

Game-Theoretic USV Patrolling for Subsea Cable Protection: A Bayesian Stackelberg Approach



Master's Thesis Report

RISK4-1



AALBORG UNIVERSITY
STUDENT REPORT

Aalborg University
Risk and Safety Management



Risk and Safety Management
Aalborg University
<http://www.aau.dk>

AALBORG UNIVERSITY

STUDENT REPORT

Title:

Game-Theoretic USV Patrolling for Subsea Cable Protection: A Bayesian Stackelberg Approach

Theme:

Master's Thesis

Project Period:

Spring Semester 2026

Project Group:

RISK4-1-S26

Participant(s):

Henri Lindroos
Mykolas Poska

Supervisor(s):

Christian D. Jørgensen
Malte von Benzon

Page Numbers: 96**Date of Completion:**

May 31, 2026

Abstract:

This thesis develops a Bayesian Stackelberg Security Game (BSSG) framework for optimising the patrol of a single Unmanned Surface Vehicle (USV) protecting subsea infrastructure in the Bornholm Basin. The DOBSS algorithm is applied to a 400-node grid incorporating real geospatial data on cable routes, seabed substrate, and vessel traffic, modelling a heterogeneous population of accidental and state-sponsored threats; the resulting coverage distribution is translated into a kinematically feasible patrol via maximum-entropy Markov routing. The central contribution is a tractable decision-support structure carrying a Stackelberg optimality guarantee given the model. A seven-dimension sensitivity analysis yields a reliability hierarchy mapped onto the underlying Strength-of-Knowledge ratings: the patrol *target* is unconditionally robust, while the specific *allocation* is conditional on two uncalibrated parameters. The equilibrium strategy improves expected defender utility by 8–9% over baselines within this case study, a gain driven by a single dominant asset and smaller than the spread from parameter uncertainty. The framework is a proof-of-concept and is not field-validated.

Contents

1	Introduction	2
1.1	Background	2
1.2	Core Problem	4
2	Problem Analysis	5
2.1	Risk and Uncertainty Foundations	5
2.2	Baltic Subsea Threat Environment	7
2.2.1	Geostrategic Significance of Baltic Maritime Corridors	7
2.2.2	Systemic Vulnerabilities and Supply-Chain Dependencies	7
2.2.3	Escalating Pattern of Subsea Sabotage in the Baltic Sea	9
2.2.4	The Shadow Fleet as a Dual-Use Threat Vector	11
2.2.5	Environmental Conditions as Threat Amplifiers	15
2.3	Autonomous Systems and Patrol Optimization Approaches	16
2.3.1	Operational Rationale	16
2.3.2	Demonstrated Capability in the Baltic Sea	16
2.3.3	Danish Industrial and Research Ecosystem	17
2.3.4	Platform Architecture and the Role of USVs	18
2.4	Existing Approaches and Methodological Gap	19
2.5	Analytical Synthesis	20
3	Problem Formulation	22
3.1	Conclusion of Problem Analysis	22
3.2	Research Question	23
3.3	Methodological Approach	23
3.4	Delimitations and Assumptions	24
3.5	Scientific Contribution	24
4	Problem Solution	25
4.1	Methodology	25
4.2	Data and Spatial Context	26
4.2.1	Location	26

4.2.2	Subsea Infrastructure	27
4.2.3	Sea Substrate	32
4.2.4	Marine Traffic	33
4.3	DOBSS Formulation and Implementation	34
4.3.1	Theoretical Foundations	35
4.4	DOBSS Implementation	37
4.4.1	Discretised Grid \mathcal{T}	37
4.4.2	Environmental Constraints and Attacker Cost	37
4.4.3	Follower Types Θ and Bayesian Threat Priors $p(\theta)$	39
4.4.4	Target Utilities U_a, U_d	39
4.4.5	Optimisation Constant M	43
4.4.6	Decision Variables x_t, q_t^θ and Auxiliary Variables z_t^θ, v^θ	43
4.5	Spatiotemporal Routing via Markov Chain Synthesis	46
4.6	Results	49
4.6.1	Adversary Target Selection	49
4.6.2	Optimal Stationary Distribution x^*	50
4.6.3	60-Hour Patrol Simulation	52
4.6.4	Results and Spatial Visualization	52
4.7	Sensitivity Analysis	53
4.7.1	Strategy Comparison: DOBSS versus Proportional and Uniform Baselines	54
4.7.2	Rationality Sensitivity	56
4.7.3	Prior Sensitivity	58
4.7.4	Monte Carlo Parameter Sensitivity and PRCC	59
4.7.5	One-at-a-Time Parameter Sensitivity	61
4.7.6	Risk-Score Weighting Sensitivity	64
4.7.7	Operational Constraint Sensitivity	66
4.7.8	Maximum-Entropy Markov Routing	67
4.7.9	Sensitivity Analysis Summary	69
5	Discussion	71
5.1	Return to the Research Questions	71
5.1.1	Sub-Question 1: Modelling a Heterogeneous Adversary Population	71
5.1.2	Sub-Question 2: Translating Coverage Probabilities into Feasible Routes	72
5.1.3	Sub-Question 3: Improvement over Baselines and Robustness Across Rationality Levels	73
5.2	The Robustness Hierarchy	73
5.3	Limitations	75
5.3.1	The One-Shot Game and Adversarial Learning	75
5.3.2	Binary Coverage and Sensor Physics	75
5.3.3	Low-Prior High-Capability Adversary (Black-Swan) Vulnerability	76

5.3.4	Static Strategy and Seasonal Dynamics	76
5.3.5	EEZ Boundary Effects and Adversary Risk Geography	77
5.3.6	Weaknesses in the Adversary Utility Function Formulations	77
5.4	Future Work	77
6	Conclusion	80
	Bibliography	82
A	Risk and Uncertainty Framework Definitions	i
B	Strength of Knowledge Assessment of the DOBSS Model and its Assumptions	iii
C	Stakeholder Analysis	v
D	Regulatory Context	viii
D.1	UNCLOS and Jurisdictional Challenges	viii
D.2	EU Binding Directives: NIS2 and CER	viii
D.3	ICPC Recommendations on Cable Protection Zones	ix
E	Adversary Typology	x
F	Implementation Code	xii

Preface

Aalborg University May 31, 2026

The authors express their sincere gratitude to their supervisors for their steadfast guidance and dedication. We are deeply grateful to our supervisor Christian D. Jørgensen for his exceptional availability, irrespective of the day of the week, and his comprehensive, insightful correspondence, which provided vital support throughout this academically challenging process. We also extend our appreciation to our second supervisor, Assistant Professor Malte von Benzon, whose domain expertise established the critical scope and boundaries required for a project constrained by time and physical realities. Furthermore, we acknowledge Mads Valentin Bram, whose mentorship initially directed us toward our supervisory team.

Our appreciation also extends to the departmental staff: Carsten, the department librarian; Mette, secretary of the Department of Energy; and Jeppe, the laboratory technician, whose technical expertise was instrumental in constructing the physical miniature model of the game. Finally, we thank Aalborg University and our peers for the exceptional education provided over the past two years within the Master of Science programme in Risk and Safety Management, an experience that has been profoundly transformative.

In accordance with the established guidelines for the use of Artificial Intelligence (AI) in this project, AI has been employed in the preparation of this report to refine and improve text originally written by the students. In addition, AI has been utilised to provide support in the development and implementation of code.

Ad nye veje

Henri Lindroos
<hlindr24@student.aau.dk>

Mykolas Poska
<mposka18@student.aau.dk>

Chapter 1

Introduction

1.1 Background

Submarine fibre-optic cables carry more than 95% of international internet and telecommunications traffic, underpinning nearly all cross-border data exchange, financial transactions, and digital communications [1, 2, 3, 4]. Alongside them, subsea power interconnectors transfer electricity between mainland grids, offshore generation facilities, and island communities, forming an indispensable component of regional energy security [5].

In Northern Europe, the Baltic Sea functions as a hub for essential data and energy corridors [6, 7]. Securing this Critical Maritime Infrastructure (CMI) is a key priority, however, attaining a thorough situational picture of the subsurface domain involves specific economic and technical challenges that differ markedly from those encountered in terrestrial or aerospace settings [8, 9].

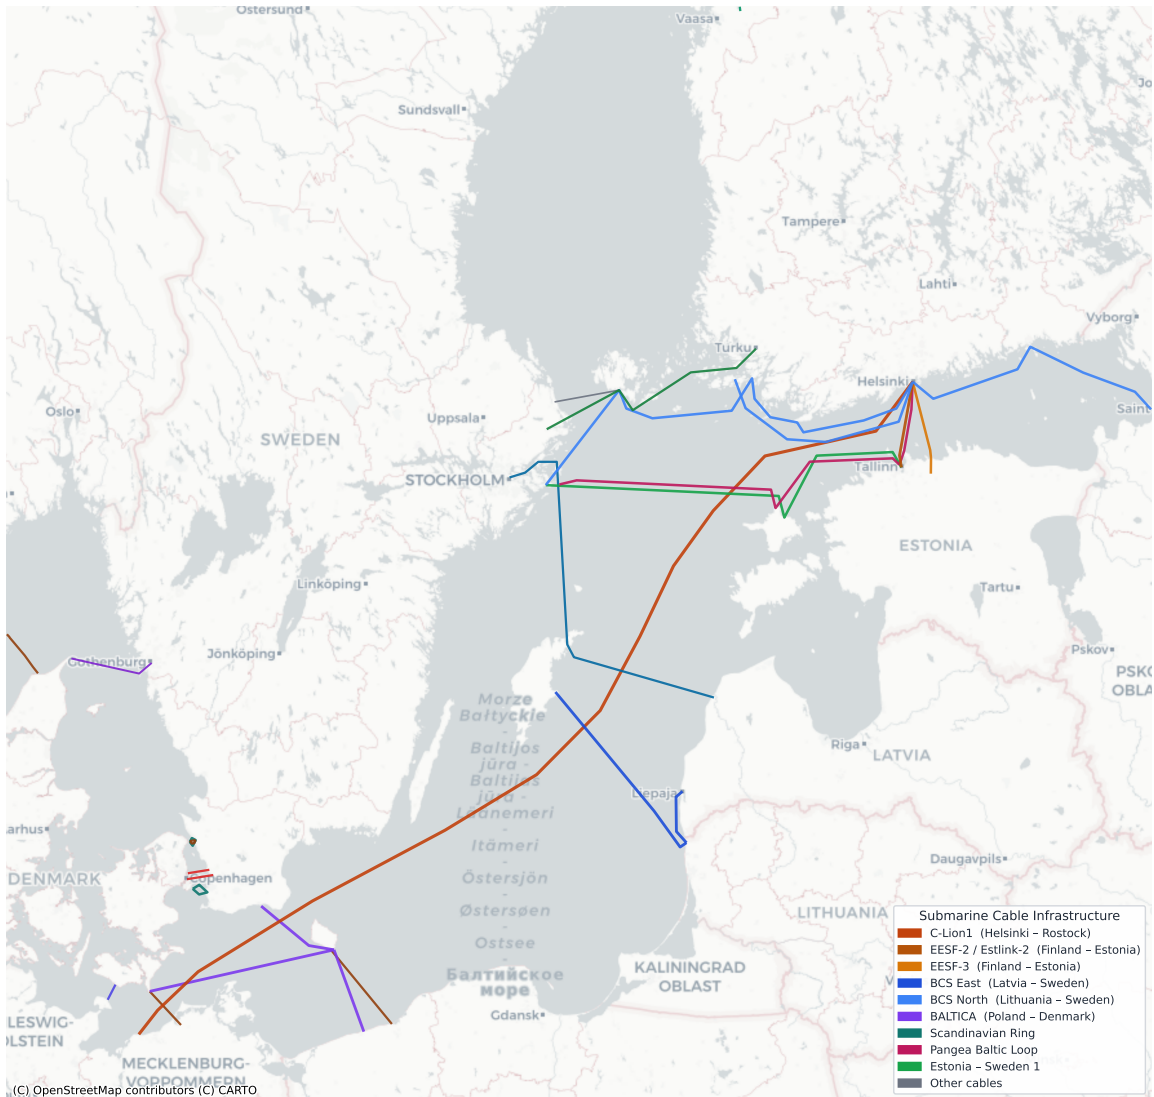


Figure 1.1: Submarine Cable Map 2026: The Baltic Sea [10]

Historically, maritime risk assessments focused primarily on accidental damage, such as unintentional anchor drags or commercial fishing gear strikes. However, the regional security landscape has fundamentally changed. Recent years have seen an increase of deliberate subsea sabotage, exemplified by the Nord Stream pipeline explosions and repeated cable severances linked to "shadow fleet" vessels operating in the Baltic Sea [11, 12]. Malicious actors increasingly weaponise standard maritime operations, such as deliberate anchor dragging, to sever cables. These hybrid tactics exploit international legal ambiguities and the natural opacity of the marine environment, allowing perpetrators to disguise intentional attacks as routine maritime accidents while inflicting severe socioeco-

conomic consequences [13, 14, 15].

Countering these adaptive, intelligent threats requires persistent Maritime Domain Awareness (MDA) [16, 17, 18]. Continuous observation of maritime corridors can transform opaque underwater activity into actionable intelligence, reducing the strategic ambiguity that adversaries currently exploit [19]. Yet, existing monitoring capabilities remain structurally inadequate. Conventional crewed naval patrols are prohibitively expensive to maintain continuously, leaving vast spatial and temporal knowledge gaps across critical maritime corridors [9]. Consequently, a severe gap exists between the need to protect exposed subsea assets and the operational reality of scarce, fragmented monitoring resources [20].

In response to this escalating vulnerability, European and international institutions have fundamentally reshaped the regulatory environment. Recent frameworks, including the EU's NIS2 and Critical Entities Resilience (CER) Directives, alongside the EU Cable Security Toolbox, have established new baseline obligations for infrastructure resilience [21, 22]. Crucially, these frameworks explicitly mandate the active, persistent monitoring of submarine cable corridors [23]. This legislative shift establishes a clear, legally grounded imperative: maritime authorities and infrastructure operators must transition from reactive incident-response models to proactive, scalable surveillance solutions capable of persistent domain awareness and threat deterrence.

1.2 Core Problem

Despite strong regulatory and strategic pressures to secure CMI, implementing effective protection in the underwater domain remains intrinsically difficult. The extensive spatial footprint of subsea assets makes continuous physical guarding impossible; conventional crewed patrols are too costly and intermittent to sustain persistent coverage; and the threat environment combines adaptive, intelligent adversaries with the legal ambiguities of international waters, rendering purely probabilistic risk models insufficient [24].

These factors create a severe operational mismatch: maritime authorities must protect a dispersed, high-value asset network with limited monitoring resources against strategic and adaptive opponents. The core challenge is therefore less about adding sensors than about the lack of a rigorous, proactive framework for allocating scarce surveillance capacity. There is a need for a structured decision-support methodology that can systematically deploy persistent monitoring assets, such as autonomous surface vehicles, to reduce strategic uncertainty, improve deterrence, and protect critical subsea infrastructure under uncertainty conditions.

Chapter 2

Problem Analysis

This chapter analyses the problem domain to identify the operational, technical, and theoretical challenges that motivate the research conducted in this thesis. The analysis proceeds from the foundational risk and uncertainty framework through the empirical Baltic Sea threat landscape to existing approaches and their limitations, progressively narrowing to the research problem formulated in Chapter 3.

2.1 Risk and Uncertainty Foundations

Assessing threats in the maritime domain requires moving away from conventional safety orientated frameworks. Standard probabilistic risk assessments typically define risk as the product of historical event frequency and associated consequences (Figure 2.1) [25]. By contrast, security focused risk science must explicitly incorporate strategic and intelligent adversaries who deliberately seek to maximise harm while exploiting contextual, environmental, and legal ambiguities to preserve plausible deniability [26]. Because such adversaries dynamically adjust their tactics in response to defensive measures, estimating the likelihood of hostile actions solely by historical data is inadequate [27, 26].

Consequently, security orientated risk science conceptualises risk not as a single probability value, but as a characterisation of the interaction between potential threat scenarios, their consequences, and the uncertainty that affects both [28].




Figure 2.1: Features of the risk concept by Aven and Thedaki [29]

Aven differentiates between *event risk*, which is the uncertainty associated with whether a hostile act will actually occur, and *vulnerability*, defined as the conditional severity of consequences if a hostile breach is successful [27]. Formally, if A' denotes the initiating threatening event, then vulnerability is characterised by the tuple $(C, U | A')$, that is, the consequences C and associated uncertainties U given that A' has occurred [30, 31]. Risk, by contrast, retains the full unconditional triplet (A, C, U) , where A encompasses the set of possible threat scenarios and U covers the uncertainty about both their occurrence and their outcomes (Figure 2.2) [32].




Figure 2.2: Risk and vulnerability [29]

The evaluation of these dimensions strongly depends on the background expertise of subject matter specialists [33]. However, in marine contexts this knowledge base is particularly limited, largely due to a chronic scarcity of direct sensory and observational data. To mitigate this deficiency, analysts rely on semi-quantitative assessments and expert judgments expressed as subjective probabilities to assess the Strength of Knowledge (SoK) that underlies their threat evaluations [34, 35, 36].

To transition from unstructured expert speculation to a systematic threat characterization, this thesis adopts Aven's modern risk science framework (A', C', Q, K) [35, 32]. Within this perspective, risk is defined in terms of the consequences (C) of an activity and the uncertainties (U) associated with those consequences [32]. The elements (A', C', Q, K) constitute the analysts' formal representation of this risk, where A' denotes the specified threat scenarios, C' the specified consequences, Q the measure of uncertainty, and K the

supporting background knowledge [35]. This conceptual structure underpins the systematic analysis of hybrid threats presently manifesting in the Baltic Sea.

2.2 Baltic Subsea Threat Environment

2.2.1 Geostrategic Significance of Baltic Maritime Corridors

The Baltic Sea hosts a network of maritime corridors of high strategic consequence, channelling international shipping and a dense concentration of critical subsea assets between Northern and Central Europe [6, 37]. Among the geographic features shaping this security environment, three island groups, Bornholm (Denmark), Gotland (Sweden), and Saaremaa (Estonia), occupy commanding positions over the region's principal sea lines of communication, giving their controlling states a disproportionate capacity to monitor and contest vessel transit through adjacent corridors [38]. The enduring strategic logic of Baltic island control is historically attested: Soviet forces occupied Bornholm from May 1945 to April 1946, a full year after Germany's capitulation, before withdrawing under sustained Danish and British diplomatic pressure [38]. The global stakes of protecting the geographical bottle-necks and their respective infrastructure are substantial: subsea cables carry most of the global data flows and underpin approximately USD\$10 trillion in daily financial transactions [9], and total EU subsea cable capacity quadrupled from 318 Tbit/s in 2010 to 3,755 Tbit/s by 2024, with the 33 newest cable systems alone accounting for 74% of operational capacity [39].

A dense network of high-capacity fibre-optic systems traverses Baltic maritime corridors on routes linking Nordic and Central European digital infrastructure [10]. The European Union's NIS2 and CER Directives require operators to safeguard both subsea infrastructure and digital assets from interference or physical destruction. Providers are required to deploy comprehensive security protocols to protect from hazards, such as ongoing threat assessments, restricted entry systems, and isolated network zones [21, 22]. Meanwhile, regional offshore energy developments in the Baltic Sea are poised to add significantly to the subsea infrastructure density; the planned commissioning of large-scale offshore wind aggregation hubs by 2030 will route high-voltage direct current cables to multiple national grids, serving millions of households [40, 41]. The co-location of irreplaceable communications and energy infrastructure within a heavily trafficked maritime corridor substantially amplifies the systemic consequences of any single deliberate disruption [42].

2.2.2 Systemic Vulnerabilities and Supply-Chain Dependencies

Beyond the physical exposure of individual cable segments, the Baltic threat environment is embedded within a set of structural dependencies that amplify the consequences of any single disruption [43, 44, 45]. The EU Expert Group's 2025 risk assessment identified that

the optical fiber used in long-haul subsea cables is manufactured exclusively by US and Japanese firms, and that the components in subsea repeaters are likewise sourced solely from US suppliers [39]. This supply-chain concentration means that accelerated cable replacement cannot be achieved independently of third-country manufacturing timelines, regardless of EU-level funding commitments.

These vulnerabilities are also present with the rise of Chinese firms like HMN Tech, which presents a supply-chain risk through infrastructure mapping and potential "back-door" access during installation. This structural dependency is now paired with a kinetic threat: Russia and China-linked vessels have been suspected of using anchor dragging to sever cables in the Baltic Sea, and Beijing's leaked patent of specialized cable-cutting devices capable of deep-sea sabotage, highlight the risk of the weaponization of maritime commercial activities into gray-zone aggression, also known as hybrid warfare [46, 47, 48, 49], which is explained in Figure 2.3.

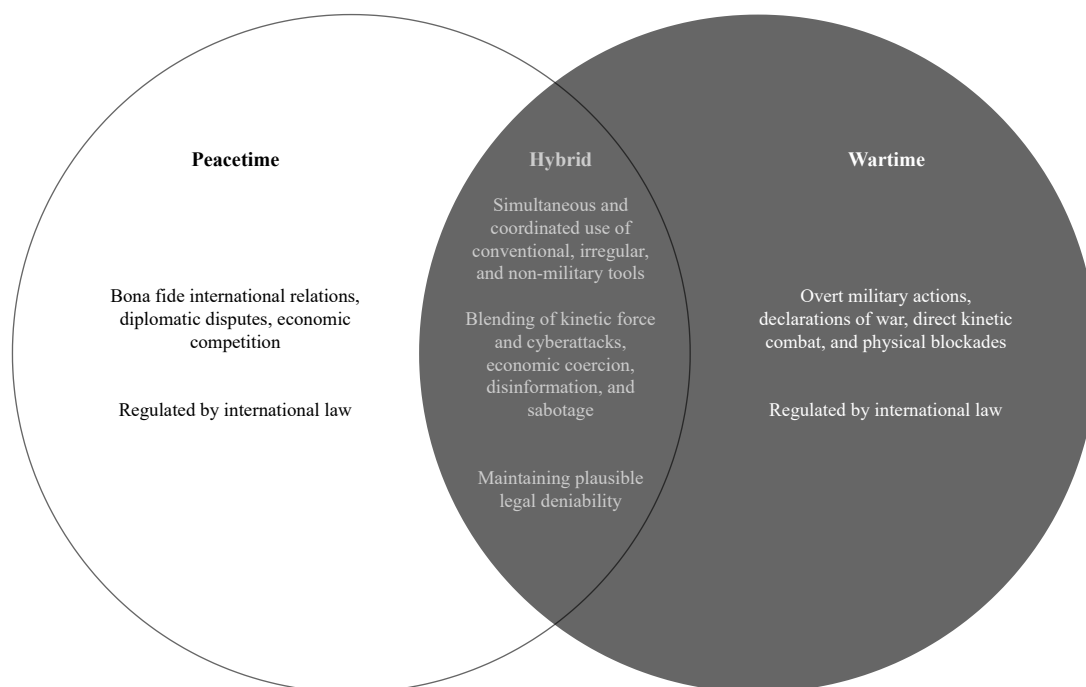


Figure 2.3: Hybrid warfare in relation to peacetime and wartime conditions

The global cable repair fleet compounds this constraint: the available pool of cable repair vessels is aging and numerically insufficient to address simultaneous multi-point failure scenarios of the type seen in the Baltic between November 2024 and January 2025, when seven cables were severed within a time period of a few months [39, 50].

2.2.3 Escalating Pattern of Subsea Sabotage in the Baltic Sea

The Baltic Sea has experienced several suspected deliberate attacks on undersea infrastructure since 2022, which Figure 2.5 visualizes for subsea cables [15]. In September of that year, three of the four conduits in the Nord Stream 1 and 2 pipeline systems were destroyed near Bornholm; Swedish forensic investigators subsequently recovered physical evidence of underwater explosive charges at the site before closing the investigation in 2024 without formal attribution [11]. This event established a strategic precedent for high-impact kinetic attacks on unmonitored subsea infrastructure in the region [45, 13].

A subsequent cluster of incidents has targeted fibre-optic data cables specifically. In October 2023, the *Newnew Polar Bear*, a Hong Kong-flagged bulk carrier, severed two undersea cables and a subsea pipeline in the Gulf of Finland while ignoring hailing from Finnish and Estonian authorities [12]. In November 2024, the Chinese-flagged bulk carrier *Yi Peng 3* was identified as the probable source of simultaneous ruptures to cables linking Germany to Finland and Lithuania to Sweden, and subsequently refused to enter Swedish waters where enforcement jurisdiction could have been exercised [12].

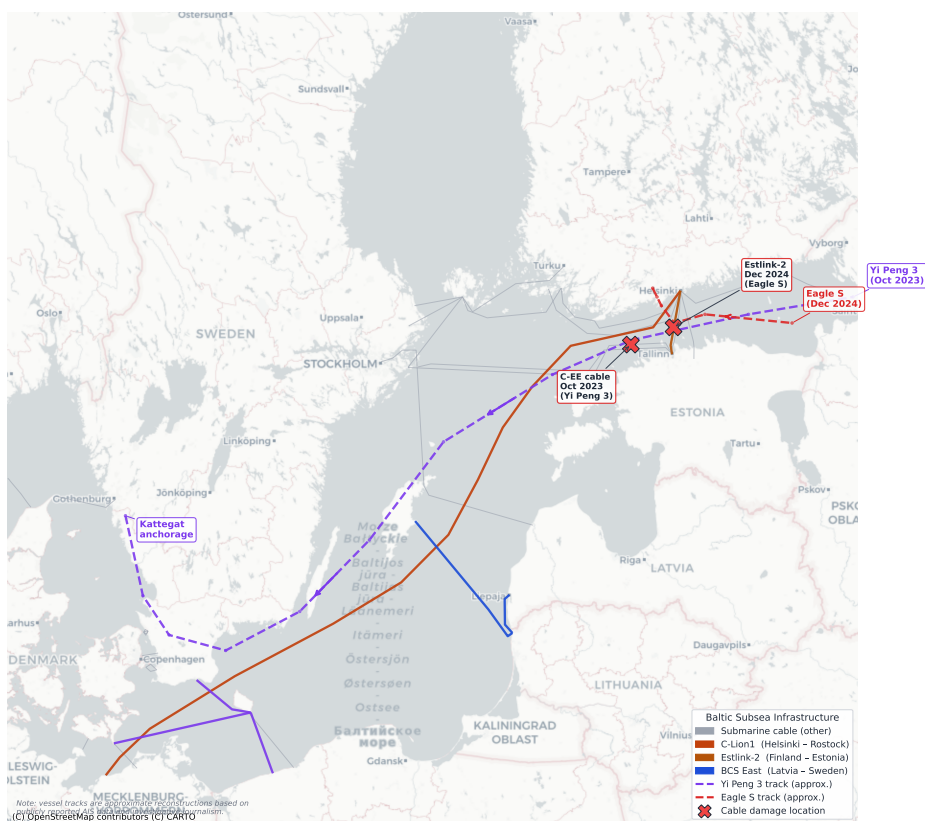


Figure 2.4: Yi Peng 3 and Eagle S incidents [12]

In late December 2024, Finnish authorities boarded the tanker *Eagle S*, a vessel linked

to the Russian shadow fleeton suspicion of severing four data cables and the Estlink 2 power interconnector between Finland and Estonia, representing the first successful enforcement seizure of a shadow fleet ship in the Baltic Sea [12, 51]. The DDIS reports that approximately ten subsea cables in the broader Baltic region have been severed since 2022, with seven of those incidents concentrated in the two-month window between November 2024 and January 2025 [50]. In May 2025, Estonian authorities detained the shadow fleet tanker *Jaguar* following suspected threats to subsea infrastructure, while Polish maritime units simultaneously observed a separate vessel conducting anomalous maneuvers near a Baltic power cable [12, 52].

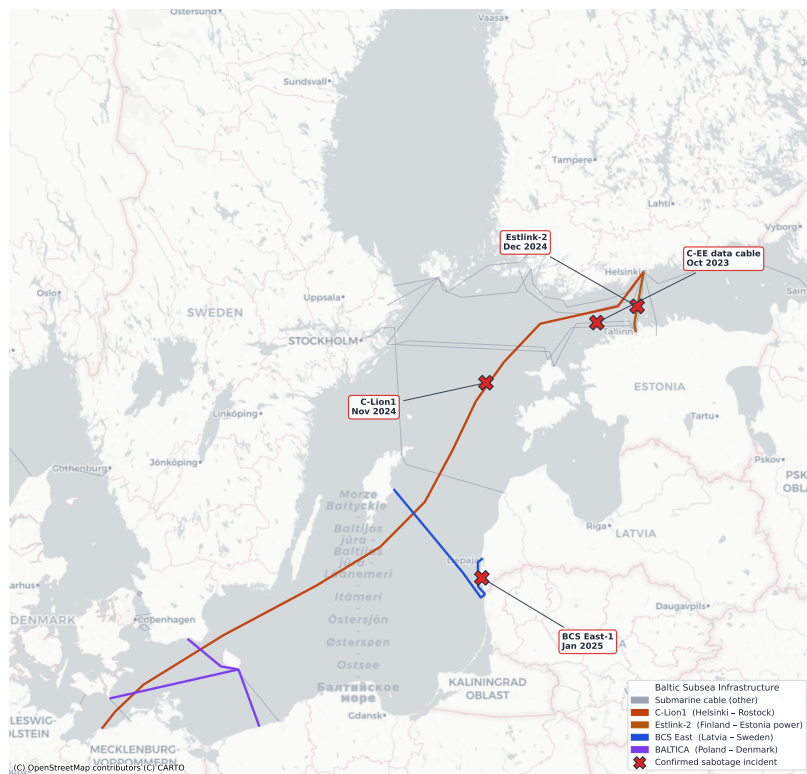


Figure 2.5: Map of Baltic Sea incidents since 2023

A summary of suspected recent grey-zone sabotage incidents, presented in Table 2.1, illustrates a distinct pattern of attack vectors and highlights the evidentiary challenges inherent in prosecuting these actors.

Date	Asset & route	Suspected vector (vessel, if identified)	Investigation status
Sept. 26, 2022	Nord Stream 1 & 2 gas pipelines (Swedish/Danish EEZ) [53]	Subsea explosive demolition; traces of explosives recovered from site.	Sweden and Denmark closed investigations in 2024 without naming suspects; German prosecutors pursuing Ukrainian suspects, one approved for handover from Italy in Nov. 2025.
Oct. 8, 2023	Balticconnector gas pipeline (Finland–Estonia) and EE-S1 telecoms cable [54]	Anchor drag: Chinese container vessel <i>Newnew Polar Bear</i> ; damaged anchor recovered from seabed.	Captain charged with "criminal damage" in Hong Kong pre-trial hearing; Estonian authorities report limited Chinese cooperation.
Nov. 17–18, 2024	BCS East-West Interlink (Lithuania–Sweden) and C-Lion1 (Finland–Germany) [55]	Anchor drag: Chinese bulk carrier <i>Yi Peng 3</i> (Russian captain); inspected by joint German/Swedish/Finnish/Danish team in the Kattegat.	No charges filed; vessel released and departed for Egypt; German minister publicly called it "probably sabotage."
Dec. 25, 2024	Estlink 2 power cable plus four telecoms cables (Finland–Estonia) [56]	Anchor drag: oil tanker <i>Eagle S</i> (Cook Islands flag, linked to Russian "shadow fleet"); seized by Finnish authorities.	Finnish court dismissed the case in Oct. 2025 (intent not proven); vessel released, no charges against owner Caravella.
Jan.–Feb. 2025	LVRTC cable Latvia–Gotland (Jan. 26); C-Lion1 Finland–Germany again (Feb. 21) [57]	Vezhen (Bulgarian-operated, Malta flag) seized for LVRTC break; Feb. break east of Gotland unattributed.	Swedish prosecutors ruled the LVRTC damage <i>accidental</i> (anchor drop in high winds) and released <i>Vezhen</i> ; Feb. incident remains open.
Dec. 31, 2025	Elisa telecoms cable (Finland–Estonia, Gulf of Finland) [58]	Anchor drag: cargo vessel <i>Fitburg</i> en route Russia–Israel, caught with anchor lowered in Finnish EEZ.	Two crew arrested, two travel-banned (crew from RU/GE/KZ/AZ); active sabotage investigation by Finnish police [59, 60].
Jan. 2, 2026	Sventoji–Liepaja telecoms cable (Lithuania–Latvia, Arelion) [61]	Suspected anchor damage; vessel docked at Liepaja boarded by Latvian police.	Criminal proceedings opened; as of Jan. 5, 2026 police found no evidence linking the boarded ship; investigation ongoing.

Table 2.1: Suspected sabotage incidents against Baltic Sea subsea infrastructure, 2022–2026.

2.2.4 The Shadow Fleet as a Dual-Use Threat Vector

A central enabling mechanism for Russian hybrid operations in the Danish maritime domain is the shadow fleet: a large assemblage of aging tankers operating under flags of convenience with opaque ownership structures and often without valid third-party in-

insurance or current classification surveys [62, 63]. These vessels transit from the eastern Baltic through the Bornholm region and onward through Danish waters, primarily transporting Russian crude oil and refined petroleum products in circumvention of Western sanctions [51]. The Danish Maritime Authority documented 292 passages by EU-sanctioned tankers through Danish waters in 2025; broader estimates including non-designated shadow fleet vessels suggest approximately 1,000 such transits annually, approaching an average of three per day [51, 64]. Monitoring data from the DMA and partner nations confirms that shadow fleet vessels were present in Danish waters on an almost daily basis throughout 2025 [65]. The EU's consolidated sanctions register covered approximately 598 such vessels as of early 2026 [65].

In addition to facilitating sanctions evasion, shadow fleet vessels pose a security risk in Danish waters. Investigations into recent Baltic subsea cable incidents have consistently revealed a recurring pattern: these ships pass near undersea infrastructure while towing their anchors along the seabed, a method that can cut cables at reachable depths while preserving plausible deniability in the absence of independent subsurface monitoring [13, 12].

Moreover, the documented coordination of shadow fleet passages with Russian military air support, observed in the Baltic Sea and in at least one transit through the Danish Straits, suggests an operational synergy incompatible with the behaviour of ordinary commercial traffic [66].

At the European level, IISS analysis documents a 246% rise in confirmed Russian infrastructure sabotage between 2023 and 2024, with a further 25 incidents across NATO in the first five months of 2025 [9]. Denmark's Global Maritime Security Strategy 2025–2028 is one part of the multifaceted mitigation measures, together with initiatives such as the Joint Expeditionary Force (JEF), which aims, in addition to other objectives, to develop surveillance capacity, including uncrewed surface and underwater vehicles [67, 68, 69, 70]. Because deterrence by punishment has demonstrated limited effectiveness against sub-threshold hybrid operations, deterrence by denial, raising adversary detection risk and operational cost, is the principal instrument available to defending states [15]. Persistent autonomous patrol is the direct operational expression of that posture [68].

The (A', C', Q, K) Framework Applied to the Shadow Fleet

The assessment of maritime risk in Baltic Sea cable corridors is categorized as follows:

- **A' (Specified Events/Threats):** The primary event of concern is the transit of "shadow fleet" tankers through Baltic maritime corridors. In 2025, the Danish Maritime Authority (DMA) recorded **292 voyages** by EU-sanctioned tankers [51]. Broader estimates suggest a total of roughly **1,000 voyages annually**, averaging three transits per day [64]. These vessels, with an average age exceeding **20 years**, often operate with opaque ownership and lack verified insurance, increasing the potential for both accidental and intentional interference with subsea assets [51, 62].

- **C' (Specified Consequences):** The primary consequences involve physical damage to subsea critical infrastructure (SCI) in the Baltic Sea, such as the severance of communication cables or power lines. Recent events in the Baltic Sea have caused severe economic consequences, as the repair process is expensive [15]. Such events can lead to regional communication blackouts and significant societal disruption [32, 35].
- **Q (Measure of Uncertainty):** Uncertainty is expressed through a subjective, knowledge-based probability (P) combined with a judgment of the strength of the supporting knowledge [35, 32].
- **K (Background Knowledge):** The knowledge base (K) for this assessment includes DMA voyage logs, vessel age profiles, AIS tracking data, and assumptions regarding the behaviour of sanctioned actors [35, 32].

Bow-tie Diagram

The bow-tie diagram, presented in Figure 2.6 is a common risk analysis tool used in risk assessments to visualize the pathway between risk sources and consequences [32].



Figure 2.6: Bow-tie diagram showing a top-event and its related risk sources and consequences, and their respective preventive barriers and mitigation barriers [32]

Evaluating Strength of Knowledge (SoK)

The probability Q on its own is an incomplete representation of risk: two numerically identical assignments may rest on background knowledge K of very different quality, and

the same decision weight should not be attached to them [34]. The extended description (Q, SoK) addresses this by attaching a qualitative judgement of how well K supports Q . Following the security-adapted criteria of Askeland et al. [34], five knowledge components are assessed (phenomenological understanding and models, data, expert agreement, assumptions, and scrutiny of K), each labelled Strong, Moderate, or Weak, with the intermediate labels Moderate-Strong and Moderate-Weak available where the criteria for a single tier are only partially met.

Table 2.2: Strength of Knowledge (SoK) characterisation for the shadow-fleet risk to Baltic subsea infrastructure, following the security-adapted criteria of Aven et al. [34, 35].

Knowledge component	Evaluation for the shadow-fleet case	SoK label
Phenomena / models	The Russian shadow fleet is identifiable as a category, but individual ownership chains and tasking relationships remain opaque; the physical capacity to drag an anchor is trivially present, while intent is contested; actor-behaviour and consequence models for hybrid interference at sea are still maturing [35].	Weak
Data	High-volume AIS and DMA traffic logs are available, but direct evidence linking specific vessels to intentional cable damage is sparse; incidents are infrequent, frame conditions vary between cases, and atypical drivers of damage are poorly characterised [35].	Moderate
Expert agreement	Regional security services broadly converge on the choke-point vulnerability of Baltic corridors and the societal criticality of subsea cables; convergence is weaker on the evidentiary threshold for attribution [15].	Strong
Assumptions	The principal explicit assumption, that shadow-fleet status correlates with elevated propensity for hybrid interference, is plausible but simplifying; no documented process has been undertaken to surface tacit assumptions underlying this link [35].	Moderate-Weak
Scrutiny of K	The knowledge base has been reviewed for "unknown knowns" relating to Russian hybrid tactics, but signals/warnings and the time-evolution of those tactics have not been examined systematically [35, 71].	Moderate
Overall SoK	One component (models) satisfies a "weak" criterion while expert agreement is strong; this places the overall classification in the intermediate band [34].	Moderate-Weak

A Moderate-Weak overall classification indicates that the probability estimate (Q , or more commonly P) is fragile: specific hostile events A may be masked by the volume of routine traffic, and the resulting numbers cannot carry the weight that a Strong classification would license. This finding motivates mitigation measures that reduce epistemic uncertainty by gathering real-time, site-specific intelligence. Two caveats apply. First, the

step from component labels to an overall label is itself a qualitative judgement: the framework prescribes no mechanical rule for combining a Weak data component with a Strong expert-agreement component [34]. Second, the assessment characterises the knowledge behind probabilities that have already been assigned, not the completeness of the event set $\{A_i\}$, a particular concern in security contexts, where adversaries seek surprise and novel attack modes may sit outside the assessment altogether [34].

While the (A', C', Q, K) framework and the resulting Strength of Knowledge assessment effectively diagnose the operational gap and theoretically justify the need for persistent surveillance, they remain conceptual tools [71, 72]. They establish *why* a structured approach to reduce uncertainty is necessary, but they do not provide the quantifiable parameters required to actively allocate operational resources. To transition from this theoretical justification to a computable decision-support strategy, the physical vulnerabilities and strategic value of the subsea assets must eventually be quantified. This necessary transition from conceptual risk assessment to a numerical framework is addressed in later chapters.

2.2.5 Environmental Conditions as Threat Amplifiers

The physical characteristics of the Baltic Sea modulate both the probability of a successful attack and the effectiveness of any surveillance response [73]. The Baltic is a semi-enclosed, low-salinity sea with pronounced seasonal variation in sea state and visibility. Winter months bring reduced optical visibility, increased wave heights, and intermittent icing in northern areas, conditions that degrade electro-optical sensor performance and raise the operational demands on both crewed and autonomous patrol platforms [74, 75, 76]. Conversely, the dense merchant traffic of summer months provides greater environmental cover for vessels engaging in anomalous maneuvers, as deviations from normal routing are harder to isolate against a high-density traffic background [51]. A recurring tactic exploiting this condition is AIS manipulation: shadow fleet vessels have been documented operating with disabled or spoofed AIS transponders, creating "dark targets" that are undetectable by standard vessel traffic monitoring but can be identified by radar- and sensor-equipped autonomous platforms [74, 12, 15, 77]. Water turbidity in the operational area, driven by suspended sediments and biological activity, limits the effective range of optical sensors to tens of metres in many conditions, reinforcing the advantage of acoustic sensing modalities, sonar and passive hydrophone arrays, for subsurface activity detection [76, 75, 78]. Finally, water depth in the Baltic Sea, one of the shallowest in the world, averages just around 50 m [79]; shallow-water segments present easier anchor-drag access but also reduce the concealment advantage of deep-water operations [13]. These environmental parameters do not only provide background context, they directly influence the susceptibility of subsea infrastructure. Therefore, any robust methodology for distributing surveillance efforts must explicitly incorporate these factors, acknowledging that characteristics such as shallow-water cable intersections and regions with intense traffic

inherently increase the local risk profile.

2.3 Autonomous Systems and Patrol Optimization Approaches

2.3.1 Operational Rationale

The surveillance deficits identified in the preceding section are structurally incompatible with crewed maritime patrol as the primary response mechanism. Manned naval vessels typically achieve operational availability of only approximately 30% over their service lifetimes, whereas autonomous surface vehicle fleets have demonstrated sustained availability exceeding 90% across comparable deployment periods [74]. USVs further address the economic constraint directly: NATO's TFX-B trials demonstrated that equivalent maritime surveillance coverage could be delivered by uncrewed systems at approximately one-third the cost of crewed frigate deployments [69]. USVs are therefore not merely a cost-saving alternative to conventional patrol; they represent a qualitatively distinct technology capable of operating continuously in environments where persistent human crewing is logistically and financially unsustainable [80].

2.3.2 Demonstrated Capability in the Baltic Sea

Several recent deployments validate USV performance under operational Baltic Sea conditions. Denmark's deployment of four Saildrone Voyager (Figure 2.7) platforms under contract with the Danish Ministry of Defence sustained 92% fleet availability over a six-month mission, covering more than 20,000 nautical miles and detecting over 170,000 individual vessel contacts, including non-cooperative targets operating without AIS transmissions [74].



Figure 2.7: Saildrone Voyager USV: A 10-meter autonomous surface vessel utilizing wind and 4 kW electric power for 180-day maritime missions. [81]

Each Voyager combines wind-assisted propulsion with diesel-electric backup, enabling endurance exceeding three months; its sensor suite encompasses a multibeam sonar to 300 m depth, a sub-bottom profiler, passive acoustic arrays, an electro-optical/infrared camera, maritime radar, and dual-redundant satellite communications [82]. During NATO TFX-B Baltic exercises, the Saildrone fleet was the only participant to successfully identify simulated threat targets within the integrated NATO command structure, demonstrating operational interoperability with Allied networks [74]. NATO's Task Force X itself, established in February 2025 during the Baltic Sentry initiative, is explicitly organized around the manned-unmanned teaming paradigm, integrating USV platforms such as the T-38 Devil Ray and Arabian Fox MAST-13 with AI-enabled sensor fusion for persistent surveillance of critical subsea infrastructure [83].

2.3.3 Danish Industrial and Research Ecosystem

A mature Danish industrial and research ecosystem supports domestic USV deployment at operational scale. TUCO Marine Group, headquartered in Faaborg, produces the ProZero Sentinel USV, an 8.2 m, diesel-electric platform with a 2,000 kg payload capacity, a range exceeding 500 nautical miles, and a maximum speed above 35 knots, alongside long-endurance hydrogen fuel cell variants designed for deployments of up to twelve months [84]. SH Defence, based in Svendborg, developed The Cube: a containerized modular mission system now adopted as the standard reconfiguration architecture for the Royal Danish

Navy's newest modular vessels, enabling any equipped ship to deploy USV, AUV, and ROV launch-and-recovery capability from a standard shipping container footprint in approximately ten minutes [85]. The Danish startup Stormborn is developing the X-WAVE USV, an all-weather uncrewed platform designed for continuous maritime surveillance and critical infrastructure protection, backed by investment from Semco Maritime [86]. At the academic level, the Technical University of Denmark leads the EU Horizon Europe LORELEI-X project, developing fault-tolerant navigation, multimodal sensing, and cyber-physical resilience for autonomous surface vehicles in collaboration with international partners [87]. The Copenhagen Orca is an USV developed by the Copenhagen Group and is especially designed for subsea critical infrastructure surveillance [88]. This ecosystem positions Denmark as both an operational test bed and an industrial contributor to the NATO-standard USV supply chain required for persistent Baltic Sea coverage.



Figure 2.8: Technical specifications of the Copenhagen Orca USV (13 m, 5.5 tons). Features a top speed >40 knots, an 825 kg payload capacity, and an operational range of 500 to 1,400 nm. Mission endurance spans from 3 to 6 weeks depending on configuration [88].

2.3.4 Platform Architecture and the Role of USVs

The choice of sensor modalities must be matched to a carrier platform capable of operating in the subsea environment [89]. *Remotely Operated Vehicles* (ROVs) are tethered, tele-operated systems that provide high-bandwidth sensor data and dexterous manipulation capability at the cost of geographic tether constraints and the need for a surface support vessel [90]. ROVs are therefore well suited to scheduled inspection and emergency response but are structurally incompatible with wide-area continuous patrol. *Autonomous Underwater Vehicles* (AUVs) are untethered and programmable, enabling pre-planned survey missions of up to many weeks [91, 92]. AUVs are typically deployed in short operational cycles followed by surface recovery for data offload and recharging, introducing

temporal gaps in coverage. Battery-life is also a limiting factor for continuous surveillance, and their speeds usually are at most 5 knots (9 km/h). Neither platform achieves the persistent, real-time surveillance posture that cable protection requires.

Unmanned Surface Vehicles (USVs) resolve this architectural tension by acting as persistent carrier platforms that bridge the surface and subsurface environments [93]. A USV retains continuous access to GNSS for precise navigation, satellite communications for real-time data relay, and solar or wind energy harvesting for extended endurance. It can simultaneously deploy towed hydrophone arrays or various sensor payloads or other platforms, thereby coordinating acoustic wide-area surveillance with point inspection capability [94]. Crucially, a USV provides the persistent visual surface presence that constitutes a deterrent signal in its own right, an attribute that purely subsurface platforms cannot offer. The integration of USV surface coverage with cable-embedded sensing systems and opportunistic AUV inspection constitutes a layered sensor architecture suited to the continuous, wide-area monitoring requirements of the cable protection problem [95, 94].

The regulatory framework governing autonomous vessel operations in the Baltic maritime domain, including UNCLOS jurisdictional constraints, the ISPS Code, the NIS2 and CER directives, and ICPC cable protection zone recommendations is detailed in Appendix D. The stakeholder analysis regarding autonomous surveillance operations is presented in Appendix C (Table C.1).

2.4 Existing Approaches and Methodological Gap

Work	Game type	Bounded rationality	Multi-asset	Real data	Maritime / USV	Kinematic routing
Pita et al. [96] (2008) ARMOR	Stackelberg	No	No	Yes (LAX)	No	No
Shieh et al. [97, 98] (2012) PROTECT	Stackelberg	No	Yes	Yes (USCG)	No	No
Pita et al. [99] (2011) GUARDS	Stackelberg	No	Yes	Yes (TSA)	No	No
Yin et al. [100] (2012) TRUSTS	Stackelberg	No	No	Yes (Transit)	No	No
Rezazadeh et al. [101] (2017) Pipeline	SSG	No	No	Simulated	No	No
Duan et al. [102] (2021) Surveillance	Stackelberg	No	No	Simulated	No	No
Yang et al. [103] (2026)	Patrol	Yes (SubQRE)	No	Simulated	No	No
This work	DOBSS	No (SSE)[†]	Yes	Yes (Baltic)	Yes	Yes

Table 2.3: Comparison with prior work on patrol and security game optimisation. [†]The core pipeline uses the Strong Stackelberg Equilibrium (fully rational attacker); QRE is applied in the rationality sensitivity analysis with x^* held fixed.

For broader surveys of the patrolling and security-game literature see Samanta et al. [104], Tokel et al. [105] and Basilico [106]; for graph-based patrolling specifically see Alpern et al. [107] and Han et al. [108].

The literature establishes that game-theoretic patrol optimization is technically mature, operationalized in real security deployments including PROTECT [98, 97] at United States Coast Guard ports and ARMOR [96] at Los Angeles International Airport. However, a systematic gap exists. No existing approach combines all of the following: (i) adversaries with heterogeneous rationality levels modelled through Quantal Response Equilibrium, (ii) geospatially grounded risk scores derived from real infrastructure data and environmental conditions, and (iii) kinematically constrained route realization for autonomous surface vehicles. The PROTECT and ARMOR systems assume fully rational adversaries and do not address the mixed-intentionality threat typology. Yang et al. [103] introduce bounded rationality in a patrol context but operate in a simulated environment without real geospatial data. None of the reviewed works addresses the specific operational context of USV-based subsea cable protection. This gap directly motivates the research problem formulated in Chapter 3.

2.5 Analytical Synthesis

The analysis conducted in this chapter reveals a convergence of strategic, operational, and informational factors that together define the security challenge addressed by this thesis. The Baltic Sea maritime domain hosts critical subsea infrastructure of irreplaceable European importance, yet this infrastructure is systematically exposed to a range of adversarial and accidental threats that exploit the fundamental opacity of the underwater environment. The escalating pattern of documented sabotage incidents, seven cable severances within a two-month window spanning late 2024 and early 2025, confirms that this is not a theoretical risk but an active operational challenge [15, 60, 109, 110, 12].

The existing security response, despite significant NATO, EU, and national investment, remains constrained by the structural limitations of crewed patrol operations, fragmented sensor coverage, and the absence of systematic methods to translate observational data into calibrated threat assessments [69, 70, 111]. The Strength of Knowledge assessment (2.2.4) demonstrates that while expert agreement on the threat landscape is high, the critical data gap at the individual vessel intent level remains open, directly degrading the quality of risk assessments and the efficiency of defensive resource allocation.

Autonomous surface vehicle systems offer a technically validated path toward closing this surveillance gap [94]. However, deploying them effectively requires more than sensor capability: it demands a principled framework for allocating finite monitoring resources against a heterogeneous adversary population in a way that accounts for strategic adaptation and bounded rationality. The adversary typology established in Appendix Section E reveals that threats range from fully rational state-sponsored actors to coerced semi-rational operators and unintentional agents, and that any effective patrol strategy

must perform well across this entire spectrum simultaneously.

The stakeholder (Appendix C) and regulatory (Appendix D) analysis confirms that the deployment of an autonomous decision-support framework is not only technically justified but legally required under evolving EU and international obligations, with active corridor monitoring explicitly designated as a priority measure. At the same time, the jurisdictional and operational delimitations established in the regulatory review define the boundaries within which any patrol optimization framework must operate.

Together, these findings point to a single analytical gap: the absence of a decision-support framework that jointly optimises patrol coverage against a heterogeneous, empirically grounded adversary population while producing kinematically feasible routes for a real autonomous surface vehicle. Chapter 3 formalises this as the central research problem and defines the scope and methodology through which it is addressed.

Chapter 3

Problem Formulation

3.1 Conclusion of Problem Analysis

The analysis conducted in Chapter 2 establishes four converging conditions that frame the central research problem of this thesis.

First, the maritime domain of the Baltic Sea contains critical subsea infrastructure important for Europe. Because these systems are physically exposed within densely trafficked sea lanes, they are highly susceptible to both intentional sabotage and accidental damage.

Second, the threat environment is inherently varied. Potential adversaries include state-sponsored entities conducting pre-operational probing, commercial actors whose activities involve anchor-drag operations, and other unintentional contributors. No single, standard threat model can adequately represent this spectrum of behaviours.

Third, there is a persistent knowledge gap regarding the intentions of individual vessels. Although expert agreement on the overall threat environment is robust, the absence of real-time visibility into subsurface activities constrains the reliability of risk evaluations and reduces the effectiveness of allocating defensive assets.

Fourth, despite the fact that autonomous surface vehicles have been operationally demonstrated for surveillance in the Baltic Sea, there is currently no decision-support framework that simultaneously accounts for a diverse, empirically grounded set of adversaries and the kinematic routing constraints of an unmanned surface vessel within a realistic geospatial maritime setting.

These four conditions are formalised in the (A', C', Q, K) risk characterisation of Chapter 2: the adversary type set $\Theta = \{\theta_1, \theta_2, \theta_3\}$ constitutes A' , the infrastructure consequence structure constitutes C' , the analyst's subjective belief about relative threat frequency is expressed as the uncertainty measure $Q = p(\theta)$, and the geospatial and intelligence inputs, whose limited observational grounding produces the knowledge gap of condition three, constitute the background knowledge K . The computational framework developed in Chapter 4 operationalises this characterisation directly, as detailed in Table 4.1.

3.2 Research Question

The preceding analysis identifies the following core research problem: current approaches to deploying autonomous surface vehicle patrol systems do not account for the strategic heterogeneity of the threats they face, leaving defensive resource allocation suboptimal. This motivates the central research question:

How can game theory be used to plan patrol strategies of Unmanned Surface Vehicles to secure critical maritime infrastructure?

Three sub-questions structure the investigation:

1. How can the diverse threat landscape, from state-sponsored actors and unintentional fishing, cargo/tanker shipping accidents, be captured within a unified framework?
2. How can theoretical coverage requirements be translated into kinematically feasible and unpredictable patrol routes for an autonomous surface vehicle?
3. To what extent does a structured, proactive patrol strategy improve expected defensive outcomes, and how robust is it against varying levels of adversary sophistication?

3.3 Methodological Approach

The research question is investigated using a computational framework based on the Stackelberg Security Game paradigm, which is developed and tested here as a proof-of-concept decision-support system.

The framework proceeds in two main stages. In the first stage, a game-theoretic algorithm is applied to a geospatially explicit representation of the case study region to derive optimal mixed coverage policies. This stage seeks to maximise expected defensive utility across all modelled adversary classes simultaneously, thus directly accounting for the heterogeneous threat environment. In the second stage, the resulting abstract coverage probabilities are converted into executable USV movement sequences. This step ensures that the derived patrol paths comply with the kinematic limitations of a real vessel while preserving strategic unpredictability.

Instead of attempting to remove the epistemic uncertainty inherent in the maritime context, the approach deliberately decomposes it. Uncertainties related to adversary type, rationality, and utility valuations are formally tested through extensive sensitivity analyses, to characterise the robustness of the strategy. In contrast, physical and kinematic constraints are treated as fixed inputs to retain computational tractability and operational clarity.

3.4 Delimitations and Assumptions

The following delimitations define the scope of the framework:

- **Proof-of-concept:** The framework is a decision-support model demonstrated on a case study. It is not an operational system and has not undergone field validation.
- **Adversary types:** The model restricts its focus to three specific adversary profiles (unintentional fishing, unintentional shipping, and strategic adversaries). Dynamic type updating and coalition adversaries are not modeled.
- **Single USV:** The framework optimises the patrol strategy of a single autonomous surface vehicle. Multi-vessel coordination is not addressed.
- **Physical threat vector:** The framework addresses physical infrastructure threats accessible to surface patrol, such as anchor-dragging. Cyber intrusions are outside the scope.
- **Static patrol strategy:** The patrol strategy is pre-computed and committed in advance. Real-time adaptive updating based on observed vessel behaviour is not implemented.
- **Static infrastructure:** Cable routes and asset risk scores are treated as fixed inputs. Dynamic changes to infrastructure or routing are not considered.
- **Deterrence and detection only:** The USV is modelled as a deterrence and detection asset. Enforcement or kinetic interdiction capability is outside scope.

3.5 Scientific Contribution

This thesis makes three principal contributions to the literature on game-theoretic patrol optimization and maritime infrastructure protection:

1. **Problem formulation:** The formulation of submarine cable protection as a security game grounded in real geospatial data, incorporating a multi-type adversary model that captures the heterogeneity of the Baltic threat landscape.
2. **Scalability and demonstration:** A proof-of-concept demonstration establishing that game-theoretic patrol optimization is computationally tractable for decision-support applications at an operationally relevant resolution.
3. **Robustness characterisation:** An empirical evaluation of the framework's patrol strategy across a wide spectrum of adversary rationality levels, from near-random to fully rational behaviour, demonstrating the conditions under which the game-theoretic approach provides consistent defensive advantage.

Chapter 4

Problem Solution

In this chapter, the technical solution developed to address the critical infrastructure vulnerabilities identified in the earlier chapters is introduced. To convert the high-level strategic problem into a practically deployable decision-support tool, this thesis formulates the Bayesian Stackelberg Security Game (BSSG). This model operates as an integrated computational framework that couples game-theoretic resource allocation with detailed environmental and spatial information, ultimately transforming abstract coverage specifications into feasible, kinematically consistent patrol trajectories for a USV.

4.1 Methodology

The game-theoretic solution underlying this framework draws on two well-established bodies of work. Patrolling games, introduced by Alpern, Morton, and Papadaki [112], model the adversarial interaction between a mobile patroller and an attacker who selects an attack location and timing to evade detection. Formulated as zero-sum games on graphs, they yield optimal randomised patrol strategies that maximise the probability of intercepting under the assumption that the attacker observes patrol patterns before committing to an attack. Stackelberg security games [113] extend this framework by explicitly modelling the strategic commitment advantage of the defender: the leader (defender) announces a mixed strategy first, after which the follower (attacker) best-responds. The Decomposed Optimal Bayesian Stackelberg Solver (DOBSS) algorithm [113] enables efficient computation of these equilibria even under attacker-type uncertainty. This paradigm has been operationalised on a scale in the PROTECT system [98], which schedules US Coast Guard patrols in major US ports and directly demonstrates the viability of game-theoretic patrol optimisation for maritime domain protection. The present project adapts and extends these foundations to the specific operational and geospatial context of autonomous USV-based submarine cable protection in the Baltic Sea [80, 114].

To put this theoretical framework into practice, the abstract elements of the game matrix must be concretely parameterised using a real-world physical infrastructure network.

Game-theoretic security optimisation presupposes a delimited domain in which targets have uneven strategic values and the defender operates under stringent resource limitations. The Bornholm Basin case study fulfils these structural conditions. As a narrow maritime corridor that hosts critical data and power transmission routes, the area constitutes a markedly asymmetric distribution of asset values, where the failure of a single component can trigger far-reaching systemic impacts. By anchoring targets, payoffs, and constraints in the observed geospatial properties of this basin, the abstract Stackelberg model is converted into a data-informed, computationally tractable decision-support instrument.

The (A', C', Q, K) risk characterisation established in Chapter 2 maps directly onto this two-phase framework, as summarised in Table 4.1.

Element	Aven framework	BSSG operationalisation
A'	Specified threat events	Adversary type set $\Theta = \{\theta_1, \theta_2, \theta_3\}$ and their action spaces (accidental anchor-drag, commercial transit, deliberate sabotage)
C'	Specified consequences	Defender utility $U_d^i(t) = -\gamma \cdot \text{AssetValue}(t)$, encoding MTTR, redundancy, and strategic weight per infrastructure node
Q	Subjective uncertainty measure	Prior distribution $p(\theta)$ over follower types, quantifying the analyst's belief about relative threat frequency given background knowledge K
K	Background knowledge	EMODnet 2023 vessel density rasters, Folk 5 substrate penetration data, DDIS threat assessments, and cable infrastructure parameters (MTTR, redundancy, strategic weight)

Table 4.1: Mapping of Aven's (A', C', Q, K) risk characterisation onto the BSSG formulation.

4.2 Data and Spatial Context

A set of concrete parameters must be defined as inputs to the BSSG model. This section introduces the spatial and environmental datasets compiled for the Bornholm Basin case study. Visual representations of these raw geospatial data establish the physical baseline of the study area prior to the application of any strategic or mathematical transformations.

To ensure proper alignment between all data points, Coordinate Reference System (CRS) EPSG:3857 is used.

4.2.1 Location

The case study centres on a 50 km x 50 km bounding box around Bornholm, anchored at coordinates 55.20° N, 14.48° E. The chosen area has been strategically selected due to the presence of multiple subsea assets, dense marine traffic, as the cargo and tanker shipping lanes pass through it and the previous occurrence of incidents, such as the Nord Stream pipeline explosion [11].




Figure 4.1: Case Study Area

4.2.2 Subsea Infrastructure

The subsea infrastructure considered in this case study comprises underwater cable systems in the Bornholm region, specifically telecommunication cables and high-voltage alternating current (HVAC) power export/interconnector cables. These assets are critical enablers of cross-border connectivity: telecommunication cables support digital communications and data traffic, while power cables transmit electricity between landing points and national grids. At the same time, their exposed linear geometry and long route lengths make them difficult to monitor continuously and susceptible to accidental damage and deliberate interference. For the BSSG model, cables are treated as spatially fixed targets whose risk is directly tied to repair time, redundancy, type of cable and physical susceptibility, furthermore it is influenced by their proximity to maritime traffic and seabed conditions. The following paragraphs summarise the key cable assets included in the Bornholm Basin case study.

Telecommunication cables

The cables shown in Figure 4.2 are the C-Lion1 and BALTICA cables. Both systems are laid in the Bornholm Basin and intersect areas of concentrated shipping activity, making them

relevant targets for the case study. Although they differ in age, architecture, and capacity, as shown in Table 4.2, they share similar exposure mechanisms, long, spatially fixed routes that are difficult to monitor and can be accessed by vessels along multiple points.




Figure 4.2: Subsea Telecommunication Cables

Attribute	C-Lion1	BALTICA
Ready for Service (RFS)	March 2016	March 1997
Current Age (as of 2026)	10 Years	29 Years
Total Route Length	1,173 km	437 km
Design Capacity	144 Tbps (Lit)	Legacy SDH / Broadband
Primary Landing Stations	Helsinki (FI), Hanko (FI), Rostock (DE)	Gedser (DK), Pedersker (DK), KoÅobrzeg (PL), Ystad (SE)

Table 4.2: Technical and Spatial Summary of Telecommunication Cables [10]

Power cables

In addition to data connectivity, the Bornholm region hosts subsea HVAC power cables that support regional energy transmission and cross-border interconnection. These cables are typically routed between coastal landing points and may incorporate near-shore protection measures, for example burial or armouring, yet they remain exposed along offshore stretches where access by surface vessels is feasible [115, 13]. From a risk perspec-

tive, power cables represent high-consequence single-point disruptions because damage can immediately curtail transfer capacity and impact grid stability. In this case study, the investigated power cables that cross or connect through Bornholm include the BEI Interconnector (Bornholm–Deutschland) and the Bornholm Cable, which are shown in Figure 4.3.

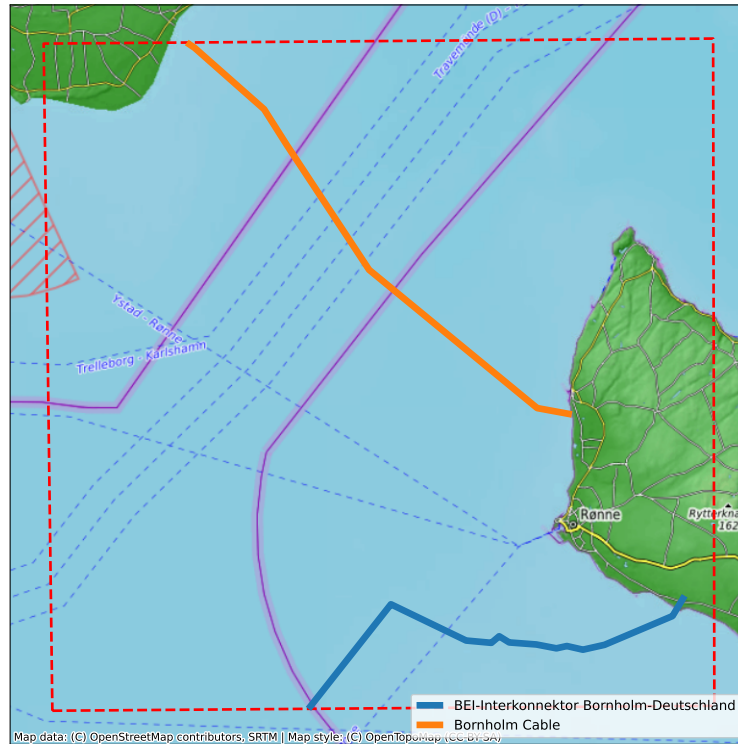


Figure 4.3: Subsea Power Cables

Attribute	Bornholm Cable (Sweden Link)	BEI Interconnector (Germany Link)
Ready for Service (RFS)	December 1980	Planned (Expected 2030)
Current Age (as of 2026)	46 Years	Under Construction / Planned
Total Route Length	50 km (43.5 km Subsea)	470 km (Total Project Area)
Design Capacity	60 MW (60 kV HVAC)	2,000 MW (525 kV HVDC)
Primary Landing Stations	Hasle (DK), Borrby (SE)	Bornholm (DK), Vierow (DE)

Table 4.3: Technical and Spatial Summary of Power Interconnectors [116, 40]

As BEI Interconnector power cable is still under construction, it is not included in the further sections and the BSSG.

Risk score calculation

With the subsea cable network in the Bornholm Basin now mapped, the next task is to quantify the utility of each individual cable so that it can be integrated into the BSSG risk score. Operationally, this utility captures generalised parameters that factor in Mean Time to Repair (MTTR), redundancy, type of cable (power or telecommunications) and armouring. This section presents the cable specific parameters that constitute the inputs to the utility function defined in Section 4.4.4.

Downtime Severity (DT_i) The duration of the outage is the dominant factor driving economic loss. This metric is expressed as MTTR. Power cables are assigned a higher Downtime Severity Factor because specialised repair vessels for these systems are much less available than fibre-optic cables [117]. As of 2023, the median time required to repair a telecommunications submarine cable was about 40 days, nevertheless, the general trend indicates that repair durations are rising. This increase may be linked to evolving maritime environmental conditions and regulatory frameworks [118]. Compared to subsea telecommunication cables, subsea power cables generally require a longer repair time. Industry statistics indicate that the typical outage duration for subsea power cables is approximately 3 months (≈ 90 days) [119]. This extended time frame is partly due to the need for project specific joints and specialised repair vessels for subsea power cable interventions [120].

Resilience Deficit (RD_i) The RD_i metric is given by 1 minus the Redundancy, where Redundancy denotes the fraction of a service that can be immediately rerouted through alternative paths in the event of a failure of the target cable [121].

The following list shows the calculations of RD_i for the individual cables listed in Tables 4.2 and 4.3.

- The Bornholm Cable is a point-to-point transmission line linking the island to Sweden. Apart from local generation sources, such as wind and thermal power plants, it serves as the island's only power connection, with no additional subsea interconnectors available [122]. Therefore, its RD is equal to 1 (as of now, once the BEI Interconnector is constructed, it can act as a redundancy).
- C-Lion1 provides a direct connection between Germany and Finland. In addition to the C-Lion1 link, Finland is served by more than 10 international subsea cables [123, 124]. During the 2024 C-Lion1 outage, network traffic in Finland experienced no noticeable, or only negligible, latency increases. Consequently, the RD value for C-Lion1 can be set to 0.1.
- BALTICA is one of the 31 internal subsea cables in the Northern Europe network. Therefore, the traffic can be rerouted to any of those in case on accidents. For this reason, the RD for the BALTICA cable is set to 0.1.

Strategic Consequence (W_i) This parameter represents the relative significance of the two cable categories. Since there are two cable types (telecommunication and power), and telecommunication is essential for digital communication, whereas power is essential for the operation of water supply, heating systems, and even telecommunication networks, their importance differs accordingly [125].

- $W_t = 1$
- $W_p = 2.5$

Physical Susceptibility (ψ_i)

In the context of infrastructure risk assessment, physical susceptibility is defined as an asset's inherent tendency to incur damage when exposed to a particular hazard. To quantify this parameter, the physical robustness of the two different cable types is analysed, this includes their armouring and diameter. The diameters of subsea power cables typically fall within the range of 7 to 21 cm [115]. In addition to having a larger diameter, these cables generally incorporate further structural layers, such as armouring and an outer sheath, that provide enhanced mechanical and environmental protection [126, 127]. The outer diameter of subsea telecommunication cables typically does not exceed about 5 cm [115]. These cables are available in several protection classes, such as lightweight (LW) and single-armoured (SA) designs [128]. Nevertheless, relative to the multiple protective layers used in subsea power cables, telecommunication cables generally provide a lower degree of mechanical protection. From these parameters, the ψ_i that is assigned to both cable types is the following:

- $\psi_t = 1$
- $\psi_p = 0.2$

Table 4.4 shows the different parameters obtained to be used as inputs for the target utility calculations.

Asset Name	Cable Type	DT_i (Days)	RD_i	W_i	ψ_i
C-Lion1	Telecommunication	40	0.1	1.0	1.0
BALTICA	Telecommunication	40	0.1	1.0	1.0
Bornholm Cable	Power	90	1.0	2.5	0.2

Table 4.4: Risk Score Calculation Parameters for Investigated Subsea Cables

To calculate the target asset value, utilised in the calculation of the utility functions in Section 4.4.4, the following equation is defined:

$$AssetValue(t) = DT_i \cdot RD_i \cdot W_i \cdot \psi_i \tag{4.1}$$

4.2.3 Sea Substrate

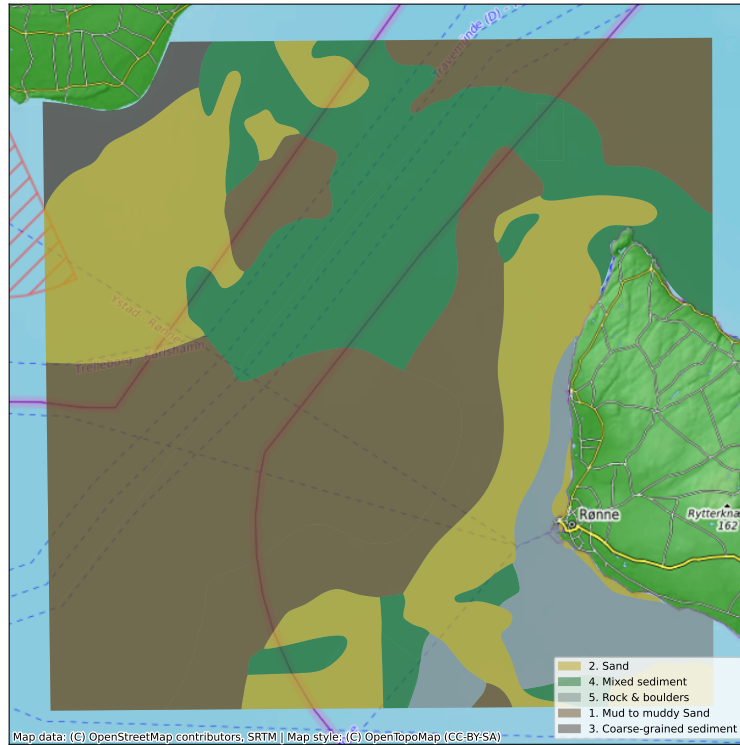


Figure 4.4: Sea Substrate

The categorisation of the sea substrate that has been used in this case is the Folk 5 classification [129]. This allows to broadly categorise the different substrates there are in the Baltic Sea. To determine the risk each of the substrate categories poses for cable sabotage, their penetration factors (f_a) are used. Table 4.5 presents the f_a for the different Folk 5 substrate types.

Folk 5 Class	Penetration (f_a)	Anchor Behavior	Cable Strike Risk (0-1)
Mud to Muddy Sand	5.0	Deep burial (high infrastructure threat)	1
Sand	≈ 1.0	Moderate burial (predictable protection)	0.4
Mixed Sediment	Variable	Bimodal (depends on mud fraction)	0.3
Coarse Sediment	< 0.5	Surface "chatter" (limited penetration)	0.15
Rock and Boulders	0.0	Surface skating (cannot reach cable)	0.05

Table 4.5: Folk 5 classification anchor penetration [130, 131, 132, 133, 134]

Within the BSSG algorithm, described in Section 4.3, the Cable Strike Risk values directly shape the attacker's cost function and expected utility. The physical characteristics

of the seabed operate as inherent defensive features, regions with elevated strike risk, such as mud or muddy sand, provide little resistance to anchors, thereby reducing the time, effort, and degree of specialised equipment needed to execute an attack. As a result, these high risk zones correspond to a very low attack cost (c_a) in the game model, which increases the attacker’s expected payoff. In contrast, hard substrates, such as rocks and boulders, impose a high cost on the attacker and sharply diminish the likelihood of a successful strike. By embedding these substrate specific risk values into the game matrix, the BSSG can more precisely determine the optimal patrol strategy of the defender, adaptively concentrating scarce maritime security assets on the most exposed portions of the subsea cable network, where environmental conditions are most advantageous to the attacker.

4.2.4 Marine Traffic

To provide an accurate representation of maritime activity for use in the BSSG algorithm, vessel density data from 2023 in the Bornholm Basin has been utilised [135]. From this raster dataset, the spatial patterns of marine traffic were extracted for three vessel categories: fishing, tanker, and cargo. To distinguish the main shipping lanes from negligible background traffic, a lower threshold of 0.2 was imposed on the density values, therefore all values below this cutoff were masked. This procedure restricts the analysis to statistically meaningful maritime routes. It should be noted that the data in this section are shown in their unprocessed form, additional transformations are introduced in Section 4.3 to render the data compatible with the BSSG algorithm. Figure 4.5 show a visual representation of the different vessel densities.

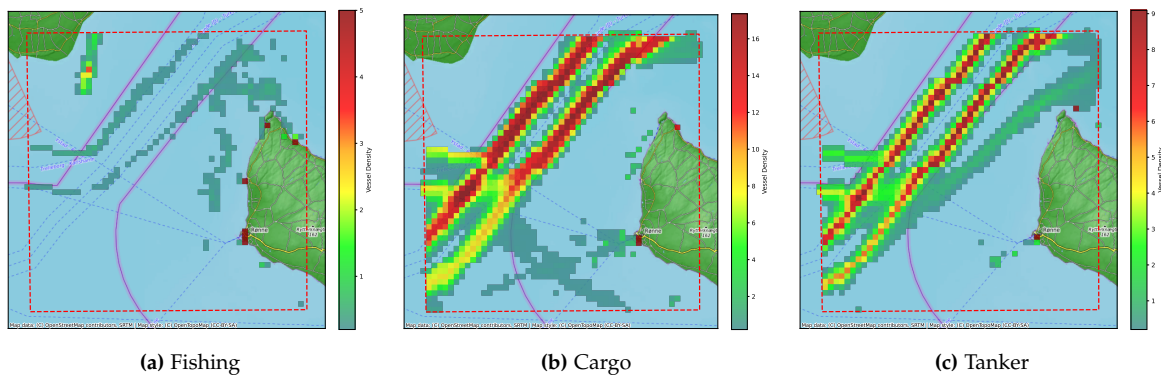


Figure 4.5: Vessel density (hours per square kilometres per month) in the case study area by vessel type.

To adapt the raster datasets to the discrete framework of the BSSG algorithm, the spatial domain is initially segmented into localised grid cells. For each discrete node, the average vessel density is derived from the underlying raster pixels. These heterogeneous values are then standardised by Min-Max normalisation: the mean density of each cell is divided by the maximum observed density throughout the study area, converting raw traffic counts

into a normalised probabilistic index from 0.0 (minimal traffic) to 1.0 (the most heavily trafficked node). After normalisation, the cargo and tanker traffic layers are combined by computing their mean element-wise, yielding a single commercial shipping layer, while the fishing density layer is retained as a separate dataset. Finally, these normalised spatial indices are directly incorporated into the adversary utility functions of the BSSG, which is described in more detail in Section 4.4.4.

4.3 DOBSS Formulation and Implementation

Having obtained the physical assets and spatial context of the Bornholm Basin, the DOBSS model can be formulated and implemented. The overall structure, parameters, and code implementation of the model are detailed below. Table 4.6 presents the notations used in the DOBSS implementation.

Symbol	Description	Case Study Value / Domain
<i>Sets and indices</i>		
T	Set of discretized target grid cells	$ T = 400$
Θ	Set of independent follower types	$\{\theta_1, \theta_2, \theta_3\}$
t	Target grid cell index	$t \in T$
θ	Attacker profile type index	$\theta \in \Theta$
<i>Decision variables</i>		
x_t	Defender's stationary mixed coverage probability	$[0, 1]$
q_t^θ	Follower best-response binary attack decision	$\{0, 1\}$
<i>Auxiliary and tracking variables</i>		
z_t^θ	Linearized utility product term ($x_t \cdot q_t^\theta$)	$[0, 1]$
v^θ	Maximum expected utility value attained by type θ	\mathbb{R}
<i>Game and environmental parameters</i>		
k	Total available physical patrol asset budget (USV count)	1
α	Residual asset damage fraction sustained under coverage	0.65
β	Deterrence multiplier reducing covered attacker reward	0.10
γ	Defender baseline loss scaling multiplier	10.0
w_{exp}	Strategic adversary traffic exposure penalty weight	0.50
$p(\theta)$	Prior probability distribution across follower types	$\sum_{\theta \in \Theta} p(\theta) = 1$
M	Big- M linearization bounding constant	Calculated dynamically
$c_a(t)$	Normalized physical attacker cost based on substrate	$[0.05, 1.00]$
λ	Attacker bounded rationality logic parameter (QRE)	$[0, \infty)$
<i>Valuations and payoff functions</i>		
$\text{AssetValue}(t)$	Combined baseline intrinsic valuation of subsea items	$[0.00, 45.00]$
$U_d^c(t), U_d^u(t)$	Defender utility functions (covered / uncovered)	\mathbb{R}^-
$U_{a,\theta}^c(t), U_{a,\theta}^u(t)$	Attacker type θ utility functions (covered / uncovered)	\mathbb{R}^+

Table 4.6: Unified case study parameterization and notation reference for the DOBSS implementation pipeline

4.3.1 Theoretical Foundations

DOBSS model

The continuous case study area is divided into a finite number of discrete target locations \mathcal{T} . The leader aims to compute an optimal mixed-strategy vector \mathbf{x} , where each component $x_t \in [0, 1]$ denotes the long-term proportion of the defensive resource time assigned to the target cell t . To linearise expected-utility terms and eliminate products between continuous defence decisions and binary attacker choices, McCormick envelope formulations are employed by defining an auxiliary variable $z_t^\theta = x_t q_t^\theta$. The global goal of the leader is then expressed as maximising the expected utility across the entire Bayesian distribution of follower types [113, 136]:

$$\max_{\mathbf{x}, \mathbf{q}, \mathbf{z}, \mathbf{v}} \sum_{\theta \in \Theta} p^\theta \sum_{t \in T} \left[z_t^\theta U_d^c(t) + (q_t^\theta - z_t^\theta) U_d^u(t) \right]. \quad (4.2)$$

In the standard security-game framework, target payoffs are modeled as functions of an intrinsic target value or risk score $R(t)$, constrained by an upper bound $R(t) \in [0, R_{\max}]$. For a generic attacker type θ , the payoffs under uncovered and covered states are specified as:

$$U_{a,\theta}^u(t) = R(t), \quad (4.3)$$

$$U_{a,\theta}^c(t) = \beta \cdot R(t), \quad (4.4)$$

while the corresponding defender payoffs are given by:

$$U_d^u(t) = -R(t), \quad (4.5)$$

$$U_d^c(t) = -\alpha \cdot R(t), \quad (4.6)$$

where the parameters $\alpha > 0$ and $\beta < 1$ quantify the reduction in defender loss and attacker gain, respectively, induced by active defensive coverage.

This optimization problem is constrained by structural resource limits and logical conditions derived from the behavioral assumptions of a Strong Stackelberg Equilibrium (SSE), in which the follower is rational and resolves ties in favor of the leader [137]. The leader's decision is subject to a finite resource budget k , while each independent follower type must choose exactly one optimal target node for intervention:

$$\sum_{t \in T} x_t \leq k, \quad x_t \geq 0 \quad \forall t \in T \quad (4.7)$$

$$\sum_{t \in T} q_t^\theta = 1, \quad q_t^\theta \in \{0, 1\} \quad \forall t \in T, \forall \theta \in \Theta \quad (4.8)$$

To preserve exact mathematical equivalence with the original mixed-integer nonlinear formulation, the linearized product variable z_t^θ is tightly constrained using the standard McCormick envelope:

$$0 \leq z_t^\theta \leq q_t^\theta, \quad z_t^\theta \leq x_t, \quad z_t^\theta \geq x_t - (1 - q_t^\theta) \quad \forall t \in T, \forall \theta \in \Theta \quad (4.9)$$

Crucially, the use of the McCormick envelope formulation introduces a zero approximation error or relaxation gap in this context. Because the attacker's target selection variable is strictly binary ($q_t^\theta \in \{0, 1\}$) and the defender's strategy is bounded within a continuous unit interval ($x_t \in [0, 1]$), the linear inequalities define a convex hull that binds the product exactly at its vertices. When $q_t^\theta = 0$, the constraints collapse to force $z_t^\theta = 0$; when $q_t^\theta = 1$, they collapse to enforce $z_t^\theta = x_t$. Consequently, the linearisation is mathematically exact, allowing the solver to reach a true global optimum without approximation artefacts.

The adversary's optimal best-response trajectory is explicitly imposed in the mixed-integer linear program through active Big-M constraints, which bound the attainable utility v^θ for each attacker type given the leader's resource allocation:

$$v^\theta \leq U_{a,\theta}^u(t) + x_t (U_{a,\theta}^c(t) - U_{a,\theta}^u(t)) + M^\theta(1 - q_t^\theta) \quad \forall t \in T, \forall \theta \in \Theta \quad (4.10)$$

$$v^\theta \geq U_{a,\theta}^u(t) + x_t (U_{a,\theta}^c(t) - U_{a,\theta}^u(t)) \quad \forall t \in T, \forall \theta \in \Theta \quad (4.11)$$

Here, the type-specific Big-M constant M is achieved by extracting the maximum absolute utility value from the entire payoff matrix and applying a scaling factor of 10:

$$M = 10 \cdot \max_{U \in \mathcal{U}} |U| \quad (4.12)$$

where, \mathcal{U} represents the set of all defined utility values in all targets and adversary types.

Spatiotemporal Markov routing model

Since the optimized static allocation vector \mathbf{x}^* implicitly assumes instantaneous repositioning of resources, it is not physically realizable for maritime platforms. To enforce kinematic feasibility, this static allocation is instead implemented as the stationary distribution of a discrete-time Markov chain, in which only targets with strictly positive coverage ($x_i^* > 0$) constitute the active state space [138]. Let A_{ij} denote the adjacency matrix with $A_{ij} = 1$ if the distance between nodes i and j can be safely traversed within one time step Δt at maximum speed v_{\max} . A stochastic transition matrix $P = (P_{ij})$ is then obtained via linear optimization:

$$\sum_{j \in T} P_{ij} = 1 \quad \forall i \in T \quad (4.13)$$

$$P_{ij} = 0 \quad \text{if } A_{ij} = 0 \quad \forall i, j \in T \quad (4.14)$$

$$\sum_{i \in T} x_i^* P_{ij} = x_j^* \quad \forall j \in T \quad (4.15)$$

Constraint (4.13) enforces row-stochasticity, ensuring that the asset always transitions to some admissible destination cell (including remaining in place), whereas (4.15) encodes exact stationarity, guaranteeing that the long-run spatial visitation frequencies of the moving asset coincide with the optimal game-theoretic allocation \mathbf{x}^* .


4.4 DOBSS Implementation

While Section 4.3.1 presents the classical and generalised mathematical structure of the DOBSS model, employing this framework as a decision-support instrument for the Bornholm Basin necessitates converting its abstract variables into measurable maritime indicators. This section describes the domain-specific adaptations applied to the baseline payoff matrices, specifies the spatial parameterisation, and explains the details of the programmatic implementation.

4.4.1 Discretised Grid \mathcal{T}

The initial step in configuring the BSSG model is to discretise the otherwise continuous study region, since the DOBSS algorithm operates on finite, explicitly defined sets of targets and actions. Although a study area could be partitioned using nodes/edges, this model adopts a uniform grid composed of 2.5×2.5 km cells. This resolution was selected as a compromise between computational efficiency and adequate spatial granularity. As a result, each grid cell represents an individual, discrete target in the game, whose overall risk and utility values are derived from the maritime infrastructure (power and telecommunication cables) and environmental characteristics located within that cell.

$$|\mathcal{T}| = 20 \times 20 = 400 \text{ cells}, \quad t_i \in \mathcal{T}, \quad i = 1, \dots, 400. \quad (4.16)$$



(a) Continuous operational environment showing critical subsea infrastructure (b) Discretization of the area into a finite game theoretic grid (\mathcal{T}) (c) Detailed view of a single 2.5×2.5 km target cell ($t \in \mathcal{T}$)

Figure 4.6: Spatial discretization process of the Bornholm case study area, translating the continuous maritime environment into a discrete set of targets for the security game model

4.4.2 Environmental Constraints and Attacker Cost

In Section 4.2.3, the cable strike risk associated with various seabed types, classified according to the Folk 5 scheme, is introduced. Within the BSSG algorithm, the seabed char-

acteristics have a direct impact on the operational costs of an attacker. The attacker cost c_a was derived from the values provided in Table 4.5. Table 4.7 presents the resulting attacker cost c_a for each sea substrate category.

Substrate Type	Cost Label	Attacker Cost c_a
Mud	Very Low	0.05
Sand	Low	0.15
Mixed Sediment	Medium-Low	0.30
Coarse-Grained Sediment	Medium-High	0.40
Rock	High	1.00
Unknown / No Data	Neutral default	0.50

Table 4.7: Seabed substrate types and corresponding normalised attacker operational costs

Nodes that do not overlap with any substrate polygon are given a neutral default value of 0.5, ensuring that the optimisation is not skewed toward terrain that is artificially easy or difficult.

The attacker cost (c_a) quantifies the physical effort and operational risk an adversary must incur to compromise seabed infrastructure and is directly linked to the penetration factor (f_a) associated with each substrate type. Within the BSSG framework, this cost is formally embedded in the attacker’s expected utility as a proportional reduction term. Specifically, the baseline reward (asset value) from successfully attacking a node (t) is multiplicatively adjusted by the factor $(1 - c_a)$. This construction guarantees that more demanding terrain proportionally reduces the effective value of a target, appropriately balancing environmental penalties against asset magnitude. As a result, it provides a realistic representation of the deterrent effect exerted by harsh environments, steering the attacker’s optimal strategy away from high-effort targets, such as cables entrenched in rock, whenever less demanding but still high-value alternatives are present.

Conversely, the defender’s capacity to defend the infrastructure is fundamentally constrained by operational limitations. In the BSSG model, the dominant constraint is the limited number of patrol assets, represented by k . This cap allows the defender to monitor only a restricted subset of the discretised study region at any moment, compelling the algorithm to focus on zones of elevated risk and value. In addition, practical routing and temporal bounds further limit the defender: any computed patrol plan must be physically feasible for maritime units within standard operational time frames, thereby ruling out patrol allocations over spatially separated areas that cannot be reached in time. In this study, the default number of patrol assets available is $k = 1$.

4.4.3 Follower Types Θ and Bayesian Threat Priors $p(\theta)$

In this case study, three distinct follower categories are considered, represented by $\Theta = \{\theta_1, \theta_2, \theta_3\}$. These follower types are associated with Bayesian prior probabilities, as shown in Table 4.8. The priors are modelling assumptions that are informed by empirical evidence but are not directly estimated from it [139]. The different types of adversaries are explained in more detail in Appendix E.

Type	Description	Prior $p(\theta)$	Dominant utility driver
θ_1	Fishing trawler	0.86	Fishing vessel density
θ_2	Cargo/tanker	0.12	Shipping lane density
θ_3	Shadow fleet / state actor	0.02	Asset value & attacker cost

Table 4.8: Bayesian prior probabilities over follower types

Within the BSSG framework, each follower type aims to optimise its own specific utility function. For unintentional actors (θ_1 and θ_2), their perceived utility is largely independent of the infrastructure itself, they behave “rationally” in the sense that they favour lucrative fishing areas or efficient shipping routes. As a result, their threat signatures align closely with historical vessel density patterns. In contrast, the malicious actor (θ_3) behaves as a conventional adversary, with a utility function explicitly designed to prioritise high value targets while remaining subject to operational constraints, such as the complications posed by the seabed substrate.

By explicitly representing these heterogeneous types, the DOBSS algorithm enables the defender to compute a robust mixed strategy that simultaneously mitigates the highly likely, low intent risk of accidental damage and preserves resilience against low probability but high consequence intentional attacks. The defender incorporates uncertainty over the realised attacker type by maximising expected utility over all possible types, with each attacker type weighted by its corresponding prior probability $p(\theta)$.

4.4.4 Target Utilities U_a, U_d

In the BSSG framework, the payoff matrices for the defender and the attacker jointly characterize the strategic environment. For each discretised target cell $t \in \mathcal{T}$, four fundamental utility terms are specified to capture the consequences of an attack under both coverage states:

- $U_d^c(t)$: The defenders expected reward (or mitigated loss) when target t is protected and becomes the object of an attack.
- $U_d^u(t)$: The defenders expected loss when target t is left unprotected and is attacked.

- $U_{a,\theta}^c(t)$: The expected loss or diminished gain for an attacker of type θ when attacking target t while it is protected.
- $U_{a,\theta}^u(t)$: The expected gain for an attacker of type θ when a strike on an unprotected target t succeeds.

In Table 4.9 the global utility parameters α , β and γ are shown.

Symbol	Description	Value
α	Residual damage fraction when defended	0.65
β	Residual attacker reward when detected	0.10
γ	Defender loss scaling multiplier	10.0

Table 4.9: Global utility parameters

The values of α , β , and γ are based on modelling assumptions and have not been empirically calibrated. They constitute a plausible, yet unvalidated, parametrisation. Consequently, a sensitivity analysis of these parameters is conducted in later Sections.

Defender Utility

The defender's utility is primarily driven by the underlying risk and value of the maritime infrastructure within the study area. If a target is attacked while unprotected, the defender incurs the maximum baseline penalty, which is formulated as:

$$U_u^d(t) = -\gamma \cdot \text{AssetValue}(t) \quad (\text{undefended: full loss}) \quad (4.17)$$

where $\text{AssetValue}(t)$ is the value of the subsea infrastructure as defined in Table 4.4.

If a target is covered by an USV, the attack is completely or partially mitigated. To model this mathematically, the parameter $\alpha = 0.65$ is introduced, representing the proportion of damage sustained despite the defender's presence. Consequently, the defender's covered utility is defined as:

$$U_c^d(t) = -\alpha \cdot \gamma \cdot \text{AssetValue}(t) \quad (\text{covered: partial loss}) \quad (4.18)$$

The code implementation of these utility functions is as follows:

```

1 # Defender Base Values (Independent of Attacker type)
2 nodes_df['Defender_Loss'] = nodes_df['Asset_Value'] * defender_multiplier
3 nodes_df['U_d_u'] = -nodes_df['Defender_Loss']
4 nodes_df['U_d_c'] = -alpha * nodes_df['Defender_Loss']

```

Attacker Utility and Environmental Modifiers

The attacker's valuation of a target varies significantly depending on their specific profile $\theta \in \Theta$. Furthermore, a deterrence parameter $\beta = 0.1$ is applied to represent the proportion of reward an attacker receives if they attempt an assault while a patrol vessel is present.

For the accidental attacker profiles (θ_1 representing fishing trawlers and θ_2 representing commercial shipping), utility is not derived from damaging the cables, but rather from executing their standard operations. Therefore, their uncovered utilities are directly driven by historical vessel density. The utility function for a fishing vessel (θ_1) is defined as:

$$U_{\theta_1,u}^a(t) = \text{DensityFishing}(t) \quad (4.19)$$

where $\text{DensityFishing}(t)$ is the mean of the normalised fishing vessel density raster as seen in Figure 4.5a.

The utility for shipping vessels (θ_2) follows a similar formulation:

$$U_{\theta_2,u}^a(t) = \text{DensityShipping}(t) \quad (4.20)$$

where $\text{DensityShipping}(t)$ is the mean of the normalised cargo and tanker vessel density rasters as seen in Figures 4.5b and 4.5c.

In code this is implemented as:

```
1 nodes_df['Density_Shipping'] = [(c + t) / 2.0 for c, t in zip(cargo_means, tanker_means)]
```

High shipping density act as camouflage and a primary utility driver. If a patrol vessel is present, their operational freedom is nominally disrupted, reducing their expected utility to:

$$U_{\theta_1,c}^a(t) = \beta \cdot \text{DensityFishing}(t) \quad (4.21)$$

$$U_{\theta_2,c}^a(t) = \beta \cdot \text{DensityShipping}(t) \quad (4.22)$$

The implementation of shipping and fishing vessel utilities is defined as:

```
1 # -- Follower 1 (Fishing Trawlers) Utilities --
2 # Utility is purely derived from fishing grounds.
3 nodes_df['U_a_theta1_u'] = nodes_df['Density_Fishing']
4 nodes_df['U_a_theta1_c'] = nodes_df['Density_Fishing'] * beta # USV deterrence works
```

```

1 # -- Follower 2 (Cargo/Tanker) Utilities --
2 # Utility is purely derived from valid shipping transit lanes.
3 nodes_df['U_a_theta2_u'] = nodes_df['Density_Shipping']
4 nodes_df['U_a_theta2_c'] = nodes_df['Density_Shipping'] * beta # USV deterrence works

```

For the shadow fleet or state actor type (θ_3), utility arises directly from the intrinsic value of the compromised asset, $\text{AssetValue}(t)$, but is offset by the attacker's operational expenditures. In this framework, the operational cost associated with θ_3 comprises both the physical seabed substrate cost $c_a(t)$, as specified in Table 4.7, and an exposure penalty determined by surrounding vessel traffic. Dense shipping activity offers essential concealment for a strategic adversary, whereas operating in sparsely trafficked regions substantially heightens the likelihood of premature detection. Accordingly, the raw exposure penalty is given by $(1 - \text{DensityShipping}(t))$. To adjust the relative impact of exposure risk compared with the physical execution difficulty, this term is scaled by a factor of 0.5. This particular heuristic, named w_{exp} coefficient was chosen to upper-bound the exposure penalty, so that it remains influential while not excessively dominating the baseline substrate cost in the reference scenarios.

To correctly normalise these penalties with respect to differing asset magnitudes, this aggregated cost is applied as a proportional reduction factor on the asset's inherent value, instead of being treated as a constant offset. The resulting uncovered utility is then defined as a function by:

$$U_{a,\theta_3}^u(t) = \text{AssetValue}(t) \cdot (1 - [c_a(t) + w_{exp}(1 - \text{DensityShipping}(t))]) \quad (4.23)$$

If the target is covered by the defender, the expected benefit from the asset drops steeply due to the high likelihood of interception, yielding a covered (protected) utility of:

$$U_{a,\theta_3}^c(t) = (\beta \cdot \text{AssetValue}(t)) \cdot (1 - [c_a(t) + w_{exp}(1 - \text{DensityShipping}(t))]) \quad (4.24)$$

The code implementation of the shadow fleet utilities is defined as:

```

1 # -- Follower 3 (Shadow Fleet) Utilities --
2 # Strategic attacker seeking Asset_Value. Low shipping density drastically increases their
  ↪ exposure cost.
3 coverage_penalty = 1.0 - nodes_df['Density_Shipping']
4 nodes_df['Attacker_Cost_theta3'] = nodes_df['Attacker_Cost'] + (coverage_penalty * 0.5)
5
6 nodes_df['U_a_theta3_u'] = nodes_df['Asset_Value'] * (1.0 -
  ↪ nodes_df['Attacker_Cost_theta3'])
7 nodes_df['U_a_theta3_c'] = (beta * nodes_df['Asset_Value']) * (1.0 -
  ↪ nodes_df['Attacker_Cost_theta3'])

```

4.4.5 Optimisation Constant M

The optimisation constant M is introduced to map the attacker’s logical decisions into a linear programming formulation. This constant functions as a mathematical switch that activates or suppresses particular bounds depending on the attacker’s binary decision variable q_t^θ . To preserve linearity, the following constraint is imposed:

$$v^\theta \leq E_{U_a}(t, \theta) + M(1 - q_t^\theta) \quad (4.25)$$

The detailed use of this formulation is discussed in Section 4.4.6.

Although M could be chosen as an arbitrarily large constant, the code implementation instead computes a safe and tight bound for M rather than using a fixed value. This is achieved by scanning the payoff matrices to identify the maximum absolute utility in the game and then applying a conservative scaling factor. The implementation is as follows:

```
1 max_abs_utility = df[util_cols].abs().max().max()
2 M = float(max_abs_utility * 10.0)
```

Defining the constant M in this manner ensures robustness with respect to variations in the input utility dataset.

4.4.6 Decision Variables x_t, q_t^θ and Auxiliary Variables z_t^θ, v^θ

The mathematical formulation of the DOBSS is implemented using the PuLP Python package, which supports MILP. It represents an asymmetric security game in which a defender distributes a limited number of USVs across multiple infrastructure nodes, while anticipating a rational, utility-maximising response from several distinct follower types (θ). To enable the linear optimisation engine (in this case, the CBC solver) to compute the theoretical interaction, two sets of decision variables and two sets of auxiliary variables are introduced.

Defender’s strategy x_t

The variable $x_t \in [0, 1]$ denotes the expected marginal probability that the defender assigns the USV to patrol a given node or target t . If the defender were constrained to a pure strategy with $x_t \in \{0, 1\}$, a rational attacker could simply learn the deterministic patrol pattern and strike only undefended nodes, achieving a 100% success rate. By instead allowing x_t to take continuous values, the defender intentionally introduces randomness and thus becomes strategically unpredictable. The code implementation of the defender’s strategy is as follows:

```

1 x = pulp.LpVariable.dicts("x", nodes, lowBound=0, upBound=1, cat='Continuous')
2 # ...
3 prob += pulp.lpSum([x[t] for t in nodes]) <= k, "Resource_Constraint"

```

Since x_t encodes coverage probabilities, the aggregate of these probabilities over the entire infrastructure graph must not surpass the total number of available physical resources. In the implementation, the values of x_t are summed over all nodes and this total is constrained to be at most k .

Attacker's strategy q_t^θ

The variable $q_t^\theta \in \{0, 1\}$ precisely maps out the follower's pure strategy. It serves as a binary indicator: if $q_t^\theta = 1$, the follower of type θ chooses to attack target t , and if $q_t^\theta = 0$, the follower does not attack t . Under the standard assumptions of a SSE, although the leader commits to a mixed strategy, the follower computes expected utilities using the publicly known attack probabilities and then selects a single target that maximises their utility. Because the game incorporates a Bayesian element through the type parameter θ , the optimal defender coverage against a set of heterogeneous adversaries can be computed. The code implementation of the attacker's strategy is as follows:

```

1 q = pulp.LpVariable.dicts("q", (theta_types, nodes), cat='Binary')
2 # ...
3 for theta in theta_types:
4     prob += pulp.lpSum([q[theta][t] for t in nodes]) == 1, f"One_Target_{theta}"

```

Linearization variable z_t^θ and McCormick Envelopes

Since the expected payoff of the game is obtained by multiplying the defense probability of node x_t with the binary attack decision q_t^θ , the resulting formulation becomes a mixed-integer nonlinear program, which is hard to scale and to solve to global optimality. To address this, an auxiliary variable z_t^θ is introduced, defined by $z_t^\theta = x_t \cdot q_t^\theta$. By applying McCormick envelope constraints, this nonlinear term is bounded through a set of linear inequalities. Crucially, because the target selection variable is strictly binary, as defined by the decision variable $q_t^\theta \in \{0, 1\}$ and the probability of coverage is bounded within the unit interval, as defined by the variable $x_t \in [0, 1]$, the convex hull defined by the envelope collapses perfectly onto the nonlinear surface. Consequently, the linearisation is mathematically exact, introducing zero approximation error or a relaxation gap. This preserves the problem as a standard MILP, allowing the CBC solver to reliably find true global optimal solutions in finite time. The following code snippet showcases the implementation:

```

1 z = pulp.LpVariable.dicts("z", (theta_types, nodes), lowBound=0, upBound=1,
  ↪ cat='Continuous')
2
3 # McCormick Envelopes to linearize z = x * q
4 prob += z[theta][t] <= x[t], f"McCormick_1_{theta}_{t}"
5 prob += z[theta][t] <= q[theta][t], f"McCormick_2_{theta}_{t}"
6 prob += z[theta][t] >= x[t] - (1 - q[theta][t]), f"McCormick_3_{theta}_{t}"

```

The implementation essentially states, that if $q_t^\theta = 0$, then *McCormick_2* forces $z_t^\theta \leq 0$. Since *lowBound* = 0, z_t^θ definitively becomes 0. If $q_t^\theta = 1$, *McCormick_1* states that $z_t^\theta \leq x_t$ and *McCormick_3* states that $z_t^\theta \geq x_t - 0$. Therefore, z_t^θ is forced to equal x_t .

Attacker's maximum expected utility v^θ

The maximum expected utility that an attacker of type θ can obtain, given the defender's mixed strategy x_t , is denoted by the continuous variable v^θ . Within an SSE, the attacker is assumed to be fully rational. After observing how the defender allocates resources, the attacker computes the expected payoff associated with attacking each possible target and then selects the target that delivers the highest payoff. Consequently, v^θ formally represents this maximum achievable utility. This variable is essential for expressing the "best-response" constraints, which guarantee that the chosen target ($q_t^\theta = 1$) is indeed the one that maximizes the attacker's expected utility. The corresponding implementation in code is given by:

```

1 v = pulp.LpVariable.dicts("v", theta_types, cat='Continuous')
2
3 # Expected utility of attacking target t for attacker theta
4 E_U_a = U_a[theta]['u'][t] + x[t] * (U_a[theta]['c'][t] - U_a[theta]['u'][t])
5
6 # Big-M Strong Stackelberg Constraints
7 prob += v[theta] >= E_U_a, f"Best_Response_LB_{theta}_{t}"
8 prob += v[theta] <= E_U_a + M * (1 - q[theta][t]), f"Best_Response_UB_{theta}_{t}"

```

In the implementation, v^θ is introduced as a continuous variable indexed only by attacker types θ , since each type is associated with a single maximal payoff value, independent of the number of targets. Two constraints are specified in the code (*Best_Response_LB* and *Best_Response_UB*). The lower-bound constraint (*Best_Response_LB*) is formulated as $v^\theta \geq E_{U_a}$ for every target t . This ensures that v^θ is at least as large as the utility of the most profitable target. For example, if target A yields utility 5 and target B yields utility 10, then v^θ must be at least 10.

The upper-bound constraint (*Best_Response_UB*) employs the Big-M formulation to connect v^θ with the attacker's decision variable q_i^θ . When $q_i^\theta = 1$, the term involving M drops out, enforcing $v^\theta \geq E_{U_a}$. Together with the lower-bound constraint, this pins v^θ to exactly match the expected utility of the selected target. When $q_i^\theta = 0$, the right-hand side becomes $E_{U_a} + M$. Since M is chosen to be a sufficiently large constant, the resulting inequality effectively becomes $v^\theta \leq \infty$, so the constraint is inactive and does not constrain the true maximum payoff.

4.5 Spatiotemporal Routing via Markov Chain Synthesis

While the DOBSS formulation described above yields the SSE, its solution (x^*) is a stationary probability distribution. This distribution specifies the optimal fraction of time a USV should allocate to guarding each node, but it does not prescribe how the vehicle should move between nodes over time. To practically operate a USV, this static distribution must therefore be converted into a concrete, spatiotemporal routing policy.

This is achieved, by modelling the system as a discrete time Markov Chain, where the active nodes (targets where $x_i^* > 0$) define the state space. The goal is to synthesise a stochastic transition matrix P , where each element P_{ij} represents the probability of the USV moving from node i to node j in a standard time step Δt (in this case 1h)

Firstly, physical parameters of the USV and time parameter for the formulation are introduced in Table 4.10.

Symbol	Description	Value
v_{max}	Maximum operational velocity of the USV	30km/h
Δt	Discrete time step interval for the Markov Chain transitions	1h

Table 4.10: Markov Routing Parameters

The Markov Chain formulation introduces a second linear programming problem, subject to the following constraints:

- **Row Stochasticity:** From any node i , the USV must move to some node j (including the option of remaining at i itself, $j = i$). Consequently, the transition probabilities in each row must sum to 1:

$$\sum_j P_{ij} = 1 \quad \forall i \quad (4.26)$$

```

1  # 3. Constraints
2  for i in active_nodes:
3      # Row stochasticity: Must go somewhere (or stay)
4      prob += pulp.lpSum([P[i][j] for j in active_nodes]) == 1.0, f"Stochastic_{i}"

```

- **Strict Stationarity (Equilibrium Preservation):** This requirement links the Markov chain representation to the Stackelberg framework. The constructed transition matrix P must ensure that, if the USV uses these transition probabilities indefinitely, the resulting stationary distribution over the nodes coincides exactly with the optimal DOBSS allocation π , where π is the normalised form of x^* . This normalisation converts the optimal game-theoretic coverage probabilities into a valid steady-state distribution that sums to one, thereby specifying the precise fraction of total patrol time that a single USV should allocate to each node. The normalisation is formulated in Equation 4.27 and the constraint is defined in Equation 4.28.

$$\pi_i = \frac{x_i^*}{\sum_k x_k^*} \quad \forall i \quad (4.27)$$

```

1  # Normalize x* to represent a true probability distribution (sum = 1) for a single
   ↪ USV
2  total_x = sum(x_dict.values())
3  pi = {k: v / total_x for k, v in x_dict.items()}
4  active_nodes = list(pi.keys()) # We only bother routing between nodes that
   ↪ require coverage

```

$$\sum_i \pi_i P_{ij} = \pi_j \quad \forall j \quad (4.28)$$

```

1  for j in active_nodes:
2  # Stationarity: sum_i (pi_i * P_ij) = pi_j
3  prob += pulp.lpSum([pi[i] * P[i][j] for i in active_nodes]) == pi[j],
   ↪ f"Stationary_{j}"

```

- **Kinematic Constraints:** In practice, the USV cannot instantaneously traverse large distances between nodes. The set of feasible transitions is restricted by an upper bound on travel distance determined by the maximum speed v_{max} and the time step Δt . If the Euclidean distance $d(i, j)$ exceeds this reachable radius, the transition is infeasible and must be assigned zero probability:

$$P_{ij} = 0 \quad \text{if } d(i, j) > (v_{max} \cdot \Delta t) \quad (4.29)$$

```

1  # Enforce kinematic constraint: If too far, P_ij = 0
2  if d > max_dist_m:
3  prob += P[i][j] == 0, f"Kinematic_{i}_{j}"

```

Objective function and fuel minimisation

Several distinct transition matrices can fulfill the previously stated constraints, therefore, the model incorporates an additional operational criterion: conservation of fuel and energy. The objective function is defined as:

$$\min \sum_i \sum_j \pi_i \cdot P_{ij} \cdot d(i, j) \quad (4.30)$$

```
1 # 4. Objective: Minimize expected transition distance to conserve fuel/energy
2 # while strictly maintaining Stackelberg Equilibrium stationarity
3 prob += pulp.lpSum([pi[i] * P[i][j] * distances[i][j] for i in active_nodes for j in
  ↪ active_nodes]), "Minimize_Fuel"
```

By minimising the expected spatial displacement while rigorously enforcing P to satisfy the Stackelberg stationarity constraint, the resulting Markov chain generates a path that remains unpredictable to an attacker.

Simulation of anti-loitering rules

When directly sampling from the transition matrix P , there are cases in which the USV repeatedly samples the self-transition P_{ii} and thus remains at a single high-priority node for long durations. To promote movement throughout the patrol area, an anti-loitering heuristic is incorporated into the physical path simulation. If the vehicle resides at any node for longer than a prescribed consecutive time threshold (in this case 2 h), the self-transition probability P_{ii} is temporarily set to zero and the remaining outgoing transition probabilities are re-normalized. This enforces continued motion through the environment while still keeping the trajectory close to the equilibrium distribution specified by the underlying Markov model. In the code implementation, it first determines if the USV has hit the time limit:

```
1 banned = []
2 # Enforce constraint: loiter_count tracks extra consecutive hours beyond the initial
  ↪ hour at a node.
3 if loiter_count >= 1: # Already stayed 1 extra turn (total 2 hours), so must move away
  ↪ from this node.
4     banned.append(current)
```

Afterwards, the mathematical zeroing of P_{ii} is implemented:

```
1 # Temporarily zero out banned nodes
```

```

2   for i, t in enumerate(targets):
3       if t in banned_nodes:
4           probs[i] = 0.0
5
6       # Re-normalize if possible
7       prob_sum = sum(probs)
8       if prob_sum > 0:
9           probs = np.array(probs) / prob_sum
10      return np.random.choice(targets, p=probs)

```

4.6 Results

4.6.1 Adversary Target Selection

Running the DOBSS model produced an optimal patrol strategy for the USV (x^*), along with predicted rational responses for each modeled adversary. The operational setting includes two non-strategic "Nature" threats, these being fishing and commercial shipping accidents and one highly strategic, utility maximising adversary, this being the Shadow Fleet.

As summarised in Table 4.11, the model designates node t_{306} as the optimal target for both θ_1 and θ_2 . Since these threats correspond to non-strategic environmental accidents, their utility is normalised to 1.0. The node t_{306} denotes a geographic location with the highest mathematical probability of unintentional incidents, influenced by dense shipping lanes or regions of intensive fishing. Collectively, these nature driven threats contribute 98% of the total expected incident probability.

In contrast, θ_3 behaves as a rational, utility maximising agent. Taking into account the sea substrate and the asset value of the subsea infrastructure, θ_3 isolates node t_{156} to maximise its expected reward ($U_a \approx 16.07$).

Adversary	Threat Profile	Prior $p(\theta)$	Target	Expected Utility (U_a)
θ_1	Fishing Accident (Nature)	0.86	t_{306}	1.0000
θ_2	Shipping Accident (Nature)	0.12	t_{306}	1.0000
θ_3	Strategic Sabotage (Strategic)	0.02	t_{156}	16.0736

Table 4.11: Bayesian Adversary Profiles and Stackelberg Best Responses

Figure 4.7 provides a graphical representation of Table 4.11. The figure depicts the three-dimensional relationship among an adversary's prior probability ($p(\theta)$, x-axis), its expected attack utility (U_a , y-axis), and the corresponding optimal USV coverage probability (x^* , encoded by bubble size). The model highlights a central Stackelberg style deterrence mechanism: it assigns no patrol effort to high-frequency, low-impact environmental

noise (θ_1 and θ_2), while concentrating defensive resources to deter the low-frequency, high-impact strategic sabotage threat (θ_3).

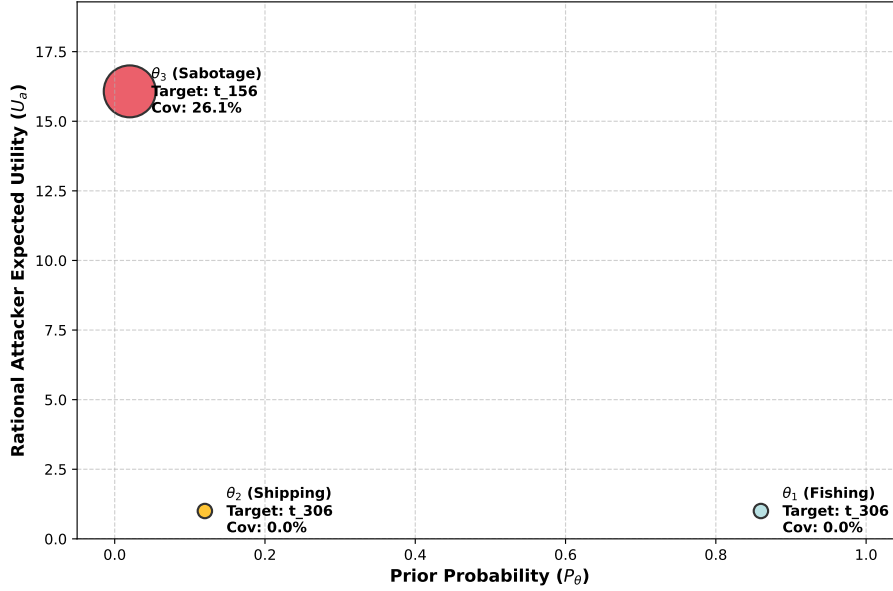


Figure 4.7: Threat Landscape and Algorithmic Patrol Response

4.6.2 Optimal Stationary Distribution x^*

The solver described in Section 4.3 yields a sparse solution. Among the 400 specified nodes, only 7 are assigned a non-zero coverage probability. The coverage probabilities for these 7 nodes are presented in Table 4.12.

Rank	Node	Coverage Probability (x^*)	Percentage
1	t_{156}	0.2607	26.07%
2	t_{251}	0.2296	22.96%
3	t_{157}	0.2008	20.08%
4	t_{271}	0.1195	11.95%
5	t_{270}	0.1192	11.92%
6	t_{176}	0.0525	5.25%
7	t_{158}	0.0177	1.77%
-	All Others	0.0000	0.00%

Table 4.12: Optimal Stationary USV Patrol Distribution (x^*)

For enhanced clarity and interpretability of the results, Figure 4.8 displays a heatmap overlaid on the study area.

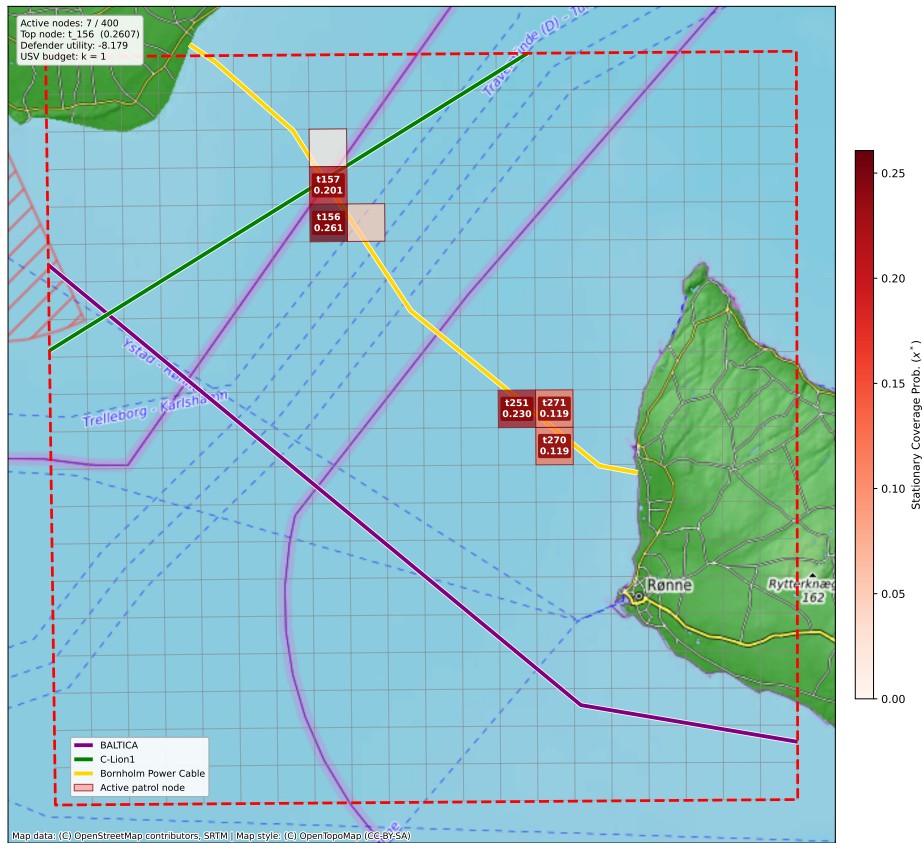


Figure 4.8: Optimal Stationary USV Patrol Distribution x^*

The resulting patrol distribution (x^*) in Table 4.12 highlights several advanced game-theoretic behaviours inherent to the DOBSS formulation, specifically regarding strategic resource allocation and adversary displacement.

Although the model identifies t_{156} as the optimal target for the strategic adversary (θ_3), the defender does not assign this node full (100%) coverage. Within the Stackelberg framework, perfectly predictable behavior is computationally discouraged: committing all defensive resources to t_{156} would simply induce a fully informed adversary, who observes the defender's stationary strategy, to shift the attack to the next unprotected yet valuable node, such as t_{251} . Consequently, the algorithm produces a mixed strategy. By protecting t_{156} at 26.07%, t_{251} at 22.96%, and t_{157} at 20.08%, the model effectively flattens the attacker's expected utility surface, inducing indifference among the primary cluster of high-value targets.

This allocation simultaneously addresses the risk of security displacement. By strongly guarding the top three priority nodes, the algorithm anticipates that a rational adversary will migrate toward secondary targets. To preempt this, secondary nodes such as t_{271} (11.95%) and t_{270} (11.92%) receive partial coverage. This secondary defensive layer ensures

that no highly valuable node within the network remains completely unprotected, thereby preventing obvious exploitative opportunities.

Under the constraint of a single USV ($k = 1$, with $\sum x \leq 1$), the defender must tolerate a nonzero level of residual vulnerability. The 0% coverage assigned to t_{306} , which accounts for 98% of expected environmental incidents, reflects a mathematically justified acceptance of background risk. The model infers that prioritizing the deterrence of deliberate, high-impact sabotage leads to a higher expected preservation of systemic utility than diverting scarce security resources to monitor largely unavoidable environmental accidents.

Whereas the attacker utilities (U_a) characterize each adversary's isolated, rational choice of target, the primary objective of the Stackelberg formulation is to maximize the defender's overall expected utility (U_d). Solving the MILP yields an optimal patrol policy x^* , which achieves a global defender utility of -8.1788.

Within this infrastructure protection setting, the defender's utility is expressed as a negative quantity to represent the mathematically expected residual loss to the network. This scalar aggregates, as a weighted average over all Bayesian priors ($p(\theta)$), both the unmitigated, high-frequency baseline damage arising from environmental incidents at t_{306} and the strongly mitigated, low-frequency strategic damage at t_{156} .

Because the DOBSS framework computes a Strong Stackelberg Equilibrium under the constraint of a single USV ($k = 1$), this utility level corresponds to the theoretical upper bound on the system's security performance. Any departure from the optimal mixed strategy in Table 4.12, for example, by disproportionately focusing on background noise or enforcing a deterministic, easily anticipated patrol would systematically increase exploitable weaknesses for the strategic adversary, pushing the expected utility to a more negative value. Consequently, $U_d = -8.1788$ constitutes the reference point for optimal risk reduction within the specified area.

4.6.3 60-Hour Patrol Simulation

Utilising the optimised transition matrix P , a 60-hour patrol simulation was conducted to project the theoretical model onto a realistic operational timeline. An operational anti-loitering constraint was incorporated to ensure the USV does not remain stationary indefinitely: if the USV occupies the same node for two consecutive hours (one hour of initial arrival followed by one hour of loitering), the self-transition probability is temporarily set to zero, compelling the vehicle to relocate.

The stochastic simulation output consists of a time-ordered, stepwise waypoint schedule that specifies the USV's position for all $t \in [0, 60]$.

4.6.4 Results and Spatial Visualization

To evaluate the simulation, the chronological schedule was translated into an edge-weighted flow and node dwell-time map, shown in Figure 4.9.




Figure 4.9: Red markers depict the node dwell-time, with the marker size scaled relative to the frequency of visits. The numerical labels denote the exact number of hours the USV spent at each node. Darker, thicker blue lines indicate frequently traversed maritime corridors, representing the highest-traffic routes derived from the Markov continuous transitions

The node diameters are scaled in direct proportion to the dwell time, offering a visual confirmation that the operational USV focuses on the high-risk nodes, identified in Figure 4.8. An illustrative portion of the hour-by-hour chronological routing is provided in Table F.1 in Appendix 6.

4.7 Sensitivity Analysis

The optimal policy produced by DOBSS is only as credible as the assumptions that generate it, and only as useful as its margin over simpler allocation rules. This chapter therefore proceeds in two parts. It first quantifies the value of the equilibrium solution relative to two intuitive baseline strategies under the fully rational attacker assumption, establishing

the headline utility improvement. It then tests whether that improvement, and the patrol corridor that produces it, persist under plausible modelling error, identifying which inputs materially drive the solution and flagging conditions under which the recommended allocation could fail or shift to a different set of nodes.

The robustness assessment examines five dimensions of uncertainty in the DOBSS result ($U_d = -8.18$, seven active patrol nodes along the Bornholm power cable corridor): attacker rationality, threat-type priors, detection and deterrence parameters (α, β, γ), risk-score construction, and operational routing constraints. For each dimension, the L_1 drift of the coverage allocation $\|\mathbf{x}_{\text{new}} - \mathbf{x}_{\text{base}}\|_1 \in [0, 2]$ and a qualitative assessment of corridor stability serve as robustness indicators; low drift classifies an input as a non-critical assumption, while large drift or active-set changes classify it as a structural assumption. Figure 4.10 presents the order of testing.

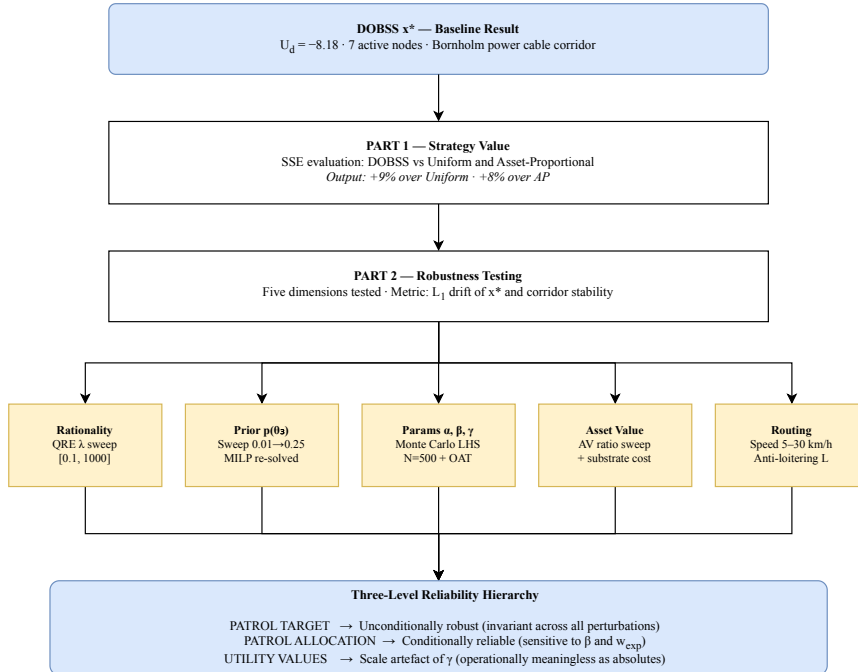


Figure 4.10: Sensitivity analysis flowchart

4.7.1 Strategy Comparison: DOBSS versus Proportional and Uniform Baselines

The value of the equilibrium solution is quantified against two intuitive baselines, evaluated under the same Strong Stackelberg Equilibrium assumption used throughout the core pipeline. The *Uniform* (UNI) strategy allocates equal coverage to every node in the operational area, $x_t^{\text{UNI}} = k/|\mathcal{T}| = 1/400 = 0.0025$. The *Asset-Proportional* (AP) strategy allocates

coverage in direct proportion to each node’s asset value, normalised to the budget $k = 1$:

$$x_t^{\text{AP}} = \frac{k \cdot \text{AssetValue}(t)}{\sum_{t' \in \mathcal{T}} \text{AssetValue}(t')} \quad (4.31)$$

AP distributes effort across all 72 nodes carrying cable infrastructure (maximum $x_{156}^{\text{AP}} = 0.036$), whereas DOBSS concentrates coverage on 7 nodes (maximum $x_{156}^* = 0.261$).

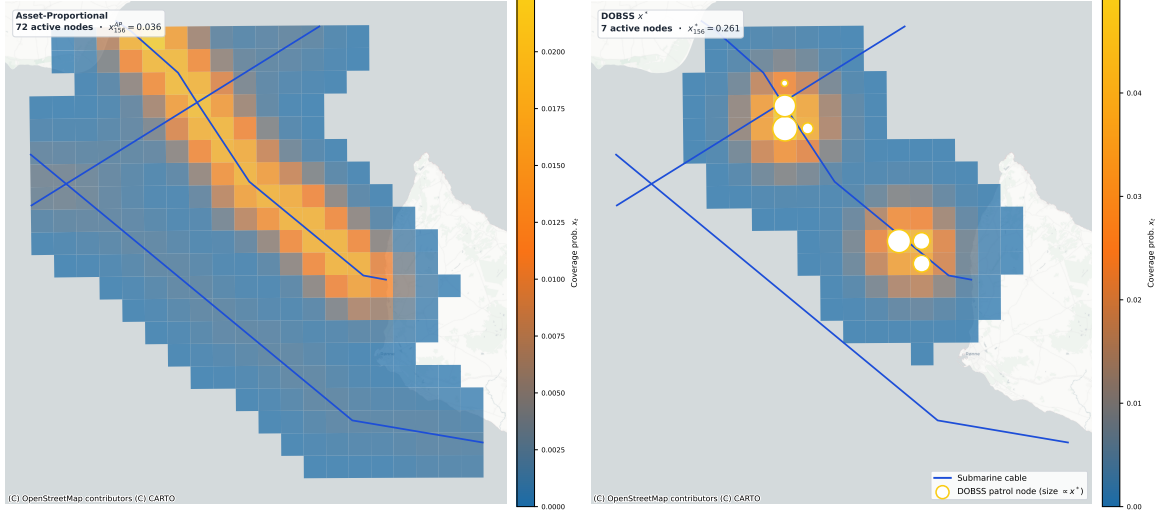


Figure 4.11: AP and DOBSS comparison

Each strategy is evaluated by computing the fully rational best response of every threat type, with ties broken in the defender’s favour, and aggregating the resulting defender utility over the Bayesian prior. The results are reported in Table 4.13.

Strategy	x_{156}	Total U_d	vs. DOBSS
Uniform	0.003	−8.992	−9.0 %
Asset-Proportional	0.036	−8.886	−8.0 %
DOBSS x^*	0.261	−8.179	—

Table 4.13: Expected defender utility under three strategies, evaluated under the Strong Stackelberg Equilibrium. x_{156} is the coverage probability assigned to the Bornholm power cable node t_{156} .

The comparison reduces to a single degree of freedom. Under fully rational attack, types θ_1 and θ_2 both target t_{306} , the peak vessel-density node, under all three strategies; because t_{306} carries zero asset value, its contribution to defender loss is exactly zero. The entire utility difference between strategies is therefore determined by coverage of the θ_3 target t_{156} : Uniform assigns $x_{156} = 0.003$, AP assigns 0.036, and DOBSS assigns 0.261. This concentration, a factor of seven above AP and two orders of magnitude above Uniform,

reduces expected defender loss by 9.0% relative to Uniform and 8.0% relative to AP. The advantage arises from anticipating the attacker’s best response and committing coverage to the single strategically dominant target rather than spreading it by asset value alone; the spatial contrast is visible in Figure 4.8, and the per-type decomposition in Figure 4.12.

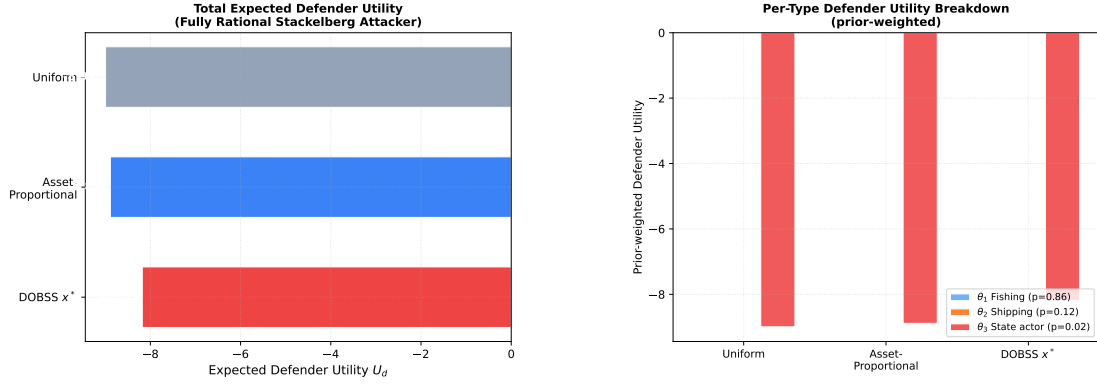


Figure 4.12: Strategy comparison under Strong Stackelberg Equilibrium. **Left:** total expected defender utility for Uniform, Asset-Proportional, and DOBSS. **Right:** prior-weighted per-type utility breakdown showing that θ_1 and θ_2 contribute zero loss under fully rational attack across all three strategies; the entire difference is driven by the θ_3 component.

4.7.2 Rationality Sensitivity

The Stackelberg equilibrium presumes that each follower deterministically chooses the utility-maximizing target. In contrast, the Quantal Response Equilibrium (QRE) model [140] relaxes this assumption by specifying that, at each node t , the attack probability is proportional to $\exp(\lambda \cdot \mathbb{E}[U_a^\theta(t)])$. Here, $\lambda \geq 0$ controls the degree of rationality: when $\lambda = 0$, the attacker behaves as a uniform random chooser, while $\lambda \rightarrow \infty$ converges to the fully rational best response. With x^* held fixed, we evaluate the prior-weighted defender expected loss over $\lambda \in [0.1, 1000]$.

λ	Regime	$\mathcal{L}_{\text{total}}$	\mathcal{L}_{θ_1}	\mathcal{L}_{θ_2}	\mathcal{L}_{θ_3}
0.5	Near-random	-38.818	-26.386	-3.690	-8.742
2.0	Human-like	-38.246	-25.789	-3.652	-8.806
5.0	Human-like	-29.344	-17.771	-2.788	-8.785
10.0	Human-like	-9.286	-0.430	-0.081	-8.776
50.0	Expert/rational	-8.780	-0.000	-0.000	-8.780
200.0	Expert/rational	-8.810	-0.000	-0.000	-8.810

Table 4.14: Defender expected loss $\mathcal{L}(\lambda)$ at key rationality checkpoints. Per-type values are prior-weighted. x^* is held fixed throughout.

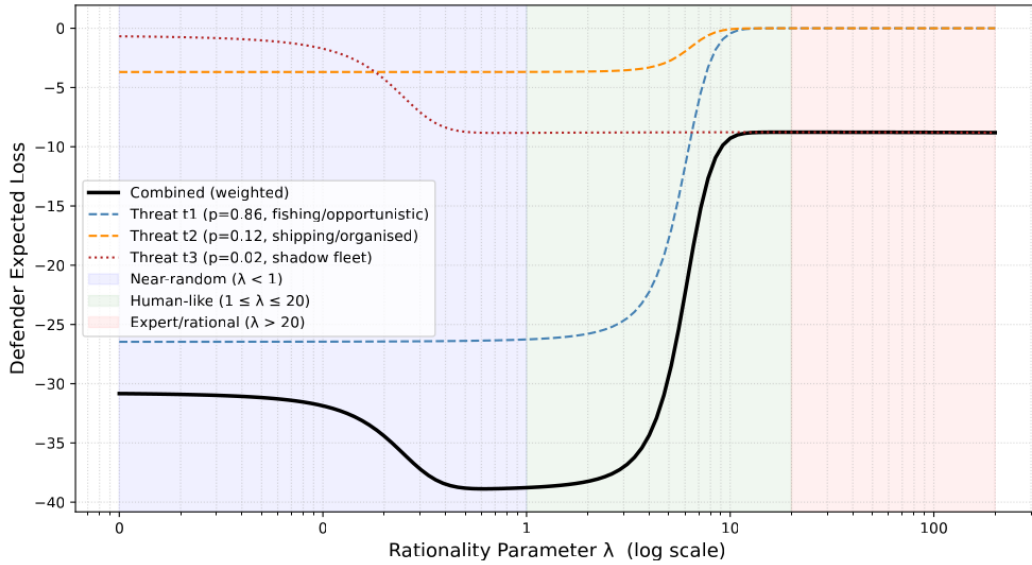


Figure 4.13: QRE sensitivity analysis (defender expected loss vs. attacker rationality, x^* held fixed)

Looking at the results of the sensitivity analysis, four findings can be seen. First, the worst-case loss occurs at near-random rationality $\lambda \approx 0.5$ (total loss ≈ -38.8), worse than the fully rational attacker (-8.81 at $\lambda = 200$). The mechanism is that a partially informed adversary avoids some covered nodes without being concentrated enough for x^* to be effective, whereas the rational attacker’s utility-maximising concentration is precisely what the Stackelberg design targets. It should be noted, however, that this worst-case partly reflects a limit of the Stackelberg framing rather than a pure sensitivity result: at $\lambda \approx 0.5$ the softmax distribution is nearly uniform over all 400 targets, making the attacker effectively random. A Stackelberg equilibrium is designed to exploit attacker rationality; against a near-random adversary the patrol has limited deterrent effect regardless of its design, because there is no strategic concentration to disrupt. Empirical QRE calibration in laboratory normal-form games typically yields $\lambda \in [0.5, 3]$ for human subjects [140, 141]; operational scenario analysis for sabotage threats would typically place intent-driven actors (state or organised-crime) in the $\lambda \geq 2$ zone where the DOBSS patrol is most effective.

Second, as λ exceeds approximately 10, the contributions of θ_1 and θ_2 to total loss collapse to near zero (-0.430 and -0.081 respectively at $\lambda = 10$, from -26.386 and -3.690 at $\lambda = 0.5$), confirming that the DOBSS strategy successfully deters the threat types representing 98% of prior probability at equilibrium.

Third, despite a prior of only $p(\theta_3) = 0.02$, the state actor contribution accounts for all residual loss at high λ , ranging from -8.74 at $\lambda = 0.5$ to -8.81 at $\lambda = 200$. Because x^* is derived under a Bayesian objective weighting θ_3 at only 2%, the strategy under-allocates coverage against a highly capable state-level actor, which is a known structural limitation

of DOBSS with low-prior high-capability types.

Fourth, total loss at $\lambda = 50\text{--}200$ stabilises near -8.78 to -8.81 , which does not fully recover the Stackelberg equilibrium objective of -8.18 . The residual gap reflects the SSE tie-breaking convention: at equilibrium, θ_3 is nearly indifferent between t_{156} and t_{158} (expected attacker utility difference of < 0.001), and the SSE assigns the attack to t_{156} — the node more favourable to the defender. Under QRE, the softmax amplifies even this marginal advantage, concentrating θ_3 on t_{158} at high λ , for which the prior-weighted defender loss is approximately -9.74 . This SSE tie-breaking sensitivity is an inherent property of solutions that require near-indifference to achieve their equilibrium value.

4.7.3 Prior Sensitivity

The Bayesian priors $p(\theta_1) = 0.86$, $p(\theta_2) = 0.12$, $p(\theta_3) = 0.02$ are one of the most subjective inputs in the model. The state-actor prior $p(\theta_3)$ is swept from 0.01 to 0.25 in 25 uniform steps, redistributing the remaining mass between θ_1 and θ_2 while preserving their baseline ratio; the full DOBSS MILP is re-solved at each prior combination. The results are presented in Table 4.15.

$p(\theta_3)$	$p(\theta_1)$	$p(\theta_2)$	Def. Utility	θ_3 Target
0.010	0.869	0.121	-4.089	t_{156}
0.020	0.860	0.120	-8.179	t_{156}
0.030	0.851	0.119	-12.268	t_{156}
0.050	0.834	0.116	-20.447	t_{156}
0.100	0.790	0.110	-40.894	t_{156}
0.150	0.746	0.104	-61.341	t_{156}
0.200	0.702	0.098	-81.788	t_{156}
0.250	0.658	0.092	-102.235	t_{156}

Table 4.15: Prior sensitivity sweep: defender utility and state-actor target at key $p(\theta_3)$ values. **Bold** marks the canonical baseline.

The preceding prior sensitivity sweep yields two key observations. First, the coverage allocation \mathbf{x}^* remains fully stable throughout the range $p(\theta_3) = 0.01$ to $p(\theta_3) = 0.25$: both the set of active nodes and the associated marginal probabilities are unchanged under all previous reweightings, with t_{156} consistently serving as the θ_3 rational target node. Second, defender utility decreases at an approximately constant rate of -4.09 for each unit increment in $p(\theta_3)$. This apparent linearity is, however, purely a consequence of the models algebraic structure rather than an empirical regularity: under linear utility aggregation with a fixed strategy, $U_d = \sum_i p(\theta_i)U_{d,\theta_i}$, so varying $p(\theta_3)$ while holding \mathbf{x}^* fixed necessarily yields exact proportional scaling. The substantive finding is that \mathbf{x}^* does not adjust in response; the linear decline itself adds no information beyond what is implied

by the models form. Nonetheless, the drop in utility indicates that absolute defensive effectiveness worsens markedly as the probability of a θ_3 adversary increases, underscoring the need for additional deterrence layers beyond a single-USV patrol.

4.7.4 Monte Carlo Parameter Sensitivity and PRCC

The utility model parameters α , β , and γ constitute the least certain inputs in the case study; consequently, prioritizing a sensitivity analysis on these parameters is essential for assessing the robustness of the DOBSS model. Their uncertainty is propagated using $N = 500$ Latin Hypercube Sampling (LHS) [142] realizations from independent uniform distributions, as specified in Table 4.16. For each sample, the complete utility matrix is recalculated and the DOBSS MILP is solved again. Two correlation measures are reported: the Spearman rank correlation ρ_S (capturing the raw monotonic relationship) and the Partial Rank Correlation Coefficient (PRCC) [143], which quantifies each parameter’s unique influence after adjusting for the others through rank-based regression.

Parameter	Description	Baseline	Lower	Upper
α	Residual damage fraction when covered	0.65	0.30	0.95
β	Residual attacker reward when detected	0.10	0.01	0.30
γ	Defender loss scaling multiplier	10.0	4.0	20.0

Table 4.16: Monte Carlo sampling configuration for utility parameters.

Parameter	Defender Utility		Active Nodes		Top Node Coverage	
	ρ_S	PRCC	ρ_S	PRCC	ρ_S	PRCC
α	-0.123*	-0.937*	+0.035	+0.055	-0.000	-0.009
β	+0.008	+0.287*	-0.834*	-0.721*	+1.000*	+1.000*
γ	-0.989*	-0.998*	+0.037	+0.069	+0.009	-0.001

Table 4.17: ρ_S and PRCC between utility parameters and key outputs ($N = 500$ LHS samples, all Optimal). Asterisk (*) denotes $p < 0.05$.

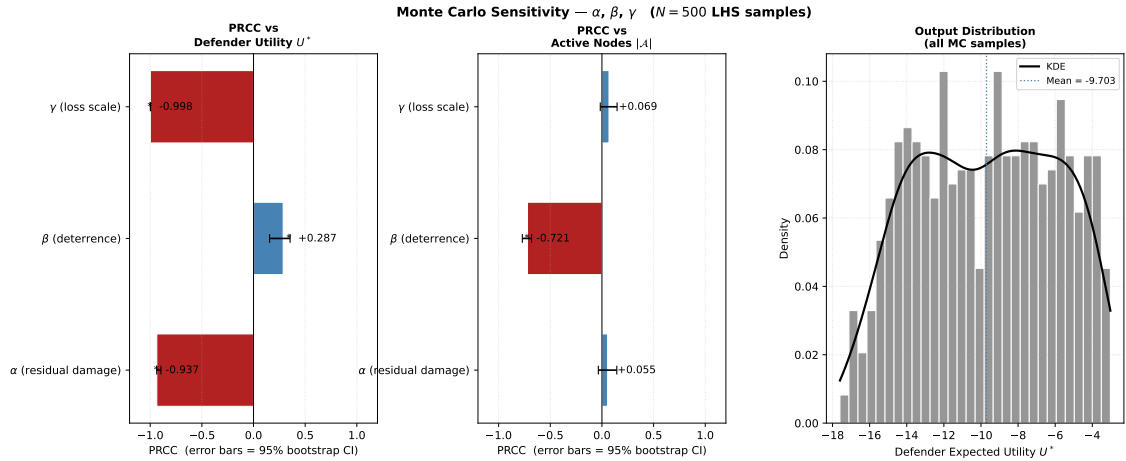


Figure 4.14: Monte Carlo sensitivity analysis ($N = 500$, Latin Hypercube Sampling). **Left** – PRCC tornado chart for defender expected utility U^* . **Centre** – PRCC tornado chart for active patrol nodes $|\mathcal{A}|$. **Right** – KDE of the U^* distribution with baseline and sample mean marked.

From the analysis, four key results are observed. First, γ almost completely governs the utility of the defender, as shown from $\rho_S = -0.989^*$, $\text{PRCC} = -0.998^*$: the loss-scale multiplier appears as a linear factor in both U_u^d and U_c^d , uniformly scaling all terms of the utility of the defender. Consequently, absolute utility magnitudes across configurations are driven almost entirely by γ , and relative strategy rankings are the most informative quantity. Second, the impact of α is almost fully obscured in the simple rank correlation, seen from $\rho_S = -0.123^*$, but becomes the second most influential driver, as presented by $\text{PRCC} = -0.937^*$ once the dominant marginal contribution of γ is removed through partial regression. This illustrates why the partial rank correlation is essential: γ explains so much variance that the independent effect of α is hidden in the raw statistic, despite the fact that α controls how much the defender gains from covering a node. Third, β perfectly explains the coverage of the top-node, shown by both ρ_S and PRCC being equal to $+1.000^*$: the attackers retained reward under coverage is the only factor that determines the probability assigned to the highest-priority node, and the total variation in the range $[0.248, 0.296]$ is attributable to changes in β . Lastly, the number of active nodes is almost entirely determined by β , as shown by $\rho_S = -0.834^*$ and $\text{PRCC} = -0.721^*$: for $\beta \in [0.01, 0.30]$ in the Monte Carlo sample, the active set ranges between 6 and 7 nodes, with mean of 6.6 and SD of 0.5, with larger values β concentrating the patrol budget on fewer nodes to maintain equilibrium. The effects of α and γ on active-set size are negligible, shown by $|\text{PRCC}| < 0.08$, both non-significant, reinforcing that the Bornholm power cable corridor is a geometric characteristic of the asset landscape rather than a consequence of the utility parameterisation.

These correlation patterns highlight a clear differentiation in the way the parameters function within the game model. The loss-scale multiplier γ serves purely as a global scaling parameter, since it enters both the covered and uncovered defender utility functions as

a uniform linear factor, its almost perfect PRCC ($\rho_S = -0.989^*$, $\text{PRCC} = -0.998^*$) is therefore mathematically predictable and does not indicate any change in strategic behaviour.

Conversely, the attacker’s residual reward multiplier β constitutes the primary strategic parameter of the security game. By modulating the fraction of utility an adversary preserves when a node is patrolled, β reshapes the relative payoff gap between defended and undefended targets. As a result, it shifts the locations of the adversary’s indifference thresholds, thereby directly determining both the number of nodes that remain active ($\text{PRCC} = -0.721^*$) and the allocation mass concentrated on the highest-priority nodes ($\text{PRCC} = +1.000^*$).

This contrast highlights that, whereas γ controls the absolute scale of the reporting coefficients, β is the operative parameter that strategically reconfigures the patrol pattern.

Another finding concerns output uncertainty. Defender utility ranges from -17.603 to -3.026 in the Monte Carlo sample, presented in Table 4.18, a spread of 14.6 units much larger than the gain from optimising over baselines, with < 0.7 units from strategy comparison. The bootstrap 95% confidence intervals ($B = 1000$ resamples) confirm the stability of the dominant effects: PRCC of γ in U^* is $[-0.998, -0.991]$, PRCC of β in active nodes is $[-0.721, -0.682]$. The near-zero effects of α and γ on active nodes, with PRCCs of $+0.055$ and $+0.069$ and CIs straddling zero, are confirmed non significant. Operational assessments should report defender utility as an interval (IQR $[-12.94, -6.48]$), rather than a point estimate, and strategy comparisons should fix γ to avoid conflating scale effects with strategic differences.

Output	Mean	Std	Min	Max	Q25	Q75
Defender Utility	-9.703	3.782	-17.603	-3.026	-12.939	-6.479
Active Nodes	6.6	0.5	6	7	6.0	7.0
Top Node Coverage	0.270	0.014	0.248	0.296	0.258	0.282

Table 4.18: Descriptive statistics of key outputs across $N = 500$ Monte Carlo samples

4.7.5 One-at-a-Time Parameter Sensitivity

One-at-a-time sweeps vary each of α , β , γ , and the asset-value heterogeneity scalar σ_{AV} individually across their plausible ranges, recording defender utility, active node count, L_1 drift from x_{base}^* , and the θ_3 attack target.

Table 4.19: OAT sweep summary: defender utility, active node count, L_1 drift of \mathbf{x}^* from baseline, and θ_3 attack target at selected parameter values. Baseline row in **bold**. $L_1 = 0$ indicates \mathbf{x}^* is unchanged; $L_1 > 0$ indicates strategy reallocation.

Param	Value	Def. Util.	Active	L_1 drift	Δ_{t_3}	θ_3 target
α	0.05	-6.771	7	0.000	-18.9	t_{156}
	0.30	-7.358	7	0.000	-18.9	t_{156}
	0.65	-8.179	7	0.000	-18.9	t_{156}
	0.95	-8.883	7	0.000	-18.9	t_{156}
β	0.00	-8.221	8	0.060	-21.0	t_{156}
	0.10	-8.179	7	0.000	-18.9	t_{156}
	0.50	-7.912	5	0.336	-10.5	t_{156}
	0.70	-7.625	3	0.618	-6.3	t_{156}
	0.90	-6.901	3	1.001	-2.1	t_{156}
	1.00	-5.850	1	1.479	0.0	t_{156}
γ	1.0	-0.818	7	0.000	-18.9	t_{156}
	10.0	-8.179	7	0.000	-18.9	t_{156}
	50.0	-40.894	7	0.000	-18.9	t_{156}
	100.0	-81.788	7	0.000	-18.9	t_{156}
σ_{AV}	0.0 (uniform)	-3.281	11	1.636	-7.6	t_{260}
	0.5	-5.651	6	0.065	-13.1	t_{156}
	1.0 (baseline)	-8.179	7	0.000	-18.9	t_{156}

DOBSS Parameter Sensitivity Analysis
 ($K = 1$, baseline: $\alpha = 0.65$, $\beta = 0.10$, $\gamma = 10$, $\sigma_{AV} = 1$)

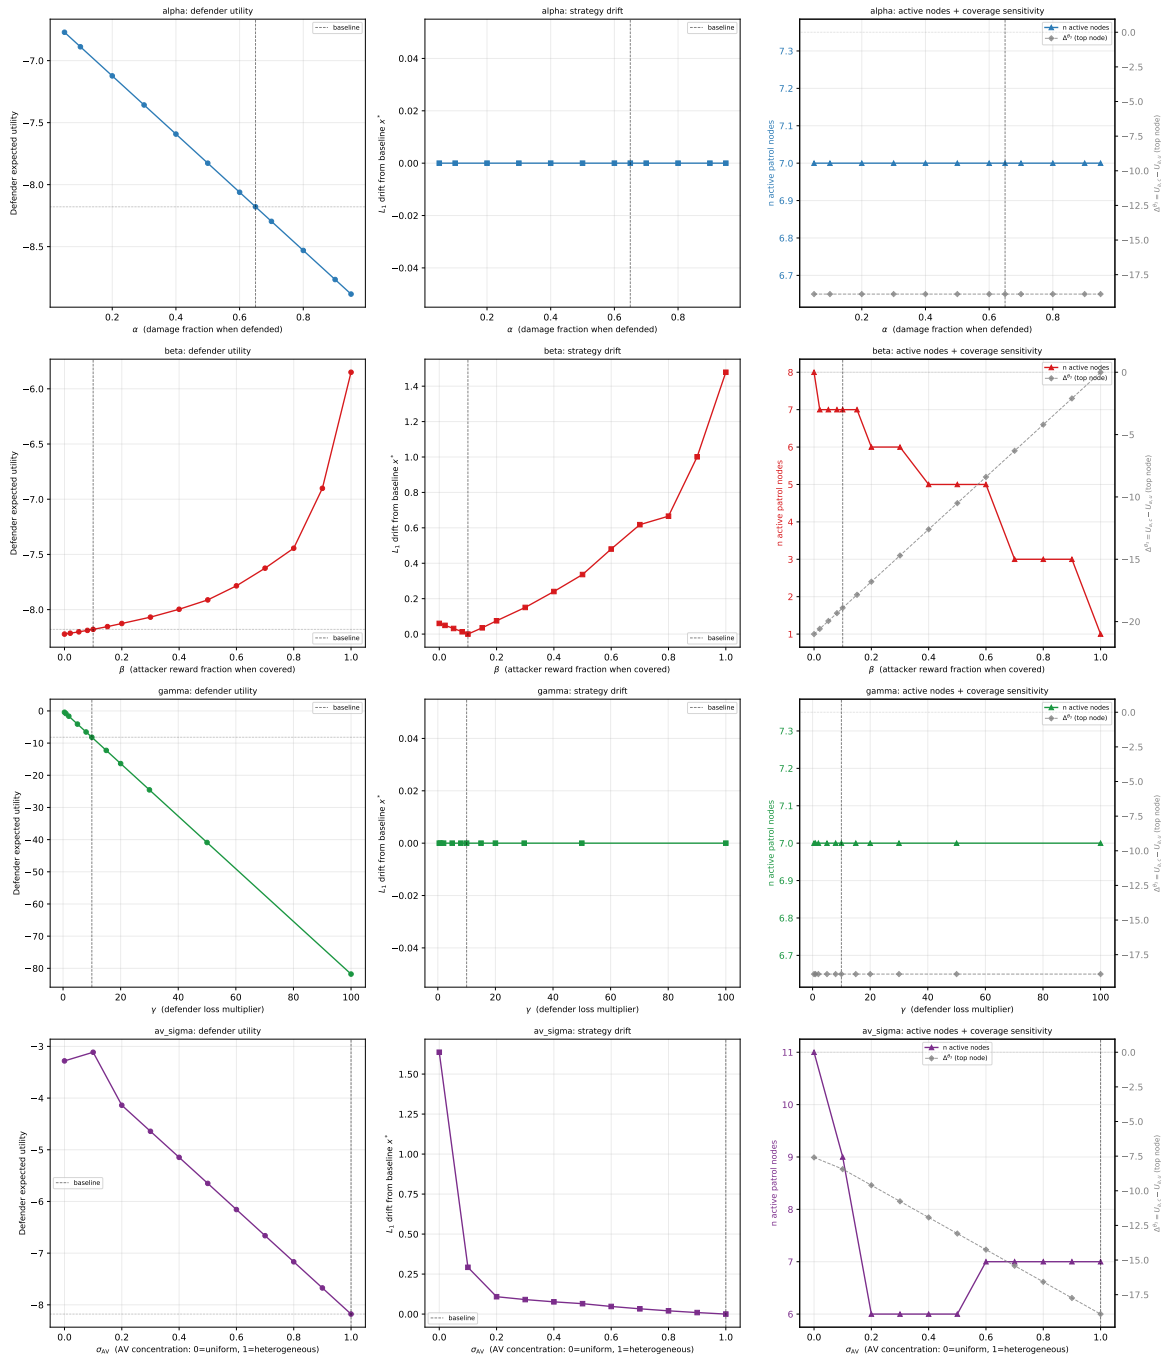


Figure 4.15: OAT parameter sensitivity sweeps. Each panel varies one parameter across its full range with all others held at baseline. **Top row:** α and β sweeps (defender utility on left axis, L_1 drift on right axis). **Bottom row:** γ and σ_{AV} sweeps. Dashed verticals mark the baseline value. The β sweep shows a monotonic contraction of the active set from 8 nodes at $\beta = 0$ to a single node at $\beta = 1$.

Looking at the results of the analysis, four observations can be seen. First, the choice of α is not critical. This is evident from the fact that the defenders utility declines approximately linearly from -6.771 to -8.883 as α varies over $[0.05, 0.95]$, covering a total span of 2.11 utility units, while L_1 remains fixed at 0 and the target θ_3 stays equal to t_{156} . Second, β causes a gradual but monotonic regime change in \mathbf{x}^* . At $\beta = 0.00$, the active set expands slightly to eight nodes ($L_1 = 0.060$); as β rises the active set contracts steadily — seven nodes at baseline, six at $\beta = 0.20$, five at $\beta = 0.40$, three at $\beta = 0.70$ — until a single node remains at $\beta = 1.00$ ($L_1 = 1.479$). The θ_3 target remains t_{156} throughout the entire sweep, confirming the power cable corridor as the dominant strategic anchor regardless of the attacker’s retained reward. For the baseline $\beta = 0.10$, the system operates in the stable, multi-node regime. Third, γ is a pure scale parameter: defender utility follows $U^* = -0.8179 \times \gamma$ exactly ($L_1 = 0$ throughout), consistent with the near-perfect PRCC of the Monte Carlo analysis. Fourth, the σ_{AV} result requires careful interpretation. At $\sigma_{AV} = 0$ (all nodes assigned the uniform asset-value mean), the active set expands to 11 nodes ($L_1 = 1.636$) and the θ_3 target shifts to t_{260} . This is not in contradiction with the Monte Carlo finding that the cable corridor is a geometric property of the asset landscape; both statements are correct but refer to different scopes. The Monte Carlo analysis holds σ_{AV} at the baseline heterogeneous landscape and varies α , β , γ , finding the corridor insensitive to those utility parameters. The OAT sweep varies σ_{AV} itself, showing that when spatial differentiation is removed, the corridor collapses. The cable corridor is therefore a consequence of the *empirical heterogeneity* of the asset landscape, specifically the Bornholm power cable’s $AV = 45$ dominating all other infrastructure, not a geometric invariant of the grid topology.

4.7.6 Risk-Score Weighting Sensitivity

The utility parameters in Sections 4.7.4 and 4.7.5 govern attacker and defender payoffs given fixed asset values; a further sweep varies the spatial risk scores that drive attacker utility directly. Two independent sweeps are run: Sweep A varies AV_{power} across $\{4, 6, 8, 10, 12, 16, 20, 30, 45, 60\}$ with $AV_{\text{fibre}} = 4$ fixed, testing when patrol shifts from the power cable corridor to the fibre cables; Sweep B varies the state-actor exposure-penalty weight $w_{\text{exp}} \in [0, 1]$ (baseline 0.50) with $AV_{\text{power}} = 45$ fixed. The full DOBSS MILP is re-solved at each point and L_1 drift from the baseline \mathbf{x}^* is recorded.

Risk-Score Weighting Sensitivity
Panel A-B: power-cable asset value ratio; Panel C-D: state-actor exposure-penalty weight

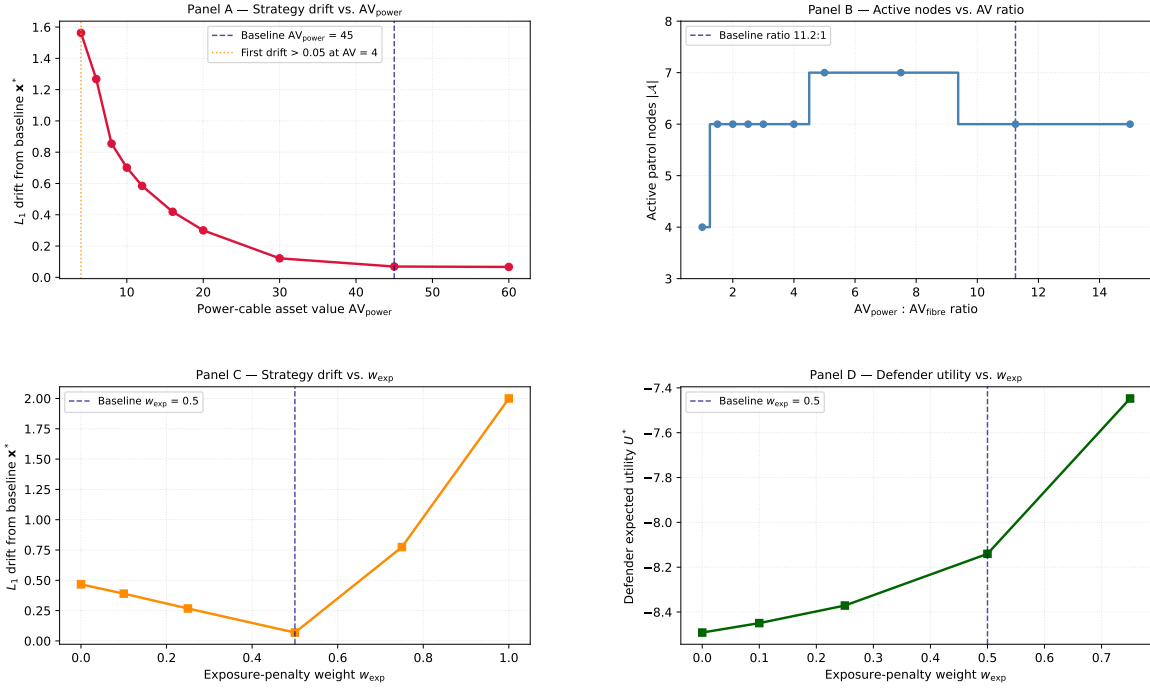


Figure 4.16: Risk-score weighting sensitivity. **Panel A** – L_1 drift of x^* as AV_{power} is reduced; dashed line marks the first value at which drift exceeds 0.05. **Panel B** – active node count vs. asset-value ratio. **Panel C** – L_1 drift vs. exposure-penalty weight w_{exp} ; navy dashed line marks the baseline. **Panel D** – defender expected utility U^* vs. w_{exp} .

Two findings emerge. First, the patrol strategy is moderately robust to AV reductions: L_1 drift remains small (< 0.15) for $AV_{power} \geq 30$ (a ratio $\geq 7.5 : 1$, corresponding to a 33% reduction from the baseline of 45), because the multiplicative θ_3 utility $U_{\theta_3}^a(t) = AV(t) \times (1 - C_{\theta_3}(t))$ preserves the relative ordering of power-cable nodes over a wide AV range. Below $AV_{power} = 20$ (5 : 1 ratio), drift exceeds 0.30 and the active-node set begins to shift; at the 1 : 1 ratio the strategy reallocates almost entirely ($L_1 = 1.56$). The baseline ratio of 11.25 : 1 sits well within the stable region. Second, the exposure-penalty weight w_{exp} is a *structurally critical assumption*: L_1 drift is near zero only at the calibrated baseline ($w_{exp} = 0.50$, $L_1 = 0.069$). Reducing w_{exp} below 0.25 expands the active set to 9–11 nodes ($L_1 > 0.27$) as lower exposure penalties make additional cable segments worth defending; increasing w_{exp} above 0.50 contracts the active set to 3 nodes at $w_{exp} = 0.75$ ($L_1 = 0.77$) and to 4 nodes at $w_{exp} = 1.0$ ($L_1 = 2.0$, maximum possible drift), as higher attacker costs concentrate the patrol budget on the few highest-value nodes. The AV ratio is therefore a *relatively robust* input, stable for ratios $\geq 7.5 : 1$; the exposure-penalty weight is a *structurally critical assumption* that requires domain validation before operational deployment.

4.7.7 Operational Constraint Sensitivity

The DOBSS optimisation produces a target coverage distribution \mathbf{x}^* ; a separate question is whether a physical USV can realise \mathbf{x}^* as the time-averaged frequency of a Markov routing solution. Two operational parameters are tested: USV speed (which limits kinematically feasible transitions in a one-hour time step) and the anti-loitering limit (which governs maximum consecutive dwell at a single node). Speed is swept across $\{5, 10, 15, 20, 25, 30\}$ km/h; the anti-loitering limit is swept across $\{1, 2, 3, 5\}$ hours with speed fixed at 30 km/h and $N = 500$ independent 60-step Monte Carlo routes simulated at each limit.

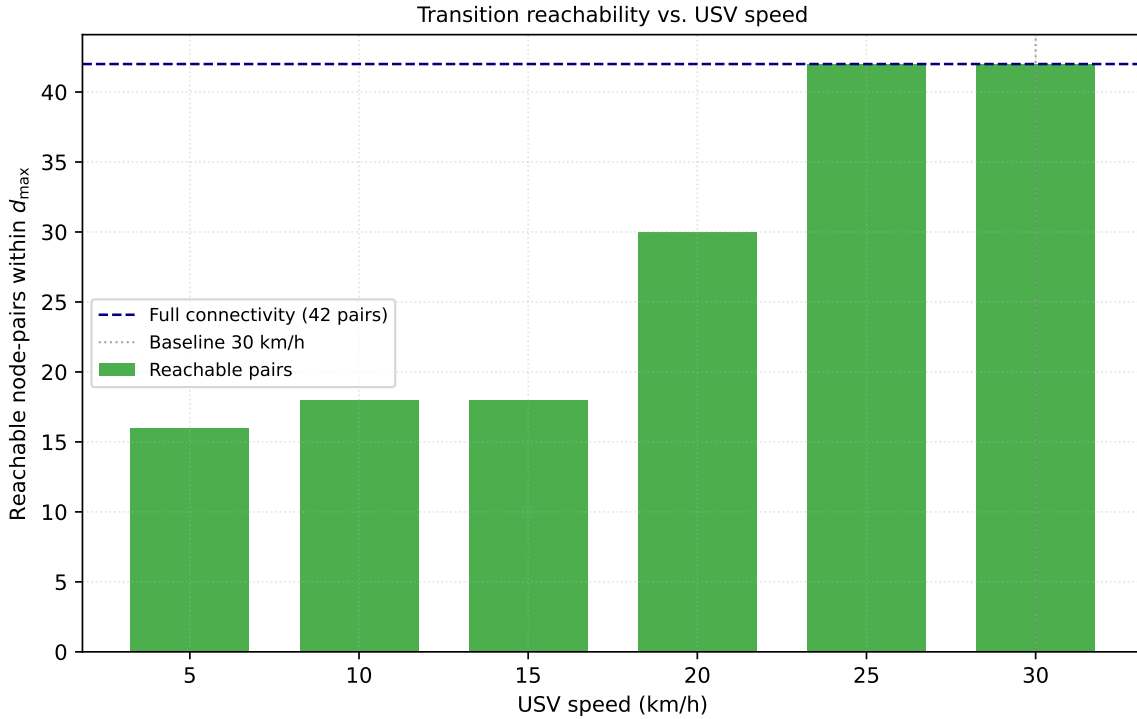


Figure 4.17: Kinematically reachable node-pairs at each speed.

Three findings follow. First, the Markov routing LP is feasible at *all six tested speeds* (5–30 km/h), an unexpected result arising from the spatial structure of the 7-node active set. The nodes partition naturally into two proximate clusters: $C_1 = \{t_{156}, t_{157}, t_{158}, t_{176}\}$ in the western cable corridor and $C_2 = \{t_{251}, t_{270}, t_{271}\}$ at the eastern cable end. Each cluster forms a path-connected subgraph even at $d_{\max} = 5$ km, and the stationarity constraint is satisfied independently within each cluster, so a globally connected chain is not required for LP feasibility. Full *inter-cluster* connectivity (all 42 node-pairs reachable within one time step) requires $d_{\max} \geq 25$ km, i.e. $v \geq 25$ km/h, driven by the maximum inter-cluster separation of 25.0 km (t_{158} – t_{270}). This provides a 17% margin below the 30 km/h baseline. Below 25 km/h the patrol operates as two independent sub-patrols, which may be

operationally acceptable but precludes cross-cluster transit within a single time step.

Second, the anti-loitering limit has *negligible effect* on coverage accuracy. Mean L_1 deviation from π is essentially identical across all four limits: $L = 1$: 0.944 ± 0.016 ; $L = 2$: 0.944 ± 0.015 ; $L = 3$: 0.946 ± 0.018 ; $L = 5$: 0.945 ± 0.016 . The near-constant $L_1 \approx 0.944$ ($\approx 47\%$ of the theoretical maximum of 2.0) indicates a structural limitation of the routing formulation: the minimum-distance LP objective produces short-hop cycles between nearest neighbours within each cluster that concentrate visit time disproportionately relative to π . The anti-loitering rule redistributes time within this suboptimal pattern but cannot correct the underlying mismatch. A convex program that maximises chain entropy directly addresses this structural limitation and is evaluated in Section 4.7.8. The baseline choice of $L = 2$ is operationally reasonable but carries no empirical accuracy advantage.

Third, the Copenhagen Orca reference platform operates at up to 40 knots (≈ 74 km/h), well above the 30 km/h assumption. Full kinematic connectivity across all 42 active-node transitions is maintained at 25 km/h (covering the maximum inter-cluster separation of 25.0 km), providing a significant speed margin below the baseline. USV speed is therefore an operational constraint on routing connectivity rather than a DOBSS model parameter; the binding constraint on coverage realisation fidelity is the routing LP formulation.

4.7.8 Maximum-Entropy Markov Routing

Section 4.7.7 identified the minimum-distance routing objective as the structural cause of $L_1 \approx 0.944$: by preferring short hops, the LP concentrates visit mass on nearest-neighbour pairs within each cluster regardless of the anti-loitering setting. This section replaces the minimum-distance objective with a *maximum-entropy* objective, asking how well the \mathbf{x}^* stationary distribution can be tracked when the routing program maximises transition unpredictability rather than minimises fuel.

Formulation. Let $\pi_i = x_i^* / \sum_k x_k^*$ denote the normalised target visit frequency at node i . The maximum-entropy Markov routing problem is:

$$\max_{\mathbf{P}} \sum_i \pi_i H(P_i) = - \sum_{i,j} \pi_i P_{ij} \log P_{ij} \quad (4.32)$$

subject to:

$$\sum_j P_{ij} = 1 \quad \forall i \quad (\text{row stochasticity}) \quad (4.33)$$

$$\sum_i \pi_i P_{ij} = \pi_j \quad \forall j \quad (\text{stationarity: } \pi \text{ is preserved}) \quad (4.34)$$

$$P_{ij} = 0 \quad \text{if } d(i, j) > d_{\max} \text{ or } i = j \quad (4.35)$$

$$P_{ij} \geq 0 \quad (4.36)$$

The objective is the π -weighted average of each row's Shannon entropy: it rewards spreading transition probability uniformly over all kinematically reachable nodes, directly counteracting the nearest-neighbour concentration produced by the minimum-distance objective. Since $-x \log x$ is concave, the problem is a convex maximisation and is solved to global optimality. The implementation uses `cvxpy` [144] with the ECOS solver [145]; the `cp.entr` atom handles the $0 \log 0 = 0$ boundary correctly.

Table 4.20: L_1 deviation from π (mean \pm SD across 500 Monte Carlo route simulations, 60 steps each, 30 km/h) for the minimum-distance LP (Section 4.7.7) and the maximum-entropy convex program.

Anti-loitering limit	Min-distance LP	Max-entropy program
$L = 1$	0.944 ± 0.016	0.188 ± 0.061
$L = 2$	0.944 ± 0.015	0.186 ± 0.060
$L = 3$	0.946 ± 0.018	0.193 ± 0.059
$L = 5$	0.945 ± 0.016	0.188 ± 0.063

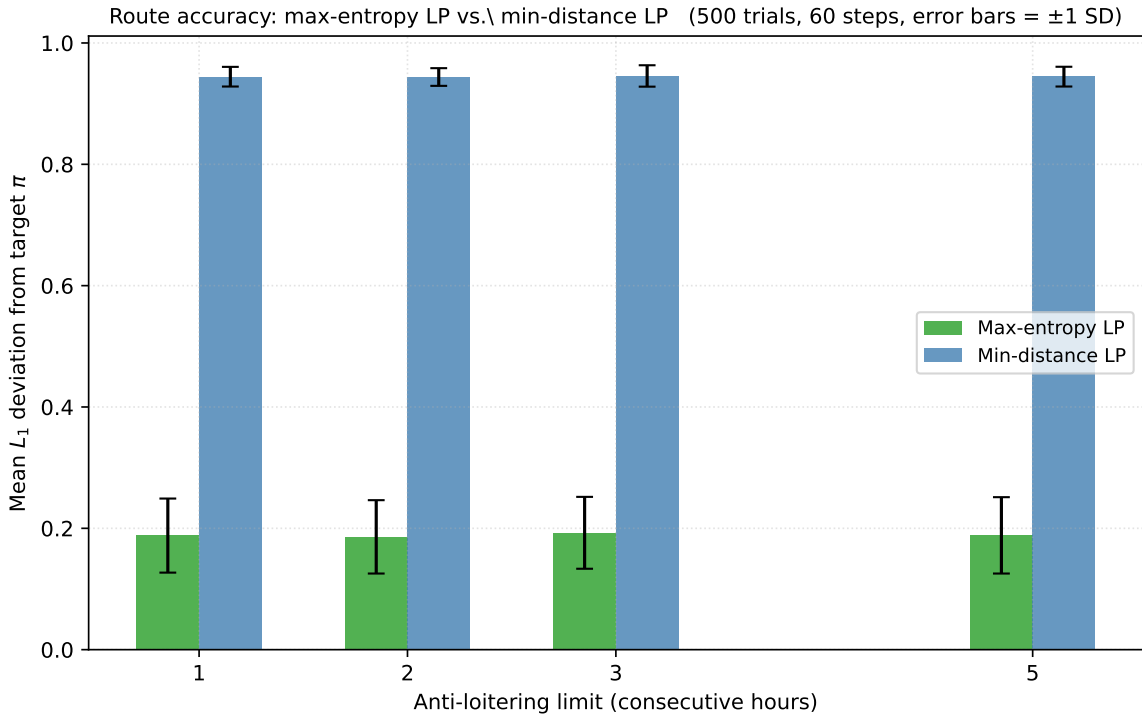


Figure 4.18: Route accuracy of the maximum-entropy program vs. the min-distance LP, measured by mean L_1 deviation from the target stationary distribution π across 500 Monte Carlo route simulations (60 steps each, 30 km/h). Error bars show ± 1 SD. The entropy program achieves $L_1 \approx 0.186$, an 80% reduction relative to the min-distance LP ($L_1 \approx 0.944$), and the improvement is robust across all tested anti-loitering limits.

Three findings follow. First, the maximum-entropy routing program reduces L_1 deviation by 80% relative to the minimum-distance baseline. At 30 km/h with $L = 2$, the entropy program achieves $L_1 = 0.186 \pm 0.060$ against 0.944 ± 0.015 , an absolute reduction of 0.758. Across all four anti-loitering limits the entropy program consistently yields $L_1 \approx 0.186\text{--}0.193$ (Table 4.20), confirming that the routing *objective*, not the loitering rule, is the governing design parameter for coverage accuracy. The residual $L_1 \approx 0.186$ ($\approx 9\%$ of the theoretical maximum of 2.0) reflects the remaining mismatch between discrete one-hour time steps and continuous target probabilities; it is substantially lower than the minimum-distance formulation achieves and within an operationally acceptable range.

Second, full inter-cluster connectivity nearly quadruples achievable entropy. At speeds of 5–15 km/h, only within-cluster transitions are kinematically reachable (16–18 node pairs) and the maximum achievable weighted entropy is $H = 0.427$. At 20 km/h (30 reachable pairs) H rises to 1.214, and at ≥ 25 km/h (all 42 pairs reachable) entropy reaches its maximum of $H = 1.522$, a 3.6-fold increase over the sub-cluster value. This reinforces the 25 km/h threshold identified in Section 4.7.7: the full-connectivity speed is critical not only for feasibility of the Markov chain problem but also for achieving maximum-unpredictability routing. At 30 km/h the expected transition distance increases from 2.8 km (min-distance LP) to 13.3 km (entropy program), confirming that the entropy objective actively selects long-range cross-cluster transitions.

Third, the anti-loitering limit remains inconsequential under entropy routing. L_1 varies by less than 0.007 across all four limits under the entropy program, mirroring the finding of Section 4.7.7. The routing objective is the dominant design parameter for coverage fidelity; the loitering rule is operationally convenient but immaterial for accuracy. Given the 80% improvement from the objective change alone, the baseline choice of $L = 2$ is retained as operationally reasonable.

The entropy routing formulation resolves the primary operational limitation of the routing phase. The same \mathbf{x}^* , the same USV speed, and the same one-hour time step are retained; only the optimisation criterion changes from minimising distance to maximising transition unpredictability. This confirms that the $L_1 \approx 0.944$ finding of Section 4.7.7 is a consequence of the routing formulation, not the game-theoretic strategy or the physical platform, and that high-fidelity realisation of \mathbf{x}^* is achievable within the existing kinematic constraints.

4.7.9 Sensitivity Analysis Summary

The patrol strategy \mathbf{x}^* rests on several uncertain inputs. A credible sensitivity analysis must test whether the core conclusion, that the Bornholm Power Cable corridor should receive concentrated patrol, is robust to plausible variations in these inputs. Seven dimensions are examined:

1. **Bounded rationality (QRE logit):** does \mathbf{x}^* remain effective if the adversary is only partially rational rather than fully strategic?

2. **Bayesian prior sensitivity:** does the patrol pattern change if the assumed probability of each threat type is revised?
3. **Monte Carlo parameter sensitivity (LHS + PRCC):** which of α , β , γ independently drives defender utility and active-node count?
4. **One-at-a-time parameter sweep:** at what individual parameter values does \mathbf{x}^* undergo regime shifts?
5. **Risk-score weighting:** does the strategy shift if the relative cable asset values or the state-actor exposure-cost formula are changed?
6. **Operational constraint sensitivity:** does the routing realisation of \mathbf{x}^* remain achievable under more conservative USV speed or anti-loitering constraints?
7. **Maximum-entropy Markov routing:** does replacing the minimum-distance routing objective with maximum-entropy routing substantially improve coverage fidelity?

Table 4.21 summarises all seven dimensions.

Table 4.21: Sensitivity analysis overview: all seven dimensions, key finding, L_1 drift range, and robustness classification. $L_1 = 0$ indicates no change in \mathbf{x}^* ; $L_1 = 2$ is maximum possible reallocation. The Monte Carlo row reports parameter influence on defender utility variance; coverage drift for the same parameters is reported in the one-at-a-time row. *Non-critical:* patrol pattern stable across the full sweep. *Structural:* large drift or regime shift observed. *Identified / Resolved:* a routing limitation surfaced by the operational-constraint sweep and closed by the maximum-entropy formulation.

Dimension	Input swept	Key finding	L_1 range	Classification
Bounded rationality (QRE logit)	$\lambda \in [0.01, 200]$	Worst-case loss at $\lambda \approx 0.5$; \mathbf{x}^* held fixed throughout sweep	n/a (\mathbf{x}^* fixed)	Non-critical
Bayesian prior sensitivity	$p(\theta_3) \in [0.01, 0.25]$	\mathbf{x}^* and active set fully stable; utility scales linearly	0.000	Non-critical
Monte Carlo (LHS + PRCC)	$\alpha, \beta, \gamma; N=500$	γ dominates utility scale (PRCC = -0.998^*); β governs active-node count (PRCC = -0.721^*)	—	β structural
One-at-a-time (OAT)	$\alpha, \beta, \gamma, \sigma_{AV}$	Monotonic active-set contraction under β ; single node at $\beta = 1.00$ ($L_1 = 1.479$); corridor collapses at $\sigma_{AV} = 0$ ($L_1 = 1.636$)	0.00–1.636	β structural
Risk-score weighting	$AV_{\text{power}} \in [4, 60]; w_{\text{exp}} \in [0, 1]$	Stable for $AV_{\text{power}} \geq 30$; w_{exp} causes near-maximum drift outside $[0.25, 0.50]$	0.07–2.00	w_{exp} structural
Operational constraint sensitivity	Speed $\in [5, 30]$ km/h; $L \in [1, 5]$	LP feasible at all speeds; min-distance LP yields $L_1 \approx 0.944$	≈ 0.944	Identified
Maximum-entropy Markov routing	Min-distance LP vs. max-entropy program	Entropy program reduces L_1 by 80% ($0.944 \rightarrow 0.186$)	0.186	Resolved

Chapter 5

Discussion

The results and sensitivity analyses presented in Chapter 4 establish that the DOBSS framework produces a structurally coherent and computationally tractable patrol strategy for the Bornholm Basin. This chapter examines what those results collectively mean: first by returning to the central research question and its three sub-questions, then by characterising the robustness hierarchy the sensitivity analysis reveals, and finally by addressing the limitations that bound the framework’s operational applicability and the extensions most likely to overcome them.

5.1 Return to the Research Questions

5.1.1 Sub-Question 1: Modelling a Heterogeneous Adversary Population

The first sub-question asks how the heterogeneous threat landscape, and unintentional agents, can be captured in a unified game-theoretic adversary model. The Bayesian Stackelberg formulation with three follower types ($\theta_1, \theta_2, \theta_3$) provides the formal answer: each type receives a utility function encoding its operational objectives, and the DOBSS MILP optimises coverage against the prior-weighted mixture simultaneously.

The sensitivity analysis adds important qualifications. The prior invariance result, discussed in Section 4.7.3 shows that \mathbf{x}^* is structurally unchanged as $p(\theta_3)$ varies from 0.01 to 0.25, confirming that the model’s spatial recommendations are not hostage to the analyst’s prior belief. The QRE analysis, presented in Section 4.7.2, shows that θ_1 and θ_2 are effectively deterred at $\lambda \geq 10$, while θ_3 accounts for all residual loss at any rationality level, confirming the three-type formulation functions as intended.

However, the model embeds a structural tension. The DOBSS objective is weighted by prior probability, which concentrates optimisation effort on the 98% case ($\theta_1 + \theta_2$). The 2% state actor (θ_3) drives the equilibrium because its utility peak at the power cable is geographically dominant, not because the Bayesian objective prioritises it. The result is that the framework inadvertently exploits a property of the spatial data rather than solving

the adversary-heterogeneity problem in its most general form. In a domain where the state-actor utility landscape were less extreme, the low prior weight on θ_3 would translate directly into under-allocation against the most capable threat. Section 5.3.3 addresses this further.

A second qualification concerns the discrete type categorisation. The boundary between accidental anchor drag (θ_1), unintentional commercial-shipping incident (θ_2), and deliberate state-directed sabotage (θ_3) is operationally indistinct: forensic attribution of Baltic infrastructure incidents to specific intent categories has proven contentious in each documented case since 2022. The three-type model is a modelling necessity rather than a claim about the real world, and its outputs should be presented to operators with this explicitly acknowledged.

5.1.2 Sub-Question 2: Translating Coverage Probabilities into Feasible Routes

The second sub-question asks how strategic coverage probabilities can be translated into kinematically feasible and unpredictable patrol routes. The Markov chain synthesis provides the formal mechanism: the stationary distribution of a discrete-time Markov chain over the active node set is constrained to equal \mathbf{x}^* , with the transition matrix solved under kinematic reachability and an anti-loitering rule enforcing within-corridor unpredictability.

The operational sensitivity analysis (Section 4.7.7) reveals that the minimum-distance routing objective is the binding constraint on coverage fidelity: the Markov LP is feasible at all tested speeds from 5 to 30 km/h, but mean L_1 deviation of empirical visit frequencies from π is approximately 0.944 across all anti-loitering settings ($\approx 47\%$ of the theoretical maximum). Short-hop cycles concentrate dwell time on nearest-neighbour pairs within each cluster, systematically under-visiting distant-but-active nodes regardless of the loitering rule. The game-theoretic component is intact: DOBSS correctly identifies the target distribution. The minimum-distance routing objective accounts for the translation loss.

Section 4.7.8 directly addresses this gap by replacing the minimum-distance objective with a maximum-entropy objective that rewards spreading transition probability uniformly over all kinematically reachable nodes. The entropy program achieves $L_1 = 0.186 \pm 0.060$ at the 30 km/h baseline, an 80% reduction from 0.944, confirming that the $L_1 \approx 0.944$ finding is a consequence of the routing formulation, not the game-theoretic strategy or the physical platform. The thesis's answer to Sub-Question 2 is therefore that the translation is both *possible* (Markov chain feasibility confirmed at all tested speeds) and, under the entropy routing formulation, *substantially faithful* (residual $L_1 \approx 0.186$, approximately 9% of the theoretical maximum).

5.1.3 Sub-Question 3: Improvement over Baselines and Robustness Across Rationality Levels

The third sub-question asks to what extent the Bayesian Stackelberg approach improves expected outcomes relative to baseline strategies, and how robust this improvement is across adversary rationality levels.

The QRE analysis addresses the rationality-robustness dimension directly. DOBSS x^* is most effective in the human-rational-to-rational regime ($\lambda \geq 10$), where the θ_1 and θ_2 contributions to total loss collapse to near zero and total loss stabilises near -8.78 to -8.81 . Notably, this does not recover the Stackelberg equilibrium objective of -8.18 ; the residual gap is a consequence of the SSE tie-breaking convention, as analysed in Section 4.7.2. The worst-case loss occurs at near-random rationality ($\lambda \approx 0.5$, total loss ≈ -38.8), a regime in which no deterrence-based patrol design is effective, because there is no strategic concentration to exploit. This worst-case is therefore a property of the utility landscape rather than a failure of the patrol design. The baseline comparison in Section 4.7.1 quantifies the improvement against heuristics directly under the same Strong Stackelberg Equilibrium assumption: DOBSS reduces expected loss by 9.0% over Uniform and 8.0% over Asset-Proportional.

The honest characterisation of model uncertainty comes from the Monte Carlo analysis. The IQR of defender utility across plausible parameter uncertainty is approximately 6.5 utility units (IQR $[-12.94, -6.48]$), emphasising that parameter calibration effort has substantially larger impact on realised utility than any strategic-choice refinement within the current model. The primary argument for DOBSS in this setting is therefore theoretical: the Stackelberg guarantee of optimality given the model, combined with its principled foundation for extensions such as fleet scaling, prior updating, and robustification. In settings with more balanced infrastructure or more ambiguous threat mixtures, the game-theoretic computation would be solving a genuinely harder allocation problem, and empirical advantages over simpler heuristics would likely be larger.

5.2 The Robustness Hierarchy

Taken together, the seven sensitivity dimensions reveal a three-level structure in the robustness of the framework's outputs.

The patrol *target*, the identification of the Bornholm power cable corridor as the concentration zone, is essentially unconditional. It is invariant to prior variation over $p(\theta_3) \in [0.01, 0.25]$, full-range OAT sweeps in α and γ , all 500 Monte Carlo parameter combinations, a 33% reduction in the power cable's assessed asset value, and all tested speed configurations. Analytically, this robustness traces to the baseline 11.25:1 asset-value ratio between the power cable and the fibre cables sitting well within the stable region identified by the asset-value sensitivity sweep: L_1 drift remains small (<0.15) for ratios $\geq 7.5:1$ and exceeds 0.30 only below 5:1. The qualitative conclusion, patrol the power cable, holds

under any perturbation the sensitivity analysis tests.

The patrol *allocation*, the specific probability distribution over active nodes, is sensitive to two inputs: β (attacker retained reward under coverage) and w_{exp} (state-actor exposure-penalty weight). β produces a monotonic spatial contraction: n_{active} declines from 7 at the baseline ($\beta = 0.10$) to 6 at $\beta = 0.20$, 5 at $\beta = 0.40$ – 0.50 , 3 at $\beta = 0.70$ – 0.90 , and 1 at $\beta = 1.00$, the boundary case where coverage confers zero deterrence ($L_1 = 1.479$). The contraction is continuous— L_1 drift rises from 0.075 at $\beta = 0.20$ to 0.336 at $\beta = 0.50$ to 1.479 at $\beta = 1.00$ —and the target node remains t_{156} throughout; there is no threshold effect or target switch. w_{exp} produces near-maximum drift in both directions from the baseline of 0.50. Both are low-SoK expert judgements without empirical calibration, meaning the specific within-corridor allocation reported throughout is conditionally reliable only to the extent that $\beta < 0.5$ and $w_{\text{exp}} \approx 0.50$ are credible.

The *absolute utility value* is almost entirely a scale artefact of γ (PRCC = -0.998^* , $S_T^\gamma \approx 1.000$). This does not affect the patrol pattern, but it means the utility numbers reported throughout are not directly interpretable as operational risk magnitudes without an explicit currency mapping for γ . Strategy rankings are the operationally meaningful quantity; absolute utility comparisons across configurations require γ to be held fixed. Rass [146] characterises this as a structural limitation of scalar security payoff games and proposes distribution-valued payoffs ordered by a stochastic ordering relation as a principled resolution—one that simultaneously removes the scale ambiguity and preserves the tail-risk information most relevant to the θ_3 scenario.

This hierarchy defines what the thesis delivers with confidence and what it delivers conditionally. The power cable corridor is the correct patrol zone. The precise allocation within that corridor is a model-conditional recommendation requiring empirical calibration of two parameters before operational deployment.

The three levels correspond directly to the Strength-of-Knowledge ratings established for the Baltic subsea threat environment in Chapter 2. The patrol *target* rests on Strong-SoK inputs: the power cable’s dominant asset value is derived from objective infrastructure parameters, MTTR, redundancy, and strategic weight, that carry high expert agreement and are insensitive to the contested prior probabilities. The patrol *allocation* is contingent on Moderate-SoK inputs: β and w_{exp} are analyst judgements explicitly rated as lacking direct observational grounding in the SoK assessment, and the sensitivity analysis confirms that the allocation is unreliable precisely where knowledge is weakest. The *absolute utility values* carry effectively Weak SoK: without an empirically calibrated currency for γ , they do not support quantitative risk comparisons across configurations. The sensitivity analysis is therefore not merely a technical robustness check but a formal propagation of epistemic uncertainty from the risk characterisation stage into the operational recommendations, closing the loop between the (A', C', Q, K) framework and the game-theoretic outputs.

5.3 Limitations

5.3.1 The One-Shot Game and Adversarial Learning

DOBSS is a one-shot Stackelberg game: the defender commits to \mathbf{x}^* in advance, and the attacker responds once. This assumption is theoretically grounded when patrol deployment precedes adversary observation. In practice, maritime patrol is a repeated process conducted over weeks and months. A patient adversary, precisely the θ_3 state actor that drives the equilibrium, can conduct pre-operational reconnaissance over multiple patrol cycles, mapping the 7-node active corridor even if the specific node visited at any moment is unpredictable within it. The anti-loitering rule introduces within-corridor unpredictability, but the corridor itself is a fixed observable signature that no within-corridor randomisation can conceal.

This is not a failure of the Stackelberg formulation; it is a scope limitation acknowledged by the framework’s designers in analogous deployments such as ARMOR at LAX and PROTECT for the US Coast Guard [147]. The static Stackelberg equilibrium provides the correct distribution to commit to; a repeated-game extension with belief updating on both sides would additionally quantify how quickly the corridor becomes exploitable and what minimum route entropy prevents it. Until such an extension is developed, the framework should be understood as specifying the correct long-run time-average patrol distribution, not as accounting for the dynamic in which a sophisticated adversary adapts its reconnaissance timeline to patrol cycle predictability.

5.3.2 Binary Coverage and Sensor Physics

The model assigns binary coverage to each 2.5×2.5 km cell: a USV either covers a cell or does not. Real sensor systems, radar, AIS, electro-optical, side-scan sonar, produce detection probability as a continuous, decreasing function of slant range, modulated by sea state, visibility, and target signature. At the baseline grid resolution, a USV at the centre of an active cell has effective radar coverage extending well into adjacent cells, which the current formulation ignores. This means that the model underestimates coverage at cells adjacent to active nodes and overestimates the sharp boundary between covered and uncovered regions.

Replacing the binary indicator with a range-dependent detection function $P(\text{detect} | d, s)$ would allow adjacent active cells to contribute partial coverage to each other [148], potentially distributing the optimal patrol more broadly across the cable corridor and reducing the vulnerability of the currently zero-coverage segments between active nodes. The McCormick linearisation used in the DOBSS MILP can accommodate a linear approximation to this function without changing the solver infrastructure, making this a tractable near-term extension.

5.3.3 Low-Prior High-Capability Adversary (Black-Swan) Vulnerability

The state actor θ_3 carries a prior of 2% yet accounts for all residual loss at high adversary rationality ($\mathcal{L}_{\theta_3} = -8.78$ at $\lambda = 50$ and -8.81 at $\lambda = 200$). This is a known structural property of Bayesian Stackelberg models with low-prior high-capability types: the objective function weights each type by its frequency, systematically under-resourcing defence against infrequent but catastrophic attacks. The three documented Baltic hybrid infrastructure incidents since 2022, Nord Stream (September 2022), Estlink-2 (December 2024), and BCS East-1 (January 2025), suggest that the θ_3 event rate may be higher than the ICPC global baseline implies, and that the true per-incident consequence is closer to the catastrophic end of the range.

In Aven’s [149] terminology, pure expected-value optimisation optimises for the mean but provides no guarantee about tail events. The finding that a near-random attacker at $\lambda \approx 0.5$ produces a fourfold loss increase relative to the rational equilibrium is one instance of this: the framework is designed for the rational-adversary case and performs poorly against scenarios that do not conform to the game-theoretic framing. A two-objective formulation that jointly minimises Bayesian expected loss and worst-case loss across threat types would allow a decision-maker to navigate the trade-off between frequency-weighted and capability-weighted optimality, explicitly quantifying the additional patrol cost required to close the θ_3 vulnerability gap. Alternatively, a minimax-regret formulation that weights θ_3 independent of its prior probability would provide a worst-case guarantee at the cost of some frequency-weighted performance. The multi-objective security game (MOSG) framework of Rass [146], developed under the EU HyRiM project, provides a direct computational path for the two-objective approach: it yields Pareto-optimal patrol distributions over separate loss dimensions without collapsing them into a single scalar, making the trade-off between frequency-weighted and worst-case loss explicit rather than embedding it in the Bayesian objective.

5.3.4 Static Strategy and Seasonal Dynamics

The framework is calibrated on 2023 annual-average vessel-density rasters and treats threat priors as fixed inputs. Baltic vessel traffic is strongly seasonal: fishing activity peaks in late spring and summer, tanker and cargo routing shifts with energy demand and Baltic weather windows, and the overall shadow-fleet transit volume is not stationary over multi-year timescales. The current \mathbf{x}^* is optimal for the annual average; it may be sub-optimal during periods of elevated threat or shifted traffic patterns.

The invariance result provides partial reassurance: \mathbf{x}^* is stable across $p(\theta_3) \in [0.01, 0.25]$, suggesting that even a significant seasonal increase in assessed state-actor risk would not alter the patrol pattern. However, the θ_1 and θ_2 priors, which jointly carry 98% of the probability mass, are derived from seasonal fishing and shipping rasters. Monthly MILP re-runs with updated EMODnet vessel density inputs are achievable within the existing computational pipeline: scaling analysis demonstrates that the 400-node MILP solves in

under 0.5 seconds, making daily re-optimisation operationally feasible.

5.3.5 EEZ Boundary Effects and Adversary Risk Geography

The model treats the USV as the sole source of detection risk over a homogeneous operational area. In practice, the adversary’s strategic space is bounded by the Danish EEZ: crossing into neighbouring states’ regional waters carries substantially higher detention risk, and the boundary transit itself is an observable event that has triggered multi-state enforcement responses in recent Baltic incidents [150]. The southern cluster of \mathbf{x}^* , closest to this boundary, therefore benefits from structural deterrents the model does not capture; coverage probabilities there should be read as upper bounds on the USV’s specific contribution.

5.3.6 Weaknesses in the Adversary Utility Function Formulations

Three related weaknesses affect the adversary model. The most fundamental concerns the inclusion of unintentional types θ_1 and θ_2 inside the Stackelberg game. Accidental agents do not adapt their behaviour in anticipation of USV coverage, so applying follower rationality to fishing vessels and commercial shipping is a structural category error. In practice, its consequence is the opposite of what the prior weights imply: because θ_1 and θ_2 optimally target high-density nodes where $\text{AssetValue} = 0$, their combined 98% prior weight contributes near-zero to the defender objective, and the MILP degenerates into optimising almost exclusively for θ_3 coverage. The fix is straightforward: restricting accidental utility to nodes where $\text{AssetValue} > 0$ would force θ_1 and θ_2 to compete over infrastructure nodes, restoring the intended balance between intentional and unintentional risk in the objective.

A second weakness concerns the θ_3 utility function, which permits negative expected utility under extreme boundary conditions, specifically when the maximum substrate penalty coincides with zero shipping density, pushing the total attacker cost above 1.0. This has not been observed to destabilise the current grid, but imposing a zero lower bound on $U_a^{\theta_3}$ would preserve strict game-theoretic rationality across all parameter configurations.

5.4 Future Work

The sensitivity analysis identifies a priority ordering for extensions based on both consequence magnitude and implementation tractability.

Empirical calibration of β and w_{exp} . The maximum-entropy routing extension (Section 4.7.8) resolved the routing translation loss, shifting the primary open question from

routing fidelity to parameter calibration. The two structural sensitivity parameters, β (attacker retained reward under coverage) and w_{exp} (state-actor exposure weight), are the highest-consequence low-SoK inputs identified by the analysis. Tabletop exercises with naval interdiction officers and coast guard commanders could yield interval estimates for β , replacing the current point assumption of 0.10 with a range that feeds directly into the Monte Carlo pipeline. For w_{exp} , engagement with Baltic maritime security assessments at an unclassified ordinal level would bound the relative deterrence weight of vessel exposure.

Uncertainty quantification on \mathbf{x}^* . The Monte Carlo analysis quantifies uncertainty in the scalar utility U^* but reports \mathbf{x}^* only at the baseline parameterisation. Bootstrap resampling of spatial input data, cable intersection fractions, substrate raster assignments, vessel density grids, would yield a distribution over coverage vectors, identifying which nodes are robustly active across all input realisations and which are marginally active (sensitive to small changes in substrate classification or vessel density). This distinction is operationally important: robustly active nodes should always be patrolled; marginally active nodes are candidates for intelligence-led conditional activation.

Seasonal prior and raster updating. Monthly MILP re-runs with updated EMODnet vessel density inputs are achievable within the existing computational pipeline. The scaling result (under 0.5 s per solve at 400 nodes) confirms that daily re-optimisation is operationally feasible. The prior invariance result suggests that the power-cable corridor will remain the spatial centre of gravity under any plausible seasonal prior shift, but seasonal reallocation within the corridor may be non-trivial when θ_1 and θ_2 vessel densities shift significantly.

Repeated-game extension. An adversary-learning extension on the defender side and observation-based updating on the attacker side would quantify the patrol observability risk and the minimum route entropy required to prevent corridor exploitation [151]. This is the most methodologically ambitious extension but addresses the most fundamental structural limitation of the one-shot formulation.

Deceptive patrol tactics. The Strong Stackelberg Equilibrium assumes that a patient θ_3 adversary can eventually reconstruct the committed coverage distribution \mathbf{x}^* through repeated observation, converting the patrol record into reconnaissance. A deceptive extension would decouple the signal the USV transmits from the strategy it executes: the defender jointly optimises an action policy π and a communication strategy σ that selectively withholds or distorts information about operational intent, degrading the informational advantage available to a state-sponsored observer. Clempner [136] establishes the computational basis for this separation in a Bayesian-Markov Stackelberg framework,

demonstrating that partial information disclosure can yield strictly higher expected defender utility than full transparency. Applied to USV patrol, the benefit is direct: reducing the intelligence yield of long-run observation degrades the θ_3 adversary's ability to preposition before a high-value attack window and complements the route-entropy objective already embedded in the maximum-entropy Markov routing stage.

Techno-Economic Assessment. Risk management inherently involves balancing financial expenditure against risk mitigation. While it is evident that deploying crewed vessels, ranging from high value capital platforms to smaller fast response craft, incurs significantly higher costs than operating autonomous systems such as Unmanned Surface Vehicles, a comprehensive techno economic analysis is required to quantify this margin. Integrating such modeling into the proposed framework represents a valuable future extension that would provide decision makers with the empirical justification needed to adopt these novel methodologies. The sensitivity analysis establishes that the normalisation constant γ is a pure utility-scale parameter (Section 4.7.9) and that absolute utility values carry no direct monetary interpretation without an empirical currency mapping. A techno-economic extension would close this gap: by eliciting monetised estimates of disruption costs for the specific Bornholm assets from infrastructure operators or domain experts, the abstract utility improvements produced by DOBSS could be expressed in euros and compared against USV procurement and operating costs to produce a break-even analysis for fleet sizing.

Partially Observable Stochastic Games. The Strong Stackelberg Equilibrium rests on a transparency assumption: the defender commits to a publicly known mixed strategy and the attacker observes it fully before responding. Both directions of this assumption are imperfect in practice. The defender cannot directly observe attacker intent or type at the moment of patrol deployment, and a sophisticated adversary reconstructs the committed distribution only gradually through repeated observation of patrol patterns. The natural generalisation that relaxes both constraints simultaneously is a Partially Observable Stochastic Game (POSG), in which each player maintains and updates a private belief state over the hidden system state rather than conditioning on a commonly known strategy profile [152, 153]. Counterfactual Regret Minimisation (CFR) and its variance-reduced variants provide tractable equilibrium-finding algorithms for this class, having demonstrated scalability to large imperfect-information games [154]. In the patrol context, a POSG formulation would allow the defender to revise its type-probability estimates from sensor observations in real time rather than treating the Bayesian prior as a fixed input, and would allow the attacker's strategy to depend on the observable patrol history, directly connecting this extension to the adversarial learning and deceptive-tactics limitations identified in Sections 5.3.1 and the preceding paragraph.

Chapter 6

Conclusion

This thesis addressed how game theory can be used to plan patrol strategies for Unmanned Surface Vehicles tasked with securing critical maritime infrastructure. The motivation arose from a concrete operational mismatch in the Baltic Sea: a documented and escalating pattern of deliberate subsea infrastructure attacks set against a monitoring capability that remains fragmented, reactive, and insufficient. Existing approaches provide no principled method for allocating scarce autonomous surveillance resources against a heterogeneous adversary population under epistemic uncertainty.

The two-phase framework developed here, a Bayesian Stackelberg Security Game solved via the DOBSS algorithm coupled to a Markov chain routing stage, provides an affirmative and tractable answer. Its primary contribution is structural rather than empirical: a decision-support architecture that carries a Stackelberg optimality guarantee given the model, grounds an abstract security game in real geospatial data, and is extended in this work to maximum-entropy routing. On the first sub-question, the three-type formulation encodes the heterogeneity of the Baltic threat landscape in a single mixed strategy that is simultaneously optimal against unintentional fishing (θ_1), unintentional shipping (θ_2), and rational state-sponsored (θ_3) adversaries. On the second, the maximum-entropy Markov routing stage realises the coverage distribution as a kinematically feasible patrol, reducing the L1 deviation from the target from 0.944 to 0.186, an 80% improvement over the minimum-distance baseline. On the third, the equilibrium is robust across the tested rationality range: θ_1 and θ_2 are deterred at $\lambda \geq 10$, while the residual loss from θ_3 stabilises at its rational-equilibrium value. Under a common Strong Stackelberg Equilibrium, DOBSS reduces expected loss by 9% over Uniform and 8% over Asset-Proportional allocation.

This advantage must be read against two findings the analysis establishes. First, the gain traces to one structural feature, the sevenfold concentration of coverage on the dominant node t_{156} , and the one-at-a-time sweep shows that the corridor collapses once asset-value heterogeneity is removed; the result is therefore a consequence of the case study's single dominant asset. Second, plausible parameter uncertainty produces an interquartile spread in defender utility of roughly six and a half units, an order of magnitude larger

than the difference between strategies. Within this model, the optimality guarantee and the foundation for extension, rather than the headline percentage, are the principal results.

The sensitivity analysis resolves these outputs into a three-level reliability hierarchy that maps onto the Strength-of-Knowledge ratings of Chapter 2. The patrol *target*, concentration along the Bornholm power-cable corridor, is unconditionally robust, invariant across all prior variations, the full ensemble of 500 Monte Carlo combinations, a 33% reduction in the cable’s assessed value, and every tested speed; it rests on Strong-SoK infrastructure parameters and an asset-value ratio exceeding the competitive threshold elevenfold. The patrol *allocation* within that corridor is conditionally reliable, sensitive to the attacker retained-reward parameter β and the exposure weight w_{exp} , both Moderate-to-Weak-SoK judgements without empirical calibration. The absolute utility *values* are scale artefacts of γ and carry no operational meaning; only strategy rankings are interpretable. The analysis thus propagates epistemic uncertainty from the (A', C', Q, K) characterisation into the operational recommendation.

The framework carries limitations that bound its applicability. It is a proof-of-concept, not a validated operational system. The one-shot game does not model adversarial learning across repeated patrol cycles, leaving the fixed corridor an observable signature that within-corridor randomisation does not conceal; the binary coverage model omits range-dependent sensor physics; the single-USV architecture excludes multi-vessel coordination; and the homogeneous-jurisdiction assumption does not represent cross-border deterrence at the study area’s southern margin. Despite a 2% prior, θ_3 accounts for all residual loss at high rationality, a property of the expected-value objective rather than of the patrol design. The multi-objective security game framework of Rass [146, 155] is identified as a direction for addressing this within the existing MILP.

The MILP solves in under half a second at 400 nodes on commodity hardware, which supports frequent re-optimisation as inputs change. The reliability hierarchy defines what the framework delivers at this stage: the power-cable corridor is the most robust output, while the within-corridor allocation remains model-conditional and contingent on calibration of β and w_{exp} . The work establishes that game-theoretic USV patrol optimisation is well-founded theoretically and computationally, and that its transition from proof-of-concept to operational decision support is gated by the calibration and adversary-learning extensions identified here.

Bibliography

- [1] Aurora Ganz, Margherita Camellini, Emmie Hine, et al. "Submarine Cables and the Risks to Digital Sovereignty". In: *Minds and Machines* 34.31 (2024). DOI: 10.1007/s11023-024-09683-z. URL: <https://doi.org/10.1007/s11023-024-09683-z>.
- [2] Lionel Carter. *Submarine Cables and the Oceans: Connecting the World*. Accessed via Google Books. Nairobi, Kenya: UNEP/Earthprint, 2009. URL: <https://books.google.dk/books?id=MTToDcTSXmAC>.
- [3] International Telecommunication Union. *Submarine Cables: Digital Resilience*. Accessed: 2025-04-28. 2025. URL: <https://www.itu.int/en/mediacentre/Pages/PR-2025-02-27-submarine-cables-summit-nigeria.aspx>.
- [4] Kristi Govella. "Undersea cables, geoeconomics, and security in the Indo-Pacific: Risks and resilience". In: *Marine Policy* 180 (2025), p. 106809. ISSN: 0308-597X. DOI: <https://doi.org/10.1016/j.marpol.2025.106809>. URL: <https://www.sciencedirect.com/science/article/pii/S0308597X25002246>.
- [5] Tobias Liebetrau and Christian Bueger. "Advancing coordination in critical maritime infrastructure protection: Lessons from maritime piracy and cybersecurity". In: *International Journal of Critical Infrastructure Protection* 46 (2024), p. 100683. ISSN: 1874-5482. DOI: <https://doi.org/10.1016/j.ijcip.2024.100683>. URL: <https://www.sciencedirect.com/science/article/pii/S1874548224000246>.
- [6] Henry M. Jackson School of International Studies. *Baltic Sea Undersea Cable Security*. <https://jsis.washington.edu/news/baltic-sea-undersea-cable-security/>. Accessed: 2025-11-26. 2024.
- [7] Katharina Buchholz. *Infographic: Baltic Sea Cable Incidents Pile Up*. Statista Daily Data. Accessed: 2026-01-11. Feb. 2025. URL: <https://www.statista.com/chart/33892/damage-to-underwater-cables-and-pipelines-in-the-baltic-sea/>.
- [8] Christian Bueger and Tobias Liebetrau. "Critical maritime infrastructure protection: Whats the trouble?" In: *Marine Policy* 155 (2023), p. 105772. ISSN: 0308-597X. DOI: <https://doi.org/10.1016/j.marpol.2023.105772>. URL: <https://www.sciencedirect.com/science/article/pii/S0308597X23003056>.

- [9] Charlie Edwards and Nate Seidenstein. *The Scale of Russian Sabotage Operations Against Europe's Critical Infrastructure*. Research Paper. International Institute for Strategic Studies, 2025.
- [10] TeleGeography. *Submarine Cable Map 2026*. <https://submarine-cable-map-2026.telegeography.com/>. Interactive map depicting 694 active or under-construction submarine cable systems and 1,893 landings. Last updated around February 2026. Sponsored by Telecom Egypt. Licensed under CC BY-SA 4.0. TeleGeography, 2026. URL: <https://submarine-cable-map-2026.telegeography.com/> (visited on 02/23/2026).
- [11] Yang Shen et al. "Analysis of Nord Stream explosions using seismic recordings". In: *Applied Geophysics* 20.3 (2023), pp. 316–323.
- [12] Carnegie Endowment for International Peace. *The Baltic Sea at a Boil: Connecting the Shadow Fleet and Episodes of Subsea Infrastructure Sabotage*. <https://carnegieendowment.org/europe/research/2025/06/baltic-russia-maritime-cable-sabotage>. Accessed: 2026-04-16. June 2025.
- [13] International Cable Protection Committee. *Damage to Submarine Cables from Dragged Anchors*. ICPC Viewpoints. Accessed: 2024-05-20. 2024. URL: <https://www.iscpc.org/publications/icpc-viewpoints/damage-to-submarine-cables-from-dragged-anchors/>.
- [14] Per Andersen, Kenneth Øhlenschläger Buhl, and Rasmus Dahlberg. "Beredskabsjura som begreb: Hybridkrig under overfladen". In: *Økonomi & Politik* 98.2 (2025). Special issue: Samfundssikkerhed under opbrud, pp. 55–69. DOI: 10.7146/okonomi-og-politik.v98i2.157884. URL: <https://tidsskrift.dk/okonomi-og-politik/article/view/157884>.
- [15] G Alexander Crowther. "The baltic sea region at an inflection point". In: *Prism* 10.2 (2023), pp. 6–17.
- [16] Dimitrios Eleftherakis and Raul Vicen-Bueno. "Sensors to Increase the Security of Underwater Communication Cables: A Review of Underwater Monitoring Sensors". In: *Sensors* 20.3 (2020), p. 737. DOI: 10.3390/s20030737. URL: <https://www.mdpi.com/1424-8220/20/3/737>.
- [17] Giovanni Soldi et al. "Monitoring of underwater critical infrastructures: The Nord Stream and other recent case studies". In: *arXiv preprint arXiv:2302.01817* (2023).
- [18] Steven C Boraz. "Maritime domain awareness: Myths and realities". In: *Naval War College Review* 62.3 (2009), pp. 137–146.
- [19] Ash Rossiter. "Cable risk and resilience in the age of uncrewed undersea vehicles (UUVs)". In: *Marine Policy* 171 (2025), p. 106434.

- [20] Sumin Ryu and Sanghyuck Han. "Environment-Aware Multi-Sensor Fusion for Maritime Domain Awareness: A Comprehensive Review". In: *Korean Journal of Remote Sensing* 41.6 (2025), pp. 1225–1250.
- [21] European Parliament and Council of the European Union. *Directive (EU) 2022/2555 of the European Parliament and of the Council on measures for a high common level of cybersecurity across the Union (NIS2 Directive)*. Directive (EU) 2022/2555. L 333/80, 27 December 2022. Official Journal of the European Union, 2022. URL: <https://eur-lex.europa.eu/eli/dir/2022/2555/oj>.
- [22] European Parliament and Council of the European Union. *Directive (EU) 2022/2557 of the European Parliament and of the Council on the resilience of critical entities and repealing Council Directive 2008/114/EC (CER Directive)*. Directive (EU) 2022/2557. L 333/164, 27 December 2022. Official Journal of the European Union, 2022. URL: <https://eur-lex.europa.eu/eli/dir/2022/2557/oj/eng>.
- [23] European Commission. *Submarine Cable Security Toolbox and Cable Projects of European Interest*. European Commission, Directorate-General for Communications Networks, Content and Technology. Feb. 2025. URL: <https://digital-strategy.ec.europa.eu/en/library/submarine-cable-security-toolbox-and-cable-projects-european-interest>.
- [24] Demetrios N Tsailas. "Risks And Threats In The 21st Century Maritime Security". In: *Security Science Journal* 6.1 (2025), pp. 106–144.
- [25] Barry Charles Ezell et al. "Probabilistic risk analysis and terrorism risk". In: *Risk Analysis: An International Journal* 30.4 (2010), pp. 575–589.
- [26] Terje Aven. "Probabilities and background knowledge as a tool to reflect uncertainties in relation to intentional acts". In: *Reliability Engineering & System Safety* 119 (2013), pp. 229–234.
- [27] Louis Anthony Cox Jr. "Some limitations of Risk= Threat× Vulnerability× Consequence for risk analysis of terrorist attacks". In: *Risk Analysis: An International Journal* 28.6 (2008), pp. 1749–1761.
- [28] Terje Aven. "On how to define, understand and describe risk". In: *Reliability Engineering & System Safety* 95.6 (2010), pp. 623–631.
- [29] Terje Aven and Shital Thekdi. *Risk Science: An Introduction*. 1st. London: Routledge, 2021. ISBN: 9781003156864. DOI: 10.4324/9781003156864. URL: <https://doi.org/10.4324/9781003156864>.
- [30] Terje Aven. "A unified framework for risk and vulnerability analysis covering both safety and security". In: *Reliability engineering & System safety* 92.6 (2007), pp. 745–754.

- [31] Terje Aven. "On some recent definitions and analysis frameworks for risk, vulnerability, and resilience". In: *Risk Analysis: An International Journal* 31.4 (2011), pp. 515–522.
- [32] Terje Aven. "On some foundational issues concerning the relationship between risk and resilience". In: *Risk Analysis* 42.9 (2022), pp. 2062–2074.
- [33] Douglas Landoll. *The security risk assessment handbook: A complete guide for performing security risk assessments*. CRC press, 2021.
- [34] Tore Askeland, Roger Flage, and Terje Aven. "Moving beyond probabilities—Strength of knowledge characterisations applied to security". In: *Reliability Engineering & System Safety* 159 (2017), pp. 196–205.
- [35] Sasan Zarghooni-Hoffmann and Terje Aven. "A risk science perspective on how to evaluate the quality of intelligence assessments". In: *Intelligence and National Security* 40.2 (2025), pp. 219–239. DOI: 10.1080/02684527.2024.2432769.
- [36] Terje Aven. "Improving the Foundation and Practice of Uncertainty Analysis: Strengthening Links to Knowledge and Risk". In: *Knowledge in Risk Assessment and Management* (2018), pp. 103–125.
- [37] Danish Maritime Authority. *Safety of Navigation*. <https://www.dma.dk/safety-at-sea/safety-of-navigation>. Accessed: 2026-04-16. 2025.
- [38] Zdzislaw Sliwa, Hans Helseth, and Viljar Veebel. "The Baltic Sea islands and their impact on the regional security". In: *Centrum Balticum, BSR Policy Briefing series* (2022).
- [39] Submarine Networks. *EU Publishes Landmark Report and Funding for Cable Hubs*. <https://www.submarinenetworks.com/en/nv/insights/eu-publishes-landmark-report-and-funding-for-cable-hubs>. Accessed: 2026-04-16. 2025.
- [40] Energinet and 50Hertz. *Bornholm Energy Island*. <https://bornholmenergyisland.eu/en/>. Accessed: 2026-04-14. 2026.
- [41] Danish Energy Agency. *Bornholm Energy Island*. <https://ens.dk/en/energy-sources/bornholm-energy-island>. Accessed: 2026-04-16. 2025.
- [42] C Kavanagh, J Franken, and W He. *Achieving Depth: Subsea Telecommunications Cables as Critical Infrastructure*. Geneva: UNIDIR, 2025.
- [43] Adriana Ávila-Zúñiga-Nordfeld. *Coping with Sabotage and Seabed Security Threats in the Baltic Sea: A Regional Maritime Security Policy*. Hague Centre for Strategic Studies (HCSS), 2025.
- [44] Julian Pawlak. *Re-Thinking War in the Baltic Sea*. German Institute for Defence and Strategic Studies, 2025.
- [45] Christine Agius, Emil Edenborg, and Sanna Strand. "From a Sea of Peace to a NATO Lake: Gendered Constructions of Baltic Islands". In: *Critical Perspectives on NATO*. Bristol University Press, 2026, pp. 29–48.

- [46] Erin L. Murphy and Matt Pearl. *Chinas Underwater Power Play: The PRCs New Subsea Cable-Cutting Ship Spooks International Security Experts*. <https://www.csis.org/analysis/chinas-underwater-power-play-prcs-new-subsea-cable-cutting-ship-spooks-international>. Accessed: May 15, 2026. 2025.
- [47] Scott Savitz. *Uncrewed Maritime Vessels: Shaping Naval Power in Hybrid Threat Operations*. Working Paper 34. Helsinki, Finland: European Centre of Excellence for Countering Hybrid Threats (Hybrid CoE), Oct. 2024. URL: <https://www.hybridcoe.fi/publications/hybrid-coe-working-paper-34-uncrewed-maritime-vessels-shaping-naval-power-in-hybrid-threat-operations/>.
- [48] Erik Reichborn-Kjennerud and Patrick Cullen. *What is Hybrid Warfare?*. JSTOR, 2022.
- [49] Pierre Thévenin et al. *Protecting Maritime Infrastructure from Hybrid Threats: Legal Options*. Research Report 14. Helsinki, Finland: European Centre of Excellence for Countering Hybrid Threats (Hybrid CoE), Mar. 2025. URL: <https://www.hybridcoe.fi/publications/hybrid-coe-research-report-14-protecting-maritime-infrastructure-from-hybrid-threats-legal-options/>.
- [50] Danish Defence Intelligence Service. *Assessment of the Hybrid Threat Against Denmark*. Tech. rep. Accessed: 2026-04-16. Danish Defence Intelligence Service (FE), Oct. 2025. URL: <https://www.fe-ddis.dk/globalassets/fe/dokumenter/2025/trusselsvurderinger/-assessment-of-the-hybrid-threat-against-denmark.pdf>.
- [51] Danish Maritime Authority. “New Data on EU Sanction-Designated Tankers in Danish Waters for 2025”. In: *AFP via GCaptain* (2026). URL: <https://gcaptain.com/denmark-records-292-russian-shadow-fleet-tankers-passing-through-danish-straits/>.
- [52] Andrius Sytas. “Estonian navy says it tried to detain one of Russian ‘shadow fleet’ in Baltic Sea”. In: *Reuters* (2025). URL: <https://www.reuters.com/world/europe/estonian-navy-says-it-tried-detain-one-russian-shadow-fleet-baltic-sea-2025-05-15/> (visited on 05/21/2026).
- [53] Reuters. *What is known about the Nord Stream gas pipeline explosions*. Accessed: 2026-05-21. 2025. URL: <https://www.reuters.com/world/europe/what-is-known-about-nord-stream-gas-pipeline-explosions-2025-08-21/>.
- [54] Kathryn Armstrong and Vishala Sri-Pathma. *Finland investigates suspected sabotage of Baltic-connector gas pipeline*. Accessed: 2026-05-21. 2023. URL: <https://www.bbc.com/news/world-europe-67070389>.
- [55] Reuters. *Sweden asks Chinese ship Yi Peng 3 to move to Swedish waters, PM says*. Accessed: 2026-05-21. 2024. URL: <https://www.reuters.com/world/sweden-asks-chinese-ship-yi-peng-3-move-swedish-waters-pm-says-2024-11-26/>.

- [56] Reuters. *Finland charges Eagle S tanker captain, officers over cable cuts*. Accessed: 2026-05-21. 2025. URL: <https://www.reuters.com/markets/commodities/finland-charges-eagle-s-tanker-captain-officers-over-cable-cuts-2025-08-11/>.
- [57] Andrius Sytas and Johan Ahlander. *Baltic undersea cable damaged by external influence Sunday, Latvian broadcaster*. Accessed: 2026-05-21. 2025. URL: <https://www.reuters.com/world/europe/baltic-undersea-cable-damaged-by-external-influence-sunday-latvian-broadcaster-2025-01-26/>.
- [58] Reuters. *Finnish police release Russia-linked ship held in cable sabotage case*. Accessed: 2026-05-21. 2026. URL: <https://www.reuters.com/business/media-telecom/finnish-police-release-russia-linked-ship-held-in-cable-sabotage-case-2026-01-12/>.
- [59] Finnish Customs. *Finnish Customs has concluded the preliminary inquiry into a suspected sanctions violation by the crew of the vessel Fitburg*. Tulli Press Release. Published 07.01.2026. Jan. 2026. URL: <https://tulli.fi/en/-/finnish-customs-has-concluded-the-preliminary-inquiry-into-a-suspected-sanctions-violation-by-the-crew-of-the-vessel-fitburg>.
- [60] Helsinki Police Department. *Criminal Investigation into the Fitburg Vessel to Continue for Weeks*. Poliisi Newsroom. Published 05.01.2026. Jan. 2026. URL: <https://poliisi.fi/en/-/criminal-investigation-into-the-fitburg-vessel-to-continue-for-weeks>.
- [61] Finnish Navy. *Post regarding authorities' control of the vessel Fitburg in the Gulf of Finland*. X (formerly Twitter). Published 31.12.2025. Dec. 2025. URL: <https://x.com/Navyfi/status/2008142473076641892>.
- [62] SP Global Market Intelligence. "Maritime shadow fleet Formation, operation and continuing risk for sanctions compliance". In: (2025). URL: <https://www.spglobal.com/market-intelligence/en/news-insights/research/maritime-shadow-fleet-formation-operation-and-continuing-risk-for-sanctions-compliance-teams-2025>.
- [63] Sabine Knapp et al. "Maritime Risk Assessment for Portugal, Sweden and the Baltic Sea with an emphasis on shadow fleets". In: (2025).
- [64] S. Raghunandan. "Denmark is fed up with Russia's shadow fleet". In: *The World from PRX* (2024). URL: <https://theworld.org/stories/2024/12/24/denmark-is-fed-up-with-russias-shadow-fleet>.
- [65] Euronews. "Russia's 'shadow fleet' of tankers reported in Danish waters almost daily in 2025". In: *Euronews* (Feb. 2026). Accessed: 2026-04-16. URL: <https://www.euronews.com/2026/02/11/russias-shadow-fleet-of-tankers-reported-in-danish-waters-almost-daily-in-2025-denmark-say>.

- [66] European Security & Defence. “Denmark in Putin’s Crosshairs”. In: *European Security & Defence* (2025). Accessed: 2026-04-16. URL: <https://european-security.com/en/denmark-in-putins-crosshairs/>.
- [67] Danish Ministry of Defence. *Global Strategy for Maritime Security 2025–2028*. Tech. rep. Accessed: 2026-04-16. Danish Ministry of Defence, Jan. 2025. URL: <https://www.fmn.dk/globalassets/fmn/dokumenter/2025/-global-strategy-for-maritime-security-1-.pdf>.
- [68] NATO. *NATO Allies join forces to enhance the security of critical undersea infrastructure*. NATO News. Dec. 2024. URL: <https://www.nato.int/en/news-and-events/articles/news/2024/12/10/nato-allies-join-forces-to-enhance-the-security-of-critical-undersea-infrastructure>.
- [69] NATO Allied Command Transformation. *Task Force X-Baltic*. <https://www.act.nato.int/activities/tfx-b/>. Accessed: 2026-04-16. 2025.
- [70] Forsvaret and Danish Armed Forces. *Joint Expeditionary Force (JEF)*. Last updated March 18, 2024; Accessed: May 21, 2026. 2024. URL: <https://www.forsvaret.dk/en/roles-and-responsibilities/international-cooperation/jef/>.
- [71] Terje Aven. “Review and Discussion of Types of Risks in View of Different Risk Perspectives”. In: *Risk Analysis* 45.10 (2025), pp. 3276–3285.
- [72] Terje Aven. “On Why Practice Needs Generic Guidance on How to Define and Understand the Concept of Risk”. In: *Risk Analysis* 45.10 (2025), pp. 2964–2973.
- [73] James C Kinsey, Ryan M Eustice, and Louis L Whitcomb. “A survey of underwater vehicle navigation: Recent advances and new challenges”. In: *IFAC conference of manoeuvring and control of marine craft*. Vol. 88. Lisbon. 2006, pp. 1–12.
- [74] Saildrone. *6 Months, 92% Uptime: Saildrone Persistent On-Station Force in the Baltic Sea*. <https://www.saildrone.com/news/6-months-92-uptime-saildrone-persistent-on-station-force-baltic-sea>. Accessed: 2026-04-16. 2025.
- [75] Liam Paull et al. “AUV Navigation and Localization: A Review”. In: *IEEE Journal of Oceanic Engineering* 39.1 (2014), pp. 131–149. DOI: 10.1109/JOE.2013.2278891.
- [76] Hemani Kaushal and Georges Kaddoum. “Underwater Optical Wireless Communication”. In: *IEEE Access* 4 (2016), pp. 1518–1547. DOI: 10.1109/ACCESS.2016.2552538.
- [77] Andrej Androjna et al. “AIS data manipulation in the illicit global oil trade”. In: *Journal of marine science and engineering* 12.1 (2023), p. 6.
- [78] Fomekong Fomekong Rachel Merveille et al. “Advancements in Sensor Fusion for Underwater SLAM: A Review on Enhanced Navigation and Environmental Perception”. In: *Sensors* 24.23 (2024). ISSN: 1424-8220. DOI: 10.3390/s24237490. URL: <https://www.mdpi.com/1424-8220/24/23/7490>.

- [79] Matti Leppäranta and Kai Myrberg. *Physical Oceanography of the Baltic Sea*. Springer Praxis Books. Berlin, Heidelberg: Springer, 2009, p. 378. ISBN: 978-3-540-79702-9. DOI: 10.1007/978-3-540-79703-6.
- [80] Kjell Hausken, Jonathan W Welburn, and Jun Zhuang. “A Review of Game Theory and Risk and Reliability Analysis in Infrastructures and Networks”. In: *Reliability Engineering & System Safety* (2025), p. 111123.
- [81] Saildrone, Inc. *Saildrone Launches the Future of Maritime Surveillance in the Baltic Sea*. Saildrone News. Accessed: 2026-05-30. June 2025. URL: <https://www.saildrone.com/news/saildrone-launches-the-future-of-maritime-surveillance-in-the-baltic-sea>.
- [82] Army Recognition. “Denmark Starts First Mission with US-Made Saildrone Voyager Maritime Drones to Secure Baltic Sea”. In: *Army Recognition* (2025). Accessed: 2026-04-16. URL: <https://www.armyrecognition.com/news/navy-news/2025/denmark-starts-first-mission-with-us-made-saildrone-voyager-maritime-drones-to-secure-baltic-sea>.
- [83] Naval News. “NATO Establishes New Task Force X to Counter Underwater Threats”. In: *Naval News* (Feb. 2025). Accessed: 2026-04-16. URL: <https://www.navalnews.com/naval-news/2025/02/nato-establishes-new-task-force-x-to-counter-underwater-threats/>.
- [84] Uncrewed Systems Technology. *TUCO Marine Group ProZero USVs*. <https://www.uncrewed-systems.com/tuco-marine-group-prozero-usvs/>. Accessed: 2026-04-16. 2025.
- [85] SH Defence. *The Cube: Containerized Multi-Mission Modules*. <https://shdefence.com/the-cube/>. Accessed: 2026-04-16. 2025.
- [86] Semco Maritime. *Stormborn: All-Weather Unmanned Surface Vehicle for Maritime Surveillance*. <https://www.semcomaritime.com/news/stormborn>. Accessed: 2026-04-16. 2025.
- [87] Technical University of Denmark. *Autonomous Marine Robotics to Save Lives at Sea: LORELEI-X*. <https://electro.dtu.dk/newsarchive/2025/02/autonomous-marine-robotics-to-save-lives-at-sea>. Accessed: 2026-04-16. Feb. 2025.
- [88] Maritime Danmark. “Hemmelig dansk drone afsløret på våbenmesse”. Danish. In: *Maritime Danmark* (2025). Accessed: May 18, 2026. URL: <https://www.maritimedanmark.dk/hemmelig-dansk-drone-afsloret-pa-vabenmesse>.
- [89] Andrew Reid. “ROV Market Prospects”. In: Douglas-Westwood, Aberdeen, 2013.
- [90] Ioseba Tena. “Automating roV operations in aid of the oil & gas offshore industry”. In: *SeeByte Whitepaper* (2011), pp. 1–9.

- [91] Dong Ma et al. "The state of the art in key technologies for autonomous underwater vehicles: a review". In: *Engineering* (2025). ISSN: 2095-8099. DOI: 10.1016/j.eng.2025.08.002. URL: <https://www.sciencedirect.com/science/article/pii/S209580992500445X>.
- [92] Sigurd Klemmensen et al. "COALA - A NOVEL APPROACH TO OFFSHORE TUBULAR INSPECTION AND CLEANING". In: *OCEANS 2023, Limerick*. 2023, In press. ISBN: 9781665416894.
- [93] Fredrik Fogh Sørensen et al. "A Quantitative Parametric Study on Output Time Delays for Autonomous Underwater Cleaning Operations". English. In: *Journal of Marine Science and Engineering* 10.6 (6 June 2022), pp. 1–26. ISSN: 2077-1312. DOI: 10.3390/jmse10060815.
- [94] Xiangen Bai et al. "A review of current research and advances in unmanned surface vehicles". In: *Journal of Marine Science and Application* 21.2 (2022), pp. 47–58.
- [95] Fredrik Fogh Sørensen et al. "Optical and Acoustic Imaging Comparison in a Controlled Underwater Environment". In: *Proceedings of 2023 OCEANS Conference & Exposition, Oceans 2023 Limerick*. IEEE, Feb. 2023, In press.
- [96] James Pita et al. "Using game theory for Los Angeles airport security". In: *AI magazine* 30.1 (2009), pp. 43–43.
- [97] Eric Shieh et al. "PROTECT: An application of computational game theory for the security of the ports of the United States". In: *Proceedings of the AAI Conference on Artificial Intelligence*. Vol. 26. 1. 2012, pp. 2173–2179.
- [98] Bo An et al. "PROTECT: A Deployed Game-Theoretic System for Strategic Security Allocation for the United States Coast Guard". In: *AI Magazine* 33.4 (2012), pp. 96–110. DOI: 10.1609/aimag.v33i4.2401. URL: <https://ojs.aaai.org/aimagazine/index.php/aimagazine/article/view/2401>.
- [99] James Pita et al. "GUARDS - Game Theoretic Security Allocation on a National Scale". In: *Proceedings of the 10th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*. 2011, pp. 37–44.
- [100] Zhengyu Yin et al. "TRUSTS: Scheduling randomized patrols for fare inspection in transit systems using game theory". In: *AI magazine* 33.4 (2012), pp. 59–59.
- [101] Amirali Rezazadeh et al. "Optimal patrol scheduling of hazardous pipelines using game theory". In: *Process Safety and Environmental Protection* 109 (2017), pp. 242–256. ISSN: 0957-5820. DOI: <https://doi.org/10.1016/j.psep.2017.03.039>. URL: <https://www.sciencedirect.com/science/article/pii/S0957582017301052>.
- [102] Xiaoming Duan, Dario Paccagnan, and Francesco Bullo. "Stochastic strategies for robotic surveillance as stackelberg games". In: *IEEE Transactions on Control of Network Systems* 8.2 (2021), pp. 769–780.

- [103] Hao-Tsung Yang et al. "Patrol Security Game: Defending Against Adversary with Freedom in Attack Timing, Location, and Duration". In: *ACM Transactions on Cyber-Physical Systems* 10.2 (2026), pp. 1–26.
- [104] Sukanya Samanta, Goutam Sen, and Soumya Kanti Ghosh. "A literature review on police patrolling problems". In: *Annals of Operations Research* 316.2 (2022), pp. 1063–1106.
- [105] Yusuf Ihsan Tokel and Jun Zhuang. *Optimizing Security Patrolling Strategies: A Cross-Domain Review of Mathematical Models and Applications*. Springer Nature, 2026.
- [106] Nicola Basilico. "Recent trends in robotic patrolling". In: *Current Robotics Reports* 3.2 (2022), pp. 65–76.
- [107] Steve Alpern et al. "Continuous Patrolling Games". In: *Operations Research* 70.6 (2022), pp. 3076–3089. DOI: 10.1287/opre.2022.2346. eprint: <https://doi.org/10.1287/opre.2022.2346>. URL: <https://doi.org/10.1287/opre.2022.2346>.
- [108] Jinpeng Han et al. "Digger: A Graph Contraction Algorithm for Patrolling Games". In: *IEEE Transactions on Reliability* 73.2 (2024), pp. 1104–1116. DOI: 10.1109/TR.2023.3329958.
- [109] NATO. *NATO launches Baltic Sentry to increase critical infrastructure security*. NATO News. Jan. 2025. URL: <https://www.nato.int/en/news-and-events/articles/news/2025/01/14/nato-launches-baltic-sentry-to-increase-critical-infrastructure-security>.
- [110] European Commission. *New York Joint Statement on the Security and Resilience of Undersea Cables in a Globally Digitalized World*. Official Policy Publication. Sept. 2024. URL: <https://digital-strategy.ec.europa.eu/en/library/new-york-joint-statement-security-and-resilience-undersea-cables-globally-digitalized-world>.
- [111] European Commission and High Representative of the Union for Foreign Affairs and Security Policy. *Joint Communication to the European Parliament and the Council: EU Action Plan on Cable Security*. Joint Communication JOIN(2025) 9 final. Brussels: European Union, Feb. 2025. URL: <https://eur-lex.europa.eu/legal-content/EN/TXT/HTML/?uri=CELEX:52025JC0009>.
- [112] Steve Alpern, Alec Morton, and Katerina Papadaki. "Patrolling Games". In: *Operations Research* 59.5 (2011), pp. 1246–1257. DOI: 10.1287/opre.1110.0983. URL: <https://pubsonline.informs.org/doi/abs/10.1287/opre.1110.0983>.
- [113] Praveen Paruchuri et al. "Playing Games for Security: An Efficient Exact Algorithm for Solving Bayesian Stackelberg Games". In: *Proceedings of the 7th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2008)*. Estoril, Portugal, 2008, pp. 895–902. URL: <https://dl.acm.org/doi/10.5555/1402298.1402348>.

- [114] Laobing Zhang and Genserik Reniers. “Applying a Bayesian Stackelberg game for securing a chemical plant”. In: *Journal of loss prevention in the process industries* 51 (2018), pp. 72–83.
- [115] Joint Research Centre. *Subsea cables: how vulnerable are they and can we protect them?* https://joint-research-centre.ec.europa.eu/jrc-explains/subsea-cables-how-vulnerable-are-they-and-can-we-protect-them_en. Accessed: 2026-04-14. European Commission, 2025.
- [116] Jacob Østergaard and John Nielsen. “THE BORNHOLM POWER SYSTEM An overview”. In: 2010. URL: <https://api.semanticscholar.org/CorpusID:1621246>.
- [117] Calash. *Subsea Infrastructure Failures*. Tech. rep. Accessed: 2026-04-14. Calash, 2025. URL: <https://www.calash.com/wp-content/uploads/2025/07/Subsea-Infrastructure-Failures.pdf>.
- [118] Recorded Future Insikt Group. *Threat Analysis: TA-2025-0717*. Tech. rep. TA-2025-0717. Accessed: 2026-04-16. Recorded Future, 2025. URL: <https://assets.recordedfuture.com/insikt-report-pdfs/2025/ta-2025-0717.pdf>.
- [119] Office of Gas and Electricity Markets (Ofgem). *Integrated Transmission Planning and Regulation (ITPR) project: Draft conclusions*. Tech. rep. Accessed: 2026-04-16. Ofgem, 2014. URL: https://www.ofgem.gov.uk/sites/default/files/docs/2014/04/isp_offshore_tc_0.pdf.
- [120] Europacable. *Introduction to HVDC Subsea Cables*. Tech. rep. Accessed: 2026-04-17. Europacable, 2012. URL: https://europacable.eu/wp-content/uploads/2021/01/Introduction-to-HVDC-Subsea-Cables-16-July-2012_.pdf.
- [121] Martin Murphy. *Redundancy and Resiliency: Securing Subsea Infrastructure*. Tech. rep. Accessed: 2026-04-16. Center for Strategic and International Studies (CSIS), 2025. URL: https://csis-website-prod.s3.amazonaws.com/s3fs-public/2025-11/251120_Murphy_Redundancy_Resiliency.pdf.
- [122] Østkraft. *Denmark: The Bornholm Power System – An Overview*. Tech. rep. Accessed: 2026-04-16. Global Islands Network, 2011. URL: http://www.globalislands.net/greenislands/docs/denmark_the_bornholm_power_system_an_overview.pdf.
- [123] João Tomé, David Belson, and Jorge S. Castro. *Resilient internet connectivity after the Baltic Sea cable cuts*. Cloudflare Blog. Accessed: 2026-04-16. 2024. URL: <https://blog.cloudflare.com/resilient-internet-connectivity-baltic-cable-cuts/>.
- [124] Telegeography. *Submarine Cable Map 2025*. <https://submarine-cable-map-2025.telegeography.com/>. Accessed: 2025-11-04. 2025.

- [125] Martin Hromada, David Rehak, and Ludek Lukas. "Resilience Assessment in Electricity Critical Infrastructure from the Point of View of Converged Security". In: *Energies* 14.6 (2021). ISSN: 1996-1073. DOI: 10.3390/en14061624. URL: <https://www.mdpi.com/1996-1073/14/6/1624>.
- [126] Nexans Norway. *Product Datasheet: 10562299*. Accessed: 2026-04-17. Nexans. 2026. URL: <https://www.nexans.no/.rest/catalog/v1/product/pdf/10562299>.
- [127] Prysmian Group. *525 kV HVDC Submarine Cable System with XLPE Insulation: Datasheet*. Accessed: 2026-04-17. Prysmian Group. 2023. URL: <https://www.prysmian.com/staticres/525-kv-hvdc-new-cable-systems/documents/XLPE-DATASHEET---Submarine.pdf>.
- [128] ABB Power Grids. *5SNA 1200G450300 HiPak IGBT Module Datasheet*. Accessed: 2026-04-17. ABB Switzerland Ltd, Semiconductors. 2013. URL: https://www.ic72.com/pdf_file/o/611053.pdf.
- [129] Geological Survey of Finland (GTK) and EMODnet Geology. *EMODnet Geology, Seabed Substrate (Multiscale, Folk 5 classification)*. Accessed: 2026-04-27. 2021. URL: <https://emodnet.ec.europa.eu/geonetwork/srv/eng/catalog.search#/metadata/67602462-9e00-40e6-98d5-1f560a010855>.
- [130] Ellie Moore, Stuart Haigh, and Geoffrey Eichhorn. "Anchor penetration depth in sandy soils and its implications for cable burial". In: *Ocean Engineering* 235 (Sept. 2021), p. 109411. DOI: 10.1016/j.oceaneng.2021.109411.
- [131] YU Sharif et al. "Comparison of the behaviour of a drag embedment anchor using 1-g and centrifuge scale model testing". In: *Proc. 5th Eur. Conf. Phys. Modeling Geotech.(ECPMG)*. 2024, pp. 1–6.
- [132] Lawrence Sanford. "Modeling a dynamically varying mixed sediment bed with erosion, deposition, bioturbation, consolidation, and armoring". In: *Computers Geosciences* 34 (Oct. 2008), pp. 1263–1283. DOI: 10.1016/j.cageo.2008.02.011.
- [133] Environment Agency. *Influence of Permeability on the Performance of Shingle and Mixed Beaches: Technical Report*. Tech. rep. SC060026. Accessed: 2026-04-17. Department for Environment, Food & Rural Affairs (DEFRA), 2021. URL: https://assets.publishing.service.gov.uk/media/602d08ccd3bf7f721c13a3b5/Influence_of_permeability_on_the_performance_of_shingle_and_mixed_beaches_technical_report.pdf.
- [134] Mark Severy and Arne Jacobson. *California North Coast Offshore Wind Studies: Subsea Cable Landfall and Onshore Main Transmission*. Tech. rep. 2020-OSW-R17. Accessed: 2026-04-17. Schatz Energy Research Center, 2020. URL: <https://schatzcenter.org/pubs/2020-OSW-R17.pdf>.

- [135] Cogea and EMODnet Human Activities. *Vessel Density Map*. Dataset based on AIS data from CLS and ORBCOMM. 2024. URL: <https://emodnet.ec.europa.eu/geonetwork/srv/eng/catalog.search#/metadata/0f2f3ff1-30ef-49e1-96e7-8ca78d58a07c>.
- [136] Julio B. Clempner. “Learning Deceptive Tactics for Defense and Attack in Bayesian-Markov Stackelberg Security Games”. In: *Mathematical and Computational Applications* 30.2 (2025). ISSN: 2297-8747. DOI: 10.3390/mca30020029. URL: <https://www.mdpi.com/2297-8747/30/2/29>.
- [137] Christian Kroer, Gabriele Farina, and Tuomas Sandholm. “Robust Stackelberg equilibria in extensive-form games and extension to limited lookahead”. In: *Proceedings of the AAAI Conference on Artificial Intelligence*. Vol. 32. 1. 2018.
- [138] Geraldo Lima Filho et al. “A Novel Bias-TSP Algorithm for Maritime Patrol”. In: *IEEE Access PP* (Jan. 2023), pp. 1–1. DOI: 10.1109/ACCESS.2023.3252013.
- [139] Douglas Burnett and Lionel Carter. “International Submarine Cables and Biodiversity of Areas Beyond National Jurisdiction: The Cloud Beneath the Sea”. In: *Brill Research Perspectives in the Law of the Sea* 1 (July 2017), pp. 1–72. DOI: 10.1163/24519359-12340002.
- [140] Jiarui Gan et al. “Robust stackelberg equilibria”. In: *Mathematical Programming* (2025), pp. 1–41.
- [141] Jacob K. Goeree, Charles A. Holt, and Thomas R. Palfrey. *Quantal Response Equilibrium: A Stochastic Theory of Games*. Princeton, NJ: Princeton University Press, 2016.
- [142] M. D. McKay, R. J. Beckman, and W. J. Conover. “A Comparison of Three Methods for Selecting Values of Input Variables in the Analysis of Output from a Computer Code”. In: *Technometrics* 21.2 (1979), pp. 239–245. DOI: 10.1080/00401706.1979.10489755.
- [143] Simeone Marino et al. “A Methodology for Performing Global Uncertainty and Sensitivity Analysis in Systems Biology”. In: *Journal of Theoretical Biology* 254.1 (2008), pp. 178–196. DOI: 10.1016/j.jtbi.2008.04.011.
- [144] Steven Diamond and Stephen P. Boyd. “CVXPY: A Python-Embedded Modeling Language for Convex Optimization”. In: *Journal of Machine Learning Research* 17.83 (2016), pp. 1–5. URL: <https://arxiv.org/abs/1603.00943>.
- [145] Alexander Domahidi, Eric Chu, and Stephen P. Boyd. “ECOS: An SOCP Solver for Embedded Systems”. In: *European Control Conference (ECC)*. 2013, pp. 3071–3076. DOI: 10.23919/ECC.2013.6669541.
- [146] Stefan Rass and Stefan Schauer, eds. *Game Theory for Security and Risk Management: From Theory to Practice*. Static & Dynamic Game Theory: Foundations & Applications. Birkhäuser, 2018. DOI: 10.1007/978-3-319-75268-6.

- [147] James Pita et al. “Deployed ARMOR protection: The application of a game-theoretic model for security at the Los Angeles International Airport”. In: *Proceedings of the 7th International Joint Conference on Autonomous Agents and Multiagent Systems (AA-MAS), Industry Track*. 2008, pp. 125–132.
- [148] Lawrence D. Stone. *Theory of Optimal Search*. 2nd. Arlington, VA: Military Applications Section, Operations Research Society of America, 1989.
- [149] Terje Aven. “Identification of safety and security critical systems and activities”. In: *Reliability Engineering & System Safety* 94.3 (2009), pp. 404–411. DOI: 10.1016/j.res.2008.04.001.
- [150] Henrik Ringbom. “New threatsold rules: Law of the sea issues raised by suspected attacks on submarine infrastructure in the Baltic Sea”. In: *Ocean Development & International Law* 56.3 (2025), pp. 390–414.
- [151] Yudong Hu et al. “Modeling opponent learning in multiagent repeated games”. In: *Applied Intelligence* 53.13 (2023), pp. 17194–17210.
- [152] Eric A. Hansen, Daniel S. Bernstein, and Shlomo Zilberstein. “Dynamic Programming for Partially Observable Stochastic Games”. In: *Proceedings of the Nineteenth National Conference on Artificial Intelligence*. AAAI-04. AAAI Press, 2004, pp. 709–715. DOI: VERIFY.
- [153] Tyler Becker and Zachary Sunberg. “Bridging the gap between partially observable stochastic games and sparse POMDP methods”. In: *arXiv preprint arXiv:2405.18703* (2024).
- [154] Martin Zinkevich et al. “Regret Minimization in Games with Incomplete Information”. In: *Advances in Neural Information Processing Systems*. Vol. 20. Curran Associates, Inc., 2007. DOI: VERIFY.
- [155] Stefan Rass and Stefan Schauer. “Game theory for security and risk management”. In: *Springer International Publishing*. doi 10 (2018), pp. 978–3.
- [156] Terje Aven. “Risk assessment and risk management: Review of recent advances on their foundation”. In: *European journal of operational research* 253.1 (2016), pp. 1–13.
- [157] Terje Aven. “How some types of risk assessments can support resilience analysis and management”. In: *Reliability Engineering & System Safety* 167 (2017), pp. 536–543.
- [158] Tom McLeod Logan et al. “Risk science offers an integrated approach to resilience”. In: *Nature Sustainability* 5.9 (2022), pp. 741–748.
- [159] Ronald K Mitchell, Bradley R Agle, and Donna J Wood. “Toward a theory of stakeholder identification and salience: Defining the principle of who and what really counts”. In: *Academy of Management Review* 22.4 (1997), pp. 853–886.
- [160] Stefan Olander. “Stakeholder impact analysis in construction project management”. In: *Construction Management and Economics* 25.3 (2007), pp. 277–287.

- [161] United Nations. *United Nations Convention on the Law of the Sea*. Accessed: 2025-04-28. 1982. URL: https://www.un.org/depts/los/convention_agreements/texts/unclos/unclos_e.pdf.
- [162] Yen-Chiang Chang, Chao Zhang, and Nannan Wang. "The international legal status of the unmanned maritime vehicles". In: *Marine Policy* 113 (2020), p. 103830. ISSN: 0308-597X. DOI: <https://doi.org/10.1016/j.marpol.2020.103830>. URL: <https://www.sciencedirect.com/science/article/pii/S0308597X19304774>.
- [163] International Maritime Organization. *Convention on the International Regulations for Preventing Collisions at Sea*. Accessed: 2025-04-28. 1972. URL: <https://www.imo.org/en/OurWork/Safety/Pages/Preventing-Collisions.aspx>.
- [164] International Cable Protection Committee. *Convention for the Protection of Submarine Telegraph Cables*. 1884. URL: https://iscpc.org/information/Convention_on_Protection_of_Cables_1884.pdf.
- [165] Graham Evans. *ICPC Activities Affecting HSSC*. Report HSSC16-07.10A. Presented at the 16th Meeting of the IHO Hydrographic Services and Standards Committee (HSSC-16). Tokyo, Japan: International Cable Protection Committee (ICPC), May 2024. URL: https://iho.int/uploads/user/Services%20and%20Standards/HSSC/HSSC16/HSSC16_2024_07.10A_EN_ICPC%20activities%20affecting%20HSSC.pdf.
- [166] International Cable Protection Committee (ICPC). *Government Best Practices for Protecting and Promoting Resilience of Submarine Telecommunications Cables*. Recommendation. International Cable Protection Committee, 2024. URL: <https://www.iscpc.org/>.

Appendix A

Risk and Uncertainty Framework Definitions

Table A.1 summarises the core variables of Aven's risk framework and their relevance to the Baltic maritime security context addressed in this thesis.

Concept	Notation	Definition	Baltic maritime relevance
Risk	(A, C, U)	Uncertainty about, and severity of, the consequences of an activity with respect to what humans value [28]	Umbrella concept capturing both deliberate sabotage and accidental damage to critical subsea infrastructure
Actions / events	A	Hazardous actions or initiating events that may affect the system	Anchor-drag manuvres, explosive demolition, AIS-dark transits in cable corridors
Consequences	C	Outcomes resulting from A , expressed in terms of human values	Telecommunications outage, energy supply disruption, cascading economic and security costs
Uncertainty	U	Epistemic and aleatory lack of knowledge about A and C [156]	Opacity of the subsurface environment; unverifiable intent of surface vessels operating in shipping lanes
Risk description	(C', Q, K)	Analyst-produced characterisation via predicted consequences C' , associated probabilities Q , and supporting background knowledge K [32]	Operational framework for translating patrol sensor data into quantified coverage priorities
Probability	Q	Subjective measure of likelihood, conditional on K	Estimated probability that a vessel executes an anchor-drag given observed loitering behaviour
Background knowledge	K	Totality of data, models, expert judgement, and assumptions underpinning Q and C'	AIS records, DDIS threat assessments, incident databases, substrate and depth data
Strength of Knowledge	SoK	Qualitative rating of the robustness and completeness of K (Strong / Moderate / Weak) [34]	Assessed as Moderate overall: strong expert consensus on the threat landscape, weak observational coverage of individual vessel intent
Vulnerability	—	Conditional severity of consequences given that a threat event materialises [32]	Determined by cable burial depth, redundancy paths, and proximity of landing stations to potential attack vectors
Resilience	—	Capacity of the system to absorb disruption, maintain essential functions, and recover to an acceptable operational state [157, 158]	Governed by cable repair vessel availability, MTTR, and inter-operator redundancy agreements

Table A.1: Risk and uncertainty framework definitions and their relevance to Baltic subsea infrastructure protection.

Appendix B

Strength of Knowledge Assessment of the DOBSS Model and its Assumptions

Table B.1 presents the Strength of Knowledge (SoK) assessment for the five principal input categories of the DOBSS model, evaluated against the framework of Askeland et al. [34]. The purpose of this assessment is to characterise the epistemic confidence behind each model input *before* the sensitivity analysis, providing a principled basis for understanding which outputs are reliable and which are model-conditional. The five inputs are assessed on the standard dimensions of reasonability of assumptions, data relevancy, expert agreement, phenomenological understanding, and examination of knowledge, collapsed here into a single per-input rating for operational clarity. The resulting ratings map directly onto the three-level robustness hierarchy established in Section 4.7: Strong-SoK inputs produce the unconditionally robust patrol target, Moderate-SoK inputs produce the conditionally reliable patrol allocation, and Weak-SoK inputs produce utility values that carry no direct operational interpretation.

Table B.1: Strength of Knowledge assessment for the principal inputs of the DOBSS model, following Aske-land et al. [34]. The rightmost column identifies the corresponding sensitivity dimension in Chapter 4.

Model Input	Rating	Justification	Sensitivity ref.
Spatial data (AIS density, substrate, cable routes)	Strong	EMODnet AIS rasters, TeleGeography cable routes, and Folk 5 substrate classification are publicly available, peer-reviewed, and cover the full operational resolution. No material gaps exist in the spatial input layer. [135, 10]	Section 4.7.5
Asset values (AV: power = 45, telecom = 4)	Moderate	Derived from objectively measurable parameters: MTTR, redundancy, and strategic weight. The 11.25:1 ratio is grounded in the Bornholm cable’s 90-day MTTR and zero-redundancy status, but the mapping to utility units involves analyst judgement without independent validation. [13]	Section 4.7.5
Threat-type priors $p(\theta_1)=0.86$, $p(\theta_2)=0.12$, $p(\theta_3)=0.02$	Moderate	Prior ordering is strongly supported by ICPC incident statistics and Baltic intelligence assessments [50, 9], but exact values are analyst judgements. The prior sweep confirms \mathbf{x}^* is spatially invariant over $p(\theta_3) \in [0.01, 0.25]$, limiting the operational consequence of prior uncertainty.	Section 4.7.3
Utility parameters α , β , γ	Weak	All three are expert point estimates without empirical calibration. β (attacker retained reward) and γ (loss scaling) have no measurement analogue in the maritime domain. Monte Carlo PRCC confirms both exceed $ 0.7 $ in output influence; β drives active-node count and γ drives absolute utility scale. [143]	Section 4.7.4
Attacker rationality ($\lambda \rightarrow \infty$, SSE)	Weak	The SSE fully-rational assumption is plausible for θ_3 but unvalidated against observed Baltic sabotage behaviour. Empirical QRE calibration yields $\lambda \in [0.5, 3]$ for human subjects [141], and SSE tie-breaking introduces a residual gap between the MILP objective and high- λ QRE evaluation. [140]	Section 4.7.2
Overall SoK	Moderate	Strong spatial data and a robust asset-value ordering anchor the patrol target unconditionally. Weak calibration of α , β , γ , and λ makes the patrol allocation and absolute utility values model-conditional outputs requiring empirical calibration before operational deployment.	—

Appendix C

Stakeholder Analysis

Stakeholder Identification and Attributes

Nine primary stakeholder groups are identified using Mitchell's Saliency Model [159]. Each group is assessed on three attributes: *Power* (P) the ability to impose the stakeholder's will on the project; *Legitimacy* (L) the normative appropriateness of the stakeholder's involvement; and *Urgency* (U) the degree to which the stakeholder's claims require immediate attention. Each attribute is rated on a scale of 1 (low) to 3 (high).

SII Methodology

Olander's Stakeholder Impact Index [160] extends Mitchell's salience assessment by weighting stakeholder influence against their attitude toward the project. Three quantities are computed for each stakeholder i :

$$A_i = P_i + L_i + U_i \quad (\text{C.1})$$

$$\text{ViII}_i = \frac{A_i}{\sum_{j=1}^n A_j} \quad (\text{C.2})$$

$$\text{SII}_i = \text{ViII}_i \times \text{Att}_i \quad (\text{C.3})$$

where A_i is the salience score, ViII_i is the normalised salience weight (Visibility \times Impact \times Influence Index), $\text{Att}_i \in \{1, 2, 3\}$ is the assessed stakeholder attitude toward autonomous SCI monitoring (1 = opposed, 2 = neutral, 3 = supportive), and the Project Impact Index is $\text{PII} = \sum_i \text{SII}_i$.

Quantitative Assessment and SII Output

Stakeholder	P	L	U	A_i	$ViII_i$	Att_i	SII_i
Danish Maritime Authority (DMA)	3	3	3	9	0.150	3	0.45
Danish Ministry of Defence (MoD)	3	3	3	9	0.150	3	0.45
Danish Defence Intelligence Service (DDIS)	2	3	3	8	0.133	3	0.40
Energinet / Submarine Cable Operators	2	3	3	8	0.133	2	0.27
NATO Allied Maritime Command (MARCOM)	2	2	2	6	0.100	2	0.20
European Commission (DG CONNECT)	2	3	2	7	0.117	2	0.23
Danish Parliament (Folketing)	3	2	1	6	0.100	1	0.10
Danish Fishing Industry	2	1	1	4	0.067	1	0.07
Environmental NGOs	1	1	1	3	0.050	1	0.05
Total				60	1.000		2.22

Table C.1: Stakeholder Impact Index (SII) assessment for the autonomous USV patrol framework.

The aggregate Project Impact Index of $PII = 2.22$ on a scale of 1–3 indicates a *favorable* stakeholder environment for deployment of an autonomous SCI monitoring capability.

Stakeholder Typology and Strategic Priority

Applying Mitchell’s typology to the SII results, stakeholders are classified into three strategic tiers:

- **Definitive stakeholders** (high P, L, U; top priority): The Danish Maritime Authority, Ministry of Defence, and DDIS each score $A_i = 8–9$ and hold supportive attitudes. Their regulatory mandate, enforcement authority, and intelligence function make them essential partners for operational deployment. Sustained engagement and co-development of data-sharing protocols is the recommended management strategy.
- **Dominant stakeholders** (high P and L; active management required): Submarine cable operators and the EU Commission hold significant legitimacy and power. Cable operators’ neutral-to-positive attitude reflects cost-sharing concerns balanced against direct protective benefit; EU Commission engagement is driven primarily by regulatory compliance obligations. Formal consultation within the EU Cable Security Toolbox process is the recommended interface.
- **Latent and marginal stakeholders** (lower overall salience; monitoring sufficient): NATO MARCOM and the European Parliament are supportive or neutral but less immediately engaged. The Danish fishing industry and environmental NGOs hold

low salience scores and opposing or neutral attitudes; their concerns relate primarily to operational constraints in shared maritime corridors and should be addressed through transparent communication about USV sea-keeping behaviour and COLREG compliance.

Appendix D

Regulatory Context

The stakeholder analysis for this project, conducted using Mitchell's Saliency Model [159] and Olander's Stakeholder Impact Index [160], is presented in Appendix C. The regulatory framework below establishes the legal and normative context within which the autonomous patrol system operates.

D.1 UNCLOS and Jurisdictional Challenges

The United Nations Convention on the Law of the Sea (UNCLOS) [161] serves as the primary legal constitution for the world's oceans. For the Baltic maritime domain, a region characterized by overlapping Exclusive Economic Zones (EEZ) and high-density international shipping lanes, UNCLOS defines the rights and obligations of the coastal state and third-party vessels.

Under Articles 58 and 79, all states enjoy the right to lay and maintain subsea cables in the EEZ and on the continental shelf. However, the legal status of Unmanned Surface Vehicles (USVs) remains an area of evolving interpretation [162]. Any deployed USV must adhere to the International Regulations for Preventing Collisions at Sea (COLREGs) [163], which govern navigational conduct and spatial constraints. The jurisdictional ambiguity of intentional damage in international waters necessitates that the security strategy focuses on deterrence and domain awareness rather than active kinetic interception, aligning with the "due regard" principle of UNCLOS.

D.2 EU Binding Directives: NIS2 and CER

Two EU directives constitute the primary binding regulatory framework for operators of submarine cable infrastructure and form the legislative basis for the security obligations that this project's patrol framework is designed to support. The Network and Information Security Directive (NIS2, Directive (EU) 2022/2555 [21]) extends mandatory cybersecu-

rity risk-management and incident-reporting obligations to operators of essential entities; proposed amendments from January 2026 would formally classify submarine data transmission infrastructure (SDTI) encompassing submarine cables, landing stations, and associated terrestrial infrastructure as a highly critical sector under Annex I. The Critical Entities Resilience Directive (CER, Directive (EU) 2022/2557 [22]) complements NIS2 by requiring Member States to identify critical entities, conduct national risk assessments, and adopt resilience strategies explicitly addressing submarine cable infrastructure, and by introducing the Cable Projects of European Interest (CPEI) category. Together, these directives are being transposed into Danish law, creating enforceable obligations for cable operators to implement active monitoring precisely the capability that the autonomous patrol system developed in this project delivers. The Cable Security Toolbox adopted in February 2026 [23] operationalises these obligations into six strategic and four technical measures, with active corridor monitoring designated as a priority technical measure under both the prevention and detection resilience stages.

D.3 ICPC Recommendations on Cable Protection Zones

To mitigate external human aggression threats such as bottom-trawling and anchoring, the International Cable Protection Committee (ICPC) [164, 165] advocates for the implementation of Cable Protection Zones (CPZs) [166]. A CPZ is a formalized maritime corridor where high-risk seabed activities are legally restricted. Crucially, while the ICPC supports protective spatial separation, it explicitly cautions against the regulatory funneling of multiple cables into mandatory, narrow corridors, as such clustering creates a single point of failure and increases vulnerability to localized catastrophic events.

Appendix E

Adversary Typology

The Bayesian Stackelberg Security Game formulated in this project models a set of distinct follower (attacker) types $\Theta = \{\theta_1, \theta_2, \theta_3\}$, each with a prior probability $p(\theta)$ and a characteristic rationality level λ^θ . Grounding these abstract types in the threat landscape established above is essential for model validity: the prior probabilities and rationality parameters must reflect empirical evidence rather than arbitrary assumptions. The three types are defined as follows.

θ_1 : Unintentional Fishing Actor ($\lambda \approx 0, p(\theta_1) = 0.86$) This type represents commercial fishing vessels whose utility is derived entirely from access to productive fishing grounds. Risk to subsea infrastructure arises from incidental anchor-drag, gear snagging, and poor situational awareness near cable corridors — the dominant accidental damage mechanism in global cable incident records [13]. Such actors have no awareness of patrol patterns and do not adapt strategically to detection risk. The near-zero rationality parameter reflects the absence of deliberate target selection. The prior $p(\theta_1) = 0.86$ reflects the historically high frequency of fishing-related cable incidents in the Baltic region and makes this the highest-frequency threat category by prior weight.

θ_2 : Unintentional Commercial Shipping Actor ($\lambda \approx 0, p(\theta_2) = 0.12$) This type represents commercial cargo and tanker operators whose utility is determined by efficient transit through established shipping lanes. Infrastructure damage is accidental, arising from anchor operations, drift in adverse weather, or navigational inaccuracy — not from deliberate intent to harm cables. Like θ_1 , this type generates risk independently of the patrol strategy. The prior $p(\theta_2) = 0.12$ reflects the significant but lower contribution of general commercial shipping to Baltic cable incidents relative to fishing activity.

θ_3 : Strategic State-Sponsored Actor ($\lambda \rightarrow \infty, p(\theta_3) = 0.02$) This type represents a fully rational, goal-maximizing adversary whose utility is derived from the strategic value of the

compromised asset, offset by physical execution costs and operational exposure risk. It encompasses both state-directed kinetic operations — corresponding to Russian GRU-linked or GUGI-coordinated activity documented in the Bornholm area [66, 50] — and coerced shadow fleet operators acting under state direction to execute anchor-drag manuvres, as documented in the *Eagle S* and *Yi Peng 3* incidents [12, 51]. Both sub-categories share the defining characteristic that target selection is deliberate and responsive to the defender’s patrol strategy, motivating the Stackelberg formulation. The DDIS HIGH classification for Baltic sabotage risk and forensic evidence from the Nord Stream investigation support the inclusion of this type [50, 11]. Despite the documented frequency of shadow fleet transits (292 EU-sanctioned passages through Danish waters in 2025 [51]), the prior $p(\theta_3) = 0.02$ reflects the low frequency of confirmed deliberate sabotage events relative to accidental incidents. This type accounts for all residual defender loss at high adversary rationality and is the primary driver of the Stackelberg equilibrium.

The specific prior probabilities $p(\theta_1), p(\theta_2), p(\theta_3)$ — subject to $\sum_{\theta} p(\theta) = 1$ — are specified and subjected to sensitivity analysis in the Problem Solution chapter. Table E.1 summarises the typology.

Type	Real-world analogue	Rationality λ	Evidence basis
θ_1	Fishing vessel / unintentional	≈ 0 (random)	$p = 0.86$; historical accidental cut rate [13]
θ_2	Commercial shipping / unintentional	≈ 0 (random)	$p = 0.12$; shipping-lane incident data [13]
θ_3	State-sponsored / shadow fleet actor	$\rightarrow \infty$ (fully rational)	$p = 0.02$; DDIS HIGH; Nord Stream, <i>Eagle S</i> , <i>Yi Peng 3</i> [50, 11, 12]

Table E.1: Adversary Typology for the BSSG Model.

Appendix F

Implementation Code

```
1 import geopandas as gpd
2 import pandas as pd
3 import numpy as np
4 import rioarray
5 import os
6
7 # We will load the data variables from the existing shapefileplot.py
8 from shapefileplot import grid_gdf_metric, cables_clipped, substrate_clipped, data_path,
   → get_substrate_color_mpl, bbox_polygon
9
10 # =====
11 # PHASE 1: DOMAIN DEFINITION & RISK-UTILITY TRANSLATION
12 # =====
13 alpha = 0.65 # Proportion of damage remaining when defended
14 beta = 0.1 # Proportion of attacker reward remaining when defended
15 defender_multiplier = 10.0 # Scaling factor for the non-zero-sum defender loss
16
17 print(f"Loaded {len(grid_gdf_metric)} grid nodes and {len(cables_clipped)} cables (including
   → OSM changes) from shapefileplot.")
18
19 # Initialize the discrete nodes dataframe
20 nodes_df = pd.DataFrame(grid_gdf_metric.drop(columns='geometry'))
21 nodes_df['Asset_Value'] = 0.0
22 nodes_df['Traffic_Density'] = 0.0
23 nodes_df['Attacker_Cost'] = 0.0
24
25 # 1. Map Asset Value to Nodes (Using SUM to account for multiple cables)
26 cables_metric = cables_clipped.to_crs(epsg=32633)
27 grid_cables_sjoin = gpd.sjoin(grid_gdf_metric, cables_metric, how='left',
   → predicate='intersects')
```

```

28
29 # Extract numeric values
30 def extract_numeric(val):
31     if pd.isna(val): return 0.0
32     return float(val.split(' ')[0])
33
34 grid_cables_sjoin['Asset_Value_num'] =
35     ↪ grid_cables_sjoin['Target_Value'].apply(extract_numeric)
36 # Aggregate (sum value if multiple cables in one cell)
37 sum_assets = grid_cables_sjoin.groupby('Node_ID')['Asset_Value_num'].sum().reset_index()
38 nodes_df = nodes_df.merge(sum_assets, on='Node_ID', how='left')
39 nodes_df['Asset_Value'] = nodes_df['Asset_Value_num'].fillna(0.0)
40 nodes_df.drop(columns=['Asset_Value_num'], inplace=True)
41
42 print(f"Max Asset Value located in a single node: {nodes_df['Asset_Value'].max()} (Verifying
43     ↪ OSM Power Cable inclusion)")
44
45 # 2. Map Substrate Cost to Nodes
46 substrate_metric = substrate_clipped.to_crs(epsg=32633)
47 grid_substrate_sjoin = gpd.sjoin(grid_gdf_metric, substrate_metric, how='left',
48     ↪ predicate='intersects')
49
50 def extract_cost(val):
51     if pd.isna(val) or val == 'Unknown': return 0.5 # mid cost if unknown
52     return float(val.split(' ')[1])
53
54 grid_substrate_sjoin['Attacker_Cost_num'] =
55     ↪ grid_substrate_sjoin['Attacker_Cost'].apply(extract_cost)
56 avg_costs =
57     ↪ grid_substrate_sjoin.groupby('Node_ID')['Attacker_Cost_num'].mean().reset_index()
58 nodes_df = nodes_df.merge(avg_costs, on='Node_ID', how='left')
59 nodes_df['Attacker_Cost'] = nodes_df['Attacker_Cost_num'].fillna(0.5)
60 nodes_df.drop(columns=['Attacker_Cost_num'], inplace=True)
61
62 # 3. Map Multi-Domain Traffic Density to Nodes (PHASE 2: ADVERSARY PROFILING)
63 def get_raster_means(tif_filename):
64     tif_path = os.path.join(data_path, tif_filename)
65     means = []
66     if os.path.exists(tif_path):
67         try:
68             with rioarray.open_rasterio(tif_path) as rds:
69                 # 1. Clip the massive raster to the WGS84 bounding box first
70                 clipped_main = rds.rio.clip([bbox_polygon], crs="EPSG:4326", drop=True)
71                 # 2. Reproject ONLY the small bounded area to metric
72                 clipped_main = clipped_main.rio.reproject("EPSG:32633")

```

```

68
69         for geom in grid_gdf_metric.geometry:
70             try:
71                 # 3. Clip individual cells from the small raster
72                 clipped_cell = clipped_main.rio.clip([geom], crs="EPSG:32633",
73                 ↪ drop=True)
74                 data = clipped_cell.values
75                 valid_data = data[(data >= 0) & (data < 1e10)]
76                 means.append(float(valid_data.mean()) if len(valid_data) > 0 else
77                 ↪ 0.0)
78             except Exception:
79                 means.append(0.0)
80     except Exception as e:
81         print(f"Error processing {tif_filename}: {e}")
82         means = [0.0] * len(grid_gdf_metric)
83     else:
84         print(f"Warning: {tif_filename} not found.")
85         means = [0.0] * len(grid_gdf_metric)
86
87     # Normalize [0, 1]
88     max_val = max(means) if max(means) > 0 else 1.0
89     return [m / max_val for m in means]
90
91 print("Calculating Traffic Density Priors (this may take a moment)...")
92 nodes_df['Density_Fishing'] = get_raster_means("vesseldensity_fishing_2023.tif")
93 cargo_means = get_raster_means("vesseldensity_cargo_2023.tif")
94 tanker_means = get_raster_means("vesseldensity_tanker_2023.tif")
95
96 # Combine Cargo and Tanker into general 'Shipping' layer for theta_2 and theta_3
97 nodes_df['Density_Shipping'] = [(c + t) / 2.0 for c, t in zip(cargo_means, tanker_means)]
98
99 # 4. Phase 2: Bayesian Adversary Parameterization
100 # Priors summing to 1.0
101 P_theta1 = 0.86 # Fishing Accident
102 P_theta2 = 0.12 # Shipping Accident
103 P_theta3 = 0.02 # Shadow Fleet Sabotage
104
105 # Store priors in separate numeric columns for easier programmatic use
106 nodes_df['P_theta1'] = P_theta1
107 nodes_df['P_theta2'] = P_theta2
108 nodes_df['P_theta3'] = P_theta3
109
110 # Defender Base Values (Independent of Attacker type)
111 nodes_df['Defender_Loss'] = nodes_df['Asset_Value'] * defender_multiplier
112 nodes_df['U_d_u'] = -nodes_df['Defender_Loss']

```

```

111 nodes_df['U_d_c'] = -alpha * nodes_df['Defender_Loss']
112
113 # -- Follower 1 (Fishing Trawlers) Utilities --
114 # Utility is purely derived from fishing grounds.
115 nodes_df['U_a_theta1_u'] = nodes_df['Density_Fishing']
116 nodes_df['U_a_theta1_c'] = nodes_df['Density_Fishing'] * beta # USV deterrence works
117
118 # -- Follower 2 (Cargo/Tanker) Utilities --
119 # Utility is purely derived from valid shipping transit lanes.
120 nodes_df['U_a_theta2_u'] = nodes_df['Density_Shipping']
121 nodes_df['U_a_theta2_c'] = nodes_df['Density_Shipping'] * beta # USV deterrence works
122
123 # -- Follower 3 (Shadow Fleet) Utilities --
124 # Strategic attacker seeking Asset_Value. Low shipping density drastically increases their
    ↪ exposure cost.
125 coverage_penalty = 1.0 - nodes_df['Density_Shipping']
126 nodes_df['Attacker_Cost_theta3'] = nodes_df['Attacker_Cost'] + (coverage_penalty * 0.5)
127
128 nodes_df['U_a_theta3_u'] = nodes_df['Asset_Value'] * (1.0 -
    ↪ nodes_df['Attacker_Cost_theta3'])
129 nodes_df['U_a_theta3_c'] = (beta * nodes_df['Asset_Value']) * (1.0 -
    ↪ nodes_df['Attacker_Cost_theta3'])
130
131 # Save Results
132 nodes_df.to_csv("dobss_utilities_phase2.csv", index=False)
133 print("\n=== Phase 2 Complete ===")
134 print("Bayesian matrix parameterized and saved to dobss_utilities_phase2.csv")
135
136 # Sort by Asset_Value descending to show the nodes directly incorporating the Bornholm power
    ↪ cable
137 print("\nTop 10 Nodes by Asset Value (Includes OSM OpenStreetMap Interconnection):")
138 print(nodes_df.sort_values(by='Asset_Value', ascending=False)[['Node_ID', 'Asset_Value',
    ↪ 'Density_Fishing', 'Density_Shipping', 'U_d_u', 'U_a_theta3_u']].head(10))

```

```

1 import pandas as pd
2 import pulp
3
4 # Load Phase 2 Utilities
5 df = pd.read_csv("dobss_utilities_phase2.csv")
6
7 # Parameters
8 k = 1.0 # Number of USVs available
9 # Big-M constraint for rational bounds. Compute a safe upper bound from the

```

```

10 # actual utility values in dobss_utilities_phase2.csv to keep the model valid
11 # if the data changes. We look at all columns whose names start with "U_",
12 # take the maximum absolute utility, and scale it by a factor to ensure that
13 # any utility difference used in the constraints is safely bounded.
14 util_cols = [col for col in df.columns if col.startswith("U_")]
15 if util_cols:
16     max_abs_utility = df[util_cols].abs().max().max()
17     # Use a conservative multiplier to keep M large enough without being excessive.
18     M = float(max_abs_utility * 10.0)
19 else:
20     # Fallback to the original manually chosen Big-M if no utility columns are found.
21     M = 100.0
22
23 # Prior probabilities
24 priors = {'t1': 0.86, 't2': 0.12, 't3': 0.02}
25 theta_types = ['t1', 't2', 't3']
26 nodes = df['Node_ID'].tolist()
27
28 # Prepare Utility Dictionaries for instantaneous O(1) mathematical lookup
29 U_d_u = dict(zip(df['Node_ID'], df['U_d_u']))
30 U_d_c = dict(zip(df['Node_ID'], df['U_d_c']))
31
32 U_a = {}
33 for theta in theta_types:
34     # We dynamically extract theta columns: U_a_theta1_u, etc.
35     u_a_u_col = f'U_a_theta{theta[-1]}_u'
36     u_a_c_col = f'U_a_theta{theta[-1]}_c'
37     U_a[theta] = {
38         'u': dict(zip(df['Node_ID'], df[u_a_u_col])),
39         'c': dict(zip(df['Node_ID'], df[u_a_c_col]))
40     }
41
42 # =====
43 # PHASE 3: DOBSS MILP FORMULATION
44 # =====
45 print("Formulating Deployed-to-Observed Bayesian Stackelberg Security Game...")
46 prob = pulp.LpProblem("DOBSS_USV_Patrol", pulp.LpMaximize)
47
48 # Decision Variables
49 # x_t: Continuous [0, 1] coverage probability for node t
50 x = pulp.LpVariable.dicts("x", nodes, lowBound=0, upBound=1, cat='Continuous')
51
52 # q_t^theta: Binary {0, 1} attacker strategy choice
53 q = pulp.LpVariable.dicts("q", (theta_types, nodes), cat='Binary')

```

```

54
55 # z_t^theta: Linearized bounded product of (x_t * q_t^theta)
56 z = pulp.LpVariable.dicts("z", (theta_types, nodes), lowBound=0, upBound=1,
    ↪ cat='Continuous')
57
58 # v^theta: Attacker expected utility (the rational equilibrium we calculate against)
59 v = pulp.LpVariable.dicts("v", theta_types, cat='Continuous')
60
61 # 1. Objective Function
62 # Maximize Defender Expected Utility
63 prob += pulp.lpSum([
64     priors[theta] * pulp.lpSum([
65         z[theta][t] * U_d_c[t] + (q[theta][t] - z[theta][t]) * U_d_u[t]
66         for t in nodes
67     ])
68     for theta in theta_types
69 ]), "Expected_Defender_Utility"
70
71 # 2. Constraints
72 # Defender Physical Resources (k USVs)
73 prob += pulp.lpSum([x[t] for t in nodes]) <= k, "Resource_Constraint"
74
75 for theta in theta_types:
76     # Attacker must rationally target exactly one node
77     prob += pulp.lpSum([q[theta][t] for t in nodes]) == 1, f"One_Target_{theta}"
78
79     for t in nodes:
80         # McCormick Envelopes to linearize z = x * q
81         prob += z[theta][t] <= x[t], f"McCormick_1_{theta}_{t}"
82         prob += z[theta][t] <= q[theta][t], f"McCormick_2_{theta}_{t}"
83         prob += z[theta][t] >= x[t] - (1 - q[theta][t]), f"McCormick_3_{theta}_{t}"
84
85         # Expected utility of attacking target t for attacker theta
86         E_U_a = U_a[theta]['u'][t] + x[t] * (U_a[theta]['c'][t] - U_a[theta]['u'][t])
87
88         # Big-M Strong Stackelberg Constraints
89         prob += v[theta] >= E_U_a, f"Best_Response_LB_{theta}_{t}"
90         prob += v[theta] <= E_U_a + M * (1 - q[theta][t]), f"Best_Response_UB_{theta}_{t}"
91
92 # 3. Solve Phase
93 print("Solving MILP structure via CBC solver...")
94 prob.solve(pulp.PULP_CBC_CMD(msg=False))
95
96 print(f"Mathematical Status: {pulp.LpStatus[prob.status]}")

```

```

97
98 # 4. Extract Stationary Output
99 results = []
100 for t in nodes:
101     if x[t].varValue is not None and x[t].varValue > 0.0001:
102         results.append({
103             'Node_ID': t,
104             'Coverage_Probability_x': round(x[t].varValue, 4)
105         })
106
107 df_res = pd.DataFrame(results).sort_values(by='Coverage_Probability_x', ascending=False)
108 df_res.to_csv("dobss_optimal_strategy.csv", index=False)
109
110 print("\n--- Optimal Stationary USV Patrol Distribution (x*) ---")
111 print(df_res.head(10))
112
113 print("\n--- Predicted Rational Attacker Action (q*) ---")
114 for theta in theta_types:
115     targets = [t for t in nodes if q[theta][t].varValue is not None and q[theta][t].varValue
116     → > 0.5]
117     target = targets[0] if targets else "None"
118     print(f"Threat {theta[-1]} (p={priors[theta]:.2f}) optimally targets Node: {target} |
119     → Saboteur Eq: {round(v[theta].varValue, 4) if v[theta].varValue is not None else 0}")

```

```

1 import pandas as pd
2 import geopandas as gpd
3 import numpy as np
4 import pulp
5 import random
6 import matplotlib.pyplot as plt
7 import contextily as ctx
8 from shapely.geometry import LineString
9
10 from shapfileplot import grid_gdf_metric, cables_clipped, bbox_wgs84
11
12 # =====
13 # PHASE 4: SPATIOTEMPORAL ROUTING (Markov Chain Synthesis)
14 # =====
15 # Parameters
16 v_kmh = 30.0
17 dt_hours = 1.0
18 max_dist_m = v_kmh * dt_hours * 1000.0 # 30,000 meters
19

```

```

20 print(f"Synthesizing Markov Transition Matrix P for {v_kmh} km/h over {dt_hours}h steps...")
21 print(f"Max reachability radius: {max_dist_m/1000} km per step")
22
23 # 1. Load the Stationary Distribution x*
24 df_res = pd.read_csv("dobss_optimal_strategy.csv")
25 x_dict = dict(zip(df_res['Node_ID'], df_res['Coverage_Probability_x']))
26
27 # Normalize x* to represent a true probability distribution (sum = 1) for a single USV
28 total_x = sum(x_dict.values())
29 pi = {k: v / total_x for k, v in x_dict.items()}
30 active_nodes = list(pi.keys()) # We only bother routing between nodes that require coverage
31
32 # 2. Compute Distances and Adjacency
33 centroids = {}
34 for idx, row in grid_gdf_metric.iterrows():
35     if row['Node_ID'] in active_nodes:
36         centroids[row['Node_ID']] = row.geometry.centroid
37
38 prob = pulp.LpProblem("USV_Markov_Routing", pulp.LpMinimize)
39
40 # Decision Variable: P_ij (Transition probability from i to j)
41 P = pulp.LpVariable.dicts("P", (active_nodes, active_nodes), lowBound=0, upBound=1,
42     ↪ cat='Continuous')
43
44 distances = {}
45 for i in active_nodes:
46     distances[i] = {}
47     for j in active_nodes:
48         d = centroids[i].distance(centroids[j])
49         distances[i][j] = d
50
51     # Enforce kinematic constraint: If too far, P_ij = 0
52     if d > max_dist_m:
53         prob += P[i][j] == 0, f"Kinematic_{i}_{j}"
54
55 # 3. Constraints
56 for i in active_nodes:
57     # Row stochasticity: Must go somewhere (or stay)
58     prob += pulp.lpSum([P[i][j] for j in active_nodes]) == 1.0, f"Stochastic_{i}"
59
60 for j in active_nodes:
61     # Stationarity: sum_i (pi_i * P_ij) = pi_j
62     prob += pulp.lpSum([pi[i] * P[i][j] for i in active_nodes]) == pi[j], f"Stationary_{j}"

```

```

63 # 4. Objective: Minimize expected transition distance to conserve fuel/energy
64 # while strictly maintaining Stackelberg Equilibrium stationarity
65 prob += pulp.lpSum([pi[i] * P[i][j] * distances[i][j] for i in active_nodes for j in
    ↪ active_nodes]), "Minimize_Fuel"
66
67 prob.solve(pulp.PULP_CBC_CMD(msg=False))
68 print(f"Markov Matrix Status: {pulp.LpStatus[prob.status]}")
69
70 # Extract Matrix
71 P_matrix = {i: {j: P[i][j].varValue for j in active_nodes} for i in active_nodes}
72
73 # =====
74 # SIMULATE 60-HOUR PATROL ROUTE
75 # =====
76 print("\nSimulating 60-Hour Patrol for USV-1 with Anti-Loitering Rules...")
77
78 # Helper to pick next node based on probabilities
79 def get_next_node(current_node, p_mat, banned_nodes=[]):
80     targets = list(p_mat[current_node].keys())
81     probs = [p_mat[current_node][t] if p_mat[current_node][t] is not None else 0 for t in
    ↪ targets]
82
83     # Temporarily zero out banned nodes
84     for i, t in enumerate(targets):
85         if t in banned_nodes:
86             probs[i] = 0.0
87
88     # Re-normalize if possible
89     prob_sum = sum(probs)
90     if prob_sum > 0:
91         probs = np.array(probs) / prob_sum
92         return np.random.choice(targets, p=probs)
93     else:
94         # Failsafe: if math forces us into a corner, try to move to a non-banned node
95         safe_targets = [t for t in targets if t not in banned_nodes]
96         if safe_targets:
97             return np.random.choice(safe_targets)
98         # If all immediate targets are banned, broaden the search to all nodes in the chain
99         all_nodes = list(p_mat.keys())
100        fallback_targets = [t for t in all_nodes if t not in banned_nodes and t !=
    ↪ current_node]
101        if fallback_targets:
102            return np.random.choice(fallback_targets)
103        # Ultimate fallback: no non-banned alternative exists anywhere; stay in place
104        return current_node

```

```

105
106 # Start USV-1 at the most critical node
107 current = df_res.iloc[0]['Node_ID']
108 route = [current]
109
110 loiter_count = 0
111 for step in range(60):
112     banned = []
113     # Enforce constraint: loiter_count tracks extra consecutive hours beyond the initial
114     ↪ hour at a node.
115     if loiter_count >= 1: # Already stayed 1 extra turn (total 2 hours), so must move away
116     ↪ from this node.
117         banned.append(current)
118
119     next_node = get_next_node(current, P_matrix, banned_nodes=banned)
120
121     if next_node == current:
122         loiter_count += 1
123     else:
124         loiter_count = 0
125
126     route.append(next_node)
127     current = next_node
128
129 print(f"Example 60h Route: {' -> '.join([r.rsplit('_', 1)[-1] for r in route[:10]])} ...
130 ↪ (truncated)")
131
132 # Plotting the route
133 fig, ax = plt.subplots(figsize=(10, 10))
134 plt.title("60-Hour Simulated Route Map (USV-1)", fontsize=16)
135
136 # Convert route to wgs84 for Contextily
137 route_points = []
138 for r in route:
139     pt_metric = centroids[r]
140     # Temp GeoSeries to cleanly project the point
141     temp_gdf = gpd.GeoDataFrame(geometry=[pt_metric], crs="EPSG:32633").to_crs(epsg=3857)
142     route_points.append(temp_gdf.geometry.iloc[0])
143
144 # Plot bounding box
145 bbox_wgs84.to_crs(epsg=3857).plot(ax=ax, facecolor="none", edgecolor="red", linewidth=2,
146 ↪ linestyle="--", zorder=10)
147
148 # Create line out of route
149 if len(route_points) > 1:

```

```

146     line = LineString(route_points)
147     gpd.GeoSeries([line], crs="EPSG:3857").plot(ax=ax, color='blue', linewidth=1.5,
    ↪ alpha=0.7, zorder=5)
148
149 # Plot nodes hit
150 gpd.GeoSeries(route_points, crs="EPSG:3857").plot(ax=ax, color='red', markersize=20,
    ↪ zorder=6)
151
152 ax.set_xticks([])
153 ax.set_yticks([])
154 ctx.add_basemap(ax, source=ctx.providers.OpenTopoMap, zorder=0)
155
156 fig.savefig("dobss_60h_simulated_route.pdf", bbox_inches='tight')
157 plt.close(fig)
158 print("Simulation complete. Route map saved to 'dobss_60h_simulated_route.pdf'")

```

Table F.1: Hour-by-hour patrol schedule (representative segment).

Hour	Node
00	t_{156}
01	t_{156}
02	t_{271}
03	t_{271}
04	t_{157}
05	t_{157}
06	t_{158}
07	t_{158}
08	t_{270}
09	t_{270}
10	t_{156}
11	t_{156}
12	t_{158}
13	t_{158}
14	t_{157}
15	t_{157}
16	t_{251}
17	t_{251}
18	t_{158}
19	t_{158}
20	t_{156}

Continued on next page

Hour	Node
21	t_{156}
22	t_{176}
23	t_{176}
24	t_{156}
25	t_{156}
26	t_{158}
27	t_{158}
28	t_{271}
29	t_{271}
30	t_{158}
31	t_{158}
32	t_{271}
33	t_{271}
34	t_{156}
35	t_{156}
36	t_{271}
37	t_{271}
38	t_{270}
39	t_{270}
40	t_{251}
41	t_{251}
42	t_{157}
43	t_{157}
44	t_{158}
45	t_{158}
46	t_{270}
47	t_{270}
48	t_{156}
49	t_{156}
50	t_{176}
51	t_{176}
52	t_{158}
53	t_{158}
54	t_{271}
55	t_{271}
56	t_{158}
57	t_{158}
58	t_{251}

Continued on next page

Hour	Node
59	t_{251}
60	t_{271}
