

7.3.1 REGARDING THE MODEL

- A first natural extension of the model would be to make the cost and performance of each defensive strategy depend on the station and on the time slot in which it is deployed. In the current formulation, a strategy has a single cost and a single effectiveness score, regardless of where and when it is applied. In reality, implementing a given measure is not equally expensive in every station: a large multimodal hub with several access points requires more personnel and equipment than a small suburban station, and a zone that contains many platforms is more costly to cover than a compact one. The same logic applies to time, since night shifts and weekends usually involve higher labor costs than regular daytime hours. Incorporating this spatial and temporal dependence would make the model more realistic and would probably reinforce the asymmetries already observed between central and peripheral stations.
- A second direction concerns the theoretical foundations of the resolution method. The bisection procedure on the ratio variable r converges very reliably in all the instances tested, but a formal mathematical proof of its convergence toward the optimal solution has not been established in this thesis. Providing such a proof, possibly by relating the procedure to the broader family of Dinkelbach-type algorithms for fractional programming, would strengthen the methodological contribution of the work and clarify the conditions under which the approach is guaranteed to succeed.

-
- A third extension would consist in adding a further step to the game in order to make it truly evolutive. In the current setup, the defender allocates resources and the attacker responds according to a Quantal Response distribution, but the interaction stops there. A more realistic model would let the attacker observe the defender's allocation and adapt by shifting attention toward undefended or lightly defended stations, which could in turn prompt the defender to reallocate resources, and so on. Studying the dynamics and the equilibria of such an iterated game would bring the framework closer to the spirit of evolutionary game theory that motivates the title of this thesis.
 - Another promising future line of research would be to incorporate historical information from past incidents through artificial intelligence techniques. Such information need not be limited to terrorist attacks alone: it could also include other types of threats, such as cyberattacks, acts of sabotage, or hybrid threats that combine several attack vectors. The goal would be to enable the model to learn from patterns observed in the past (types of targets, timing, methods that have historically been favored), and to integrate this accumulated experience directly into the dynamics of defensive strategy selection. In the current formulation, the attacker's behavior is governed solely by the Quantal Response model and the parameters estimated for the network, without any feedback from real-world events. By coupling this game-theoretic framework with a learning component, the defender's resource allocation could become more adaptive and more sensitive to the evolving threat landscape. Over time, the model would not only react to the structural importance of each station, but also adjust its priorities according to the historical frequency and nature of the threats observed, leading to a defense strategy that improves as new data becomes available.
 - Finally, the way network perturbations are handled could be improved by propagating them across time. At the moment, the cost of network delays is computed within a single time slot, but a disruption occurring during the morning peak is likely to affect passenger flows and operations for several hours afterwards, possibly throughout the entire day or even the rest of the week. Modeling this temporal propagation would give a more accurate picture

of the true economic and social cost of an attack, and would probably increase the relative value of stations whose disruption has long-lasting consequences on the network.

7.3.2 REGARDING THE CASE STUDY

- On the case study side, a first improvement would be to refine the passenger flow model so that it better reflects the diversity of operational situations encountered in practice. The current distribution is assumed to be constant across days, but it is well known that bank holidays, large events, school vacations, and weekends produce significantly different patterns, often with more people concentrated in the central stations and fewer commuters from the periphery. Incorporating these variations would allow the model to be tested on a wider range of realistic scenarios and would probably reveal that the optimal allocation is itself time-varying at a longer scale than the one currently considered.
- A second avenue would be to scale up the network and to study how the size and structure of the system influence the results. The case study presented here focuses on ten representative stations of the Madrid Cercanías network, which is already enough to highlight the main mechanisms at play, but it would be valuable to see how the model behaves when applied to the full network, or to networks of different cities with different topologies. Such an analysis would clarify whether the insights obtained here generalize, and it would also give a better idea of the computational limits of the approach.
- A third and closely related direction would be to investigate more systematically the differences between central and peripheral stations. In all the experiments conducted, the two central stations turned out to be the most valuable and the first to be protected, which is intuitive but also somewhat dependent on the structure of the chosen subnetwork. Looking at a larger and more diverse set of stations would make it possible to characterize more precisely the role of centrality, to identify situations in which peripheral stations become competitive targets, and to understand under which conditions the hierarchy observed in this thesis still holds.

8. APPENDIX: SOCIAL IMPACT

The main Sustainable Development Goals (SDGs) addressed by this project are the goals nº16, 9 and 11: *Peace, Justice and Strong Institutions, Industry, Innovation and Infrastructure, and Sustainable Cities and Communities*.

This alignment is further emphasized by the target indicators that directly reflect the purpose of this work, in Figure 45: Alignment with SDGs.

Target	Indicator	Alignment of this work	
16.1	Significantly reduce all forms of violence and related death rates everywhere	This work aims to reduce terrorist attacks, lowering thereby violence and death toll	
16.3	Promote the rule of law at the national and international levels	This work aims to help law enforcement in their duty	
9.1	Develop quality, reliable, sustainable and resilient infrastructure	This work contributes to resilience by helping railways operators mitigating risks	
9.4	Upgrade infrastructure to make it more sustainable and resource-efficient	This work promotes efficient use of limited security and maintenance resources	
11.2	Provide access to safe, affordable, accessible and sustainable transport systems	Protect railways from attacks ensures continuous and safe mobility	
11.5	Reduce the number of people affected by disasters	This work aims to reduce terrorist attacks, which are human-made disasters	

Figure 45: Alignment with SDGs

References

- [1] C.-H. Lin, J. Arcos-Pumarola, and N. Llonch Molina, "Tourism safety on train systems: A case study on electronic word-of-mouth in Spain, Italy and Greece," *Secur. J.*, vol. 37, pp. 1033–1059, Dec. 2023, doi: 10.1057/s41284-023-00405-1.
- [2] A. Rezazadeh, L. Talarico, G. Reniers, V. Cozzani, and L. Zhang, "Applying game theory for securing oil and gas pipelines against terrorism," *Reliab. Eng. Syst. Saf.*, vol. 191, p. 106140, Nov. 2019, doi: 10.1016/j.ress.2018.04.021.
- [3] S. Lakshminarayana, E. V. Belmega, and H. V. Poor, "Moving-Target Defense Against Cyber-Physical Attacks in Power Grids via Game Theory," *IEEE Trans. Smart Grid*, vol. 12, no. 6, pp. 5244–5257, Nov. 2021, doi: 10.1109/TSG.2021.3095083.
- [4] A. Sanjab and W. Saad, "On bounded rationality in cyber-physical systems security: Game-theoretic analysis with application to smart grid protection," in *2016 Joint Workshop on Cyber-Physical Security and Resilience in Smart Grids (CPSR-SG)*, Apr. 2016, pp. 1–6. doi: 10.1109/CPSRSG.2016.7684101.
- [5] F. M. D. Fave *et al.*, "Game-Theoretic Patrolling with Dynamic Execution Uncertainty and a Case Study on a Real Transit System," *J. Artif. Intell. Res.*, vol. 50, pp. 321–367, Jun. 2014, doi: 10.1613/jair.4317.
- [6] K. Hunt and J. Zhuang, "A review of attacker-defender games: Current state and paths forward," *Eur. J. Oper. Res.*, vol. 313, no. 2, pp. 401–417, Mar. 2024, doi: 10.1016/j.ejor.2023.04.009.
- [7] 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*, Stockholm, Sweden: International Joint Conferences on Artificial Intelligence Organization, Jul. 2018, pp. 5494–5501. doi: 10.24963/ijcai.2018/775.
- [8] A. Yolmeh and M. Baykal-Gürsoy, "Urban rail patrolling: a game theoretic approach," *J. Transp. Secur.*, vol. 11, no. 1–2, pp. 23–40, Jun. 2018, doi: 10.1007/s12198-018-0187-z.
- [9] F. Fang *et al.*, "Deploying PAWS: Field Optimization of the Protection Assistant for Wildlife Security," *Proc. AAAI Conf. Artif. Intell.*, vol. 30, no. 2, pp. 3966–3973, Feb. 2016, doi: 10.1609/aaai.v30i2.19070.
- [10] G. Yang, R. Poovendran, and J. P. Hespanha, "Adaptive Learning in Two-Player Stackelberg Games with Application to Network Security," Jan. 08, 2021, *arXiv*: arXiv:2101.03253. doi: 10.48550/arXiv.2101.03253.
- [11] B. Zou, Y. Wang, C. Liu, M. Dai, Q. Du, and X. Zhu, "Generation of security system defense strategies based on evolutionary game theory," *Nucl. Eng. Technol.*, vol. 56, no. 9, pp. 3463–3471, Sep. 2024, doi: 10.1016/j.net.2024.03.043.
- [12] R. Yang, A. X. Jiang, M. Tambe, and F. Ordo, "Scaling-Up Security Games with Boundedly Rational Adversaries: A Cutting-Plane Approach".
- [13] K. Liu, "Lemke and Howson Algorithm".
- [14] H. Jin *et al.*, "Evolutionary game decision-making method for network attack and defense based on regret minimization algorithm," *J. King Saud Univ. - Comput. Inf. Sci.*, vol. 35, no. 3, pp. 292–302, Mar. 2023, doi: 10.1016/j.jksuci.2023.01.018.

-
- [15] Bier, Vicky, Oliveros, Santiago, Samuelson, and Larry, “Choosing What to Protect: Strategic Defensive Allocation against an Unknown Attacker,” *J. Public Econ. Theory*, vol. 9, pp. 563–587, Mar. 2007, doi: 10.1111/j.1467-9779.2007.00320.x.
 - [16] R. Yang, F. Ordonez, and M. Tambe, “Computing Optimal Strategy against Quantal Response in Security Games”.
 - [17] R. D. McKelvey and T. R. Palfrey, “Quantal Response Equilibria for Normal Form Games,” *Games Econ. Behav.*, vol. 10, no. 1, pp. 6–38, Jul. 1995, doi: 10.1006/game.1995.1023.
 - [18] A. Balakrishnan and S. C. Graves, “A composite algorithm for a concave-cost network flow problem,” *Networks*, 1989, doi: 10.1002/net.3230190202.
 - [19] K. Croxton, B. Gendron, and T. Magnanti, “A Comparison of Mixed-Integer Programming Models for Non-Convex Piecewise Linear Cost Minimization Problems,” *Manag. Sci.*, vol. 49, pp. 1268–1273, Sep. 2003, doi: 10.1287/mnsc.49.9.1268.16570.
 - [20] J. P. Vielma, S. Ahmed, and G. Nemhauser, “Mixed-Integer Models for Nonseparable Piecewise-Linear Optimization: Unifying Framework and Extensions,” *Oper. Res.*, vol. 58, no. 2, pp. 303–315, Apr. 2010, doi: 10.1287/opre.1090.0721.
 - [21] A. Charnes and W. W. Cooper, “Programming with linear fractional functionals,” *Nav. Res. Logist. Q.*, vol. 9, no. 3–4, pp. 181–186, 1962, doi:10.1002/nav.3800090303.
 - [22] “Conjuntos de datos - Renfe Data.” Accessed: May 31, 2026. [Online]. Available: <https://data.renfe.com/dataset>
 - [23] “Datos anuales - Ayuntamiento de Madrid.” Accessed: Jun. 01, 2026. [Online]. Available: <https://www.madrid.es/portales/munimadrid/es/Inicio/El-Ayuntamiento/Estadistica/Areas-de-informacion-estadistica/Mercado-de-trabajo/Afiliaciones-a-la-Seguridad-Social/Datos-anuales/?vgnextfmt=default&vgnextoid=9c7d29bc75a80510VgnVCM1000000b205a0aRCRD&vgnnextchannel=f26a62a006986210VgnVCM2000000c205a0aRCRD>
 - [24] C. R. de T. de Madrid, “Consorcio Regional de Transportes de Madrid - CRTM Inicio.” Accessed: Jun. 01, 2026. [Online]. Available: <https://www.crtm.es/conocenos/enlaces-y-documentacion/edm2018/>
 - [25] “Adif invierte 7,5 millones de euros en realizar mejoras en la estación de Sol,” Adif. Accessed: Jun. 01, 2026. [Online]. Available: <https://www.adif.es/w/adif-invierte-7-5-millones-de-euros-en-realizar-mejoras-en-la-estaci%C3%B3n-de-sol>
 - [26] “Adif invierte 2,72 M€ para mejorar las estaciones de Parla y de Embajadores,” Adif. Accessed: Jun. 01, 2026. [Online]. Available: <https://www.adif.es/w/adif-invierte-2-72-m%E2%82%AC-mejorar-estaciones-parla-embajadores>
 - [27] “https://www.adif.es/documents/20124/1692783/Plan_Integral_Madrid.pdf.” Accessed: Jun. 01, 2026. [Online]. Available: https://www.adif.es/documents/20124/1692783/Plan_Integral_Madrid.pdf
 - [28] F. Reinares, “Lo que el 11-M nos sigue enseñando sobre la incesante amenaza del terrorismo yihadista en Occidente”.
 - [29] “New report: major developments and trends on terrorism in Europe in 2024,” Europol. Accessed: Jun. 01, 2026. [Online]. Available: <https://www.europol.europa.eu/media-press/newsroom/news/new-report-major-developments-and-trends-terrorism-in-europe-in-2024>
 - [30] D. interior Ministerio, “2 millones de euros es el valor estimado por evitar o prevenir un fallecimiento en siniestro de tráfico.” Accessed: Jun. 02, 2026. [Online]. Available:

- <https://www.dgt.es/comunicacion/notas-de-prensa/2-millones-de-euros-es-el-valor-estimado-por-evitar-o-prevenir-un-fallecimiento-en-siniestro-de-trafico/>
- [31] “Tablas_indemnizatorias_Baremo_2024.pdf.” Accessed: Jun. 03, 2026. [Online]. Available: https://dgsfp.mineco.gob.es/es/Regulacion/DocumentosRegulacion/Tablas_indemnizatorias_Baremo_2024.pdf
- [32] Directorate-General for Mobility and Transport (European Commission) *et al.*, *Handbook on the external costs of transport: version 2019 – 1.1*. Publications Office of the European Union, 2020. Accessed: Jun. 03, 2026. [Online]. Available: <https://data.europa.eu/doi/10.2832/51388>
- [33] “2024_InfAudit+CCAA+InfGestion+EINF_CONSO_FIRMADAS.pdf.” Accessed: Jun. 03, 2026. [Online]. Available: https://www.renfe.com/content/dam/renfe/es/Grupo-Empresa/Gobierno-corporativo-y-transparencia/cuentas-anuales/grupo-renfe/2024_InfAudit+CCAA+InfGestion+EINF_CONSO_FIRMADAS.pdf
- [34] “Renfe Operadora.” Accessed: Jun. 03, 2026. [Online]. Available: <https://grupo.renfe.com/es/es/gobierno-y-transparencia/informacion-economica/renfe-operadora>