UNIVERSITÄT PASSAU
Fakultät für Informatik und Mathematik

RISK-BASED SECURITY MANAGEMENT IN CRITICAL INFRASTRUCTURE ORGANIZATIONS

ALI ALSHAWISH

A Thesis submitted for

Doctoral Degree

Chair of Computer Networks and Computer Communications
Faculty of Computer Science and Mathematics
University of Passau

October 20, 2021

REVIEWERS:
Prof. Dr.-Ing. Hermann de Meer
Chair of Computer Networks and Computer Communications
Faculty of Computer Science and Mathematics
University of Passau
Innstraße 43
D-94032 Passau, Germany
hermann.demeer@uni-passau.de
http://www.net.fim.uni-passau.de/

Prof. Dr.-Ing. Stefan Rass
Secure and Correct Systems Lab
Linz Institute of Technology
Johannes Kepler University Linz
and
System Security Group
Alpen-Adria-University of Klagenfurt
Altenberger Straße 69
A-4040 Linz, Austria
stefan.rass@aau.at
stefan.rass@jku.at
https://www.syssec.at/en/team/rass

Prof. Dr.-Ing. Jana Dittmann
Research Group for Advanced Multimedia and Security
Department of Computer Science
Otto-von-Guericke-University of Magdeburg
Universitätsplatz 2
D-39106 Magdeburg, Germany
jana.dittmann@iti.cs.uni-magdeburg.de
http://wwwiti.cs.uni-magdeburg.de/%7Ejdittman/

ABSTRACT

Critical infrastructure and contemporary business organizations are experiencing an ongoing paradigm shift of business towards more collaboration and agility. On the one hand, this shift seeks to enhance business efficiency, coordinate large-scale distribution operations, and manage complex supply chains. But, on the other hand, it makes traditional security practices such as firewalls and other perimeter defenses insufficient. Therefore, concerns over risks like terrorism, crime, and business revenue loss increasingly impose the need for enhancing and managing security within the boundaries of these systems so that unwanted incidents (e. g., potential intrusions) can still be detected with higher probabilities. To this end, critical infrastructure organizations step up their efforts to investigate new possibilities for actively engaging in situational awareness practices to ensure a high level of persistent monitoring as well as on-site observation.

Compliance with security standards is necessary to ensure that organizations meet regulatory requirements mostly shaped by a set of best practices. Nevertheless, it does not necessarily result in a coherent security strategy that considers the different aims and practical constraints of each organization. In this regard, there is an increasingly growing demand for risk-based security management approaches that enable critical infrastructures to focus their efforts on mitigating the risks to which they are exposed. Broadly speaking, security management involves the identification, assessment, and evaluation of long-term (or overall) objectives and interests as well as the means of achieving them.

Due to the critical role of such systems, their decision-makers tend to enhance the system resilience against very unpleasant outcomes and severe consequences. That is, they seek to avoid decision options associated with likely extreme risks in the first place. Practically speaking, this risk attitude can significantly influence the decision-making process in such critical organizations. Towards incorporating the aversion to extreme risks into security management decisions, this thesis investigates thoroughly the capabilities of a recently emerged theory of games with payoffs that are probability distributions. Unlike traditional optimization techniques, this theory provides an alternative decision technique that is more robust to extreme risks and uncertainty. Furthermore, this thesis proposes a new method that gives a decision maker more control over the decision-making process through defining loss regions with different importance levels according to people's risk attitudes. In this way, the static decision analysis used in the distribution-valued games is transformed into a dynamic process to adapt to different subjective risk attitudes or account for future changes in the decision caused by a learning process or other changes in the context.

Throughout their different parts, this thesis shows how theoretical models, simulation, and risk assessment models can be combined into practical solutions. In this context, it deals with three facets of security management: allocating limited security resources, prioritizing security actions, and tweaking decision making. Finally, the author discusses experiences and limitations distilled from this research and from investigating the new theory of games, which can be taken into account in future approaches.

ACKNOWLEDGMENTS

I am deeply indebted to many persons who have, in various ways, contributed to me being able to complete this thesis.

First and foremost, I would like to express my sincere gratitude to my primary supervisor Prof. Hermann de Meer for providing me the opportunity to complete my PhD study at the University of Passau, as well as for his patience and continuous support. His guidance helped me in all the time of research and writing of this thesis.

My sincere thanks also go to my second supervisor Prof. Stefan Rass for the in-depth perspectives and discussions he shared with me during my research. Without his invaluable support and feedback, it would not have been possible to conduct this research.

I would like to thank Prof. Jana Dittmann for kindly agreeing to be the external reviewer of this thesis.

I would like to thank my colleagues at the chair of Computer Networks and Computer Communications (CNACC) for their support. Special thanks go to Silvia Lehmbeck for her encouragement, which pushed me through the difficult time in my PhD journey, Michael Niedermeier for taking the time to review my thesis, and Mohamed Amine Abid for his help in achieving my goal and for his constructive and valuable feedback.

I gratefully acknowledge Project HyRiM (Hybrid Risk Management for Utility Providers) for providing the context to discuss some parts of the work presented in this thesis.

My sincere gratitude and appreciation go to my best friends, Ammar Alyousef and Waseem Mandarawi, for being always by my side through successes and failures.

I would like to thank my parents and my sisters for supporting me spiritually throughout writing this thesis and in all aspects of my life.

Last but not least, I would like to extend my deepest gratitude to my wife, Raneem, and my children, Hasan, Shahd and Mohamad, who always give me the strength and confidence in myself that I need. I am fortunate to have you in my life; thank you for all the support and constant patience over the past years!

CONTENTS

LIST OF FIGURES

LIST OF TABLES

AC	Attack Complexity
APT	Advanced Persistent Threat
AV	Attack Vector
AWS	Amazon Web Services
B2B	Business to Business
B2C	Business to Customer
BSI	Bundesamt für Sicherheit in der Informationstechnik
CCTV	Closed Circuit Television
CI	Critical Infrastructure
CIA	Confidentiality, Integrity, and Availability
CTMC	Continuous Time Markov Chain
CVSS	Common Vulnerability Scoring System
DER	Distributed Energy Resources
DHS	Department of Homeland Security
DSL	Digital Subscriber Line
HMI	Human Machine Interface
IAEA	International Atomic Energy Agency
ICT	Information and Communications Technology
IED	Intelligent Electronic Device
IoT	Internet of Things
IPRMA	Iterated Prioritization of Risk Mitigation Actions
ISMS	Information Security Management Systems
ISO	International Organization for Standardization
IT	Information Technology
KPI	Key Performance Indicator
M2M	Machine to Machine
MTTC	Mean Time To Compromise
MTU	Master Terminal Unit
NAC	Network Access Control
NIST	National Institute of Standards and Technology
NSA	National Security Agency
NVD	National Vulnerability Database
OT	Operational Technology
PPD	Presidential Policy Directive
PTTC	Path Time To Compromise
RTU	Remote Terminal Unit
SCADA	Supervisory Control And Data Acquisition
TTC	Time To Compromise
TTTC	Transition Time To Compromise
UAV	Unmanned Aerial Vehicle
VAR	Volt-Amp Reactive
VPN	Virtual Private Network

Part I

SECURITY MANAGEMENT

This part provides a general introduction to the whole thesis. It shows that complete security, known as the absence of threats, is not an attainable objective, and there is a need for risk-based security management to maximize the benefits of security efforts. Part I provides a list of the overall contributions of the thesis, followed by background information and a discussion of related research areas. Besides, a methodological approach for security management in critical infrastructures is presented in this part and instantiated in Part II and Part III.

Publication references*:

- Ali **Alshawish**, Mohamed Amine Abid, and Hermann de Meer. "Quasi-purification of mixed game strategies: Sub-optimality of equilibria in security games." *Computers & Security* 87:101575, 2019.
- Ali **Alshawish**, Mohamed Amine Abid, Hermann de Meer, Stefan Schauer, Sandra König, Antonios Gouglidis, and David Hutchison. "G-DPS: A game-theoretical decision-making framework for physical surveillance games." In *Game Theory for Security and Risk Management*, pp. 129-156. Birkhäuser, Cham, 2018.
- Ali **Alshawish**, Mohamed Amine Abid, Zhiyuan Sui, Hermann de Meer, Antonios Gouglidis, and Stefan Rass. "HyRiM Deliverable 4.3: How to Enhance Perimeter Security Using New Surveillance Technologies." 2017.

* The research work (including some ideas and figures) from these papers, which is included in Part I, was carried out and documented by the author of this thesis.

1 INTRODUCTION

The vast majority of Critical Infrastructures (CIs), especially utility networks such as power, transportation, water, and gas networks, have been working (operational) for several decades. They represent the main pillars for national economy and prosperity, as they provide essential services and fundamental networks upon which we all are extensively dependent and tightly linked. Advances of some infrastructures, in particular Information and Communications Technology (ICT), have a significant impact on other sectors through providing new opportunities and mechanisms for improving business efficiency and managing complex operations [8]♮[1].

Given the high reliance upon such infrastructures, their outages, inadequate service supply, or temporal disturbances can adversely impact the overall quality of life, public safety and security, or even nations' progressiveness and competitiveness. Besides potential delays in delivering vital services, disruptions of CIs are most likely associated with high costs of damage recovering. Moreover, the damage is not always limited to the affected area but can occasionally propagate into other sectors and areas through cascading and escalating failure mechanisms [8]♮. Therefore, the protection of CIs has been addressed as a national priority in many countries. Nowadays, the resilience of CIs attracts global attention and a great deal of interest in the industrial and scientific world.

Intuitively, protection implies the state of keeping the valuable assets, which an organization owns, manages, or controls, from being damaged, stolen, or lost [8]♮. To this end, most organizations surround their valuable assets with a secure perimeter, where everything in it is allegedly protected. However, the development and deployment of security controls are becoming more complex and expensive than ever before, while penetration means are getting cheaper and publicly available [116]. Therefore, risk-based security management becomes an integral part of CIs' protection programs since it aims at coordinating and balancing various security efforts towards reducing the risk exposure of such systems.

1.1 OVERVIEW OF SECURITY MANAGEMENT

The increasing connectivity and complexity of CIs make traditional solutions, such as antivirus software or firewalls insufficient to ensure security. Modern attackers are more organized and sophisticated than ever before. They develop adaptive strategies to attack weak points in an organization's defense. Mobile and outsourcing technologies, for example, have expanded CI boundaries. Thus, cybercriminals seek to compromise mobile devices of employees and contractors before infiltrating more sensitive (trusted) areas housing more valuable assets. Therefore, security has to be managed and enforced not only at the CI's perimeter but also within the system's boundaries.

1 Throughout this document, "♮" symbol is used to mark references authored or co-authored by the author of this thesis.

Traditional security can be characterized as a matter of best-practice measures, which are designed and configured towards protecting well-defined goals such as Confidentiality, Integrity, and Availability (CIA) [97, 236]. In practice, the question of which security controls should be in place is usually solved by means of security standards and guidelines to meet some compliance obligations [25, 96, 100, 184]. However, compliance with standards and regulations, which do not provide tailored implementation procedures and details and are designed to ensure specific objectives (usually determined by regulatory bodies), can fail to produce a coherent and long-term protection plan for CIs without assessment of their risk exposure in a comprehensive manner. Therefore, a "one-size-fits-all" solution to choose, prioritize, and deploy security controls leads to insufficient protection [76]. Recently, organizations increasingly adopt risk-based management approaches, which enable conducting a comprehensive assessment of risks in their current operational environments. Risk-based approaches allow organizations to focus their efforts on the risks that are more significant to their operations, thereby maximizing security resource efficiency. In this thesis, security management pays particular attention to how to configure and optimize security operations through changing the focus to (or aligning best-practice approaches with) security risk.

The core of a risk-based security management approach involves a decision-making process. It uses risk as the basis for performance assessments and captures offensive-defensive interactions between involved agents (e.g., a defender and potential attacker) towards making risk-informed strategic decisions. BSI[2] standard 100–1 on Information Security Management Systems (ISMS) states that

> *"practical experience has shown that optimizing information security management frequently improves information security more effectively and lastingly than investing in security technology."* [22].

This thesis is aligned with this principle since its primary purpose is an endeavor

- to provide the means to bridge the gap that standards and guidelines can have regarding the implementation of security controls and
- to support and guide a system's defender in managing security investments and priorities.

In brief, security management problems are concerned with how to create a coherent security strategy that leaves organizations well-positioned in the continual race against potential adversaries. It defines a process of controlling and coordinating security practices such as conducting random patrols or regular spot-checks to prevent or deter potential intrusions as well as prioritization of vulnerability remediation actions. Moreover, security management as a broader concept is not equal but closely related to several concepts and methodologies such as risk management, resource allocation, mechanism design, and decision theory (see Section 4.3 for further details).

1.2 CHARACTERISTICS OF SECURITY MANAGEMENT PROBLEMS

While the goal of complete security is still – and will most probably continue to be – unattainable, security management approaches seek to maximize the benefits of

2 The German Federal Office for Information Security (Bundesamt für Sicherheit in der Informationstechnik (BSI))

available security resources through rendering potential attacks non-economic and thus not meaningful to be mounted [13]♮. In this regard, a system defender has to configure available security controls in a way that minimizes the benefits received by a potential attacker upon a successful attack. Thus, decisions related to security management have to be made in a *competitive environment*, taking into account interactions between *involved agents* (or players including defenders and attackers) as well as the potential presence of *multiple decision objectives*. To guarantee an adequate level of security, an organization must consider its resources and processes in a comprehensive approach that can balance the security risks to which an organization is exposed with other aims and security goals. There can be different objectives and sub-objectives corresponding to the different levels and units of an organization. In practice, objectives can significantly differ from one organization to another.

When thinking of possible players' actions, numerous security standards and guidelines are available and constitute the basis on which organizations identify their own catalog of potential attack plans and proper measures to mitigate the impact thereof. In addition to standards, systematic system analysis and involvement of experts with different domains of expertise can provide valuable information to identify both *offensive and defensive actions* [8]♮.

Besides the characteristics mentioned above, there are several sources of *uncertainty* affecting security decisions. They can include, just to name a few, dynamic nature of the system at hand, practicalities and imperfections of the applied security plans, unforeseen external events, lack of information about attackers' types and incentives, among others [8]♮. Throughout this thesis, two types of uncertainties are accounted for:

- Type I: This type of uncertainty refers to *randomness of consequences* that an action has. That is, the impact of an action fits along a spectrum ranging from consequences that are low or associated with low damage or loss to those that are deemed severe and they must be avoided even if they are rare. This type addresses mainly the integration of variability and uncertainty into risk assessments[3]. It is, therefore, impossible to precisely determine the future (state) outcome of an action or situation, which typically gives sufficient grounds for making decisions.

- Type II: This type of uncertainty refers to a complete *absence of information*. In the realm of security, this type emerges when reliable information about preferences and revenues of potential adversaries is involved in decision-making processes. Commonly, potential attack strategies can be identified based on analyzing the system of interest and available domain-knowledge. Any assumptions, however, on the different attackers' behaviors and intentions (i. e., an attacker's preferences on which action is more likely to happen) may be wrong and can significantly affect the final results [8]♮. Due to the existence of different kinds of potential attackers, there is usually no reasonable information on what could be their expected payoffs from attacking CIs. Therefore, defending agents seek to optimize their behaviors under *uncertainty of the attacker type*. This optimized behavior or configuration is called a *security strategy* [171, 218].

When making decisions, it is crucial to understand that this process does not depend only on what the pure (objective) outcomes of actions might be, but also on how the

3 Being measurable is a fundemental assumption to manage risk. Otherwise, management processes are not able to properly improve (here, minimize) risk; "If you cannot measure it, you cannot improve it"[155].

involved decision-maker (subjectively) values these actions according to his interests and preferences. Due to the crucial role of CI organizations and potentially catastrophic effects of their failures, preferences of involved defenders and riskiness of actions are closely related. That is, decision-makers involved in the protection operations of critical systems are *sensitive to extreme risks* and tend to enhance the system resilience against extreme events. For that purpose, they seek to avoid strategies and decision alternatives associated with likely severe consequences. Practically, this risk attitude guides the decision-making process in such critical organizations, and hence the security management process as well [8]♮. Security management in CIs adopts *the principle of extreme risk minimization*, rather than the traditional principle of utility maximization, in which decision-makers seek to maximize their long-term (expected) payoffs [217]. Resilience engineering states that complex socio-technical systems, such as nuclear power plants and other critical infrastructure systems, must be designed and configured to cope with everyday situations, including accidents [90]. This thinking justifies the necessity to enhance the resilience and preparedness of critical systems to withstand (account for) extreme conditions and worst-case scenarios. Nonetheless, each organization or country might have different philosophies and regulations for dealing with severe and extreme risks [201]. For this specific reason, the very low probability of worst-case scenarios (extreme events) can cause a significant discrepancy between decisions based on pessimistic views and those based on positive ones. To further elaborate on the importance of extreme events in making decisions in CIs, an overview of the Fukushima nuclear disaster is included in Section 2.2.2.

In summary, security management problems can be modeled as decision-making problems that have the following characteristics or requirements to be accounted for:

Req1 Multiagent (more precisely, 2-agent) competitive environment

Req2 Interactions between offensive and defensive actions

Req3 Multi-objectiveness

Req4 Decision under uncertainty: random consequences

Req5 Decision under uncertainty: unknown type of attacker

Req6 Attitude toward extreme risks

1.3 CONTRIBUTIONS OF THIS THESIS

To address the challenges as mentioned earlier, this thesis investigates the limitations of existing decision-support approaches that are based on classical game-theoretical models. It turns out that classical game models alone would not be sufficient to address **Req4** and **Req6**. Therefore, this thesis relies on generalized game models with distribution-valued payoffs that allow integrating the identified challenges into the decision-making process. To this end, a methodological approach is presented that supports defenders of CIs to assess possible decision alternatives towards finding proper security strategies. In addition, a novel comparative analysis is provided to evaluate solutions of classical and distribution-valued games. As a further improvement to distribution-valued games, this thesis proposes a more dynamic decision analysis process that takes into account individual and subjective risk attitudes.

The presented methodological approach integrates concepts and principles from risk management, decision theory, and game theory. This thesis explains thoroughly the application of this approach on two security problems in physical and cybersecurity domains. Considering these points, the exposition of this thesis falls into three main parts: Part I – Security Management, Part II – Physical Security Management, and Part III – Cybersecurity Management. Aligned with those three parts, the contributions of this thesis are listed in the following:

PART I: SECURITY MANAGEMENT

EXTENDED PERIMETER — The perimeter security model, which aims at building a hard shell around valuable assets, is still one of the most widely adopted security practices. However, with the ongoing paradigm shift of CIs to be more collaborative and dynamic, the boundaries of these systems that provide the distinction between the trusted inside and the untrusted outside tend to vanish or to turn into a fuzzy concept hard to define. Therefore, this thesis provides a comprehensive discussion on this phenomenon and shows how CIs can extend beyond their traditional boundaries, thereby paving the way for attackers to undetectably get into trusted inside environments. Besides the definition of the extended perimeter, it is pointed out in this thesis that some assumptions underlying perimeter security models are not necessarily valid in the context of CIs.

METHODOLOGICAL APPROACH — This thesis presents a methodological approach for risk-based security management in CIs and its closely related techniques and methods. The presented approach relies on a recently emerged theory of games with payoffs that are random variables described by entire probability distributions. This theory allows representing security problems as formal game models and integrating the identified challenges into the decision-making process. The approach breaks down into smaller and manageable steps to support defenders of CIs to assess possible security choices towards finding proper security strategies.

PART II: PHYSICAL SECURITY MANAGEMENT

PHYSICAL SURVEILLANCE GAMES — Having dynamic and mobile surveillance strategies is highly important to maintain situational awareness even within a CI system so that potential intruders can still be detected. In this thesis, the concept of physical surveillance games is introduced to address scenarios in which mobile agents perform repetitive spot-checks within CI boundaries to improve flexibility and intrusion detection probabilities. Being an instance of the generalized model of security management games, physical surveillance games integrate the uncertainty inherent to surveillance applications into the payoff structure. These games seek to bridge the gap between defining a sophisticated theoretical model and practically instantiating it. Physical surveillance games have several important real-life applications, such as physical border patrolling, public transit security, fare enforcement planning, among others.

ENTROPY-BASED PRIVACY MODEL — Surveillance practices such as random spot-checks could pose a serious threat to location privacy. Having access to

location traces, attackers can infer the identities of employees or plan further (targeted) attacks on their organizations. Therefore, this thesis presents an entropy-based model to assess the impact of different security inspection strategies on the preservation of employees' locations. The model employs the technique of Continuous Time Markov Chains (CTMCs) to quantify location privacy as a function of time. CTMCs enable the developed model to address the observation that the importance of location traces can fade over time.

SIMULATION MODEL FOR PHYSICAL SURVEILLANCE GAMES — The actual problem of optimal surveillance is diversified and involves multiple objectives to be satisfied. To achieve that, suitable and comprehensive assessments of the effectiveness of different surveillance configurations have to be performed. In this thesis, a simulation model for physical intrusion problems in CIs is developed. The model allows establishing a faithful image of the physical environment of the facility of interest, the deployed personnel and their behavior, as well as the potential attacks that may occur.

COMPARATIVE ANALYSIS — One of the main contributions of Part II is the demonstration of how simulation, physical understanding of CIs, and theoretical models can be combined toward a practical solution. In this regard, a quasi-purification method is presented in this thesis to facilitate interpreting and hence implementing obtained game-theoretical decisions in practice. Furthermore, a novel comparative analysis is conducted to achieve a better understating of the differences between classical security games with scalar-valued payoffs and generalized games with distribution-valued payoffs. Given the analysis results, it turns out that optimality of a defense is not the same as optimizing a security score[4], since the means by which security is quantified and optimized plays a much deeper role than intuitively expected.

PART III: CYBERSECURITY MANAGEMENT

PRIORITIZATION FRAMEWORK — An ever-increasing reliance on ICT in CIs, makes them vulnerable to cyber threats and risks. In CIs, however, resolving all vulnerabilities at once could seem like an insuperable hurdle for patch management teams due to several technical and economic constraints. Therefore, they need to prudently assess priorities and make a decision on the importance of possible remediation activities in order to implement them more effectively. This thesis introduces a prioritization framework based on the introduced methodological approach for security management to assist the defenders of CIs in making risk-informed decisions on the action priorities. Technically, remediation actions are prioritized through successively playing a chain of cybersecurity games towards minimizing the risk of compromise.

TIME-TO-COMPROMISE MODEL — In this thesis, cybersecurity risk is assessed based on a Time To Compromise (TTC) metric. TTC-based risk estimates deliver insights into a system's robustness against technical vulnerabilities given different remediation actions. The developed risk estimator integrates a generalized

4 A security score is a static number used to describe the security posture of a system or the performance of a security strategy.

stochastic TTC model with Monte Carlo simulation techniques to consider several challenges, including inherent prediction uncertainty, interdependencies among network components, different attackers' skill levels, and public vulnerability and exploit information.

TWEAKABLE STOCHASTIC ORDER — Unlike traditional game-theoretical models that set utility values so as to reflect a person's choices as accurately as possible, this thesis pays more attention to the ordering relation itself upon which rational behavior is defined. This leads to the introduction of a tweakable stochastic order, which is defined on random variables such as loss distributions known in actuarial science. This ordering relation can be adapted to individual risk attitudes of game players. In this respect, a dynamic decision analysis is defined by using the shape (i. e., convexity or concavity) of subjective utility weighting functions to partition loss distributions according to the segments that influence the players' decision-making under risk.

1.4 THESIS STRUCTURE

The thesis structure is depicted in Figure 1.1, where the interplay of the different parts and chapters is visualized. As mentioned above, this thesis consists of three fundamental parts.

Part I – Security Management – encompasses four chapters: Chapter 1 is an introductory chapter and gives an overview of the whole thesis. In Chapter 2, background information and a review of related work in the respective areas are given. Chapter 3 investigates the extended perimeter phenomenon in CIs more deeply. Several parts of Chapter 3 are based on the research work appeared in [8]♮. Then, the methodological approach for risk-based security management in CIs is presented in Chapter 4. The content of Chapter 4 is based on the research work published in [2, 8, 10]♮.

Part II – Physical Security Management – encompasses three chapters: Chapter 5 explains the notion of surveillance as well as discusses potential limitations and challenges. Then, it introduces models of physical surveillance games and proposes an entropy-based model for quantifying location privacy. The content of Chapter 5 is based on the research work published in [5, 8, 10, 11]♮. In Chapter 6, a use case of a nuclear power plant is adopted to demonstrate the application of physical surveillance games. It explains how a defender can choose among several possible resource allocations, and relies on game theory for an optimal choice. The developed simulation model for physical intrusion problems is described in Chapter 6, as well. The content of Chapter 6 is based on the research work published in [1, 2]♮. The obtained decisions are then evaluated in Chapter 7 using a four-dimensional comparative analysis. The content of Chapter 7 is based on the research work published in [2]♮.

Part III – Cybersecurity Management – encompasses three chapters: In Chapter 8, a generalized stochastic TTC model is proposed, and a prioritization framework is described in detail. The content of Chapter 8 is based on the research work published in [3, 4, 6, 7]♮. Chapter 9 demonstrates and evaluates the application of the proposed framework in prioritizing remediation actions available to upgrade an electricity provider network according to their risk mitigation effects. The content of Chapter 9 is based on the research work published in [7]♮. In Chapter 10, the idea and method of tweakable

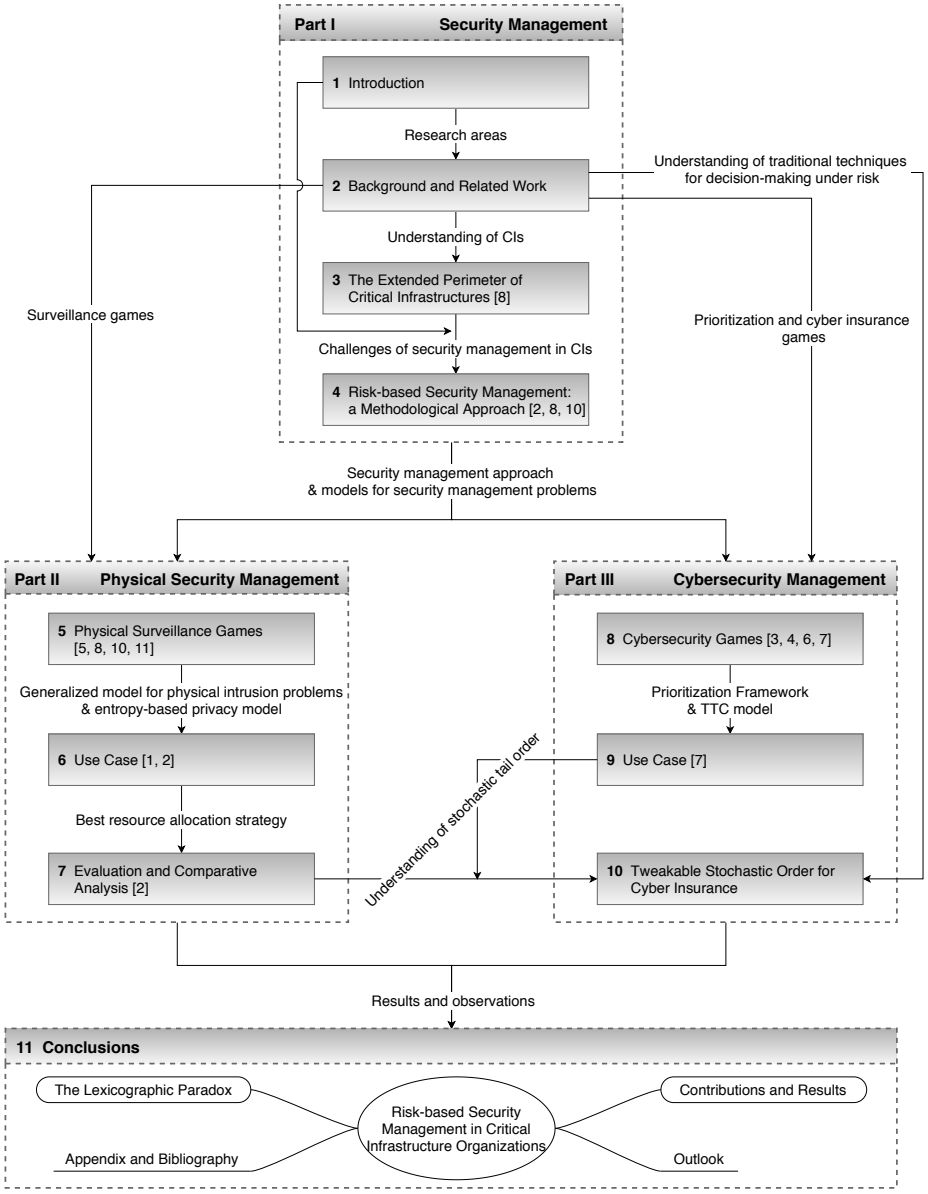

Figure 1.1: Thesis structure and chapter dependencies

stochastic order are presented and illustrated using a bimatrix cyber insurance game model. The proposed tweakable order enables having an auditing policy tailored to the customer's risk attitude so that the insurance can act more informed and accurate on the detection of fraud.

Finally, Chapter 11 concludes this thesis and recaps the main results thereof. Additionally, an outlook is provided to indicate potential future research directions based on the outcomes of this thesis.

2 BACKGROUND AND RELATED WORK

This thesis aims to combine several perspectives and techniques to approach security management problems in CIs. These include game-theoretical models, risk management perspectives, and valuation techniques of choice options. Therefore, there are several research areas connected to this thesis. This chapter gives an introduction and discussion of related work in the following research topics: (i) critical infrastructure systems, (ii) risk and decision making, as well as (iii) game theory for security.

2.1 CRITICAL INFRASTRUCTURE SYSTEMS

In our modern society, daily life is more and more dependent on a set of key resources and services, including electricity, water supply, transportation, communication, and many others. Due to the enormous importance of such services and their immediate influence on the economic prosperity and well-being of nations, the complex socio-technical systems that deliver those services have been recognized as *critical infrastructures* in their own right [118].

Generally speaking, there is a lack of one commonly accepted definition of CIs. Different countries adopt, therefore, different definitions and interpretations, which ultimately result in dissimilarities in the identification of their individual CI sectors. In Germany, for instance, the National Strategy for Critical Infrastructure Protection defines CIs as "*organizations or institutions with major importance for the public good, whose failure or damage would lead to sustainable supply bottlenecks, considerable disturbance of public security or other dramatic consequences*" [61]. Building on this definition, the Federal Government categorizes 9 CI sectors: energy; transport and traffic; water; finance and insurance; food; information technology and telecommunications; media and culture; health; as well as government and public administration [211].

In the US, CIs are defined as "*systems and assets, whether physical or virtual, so vital to the United States that the incapacity or destruction of such systems and assets would have a debilitating impact on security, national economic security, national public health or safety, or any combination of those matters*" [212]. Thus, the Presidential Policy Directive (PPD)-21 identifies 16 CI sectors: chemical; commercial facilities; communications, critical manufacturing; dams; defense industrial base; emergency services, energy, financial services, food and agriculture; government facilities; healthcare and public health; information technology; nuclear reactors, materials, and waste; transportation systems; as well as water and wastewater systems [208].

It is worth noting that although CI definitions vary across countries, they agree and share the same focus on the vital role of CIs on the society and the devastating impact of their disruptions [190]. As a result, the protection of CIs has been considered in many countries as a national priority, which requires joint cooperation between public and private sectors [73]. Indeed, the interest in *planned* CI protection prominently featured two decades ago. Through those years, it turns out that 100% security is an impossible and impractical objective for CI protection due to sheer size, complexity,

and (inter)dependencies of CIs [118]. Therefore, there is recently a shift in perspective towards strengthening resilience[1] and security through risk-based decision making [53, 118, 190]. The Department of Homeland Security (DHS) defines risk-informed decision making as "*determination of a course of action predicated on the assessment of risk, the expected impact of that course of action on that risk, as well as other relevant factors*" [50]. That is, risk assessments are the primary driver for making security decisions in CIs. In close connection with this insight, security is viewed as "*reducing the risk to critical infrastructure by physical means or defens(ive) cyber measures to intrusions, attacks, or the effects of natural or man-made disasters.*" [208]. Thus, it is fundamental to explain the concept of risk and other key risk-related terms, as shown in Section 2.2.

2.2 RISK AND DECISION MAKING

This section starts with the definitions of "risk" and "risk attitude", and then discusses the common methods for decision-making under risk.

2.2.1 *What is risk?*

Although the word *risk* is prevalent and widely used, it represents an ambiguous term that still lacks a precise definition and a broadly accepted meaning. Informally and in daily life activities, the notion of risk is used by different people to indicate different things and meanings [198]. The following three questions, for example, use the same word risk but to clearly express different meanings:

1. Which *risk* should be top-ranked or remediated first? Here, risk refers to an **adverse event or threat** such as cyber intrusions, theft and burglary, as well as phishing and social engineering attacks. More specifically, risk is used to describe things (states, events, or persons) regarded as potential sources of (or involving exposure to) danger and loss.

2. What is the *risk* of getting compromised via Internet access? Here, the word risk is used to refer to **the probability or chances** that an adverse event or unpleasant thing (i. e., getting compromised) will occur.

3. What is the *risk* of using unpatched vulnerable components or unsupported operating systems? Here, risk refers to **the possible impact or consequences** of an adverse event (i. e., having an unsupported operating system in the network) such as loss and damages. It is worth noting that the event is defined, but its impact is uncertain (here, in the sense of variability) and can accept any value in a range of possible values.

Consulting several dictionaries makes it more evident that there is a lack of a formally unified and overarching definition of risk. Merriam-Webster Dictionary defines risk as "1) a possibility of loss or injury, and 2) someone or something that creates or suggests a hazard" [179]. Cambridge dictionary defines risk as "1) the possibility of something bad

1 The directive PPD-12 defines resilience as "*the ability to prepare for and adapt to changing conditions and withstand and recover rapidly from disruptions...it includes the ability to withstand and recover from deliberate attacks, accidents, or naturally occurring threats or incidents.*" [208].

happening; 2) something bad that might happen (e. g., former employees are a security risk); 3) danger; 4) possiblity of danger, defeat, or loss; 5) someone or something that could cause a problem or loss" [181].

Practically, equating risk with adverse events is not appropriate because events are very abstract and cannot be quantified or assessed. Moreover, relying merely on the likelihood or potential outcome of a hazard is not enough as well because some rare events can be associated with possibly catastrophic outcomes, while some high-frequency hazards might lead to easily bearable consequences.

When thinking of quantification, it is essential to mention that risk is closely related to the notion of uncertainty, yet they are not interchangeable. One of the early definitions that stresses the link between risk and uncertainty is provided by Knight (1921), who introduced risk as a quantifiable uncertainty [110]. Furthermore, Hillson and Murray-Webster define risk as "the uncertainty that matters" [88]. Hence, uncertainty is essential to make a situation or decision appears risky [228]. Nevertheless, "A risk is not an uncertainty where neither the probability nor the mode of occurrence is known" [180]. On the contrary, risk involves that the probability of a variable (such as compromising a computer network) is known but when a mode of occurrence or the actual value of the occurrence (whether the compromise will occur in a particular time period) is not [180].

Aligned with the latter perspective, the National Institute of Standards and Technology (NIST) introduces risk as a compound notion of those three factors, namely event, outcome, and occurrence probability. NIST defines risk as "*the potential for an unwanted outcome resulting from an incident, event, or occurrence, as determined by its likelihood and the associated impacts*" [204]. Therefore, risk is typically characterized using two quantities [203]: (i) the magnitude (or severity/impact) of the various adverse outcomes that can potentially result from an event, and (ii) the occurrence probabilities of these respective outcomes. Throughout this work, probability distributions are used as integrative riskiness models in which the probabilities and adverse consequences of events are combined.

It is worth mentioning that risk cannot be identified or measured without being linked to some kind of objective. This can be easily derived from the definition of risk as "the uncertainty that matters or affects one or more objectives" [88]. Similarly, International Organization for Standardization (ISO) standards (e. g., 31000 [99], 27000 [97], and ISO Guide 73 : 2009 [101]) define risk as "the effect of uncertainty on objectives"[2]. Linking risk to an objective plays a crucial role in defining the type of risk, such as financial risk, reputation risk, security risk, among others.

2.2.2 *What is risk attitude?*

Risk attitude is defined as the "chosen response to the perception of uncertainty that matters" [87]. Hence, it derives the reaction and behavior of each person based on how he perceives the environment or situation. In this regard, one needs to differentiate

2 This definition bears two sides of risks, which are threats and opportunities. In this work, the focus is laid on minimizing threats. Nevertheless, if opportunities are considered, one can simply use the same techniques but to maximize them. The difference between opportunities and threats is the sign of the impacts. In a competitive two-player game model, for example, if one player seeks to minimize threats, then the opponent player seeks to maximize his opportunities. Basically, there are four responses to negative risks which are avoid, transfer, reduce, and accept. Their counterparts of positive risks are exploit, share, enhance, and accept.

between risk attitude and risk perception. The latter refers to "people's judgments and evaluations of hazards they (or their facilities, or environments) are (or might be) exposed to". The former, however, describes "people's intentions to evaluate a risk situation in a favorable or unfavorable way and to act accordingly" [183].

Risk attitudes exist at individual, group, corporate, and national levels [87]. That is, different individuals or organizations may adopt different actions for the same risky situation. Hence, risk attitudes add another layer of complexity into decision-making processes in CIs.

While risk perception is the interpretation of a world/situation/option based on knowledge and believes, risk attitude steers the evaluation process and determines the preferences of individuals or groups. Differences in attitudes towards risk can be caused by several factors such as different levels of knowledge and experiences held by involved participants, different organizational risk-appetite and risk-tolerance statements, or different adopted assumptions [63, 64, 224–227]. In general, humans tend to hold *domain-specific attitudes towards risks* [70, 224]. That is, properties of the investigated systems or environments play a crucial role in the process of decision-making and formulating the preference relations. As a result, decision-making under risk is a context-dependent process which requires careful alignment of the choice behavior with context-related preferences. Critical infrastructures, for example, are very vital systems that need a higher reduction in risk, in particular, high-impact (or extreme) risk before they become safe enough [31, 90]. Therefore, this thesis discusses integrating risk attitudes into the decision process itself, thereby seeking to increase the satisfaction and quality of delivered decisions. For the sake of explaining the importance of extreme risks in making decisions in CIs, the following section gives an overview of the Fukushima nuclear disaster.

Fukushima-Daiichi Nuclear Power Plant Disaster

In 2011, a massive earthquake with an intensity of 9.0 on the Richter scale (the fourth-largest earthquake recorded since 1900) struck off the North Pacific coastal areas of Japan. The earthquake disrupted several CIs, including electricity, gas, water, and transportation. Due to a massive reduction (or loss) of electricity supply, the emergency backup generators at the nearby Fukushima nuclear power plant kicked in, in an attempt to keep nuclear reactors under control and to shut them down properly. However, shortly after the earthquake, gigantic tsunami waves hit the plant, and the backup generators located under the ground were flooded, leading to a station blackout. This occurrence disabled the core cooling system resulting in a partial meltdown at the reactors, which in turn caused the release of an enormous quantity of radioactive materials into the environment [90, 112].

Before delving deeper into possible causes, it is worth mentioning that nuclear power plants typically undergo strict review and inspection practices to verify their compliance with safety regulations, standards, and guidelines. This fact raises the question of whether the standards were faulty or the designer and experts of the plant made improper design decisions. As discussed in Section 1.1, standards and guidelines help organizations to identify their own catalogs of measures to ensure security and safety. Nevertheless, they may leave detailed implementation procedures and operation policies to compliance and design teams, who have the necessary skills and knowledge about the nature of the environment surrounding their organizations.

One reason for Fukushima disaster can be the reliance upon improper techniques. They include approaches that use single values of goal evaluations or risk assessments (e.g., mean value) as design criteria. Expected utility theory or cost-effective approaches can underestimate or ignore extreme events, which might have the lowest probability and most significant magnitude of the impact. This can result in a misconfigured or even under-dimensioned (critical) system. For example, the perimeter wall at Fukushima nuclear plant is designed to withstand the average impact (probably with some margin) of tsunami waves with 5.7 meters in height. However, the height of the waves struck the plant in 2011 exceeded 14 meters [201]. Since the effects of rare events are usually not well-understood, their assessment processes are definitively associated with significant uncertainties. Nevertheless, the importance of uncertainty (i. e., as an indicator of possible catastrophic conditions or consequences) is neglected by taking average numbers as a main design criterion. Hence, events with very low probability were not considered in the design basis of the nuclear power plant. It is definitively intolerable to ignore severe conditions while designing or operating critical systems just because they are associated with (very) low probabilities [90], as evidenced by Fukushima disaster. Safety guidelines stress the importance of probabilistic models for risk assessments to establish the basis for suitable design (see [57] and the references cited therein).

2.2.3 *Decision-making under risk*

Decision theory offers several techniques to decide on the preference among a set of risky choices. Such techniques involve the use of a proper evaluation tool that allows comparison among the choices. The vast majority of risk-based decision-making problems are solved using scalar (expected) risk values. On the one hand, the single-value representation has a positive smoothing effect by removing unimportant perturbations from the decision analysis process, thereby significantly simplifying the comparison process [182]. But, on the other hand, such representation is associated with a loss of information that can lead to improper conclusions and decisions. Moreover, scalar-valued evaluation techniques strongly imply that all decision-makers share the same perception and attitude towards risks. However, this assumption is not always realistic and has been violated in practice, as shown in [46, 122, 158, 210]. As discussed earlier in this thesis, decision-makers involved in critical infrastructure protection are more concerned with avoiding extreme risk situations in order to improve the resilience and preparedness of their systems. In other (less critical) contexts, however, decision-makers might be indifferent between choices with equal expected values regardless of their riskiness.

A key challenge is, therefore, how to model decision-makers' preferences over uncertain outcomes. Let us approach the problem abstractly from a general perspective. Suppose a fixed set of possible outcomes (e. g., states or severity ratings) $\mathcal{C}$, an option X is a function assigning to each outcome $a \in \mathcal{C}$ a probability value $p(a)$ such that $\sum_{a\in\mathcal{C}} p(a) = 1$. For the sake of simplicity, $X = \{(x_i, p_i) | i = 0, \cdots, n-1; x_0 < x_1 < \cdots < x_{n-1}; p_i = p(x_i); \sum_{i=0}^{n-1} p_i = 1\}$ where only finitely many outcomes $x_i \in \mathcal{C}$ have non-zero probability p_i. Suppose a decision-maker is faced with two risky options described by random variables X and Y (defined likewise). The challenge is how to choose the better option (X or Y), given that the decision-maker does not know the

future realization of the actual random consequence of what action he takes. In order to model the preference relation between these options, several methods can be leveraged.

EXPECTED VALUE THEORY: This theory relies on evaluating and comparing the different options using their expected values. The expected value is the sum of the value of each possible outcome multiplied by its associated occurrence probability; i.e., it corresponds to the mathematical expectation. The roots of the notion of "expected value" date back to the 17^{th} century (invented by Blaise Pascal).

$$E(X) = \sum_{i=0}^{n-1} p_i x_i \tag{2.1}$$

Equation (2.1) aims to evaluate and represent uncertain options such as X and Y using single values $E(X)$ and $E(Y)$, respectively. This theory states that the option associated with the highest total expected value should be chosen. The idea of judging an option by combining its possible outcomes with the probability of each of these had become the essence of some highly valuable theories such as expected utility theory [217] and prospect theory [106, 210].

EXPECTED UTILITY THEORY: Being applied in the context of decision analysis, the aforementioned expected value theory had been criticized for not accounting for the characteristics and attitudes of decision-makers towards the real possible outcomes. Daniel Bernoulli showed in his solution approach of the St. Petersburg paradox[3] that the evaluation of an uncertain option should not be based on its real possible outcomes, but rather on the utilities they yield [28]. Therefore, judging an option using the expected utility theory involves the definition of a utility function, which assigns a real number (aka utility) to each of the possible option outcomes. The utility $u(x_i)$ of an outcome x_i is a subjective value expressing how valuable this specific outcome is to the respective decision-maker.

$$E(u(X)) = \sum_{i=0}^{n-1} p_i u(x_i) \tag{2.2}$$

Von Neumann and Morgenstern's theory of expected utility states that the option associated with the highest expected utility value (as defined in Equation (2.2)) is preferred over other options even if it does not provide the highest expected value [217]. Concerning the limitation of this theory, Kahneman and Tversky showed in their empirical study of economic decisions, how human behavior can deviate from the expected-utility maximization paradigm [106]. They argue that expected utility theory disregards several biases, including, but not limited to, (i) people tend to overweight small probabilities, but underweight large ones; and (ii) people perceive losses and gains differently and they assign more weight on the losses in case of equal gains. Therefore, they introduced prospect theory as an attempt to justify the observed deviations from the assumption of expected utility maximization.

3 Suppose a coin toss game, in which a player wins 2^i €, if heads appear on the first i tosses and tail on the $i+1$ toss. The expected (win) value of this game is unbounded. Players should not, therefore, be reluctant to pay any price to enter such a game, which is not the case in reality. St. Petersburg paradox addresses such a situation where using only expected values as a decision criterion fails to predict the actions of actual players [178].

PROSPECT THEORY: This theory provides a descriptive approach for human decisions under risk and uncertainty. It relies on (i) defining a (subjective) reference point to classify possible outcomes into (relative) losses and gains; (ii) defining two different utility functions u^+ and u^- for gains and losses; and (iii) defining a probability weighting function π, which transforms the objective (real) probability to support the observation that small probabilities are overrated, and large probabilities are underestimated. The formula that prospect theory assumes to evaluate decision options is given by Equation (2.3):

$$V(X) = V(X^+) + V(X^-) = \sum_{x_i \geqslant 0} \pi(p_i) u^+(x_i) + \sum_{x_i < 0} \pi(p_i) u^-(x_i) \tag{2.3}$$

In 1992, Tversky and Kahneman introduced an extension of prospect theory, called cumulative prospect theory [210]. This extension proposes to use a different probability weighting function to account for the empirical observation that people overrate low and underrate large probabilities, related to extreme but rare, and non-extreme but frequent events. Several probability weighting functions have been proposed to enhance the descriptive models of decision under risks. In this regard, the probability weighting function derived by Drazen Perlec has attracted widespread attention, since it is based on an axiomatic foundation and consistent with many existing empirical evidences [157].

In summary, all decision models, based on the theories mentioned above, seek to evaluate each uncertain option through combining its different possible outcomes and their respective occurrence probabilities into a single representative value. The evaluation process may involve some interesting transformation functions such as utility functions and probability weighting functions. These functions seek to model the non-linear behaviors and preferences of people with respect to the real outcomes and probabilities of decision options. They are usually derived based on a set of existing empirical observations. However, traditional evaluation techniques merge and combine different probabilities and severity levels into a single smoothed quantity concealing possible indications to extreme risk occurrences, which are significantly important for avoiding severe and perhaps irreversible consequences. Expected risk values would equate a low-probability high-impact option with a high-probability low-impact option, which are obviously not the same thing [107]. Thus, a single number is not enough to communicate risks properly [107]; risk is not the average value of a loss distribution but rather the whole distribution function.

2.3 GAME THEORY FOR SECURITY

Game theory is a field of study that was applied initially in economics to represent theories and to model conflict and cooperation between rational decision-makers [121]. It enables analyzing diverse decision problems using games, which are well-defined mathematical models with a set of essential elements, namely players, actions, and payoffs (cf. a detailed game-theoretical model in Section 4.1.1). Game solutions are basically predictions on strategies adopted by involved players to play the game, such as Nash equilibrium and its refinements [75, 141]. Hence, they play a vital role in the design of defense mechanisms and risk management under adversarial conditions. Recently, game theory has seen many diverse applications, including security and

resource allocation problems [20, 60, 104, 125, 161, 169, 173, 196]. For the connection between game theory and security, two books covering the topics are [172, 207]. To narrow down the focus of the analysis of related work, the discussion here is limited to the application of game theory to research topics connected to this thesis, which are physical surveillance, patch prioritization, and cyber insurance problems, as described in the remaining sections in this chapter.

2.3.1 *Surveillance games*

The application of game theory to optimize surveillance has been subject to a considerable amount of prior work. The general cops-and-robbers game has been studied in a variety of different forms. They include asking for the minimal number of cops to catch one (or more) robber(s) [80], relating structural properties of the graph to winning strategies for either party [17], or discussing the benefit of (in)visibility for either player [108]. Given the vast amount of available research, this thesis refers interested readers to surveys such as [19, 65] as well as the references in the cited literature.

Further interesting applications closely related to the physical surveillance games introduced in this thesis involve observing evading targets [30], optimal surveillance resource allocation under imperfect information for the attacker [21], sensor and mobile ad hoc network surveillance [82, 206], purely camera-based pursuit-evasion models [197], or the more general area of counter-terrorism [230], to mention only a few. Furthermore, several Stackelberg games have been employed to establish randomized patrol schedules towards ensuring security and fare enforcement in public transportation systems [39, 52, 105, 214].

Those approaches usually assumes perfectness of payoffs (even though not necessarily assuming perfect information) and their focus is purely laid on the game-theory side. As a result, they leave out the specifics and limitations of surveillance systems that can dramatically change the gameplay due to their imperfections. Different from most previous work, this thesis relies on generalized games over abstract spaces to deal with imperfect and uncertain payoffs (the entire theory is put forward in [163, 164, 167] and applied in [1, 2, 9][♮]). Those games work with whatever number of surveillance people are available and their computed security strategies are further used as resource sharing rules in practice.

Practicality of security strategies

The meaning of mixed strategies and the "practicality" of defining rationality as utility maximization or loss minimization has ever since its proposal been subject to controversial discussion. Assessments of equilibria against plausibility or observability in practice leads to refinements of the Nash equilibrium, inducing uncertainty in different forms. Among the most popular such notions are trembling hands equilibria [188] (where players deviate from their intended action at random), disturbed equilibria [83] (adding random noise to the received outcome) and the quantal response equilibrium [127] (where players choose their actions without assurance that their choice is optimal). This thesis is similar to these in considering the "quality" of a standard Nash equilibrium against alternative concepts of optimal behavior.

Given that games over abstract spaces are a relatively new concept, this thesis is the first of its kind, comparing two different ways of expressing the outcomes for the same game. As such, it is loosely related to work on *bounded rationality*. This term summarizes an entire area of research mainly devoted as to why people do not act utility maximizing in reality. The games considered in this work are no exception to this but bear a novel interpretation of mixed strategies as resource sharing rules. Unlike in a classical game, where the mixed strategy tells the (asymptotic) frequency of choosing an action, the defender can – in the games described – take the probability portions as fractions to allocate resources to play all actions *at the same time*.

The findings in this work relate to matters of bounded rationality in the sense of seeking an explanation as to why a classical Nash equilibrium may need refinement or is not always accurately observed in practice. The simplest explanation for such deviations from a utility-maximizing optimum may be having had "simply the wrong utility model", or more sophisticated in offering alternative such models. Among the candidates are level k thinking [140] or methods involving "hidden" incentives (like the cost to change a strategy [170]). The excellent essay of [205] discusses a huge body of literature on alternatives to the expected utility model, among them also stochastic dominance theories such as where games over stochastic orders would fall into. For example, [32] presents a model allowing for stochastic outcomes under a very specific error shape (Gaussian with zero mean). In contrast, the theory employed here [168] makes no such normality assumption (and is as such non-parametric).

2.3.2 *Prioritization games*

Among recent research activities on enhancing the cybersecurity of CIs, Shelar et al. propose a game-theoretical model to optimize the security strategy of electricity distribution networks [194]. They consider a specific adversary model, in which false data injection attacks are used to compromise vulnerable Distributed Energy Resources (DER) nodes. With regard to vulnerability patch management, the authors in [72, 124] combine game theory principles and vulnerability scoring techniques to prioritize vulnerabilities based on assessed severity indicators. Such approaches are vulnerability-centric; that is, their decisions always dictate that the vulnerability with the highest severity score should be resolved first. Such decisions are, however, not necessarily the best response in terms of minimizing risk. Suppose all devices in a network are affected by the same severe vulnerability like CVE-2017-0144 with the severity rating of 8.1 *HIGH* (CVSS v3.0) [40]. In this case, all devices – regardless of their characteristics or location on the network – are at high risk of being compromised and have the same priority to be patched first. This decision is not always actionable, thereby extremely confusing an involved security team. Thus, vulnerability prioritization that is naturally severity-based is not adequate for patch prioritization processes, which seeks to reduce the risk of compromise in the first place.

In [150], Panaousis et al. discuss applying game theory to advise security managers on how to invest in security controls optimally. Their game-theoretical model, however, assumes deterministic assessments (scalar-valued payoffs). That is, it does not account for inherent prediction uncertainties. Additionally, prioritization decisions made by existing frameworks do not consider the specified risk attitude of the decision-makers

involved in the protection of CIs. In classical game models, extreme risks may still be undesirably probable though the average risk has been optimized.

Another limitation of existing prioritization practices is that they depend heavily on qualitative judgments, which are typically highly subjective. Thus, they might lead to improper decisions significantly biased by individual perspectives. Such decisions could be influenced by an inaccurate interpretation of a system state caused by a forced consensus of the judgments as well as disregard of diversity. Towards mitigating this issue, this thesis employs the TTC security metric to quantify the risk of compromise. Security metrics such as TTC have attracted significant attention from the research community as a means to assess and prioritize various security risks as well as defense strategies. Among the earliest works of modeling and applying TTC metric are [117, 128, 129]. In [128, 129], McQueen et al. propose a basic model for estimating the time to compromise a specific control system. The model is leveraged to calculate the shortest path (in terms of its time) to reach and damage a target node of a system of interest. This model has been originally designed to provide estimates of the risk associated with potential attacks against critical elements of electric power systems, which are Supervisory Control And Data Acquisition (SCADA) control systems. In [117], Leversage el al. employ the same TTC model to estimate Mean Time To Compromise (MTTC) values of different systems and mitigation strategies used to enhance the security of SCADA systems. More recent research work such as [146, 234] proposes new models for estimating MTTC values of different security solutions and configurations applied in CI environments. They involve the use of vulnerability-based attack graphs. Each vulnerability represents a state in the final graphical model and has its own MTTC value. Ultimately, the final MTTC estimate is computed based on the MTTC values of the states and their Common Vulnerability Scoring System (CVSS)-driven probabilities. In [234], the MTTC metric is modified to evaluate the reliability of power systems using the IEEE RTS79 as a test system. The presented results show that the power system becomes less reliable with the increased rate of successful attacks on the cyber components. The main limitations of existing TTC models are threefold. Firstly, these models yield merely single-point TTC estimates. Such estimates do not account for the uncertainty, ambiguity, and variability of involved observational data. Further, they can convey misleading indications of extreme risks due to aggregation. Thus, they can not ensure robust and accurate risk measures. Secondly, the models shown in [128, 129] do not address the characteristics of potential zero-day vulnerabilities explicitly. Thirdly, the models in [146, 234] use vulnerability-based attack graphs, which suffer from the state explosion problem, where the size of the state space becomes quickly unmanageable. This can significantly limit the applicability of the models in real-world scenarios. To alleviate these challenges, the risk of compromise is quantified in this thesis using a developed TTC estimator that has the following features: (i) simple and easy to understand, even for non-professionals; (ii) practical through the use of asset-centric compromise graphs instead of vulnerability-centric attack graphs; and (iii) addressing the inherent uncertainty and variability of involved observational/statistical data using Monte Carlo simulation techniques. Therefore, the obtained TTC-based risk estimates are comprehensive and convey rich information on the two primary dimensions of risk descriptors, i. e., risk impact levels and their occurrence probabilities. The developed risk estimator can be leveraged to give indications on system robustness against not only technical vulnerabilities but also social and organizational factors. However, for

the sake of simplicity, the underlying TTC model presented in this thesis is limited to only software (technical) vulnerabilities.

2.3.3 *Cyber insurance games*

With the advent of advanced persistent threats and the growing illegal business models known as "cybercrime as a service", cyber-insurance has become a matter of increasing interest. Indeed, early work [109, 115] related to externality effects of cybersecurity in networks, while some more recent papers considered insurance more explicitly, such as [148, 149]. This thesis adds to this line of research the possibility to account also for moral hazards [91, 92], i.e., the willingness to take more risks because insurance is in place. Such a mindset can be reflected in the shape of the utility weighting functions. Generally, however, cyber insurance is a question of investment into cybersecurity [14, 36], and as such has received the most attention in the past concerning investment optimization and pricing of insurance for security; cf. [33, 84, 232] among others. Common to these past approaches is their focus on the nature of the threat and its implications. They paid, in contrast, only little or no attention to the subjective behavior of actors. This motivates proposing the *tweakable stochastic order* to account for this, which at the same time can simplify related game-theoretic models on the grounds of a more complex ordering than over the real values. Essentially, this is a challenge of decision making under uncertainty and risk. A notable contribution along these lines has been made by [46], which follows a different direction of seeking explanations of risk attitudes, while the focus here is on how to make such attitude models part of a preference ordering.

While many game-theoretic models for cyber insurance try to set utility values so as to reflect a person's choices as accurately as possible, bounded rationality research has shown that many such attempts failed. This motivates paying more attention to the ordering relation itself upon which rational behavior is defined, leading to the introduction of a *tweakable stochastic order*. This is a total ordering relation defined on random variables such as loss distributions known in actuarial science, which can be adapted to individual risk attitudes of players in an insurance game model. Therefore, this thesis proposes a new view on game-theoretic models for cyber insurance, by incorporating subjective risk attitudes into the choice preference rules of players, rather than into the payoffs.

2.4 ASSUMPTIONS

The different assumptions underlying this thesis will be explained and justified in the context of their respective sections. However, it is worth mentioning that the whole analysis conducted in this thesis rests on the assumption that "risk is quantifiable". This assumption is consistent with the above discussion in Section 2.2.1 that "risk involves that the probability and the mode of occurrence (i. e., consequences) are known, but the actual mode of occurrence is not". That is, security management results have to be interpreted and validated on the grounds of risk assessment samples available in each scenario.

One aim of this thesis is to study the possibility of reducing the chance of extreme risks or equivalently reduce the likelihood of the highest category of possible consequences in

the first place. Therefore, this thesis is not concerned with predicting *unexpected* extreme risks that are not covered by available assessments. To address the latter point, extreme value theory provides a wide range of statistical methods and approximation approaches that allow predicting the probability of extreme events that have not been formerly observed through extrapolating existing data [43, 145]. Nevertheless, if probability distributions of unexpected risks are available, then the methodological approach presented in this thesis can be used to infer the best decisions towards minimizing those risks.

3 THE EXTENDED PERIMETER OF CRITICAL INFRASTRUCTURES

Most CI systems have been operational for several decades, and hence they have been designed with neither widespread connectivity nor adequate security in mind. Gradually, owners and operators of CIs have started realizing the vital importance of securing and protecting their systems from accidental and intentional occurrences and against a broad spectrum of potential attackers ranging from amateur (cyber) criminal to advanced terrorist and state-sponsored attackers [8]♮. To date, the perimeter security model is still one of the most widely adopted security practices in these systems. It is obviously believed that building a protective hard shell around assets and resources of interest is adequate to keep potential danger out of the system. In this chapter, the perimeter security model is studied in light of current communication and business paradigms. This study shows that perfect security is not attainable, and managing security efforts and resources towards mitigating risks is, therefore, a much more practical possibility. Several parts of this chapter are based on previous research work appeared in [8]♮.

3.1 BEYOND TRADITIONAL BORDERS

The ongoing paradigm shift of CIs to be more collaborative and dynamic involves the adoption of emergent technologies and communication patterns, as well as new organizational structures and management models. For a long time, a solid and robust security perimeter has been deemed as a vital solution for an adequate level of protection. However, due to the increased rate of interdependencies and collaboration within and throughout today's complex systems, it has become very porous, thereby gradually failing to fulfill its core mission of keeping risks to the infrastructure's assets to a minimum.

Intuitively, protection implies the state of keeping the valuable assets, which an organization owns, manages, or controls from being damaged, stolen, or lost [8]♮. This, in turn, demands that those assets are well identified or easily identifiable. Defining the perimeter of an organization can immensely simplify the process of identifying the valuable resources not only for effective protection strategies but also for preparing proper coordination and management plans [8]♮. Clear boundaries will serve to identify the scope of responsibilities and activities to absorb potential disturbances and to recover from failures and disasters. Given the fact that CIs are increasingly becoming interconnected and interdependent, their perimeter structure becomes more vital to recognize different linkages and interconnections towards avoiding damages caused by other interconnected systems.

CIs and their respective assets and facilities usually tend to feature a sheer size spreading over large areas and across long distances connecting regions, which are geographically far apart. These systems can also cross regional and national boundaries passing through different environments and different environmental circumstances [8]♮. These environments can be characterized by different national, global, personal, orga-

nizational, business, and operational aspects. As a consequence, the large geographic span of such systems can solely responsible for making the process of protecting the respective perimeter an extremely challenging task.

CIs exploit ongoing advances in ICT for supporting the control and automation of their processes. The cyber and computational elements can include control systems of physical infrastructures as well as business and corporate network infrastructures. These elements aim at enhancing the ability to adapt to changes rapidly and subsequently to achieve resilient operation performance. Besides, CI systems can extend beyond their physical borders to include other entities such as vendors, business partners, service providers, or even costumers. As a result, various systems, which were previously isolated from each other by clear and well-defined boundaries, might be spontaneously and seamlessly integrated into one system crossing the boundaries marked by their traditional individual perimeters [8]‡.

3.2 DE-PERIMETERISATION VERSUS RE-PERIMETERISATION

Agility and flexibility are vital characteristics of recent CIs, which are evolving in a complex and uncertain business environment. As a result, the process of blurring or breaking down boundaries threatens such organizations more and more. This process is referred to as "De-perimeterisation", initially coined by a former chief security researcher at UK's Royal Mail Group - Jon Mescham [113]. Afterward, this term was adopted and promoted by Jericho Forum, an international working group hosted by Open Group, established to deal with the challenges associated with surviving in a network without boundaries. Therefore, Jericho Forum refers to de-perimeterisation as a concept or strategy that uses a mixture of inherently-secure protocols and components to protect organization data and systems rather than the reliance on a security perimeter. Simultaneously, the term de-perimeterisation is used to describe the process of a gradual dissolve of an organization's security perimeter, focusing only on the cyber world and data protection [8]‡.

Due to the current trend of a networked world, it is likely that organizations do not have their own Information Technology (IT) infrastructures or even do not have any control over it. Nowadays, complex systems are built on top of other systems and probably communicate with other ones [213]. Therefore, unknown and obscured connections and dependencies are currently not uncommon. This becomes inevitable in particular with the emergence of new business paradigms and technologies such as Business to Business (B2B), Business to Customer (B2C), Machine to Machine (M2M), cloud computing, virtualization, Internet of Things (IoT) and mobile internet. Having a perimeter as a protective measure would significantly impact the level of connectivity and hence strongly impede the envisaged business growth and collaboration. In other words, a security perimeter would act more in a blocking manner rather than in a facilitating or enabling one with respect to business objectives [8]‡.

Therefore, several technical and organizational mechanisms play a crucial role against the perimeter security model, such as increased network capacity, increased assets' mobility, transferring data, business and employee dynamics, service-oriented application, individual empowerment, among others [48]. Nevertheless, several other forces are opposing this trend, including the need for accountability, privacy, reliability, safety, and security that places an inevitable demand due to the rapidly increasing of secu-

rity incidents in critical infrastructures [8][1]. These forces are pushing on the opposite end against de-perimeterisation, leading to re-perimeterisation [213]. In other words, valuable assets have to be surrounded by a protective perimeter. This conflict between de-perimeterisation and re-perimeterisation motivates paying more attention to understand the security perimeter in CIs.

3.3 ASSUMPTIONS UNDERLYING PERIMETER SECURITY MODELS

Kevin Mitnick, a computer security consultant, stated that "*it is naive to assume that just installing a firewall is going to protect you from all potential security threats. That assumption creates a false sense of security, and having a false sense of security is worse than having no security at all*" [126]. In practice, security decisions depend mainly on both perception (feeling) of security and reality of security (probabilities and mathematical calculations). Both aspects are deeply related, but they are not the same [186]. Despite the mismatch between perception and reality of security, they are eventually affected by the same assumptions. Security systems' designers/analysts/users get used to implicitly (or explicitly) make assumptions when they design/evaluate/choose a security system or decision. Over time, most of these assumptions become obsolete, forgotten, or even invalid.

It is not uncommon that working environments of several organizations steadily change in response to various internal, external, business or regulatory forces. However, most of their adopted security solutions do not undergo a parallel revalidation process to figure out the new situation changes and to revalidate the underlying assumptions. People sometimes make assumptions intuitively. Most of the time, security decisions are made without even realizing the basic assumptions underlying the investigated system. Additionally, security teams depend first and foremost upon conventional and best practices, even without validating their applicability in the new environment. Since CI protection has different dimensions, this section focuses primarily on a set of assumptions that are closely related to perimeter security models. It is noteworthy to point out that some of these assumptions are not necessarily valid in modern CI systems.

DANGEROUS OUTSIDE AND SAFE INSIDE: It is still widely adopted that the inside of an organization is a trusted environment and not prone to danger and attacks like the open outside. For that reason, the internal zone demarcated by a protective wall (e. g., using a firewall or an intrusion detection system) is usually equipped with minimum or even zero levels of protection. Thus, once an adversary gets past the protective wall, he will be able to act inside in an unimpeded manner. Most organizations adopt the perimeter model for security since protecting one entity (i. e., perimeter) is a more manageable and economical option in comparison with securing large internal networks and subsystems.

COMPLETE AWARENESS OF THE LOCATION OF THE ASSETS: In fact, static assets would enormously simplify the design and deployment of monitoring and protection activities. Accessing the assets only during the working hours and only from inside the organization will extremely help to have a physical oversight of the accessed and accessing devices as well as the involved employees. Even with limited and well-

defined movement directions (e. g., unidirectional data flow from industrial assets to the control room), perimeter-centric solutions are still deemed effective. Nowadays, with the prevalence of cloud services and bidirectional communications for both data and control, the ability of conventional perimeter solutions to keep adversaries out is undoubtedly questionable. Any internet-facing device can be compromised to initiate a connection to target systems. Mobility of accessing or even accessed assets will basically lead to a lack of control.

OBEDIENT EMPLOYEES: Security administrators and policymakers assume that organizations' employees will literally adhere to prescribed security policies. They often overlook the fact that human factors[1] and potential administrative changes might conflict with the prescribed security policy. In daily life settings, people are more interested in getting their job done. If security measures and security perimeter are slowing them down or limiting them, they will try to find ways to circumvent such controls. Currently, many employees use cloud-based services for collaboration as well as business data storage and sharing without any specific policies in place. Lack of staff training and lack of security policy awareness are essential factors for the increased rate of security breaches. Shey et al. reported in a study of information workers in North America and Europe that "only 42% of the employees have received training on how to stay secure at work, and only 57% indicated that they were aware of their organization's current security policies" [195].

CONTINUOUS OPERATION OF THE PERIMETER: Under some circumstances (e. g., a firewall's underperformance), security managers will be put under stress. They might be asked to simplify the procedures in order to speed up the communication process. In case of attacks or broken firewalls, system administrators could reconfigure the firewall or open some unprotected pathways to some critical systems. As a result, the perimeter defense is not always operational and effective at keeping the external danger away.

OWNERSHIP OF THE ASSET: It implies that there is a single owner or operator of the asset. Thus, it is possible to limit the access of third party entities, both technically and organizationally. However, organizations, in particular CIs, increasingly cooperate to achieve benefits such as cost saving, agility, load balancing, responsiveness, interoperability, and jointly coping with the increased demand on their services and products. This cooperation eventually imposes that access patterns of individual assets and resources have to be modified, allowing for shared access and usage. As a result, CI organizations have gradually started to lose the full control over their assets. Another consequence is that lack of clarity of roles and responsibilities within individual systems and among interdependent systems will certainly contribute to incident response delay.

ISOLABLE SYSTEMS: Due to the high reliance on CI systems, most of them are mainly managed and controlled by cyber systems. In order to avoid the potential mass destruction caused by cyber-attacks, it is still common that critical systems have to be disconnected from the Internet or any open network, creating so-called "air-gapped systems". An air-gapped system is physically segregated and, therefore, unable to connect with other systems and networks. However, the main concern is whether all

1 Human factors refer to situations when human actions or errors lead to a successful attack or damage.

control components of the system of interest are truly air-gapped. CI systems are complex systems (or system of systems). Hence, it is not uncommon to contain some control (sub)systems with unintended and unprotected network connections (such as wireless connections) bridging the alleged air-gap and expose the whole system to vulnerabilities. This raises the concern of whether it is possible at all to truly isolate systems in the era of IoT, mobile Internet, and cloud computing. Stuxnet is one of the most famous attacks that could breach the protective perimeter and bridge the air-gap around Iranian nuclear facilities. The air-gapped systems were infiltrated by introducing an infected USB flash drive into the trusted zone. AirHopper [79], in turn, uses another technique to bridge the air-gap in the opposite direction (i. e., data exfiltration) between physically isolated computer and nearby mobile phone using electromagnetic waves.

FIXED-SIZE PERIMETER: As CIs grow up, the value of their assets and constituent elements will increase, as well. The first consequence of this growth is that the number of valuable and critical systems behind the organization's perimeter will alike increase. Although the importance of a security perimeter is proportional to the number and value of the systems behind, its performance and effectiveness are as much inversely proportional to this number. Gray et al. have referred to this phenomenon by "perimeter protection paradox" [77]. Because of the large number of systems behind the perimeter, the function of the perimetric security measures will be significantly impaired to facilitate the operation of the various system applications communicating through. Moreover, the perimeter will gradually turn into a bottleneck and single point of failure, unable to handle an increasingly growing load and enormous requests through it.

CENTRALIZED SYSTEM: For a long time, organizations' assets and resources have been distributed on single or few well-known sites. The vast majority of users (employees) accessing those resources were physically inside the organization. Therefore, the infrastructure can be safely protected by a solid surrounding security perimeter. Being inside the perimeter was a condition to cause damage and launch attacks. However, with the advent of new technologies and communication paradigms, CIs extend accordingly to include highly distributed frameworks.

ROUTINE ASSESSMENT: The routine assessment of system performance and vulnerabilities is a vital step of protection and risk management processes. Nevertheless, many security teams overlook the practice of testing their infrastructures (e. g., regular penetration testing). As a result, they miss the opportunity of proactively discovering exposures and hence leave their most valuable assets at risk. The routine assessment is one of the necessary foundations on which almost every security practice is built. This activity has to be scheduled and conducted consistently, not only once during the design process.

CLEAR DISTINCTIONS BETWEEN PRIVATE LIVE AND WORK LIVE: Currently, the majority of the workforce belongs to the generation that has grown up with the Internet and ubiquitous connectivity as inherent elements of their daily life, so-called Millennials [116]. They intensively depend on technology and personal communication devices. Thus, they interact with their environment differently from the previous workforces. Since they expect the connectivity everywhere, their productivity relies almost entirely

on technology and available network access. The perception of their roles, time, and space is different, which has significantly affected their way of thinking regarding outside and inside as well as private and business. From their perspectives, private and business spaces are interleaved in such a way that they can do some work outside their prescribed working hours and some non-work issues during their working time [116]. Furthermore, their sense of privacy and security is different. They might, therefore, work on some sensitive business information at a public place. They can allow that their locations are tracked and profiled via their companioned devices, which is a significant concern for some specific organizations such as police [85].

CLEAR EMPLOYMENT RELATIONSHIP: Until recently, workforces were employed by one organization and set within its facilities. This practice has currently changed, and some employees can be members of several organizations at the same time [67]. This fact makes it challenging to understand the normal employee behavior and to determine whether an employee's actions are driven by his role inside or outside a specific organization.

BEST-PRACTICE-BASED SECURITY STRATEGY: Many organizations follow in their security policies and strategies best-practice measures. As a consequence, the security strategies of these organizations are becoming predictable and certain, more and more. Anticipating the security posture of a system will give attackers an important opportunity to stay ahead in the security game. In other words, having the same security solutions results in protective monocultures where the systems share the same threats and weaknesses. Therefore, an attacker will invest and focus on one tactic to collect information about a single system, knowing that it can be easily leveraged to attack any other similar system.

PROPRIETARY PROTOCOLS AND APPLICATIONS: For a long time, it has been a common practice to use dedicated proprietary protocols and purpose-built software for operating systems and devices produced by a single manufacturer. The main reason for such practice is that most of the deployed devices are functionally limited concerning computing and communication capabilities. Therefore, using lightweight dedicated software and protocol is undoubtedly more preferable, especially in case of strict industrial operation conditions (e. g., real-time monitoring and control). The second reason is the prevalent belief in the industrial security society that proprietary systems are obscure, and their security is wholly ensured as attackers are unaware of the systems' mechanisms, designs, and implementations. However, this belief is no more valid, in particular with multi-vendor systems, in which open protocols are widely adopted for the sake of seamless operation and networking. From the perspective of infrastructure operators, increasing systems' agility and configurability using general-purpose entities outweighs the benefit of using purpose-built solutions that could limit failure propagation to only parts of the system that share similar exploits. In contrast to proprietary protocols, standard protocols are well described in the public domain. Their weaknesses and exploits are publicly known as well. General-purpose protocols will significantly impair the effectiveness of perimeter-centric security measures since they are no more able to monitor and filter unwanted traffic as it could be easily hidden and encapsulated in other allowed data flow.

THE SYSTEM IS NOT UNDER RISK: This belief stems from the assumption that the system is well designed and segregated from outside, where all sources of damage present. This assumption can be further reinforced if the system has never failed before. In such a situation, the incentive to invest in alternative (more advanced and costly) security approaches is considered small. Moreover, the feeling of being safe could sometimes result in a tendency to reduce existing precautions for unmeasurable security and safety to achieve some measurable financial benefits [119]. For example, most CI owners and operators could not clearly and distinctly perceive the risks of interconnections and mutual dependencies among their systems since these risks could not become apparent unless something goes wrong. In 2008, for example, the International Atomic Energy Agency (IAEA) had expressed concern about the ability of Japan's nuclear power plants to withstand strong earthquakes. The design basis of the Fukushima plant is to withstand an earthquake of magnitude 7.0, while the recorded earthquakes exceeded this limit in some cases. However, Japanese safety experts believed that "all safety analyses were appropriately conducted" [154]. After the incident in 2011, some (overconfident) experts argued that their systems have been operating safely for years, and their safety or security decisions are therefore robust. This reasoning, however, is utterly flawed since the absence of failure cannot prove that the precautions were correct or even sufficient [90].

3.4 WHAT IS A SECURITY PERIMETER?

It is worth looking at some definitions of the perimeter in an attempt to understand its value for existing systems and networks[2]. The Cambridge Dictionary, for example, defines the term perimeter as "The outer edge of an area of land or the border around it" [55]. While the American Heritage Dictionary defines it as "A defended boundary protecting a military position" [54]. Obviously, both definitions are subject to interpretations varying according to the considered context. Nevertheless, one can at least figure out the basic functionality of a perimeter from these definitions, namely demarcation and protection. Demarcation refers to the process of setting or marking boundaries or limits [54]. In turn, protection is associated with the preservation from harm, destruction, and loss as well as unwanted activities. Combining both aspects leads to the conclusion that a perimeter is the outer boundary that protects the inside holdings from the outside danger. This conclusion is consistent with the definition: "a security perimeter is a technical solution to protect assets from negative influences originating in its environment" [48]. According to this definition, sources of threats and negative influences to an arbitrary inside asset are located mainly in its surrounding outside environment. The perimeter is, therefore, a borderline, at which the inside ends and the outside begins. Hence, it is the first entity that any external entity from the outside has to come into contact with, prior to infiltration into the inside. Depending on this discussion, this term can be perceived as follows:

A security perimeter is any (simple or composite) entity that surrounds physically or virtually a facility with its various assets, and through it, the contact (communication paths) between the internal (inside) world, i. e., the assets to be protected, and the external (outside) world will be established and facilitated.

2 The content of Section 3.4 and its subsections is based on the work appeared in [8]‡.

3.4.1 *Structure of a security perimeter*

Based on the above definition, a security perimeter can be presented as a component or a (sub)system that is responsible, firstly, for maintaining the required isolation and separation between the internal private zone and the external public world (i. e., the perimeter encompasses the solid defending wall surrounding the inside). Secondly, it is responsible for provisioning the appropriate paths of communication from and to the internal zone (i. e., the perimeter embraces the doors allowing access to the inside). Thus, the core components of a security perimeter would be:

- A security wall: it surrounds the system to be protected, tearing apart the area (space) into two areas, inside and outside. The inside refers to the internal area behind the wall and towards the target assets. In contrast, the outside refers to the external area, which starts at the wall and spans away from the target. Thereby, the wall maintains the inaccessibility of the inside from the outside and vice versa. Therefore, its function is always to prevent access regardless of the circumstances.
- A security door(s): it represents a controlled and monitored pathway for entering and leaving the internal zone. In the case of unauthorized entities, it integrates to the defending wall to maintain isolation (closed mode). In contrast, in the case of legitimate entities, it provides a way to circumvent the surrounding solid wall (open mode). In other words, a door is (should be) merely a wall for those without a proper key. Through doors, diverse permissible communication paths from and to the internal zone are provisioned and mediated

There is another element, which is unwanted but most likely not entirely avoidable, called a security hole. It refers to a location where the wall fails to fulfill its mission in maintaining inaccessibility. Security holes are attractive elements for potential intruders since they are always-open doors with open (unauthorized) access connections.

3.4.2 *Nature and function of security perimeter*

A security perimeter is responsible for keeping adversaries and malicious entities away, on the one hand, and preventing or reducing loss, leakage and theft of assets and resources such as sensitive data, on the other hand. Simultaneously, the perimeter plays an essential role in protecting the external world from damages stemming from the internal world. The perimeter, as the primary provider of contact points, defines the components that constitute an organization's attack surface. These components have to be appropriately controlled and managed in order to achieve the envisaged protection and security. The contact points (paths) are typically provisioned by door-entities that are separated from the protected assets. They include, for example, firewalls, Virtual Private Network (VPN) gateways, Network Access Control (NAC) devices in the cyber world, as well as gates for authorized ingress and egress to a respective facility, biometric access control systems and security guards in the physical world. These solutions are primarily employed to provide valuable protection and control capabilities. They are deployed at the borders of the protected system serving as the first line of contact and defense. Nevertheless, some system elements can immediately communicate with external entities without any mediation through the deployed security perimeter,

including control devices with wireless internet connection, resources delivery network infrastructure, mobile devices, human resources, or even data itself. Control devices with data or internet connections can provide paths for data to move and migrate or provide remote access to the internal network. Delivery and transportation grids of provisioned resources, such as water and gas pipeline or power distribution networks, are ultimately difficult to be confined behind a dedicated perimeter infrastructure since they often extend for very long distances. Mobile devices and human resources cannot be consistently confined behind static and predefined perimeter structure, as well. Data can likewise have an unmediated contact with the external untrusted environment since it might be exported and exchanged with other organizations or individuals for the purposes of support and collaboration. In this way and according to the previous definition, these elements have to be immediately moved to be part of the perimeter since they are directly in contact with the external world entities and have to be appropriately controlled, monitored, and managed. Otherwise, these elements will result in numerous holes in the protective perimeter, rendering it more porous and less effective. It is not difficult to imagine the number of potential holes in an organization's perimeter with one thousand employees in a workplace, and everyone is equipped with an individual tablet or mobile phone. These devices can easily bridge the gap between external and internal zones through the available simultaneous access to a corporate network and public network in the same entity. In order to avoid any potential confusion and to consider the new dynamic nature of the perimeter due to the ongoing technological development, this thesis divides the perimeter of CIs into two subgroups, *non-extended and extended perimeter*. The former would refer to the diverse entities that aim at ensuring security and safety of a system, while the latter would oppositely refer to potential attack vectors. Hence, there is a vital need to pay special attention to extended perimeter components.

3.4.3 *Non-extended perimeter components*

Traditionally, a security perimeter is usually designed and implemented with the D5 strategy in mind; standing for Demarcation, Deter, Detect, Delay, and Defend [138]. Demarcation refers to the process of creating virtual and physical boundaries around a facility's assets, such as buildings and data. In the context of physical security, these boundaries should be visible to avoid innocent boundary crossings and to simplify identifying hostile intentions. The goal of deterrence is to create an unattractive environment for potential adversaries, thereby reducing an adversary opportunity to commit an attack unobserved. However, deterrent efforts are not enough to keep adversaries out. Therefore, it is of vital importance to detect unwanted activities and to delay potential perpetrators long enough to allow security forces or first responders to intercept and defend by denying access to the internal critical assets and resources. In the context of physical and cybersecurity, a security perimeter can include fences, walls, monitoring points, entrance gates or doors, vehicle barriers, security lighting, landscaping, surveillance systems, alarm systems, guards, firewalls, routers, access control devices, intrusion detection systems, among others. All of these measures have been used to ensure that any contact with critical assets is authorized by a predefined perimeter before taking place within the internal zone. Due to the apparent static nature and

predictable placement of these mechanisms, this thesis refers to them as *non-extended perimeter*.

3.5 IDENTIFYING EXTENDED PERIMETER COMPONENTS OF CI SYSTEMS

In contrast to non-extended perimeter components that provide diverse monitoring and controlling functions, there are other entities and system elements, which belong to the infrastructure to be protected. However, they are directly accessible from the external world without any intermediate mechanisms controlling and managing the communication paths [8]♮. As a result, these entities constitute an extension to the standard (non-extended) perimeter since they can provide unobserved communication paths with the outside. While some of these elements can presumably offer a certain level of security, the vast majority of them pose severe threats to the system of interest. Broadly speaking, an organization's extended perimeter encompasses all components that have the potential for circumventing and bridging the air-gap maintained by the non-extended perimeter. They are ordinarily legitimate components allowed to access the inside environment and hence can serve as potential carriers of infection for other interconnected entities. The extended perimeter components are almost involved in two roles: (i) operational role as a part of the business and operational process, and (ii) perimetric role due to the ability to communicate with the external world immediately. This thesis identifies basically four major elements of an extended perimeter as detailed in the following. The following subsections of Section 3.5 are based on the work appeared in [8]♮.

3.5.1 *Unattended infrastructures*

This group encompasses all elements which are not explicitly or implicitly behind any non-extended perimeter infrastructure, and they are obviously situated in a potentially hostile environment. For examples:

- Transportation grids for the provided resources: examples of such elements are water and gas pipelines as well as power transmission and distribution networks. Their share scale of spread makes it almost impossible to deploy and manage a typical defending perimeter around them. Unattended on-field control stations are also elements of this group.
- Control system devices: These devices are increasingly outfitted with networking capabilities to allow remote access and configuration. Therefore, they are directly accessible from outside through modem access, Digital Subscriber Line (DSL) technology, wireless access, or VPN tunnels that can be easily exploited to bypass installed firewalls because of its encrypted nature.
- Mobile delivery entities: fuel delivery vehicles, as well as bulk oil delivery tanker trucks and vessels, are excellent examples of these elements. For attackers, these entities are relatively easy targets with high potential of disruption to the respective supply chain. Adversaries can also leverage them against other potential targets due to their mobile nature.

These elements represent a serious concern to CI operators. Their specific nature makes potential attacks more dramatic and devastating. For example, the economic loss and

environmental effects associated with an explosion of a major oil or gas pipeline (of hundreds of kilometers) could be really of high-order dimensions.

3.5.2 *Trends and technology populism*

Traditionally, an organization's IT department is solely responsible for planning and choosing the organization's IT infrastructure as well as provisioning employees with the necessary hardware and software they need for their daily job. Therefore, the training budget is an essential and inevitable part of any system planning process since employees have to be trained on how to use these tools and applications to ensure conveying skills necessary for productivity. Nowadays, employees' productivity is, however, increasingly dependent on the use of personally owned technologies and systems, in particular for the tech-savvy workforce who grew up with technologies. The key enablers for this trend are usability in the sense of ease-of-use of owned systems and IT self-supporting, which outweigh basic security requirements from the employees' perspectives. The trend of bringing tools and applications designed to be used at home for personal usage to workplaces for further business purposes is referred to as *technology populism* [153]. Online collaboration, ubiquitous internet, social networking, and mobile devices are all ways to increase employee happiness and productivity. Yet, technology populism brings security threats and risks alongside its benefits, such as:

- The organization does not own the whole IT infrastructure used in its business operations, and hence they do not have a valid assurance of quality-of-service.
- IT managers are not able to train and to provide support for every tool and application used by the employees. This results in creating a very complex support environment and a lack of centralized administration.
- Security managers have neither full control nor full awareness of the current set of applications and information resources used by the employees. Hence, it is challenging to identify the organization's exposure level and security posture.

Technology populism exposes the internal system to numerous insecure connections that can be easily exploited by a criminal to get privileged access, such as using public information sharing platforms, infected laptops, or mobile phones. Employees' mobile phones can easily move between domains of different security levels bridging the gap provisioned by a non-extended perimeter. Contractors and vendors, in their turn, bring their devices and machines to perform on-field maintenance operations, and they have, therefore, to be connected to the internal network. Mobile communication and computing devices include portable media (i. e., flash memory devices or portable hard disk drives), laptops, tablets, and smartphones. These devices, with or without Internet or wireless connections, have the ability to connect and bridge two different networks. For example, if a laptop can be configured as an open access point, it can readily and spontaneously provide a connection between an internal corporate network and a remote infected host. Mobile devices, even without wireless communication, can provide unchecked paths since they can move between zones of different security levels, such as home and corporate networks. These different means have enabled not only legitimate users but also adversaries to remotely access assets and resources and to circumvent the deployed (non-extended) perimeter easily.

3.5.3 *Outsourcing*

Lack of skilled personnel, limited resources capabilities, as well as high costs for operations and maintenance, are all driving forces behind leveraging outsourced functions and resources. Consumers and end-users increasingly demand online services and ubiquitous access to their data and transactions. Thus, in order to ensure an adequate level of performance and efficiency, many organizations and CI operators decided to outsource part of their business functions. Outsourcing some functions to external partners will allow further focusing on core assets and valuable resources. Thereby, a system management process will be performed more effectively and efficiently. Furthermore, outsourcing will play an important role in the cost reduction process and in reaching a skilled and competitive workforce over the wide world. It also enables flexible and on-demand usage of up-to-date technology with a minimum set of control, management, and maintenance activities. As a result, organizations and CI operators are increasingly engaged with outsourcing service providers to ensure agility and flexibility of their systems cost-effectively and to increase responsiveness to steadily changing business and environment conditions as well as increased demand for their services. To ensure quality of service delivery and business continuity, an external outsourcing partner is usually provided with access to the internal organization infrastructure network. This will open the door for more weaknesses and threats, providing potentials intruders with unprecedented access to other valuable assets.

In this regard, the main risk stems from the lack of control over resources and the potential discrepancy between priorities of outsourcing service providers and their clients. This mismatch or disharmony between both sides can easily result in exposing critical assets to unauthorized access that could impact their confidentiality, integrity, and availability and consequently safety, reliability, and availability of the whole CI systems. Snowden disclosures, for example, have revealed that National Security Agency (NSA) had access to data stored at some American cloud-based services and servers, such as Google and Yahoo. The NSA spied on users through (i) collecting plenty of emails, contact lists, and search content, as well as (ii) tracking and mapping locations via mobile phones. Basically, encryption is used to create a protective perimeter around valuable and outsourced data. In this case, data cannot be separated from its perimeter. Therefore, any owned or collected data, which is stored, replicated to, or processed on off-premise resources, exchanged with other parties, even if it is encrypted, is part of the organization's extended perimeter.

3.5.4 *Human factor*

The increased adoption of national and international collaboration and partnership models justifies the constant tendency of such systems to extend beyond their conventional physical existence to include other entities such as vendors, business partners, service providers, or even customers. Consequently, it is now a prevalent practice to see different external entities within a system complex, such as temporary workers, interns, independent contractors and subcontractors, or visitors. Even if access to sensitive industrial zones is tightly controlled at the borders, the freedom of movement is most likely ensured behind the borders for all regular employees and temporary workers. Potential adversaries can exploit the dynamic nature of the systems as well as the lack

of proper human resources management strategy to cause damage. For example, an infected contracted programmer can become an unknowing carrier for malware. He can inadvertently infect air-gapped systems via transferring data from his own laptops and external hard drives.

Basically, human resources cannot be consistently confined behind the static and predefined perimeter. Moreover, entities, such as employees, can also exploit their knowledge and privileged access to open paths of communication between inside and outside bypassing installed access control mechanisms. Lack of risk awareness of employees can lead to many disastrous results. Furthermore, persons and their desires and motivations are the most difficult components to be constrained.

Visual summary

Figure 3.1 illustrates those above mentioned extended and non-extended perimeter components of CIs. The non-extended part of the perimeter represents (i) the wall, which maintains the segregation with the surrounding hostile environment (e. g., via fences, physical segregation of devices, virtual isolated domain); and (ii) doors, which provide a controlled and monitored access to the internal world for legitimate entities (e. g., using firewall, NAC gateway). Monitoring and surveillance systems of the non-extended perimeter provide twofold security enhancement, namely (i) increasing the deterrent effect of the perimeter, thereby discouraging an attacker from committing undesired activities and crimes; and (ii) maintaining a wide-area situation awareness to boost the system responsiveness [11]. Broadly speaking, the non-extended perimeter is responsible for controlling and providing authorized communication paths (green arrows in Figure 3.1) with the internal protected systems. In contrast, the extended perimeter part would create holes (red lines in Figure 3.1) in the respective defending wall impairing the security posture of the own system and allowing adversaries to bypass the traditional security measures.

3.6 CLASSIFICATION OF SECURITY INCIDENT CAUSES

To further understand potential security risks posed by an extended perimeter of a CI system, the following four real security incidents have been selected. These incidents have occurred during the last years. They were mostly known for their significant effects in terms of significant financial and reputational losses. Table 3.1 shows which components of the extended perimeter have considerably contributed to the respective incidents.

- **Stuxnet:** In 2010, Stuxnet malware was discovered. It has targeted Iran's nuclear enrichment program. This malware reaches its targets like industrial control devices using human vectors as well as infected devices and files. A contractor was infected first and then became an unknowing carrier for the malware. Inadvertently, he infected the air-gapped systems via transferring data from his own laptops and flash drives [59].
- **Metcalf sniper attack:** On 16 April 2013, the Metcalf transmission substation located outside of San Jose, California, was attacked using rifles. The attack resulted in severely damaging 17 giant electrical transformers, damage cost of

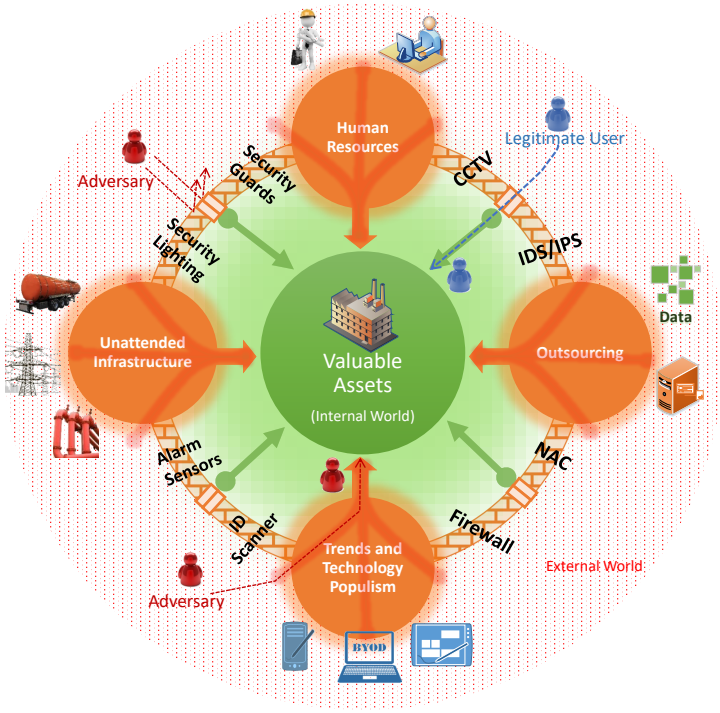

Figure 3.1: Overview of extended and non-extended perimeter components of CIs [8]♮

millions of dollars, and 27 days for repair and recovery [200]. Before opening fire on the electrical station, the communication service to the surrounded area was also knocked out by cutting telecommunication cables in an adjacent unattended underground vault. It was one of the most significant attacks involving a power grid that has ever happened.

- **NSA Breach:** In 2013, Edward Snowden, a former NSA contractor, leaked classified materials about NSA's secret surveillance program and its tremendous volume of gathered surveillance data. He obviously exploited his top-secret access to breach one of the best-secured organizations in a hidden manner [209].

- **AWS outage:** In June 2012, a power outage affected the availability of an Amazon Web Services (AWS) data center in northern Virginia. The outage resulted in service interruption of all services hosted in the affected data center, in particular cloud-based businesses such as Quora and Pinterest [133].

Table 3.1: Security incidents caused by extended perimeter elements[a]

INCIDENTS	UNATTENDED INFRASTRUCTURE	HUMAN FACTOR	OUTSOURCED RESOURCES	TECHNOLOGY POPULISM
Stuxnet		x		x
Metcalf attack	x			
NSA breach		x		
AWS outage			x	

[a] This table is not intended to be exhaustive, but rather to give some relevant examples.

3.7 SUMMARY

The extended perimeter a CI organization consists of all objects that are (i) entirely or partially owned or controlled by the organization; (ii) directly or indirectly impact the business and operational processes, and (iii) not fully isolated from the outside world by a well-defined non-extended perimeter. Potential attackers rely on the static nature of non-extended perimeter mechanisms and the predictable placement of their devices to sneak into a system of interest. To achieve that with less effort, they can exploit hidden paths provisioned by the extended perimeter components of the respective system. Undetectability within the system complex will give an attacker an excellent opportunity to reconnaissance the target system, gather some sensitive information, and probably cover the tracks of ongoing attacks, too. Therefore, securing CI systems can never be 100% effective, and hence it is essential to manage (i. e., optimize) security efforts and resources not only at the borders but also within the system complex towards reducing security risks.

4

RISK-BASED SECURITY MANAGEMENT: A METHODOLOGICAL APPROACH

As explained in Chapter 1, security management problems are concerned with how to create a coherent security strategy that leaves CI organizations well-positioned in the security game against potential adversaries. Therefore, they call for a process of orchestrating security efforts and optimizing the allocation of limited security resources in a way that reduces the impact and likelihood of security risks. This chapter is dedicated to addressing this point by introducing a methodological approach for security management in CIs that uses game theory to model the natural competition between defenders and potential attackers. Hence, it describes a game-theoretical approach that takes into account the different requirements of security management problems as described in Section 1.2. The content of Chapter 4 is based on the research work published in [2, 8, 10]♮.

4.1 GAME-THEORETIC APPROACH FOR RISK-BASED SECURITY MANAGEMENT

This section starts with sketching some game theory principles to explain necessary mathematical notations, before going into detail on the game model used in this thesis. The content of this section is based on the research work published in [2]♮.

4.1.1 *Game theory principles*

With regards to the components of game-theoretical models, almost all of game theory considers individual action (strategy) spaces $\mathcal{A} = \{A_1, \ldots, A_n\}$ for distinct players (decision-makers) named $\mathcal{I} = \{1, 2, \ldots, n\}$, each of which receives a value (also called payoff or utility) $u_i(a_i, \mathbf{a}_{-i}) \in \mathbb{R}$, depending on its own action $a_i \in A_i$, and the compound actions of its opponents; here denoted by the symbol $\mathbf{a}_{-i} \in \prod_{\substack{j=1 \\ j \neq i}}^{n} A_j =: A_{-i}$ (the subscript $-i$ consistently with the literature denotes "all coordinates except the i-th", and vectors are denoted as bold-face lower case letters) [2]♮. Thus, besides the players and their strategies, the third major component of a game-theoretical model is a utility function[1] that defines the players' preferences in terms of establishing a ranking of the different decision alternatives (i.e., strategies). That is, given a utility function u on $\mathbb{R}$ and two actions x and y, the preference relation $\preceq$ is defined by the condition:

$$x \preceq y \quad \Longleftrightarrow \quad u(x) \leqslant u(y)$$

1 Some literature on game theory explicitly differentiates a utility function from a consequence function. The latter, defined as $g_i : A_i \to C$, associates each action from the set A_i with a consequence from a set of possible consequences C. A utility (or loss) function, $u : C \to \mathbb{R}$, assigns a value for each consequence, thereby defining a preference relation $\preceq$ on the set C.

In classical game theory, the utility function is for the i-th player a mapping $u_i : A_i \times A_{-i} \to \mathbb{R}$, and the rational[2] player's objective is to optimize, say minimize, u_i against what all other players do. Therefore, a rational decision-maker i choose an optimal action a_i^* such that:

$$a_i^* \in \arg\min_{a_i \in A_i} u_i(a_i, \mathbf{a}_{-i})$$

That is, when other players' use a compound strategy $\mathbf{a}_{-i}$, player i can select an action a_i^* that minimizes its utility function. The action a_i^* is called the best response of player i:

$$u_i(a_i^*, \mathbf{a}_{-i}) \leqslant u_i(a_i, \mathbf{a}_{-i}) \quad \forall a_i \in A_i$$

To study the interaction among players, the notion of equilibrium in games is introduced by the mathematician John Nash [141]. A *pure strategy Nash equilibrium* of a strategic (non-cooperative) game $\langle \mathcal{I}, \mathcal{A}, (u_i)_{i \in \mathcal{I}} \rangle$ is the combination of all players' best response strategies a_i^* (one for each player, given what the other players do). In other words, Nash equilibrium represents a game solution in which no individual player obtains better payoff by changing only its own strategy. Formally, pure strategy Nash equilibrium is referred to as a strategy profile $\mathbf{a}^* = (a_1^*, \ldots, a_n^*) \in \prod_{i \in \mathcal{I}} A_i$.

In many concrete instances of a game, the optimum will not be attained within the action space A_i. This happens when players have to be unpredictable in their play. For example, if the game is coin flipping between two persons, then player 1 has two strategies in its action space $A_1 = \{heads, tails\}$. Assume that the player loses if the opponent correctly guesses how the coin comes up, so the action space for player 2 is also $A_1 = A_2$. But if any of the actions of player 1 were optimal, then player 2 can always win the game by taking exactly that (known) guess. Another simple example is penalty kicks in soccer, where any player's success relays upon his action being unpredictable. If the goalkeeper knows to which side the kicker will shoot he will always choose that side to defend, and vice versa [49]. Hence, there is no optimal "pure" strategy among the actions, and players have to randomize their calls.

This amounts (for all players) to choosing actions according to an optimized probability distribution $\mathbf{a}_i^* \in \Delta(A_i) = \{(p_1, \ldots, p_{|A_i|}) : p_k \geqslant 0 \ \forall k \in \{1, \cdots, |A_i|\}, \sum_{k=1}^{|A_i|} p_k = 1\}$. Here, $\Delta(A_i)$ is called the *simplex* over the set A_i and assume A_i to be *finite* and static hereafter[3], i. e., remain constant (unchanged) over time or repetitions of the game. This technical change convexifies the space A_i and, at the same time, induces the need to redefine what a "best" action in $\Delta(A_i)$ would be. Essentially, an element $\mathbf{a}_i^* \in \Delta(A_i)$ is called an *optimal randomized choice rule*, if $E_{\mathbf{a}_i^*}(u_i)$ is minimized; that is, the *expected loss* according to the randomized choice of actions based on the distribution $\mathbf{a}_i^*$ should be optimal, given what the other players do. Adopting this convention for all players, one gets a *mixed strategy Nash equilibrium* in the game to be a set of simultaneously optimal choice rules (or mixed strategy profile) $(\mathbf{a}_i^*, \ldots, \mathbf{a}_n^*) \in \prod_{i \in \mathcal{I}} \Delta(A_i)$ so that each player minimizes its own loss as $\forall i \in \mathcal{I}$:

$$E_{\mathbf{a}_i^*} u_i(\mathbf{a}^*_{-i}) \leqslant E_{\mathbf{a}_i} u_i(\mathbf{a}^*_{-i}) \qquad \forall \mathbf{a}_i \in \Delta(A_i)$$

2 Players acts rationally in terms of selecting the option that delivers them the best payoff (i. e., higher benefit or less loss), given their own beliefs about the behavior of their opponents.

3 More general versions of games that admit a change of the utility during repetitions of the game are more complex to treat and are outside the scope of this work.

That is, no unilateral deviation from the optimum $\mathbf{a}_i^*$ could further reduce the losses for the i-th player when the game is played at equilibrium by all players. This convention implicitly assumes that the game is *repeated* (for infinity, since the optimization is over the long-run average payoffs) and that actions are chosen afresh between independent repetitions of the gameplay. At the same time, the randomization of strategies ensures the existence of equilibria, which without repetition would need to exist within the pure strategies, and there are many counterexamples (in addition to the above) of games that lack such equilibria in pure strategies.

4.1.2 *Security management games*

Based on the description of security management problems in Section 1.2 and especially **Req1** and **Req2**, the game model presented in Section 4.1.1 can be simplified to a two-player non-cooperative game, in which a defender $\mathcal{D}$ (player 1) engages in a competition against an attacker $\mathcal{A}$ (player 2), who seeks to cause maximal damage to the defender. The latter abstracts all external adversaries that seek to benefit from a system's weaknesses to cause damage. In contrast, $\mathcal{D}$ abstracts any decision-maker (e. g., chief security officer or patch management operation team) seeking to minimize the risk of compromising and damaging the respective system [7]$^{\natural}$. $SP_{\mathcal{D}} = \{d_i\}$ denotes a finite set of the security actions (e. g., vulnerability remediation activities, security inspection schedules) the defender can perform to defend the system in question. Additionally, the set $SP_{\mathcal{A}} = \{a_i\}$ represents the potential ways the attacker can use to compromise and cause damage to the system[4] [7]$^{\natural}$. The cardinality of $SP_{\mathcal{D}}$ and $SP_{\mathcal{A}}$ are n and m, respectively.

To deal with the lack of reliable information about type and payoffs of a potential attacker (see **Req5**), minimax principle[5] comes into play [142, 221]. This principle is used to optimize the defender's decisions against worst-case attack scenarios. That is, CI organizations should be prepared for the worst, no matter what happens after the decision is made [7]$^{\natural}$. In this case, it is assumed that whenever a player gains some benefits, the other player must lose the same amount. That is, the players' payoffs are sign-opposite defined as $u_2 := -u_1$, and the game is therefore *zero-sum*. For zero-sum games, the utility function u_1 can be represented by a matrix A (simply telling the loss under every possible combination in $SP_{\mathcal{D}} \times SP_{\mathcal{A}}$) [2]$^{\natural}$. Furthermore, the expected loss under randomized choices $\mathbf{x} \in \Delta(SP_{\mathcal{D}})$ and $\mathbf{y} \in \Delta(SP_{\mathcal{A}})$ takes the particularly simple form $E_{(\mathbf{x},\mathbf{y})}(u_1) = \mathbf{x}^T \cdot A \cdot \mathbf{y}$ (cf. Section 4.1.2.1 for more details on benefits on the zero-sum game assumption).

In real-world applications, actions may not work out as expected or come with intrinsically probabilistic consequences, as **Req4** expressly states. It is not difficult to generalize the mapping u_i from targeting $\mathbb{R}$ into more general loss descriptions, such as probability distributions [2]$^{\natural}$. For the sake of clarity, suppose that either by simulation or by other means of assessments (expert domain knowledge, crowd sourcing, penetration testing, etc.), decision-makers have obtained a collection of data dat_{ij} that refers to the effectiveness of defense strategy d_i against attack strategy a_j. This information may include indicators like detection events, correct incident recognition,

4 For the sake of clarity, the author changes the symbols A_1 and A_2 used in Section 4.1.1 by $SP_{\mathcal{D}}$ and $SP_{\mathcal{A}}$, respectively.

5 Minimax principle is an optimality principle for a two-person zero-sum game.

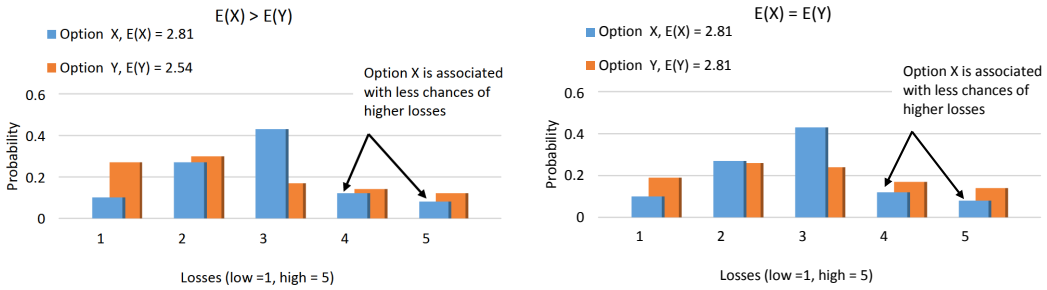

(a) different expectations (b) equal expectations

Figure 4.1: Comparison of two loss distributions; casting uncertain payoffs into expected quantities fails to capture the preferences of CI defenders

correct classification, or similar. From this data, the decision-makers can construct the payoff matrix $A = (A_{ij})_{(i,j=1)}^{(n,m)}$ by specifying probability distributions as payoffs instead of single numbers. An easy (non-parametric) choice is kernel density estimates F_{ij}, based on $data_{ij}$, which make the random payoff A_{ij} to be

$$A_{ij} \sim F_{ij}(data_{ij}).$$

However, the collection of distributions has no natural order on it such as $\mathbb{R}$, so finding a sound replacement for $\leqslant$ in the above equilibrium conditions (cf. Section 4.1.1) is nontrivial.

In classical settings of game theory, utility functions define a mapping from the action space of players into comparable payoffs, typically real numbers. In light of this, the first possibility to deal with uncertain actions is to replace each distribution-valued payoff by its first moment (i.e., average value). This approach is to a great extent compatible with the known formula for quantifying risk mostly used in standards, which is (*risk* = *consequence* $\times$ *likelihood*) [97, 99, 168].

This restores things back to classical game settings at the price of losing important features of the (comprehensive) distribution-valued payoffs, such as the riskiness of the actions in the sense of the catastrophic consequences they might have. That is, decisions made by classical game models overlook such pieces of information, and therefore **Req6** is not well accounted for. To explain this point, consider the payoff distributions X, Y depicted in Figure 4.1. In Figure 4.1a, casting the uncertain payoffs into scalar expected quantities leads to the decision that Y is the preferable option since the loss expectation of Y is less than that of X. The second example depicted in Figure 4.1b shows that both options are equally preferable as the expectations of X and Y are equal. However, it is clear in both examples of Figure 4.1 that X ensures less chances of higher losses than Y, and a decision-maker involved in CI protection would, therefore, prefer X to Y.

To preserve all the features provided in probability distributions, there is a need to play games directly with the distribution-valued payoffs rather than having to convert them into "representative" real numbers. Stochastic orders [191] are ordering relations between probability distributions, which in particular allow comparing uncertain or vague objects. The "standard stochastic order" is defined by putting two random variables X, Y in order $X \leqslant_{st} Y$ if and only if $\Pr(X > t) \leqslant \Pr(Y > t)$ for all $t \in (-\infty, \infty)$. It can be shown that this ordering relation holds if $E(u(X)) \leqslant E(u(Y))$ for all nondecreasing functions u, which connects the stochastic order naturally to the aforementioned utility

functions. However, less obvious is the fact that the standard stochastic order is not generally total, i.e., the condition on the expectations may not apply in any direction for specific pairs of random variables (the same goes for many other orders too, though for different reasons) [2, 8]♮.

Nevertheless, there is a stochastic ordering $\preceq$ for security applications designed in [168] to model the preferences of CI defenders. The standard stochastic order is not demanded on the entire support, but only on the tails of the distributions. This requirement is also given for other orders, defining an entire class of *tail orders*. Most importantly for game theory, the $\preceq$ order is *total* on the set $\mathcal{F}$ of probability distributions that satisfy the following regularity conditions: a random variable X, respectively its distribution function F_X, is contained in the family $\mathcal{F}$ if and only if conditions R1, R2, and R3 hold [175]:

R1: $X \geqslant 1$

R2: X is bounded from above

R3: the distribution F_X has a continuous density function with respect to either the counting or Lebesgue measure and is piecewise polynomial over a finite partition of a compact interval $\subset [1, \infty)$.

Conditions R1 and R2 basically restrict X to have a compactly supported measure, which, for finite games, can be assumed without loss of generality (as one can just shift all loss functions u_i into the range $[1, \infty)$ without strategically altering any equilibria). Condition R3 is a technical one that assures, among others, efficient (algorithmic) decidability of the order. (see [168, 175] for respective proofs).

In brief, let the random variables X and Y, represent payoffs in the matrix structure, and assume that both fulfill the aforementioned conditions of having probability distributions that (i) are supported on a compact set $[a, b] \subset [1, \infty)$ and (ii) have piecewise polynomial densities over a finite partition of $[a, b]$. In this case, X and Y can be uniquely represented as hyperreal numbers[6] using their moment sequences $(m_X(k))_{k\in\mathbb{N}}$ and $(m_Y(k))_{k\in\mathbb{N}}$ where $m_X(k) = E(X^k)$ and $m_Y(k) = E(Y^k)$ are the k-th moments of X and Y, respectively. And the preference relation $\preceq$ is defined based on which moment sequence will eventually dominate starting at some point, as follows:

$$X \preceq Y \iff \exists K \in \mathbb{N} : \forall k \geqslant K : m_X(k) \leqslant m_Y(k)$$

Under this embedding of distributions into ${}^*\mathbb{R}$, we can play the game "as usual", only bearing in mind that the gameplay itself is now over a new algebraic structure. Things are, however, greatly simplified in the sense that there is no need to deal with hyperreal arithmetic, if the two distributions are categorical[7] as they can be compared by looking

6 Hyperreal space (${}^*\mathbb{R}$) is an extension of standard real numbers $\mathbb{R}$ that includes additionally infinite and infinitesimal quantities. Hyperreal numbers are also known as nonstandard reals and form a totally ordered field. The relation between $\mathbb{R}$ and ${}^*\mathbb{R}$ is defined using the transfer principle, which states that any theorem (e. g., the existence of equilibria) that is true for real numbers remains valid if it is extended to the hyperreal numbers.

7 In the context of risk management, it is recommended to assess risks based on categorical scales. Here, both the potential outcomes and the probabilities are classified according to a fixed number of predefined categories (e.g., "low" "medium" and "high"). These categories can be defined differently depending on the security objective and have different semantics [174].

at their tails [175]. Specifically, the tail order is equivalent to a *lexicographic* comparison of the probability masses of two categorical distributions taken from right to left.

The only remarkable characteristic of ${}^*\mathbb{R}$ is that its elements form a totally ordered field. Hence, the totality of the stochastic order $\preceq$ on the set $\mathcal{F}$ can be directly transferred from the hyperreal space of the representative moment sequences. In light of this, it is possible to reconstruct an entire theory of games based on (any) total stochastic order, such as $\preceq$ relation [167]. Since this works on probability distributions, one can safely let the loss functions u_i come up "distribution-valued". At the same time, all the definitions from before remain intact (only having the numeric $\leqslant$ replaced by the stochastic $\preceq$ relation).

4.1.2.1 *Uncertainty: Zero-sum and Bayesian games*

Despite the zero-sum constraint not necessarily being an accurate model for a real competition, it nevertheless is a provably correct worst-case assumption. That is, if the defender simply *assumes* the opponent to have opposite intentions than himself, whatever the opponent really does can cause only less damage than expected under the zero-sum assumption (as a consequence of the definition of equilibrium and unilateral deviation from it; cf. [18]). The zero-sum assumption is clearly pessimistic, and the defender could improve its situation upon any knowledge about the adversary's intentions. However, adversary modeling is a difficult task, and the resulting hypotheses may be unreliable. More importantly, the assurances obtained from a zero-sum assumption remain intact even under (unnoticeable) changes of the adversary's mind regarding its attack aims. This adds some "robustness" to the model against differently incentivized adversaries, as long as they all have the same action space. If a set of "plausible and likely" different kinds of adversaries can be identified, then Bayesian games [60] can be used as an improvement [2]$^\natural$. For each type of attacker, the game would need an accurate payoff structure (imposing considerable modeling efforts). In the absence of any such reliable knowledge, zero-sum games remain robust and easier to use alternative models. This simpler model only requires knowledge of own investment (equivalently own losses if an asset were stolen, damaged, or manipulated), which is a much more reliable piece of information than anything that can be presumed about the attacker.

4.1.2.2 *Solution Concept: (Multi-goal) security strategy*

As explained in Section 4.1.1 and Section 4.1.2, besides classical games, there is a new class of distribution-valued games. Concisely, the two classes differ in their utility/loss functions for the players [2]$^\natural$:

Classical game: $u_i : A_i \times A_{-i} \to (\mathbb{R}, \leqslant)$
Distribution-valued game: $u_i : A_i \times A_{-i} \to (\mathcal{F}, \preceq)$

In classical models, the traditional expected utility maximization (or its logical equivalent of expected loss minimization) is applied, in which the players are indifferent between choices with equal expected payoffs (losses) even if one choice is riskier. In contrast, distribution-valued games are built based on payoffs being random variables represented by their entire probability distributions, rather than just real numbers (i.e., averages). They aim to integrate the uncertainty into the decision-making process and

hence to model the attitude and sensitivity of the players towards extreme risks, thereby fulfilling **Req6**.

Solving those games deliver security strategies that assure the best behaviors of the defender under uncertainty of the attacker [171]. The computation of security strategies involves the computation of Nash equilibria in a classical game and lexicographic Nash equilibria in a distribution-valued one. In practice, a distribution-valued game delivers a lexicographic Nash equilibrium, where a deviation will indirectly cause losses for the deviating player in regards of a more important payoff dimension [175]. That is, if the defender deviates from the lexicographic Nash equilibrium, extreme events become more probable. Further explanation of the difference between lexicographic and conventional Nash equilibria is provided by the lexicographic paradox described in Section 11.1. The "expected loss" in both games is defined in the same way as $E(u_i) = (\mathbf{x}^*)^T A \mathbf{y}^*$, where $\mathbf{x}^* \in \Delta(SP_{\mathcal{D}})$ is the defender's best behavior (i. e., the security security strategy), and $\mathbf{y}^* \in \Delta(SP_{\mathcal{A}})$ is the worst-case adversarial mixed strategy in the game. For distribution-valued games, A is a matrix of distribution functions, so that $E(u_i)$ comes up as another distribution function (soundly defined since the bilinear form directly boils down to the law of total probability, assuming that the players take actions independently).

The benefit of tail orders lies in their effect of doing the optimization by "shifting mass" rather than optimizing a single statistic. To see this, consider a classical game over a numeric $\leqslant$-order, as opposed to playing a distribution-valued game over the stochastic $\preceq$-order. In both games, if the action choices are randomized, let the optimal outcome (loss) be a random variable $X_{\text{classical}}$, resp. X_{dist}. The classical game optimizes $E(X_{\text{classical}})$ by numerically minimizing it. The distribution-valued game optimizes X_{dist} directly by minimizing the tail masses. The difference lies in the "ignorance" of other characteristics than the expectation in the classical case. Knowing that $E(X_{\text{classical}})$ is minimal does not tell anything about the variance or mass in the tails of the distribution of $X_{\text{classical}}$. Hence, extreme losses may still be undesirably likely though the average loss has been minimized. Stochastic tail orders have the appeal of minimizing the likelihood of extreme losses (and hence disappointments), at the tradeoff of having perhaps a higher average loss. For applications in security, this amounts to a pessimistic view on the worst that can happen, leaving the "average case" as a matter for the (standard) business continuity management.

In most practical domains, the defender seeks to optimize multiple objectives simultaneously (see **Req3**). For example, the operator of a CI system needs to improve the safety and security inside the system by increasing the inspection and monitoring activities. However, employees generally tend to prefer less monitoring and more freedom in their workplaces (detailed definitions of similar decision objectives follow in Part II). Thus, the optimal decision needs to be made in the presence of trade-offs between multiple (negatively correlated) objectives. Figure 4.2 shows a multiobjective zero-sum game model, in which each objective k is represented as a distinct (distribution-valued) payoff matrix $A^{(k)} \in \mathcal{F}^{n \times m}$.

In multiobjective security games, the solution concept is, therefore, a multi-goal (or Pareto-efficient) security strategy, and the ordering is applied per coordinate [120]. Pareto-optimality means that any unilateral deviation from the equilibrium will result in a degeneration of at least one objective for the deviating player. Let $g \geqslant 1$ be the number of objectives of a multiobjective zero-sum security game, in which player 1 (the defender)

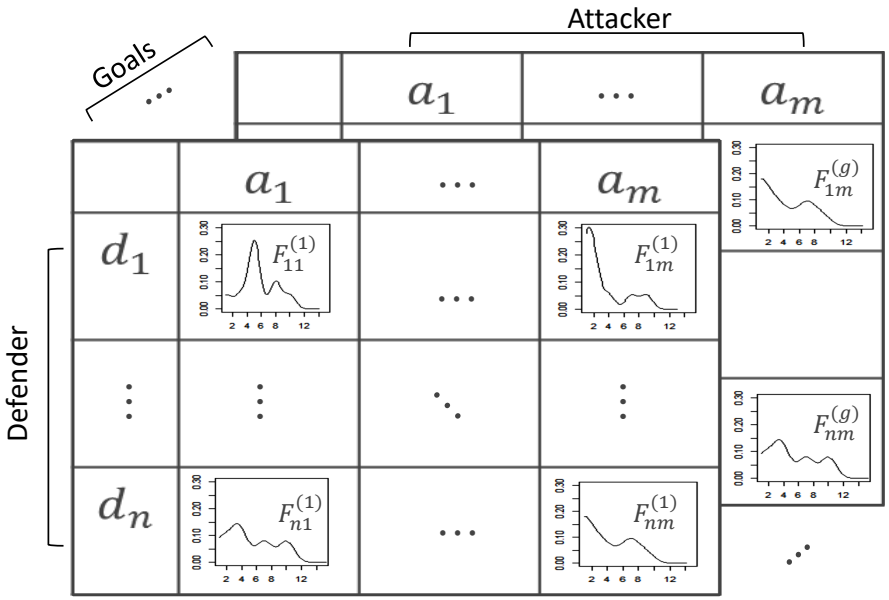

Figure 4.2: Multiobjective zero-sum security management game model using distribution-valued payoffs

owns g distinct payoff functions $u_{1,1}, \ldots, u_{1,g}$. The function $u_{1,k}$ denotes the payoff of player 1 on objective $k \in \{1, \ldots, g\}$. Algorithmically, Pareto-efficient security strategies in multiobjective games can be found through scalarizing these games into corresponding single-objective games [120]. That scalarization is nothing else than a weighted sum of all objective functions, where the weights can be set to reflect priorities of each objective, under the sole constraint of the weights to be all strictly positive. For a zero weight, one can simply exclude the respective goal from the analysis completely. Observe the neat side-effect here: the scalarization induces a set of variables for theoretical reasons, yet these variables have a perfectly meaningful practical use in the specification of the importance of each goal. This is an independent benefit of the particular method applied here to compute multi-objective optimal defense strategies. Having defined the objective weights satisfying the condition that $\sum_{k=1}^{g} w_k = 1$, the payoff function of player 1 after scalarization is defined as follows:

$$\text{minimize} \longrightarrow u_1 = w_1 \cdot u_{1,1} + \ldots + w_g \cdot u_{1,g}$$

The interested reader is referred to [120] for further details on the definition of other payoff functions and the whole transformation process of any multiobjective game. Section 4.2 explains thoroughly how this new class of game models can be exploited to support security management processes in CI systems.

4.2 A METHODOLOGICAL APPROACH FOR SECURITY MANAGEMENT

The core of security management involves a decision-making process, in which an involved decision-maker assesses possible choices and alternatives towards finding optimal configurations and rules. Towards streamlining security management operations, the approach presented in this thesis breaks down into six smaller and manageable steps. They are context establishment, identification of strategies, identification of goals,

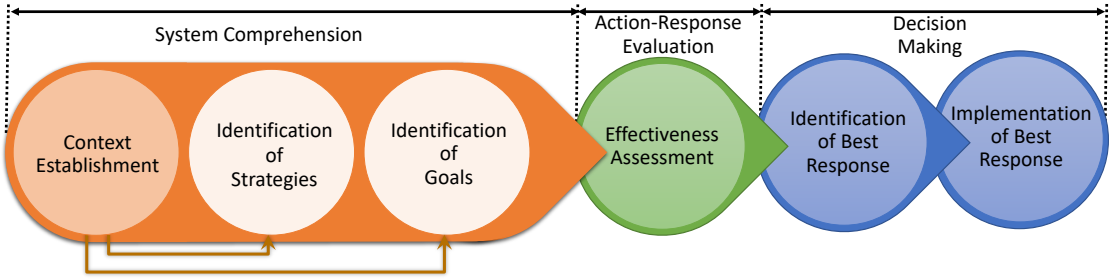

Figure 4.3: Methodological approach for risk-based security management

effectiveness assessment, identification of best response, and implementation of best response. This approach ensures a systematic workflow and a seamless integration between different techniques and principles such as risk management, decision theory, game theory, among others. Figure 4.3 shows that those six steps can be grouped into three successive phases as follows:

1) System Comprehension: This phase seeks answers to the following questions:
   - what is the context of our analysis?;
   - what are the action sets available to the involved actors given the identified context?; and
   - what are the objectives of our analysis given the identified context?
2) Action-Response Evaluation: This phase relies on the output of the former phase to assess the outcomes of the different actions with respect to the identified objectives under the current system settings.
3) Decision Making: This phase aims to identify the best action to be implemented in the future based on the assessment results of the identified actions and preferences of the involved decision-makers.

A schematic representation of the approach is depicted in Figure 4.4, while the respective steps are described in detail in Section 4.2.1 up to Section 4.2.6.

4.2.1 *Context establishment*

The first step aims at understanding the system and the environment of interest. This involves, just to name a few:

- identifying the boundaries of the environment and hence the overall scope of the security management process;
- identifying the different agents (individuals or aggregated entities) involved in (or has an effect on) making decisions (i. e., the defender and potential attacker);
- identifying the different functions, units, processes and resources relevant to the system under investigation and the connections among them;
- identifying possible exposures to risks using techniques such as vulnerability assessment or organizational architecture analysis; and

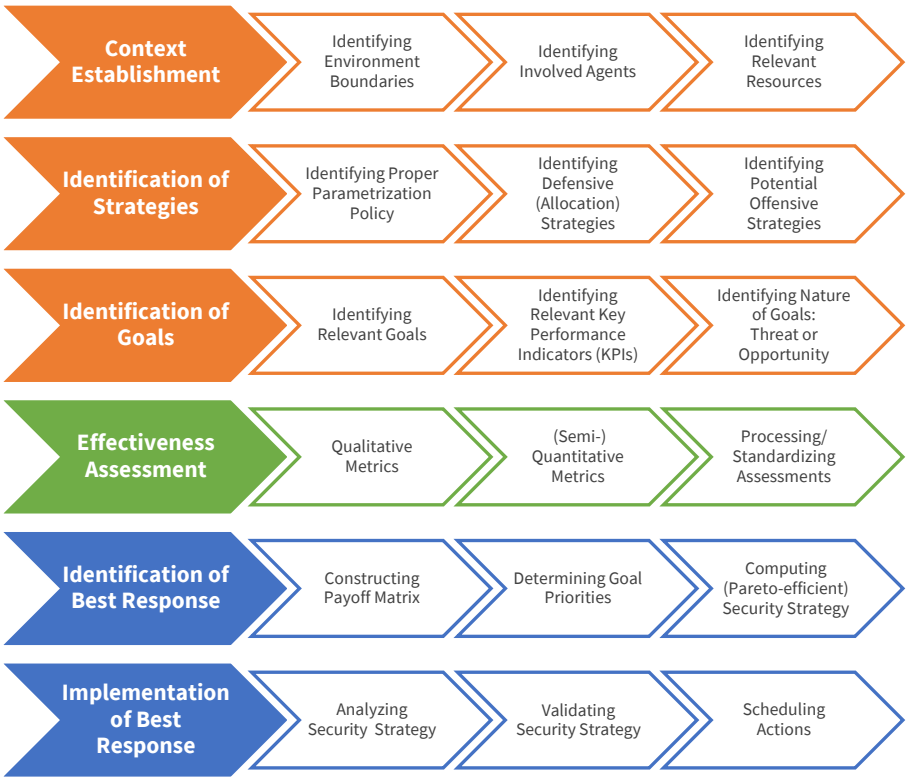

Figure 4.4: Schematic representation of the different steps of security management approach (adapted from [10])

- identifying a potential target component(s) or area(s) that matters most to the system of interest.

Note that "context establishment" is a prerequisite step for all other steps[8], even within the first phase, "system comprehension", as illustrated in Figure 4.3. In this step, a comprehensive system analysis has to be performed. This practice usually dictates the involvement of many experts with different domains of expertise. The knowledge collaboratively acquired from several experts can be further vetted to determine its accuracy and usefulness. Therefore, incorporating the expertise of several experts has positive effects concerning (i) knowledge completeness, as well as (ii) quality and reliability of the acquired knowledge.

4.2.2 *Identification of strategies*

This step involves the identification of possible measures, configurations, layouts, and operational patterns available to agents identified in the course of the preceding step (cf. context establishment in Section 4.2.1). Therefore, strategies refer to what can be done by involved agents to accomplish their intended goals. At this step, a proper parameterization policy can be defined, if possible. Strategy parameterization refers to the process of describing strategies in terms of independent parameters reflecting (tunable) characteristics and properties of each strategy. Parameterization facilitates developing a shared understanding of the broad spectrum of potential strategies and also enables resource allocation through adjusting parameter values. Distinct parameter settings will then correspond to distinct strategies. Although the parameterization process is mostly driven by respective application and use cases, it represents the basis for threat intelligence sharing in an automated and standardized way [2]♮. Broadly speaking, this step is a purely technical issue and based on domain knowledge about the infrastructure, enterprise, premises, or environment at hand [111][10]♮. In [13]♮, there are several examples of defense/attack strategies in APT-like intrusion avoidance games. An intruder can gain unauthorized access to high-security zones in a critical facility using strategies such as forged/stolen ID card, forged/stolen third party ID card, or a valid visitor ID cards with tailgating authorized employees to enter a secured zone. A defender, in turn, has several options to thwart an intrusion, including awareness training of employees, ID check enforcement, visitor and third party back check, as well as reporting stolen ID card.

4.2.3 *Identification of goals*

This step aims at identifying the different operational, legal, organizational, or technical goals and their relevant Key Performance Indicators (KPIs)[9] [7]♮. Utilizing optimization techniques, decision-makers seek to find the best configurations or choice rule that can keep the balance between all identified goals. Security management problems can involve several conflicting (in the sense of a negative correlation) goals that need to be optimized simultaneously, including costs, caused damage, privacy issues, or

8 Steps "identification of strategies" and "identification of goals" can be performed in any arbitrary order.
9 Key performance indicators measure the degree of achievement or fulfillment of identified goals and interests.

employee comfort and productivity [1, 5, 9]♮. Therefore, it is necessary to identify which goals must be maximized (i. e., perceived as opportunities) and which to be minimized (i. e., perceived as threats). This complies with risk management processes that seek to optimize both sides of risk, i. e., maximizing opportunities and minimizing threats.

4.2.4 *Effectiveness assessment*

In this step, the effectiveness or performance of each known configuration (i.e., each strategy) identified in Section 4.2.2 is determined with regards to all goals identified in Section 4.2.3. Given the wide variety of possible goals, there are different assessment methodologies of the various scenarios. Therefore, action-response models have to be defined to understand the possible consequences (i. e., losses) of the different actions under the current system configurations. Such models leverage different qualitative, quantitative, or semi-quantitative assessment techniques such as mathematical models, simulation, eliciting expert judgments, or using historical data [7]♮.

Action-response models can involve the use of a wide variety of observed and statistical data. That is, significant uncertainty and variability are associated with such data and can have a serious impact on the assessment process. In this case, single-point estimates can fail to communicate comprehensive effectiveness assessments to decision-makers. In the context of CI protection, for example, decision-makers are concerned with the mitigation of extreme risk. Hence, they prefer security measures that prevent (or minimize) the occurrence of a catastrophic loss.

Therefore, throughout this thesis, each action scenario undergoes several assessment iterations to address inherent variability and uncertainty of input data as well as dynamic system responses. Subsequently, the outcomes of all iterations are merged using several techniques (e. g., frequency histogram, kernel density estimation, or the maximum entropy method) to generate the final assessment distribution function [7]♮. Distribution-valued assessments provide essential information for better-informed decision making under uncertainty.

Simulation is one of the standard assessment practices that deliver comprehensive goal assessments. However, simulation is not feasible for all goals of interest, especially when the response dynamics, such as employee satisfaction and social response, are unknown. In such scenarios, the assessment process can be performed with the aid of "soft" indicators like the degree to which end-users appreciate some measures (e. g., surveillance) or feel uncomfortable upon some activities. In this regard, empirical data coming from traditional surveys, expert and stakeholder opinions, or historical and statistical data may be employed to determine the values of such soft indicators. For example, end-users may be asked how they feel upon having installed cameras somewhere, or whether or not they would be willing to have their own devices become part of surveillance infrastructure. Even if a user consents, a surveillance device (e.g., a mobile device) may not always be connected, may be out of power, among others, which adds an intrinsic element of randomness to the outcome in every scenario [111][10]♮.

In either case, all assessments obtained in this step are combined into a categorical (or continuous) probability distribution to avoid any loss of relevant information. In this way, conflict resolution and consensus problems can be avoided, which arise when collected data (e. g., expert opinions) is diverging and has to be aggregated into a single

representative value. Moreover, these distributions must be constructed/processed under the following constraints:

1. All assessments are made on the same scale. This makes payoffs comparable as well as standardizes the taxonomies in which outcomes of actions are expressed. This constraint is required for the multiobjective optimization to work. Numeric indicators are thus discretized onto a common categorical scale that all categorical indicators use as well;

2. The data source is reliable with regards to the intended goal assessment.

4.2.5 *Identification of best response (strategies)*

This step involves basically the construction of a security management game based on the model $G = \langle\{\mathcal{D}, \mathcal{A}\}, \{SP_{\mathcal{D}}, SP_{\mathcal{A}}\}, (A^{(k)} \in \mathcal{F}^{n\times m})_{k\in\{1,\ldots,g\}}, \preceq\rangle$ introduced in Section 4.1.2. In fact, the four preceding steps play a crucial role in constructing the game. More precisely, the step "context establishment" defines the game players $\{\mathcal{D}, \mathcal{A}\}$ and other context-related parameters. While the step "identification of strategies" determines the players' action sets $\{SP_{\mathcal{D}}, SP_{\mathcal{A}}\}$, the step "identification of goals" determine the $g \geqslant 1$ game objectives. Then, the assessment results delivered by the step "effectiveness assessment" should be leveraged to construct the distribution-valued payoff matrices $(A^{(k)} \in \mathcal{F}^{n\times m})_{k\in\{1,\ldots,g\}}$. In the presence of multiple goals, it is necessary to define the weight variables, which reflect the importance of each goal (cf. Pareto-efficient security strategy in Section 4.1.2.2).

Technically, the overall method to compute security strategies of multiobjective security management games is based on [165, 171]: it treats the defender $\mathcal{D}$ as "player 0", opposing a set of g opponents, each of which corresponds to a different objective (for the defender). In the competition, the defender then seeks the simultaneously best behavior against all opponents, each of which acts independently of the others and where the i-th opponent seeks to minimize player 0's payoff in the respective i-th objective. This is a so-called "one-against-all" competition, for which fictitious play is known to converge [187]. Fictitious play is a self-learning algorithm in which each player chooses his next best action based on a recorded history of all choices made by other players so far. It requires selecting maximum or minimum from a finite set of actions, which is easy once the ordering relation is defined over the payoff space. The particular setup of decomposing the physical (single) opponent into a set of hypothetical and independently acting competitors makes the resulting equilibrium in the one-against-all game a Pareto equilibrium in the original game, which pessimistically bounds the payoffs for player zero (here the defender) [162]. To solve a multiobjective zero-sum security management game G, a generalized version of the fictitious play (FP) algorithm has been implemented in R package [12][♮] (more details on the FP algorithm are included in Appendix A.1).

It is worth noting that the $\preceq$-ordering includes the $\leqslant$-order between two real values as a special case of comparing Bernoulli distributions: let $a, b \in \mathbb{R}$ be given and, without loss of generality, assume $a, b > 1$ (otherwise, we may just shift the values accordingly without changing their relative order). Choose $M > \max\{a, b\}$ so that $a/M, b/M \in (0, 1)$, and define Bernoulli distributions $(1 - a/M, a/M)$ and $(1 - b/M, b/M)$ for two random variables X, Y. Then, we have $X \preceq Y$ if and only if $a \leqslant b$ (as follows from

directly by [168, Theorem 3]). This procedure works analogously for an entire game. Simply choose M as a common scale factor to map all entries in the game matrix A into the interval $(0, 1)$ and define per element Bernoulli distributions from it to mimic a $\leqslant$-comparison via the stochastic order $\preceq$ [2]♮. Consequently, the generalized version of the fictitious play (FP) algorithm can be used to solve multiobjective classical games, too.

It worth mentioning that identification of the best response can involve constructing several distribution-valued games where the optimal decision will be reached after playing a chain of security games.

4.2.6 *Implementation of best response*

Having found a security strategy, such as optimal surveillance routes and frequencies, the daily business requires to implement the static precautions, e. g., building the surveillance system according to its optimal layout and configuration, and adhering to random reconfigurations and daily operation [111][10]♮. In some cases, the assessment process (cf. Section 4.2.4) applied to assess the effectiveness of involved players' actions, could be leveraged to analyse and validate the efficiency and feasibility of the obtained security strategies. For example, if the assessment is conducted using simulation, the game equilibrium strategy can be similarly implemented in the developed simulation environment and then contrasted with results obtained in early steps. Towards a practical implementation of obtained security strategies, remember that all we require is a certain frequency of actions to happen over repetitions of the game. For example, let us fix a time unit, say T hours, then if the equilibrium prescribes action a_1 to happen with probability p_1, this means an average of $p_1 \cdot T$ actions during a day. Taking the pauses between repetitions of action a_1 as exponentially distributed with rate parameter $1/p_1$, it is a simple matter of drawing exponentially distributed pause times to get the time when action a_1 is to be launched next. In turn, the number of actions is a Poisson distributed variable with the same rate parameter, as desired to play the equilibrium. For the other strategies, the procedure works analogously, and ultimately gives a (randomized) schedule of actions that assures the optimal frequencies as prescribed by the equilibrium [9]♮. Other possible implementations of obtained security strategies are described in Part II and Part III.

4.3 TECHNIQUES AND METHODS

The approach presented in Section 4.2 involves the application of several methods and systematic techniques in each phase to deliver intended results to other dependent phases and steps, as sketched in Figure 4.5.

1. System comprehension phase involves techniques to understand the organization's environment and scope of the analysis, such as:
    - Organization architecture analysis that delivers a detailed technical description of involved systems, tasks, activities, operations, and participating players as well as interconnections and information exchange patterns required to satisfy specific operational needs.

- Ethnographic studies that explore the social aspects within organizations and how employees behave in specific situations and to which extent security policies and methods of operations are adhered to in real life.
- Business process analysis that is broadly applied to capture interconnections between involved processes and players towards identifying opportunities for improving efficiency and effectiveness of business operations.
- Vulnerability analysis that is a systematic process of identifying and classifying weaknesses which create possible points of security compromise in an organization's environment, thereby providing knowledge and basis necessary for identifying possible exposures to risks and mitigation actions.

2. Effectiveness assessment phase can benefit from a wide range of consequences analysis techniques, such as structured mathematical models, simulation, eliciting expert opinions, as well as historical and statistical data.

3. Decision-making phase involves methodologies such as:
   - Game theory that provides a sound mathematical foundation to perfectly model the competitive situation between an organization's security team and potential adversaries towards inferring the optimal strategic decisions thereof.
   - Decision theory that offers several evaluation techniques to decide on the preference among a set of (uncertain) options.
   - Minimax and (lexicographic) Nash equilibrium that are the rules for predicting how the different players will play or adopt their strategies towards maximizing their own benefits.
   - Purification approach that discusses how equilibrium decisions can be implemented in practice, especially when the equilibrium is mixed.

Besides the techniques and principles as mentioned above, the methodological approach presented in this thesis is complaint with the generic risk management methodologies like the one defined by ISO 31000 standard [99] and its other closely aligned standards such as ISO/IEC 27005 [98]. Furthermore, it can be easily integrated into other standard frameworks such as NIST 800-37 [58] as well. The three phases of the presented approach correspond to the core steps of any standard risk management processes, namely establish the context, risk assessment, and risk treatment (cf. Figure 4.5).

In a very similar way, Figure 4.5 shows the compatibility with the situational awareness process that aims to glue past, present, and future together to enhance the security posture of CI organizations. The presented approach defines an integrated decision-making process that utilizes past knowledge and experience about the system dynamics to identify a set of technically possible offensive and defensive actions. This knowledge paves the way for constructing appropriate action-response models to assess the outcomes of these different actions and behaviors under the current system configurations in order to infer the action with the best response that has to be implemented in the future towards minimizing the risk of interest [7]♮.

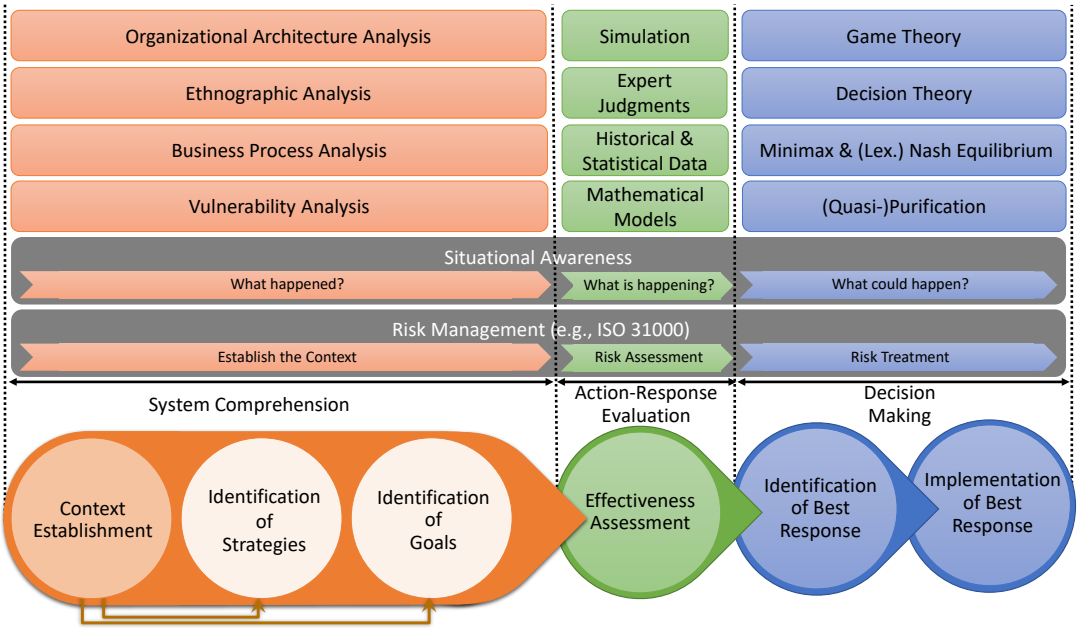

Figure 4.5: Methodological approach for security management and its closely related techniques and methods

4.4 SUMMARY

This chapter describes in detail the methodological approach that is used throughout this thesis to address security management problems. The approach is divided into manageable steps to support defenders of CIs to smoothly identify and assess possible security choices towards finding proper security strategies. As a decision-making mechanism, a generalized game model for security management is introduced, which address the different challenges of security management problems identified in Section 1.2. Unlike classical game models, this type of games accepts randomness as an inherent part of the payoffs. Furthermore, it integrates the specific risk attitude of CI's defenders into the decision-making process itself. In Part II and Part III, this thesis thoroughly explains the application of the presented approach to security problems in physical and cybersecurity domains, respectively.

Part II

PHYSICAL SECURITY MANAGEMENT

This part shows the application of the methodological approach presented in Part I to physical security problems.

Publication references*:

- Ali **Alshawish**, Mohamed Amine Abid, and Hermann de Meer. "Quasi-purification of mixed game strategies: Sub-optimality of equilibria in security games." *Computers & Security* 87:101575, 2019.
- Ali **Alshawish**, Mohamed Amine Abid, Hermann de Meer, Stefan Schauer, Sandra König, Antonios Gouglidis, and David Hutchison. "G-DPS: A game-theoretical decision-making framework for physical surveillance games." In *Game Theory for Security and Risk Management*, pp. 129-156. Birkhäuser, Cham, 2018.
- Ali **Alshawish**, Mohamed Amine Abid, and Hermann de Meer. "Game-Theoretic Optimization for Physical Surveillance of Critical Infrastructures: A Case Study." In *Game Theory for Security and Risk Management*, pp. 353-389. Birkhäuser, Cham, 2018.
- Antonios Gouglidis, Benjamin Green, David Hutchison, Ali **Alshawish**, and Hermann de Meer. "Surveillance and security: protecting electricity utilities and other critical infrastructures." *Energy Informatics* 1:15, 2018.
- Ali **Alshawish**, Mohamed Amine Abid, Stefan Rass, and Hermann de Meer. "Playing a Multi-objective Spot-checking Game in Public Transportation Systems." In *Proceedings of the 4th Workshop on Security in Highly Connected IT Systems*, pp. 31-36. ACM, 2017.
- Ali **Alshawish**, Stefan Rass, and Hermann de Meer. "A Game-theoretical Decision-making Framework for Physical Surveillance Games." In: *Workshop on Novel Approaches in Risk and Security Management for Critical Infrastructures.* Vienna, Austria, Austrian Institute of Technology AIT, 2017.
- Stefan Rass, Ali **Alshawish**, Mohamed Amine Abid, Stefan Schauer, Quanyan Zhu, and Hermann De Meer. "Physical intrusion games–optimizing surveillance by simulation and game theory." *IEEE Access* 5:16904577, 2017.
- Ali **Alshawish**, Mohamed Amine Abid, Zhiyuan Sui, Hermann de Meer, Antonios Gouglidis, and Stefan Rass. "HyRiM Deliverable 4.3: How to Enhance Perimeter Security Using New Surveillance Technologies." 2017.

* The research work (including some ideas and figures) from these papers, which is included in Part II, was carried out and documented by the author of this thesis.

5 PHYSICAL SURVEILLANCE GAMES

The content of Chapter 5 is based on the research work published in [5, 8, 10, 11]♮.

5.1 INTRODUCTION

CI systems can be characterized by closed structural environments (e.g., power plant, refinery, and airports) or open structural environments (e.g., public transit systems and nation's borders). In either case, disruptions of these systems could have widespread and devastating consequences on the national economy and public health. Therefore, objectives like safety, security, and service continuity are of utmost importance to these systems. Although there are several advanced access control techniques, which can be used to secure facilities of interest, visual monitoring and on-site observation are still indispensable practices to ensure persistent surveillance in such environments. However, covering a moderate-sized environment requires a substantial number of static cameras [5, 10]♮. This produces heavy monitoring activities for security personnel behind monitoring screens and can lead to poor efficiency due to potential fatigue [30].

The employment of mobile surveillance devices for airports and train stations can help to detect abnormal behaviors and identify potential terrorist threats [143]. This chapter, therefore, addresses mainly scenarios in which mobile agents (e. g., security guards, law enforcement officers, and police robots) can be deployed in the environment for surveillance applications. In such scenarios, while potential adversaries seek to cause maximum damage to a target infrastructure, the defenders or first responders, to the contrary, seek optimal resource allocation in an attempt to thwart adversarial plans. Mostly, security resources (mobile agents) are not adequate to track all targets at once. Thus, such resources have to be strategically assigned to maximize the benefits for the system's defenders [5, 10]♮. This problem has already been reflected in several game-theoretical models, as explained in Section 2.3.1. Existing models, however, assume a deterministic outcome of the gameplay, while such a decision-making process involves several types of uncertainties. For instance, even if a security guard and an adversary share the same site, there is a probability that the guard misses the adversary inducing randomness in the players' outcomes. Modeling this randomness based on domain knowledge usually culminates in an expected payoff for the players (e. g., a success rate for the patroller, average damage for the attacker). But, this is basically a reduction of information from the full-fledged probabilistic model (a distribution function) back to a real value [5, 10]♮.

Based on the game model presented in Section 4.1.2, *physical surveillance games* (or equivalently "physical intrusion avoidance games") can be understood as security management games that model the interaction between two players (i. e., defenders/-first responders/security personnel and potential attackers/intruders/criminals) each equipped with a finite action set (i. e., strategies). Additionally, the chance is deemed as a "hypothetical" third player that induces randomness in the real player's outcome. The security strategies of such games will deliver the defenders with optimal surveillance

policies and strategic allocation of the available resources within the environment of interest [5, 10]♮.

Regarding the general setting of physical surveillance games, they can consider large environments, e. g., an industrial complex of a utility provider, consisting of several areas/sectors/working lines of different importance and having a number of security guards, who are patrolling these areas to detect potential violations. Broadly speaking, physical surveillance games have several important real-life manifestations such as physical border patrolling, scheduling random security checkpoints, mobile robot path planning, public transit security, and fare enforcement planning [9]♮, among others. Finally, physical surveillance activities are not only useful to detect malicious intrusions but also to ensure that local safety measures such as smoke detectors and fire extinguishers are functioning properly. Such devices are useless if they are not functioning when they are needed.

5.2 OVERVIEW OF SURVEILLANCE

Surveillance is commonly described as the careful watching of objects, persons, and areas due to a crime that has happened or is expected to happen. Surveillance has been explained as "*the systematic investigation or monitoring of the actions or communications of one or more persons. Its primary purpose is generally to collect information about them, their activities, or their associates. There may be a secondary intention to deter a whole population from undertaking some kinds of activity*" [47]. The deployment layout of sensors and surveillance entities represents a vital step towards achieving the predefined goals of any surveillance system.

5.2.1 *Categorization*

Currently, a wide range of surveillance technologies is used in order to provide end-users with different levels of functionality. To identify the various surveillance technologies, a systematic literature review of surveillance technologies for the protection of CIs is conducted in [11]♮. In general, the analysis resulted in the identification of several categories of surveillance technologies, including:

- Biometrics are concerned with automated methods in order to identify or recognize the identity of a living person based on his/her physiological or behavioral characteristics [132, 222, 223].
- Visual surveillance technologies are characterized by their wide variety of technologies, e.g., video, imaging scanners, photography, satellites, or Unmanned Aerial Vehicles (UAVs) [38].
- Dataveillance technologies are mostly utilized in the context of data systems that collect personal information. This information could be used subsequently in the investigation or monitoring of the actions or communications of one or more persons [27].
- Communication surveillance is used to monitor, intercept, collect, preserve and retain information that has been communicated, relayed, or generated over communication networks to a group of recipients by a third party [176].

- Location tracking surveillance technologies are used to monitor position and movements, e.g., proximity sensing, scene analysis, and triangulation [86].
- Ubiquitous surveillance is related to the unilateral gathering of data on people in their living environment through the use of various embedded sensors [147][11]♮.

5.2.2 *Limitations*

Traditionally, the physical security of organizations has been mainly shaped by the castle (fortress) protection model, which aims at building a hard shell around a presumably trusted area [13]♮. This area encompasses different valuable assets varying from people, hardware, and software to data and information resources. Therefore, security controls have been mainly deployed and mounted at the outer boundaries of the facility of interest (see security perimeter described in Chapter 3). Due to the inflexibility and fixed installation of these systems, their deterrent effect will be considerably less [8]♮. Hence, an intruder's chance of successfully circumventing security controls located at the perimeter is significantly higher.

The current tendency of CI organizations to extend beyond their conventional borders to reach other entities such as vendors, business partners, service providers, or even customers results in having different external entities within the system complex, such as temporary workers, interns, independent contractors and subcontractors, or visitors. Even if access to the sensitive industrial zones is tightly controlled at the borders, behind the borders, the freedom of movement is almost entirely ensured for ordinary organization's personnel as well as for temporary ones [8]♮. Therefore, potential adversaries can exploit such conditions to cause more loss and damage. Undetectability within the system complex will give the adversary a good opportunity to reconnaissance the target area, to gather some sensitive information, and to probably cover the tracks of ongoing attacks, too.

Nowadays, a breach of physical security remains probable due to accidents, human errors, or even targeted attacks. It is very likely that employees behave and perform inappropriately, resulting in direct breaches of an organization's security policy. For example, adversaries can exploit the fact that issued badges of terminated employees or temporary visitors are not always timely recovered before leaving the site, and the access of stolen or lost badges is similarly not revoked on time. As a consequence, perimeter-centric physical security measures such as traditional Closed Circuit Television (CCTV) systems or access control solutions that use static devices mounted at specific locations are not adequate to detect and prevent such potential intruders [151][5, 10]♮.

As discussed in Section 5.1, full coverage of large-scale areas is very challenging. In general, surveillance coverage is strictly limited by the number of available resources such as sensors, processing devices, or human resources. Therefore, available surveillance resources have to be strategically allocated to achieve envisaged goals. Furthermore, it is highly essential to maintain situational awareness even within the system complex so that potential intruders can remain detectable, and security managers can respond timely. Having dynamic and mobile surveillance systems (or strategies) will definitely increase a system's robustness as well as increase the attack costs and complexity. This, in turn, will give CI's defenders the advantage to stay ahead of the attackers in the respective security game [11]♮.

5.3 PHYSICAL SURVEILLANCE GAMES

5.3.1 *Challenges of physical surveillance*

A core challenge is how to allocate available resources in terms of scheduling the route and frequency of patrol inspections. In some cases, the frequency of inspections needs to correlate with the value of the asset. More precisely, if highly sensitive business assets are stored at location B, while relatively less sensitive data resides at location C, then it makes sense to check B more often than C, at frequencies proportional to the value of the respective goods. Extending the problem to a whole infrastructure, calling for all-encompassing protection, quickly induces the need for a surveillance strategy which performance can be optimized.

When thinking of physical surveillance, this can be done at different locations and at different levels of granularity (e.g., ranging from quickly inspecting to thoroughly examining an area, with the latter being more time consuming) and variable rates (e.g., hourly, every two hours, or every six hours). Intruders will, in turn, react on the surveillance patterns by allocating their efforts to places that are (currently) not under surveillance.

In general, a surveillance game is essentially a simplified version of a pursuit-evasion ("cops and robbers") game [37, 152], in which the security guard is the "cop" and the intruder is the "robber". However, the issue in real-world scenarios is that an intruder may not always be detected by the surveillance system. Hence, let us here state that there is an intrinsic likelihood of missing the intruder in every round of the game, and thus for the intruder to be able to cause a certain amount of damage in the specific area. In essence, if some zones are known to be under stronger surveillance than others, the natural reaction would be to focus intrusion efforts on spots with weakest supervision and detection mechanisms. Therefore, the overall goal is to avoid damage suffered from intrusions by managing the surveillance activities accordingly. Consequently, the performance of surveillance has to be quantified in terms of damage prevention to make surveillance activities comparable.

Quantifying the damage expected from an intrusion is usually the most challenging part in a practical application of game theory in the context of physical surveillance. Obviously, it is not always possible to define the effect of a successful intrusion as a payoff being equal to the negative value of the stolen good simply because this value may be unknown or difficult to quantify. Likewise, assigning a non-negative payoff upon thwarting an intrusion is improper, as this event may not even be noticed in practice. Often, this ends up with a purely nominal quantification of both value and probability, according to fuzzy terms like "damage is high if the intruder enters a high-security area; however, this is expected only with very low probability". For setting up a game-theoretic model to optimize the surveillance system's configuration, more reliable assessments are required. The latter is achieved by querying a maximum of available sources of information and aggregating the results.

Combining the multiplicity of potential sources usually leads to a detailed and thus difficult picture to manage risk minimization. For example, cameras may raise alarms upon detection of unusual behavior, or even classify the current image sequence in terms of criticality (e.g., if a person is showing up at some place at a time when this place is supposed to be empty, or if a car remains parked when all others left the place).

This information and its classification are by themselves subject to some errors and presented to human operators to decide upon taking action or not. Additionally, a purely static surveillance system cannot avoid having dead angles or shadowed spots, so that the static surveillance data is usually combined with "dynamic" information obtained from the security staff patrolling the premises. The immediate question here is concerned with how to do the surveillance optimally, i.e., where to place the surveillance equipment, what data to collect and how often, etc. Assuming that every such choice is among finitely many alternatives only, the issue can be rephrased as a game-theoretic problem [111][10]♮. Two models of physical surveillance games are presented in Section 5.3.2 and 5.3.3. Both models are based on previous research work published in [8, 10]♮

5.3.2 *Basic model of physical surveillance games*

It is convenient to think of the infrastructure environment as a finite undirected graph $G = (V, E)$ with V being the set of nodes corresponding to physical areas (buildings, or vehicle trips in the context of public transit systems), and E the set of edges representing connection paths among them. Without loss of generality, edges can be assumed to be without surveillance, since any path (e. g., an aisle) under surveillance can be modeled as another node in the middle of the edge. More formally, if areas A and B are connected by an aisle and that aisle is under surveillance (e.g., by a camera), then it is treated as a third place C with the graph model having the edge sequence $A - C - B$, instead of the single edge $A - B$ in which the aisle would be assumed without any protection or detection mechanism. In this view, the intruder may (randomly) walk on the graph in an attempt to reach his goal (the area with the valuable business assets) while avoiding meeting the security personnel at any node. In case the intruder is captured, it gets kicked out of the area (removed from the graph), and the gameplay starts afresh again.

Putting this in a more formal way, let a single pure strategy in the standard model be a circle in the infrastructure graph G, so that the strategy space of the surveillance person is a (not necessarily minimal) set of circles $C_1, \ldots, C_n$ that spans G. Likewise, let the attacker's action set be a set of paths $P_1, \ldots, P_m$ which, without loss of generality, all end at a specific valuable target node $v_0 \in V$. In the classical version of the pursuit-evasion game, the payoff in the game would correspond to the outcome of the detection of the intruder. In this case, the game itself becomes a simple matrix-game, whose payoffs are stochastic in the sense that the payoff matrix $A = (A_{ij})_{(i,j=1)}^{(n,m)}$ is one of the Bernoulli random variables $A_{ij} \sim \mathrm{Ber}(p_{ij})$ with the semantic that:

$$A_{ij} := \begin{cases} 0 & \text{if the intruder is missed;} \\ 1 & \text{if the intruder is caught,} \end{cases} \tag{5.1}$$

in which the parameter p_{ij} tells how possible detection of the path P_j along the tour C_i is. Packing all temporal matters and detection errors into the simulation or other assessment methodologies (as discussed in Section 4.2), it is an easy yet laborious matter of working out the specific distributions. Solutions in the sense of Nash equilibria of the resulting "non-deterministic" game can be obtained in various ways. The most obvious one is to convert the matrix of random variables into a real-valued matrix by taking the expectation per element. This results in a real-valued matrix

$B = (p_{ij})_{i,j=1}^{n,m} = (E[A_{ij}])_{i,j=1}^{n,m}$ that can be treated with the entire well-known machinery of game-theory (von Neumann's minimax theorem and linear optimization [217]).

5.3.3 *Generalized model of physical surveillance games*

The basic model sketched in Section 5.3.2 deviates from reality for exactly the reasons already mentioned in Section 5.1 and Section 5.3.1 above. In real-world surveillance systems, several practicalities and imperfections can significantly result in fluctuating detection performance of these systems. There are pieces of uncertainty that must be reflected in a good model [5, 10]♮.

To describe the uncertainty stemming from these various limitations of surveillance systems, let the payoff of a physical surveillance game not to be quantified by a single number. Instead, it is described by a set of possible outcomes that either stems from simulations, surveys, or expert interviews. In any case, a real-valued payoff matrix, similar to matrix B and based on the Bernoulli random variables from matrix A in Equation (5.1) is no longer appropriate. Moreover, there is a need to resort to a more expressive categorical distribution to avoid information loss.

Putting this in a more formal way, let us assume that $T_1, T_2, \cdots, T_{Max}$ are different types of areas tagged with their respective security demands. Accordingly, let a single pure strategy in the model be a set of frequencies $f = (f_{T_1}, f_{T_2}, \cdots, f_{T_{Max}})$ representing the number of times a security guard is performing a security check in the different security demand areas, respectively. Hence, the strategy space is the collection $f_1 \dots f_n$ of all admissible (i.e., practically reasonable and doable) frequency tuples. Accordingly, the adversary's strategy space comprises paths to the set of target security zones $Z_1 \dots Z_m$, where the adversary wants to cause some damage. Suppose that either by simulation or by other means of assessments (expert domain knowledge, crowdsourcing, penetration testing, etc.), a collection of data dat_{ij} is obtained that refers to the effectiveness of defense strategy i against attack strategy j. This information may include the aforementioned indicators like detection events, correct incident recognition, correct classification, or similar. From this data, one can construct the payoff matrix $A = (A_{ij})_{(i,j=1)}^{(n,m)}$ by specifying probability distributions as payoffs instead of single numbers. An easy (non-parametric) choice is kernel density estimates F_{ij}, based on dat_{ij}, which make the random payoff A_{ij} to be

$$A_{ij} \sim F_{ij}(dat_{ij}). \tag{5.2}$$

Note that this approach can also be described in the terms introduced in Section 5.3.2, where circles $C_1, \dots, C_n$ represent the tour of the security guard and $P_1, \dots, P_m$ represent the intruder's paths. The set of frequencies $f = (f_{T_1}, f_{T_2}, \cdots, f_{T_{Max}})$ can be translated to a sequence of areas the security guard has to check, thus corresponding to a circle C_i in the infrastructure graph. On the other hand, an intruder often has to pass several security areas before he reaches his target Z. This set of areas he has to pass can be translated to an attack path, P_j, which is a strategy in the game model (determining the random outcome distributed according to F_{ij} if the defender plays its i-th move to protect).

Following the methodological approach put forth in Section 4.2, physical surveillance games can be played directly with the distribution-valued payoffs rather than having to convert them into "representative" real numbers. Moreover, it is possible to add several

more dimensions to the gameplay optimization, such as the inconvenience caused by unwanted and too frequent security checks (since they might interrupt the current work of a person or might not be possible immediately). However, the most important benefit from directly working with the distribution is gained when the Bernoulli-distribution is replaced by a more general, categorical or even continuous, distribution model over the categorical damage scale that applies to the different indicators (e.g., detection rates, privacy infringement, comfort, etc.). The physical surveillance model introduced here provides the basis for developing the simulation environment described in Section 6.5.1 and employed to assess the performance of different physical security strategies.

5.4 ENTROPY-BASED MODEL FOR QUANTIFYING LOCATION PRIVACY

Surveillance practices, such as random spot checks, could pose a severe threat to location privacy. Having access to location traces at a particular period of time, attackers can infer sensitive attributes such as identities of employees or create movement profiles that can be used to plan further (targeted) attacks on their organizations. In general, it is difficult to capture a full understanding of individuals' privacy. In some cases, however, it is possible to define metrics to quantify privacy depending on the context and nature of the application. In this work, an entropy-based model is introduced to assess the impact of different security inspection strategies on the preservation of employees' locations. This model falls into the category of uncertainty metrics. That is, it assumes that an attacker cannot make reasonable guesses based on uncertain information. Hence, highly uncertain attacker correlates with high privacy[1] [220]. The model[2] has the following ingredients:

1. $\mathcal{Z}$ is a set of (non-overlapping) areas (or zones) composing the facility to be monitored (cf. zones defined in Table 6.1). The number of areas is denoted by $|\mathcal{Z}| = z$;
2. $\mathcal{S}$ is a set of subjects (e. g., employees) present in the facility. The number of subjects is denoted by $|\mathcal{S}| = s$;
3. $CS \subseteq \mathcal{S}$ is a set of subjects being checked;
4. $\mathcal{O}$ is a set of independent events $o_{i,j}$, each of which implies that the i^{th} subject of $\mathcal{S}$ is in the j^{th} area of $\mathcal{Z}$; and
5. $Pr(o_{i,j})$ is the occurrence probability of an event $o_{i,j}$.

5.4.1 *Static model*

The model aims to provide a metric of privacy preservation after a potential leakage of location information caused by performing spot checks on different zones. Let us

1 The presented model focuses mainly on measuring the average uncertainty associated with predicting the employees' locations due to potential information leakage. Hence, measuring the accuracy and precision of the attacker's guesses is beyond the scope of this model but a subject of future investigations. In addition to Shannon entropy, min-entropy is another technique to assess the risk of privacy breach by quantifying the number of trials required to guess the information successfully [177, 199].

2 This model is based on previous research work appeared in [8]‡.

suppose that an attacker can only eavesdrop the communication channel and get some location information exchanged between mobile badge checkers and a remote server, maintaining different employee records [8]♮. In this case, the attacker does not know the original locations of all employees, and his knowledge can be limited to the total number of subjects and areas. The attacker goal is to predict the locations of all subjects with a high degree of confidence through guessing the probabilities given by the matrix presented in Equation (5.3):

$$\Pr(\mathcal{O}) = \begin{array}{c} \\ 1 \\ \vdots \\ s \end{array} \begin{array}{c} \begin{array}{ccc} 1 & \dots & z \end{array} \\ \begin{pmatrix} \Pr(o_{1,1}) & \dots & \Pr(o_{1,z}) \\ \vdots & \ddots & \vdots \\ \Pr(o_{s,1}) & \dots & \Pr(o_{s,z}) \end{pmatrix} \end{array} \tag{5.3}$$

where

- $0 \leqslant \Pr(o_{i,j}) \leqslant 1, \quad \forall i \in \mathcal{S}$ and $\forall j \in \mathcal{Z}$, as well as
- $\sum_{j=1}^{z} \Pr(o_{i,j}) = 1, \quad \forall i \in \mathcal{S}$.

Before doing any spot check or as long as a subject i is not yet checked, he can be located in any area within the facility [189].

Assumption 1 (Uniform occurrence probabilities)**.** There is an equal chance of the subject $i \notin CS$ being in any area of the facility of interest:

$$\Pr(o_{i,j}) = \frac{1}{z} \quad \forall i \notin CS \text{ and } \forall j \in \mathcal{Z}$$

Suppose that a single inspection mission is performed per area, how to measure the extent of location privacy threat to which the employees are exposed, taking into account that in one mission, multiple subjects can be checked. In this work, a metric for quantifying location privacy based on information theory (i. e., an entropy-based metric) is presented. Broadly speaking, entropy is a measure of unpredictability of information content [189, 193]. If the content is certain, the entropy is minimized (equal to zero), and the outcome can be predicted perfectly. As entropy increases, the attacker becomes more and more uncertain. This concept is utilized to measure the average level of uncertainty of the attacker about the employees' locations after observing some relevant traces.

Definition 1 (Shannon entropy)**.** The information entropy of a discrete random variable X according to Shannon can be expressed by:

$$H(X) = -\sum_{x \in X} \Pr(x) \log_2 \Pr(x) \tag{5.4}$$

In light of Equation (5.4) and using the rule of compound probability, the Shannon entropy of a system of independent events like $\mathcal{O}$ can be given by:

$$H(\mathcal{O}) = -\sum_{j_1=1}^{z} \sum_{j_2=1}^{z} \dots \sum_{j_s=1}^{z} \left(\prod_{i=1}^{s} \Pr(o_{i,j_i}) \right) \cdot \left(\log_2 \prod_{i=1}^{s} \Pr(o_{i,j_i}) \right) \tag{5.5}$$

For the sake of illustration, let $z = 2$ (two areas: A, B) and $s = 3$ (three subjects denoted by 1, 2, and 3):

Example 1. If $CS = \emptyset$ (i. e., no subject is checked): According to Assumption 1, the probability matrix of the event set $\mathcal{O}$ is

$$\Pr(\mathcal{O}) = \begin{array}{c} \\ 1 \\ 2 \\ 3 \end{array} \begin{array}{c} \begin{array}{cc} A & B \end{array} \\ \begin{pmatrix} 1/2 & 1/2 \\ 1/2 & 1/2 \\ 1/2 & 1/2 \end{pmatrix} \end{array}$$

The corresponding Shannon entropy is

$$H(\mathcal{O}) = -\sum_{j_1=1}^{2}\sum_{j_2=1}^{2}\sum_{j_3=1}^{2}\left(\prod_{i=1}^{3}\Pr(o_{i,j_i})\right)\left(\log_2\prod_{i=1}^{3}\Pr(o_{i,j_i})\right) =$$

$$-\begin{pmatrix} (\Pr(o_{1,1}).\Pr(o_{2,1}).\Pr(o_{3,1}).\log_2(\Pr(o_{1,1}).\Pr(o_{2,1}).\Pr(o_{3,1})))+ \\ (\Pr(o_{1,1}).\Pr(o_{2,1}).\Pr(o_{3,2}).\log_2(\Pr(o_{1,1}).\Pr(o_{2,1}).\Pr(o_{3,2})))+ \\ (\Pr(o_{1,1}).\Pr(o_{2,2}).\Pr(o_{3,1}).\log_2(\Pr(o_{1,1}).\Pr(o_{2,2}).\Pr(o_{3,1})))+ \\ (\Pr(o_{1,1}).\Pr(o_{2,2}).\Pr(o_{3,2}).\log_2(\Pr(o_{1,1}).\Pr(o_{2,2}).\Pr(o_{3,2})))+ \\ (\Pr(o_{1,2}).\Pr(o_{2,1}).\Pr(o_{3,1}).\log_2(\Pr(o_{1,2}).\Pr(o_{2,1}).\Pr(o_{3,1})))+ \\ (\Pr(o_{1,2}).\Pr(o_{2,1}).\Pr(o_{3,2}).\log_2(\Pr(o_{1,2}).\Pr(o_{2,1}).\Pr(o_{3,2})))+ \\ (\Pr(o_{1,2}).\Pr(o_{2,2}).\Pr(o_{3,1}).\log_2(\Pr(o_{1,2}).\Pr(o_{2,2}).\Pr(o_{3,1})))+ \\ (\Pr(o_{1,2}).\Pr(o_{2,2}).\Pr(o_{3,2}).\log_2(\Pr(o_{1,2}).\Pr(o_{2,2}).\Pr(o_{3,2}))) \end{pmatrix} \begin{array}{l} \overset{(1,2,3)}{\longrightarrow} \\ \longrightarrow (A,A,A) \\ \longrightarrow (A,A,B) \\ \longrightarrow (A,B,A) \\ \longrightarrow (A,B,B) \\ \longrightarrow (B,A,A) \\ \longrightarrow (B,A,B) \\ \longrightarrow (B,B,A) \\ \longrightarrow (B,B,B) \end{array}$$

$$H(\mathcal{O}) = -\left(\underbrace{\left(\frac{1}{2}\right)^3.\log_2\left(\frac{1}{2}\right)^3+\ldots+\left(\frac{1}{2}\right)^3.\log_2\left(\frac{1}{2}\right)^3}_{8\text{ times}}\right)$$
$$= \log_2 2^3 = 3$$

In this case, Equation (5.5) can be reduced to the form:

$$\begin{aligned} H(\mathcal{O}) &= -\sum_{k=1}^{z^s}\underbrace{\left(\frac{1}{z}\right)\ldots\left(\frac{1}{z}\right)}_{s\text{ times}}\cdot\log_2\left(\left(\frac{1}{z}\right)\ldots\left(\frac{1}{z}\right)\right) \\ &= \log_2(z^s) = H_0(\mathcal{O}), \end{aligned} \tag{5.6}$$

where $H_0(\mathcal{O})$ is the maximum value entropy and known as Hartley entropy [220].

Example 2. If $CS = \{1\}$ (i. e., subject $i = 1$ is checked): Suppose that subject $i = 1$ is checked in area A and such piece of information is leaked; that is, $\Pr(o_{1,A}) = 1$ and $\Pr(o_{1,B}) = 0$. Hence, the probability matrix changes as follows:

$$\Pr(\mathcal{O}) = \begin{array}{c} \\ 1 \\ 2 \\ 3 \end{array} \begin{array}{c} \begin{array}{cc} A & B \end{array} \\ \begin{pmatrix} 1 & 0 \\ 1/2 & 1/2 \\ 1/2 & 1/2 \end{pmatrix}, \end{array}$$

which affects Shannon entropy as well:

$$\mathrm{H}(\mathcal{O}) = \qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad (1,2,3)$$

$$-\begin{pmatrix} (\overbrace{\Pr(o_{1,1})}^{1}.\Pr(o_{2,1}).\Pr(o_{3,1}).\log_2(\overbrace{\Pr(o_{1,1})}^{1}.\Pr(o_{2,1}).\Pr(o_{3,1})))+ & \longrightarrow (A,A,A) \\ (\Pr(o_{1,1}).\Pr(o_{2,1}).\Pr(o_{3,2}).\log_2(\Pr(o_{1,1}).\Pr(o_{2,1}).\Pr(o_{3,2})))+ & \longrightarrow (A,A,B) \\ (\Pr(o_{1,1}).\Pr(o_{2,2}).\Pr(o_{3,1}).\log_2(\Pr(o_{1,1}).\Pr(o_{2,2}).\Pr(o_{3,1})))+ & \longrightarrow (A,B,A) \\ (\Pr(o_{1,1}).\Pr(o_{2,2}).\Pr(o_{3,2}).\log_2(\Pr(o_{1,1}).\Pr(o_{2,2}).\Pr(o_{3,2}))) & \longrightarrow (A,B,B) \end{pmatrix}$$

$$= \log_2 2^2 = 2$$

Let $\overline{CS} = \mathcal{S} \setminus CS$ is the set of unchecked subjects. Then, the probability matrix of the event set ($\overline{\mathcal{O}}$) associated with the unchecked subjects can be defined as follows:

$$\Pr(\overline{\mathcal{O}}) = \begin{matrix} \\ 1 \\ \vdots \\ |\overline{CS}| \end{matrix} \begin{matrix} \begin{matrix} 1 & \dots & z \end{matrix} \\ \begin{pmatrix} \Pr(\bar{o}_{1,1}) & \dots & \Pr(\bar{o}_{1,z}) \\ \vdots & \ddots & \vdots \\ \Pr(\bar{o}_{|\overline{CS}|,1}) & \dots & \Pr(\bar{o}_{|\overline{CS}|,z}) \end{pmatrix} \end{matrix}$$

The corresponding information entropy is then defined by

$$\mathrm{H}(\overline{\mathcal{O}}) = -\sum_{j_1=1}^{z}\sum_{j_2=1}^{z}\dots\sum_{j_{|\overline{CS}|}=1}^{z}\left(\prod_{i=1}^{|\overline{CS}|}\Pr(\bar{o}_{i,j_i})\right)\left(\log_2\prod_{i=1}^{|\overline{CS}|}\Pr(\bar{o}_{i,j_i})\right) \tag{5.7}$$

And under Assumption 1, it can be reduced to the form

$$\begin{aligned}\mathrm{H}(\overline{\mathcal{O}}) &= -\sum_{k=1}^{|\overline{CS}|}\underbrace{\left(\frac{1}{z}\right)\dots\left(\frac{1}{z}\right)}_{|\overline{CS}| \text{ times}}\log_2\left(\underbrace{\left(\frac{1}{z}\right)\dots\left(\frac{1}{z}\right)}_{|\overline{CS}| \text{ times}}\right) \\ &= \log_2 z^{|\overline{CS}|} = \log_2 z^{s-|CS|}\end{aligned} \tag{5.8}$$

Consequently,

- if $CS = \emptyset$: $\mathrm{H}(\overline{\mathcal{O}}) = \mathrm{H}_0(\mathcal{O})$ and $\Pr(\overline{\mathcal{O}}) = \Pr(\mathcal{O})$;
- if $CS \neq \emptyset$: $\mathrm{H}(\overline{\mathcal{O}}) < \mathrm{H}_0(\mathcal{O})$ and $\Pr(\overline{\mathcal{O}}) \neq \Pr(\mathcal{O})$.

That is, the maximum uncertainty is reached when no subject is checked. Any deviation means less uncertainty and should be then quantified in lower entropy. The extreme case occurs when all employees are checked, eliminating any uncertainty. Thus, the entropy is equal to zero, given the fact that the attacker has the potential to access location traces of the employees.

Based on the discussion above, privacy preservation of security inspection missions can be defined as the normalized Shannon entropy [114]:

$$V(\mathcal{O}) = \frac{\mathrm{H}(\overline{\mathcal{O}})}{\mathrm{H}_0(\mathcal{O})} \tag{5.9}$$

However, Equation (5.9) can be reduced under Assumption 1 as follows:

$$V(\mathcal{O}) = \frac{\log_2 z^{s-|CS|}}{\log_2 z^s} = \frac{s-|CS|}{s} \tag{5.10}$$

The larger the quantity $V(\mathcal{O})$, the greater the privacy preservation of an inspection strategy. That is:

- If $CS = \emptyset$: $V(\mathcal{O}) = 1$, which represents the maximum degree of privacy preservation.
- If $CS \neq \emptyset$: $V(\mathcal{O}) < 1$, meaning a non-null privacy breach.
- If $CS \equiv \mathcal{S}$: $V(\mathcal{O}) = 0$, which corresponds to the maximum degree of a privacy breach.

5.4.2 *Time-based model*

While the model described in Section 5.4.1 introduces the normalized Shannon entropy as a metric to quantify the impact of doing spot checks on location privacy, it relies on a static probability matrix $Pr(\mathcal{O})$. Such a matrix is only computed once a single inspection mission is finished ignoring any information from any previously conducted missions. That is, all events associated with unchecked employees follow the uniform probability distribution (cf. Assumption 1).

This section explains how $Pr(\mathcal{O})$ can be updated based on location traces stemming from current and previous missions causing continuous privacy violation. Attackers, who manage to access such traces, could guess with better confidence the location of subjects even a while after being checked (i. e., better than the mere *even* guess $1/z$). Measuring privacy breach/preservation as a function of time implies that predicting the location of an employee shortly after being checked would be with higher confidence compared to a prediction made a long time after the last check. This is due to the fact that the acquired data at a given time (right after the check) loses its value over time and slowly converges to the steady (equally likely) state. Thus, the respective probabilities of a subject i to be (in or *out* of a given area j should rather evolve over time. In reality, checks are made in an asynchronous way: repetitive and spread over time. As time ticks, the subject i, who was checked in area j at t_0, can leave this area and move to another zone. Thus, the probabilities $Pr_{t_0}(o_{i,j})$ and $Pr_{t_0}(o_{i,j'\neq j})$, which were at t_0 (when the check occurred) equal to 1 and 0 respectively, change over time until they reach a steady-state value of $1/z$ (as stated by Assumption 1) [8]♮.

Recall here that the exponential distribution is usually used to model the elapsed time between the occurrences of events in a Poisson process. That is, it is an appropriate model if the following conditions are satisfied [8]♮:

- T is the time between events, with ($T > 0$).
- If an event e_1 occurs, it will not affect the probability of occurrence of another event e_2 (events are independent of each other).
- The rate at which events occur is constant; i. e., events occur at random independent of the past, but with a known long term average rate.

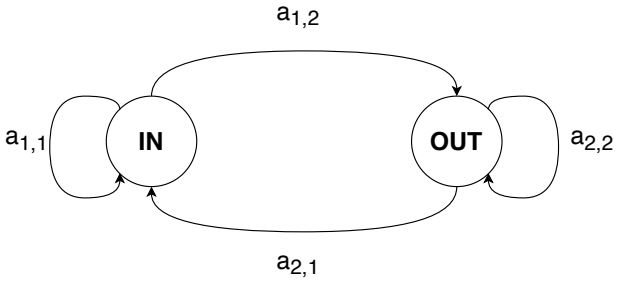

Figure 5.1: The model of subject movement from/to a given area

Suppose a consistent movement pattern[3], then the stay duration (denoted by $T_i^{in}(j)$) of a subject i in a given area j can be modeled by an exponential distribution of parameter $\lambda_{i,j}$ (i. e., $T_i^{in}(j) \sim exp(\lambda_{i,j})$)[4]. Similarly, the time spent outside an area j (denoted as $T_i^{out}(j)$) follows an exponential distribution of parameter $\mu_{i,j}$ (i. e., $T_i^{out}(j) \sim exp(\mu_{i,j})$)[5]. Otherwise, the working day can be divided into different time intervals such that for each particular interval, the following assumptions are kept correct:

Assumption 2. $\forall i \in \mathcal{S}$ and for a given area $j \in \mathcal{A}$: The stay duration of a subject i in an area j is given by the random variable $T_i^{in}(j) \sim exp(\lambda_{i,j})$. Moreover, the stay duration of a subject i outside area j is represented by the random variable $T_i^{out}(j) \sim exp(\mu_{i,j})$.

Assumption 3. There are several non-overlapping zones $z > 1$, otherwise the event $o_{i,1}$ will always be true $\forall i \in \mathcal{S}$.

Under Assumption 2, each area j can be modeled as a CTMC process of two states (cf. Figure 5.1): $IN_{i,j}$ (the subject i is in the area j) and $OUT_{i,j}$ (the subject i is out of the area j). For the sake of convenience, the generator matrix of the underlying stochastic process $A = \begin{pmatrix} a_{1,1} & a_{1,2} \\ a_{2,1} & a_{2,2} \end{pmatrix}$ is computed. This generator is an array describing the rates at which the CTMC moves between the different states (Here, $a_{1,2} = \lambda_{i,j}$ and $a_{2,1} = \mu_{i,j}$).

For small t:

- Suppose that a subject i was initially (at t_0) in the area j, he is still in the area j at time t, then:

$$P\left(X_t = IN_{i,j} | X_{t_0} = IN_{i,j}\right) \simeq Pr(T_i^{in}(j) > t) = e^{-\lambda_{i,j} t};$$

  this leads to the rate

$$a_{1,1} = \frac{d}{dt} P\left(X_t = IN_{i,j} | X_{t_0} = IN_{i,j}\right)_{|(t=0)} = -\lambda_{i,j}$$

- If the subject i was initially outside the area j and stays in the same state at time t, then:

$$P\left(X_t = OUT_{i,j} | X_{t_0} = OUT_{i,j}\right) \simeq Pr(T_i^{out}(j) > t) = e^{-\mu_{i,j} t};$$

3 That is, each individual subject (e. g., employee) behaves exactly the same way for the working day, and the behavior is not affected by special times such as lunchtime.

4 If the average sojourn time of a subject i in an area j, given by $E[T_i^{in}(j)]$, is known, then $\lambda = \frac{1}{E[T_i^{in}(j)]}$ based on the properties of exponentially distributed random variables, e. g., $E[T_i^{in}(j)] = 3\ h \equiv \lambda = 0.333\ h^{-1}$.

5 λ and μ can be perceived as the departure/arrival rate from/in a specific area, respectively.

this leads to

$$a_{2,2} = \frac{d}{dt} P\left(X_t = \text{OUT}_{i,j} | X_{t_0} | (t = 0) = \text{OUT}_{i,j}\right)_{|(t=0)} = -\mu_{i,j}$$

For the sake of simplicity, the subscript is omitted, that is $\lambda_{i,j} = \lambda$, $\mu_{i,j} = \mu$, $\text{IN}_{i,j}$ is replaced by IN, and $\text{OUT}_{i,j}$ by OUT. Consequently, the generator matrix of the presented CTMC is given by:

$$A = \begin{pmatrix} -\lambda & \lambda \\ \mu & -\mu \end{pmatrix}$$

Then, the (2×2) transition probability matrix $P(t) = \left(P_{l_1,l_2}(t)\right)$, where $P_{l_1,l_2}(t)$ is the probability that the process is in state $l_2 \in \{\text{IN}, \text{OUT}\}$ at time t given that it started in a state $l_1 \in \{\text{IN}, \text{OUT}\}$ at time t_0, can be computed using a matrix-exponential representation [102]:

$$\begin{aligned} P(t) &= e^{At} \\ &= \frac{1}{\mu+\lambda}\begin{pmatrix} \mu & \lambda \\ \mu & \lambda \end{pmatrix} + \frac{e^{-(\lambda+\mu)t}}{\mu+\lambda}\begin{pmatrix} \lambda & -\lambda \\ -\mu & \mu \end{pmatrix} \\ &= \begin{matrix} \\ \text{IN} \\ \text{OUT} \end{matrix} \begin{matrix} \begin{matrix} \text{IN} & \text{OUT} \end{matrix} \\ \begin{pmatrix} \frac{\mu}{\lambda+\mu} + \frac{\lambda}{\lambda+\mu}e^{-(\lambda+\mu)t} & \frac{\lambda}{\lambda+\mu} - \frac{\lambda}{\lambda+\mu}e^{-(\lambda+\mu)t} \\ \frac{\mu}{\lambda+\mu} - \frac{\mu}{\lambda+\mu}e^{-(\lambda+\mu)t} & \frac{\lambda}{\lambda+\mu} + \frac{\mu}{\lambda+\mu}e^{-(\lambda+\mu)t} \end{pmatrix} \end{matrix} \end{aligned}$$

Let $\delta = \mu/(\lambda+\mu)$, this yields:

$$P(t) = \begin{matrix} \\ \text{IN} \\ \text{OUT} \end{matrix} \begin{matrix} \begin{matrix} \text{IN} & \text{OUT} \end{matrix} \\ \begin{pmatrix} \delta + (1-\delta)e^{-(\lambda+\mu)t} & (1-\delta) - (1-\delta)e^{-(\lambda+\mu)t} \\ \delta - \delta e^{-(\lambda+\mu)t} & (1-\delta) + \delta e^{-(\lambda+\mu)t} \end{pmatrix} \end{matrix} \tag{5.11}$$

For small t, the transient state occupancy probabilities for the presented CTMC can be computed as follows:

$$\pi(t) = (\pi_{\text{IN}}(t), \pi_{\text{OUT}}(t)) = \pi(t_0) \times P(t), \tag{5.12}$$

where

- $\pi_{\text{IN}}(t)$ is the probability that the subject i is in the state IN (i. e., inside the area j) at time t,
- $\pi_{\text{OUT}}(t)$ is the probability that the subject i is in the state OUT (i. e., outside the area j) at t, and
- $\pi(t_0)$ describes the initial state whether the subject i was checked in the area j during his last check at time t_0 (i. e., $\pi(t_0) = (1, 0)$) or in another area $j' \neq j$ (i. e., $\pi(t_0) = (0, 1)$).

Since there is a direct path from each state to the other one, the presented Markov chain is strongly connected and thereby irreducible [35]. That is, the chain converges and enters a steady-state for a large value of t:

$$\pi.P(t) = \pi,$$

where $\pi = (\pi_{IN}, \pi_{OUT})$ is a steady-state probability distribution on the states IN and OUT. In the presented two-state Markov chain, $(\pi_{IN}, \pi_{OUT}) = [\delta, 1-\delta]$ and can be computed by solving [144]:

$$\pi A = 0$$

with the additional constraint

$$\pi_{IN} + \pi_{OUT} = 1$$

Remark. Under Assumption 1, when the steady-state is reached, then $\pi_{IN} = \frac{1}{z}$ and $\pi_{OUT} = 1 - \pi_{IN} = \frac{z-1}{z}$. Hence, $\lambda = (z-1)\mu$ (taking into account that $z > 1$ as stated in Assumption 3).

Let Δt be the approximated time required to reach the steady-state. It satisfies $\pi_{IN}(\Delta t) - \pi_{IN} \leqslant \varepsilon$, where ε is a very small quantity. This yields:

$$\Delta t \geqslant -\frac{z-1}{\lambda z} \times \ln\left(\frac{z\varepsilon}{z-1}\right).$$

Obviously, the bigger λ is, the faster is the convergence. That is, a location trace (e. g., *the subject* i *is in the area* j) loses its value faster since the expected time of stay (expected sojourn time $E[T_i^{in}(j)] = 1/\lambda$) inside the area j is smaller.

Finally, it is necessary to explain how to use transient and stationary state probability distributions to update the probability matrix $Pr_t(\mathcal{O})$ (see the static matrix defined in Equation (5.3)) required to measure privacy preservation using a metric similar to the one defined in Equation (5.9). Let each subject $i \in S$ be associated with z different CTMCs (one per area $j \in \mathcal{Z}$). This yields $z \times s$ different state machines. At time t and given t_i^0 that corresponds to the time instant when the subject i was last checked, the z different CTMCs associated with the subject i would be reinitialized as follows:

1. t_i^0 is seen as the new time origin (i. e., $t_0 = t_i^0$) for the z CTMCs of the subject i and a specific time variable Δt_i is updated such that $\Delta t_i = t - t_i^0$.

2. For $\Delta t_i < \Delta t$: (i. e., transient state probability distribution)

$$Pr_{\Delta t_i}(o_{i,j}) = \begin{cases} \delta_{i,j} + (1-\delta_{i,j})e^{-(\lambda_{i,j}+\mu_{i,j})\Delta t_i}, & \text{if } \pi_{i,j}(t_0) = (1,0) \\ \delta_{i,j} - \delta_{i,j}e^{-(\lambda_{i,j}+\mu_{i,j})\Delta t_i}, & \text{if } \pi_{i,j}(t_0) = (0,1) \end{cases}$$

3. For $\Delta t_i \geqslant \Delta t$: (i. e., steady-state probability distribution)

$$Pr_{\Delta t_i}(o_{i,j}) = \delta_{i,j}$$

4. Then, the probability matrix at time t is computed using information from all subjects as follows:

$$Pr_t(\mathcal{O}) = \begin{array}{c} \\ 1 \\ \vdots \\ i \\ \vdots \\ s \end{array} \begin{array}{c} \begin{array}{cccccc} 1 & \dots & j & \dots & z \end{array} \\ \begin{pmatrix} Pr_{\Delta t_1}(o_{1,1}) & \dots & \vdots & \dots & Pr_{\Delta t_1}(o_{1,z}) \\ \vdots & \ddots & \vdots & & \vdots \\ \dots & \dots & Pr_{\Delta t_i}(o_{i,j}) & \dots & \dots \\ \vdots & & \vdots & \ddots & \vdots \\ Pr_{\Delta t_s}(o_{s,1}) & \dots & \vdots & \dots & Pr_{\Delta t_s}(o_{s,z}) \end{pmatrix} \end{array}$$

Remark. At any time t, if a subject i was not checked previously in any area, then $t_i^0 = 0$ and $Pr_{\Delta t_i}(o_{i,j}) = \delta_{i,j} \quad \forall j \in Z$.

5. Following Equation (5.5) and Equation (5.9), privacy preservation can be measured as a function of time. In this case, $V(\mathcal{O})$ should be replaced by $V_t(\mathcal{O})$:

$$V_t(\mathcal{O}) = \frac{H_t(\mathcal{O})}{H_0(\mathcal{O})} \tag{5.13}$$

where

- $H_t(\mathcal{O})$ is the information entropy computed at time t using $Pr_t(\mathcal{O})$, and
- $H_0(\mathcal{O})$ is the maximum information entropy computed before doing any spot checks at time $t = 0$. Under Assumption 1, $H_0(\mathcal{O}) = \log_2(z^s)$.

In this work, Equation (5.13) is used to assess the privacy preservation of different inspection strategies in a simulated environment, as explained further in Section 6.5.1.

5.5 SUMMARY

This chapter presents a game-theoretic approach towards the optimization of security measures being physical surveillance systems. This will lead to a minimization of the potential damage an intruder can cause, and thus provide a strategy for risk minimization. Unlike traditional surveillance games, the presented approach assumes that the impact of surveillance systems cannot be expressed completely in a numeric utility to the involved players. Furthermore, this chapter takes the specifics of surveillance technologies into account and tailors the game-theoretic model to the specifically fuzzy terms in which the quality of the surveillance is usually expressed.

The practice of extending static surveillance infrastructures using some dynamic inspection strategies can be perceived as uncomfortable by employees in terms of their privacy. This goal is especially interesting and relevant in the use case of surveillance systems since it is tightly linked with the acceptance of the security measure by the employees. Therefore, this chapter presents an entropy-based model to assess the impact of different security inspection strategies on the preservation of employees' location privacy. This model is integrated into the simulation environment developed to assess the impact of different inspection strategies, as explained in the following case study.

6 USE CASE

The content of Chapter 6 is mainly adopted from the research work published in [2]♮.

6.1 INTRODUCTION

The management process of physical security activities involves identification and assessment of a set of resource allocation choices towards finding the best allocation pattern that minimizes the impact of potential risks. To demonstrate the application of the methodological approach put forth in Section 4.2 and the surveillance games developed in Section 5.3 to address physical intrusion problems, a use case of a nuclear power plant is adopted in this thesis.

A nuclear power plant is a CI that involves several processing units and auxiliary facilities (e. g., nuclear reactor, cooling units, pipes, and steam-turbine-driven electrical generators) [2]♮. The power plant is illustrated by the map presented in Figure 6.1. Being sensitive (i.e., the business and production processes), this infrastructure is a potential target of a wide variety of attacks. Here, the main focus is laid on studying the risk of physical intrusions. Ethnographic studies carried out in a comparable industrial complex showed that due to a prevalent belief of the security personnel that the deployed security solutions (i. e., CCTV cameras and the access control system at the entrance) are able to prevent any illegitimate access, the risk of physical intrusions is seriously underestimated [1]♮. Thus, if such infiltration occurs, it would most likely not be detected early enough due to an inadequate vigilance level. Therefore, this chapter showcases the application of physical surveillance games to manage spot-checking activities in a CI environment. Henceforth, the investigated infrastructure is referred to as *the power plant*.

6.2 CONTEXT ESTABLISHMENT

This step aims at understanding the target CI environment. A business process analysis shows that *the power plant* carries out several critical processes, including heat generation through nuclear fuel, a cooling process (using water pumped from the river next to the site as shown in Figure 6.1), and radioactive waste management operations. Such critical processes make *the power plant* a target of the following attacks:

- sabotage attacks aiming to damage critical buildings and machinery;
- vandalism attacks causing damage to the public and nearby environment through tampering some security and safety measures in place or disrupting the storage/elimination process of the nuclear waste, thereby disastrously interrupting the power generation process or polluting the surrounding area with highly dangerous/radioactive materials; and

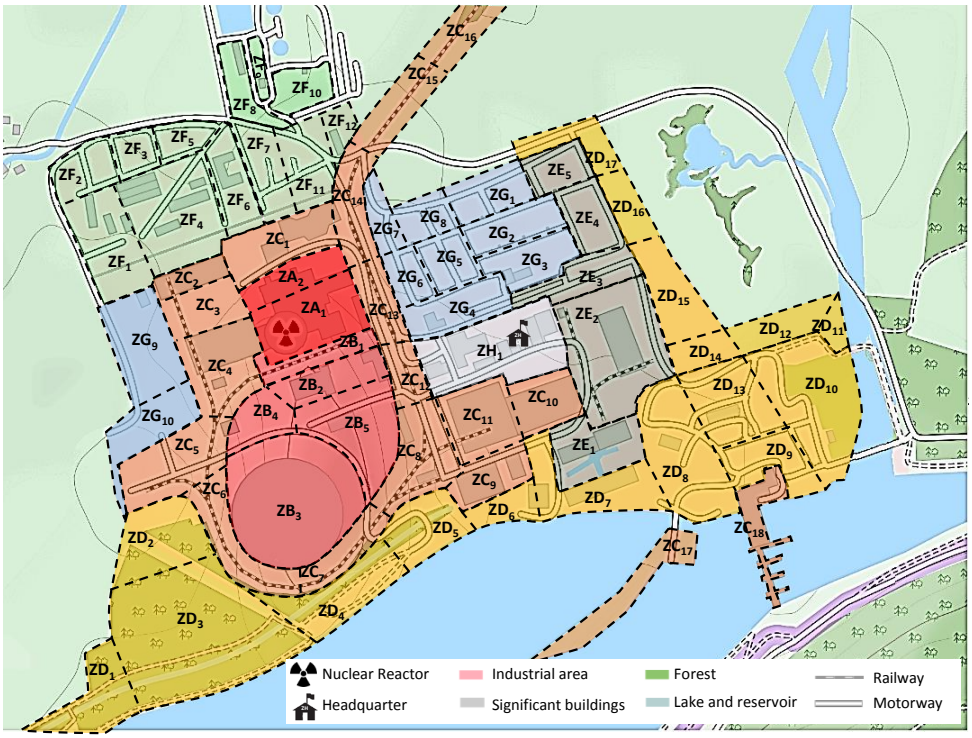

Figure 6.1: A simplified map of *the power plant* [2]♮

- espionage attacks causing leakage of sensitive information, either to a competitor or directly to the public domain resulting in damage to *the power plant*'s reputation and trust.

In *the power plant*, the authorized personnel (i.e., employees, temporary workers, visitors, and maintenance staff) are equipped with security badges storing information such as the owner's ID number, name, photo, and a list of allowed areas (zones). All data can be retrieved upon scanning the badge with a specific verification device. To overcome the limitations of static security measures (e. g., surveillance cameras), there is a vital need for a dynamic security solution. Such a solution seeks to ensure an adequate level of situational awareness within the boundaries of such large-scale CI. In this regard, *the power plant* dedicated 15 employees serving as on-demand security badge inspectors; that is, the 15 guards represent the available security resources to be allocated optimally to minimize the risk of physical intrusion. Every security guard follows a determined inspection schedule instructing him to check the identity of randomly selected persons located in a determined set of zones. For each mission, the inspector will be moving from the headquarter (pointed out as HQ in Zone ZH_1 in Figure 6.1) to a given target zone. Security guards are equipped with mobile devices capable of reading security badges and checking whether a person holding a badge is its rightful owner. The verification process of security badges is exemplified in Figure 6.2.

This study focuses on potential intrusions, in which attackers succeed in accessing *the power plant* using stolen or forged security badges. The goal of security inspections is to

- deal with an intrusion if ever it happens (i.e., the badge verification system at the entrance is already bypassed), and

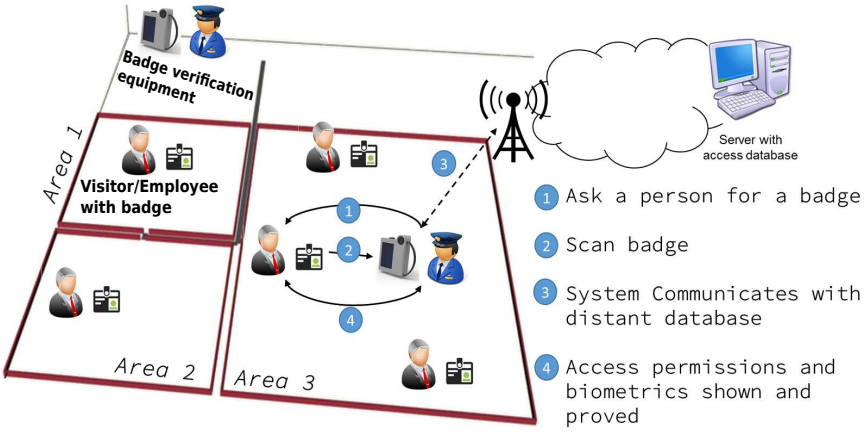

Figure 6.2: Illustration of the security badge verification process

Table 6.1: Security levels of the identified zones

ZONE LABEL	SECURITY LEVEL
ZA_1, ZA_2	16
$ZB_1, \ldots, ZB_6$	12
$ZC_1, \ldots, ZC_{18}$	10
$ZD_1, \ldots, ZD_{17}$	8
$ZE_1, \ldots, ZE_5$	5
$ZF_1, \ldots, ZF_{12}$	4
$ZG_1, \ldots, ZG_{10}$	2
ZH_1	1

- stop the attackers before causing total damage to the facility (the fixed camera network does not really help at this stage since it mostly serves as a reactive solution).

The layout of the area under surveillance (i.e., the buildings, road, perimeter, etc.) is known. The entire area is divided into several zones, as depicted in Figure 6.1. Each zone has a specific security level indicating its level of criticality, as shown in Table 6.1; the higher the level, the more critical the zone. The security level of a specific zone depends on the assets located therein (e. g., areas with control systems or a nuclear reactor such as Zone ZA_1 in Figure 6.1), on the information stored in that zone (e. g., record storage rooms such as Zone ZA_2), or on the ease of intrusion (e. g., the area around the railway such as Zone ZC_{14}). On average, *the power plant* counts 180 employees.

6.3 IDENTIFICATION OF STRATEGIES

This step aims at identifying the action spaces of the defender $\mathcal{D}$ (e. g., security mechanism designer) and the attacker $\mathcal{A}$ (i. e., potential intruders). Throughout this work, all players' strategies are expressed in terms of independent parameters describing the characteristics of each strategy. Strategy parameterization facilitates developing a

common understanding of the wide spectrum of potential strategies and also enables resource allocation through adjusting parameter values. Distinct parameter settings will then correspond to distinct strategies for a player. This step delivers a finite set of strategies that represent a set of "defense proposals" selected based on a certain amount of knowledge and experience, and hence they are likely to be realistic. This set can, therefore, include more strategies or less based on the use case environment and the needs of the security mechanism designer. Furthermore, this work avoids defining fixed and static strategies (e. g., specific patrol routes) since they usually fail in such a dynamic environment involving strategy execution uncertainty (i.e., strategies can be frequently interrupted to handle emergencies and unforeseen events). Therefore, dynamic and high-level strategies are defined to alleviate these challenges. Moreover, taking rounds on random routes or checking zones at random times creates sort of a "moving obstacle defense" [103, 235], since the intruder is confronted with additional uncertainty [1, 10]♮.

In this case study, the attacker strategies (referred to as $SP_{\mathcal{A}}$) are described as strings with syntax NIxTy where "NI" and "T" are parameter labels, followed by numeric values (x and y) concretizing the parameter: $NI \in \{5, 10, 15\}$ refers to the number of involved intruders and $T \in \{R, HSLF\}$ describes the movement pattern of the intruders and the way of targeting the zones. As identified in Section 6.2, one can distinguish between sabotage/vandalism attacks and espionage attacks. In the latter, the attacker is roaming around and targeting the zones randomly to gather more information (e. g., preparation for a cyber-attack as part of an Advanced Persistent Threat (APT)). This movement pattern is denoted by *Random* (R). The former represents a targeted attack, in which the attacker has more knowledge about the critical zones that he will target in the first place to cause maximum damage. This pattern is denoted by *Higher Security Levels First* (HSLF).

Likewise, the defender strategies (referred to as $SP_{\mathcal{D}}$) are labeled with strings that encode three parameters: $NG = 15$ refers to the number of involved security guards and this parameter is the same for all strategies; $T \in \{R, HSLF\}$ refers to the movement pattern of the inspectors where R indicates checking of random zones, and HSLF indicates that zones of higher security levels have a higher priority to be checked; and $F \in \{2, 3, 5, 8\}$ refers to the frequency of inspection missions per working day per security guard. The whole sets of the attacker/defender strategies are depicted in Table 6.2. Finally, it is worth mentioning that the various parameters were identified based on the knowledge of experts involved in the operations and management of CI systems. Those experts had been consulted within the context of the HyRiM project[1].

6.4 IDENTIFICATION OF GOALS

In the given scenario, the decision-making process involves multiple goals to be satisfied. The identified goals are *the caused damage*, *the minimum privacy preservation*, *the maximum comfort breach*, and *the detection rate*. These goals can be quantified as follows:

1 HyRiM (Hybrid Risk Management for Utility Networks), FP7 EU Project Number 608090, online: https://hyrim.net/

Table 6.2: List of the strategies considered for the defender and attacker

#	LABEL	DESCRIPTION
Defender Strategies $\|SP_{\mathcal{D}}\| = 8$:		
1D	NG**15**F**2**T***HSLF***	15 guards & freq(F) = 2 & target(T): HSLF
2D	NG**15**F**3**T***R***	15 guards & freq(F) = 3 & target(T): R
3D	NG**15**F**5**T***R***	15 guards & freq(F) = 5 & target(T): R
4D	NG**15**F**8**T***R***	15 guards & freq(F) = 8 & target(T): R
5D	NG**15**F**3**T***HSLF***	15 guards & freq(F) = 3 & target(T): HSLF
6D	NG**15**F**5**T***HSLF***	15 guards & freq(F) = 5 & target(T): HSLF
7D	NG**15**F**8**T***HSLF***	15 guards & freq(F) =8 & target(T): HSLF
8D	NG**15**F**2**T***R***	15 guards & freq(F) = 2 & target(T): R
Attacker Strategies $\|SP_{\mathcal{A}}\| = 6$		
1A	NI**5**T***R***	5 intruders & target(T): R
2A	NI**5**T***HSLF***	5 intruders & target(T): HSLF
3A	NI**10**T***R***	10 intruders & target(T): R
4A	NI**10**T***HSLF***	10 intruders & target(T): HSLF
5A	NI**15**T***R***	15 intruders & target(T): R
6A	NI**15**T***HSLF***	15 intruders & target(T): HSLF

- *Detection rate* is the ratio of detected intruders (denoted by DNI) to the total number of involved ones (NI). This goal is to be maximized, or equivalently, the defender seeks to minimize the miss rate in the detection.

$$\mathrm{DetectionRate} = \frac{\mathrm{DNI}}{\mathrm{NI}} \tag{6.1}$$

- *Caused damage* is defined as the average time spent inside the targeted zones per intruder, weighted according to the respective security levels. Formally, if NI is the number of intruders, and NA is the number of zones in *the power plant*, then the damage is understood as follows:

$$\mathrm{CausedDamage} = \frac{1}{\mathrm{NI}} \times \sum_{i=1}^{\mathrm{NI}} \sum_{j=1}^{\mathrm{NA}} \mathrm{timeSpent}(\mathrm{atk}_i, \mathrm{ar}_j) \times \mathrm{secLevel}(\mathrm{ar}_j) \tag{6.2}$$

  where $\mathrm{timeSpent}(\mathrm{atk}_i, \mathrm{ar}_j)$ represents the total time spent by the intruder atk_i inside the zone ar_j; and $\mathrm{secLevel}(\mathrm{ar}_j)$ gives the security level of the zone ar_j. Obviously, this goal is to be minimized.

- *Minimum privacy preservation* is inversely related to the maximum possible disclosure of employees' locations. Intuitively, the more frequently spot checks are performed, the more effective the system can be. However, this comes at a price: frequent checks may have an essential impact on the location privacy of the employees, especially if such information leaks. Therefore, the defender is interested

in inspection strategies that maximize the minimum level of privacy preservation. In other words, strategies that keep the maximum privacy disclosure at its minimum are more favorable. To estimate the impact of the different inspection strategies, the entropy-based privacy metric developed in Section 5.4 is employed. It relates potential location privacy breach to the number of security checks and the movement of involved employees.

- *Maximum comfort breach* is the maximum degree of discomfort experienced by the employees in *the power plant*. In fact, the more checks experienced by a worker, the more uncomfortable he will feel. However, it is still a subjective issue after how many checks a person starts feeling uncomfortable and to what extent. Thus, the assessment process relies on ethnographic studies[2] conducted while establishing the context. *The power plant* seeks to minimize the maximum perceived comfort breach in order to satisfy its employees.

It is worth noting that these goals can differ from one organization to another. The detection rate could correspond to the rate of detected criminals/intruders or anomalous behavior. Caused damage could represent physical damage, reputational damage, data theft, and many other types of loss. *The power plant* considers a surveillance setup, which can only detect but not prevent. Thus, detection seems the natural performance indicator for surveillance. Most contemporary attacks (including APTs as one example) try to remain stealthy and involve preparatory phases that do not immediately cause damage but prepare the ground to do so later. Therefore, detecting the enemy is an independent matter, and not necessarily tied to damage up to the detection. Nevertheless, these two goals tend to prefer a defense strategy with a higher frequency of inspection activities. Conversely, comfort breach and privacy preservation prefer strategies with less inspection activities since both reflect the employees' preferences: the higher the inspection frequency the higher the workers' feeling of distrust and the higher the stress in the workplace (reflected by the comfort breach) as well as the more information about the workers' locations and how they spend their time (reflected by the privacy preservation). Increasing monitoring activities might degenerate the trust relationship between the organization and the workers. Employees generally tend to prefer more freedom in the workplace and definitely less monitoring and less stress. Therefore, the delivered security decisions will provide the defender with advice on how to balance between security and the consequences of lower comfort and lack of privacy, which could be costly and time-consuming in terms of potentially high staff turnover.

6.5 EFFECTIVENESS ASSESSMENT

This step focuses on assessing the effectiveness/losses of the different strategies identified in Table 6.2. To achieve this goal, each possible combination of the strategies $(d_i, a_j) \in SP_{\mathcal{D}} \times SP_{\mathcal{A}}$ of both players needs to be evaluated with regards to all goals identified in Section 6.4. The assessment can be performed in various ways such as literature review, historical data and observations, statistical analysis, expert opinions, stockholder surveys, among others.

2 For privacy and confidentiality compliance, only the results of the performed surveys and interviews are summarized in this work (cf. Figure 6.5, as well as Equation (6.3) and (6.4)). The results have been integrated into the developed simulation environment.

In general, running large-scale intrusion scenarios for the sake of assessment turns to be very costly and unrealistic. As an alternative, simulation seems to be a proper way to reproduce each scenario several times, thereby integrating the impact of uncertainty into the assessment results. In this work, a simulation model based on the INET 3.4 framework [34], on top of OMNeT++ 5.1 simulator [215] is developed and used for simulating physical intrusions in CIs. The model allows establishing a faithful image of the physical environment of the facility of interest, the deployed personnel and their behavior, as well as the potential attacks that may occur. All applied policies, such as zone restrictions, employees' profiles, and badge checking policies, as well as behaviors of the different actors, including security guards, field workers, or intruders are reflected in the model. The simulation tool allows each scenario to run several times (results can be then presented as comprehensive distributions). It provides measurements of various KPIs to compare the different deployed strategies. Using these results, one can better assess and compare strategies or even find the optimal inspection plan [5, 10]♯.

6.5.1 *Simulation setup*

This section is devoted to giving a comprehensive overview of the simulation tool developed to assess physical surveillance strategies. The content of this section is based on the work appeared in [5, 8]♮

6.5.1.1 *Physical environment*

In the developed simulation model, the same map layout given by Figure 6.1 has to be established (in terms of the number of areas, their geographic repartition, their sizes, and the routes connecting them). In compliance with the game models specified in Section 5.3, the zones are reachable through a web of paths followed when moving from/towards any of the zones. These zones represent the smallest level of granularity of the studied site. Each zone has an attribute called *security level*, indicating the criticality of the respective area (as described in Table 6.1). All this information, i.e., paths, gates, and areas, is described in an XML file used to build and render the physical structure of the monitored environment. For the sake of illustration, Figure 6.3 provides a simple example of a three-area facility with the corresponding XML code: every area is modeled as a convex polygon with a unique identifier and a set of attributes such as "`position`", "`securityLevel`", "`hasExitPoint`", among others. Paths are modeled as a non-oriented graph $G = (V, E)$, where V is the set of vertices, and E is the set of edges as depicted in Figure 6.4a. Vertices in V represent waypoints, characterized by their geographic coordinates. Each waypoint corresponds to a particular location in the physical site such as intersections or area gates. For every pair of vertices $(v_i, v_j) \in V \times V$, an edge $e_{ij} = (v_i, v_j)$ is added to E if the two waypoints are directly connected by a path in the actual map. It is worth mentioning that it is assumed that employees can only move straight from waypoint v_i to waypoint v_j if there is a connecting edge in E. Such representation allows defining one or more weight functions (e. g., hop count or actual distance) to help selecting the best way to move from one source point (e. g., the head quarter of a security guard) to another destination point (e. g., the gate of a target area).

```
<environment>
    <!-- Areas -->
    <Area id="1" position="100 0 0 200 100 0 100 100 0 200 0 0" hasExistPoint="Yes"
        exitSize="3 2.4" exitPosition="200 50 0" ... />
    <Area id="2" position="205 0 0 300 100 0 205 100 0 300 0 0" hasExistPoint="Yes"
        exitSize="3 2.4" exitPosition="205 50 0" securityLevel="2" ... />
     <Area id="3" position="100 105 0 300 250 0 100 250 0 300 105 0" hasExistPoint="Yes"
         exitSize="3 2.4" exitPosition="200 105 0" securityLevel="10" hasFence="Yes"  ...
         />
</environment>
<map>
  <!-- paths -->
  <waypoints>
       <waypoint id="1" isGate="Yes" description="HQ" position="25 102.5 0"/>
       <waypoint id="2" isGate="No" position="200 102.5 0"/>
       <waypoint id="3" isGate="No" position="202.5 102.5 0"/>
       <waypoint id="4" isGate="No" position="202.5 50 0"/>
       <waypoint id="5" isGate="Yes" description="Gate Area 1" position="200 50 0"/>
       <waypoint id="6" isGate="Yes" description="Gate Area 2" position="205 50 0"/>
       <waypoint id="7" isGate="Yes" description="Gate Area 3" position="200 105 0"/>
  </waypoints>
       <edges e = "(1,2) (2,3) (2,7) (3,4) (4,5) (4,6)" />
</map>
```

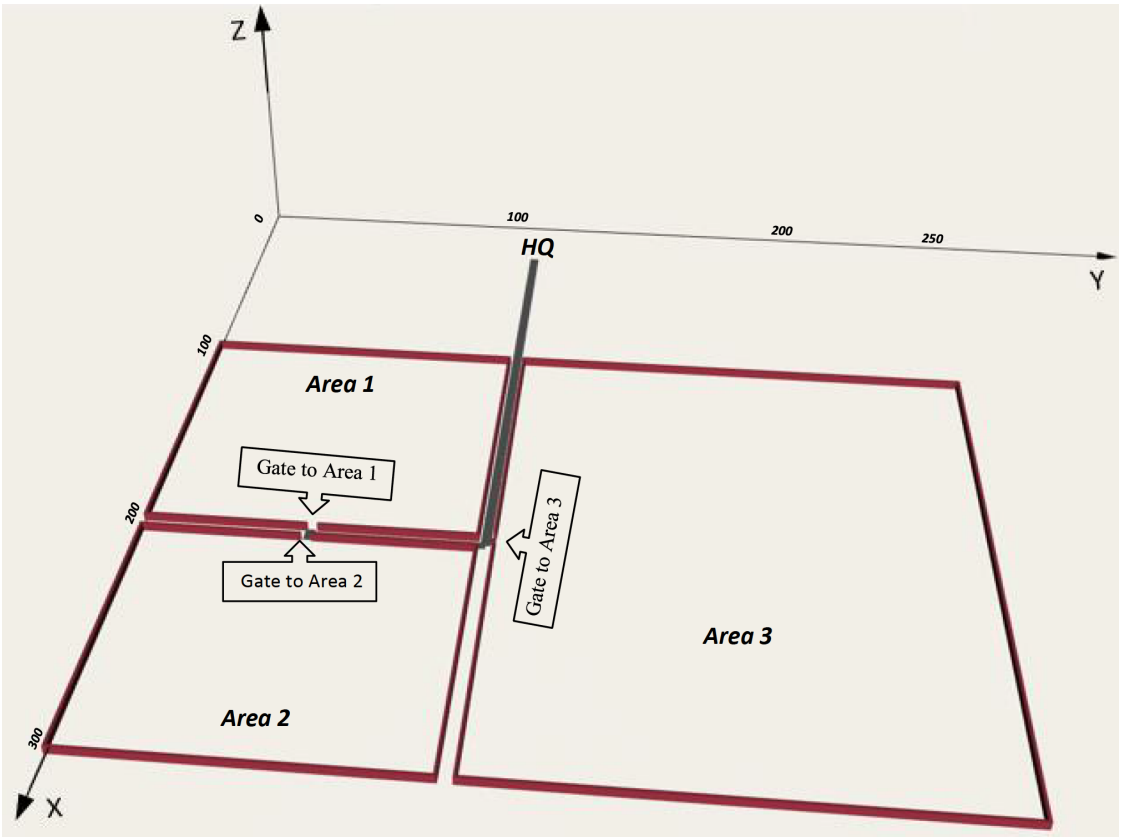

Figure 6.3: A simple three-area physical environment with its corresponding XML-based representation (adapted from [8]♮)

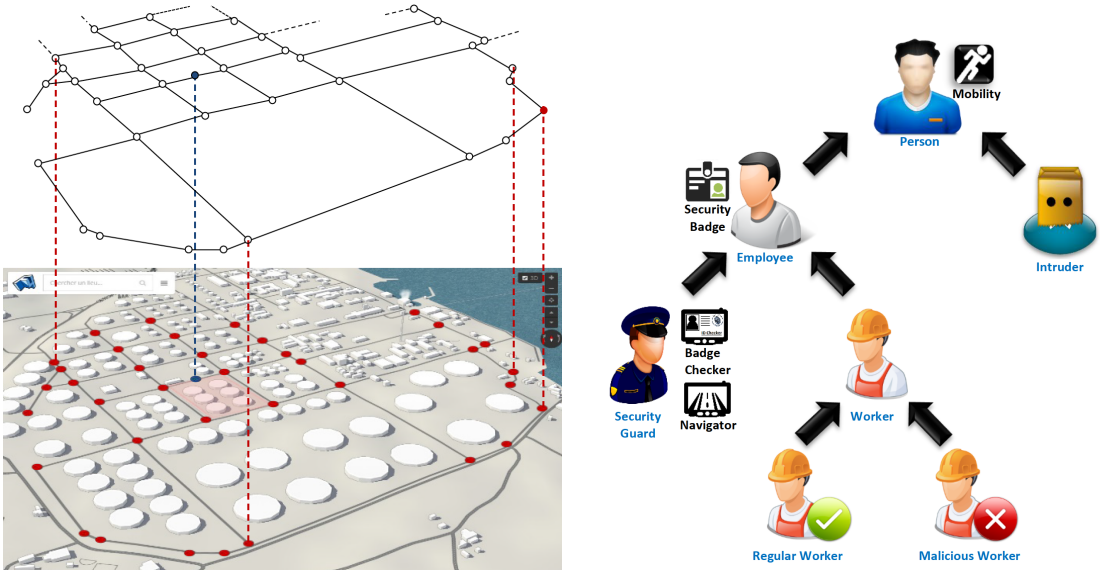

(a) Modeling of paths

(b) Modeling of actors

Figure 6.4: Modeling of paths and actors [8]♮

6.5.1.2 *Actors*

There are two main (mobile) actor categories, namely employees and intruders. An employee can be either a worker or a security guard. Employees hold valid security badges; that is, they are known to the system. In contrast, an intruder is someone from outside the facility either without a security badge or with a fake/stolen one that does not correspond to his biometrics (i. e., fingerprint or facial photo). In all these cases, he will not be recognized by the system as a regular employee. Thus, he should be caught upon the first badge check, whenever it is done, and wherever he is located inside the monitored facility [5, 10]♮.

Depending on their job, employees are allowed to access certain areas of the facility but may be denied access to others. The restrictions vary among employees. In the presented simulation model, there is a set of profiles, each of which indicates a subset of allowed areas. Using an XML file, each worker is assigned one profile, indicating the areas he can access. This information is stored in the security badges. A *regular* worker is a person who does respect areas' restrictions. On the other hand, a *malicious* worker is an employee with a valid ID card, but who intends to harm the facility physically. In this work, it is assumed that such suspicious behavior manifests in targeting areas that the malicious worker is denied to access. During a security check, a malicious worker can only be caught if he is behaving suspiciously at that time (i.e., being in a restricted area when the check takes place). Such information can be acquired from the first step (i. e., context establishment).

Conversely, intruders are not authorized to be in any of the zones of *the power plant*. An intruder may choose to remain in the zone where he is, or move from one zone to another following a given strategy (i.e., "R" or "HSLF"). At the cost of being possibly detected by a security guard, staying in the same zone means adopting a movement pattern similar to a regular employee.

On the other hand, security guards are allowed to access all areas of the facility. A special profile is then created just for them. A security guard owns two main devices: A

navigation system and a badge checker (they are virtually two separate devices, but could also be integrated into one single physical device). The navigation system serves as a mission scheduler. Inspection missions are assigned to a security guard using this device. It first indicates which zone a security guard needs to check, shows the way to reach this area, and decides the strategy to be adopted during the spot check. Figure 6.4b shows the different actors of the underlying use case.

6.5.1.3 *Security inspection mission*

Basically, a mission consists of three phases: (i) select a target area, (ii) visit the targeted area and perform a spot check, and (iii) go back to the headquarter.

FIRST PHASE: It involves choosing a target area (zone from Table 6.1) and guiding the security guard towards it. This selection is made according to a given strategy, i.e.,

- random choice (R), in which each area has the same probability of selection; or
- a choice based on the security level of the zones (HSLF): let $\mathcal{Z}$ be the set of identified zones and sl_i the security level of zone $i \in \mathcal{Z}$. In this case, the probability of selection of zone i is given by $Pr_i = \frac{sl_i}{\sum_{j \in \mathcal{Z}} sl_j}$.

The navigation device, storing the map of the whole site (i.e., areas and paths), guides the security guard initially located at the headquarter (HQ) towards the gate of the targeted area. This is done by applying any shortest path algorithm on the graph representing the paths of the modeled site, between the current position and the gate of the area to be checked. In this phase, security guards are supposed to be equipped with vehicles and thus moving at a speed of 20 km/h at most.

SECOND PHASE: It involves inspecting the area selected in the first phase. The security guard needs to walk (at a speed of 3,6 km/h; 1 m/sec, on average) all around and meet persons located in this area for an eventual badge verification. Inside an area, it is possible to apply any of the mobility models provided by the INET framework. Throughout this work, the well-known Random Waypoint mobility model [42, 94] is adopted. Basically, a mobile node uniformly generates a target position inside the polygon surface of a zone, selects a speed, and then moves towards its target. At its arrival, the node waits at its position for a randomly generated time, before reproducing the same behavior once again. Notice here that all actors are moving with respect to this same mobility model. The only difference might be the *move-wait* pattern. Workers would spend more time in the same place doing some work, then moves to another place to do some other work.

On the other hand, a security guard would spend most of the time moving from one position to another, with short waits. A malicious subject, either an intruder or a worker, would be moving like a regular worker, spending as much time waiting as he is supposed to do some harmful work. In the simulation tool, these values are in order of several minutes for workers and intruders, but a few seconds for a security guard. Moreover and for the sake of simplicity, it is assumed that the process, which describes when the workers change their zones, follow the Poisson process with the departure rate of $\lambda = 0.33\ h^{-1}$; or equivalently expected stay equal to 3 h. Of course, different values of λ can be used to model different relations between workers and zones. While

moving, a security guard will meet persons who are in the checked area. For everyone in his direct vicinity, a security guard decides to check his security badge with a given probability (by default, the probability is set to 0.5). This probability should be closely related to the security level of the area. Every selected subject remains at his current place until the check is performed. If a malicious person (i.e., intruder or worker) is detected, a *handle situation* procedure is triggered. This procedure could be of any type, like (i) calling a third party to drive the caught individual to an interrogation room, (ii) stopping the mission and driving the checked person back to the headquarter by himself, (iii) remove the malicious node from the simulation and continue the checking mission (this option is adopted in this work), or (iv) more drastically stopping the simulation.

Besides, to avoid that a security guard repeatedly checks the same person again and again during the same mission, a memory module is added to the security guard. This module, being adjustable, will control the behavior of a security guard according to three basic features: (i) how easily can he remember a new face?, (ii) how long can he keep remembering it?, and (iii) how many faces can he remember? The first feature, called *memory quality*, is a probability-like parameter to be given as input: it ranges between 0 meaning that he cannot remember anything, and 1 meaning he remembers everything. The second feature called *memory time*, which is a time duration to be given as input. It can either be a fixed duration or a distribution (e. g., a uniform distribution), which indicates how long a newly met face is remembered. Every new entry to the memory will be assigned a *memory time* value to decide when it should be forgotten. The third feature represents the size of the memory and hence called *memory size*. It is implemented as a circular buffer so that if it is full, the entry having the smallest *memory time* value would be forgotten first. Based on some expert feedback, the values of memory quality, time, and size are 0.3, uniform(30min, 2h) (i.e., uniformly distributed between 30min and 2h), and 15, respectively. The end of the second phase can be determined in several ways: It can end after (i) a time duration spent inside the area, (ii) a number of checked persons is reached, or (iii) the checking ratio goes beyond a given threshold (only possible if the number of workers inside the area is known in advance). In any of these cases, the security guard announces the end of this phase using his navigator device. Moreover, the mission proceeds to its third and final phase. In this work, every mission lasts for a duration between $\mathtt{MinDuration} = 10$ minutes and $\mathtt{MaxDuration} = 20$ minutes. A verification operation may last between 1 and 3 minutes. Moreover, the number of inspection missions is controlled through the frequency parameter F of each defense strategy defined in Table 6.2, uniformly spread over the 8 working hours.

THIRD PHASE: It only involves guiding the security guard back to the headquarter using the reverse path stored in the navigation device. The security guard needs to empty his memory because, in the upcoming missions, he should be able to re-check a person as this person could move from one zone to another at any time.

6.5.1.4 *Implementation of goals*

As explained in Section 6.4, the defender seeks to keep a balance between the detection rate, caused damage, privacy preservation, and comfort breach. Some goals can be easily integrated into the developed simulation model, such as detection rate and caused

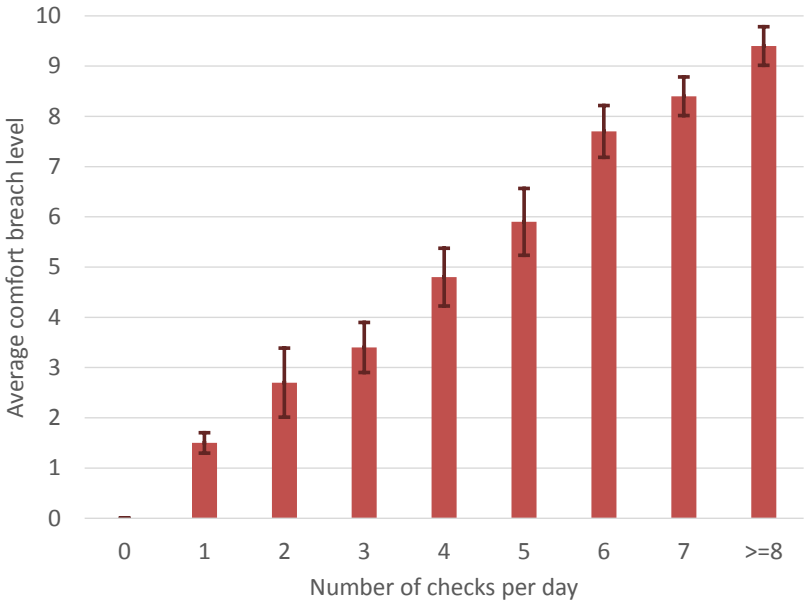

Figure 6.5: Average comfort breach with 95% confidence intervals

damage by applying Equation (6.1) and Equation (6.2), respectively. In contrast, other goals, like privacy preservation and comfort breach, need more attention.

Surveys and ethnographic methods are very valuable to gain insights into assessing the subjective feeling of workers regarding their comfort breach caused by repetitive security checks. Through a questionnaire conducted within the context of the HyRiM project, 35 employees were asked about their feelings (scored between 0 and 10, where 0 is total comfort preservation and 10 means a maximum comfort breach) if they get checked $0, 1, 2, \ldots, \geqslant 8$ times a day. The collected data is summarized with respect to the number of checks per day and depicted in Figure 6.5. To integrate this piece of information into the simulation model, a multivariate Gaussian [16] of 9 dimensions $(\mu \in \mathbb{R}^9, \Sigma \in \mathbb{R}^{9\times 9})$ is established as a primary generator of non-satisfaction in the sense of the degree of discomfort. The respective mean vector μ and the covariance matrix Σ are represented in Equation (6.3) and Equation (6.4), respectively. In this way, it is possible to create as many workers as needed with different subjective comfort breach measures, but following the same general shape as the one shown in Figure 6.5.

$$\mu = (0,\ 1.5,\ 2.7,\ 3.4,\ 4.8,\ 5.9,\ 7.7,\ 8.4,\ 9.4) \tag{6.3}$$

$$\Sigma = \begin{array}{c} \#\text{checks} \\ 0 \\ 1 \\ 2 \\ 3 \\ 4 \\ 5 \\ 6 \\ 7 \\ \geqslant 8 \end{array} \begin{array}{c} \begin{array}{ccccccccc} 0 & 1 & 2 & 3 & 4 & 5 & 6 & 7 & \geqslant 8 \end{array} \\ \begin{pmatrix} 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0.35 & 1.2 & 0.85 & 1 & 0.95 & 0.6 & 0.3 & 0.3 \\ 0 & 1.2 & 4 & 3 & 3.3 & 3.8 & 2.6 & 1.5 & 1.5 \\ 0 & 0.85 & 3 & 2.1 & 2.4 & 2.7 & 1.9 & 1.25 & 1.25 \\ 0 & 1 & 3.3 & 2.4 & 2.8 & 2.7 & 1.85 & 0.9 & 0.9 \\ 0 & 0.95 & 3.8 & 2.7 & 2.7 & 3.75 & 2.8 & 1.8 & 1.8 \\ 0 & 0.6 & 2.6 & 1.9 & 1.85 & 2.8 & 2.25 & 1.55 & 1.55 \\ 0 & 0.3 & 1.5 & 1.25 & 0.9 & 1.8 & 1.55 & 1.25 & 1.25 \\ 0 & 0.3 & 1.5 & 1.25 & 0.9 & 1.8 & 1.55 & 1.25 & 1.25 \end{pmatrix} \end{array} \tag{6.4}$$

To assess the different defense strategies in terms of privacy preservation, the model presented in Section 5.4 based on information theory and Markov chains is integrated into the developed simulator. The model looks at location privacy as the capacity of an attacker, in the worst-case scenario, to estimate the employees' positions with high confidence at a given instant. It further allows capturing the decreasing significance of leaked location information over time. As explained in Section 5.4, privacy is assessed in each simulation run using Equation (5.13) at one-minute intervals towards identifying the required minimum privacy preservation quantity.

6.5.2 *Assessment results*

As explained in Section 4.2.4, multiobjective game models demand that all assessment outcomes should be comparable through defining a common scale[3] over all identified goals. Each (defender, attacker) strategy combination is assessed using 100 simulation runs, which is the size of each assessment sample. Afterward, the simulation results are categorized based on a popular scale in the risk community, which is the five-category scale {1: *Very Low*, 2: *Low*, 3: *Medium*, 4: *High*, 5: *Very High*}. It is worth noting that this work follows a hybrid (semi-quantitative) risk assessment method [192] that allows assessing and comparing risks and various strategies in a more rigorous way benefiting from both quantitative and qualitative risk assessment approaches. Categorical labeling is the core of a hybrid risk assessment process and a requirement of the developed game solver [12]$^{\natural}$. This involves that all assessments should be ranked using the same (organizationally) predefined rating scale. Throughout this work, the categorization process is simply performed by dividing the resulting assessments into the five risk categories that span the numeric range of the respective goals, thereby constructing loss distributions from the assessment results. The distributions depicted in Figure 6.6, Figure 6.7, Figure 6.8, and Figure 6.9 form the payoff (matrices) of the game concerning the four identified goals, which are maximum comfort breach, minimum privacy preservation, caused damage, and detection rate, respectively. Each matrix element (i, j) represents the assessment results of the combination $\left(d_i, a_j\right) \in SP_{\mathcal{D}} \times SP_{\mathcal{A}}$ with respect to a particular goal [2, 8]$^{\natural}$.

3 For technical reasons, the used scale should not contain values smaller than 1 (cf. regularity condition R1 from Section 4.1.2).

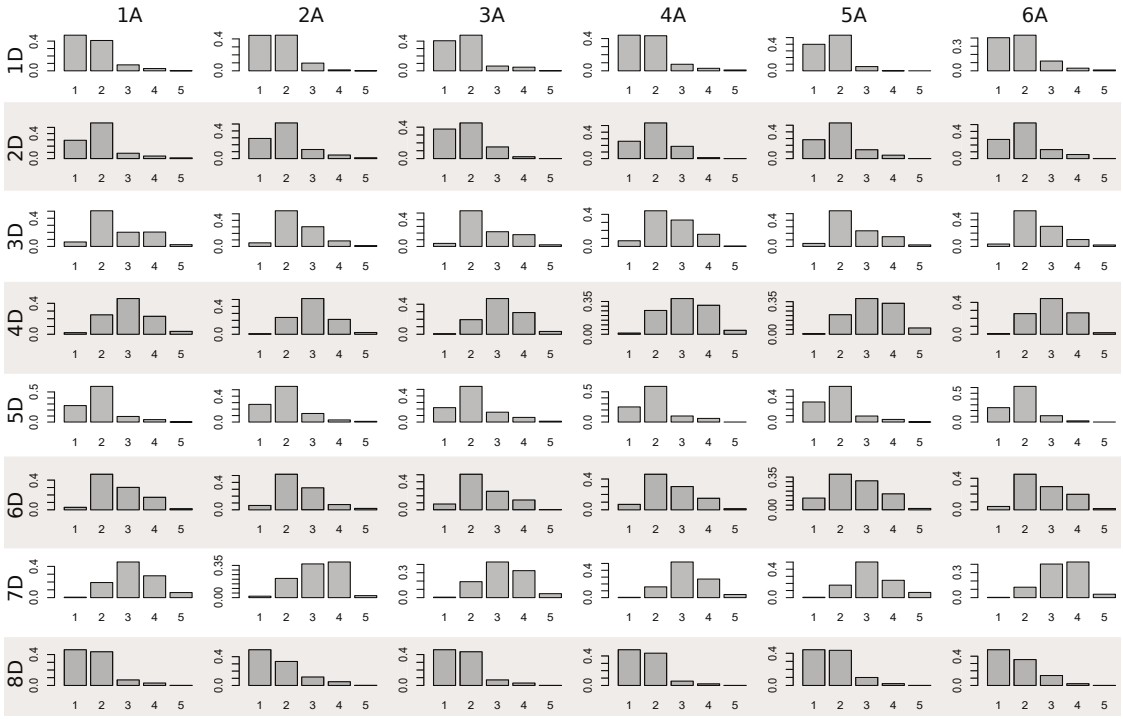

Figure 6.6: The 8×6 payoff matrix of the physical surveillance game with respect to "max comfort breach" [2]♮

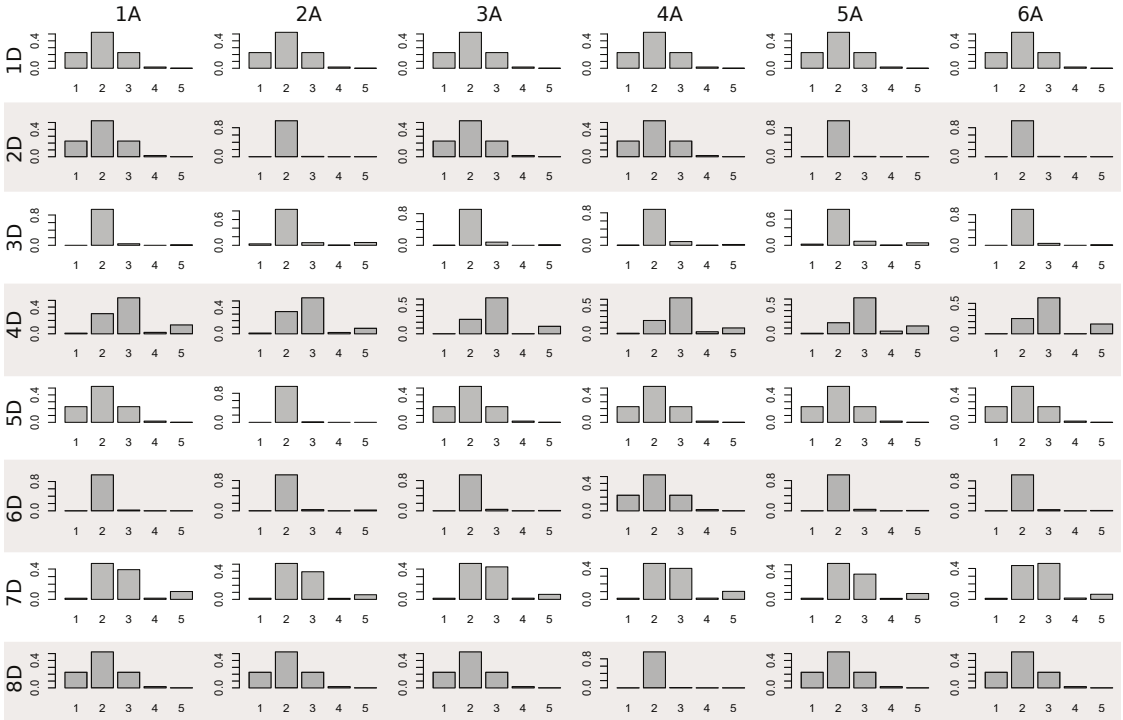

Figure 6.7: The 8×6 payoff matrix of the physical surveillance game with respect to "min privacy preservation" [2]♮

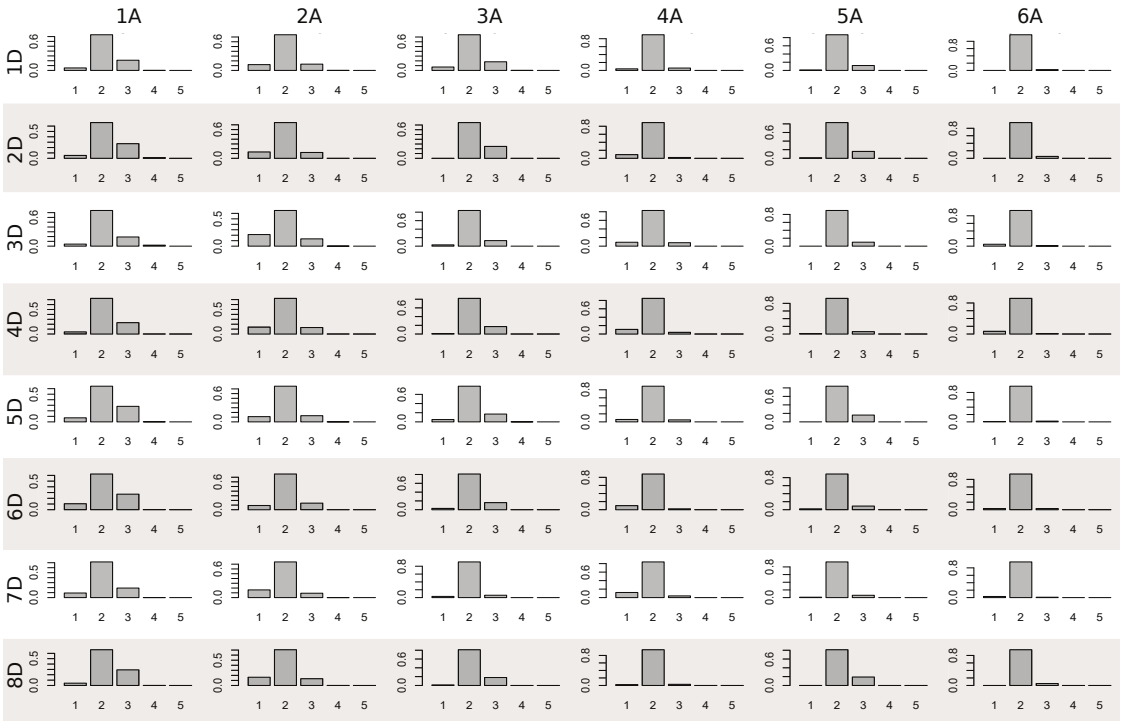

Figure 6.8: The 8×6 payoff matrix of the physical surveillance game with respect to "caused damage" [2]♮

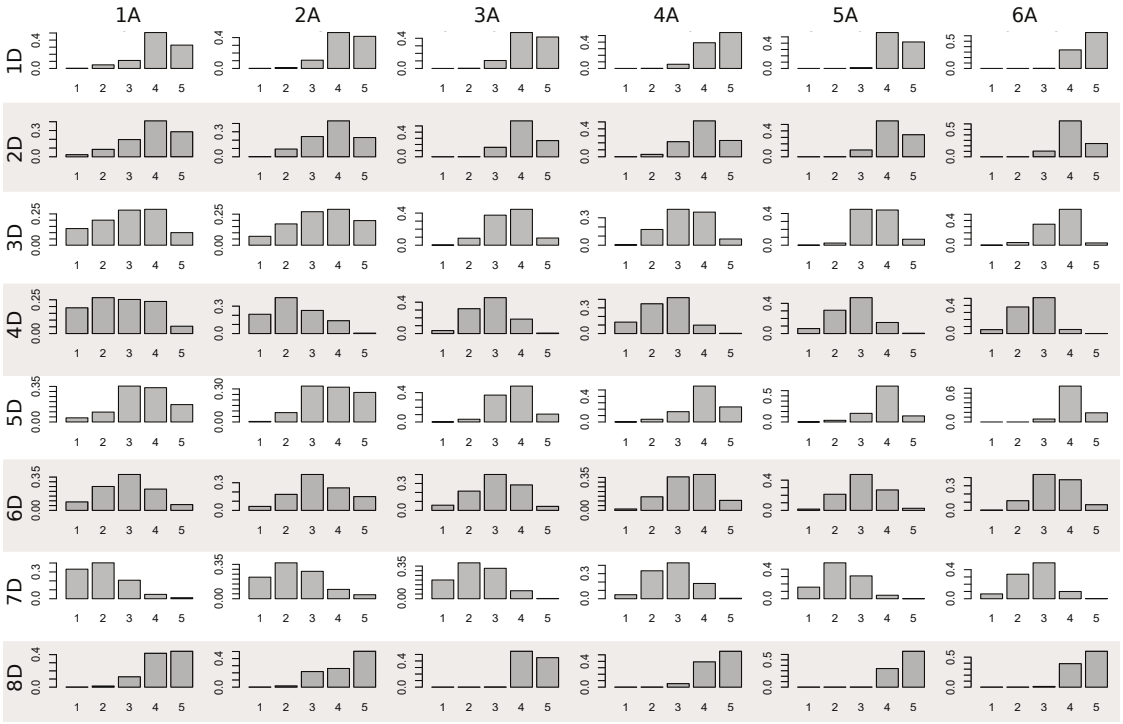

Figure 6.9: The 8×6 payoff matrix of the physical surveillance game with respect to "detection rate" [2]♮

Remark. As described in Section 4.1, the game model is designed for risk-based security management. Hence, it is generally assumed a minimizing player 1 (defender) and the

payoffs are thought of as losses. Therefore, any assessments relevant to a goal to be maximized have to be (sign-)changed before categorization to construct corresponding empirical loss distributions.

6.6 IDENTIFICATION OF BEST RESPONSE

This step involves mainly the process of constructing the game $G_{dist} = \langle \{\mathcal{D}, \mathcal{A}\}, \{SP_{\mathcal{D}}, SP_{\mathcal{A}}\}, (A^{(k)} \in \mathcal{F}^{8\times 6})_{k\in\{1,\dots,4\}}, \preceq \rangle$ using the output of the former steps and computing the security strategy (more specifically, the Pareto-efficient security strategy) as described in Section 4.2.5. In this work, all goals are taken as equally important (i.e., uniform weighting in the scalarization of the game when the equilibrium is computed $w_1 = w_2 = w_3 = w_4$ [120]). Figure 6.10 summarizes the game solution. It shows a nontrivial *optimal choice rule* describing how the defender should randomly choose among available pure actions; the obtained rule is called the mixed security strategy. In more detail, the game output dictates that best configurations can be reached when strategies 4D, 6D, and 7D are applied (or combined) with the respective probabilities 0.1, 0.768, and 0.132, while a practitioner can abandon other strategies. Further details on the implementation of the obtained mixed strategy follow in Chapter 7.

Concerning player 2 (the attacker), the game delivers a worst-case attack strategy for every goal, since multiobjective security games are solved by means of corresponding "one-against-all" competitions, where the defender plays against all (hypothetical) adversaries (cf. Section 4.2.5 and [187] for more details). These strategies are depicted in Figure 6.11, Figure 6.12, Figure 6.13, and Figure 6.14 together with the loss distributions expected by the defender if these worst-case attacks is applied and the defender follows his obtained security strategy illustrated in Figure 6.10. For instance, when thinking in terms of detection rate, the attacker can cause maximal loss (or equivalently lowest detection rate) by mainly choosing attack strategy 4A in which 5 intruders are moving randomly in the studied environment. It worth noting that the optimal attack strategies are different between the four goals, meaning that the attacker can never cause maximal (extreme) loss in all four goals at the same time. That is, the computed loss distributions over the four goals, as depicted in Figure 6.11, Figure 6.12, Figure 6.13, and Figure 6.14 are very pessimistic and reality should look much better (expectedly).

6.7 SUMMARY

In this chapter, physical surveillance games are illustrated by examining a case study of nuclear power plant considering a physical intrusion threat. That threat poses a great concern to such critical systems, no matter whether enacted by opportunistic or targeted attacks. To maximize the benefits of available security resources, the security management approach presented in Chapter 4 is applied throughout this chapter. After establishing the context of the target infrastructure, sets of attack and defense strategies are defined, and a simulation model is developed to assess the incurred risk on the basis of four goals. The game result is a nontrivial mixed security strategy. The results indicate that three identified defense strategies have to be combined while other strategies do not contribute to reducing the risk within the organization, given the identified attacks. Further details on the implementation of the obtained security strategy follow in Chapter 7.

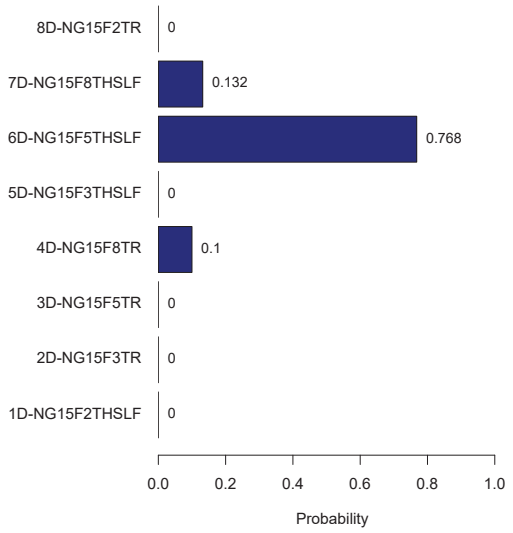

Figure 6.10: The mixed security strategy of G_{dist} [2]♮

Implications for general security resource sharing problems

In fact, concerns over risks like terrorism, crime, and business revenue loss impose the need for enhancing situational awareness inside the boundaries of CIs, with a considerably high level of persistent monitoring and surveillance as well as on-site observation activities. This is one of the key practical drivers presented by the proverbial power plant. Practices, such as conducting random patrols or regular spot-checks to prevent or deter potential violations, are strictly limited by (i) the number of available security resources, (ii) the uncertainty associated with patrol schedule execution as well as (iii) the ability of potential opponents to predict or observe the defenders' presence patterns [9]♮. The physical surveillance game model studied in Part II takes all these factors into account. Therefore, it provides practitioners with security strategies that keep a balance between several goals. Ultimately, the practitioners use the game advice to generate (randomized) schedules of actions that assure the optimal frequencies.

Regarding other application domains, the presented framework can be further applied to ensure security and business continuity of airports/public transportation systems that requires the physical presence of security staff to mitigate and control accidents and to perform law enforcement functions. These systems are at significant risk of crime since they gather large quantities of people in both time and space. In [9]♮, the author of this thesis showed how to carry out spot checks in public transportation systems to address risks facing proof-for-payment-based transportation systems, namely risks of fare evasion. A physical surveillance game is constructed to model the interactions between the involved competitive entities (i. e., inspectors/security officials and criminals/fare dodgers). The model enables integrating measurements from heterogeneous natures (e. g., statistics, expert opinions, or simulation results) towards finding optimal cost-effective fare-enforcement plans.

In more general terms, security resource allocation problems can be approached as follows: One can divide available resources into meaningful units and suppose some practically feasible allocations to define the gameplay. This would correspond to the aforementioned "defense proposals" or "educated guesses" performed by the involved

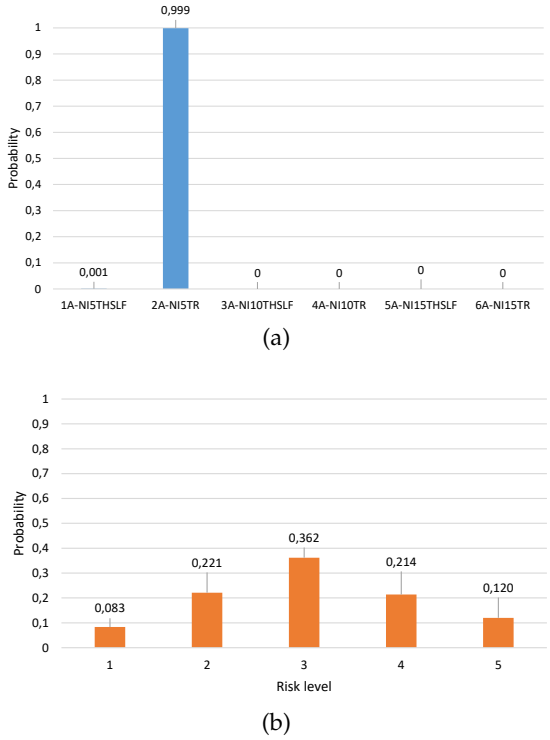

Figure 6.11: G_{dist} (a) Worst-case attack for the goal "detection rate"; (b) maximum assured loss with respect to the goal "detection rate"

security experts. For instance, this could mean the number of sensors placed in an energy grid to detect anomalies. To this end, one would first divide the energy network into (logically) meaningful areas, e.g., network segments, physical areas, or similar, in which a number of anomaly-detection sensors can be placed. Then, the game can be used to refine an allocation by simulating attacks under the so-constructed strategies, to get an optimized mix of the strategies, each of which corresponds to the placement of, say n_k sensors into area k. The decision-maker can then "quasi-purify" the mix by relocating sensors from one area to the other, according to the equilibrium in the game [2]♮. The latter point is thoroughly discussed in Chapter 7.

The simulation model introduced in this chapter has been further employed to optimize surveillance infrastructures of cyber-physical systems against coordinated APT-like attacks, as shown in [13]♮. The analysis results provide fundamental principles to design defense-in-depth mechanisms that yield security strategies to protect CIs [13]♮.

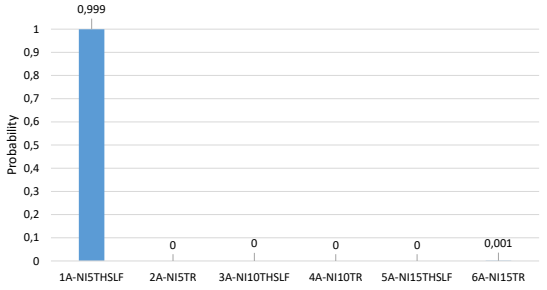

(a)

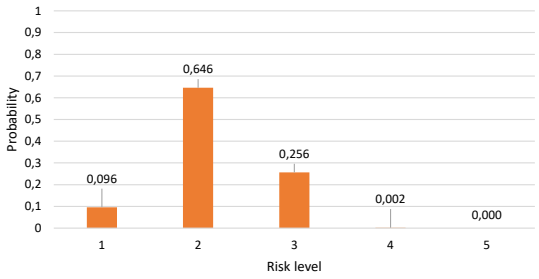

(b)

Figure 6.12: G_{dist} (a) Worst-case attack for the goal "caused damage"; (b) maximum assured loss with respect to the goal "caused damage"

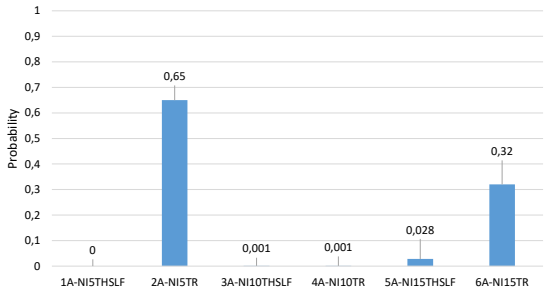

(a)

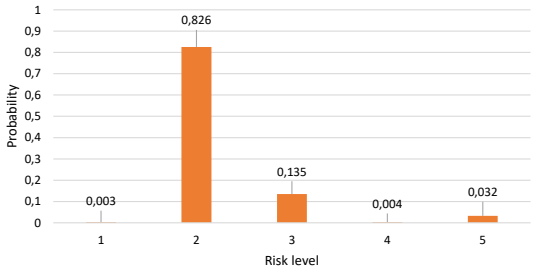

(b)

Figure 6.13: G_{dist} (a) Worst-case attack for the goal "minimum privacy preservation"; (b) maximum assured loss with respect to the goal "minimum privacy preservation"

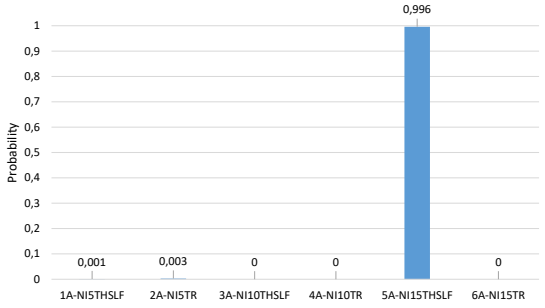

(a)

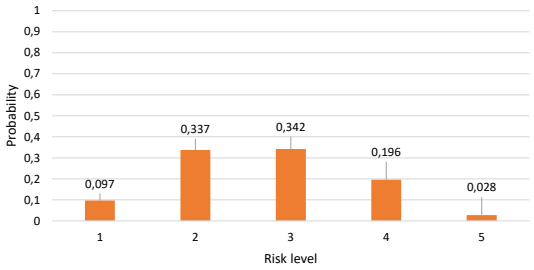

(b)

Figure 6.14: G_{dist} (a) Worst-case attack for the goal "maximum comfort breach"; (b) maximum assured loss with respect to the goal "maximum comfort breach"

7 EVALUATION AND COMPARATIVE ANALYSIS

The content of Chapter 7 is based on the research work published in [2]♮.

7.1 INTRODUCTION

Security resources, such as security personnel and surveillance devices, are scarce and usually expensive. To address this challenge, Chapter 6 explains how a defender can choose among several possible resource-allocations, and relies on game-theory for an optimal choice. For the sake of evaluation, Figure 7.1 illustrates the methodology applied to achieve a better understanding of the differences between the new class of generalized games in which payoffs are probability distributions and its counterpart of classical (real-valued) games. The evaluation process consists of the following steps:

1. Constructing a classical game $G_{classical}$ of the presented surveillance problem
2. Converting the game solutions into consistent security resource allocations using the quasi-purification process
3. Assessing the effectiveness of the purified strategies
4. Analyzing the results across the following evaluation dimensions:
   - Mixed-strategy extended game
   - Closeness to the ideal point
   - Graphical comparison
   - Disappointment rate

The individual steps are described in detail in the following sections.

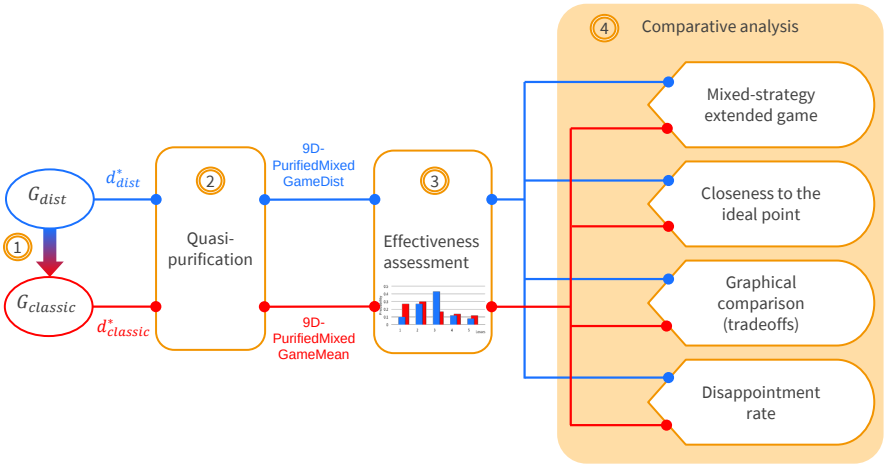

Figure 7.1: Evaluation methodology for resource allocation based on game-theoretical advice

7.2 CLASSICAL GAME MODEL

The first step is to construct a classical zero-sum game of the physical surveillance problem presented in Chapter 6 (more details on classical games can be found in Section 4.1). Let us refer to the classical game as $G_{classical} = \langle\{\mathcal{D}, \mathcal{A}\}, \{SP_{\mathcal{D}}, SP_{\mathcal{A}}\}, (A^{(k)} \in \mathbb{R}^{8\times 6})_{k\in\{1,\dots,4\}}, \leqslant\rangle$. That means:

- The action sets of the defender $\mathcal{D}$ and the attacker $\mathcal{A}$ are those defined in Table 6.2, which are $SP_{\mathcal{D}}$, $SP_{\mathcal{A}}$, respectively.
- The game goals are as defined in Section 6.4: caused damage, minimum privacy preservation, maximum comfort breach, and detection rate.
- The payoff matrices $A^{(k)}$, one for each goal, include scalar-valued payoffs, which are the arithmetic mean values of the distribution-valued payoffs in the corresponding matrices of G_{dist}. In this regard, Table 7.1, Table 7.2, Table 7.3, and Table 7.4 show the payoff structures of $G_{classical}$ with respect to maximum comfort breach, minimum privacy preservation, caused damage, and detection rate, respectively.
- The game is played over a numeric $\leqslant$-order, and the players are indifferent between choices with equal expected payoffs (losses) even if one choice is riskier.

The security strategy of $G_{classical}$ is illustrated in Figure 7.2. Analogous to G_{dist}, $G_{classical}$ delivers a mixed defense strategy. The mixed strategy of $G_{classical}$ is, however, different from the one of G_{dist}. Following $G_{classical}$, the defender should randomly choose between 5D and 6D strategies according to the probabilities 0.534 and 0.466, respectively. In contrast, G_{dist} has an equilibrium when strategies 4D, 6D, and 7D are combined, as illustrated in Figure 6.10. Although both games are basically constructed using the same underlying assessment results, they delivered two different mixed strategies. Therefore, it is important to investigate which game model is better especially with respect to defender satisfaction.

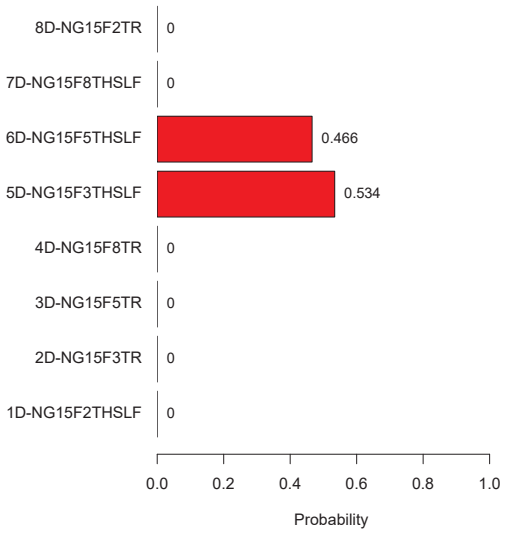

Figure 7.2: The mixed security strategy of $G_{classical}$ [2]

Table 7.1: The 8×6 payoff matrix of $G_{classical}$ w.r.t. "max comfort breach" [2]♮

Strategy No.	1A	2A	3A	4A	5A	6A
1D	1.64	1.66	1.74	1.71	1.66	1.79
2D	1.89	1.96	1.8	1.94	1.94	1.96
3D	2.62	2.45	2.6	2.57	2.56	2.54
4D	3.02	3.01	3.16	3.12	3.25	3.03
5D	1.89	1.94	2.09	1.97	1.85	1.9
6D	2.65	2.47	2.47	2.57	2.56	2.69
7D	3.21	3.19	3.22	3.2	3.21	3.38
8D	1.65	1.7	1.65	1.6	1.68	1.67

Table 7.2: The 8×6 payoff matrix of $G_{classical}$ w.r.t. "min privacy preservation" [2]♮

Strategy No.	1A	2A	3A	4A	5A	6A
1D	2	2	2	2	2	2
2D	2	2.01	2	2	2.01	2.01
3D	2.07	2.24	2.11	2.15	2.25	2.08
4D	2.98	2.84	3.01	3	3.13	3.07
5D	2	2.01	2	2	2	2
6D	2.02	2.09	2.07	2	2.07	2.06
7D	2.73	2.61	2.65	2.74	2.64	2.69
8D	2	2	2	2.01	2	2

Table 7.3: The 8×6 payoff matrix of $G_{classical}$ w.r.t. "caused damage" [2]♮

Strategy No.	1A	2A	3A	4A	5A	6A
1D	2.16	2.01	2.11	2.02	2.11	2.02
2D	2.24	1.99	2.25	1.93	2.15	2.05
3D	2.19	1.92	2.1	1.99	2.1	1.96
4D	2.19	1.99	2.16	1.93	2.05	1.94
5D	2.21	2.02	2.12	1.99	2.16	2.01
6D	2.17	2.05	2.13	1.92	2.07	2
7D	2.1	1.93	2.03	1.92	2.05	1.98
8D	2.25	1.97	2.17	2.01	2.19	2.05

Table 7.4: The 8×6 payoff matrix of $G_{classical}$ w.r.t. "detection rate" [2]♮

Strategy No.	1A	2A	3A	4A	5A	6A
1D	4.13	4.3	4.32	4.49	4.41	4.66
2D	3.89	3.81	4.11	3.96	4.24	4.14
3D	3	3.38	3.54	3.32	3.57	3.61
4D	2.65	2.32	2.8	2.48	2.71	2.57
5D	3.51	3.8	3.67	4	3.89	4.13
6D	2.91	3.29	3.04	3.43	3.08	3.39
7D	1.97	2.36	2.29	2.76	2.24	2.63
8D	4.31	4.28	4.45	4.51	4.66	4.59

Table 7.5: Quasi-purification (parameterization) of mixed strategies [2]♮

Prob. Dist.	Strat. No.	#Guards (NG)					Add. parameters: (F)req. & (T)arget
		Quota	Q_I	%	Q_F	Total $Q_I + Q_F$	
9D-PurifiedMixedGameDist							
76.8%	6D	11.52	11	0.52	1	12	F = 5 & T: HSLF
13.2%	7D	1.98	1	0.98	1	2	F = 8 & T: HSLF
10%	4D	1.5	1	0.5	0	1	F = 8 & T: R
9D-PurifiedMixedGameMean							
53.4%	5D	8.01	8	0.01	0	8	F = 3 & T: HSLF
46.6%	6D	6.99	6	0.99	1	7	F = 5 & T: HSLF

7.3 QUASI-PURIFICATION AND EFFECTIVENESS ASSESSMENT

The security strategies depicted in Figure 6.10 and Figure 7.2 aim at assisting the defender in identifying the optimal defense plan. However, if there is no purely optimal defense action delivered, then the definition of a new defense action is possible through proper handling of the achieved (mixed) defense strategy; this work calls this process "*quasi-purification*"[1] of mixed strategies. While this process might not be possible in some models, the specific characteristic of the identified defense strategies that they are expressed in terms of parameters, facilitate the quasi-purification through a proper re-parametrization of the pure defense strategies. That is, if the resources spent on different pure strategies are "non-atomic" (partly transferable) and can be reassigned to play multiple pure strategies at the same time by proper resource-sharing, the game can be extended through adding the mixed security strategy in pure form.

Throughout this work, a quasi-purified mixed strategy is obtained through adjusting the amount of security resources of each pure strategy (number of security guards, being a *parameter* there) according to the probability distribution prescribed by the security strategies of the respective games. It is worth noting that the process of allocating security resources of the quasi-purified mixed strategy is derived from "the Hamilton method" (aka the largest remainder method) [69]: suppose that the pure strategy $d_i \in \mathrm{SP}_{\mathcal{D}}$ (this work adopt the position of the first player, the defender, here, w.l.o.g.), being an integer assignment of security guards, would be played with an optimized likelihood p_i as found in $\delta^* \in \Delta(\mathrm{SP}_{\mathcal{D}})$. The optimal assignment would then come to $p_i \cdot d_i$, which can be fractional, i.e., prescribe a non-integer number of guards. The Hamilton method resolves this by assigning the integer part of $\lfloor p_i \cdot d_i \rfloor$ first, thus leaving some resources unallocated for the moment. These unused guards are then assigned to those strategies having the maximal fractional remainder $p_i \cdot d_i - \lfloor p_i \cdot d_i \rfloor$, sorted in descending order for $i = 1, 2, \ldots, |\mathrm{SP}_{\mathcal{D}}|$, each of which getting one further unit (guard) assigned, until the remainder of unallocated resources is used up. The whole method is illustrated in Table 7.5, which explains a possible definition of the new defense strategies of the classical and distribution-valued games, called "9D-PurifiedMixedGameMean" and "9D-PurifiedMixedGameDist", respectively.

1 The term quasi-purification was chosen to avoid confusion with the concept of purification, as put forth in Harsanyi's celebrated purification theorem [83].

Alternatively, yet not further explored in this work, one can consider a part-time assignment of resources to one strategy or another. For example, if d_i and d_j both have a 0.5 fractional part, then a guard can spend half its time acting out action d_i and the other half of its time doing d_j, thus effectively realizing the "0.5" allocation of itself as a resource. For other contexts, the division into fractions can be more straightforward, such as computing time and cloud resources. Such a part-time allocation may, however, be impractical, e.g., if the effort of changing configurations (switching the team, commuting between different workplaces, or similar) is too high. Game models, with a few exceptions [170], usually do not account for this kind of cost.

Afterward, the developed simulation framework presented in Section 6.5 is leveraged to assess the effectiveness of the quasi-purified strategies against the identified attacker strategies with respect to the four specified goals. The results have been collected from 100 runs for each scenario and transformed into suitable payoff structures of each game model (i.e., scalar and distribution-valued payoffs).

7.4 A COMPARATIVE ANALYSIS

Before delving into the possibilities to assess the effectiveness of the game-theoretical outcomes, it is of vital importance to look at the interpretation of these outcomes. The two apparent interpretations, as discussed by [231], are the descriptive interpretation and the normative interpretation. The latter is concerned with providing the players with advice on how to act better in game-similar situations, while the former is mainly concerned with predicting the actual players' behaviors. Since the normative interpretation properly fits the settings of security management games, this chapter seeks to examine whether the game-theoretical advice helps the defender to make better decisions than what he might otherwise have made (referring to the defender's pure strategies shown in Table 6.2). Furthermore, it will examine which decision would better satisfy the expectations of the defender [2]$^\natural$. The analysis pursued in this chapter examines four evaluation dimensions towards achieving some evidence on the quality of the delivered game-theoretical advice.

7.4.1 *First dimension: mixed-strategy-extended games*

If the optimum exists only in randomized strategies,then the defender needs to "purify" the resource assignment as described in Section 7.3, hoping to retain the best protection [2]$^\natural$. That is, the new defense actions defined in Table 7.5 should outperform all previous defenses (cf. Table 6.2). Utilizing game theory principles, the validity of this expectation is studied experimentally.

Therefore, a distribution-valued extended game $G_{distExt}$ is constructed by adding the newly defined and empirically assessed strategy "9D-PurifiedMixedGameDist" to the defender's action space of G_{dist} and its assessment results into the distribution-valued payoff matrices as depicted in Figure 7.3, Figure 7.4, Figure 7.5, and Figure 7.6.

Similarly, a classical extended game $G_{classicalExt}$ is constructed through integrating the strategy "9D-PurifiedMixedGameMean" into the original $G_{classical}$. The 9×6 payoff matrices of $G_{classicalExt}$ are depicted in Table 7.6, Table 7.7, Table 7.8, and Table 7.9.

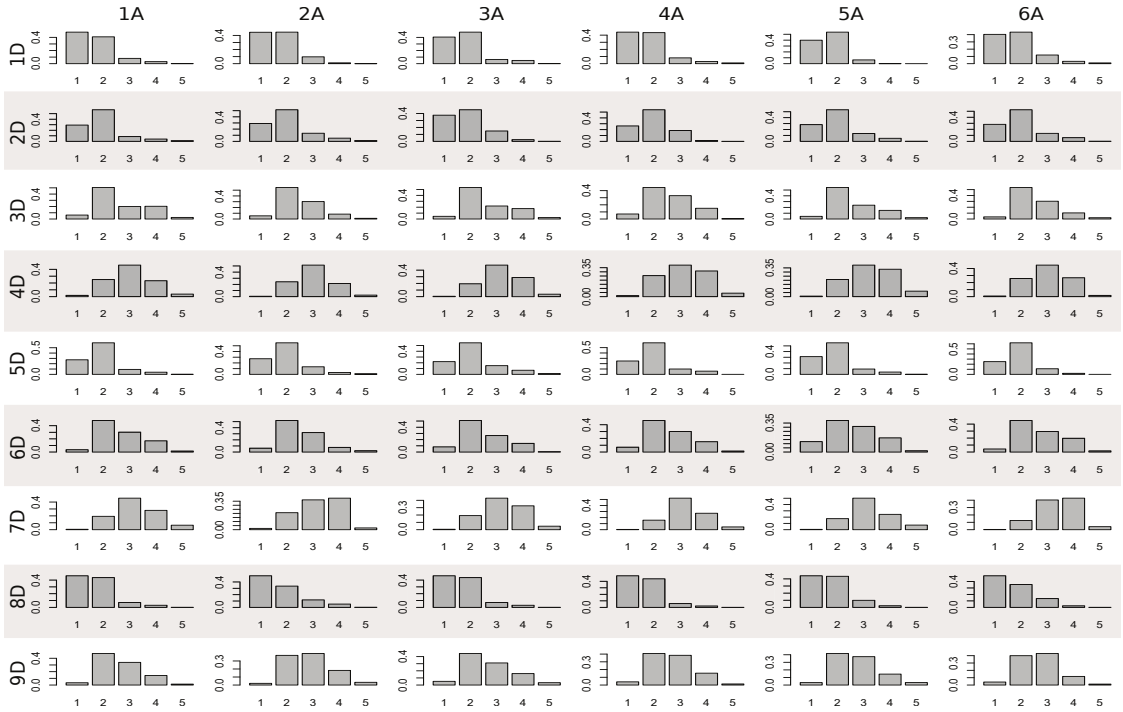

Figure 7.3: The 9×6 payoff matrix of G_{distExt} w.r.t. "max comfort breach" [2]♮

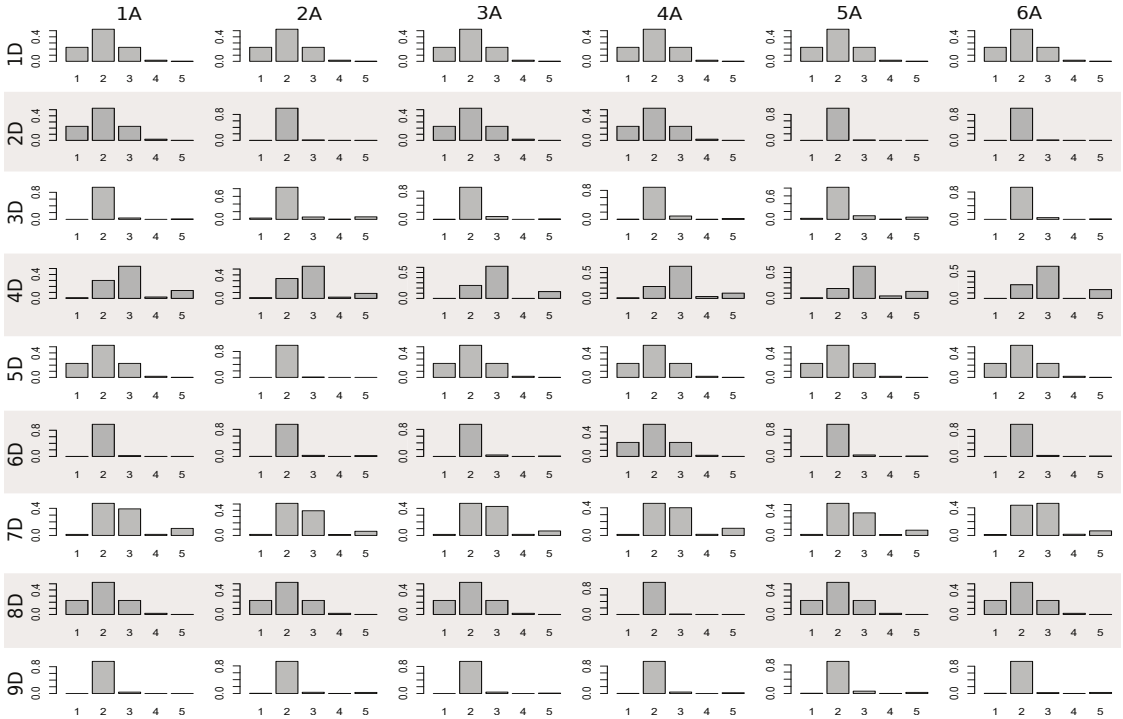

Figure 7.4: The 9×6 payoff matrix of G_{distExt} w.r.t. "min privacy preservation" [2]♮

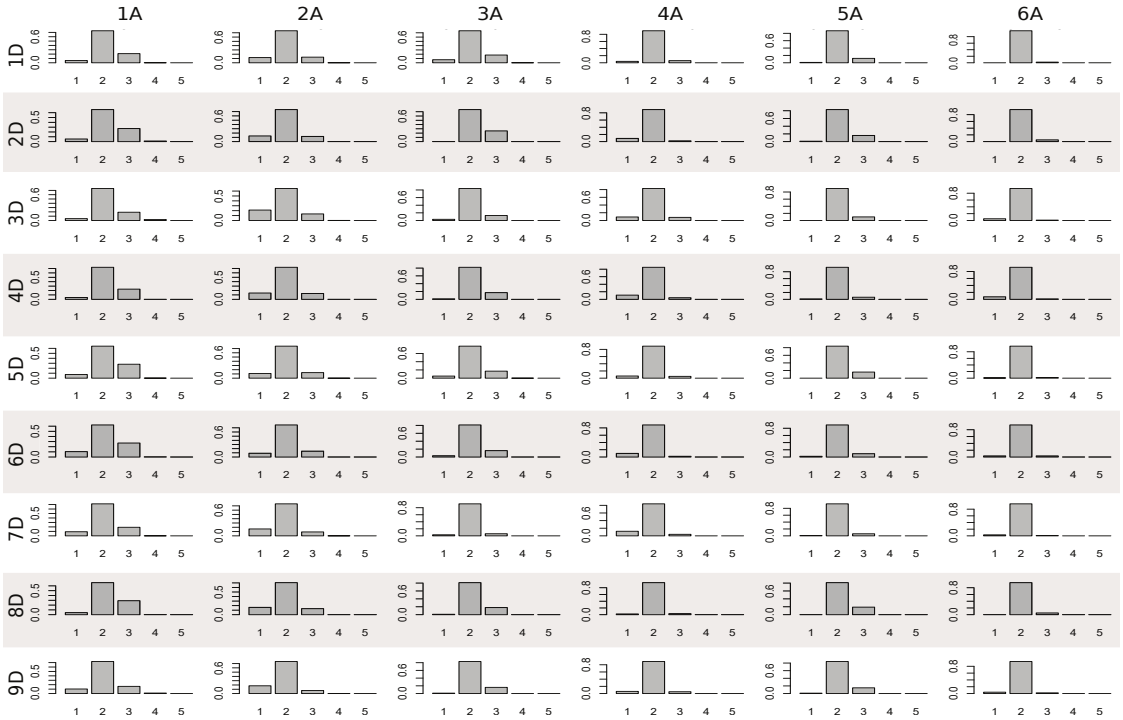

Figure 7.5: The 9×6 payoff matrix of G_{distExt} w.r.t. "caused damage" [2]♮

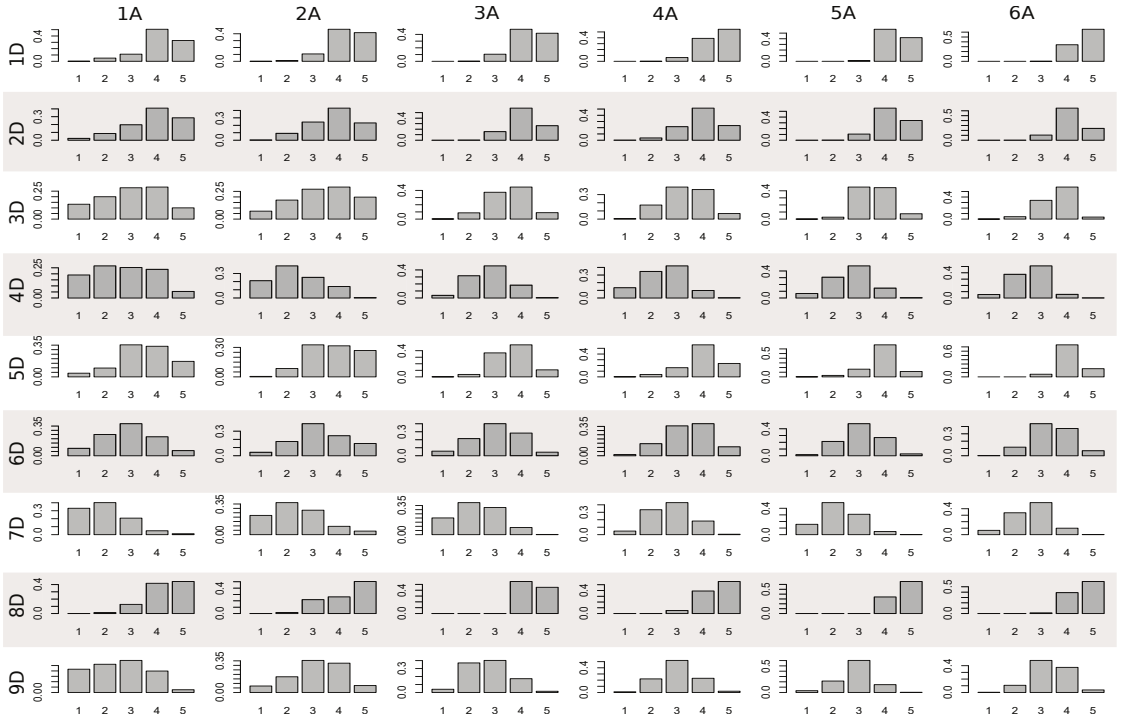

Figure 7.6: The 9×6 payoff matrix of G_{distExt} w.r.t. "detection rate" [2]♮

Table 7.6: The 9×6 payoff matrix of $G_{classicalExt}$ w.r.t. "max comfort breach" [2]♮

Strategy No.	1A	2A	3A	4A	5A	6A
1D	1.64	1.66	1.74	1.71	1.66	1.79
2D	1.89	1.96	1.8	1.94	1.94	1.96
3D	2.62	2.45	2.6	2.57	2.56	2.54
4D	3.02	3.01	3.16	3.12	3.25	3.03
5D	1.89	1.94	2.09	1.97	1.85	1.9
6D	2.65	2.47	2.47	2.57	2.56	2.69
7D	3.21	3.19	3.22	3.2	3.21	3.38
8D	1.65	1.7	1.65	1.6	1.68	1.67
9D	2.36	2.29	2.31	2.29	2.02	2.46

Table 7.7: The 9×6 payoff matrix of $G_{classicalExt}$ w.r.t. "min privacy preservation" [2]♮

Strategy No.	1A	2A	3A	4A	5A	6A
1D	2	2	2	2	2	2
2D	2	2.01	2	2	2.01	2.01
3D	2.07	2.24	2.11	2.15	2.25	2.08
4D	2.98	2.84	3.01	3	3.13	3.07
5D	2	2.01	2	2	2	2
6D	2.02	2.09	2.07	2	2.07	2.06
7D	2.73	2.61	2.65	2.74	2.64	2.69
8D	2	2	2	2.01	2	2
9D	2	2.01	2	2	2.03	2

Table 7.8: The 9×6 payoff matrix of $G_{classicalExt}$ w.r.t. "caused damage" [2]♮

Strategy No.	1A	2A	3A	4A	5A	6A
1D	2.16	2.01	2.11	2.02	2.11	2.02
2D	2.24	1.99	2.25	1.93	2.15	2.05
3D	2.19	1.92	2.1	1.99	2.1	1.96
4D	2.19	1.99	2.16	1.93	2.05	1.94
5D	2.21	2.02	2.12	1.99	2.16	2.01
6D	2.17	2.05	2.13	1.92	2.07	2
7D	2.1	1.93	2.03	1.92	2.05	1.98
8D	2.25	1.97	2.17	2.01	2.19	2.05
9D	2.17	2.1	1.94	1.99	2.06	2.01

Table 7.9: The 9×6 payoff matrix of $G_{classicalExt}$ w.r.t. "detection rate" [2]♮

Strategy No.	1A	2A	3A	4A	5A	6A
1D	4.13	4.3	4.32	4.49	4.41	4.66
2D	3.89	3.81	4.11	3.96	4.24	4.14
3D	3	3.38	3.54	3.32	3.57	3.61
4D	2.65	2.32	2.8	2.48	2.71	2.57
5D	3.51	3.8	3.67	4	3.89	4.13
6D	2.91	3.29	3.04	3.43	3.08	3.39
7D	1.97	2.36	2.29	2.76	2.24	2.63
8D	4.31	4.28	4.45	4.51	4.66	4.59
9D	3.26	3.3	3.61	3.68	3.41	3.84

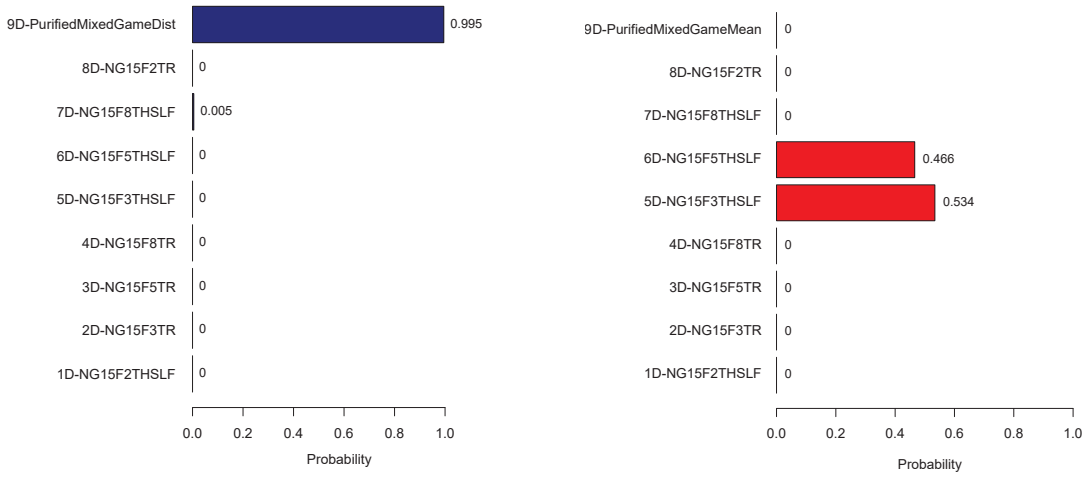

(a) $G_{distExt}$ (b) $G_{classicalExt}$

Figure 7.7: The optimal defense strategies of (a) distribution-valued and (b) classical extended games; validation of the purified mixed strategies 9D-PurifiedMixedGameDist and 9D-PurifiedMixedGameMean [2]♮

Indeed, one would expect the new defense strategy to be a trivial equilibrium in the resulting game. As expected, this will hold for the distribution-valued game, but quite surprisingly, fail in the classical one, as depicted in Figure 7.7. More precisely, while the new strategy "9D-PurifiedMixedGameDist" is the most effective defense action with almost 99.5% over all other defense strategies of the distribution-valued extended game $G_{distExt}$, the comparable "9D-PurifiedMixedGameMean" does not appear at all in the security strategy of the classical extended game $G_{classicalExt}$. Therefore, the results shown in Figure 7.7 deliver the first evidence that distribution-valued game solutions can outperform their classical counterparts (at least in this application context). Furthermore, Figure 7.7 can give the impression that distribution-valued games are more accurate than classical games in predicting what could happen in a realistic environment. Such an interpretation would definitely demand further investigations.

7.4.2 *Second dimension: distance measures (closeness to the ideal point)*

The presented physical surveillance games involve the consideration of 4 goals that need to be optimized simultaneously (cf. Section 6.4). These goals may indeed be conflicting (in the sense of a negative correlation). As described in Section 6.5.2 and without loss of generality, the payoffs of the examined games are presented as losses, and player 1 is technically minimizing all goals. Given this setting, each defense action $d_i \in SP_{\mathcal{D}}$ can be represented in the objective space using a best (in the sense of minimal) expected loss vector $f(d_i) = (f_1(d_i), \ldots, f_4(d_i))$, where

$$f_k(d_i) = \min_{a_j \in SP_{\mathcal{A}}} u_{1,k}(d_i, a_j) \quad \forall k \in \{1, \ldots, 4\}$$

Then, the outcome of the classical and distribution-valued game can be evaluated by comparing the distance to a reference point (vector) in the objective space. This reference point represents the aspiration level of the defender (i.e., the desired objective values). The preferred solution is the most satisfactory one depending on the defender's

expectation and, therefore, the closest one to the reference point. Throughout this work, the reference point is recognized as the ideal (utopia) point that would be achieved by optimizing (here, minimizing) each objective individually [78, 123]. Thus, the ideal point is represented by a vector $\alpha = (\alpha_1, \ldots, \alpha_4)$, where

$$\alpha_k = \min_{d_i \in SP_{\mathcal{D}}, a_j \in SP_{\mathcal{A}}} u_{1,k}(d_i, a_j) \quad \forall k \in \{1 \ldots 4\}$$

The ideal point is in general unattainable if the goals are negatively correlated (i. e., conflicting), which is the case here, as explained in Section 6.4. Given the assessment results of the studied scenario, the ideal point is $\alpha = (1, 2, 1, 1)$, obtained by optimizing each goal in isolation, i. e., independent of all other goals. This is the best outcome possible for each goal, *disregarding* all existing interdependencies between goals.

Afterward, the distance between the consequence of the defender's action d_i (the actual outcome) and the reference point α (the ideal outcome) is computed using the L_p-norm on $\mathbb{R}^4$ (Minkowski distance):

$$d_{L^p}(d_i, \alpha) = \|f(d_i) - \alpha\|_p = \left(\sum_k |f_k(d_i) - \alpha_k|^p \right)^{\frac{1}{p}}, \text{ with } p \in \{1, 2, \infty\} \tag{7.1}$$

Well-known norms can be obtained for different values of p. The case for $p = 1$ is the Manhattan distance, $p = 2$ is the Euclidean distance, and $p = \infty$ is the Chebyshev distance, which measures the maximum component. The latter is perhaps the easiest to interpret in practice as being a uniform worst-case measure for all components. Figure 7.8a and 7.8b depict the evaluation results of the different mixed and pure defense strategies in terms of the three distance metrics described in Equation (7.1).

Figure 7.8a shows that the outcome of the mixed defense strategy delivered by the distribution-valued game is closer to the ideal point α than the outcome of the classical game mixed strategy. That is, the distribution-valued game's advice delivered a more satisfactory decision for a realistic world in terms of the aspired objective values. Furthermore, Figure 7.8b confirms this finding by presenting the distances of the outcomes of the 8 pure strategies (cf. Table 6.2) as well as the two quasi-purified mixed strategies to the ideal point. It shows that the distribution-valued game's defense strategy, in contrast to classical game strategy, is the closest under all distances.

7.4.3 *Third dimension: graphical comparison*

To correlate and contrast the performance of the two quasi-purified mixed strategies over the four identified goals, a graphical comparison method is pursued that involves representing the consequences of the strategies using a radar chart as depicted in Figure 7.9a. The technique of radar chart provides esstential insights into the outcomes of the strategies with respect to the *individual* goals [62], thereby illustrating the strategies' similarities and differences on multiple variables (i.e., goals) in a single two-dimensional graph [185]. The results in Figure 7.9a reveal that no strategy is better than another in every goal. That is, the investigated game models behave differently from each other in achieving tradeoffs among the different goals. Thus, the areas of the polygons depicted by each strategy's loss vector are leveraged to reveal the differences between the strategies; smaller areas indicate better tradeoff among the various goals. Figure 7.9b

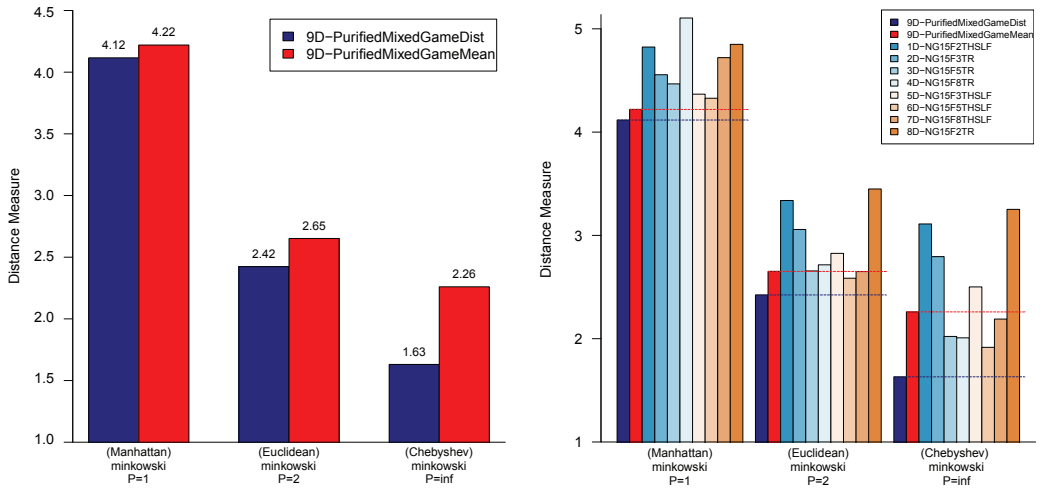

Figure 7.8: Distance to the ideal point of (a) quasi-purified mixed strategies of both classical and distribution-valued games; (b) 8 pure defense strategies as well as 9D-PurifiedMixedGameDist and 9D-PurifiedMixedGameMean [2][♮]

shows that the distribution-valued game strategy has slightly less area (better fit) than classical game strategy and, therefore, more preferable by the defender[2]. Furthermore, the findings depicted in Figure 7.9b give another evidence that distribution-valued games may achieve better tradeoffs than classical games and the pure strategies (defined in Table 6.2) do in this context.

7.4.4 *Fourth dimension: disappointment rate*

Besides the utility (loss) satisfaction measured by the distance metrics introduced in Section 7.4.2, the analysis process examines here a measure of psychological satisfaction experienced by the defender after making decisions in the sense of disappointment. This idea dates back to the 1980s already (see [44, 45] and references therein), but is clearly a natural measure of "goodness" of an equilibrium. People may feel disappointed when the actual impact of their chosen action is worse than their prior expectations. Therefore, they tend to anticipate the potential disappointment and make their decisions towards minimizing it [26]. In the presented security games, if a player knows an average loss $E(u_i)$[3], then this expected value may truly be what player i *expects* to get in every round. However, as being an average, there will be rounds that cause a loss $< E(u_i)$ and other rounds where the loss is $> E(u_i)$. The latter would disappoint the player (as the game seems to do not keep the promise), and player i could act towards minimizing the disappointment rate $d = \Pr(u_i > E(u_i))$ too, besides minimizing losses. For security, the expectation is what decision-makers prepare for, say by buying insurance, and the disappointment rate tells the chance for this preparation to be "too weak", i.e.,

2 Having several polygons in a radar chart makes it difficult to read. Therefore, only the mixed strategies are placed in Figure 7.9a, while Figure 7.9b shows the polygon areas of the entire set of the defender's strategies described in Table 6.2

3 For randomized actions x, y and a payoff matrix A, the expected loss in both games is defined in the same way: $E(u_i) = \mathbf{x}^T A \mathbf{y}$ as stated in Section 4.1.2.2.

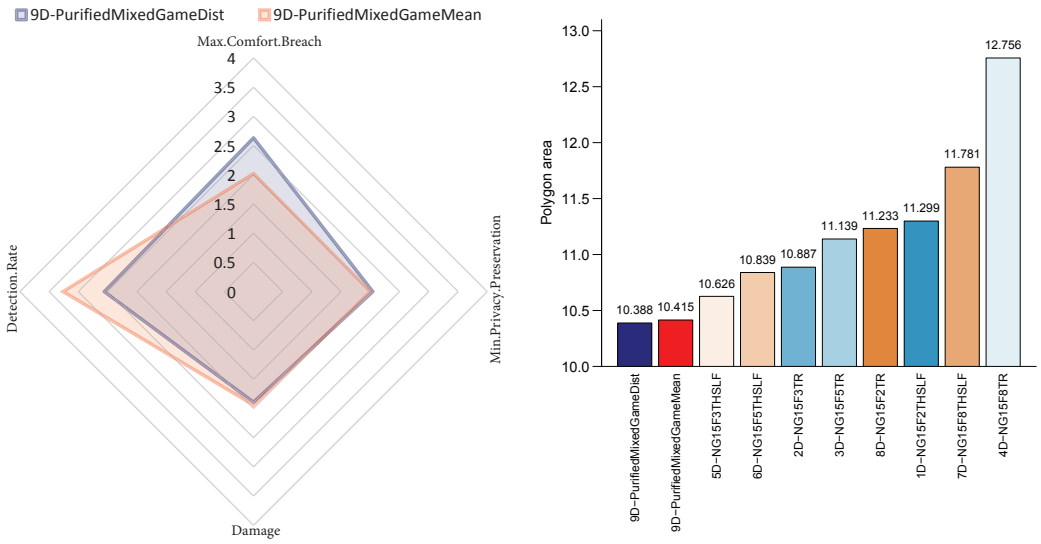

(a) (b)

Figure 7.9: (a) Radar chart of the consequences of the mixed strateges "9D-PurifiedMixedGameMean" and "9D-PurifiedMixedGameDist"; (b) The poygon areas of all defense strategies when their assessment results are represented using the radar chart [2]♮

the preparation is insufficient and the damage could be irrecoverable. Irrespectively, subjective disappointments may incentivize the defender to deviate from the optimal defense policy, thus having another weakening effect on the overall security. For these two reasons, disappointment rates appear highly relevant in security applications.

Due to the negative impact of disappointment on the satisfaction level of the chosen security measure, it is clearly a natural measure of "goodness" of the security strategies in the studied games. From this perspective, the game outcome that gives rise to a higher disappointment rate is less favorable. In distribution-valued games played over the stochastic order $\preceq$, the outcome (equilibrium) random loss X^* directly tells us the disappointment rate as $d = \Pr(X^* > E(X^*))$, which can be computed from the optimal loss distribution that the game delivers directly (by construction). Classical games do not deliver this information (as being focused on the expectation only), and an estimate of d would have to be approximated by the law of large numbers and running many repetitions of the game (counting the disappointing rounds).

Figure 7.10a depicts promises of G_{dist} and $G_{classical}$ games, which are the expected values of the four considered goals if the defender plays the optimal defense strategies depicted in Figure 6.10 and Figure 7.2. As shown, both game models give approximately close promises with respect to privacy preservation and damage goals, while they show higher differences with respect to comfort breach and detection rate. To compare the impact of the decisions of both games, Figure 7.10b depicts the disappointment rates resulted from the actual performance of 9D-PurifiedMixedGameDist and 9D-PurifiedMixedGameMean against the 6 attacker strategies and with respect to the 4 individual goals. The results show that the disappointment rate resulted from G_{dist} is equal or less than the one resulted from $G_{classical}$ over the respective goals. When the rates of both privacy preservation and damage are the same, this tells us that both games share the same number of disappointing scenarios with respect to these two

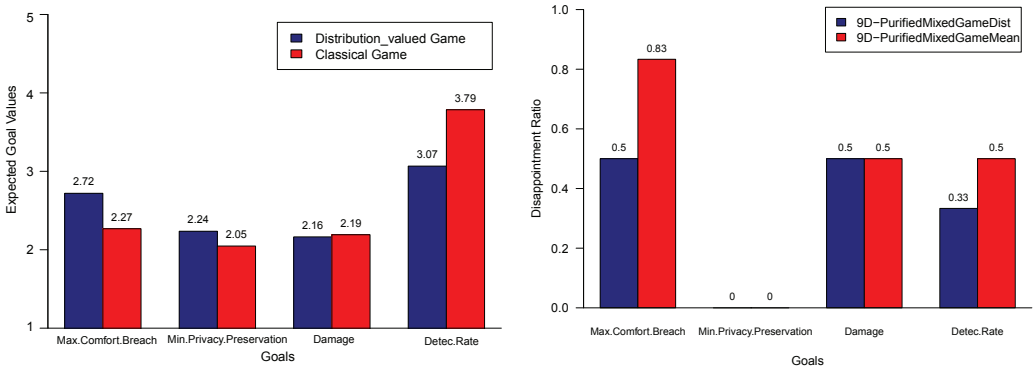

(a) (b)

Figure 7.10: (a) The expected goal values if the defender adheres to optimal defense strategies: distribution-valued game vs. classical game; (b) Disappointment rate: distribution-valued game vs. classical game [2]♮

goals. Otherwise, the distribution-valued game led to less disappointment. One reason for this difference may be that in playing the game using a stochastic *tail* order (as done by using $\preceq$ in G_{dist}), the game minimizes the mass located in the tail region (t_0, ∞) for a value t_0 that depends on the two distributions compared. Depending on how much this region overlaps with the interval $(E(X^*), \infty)$, this incurs some control over the disappointment rate, too (without additional efforts). In agreement with this line of argumentation, a recent research work [219] discusses improving the accuracy of security game models through including the concept of disappointment aversion as an additional decision objective and compute the equilibria accordingly.

7.5 SUMMARY

To deal with the uncertainty in the assessments, two game-theoretic models of surveillance optimization are investigated in Part II. In classical game settings, one can average over all assessments by computing the arithmetic mean (i.e., optimization using scalar-valued security scores). In the distribution-valued games, the assessment samples are described in the space of probability distributions to avoid the loss of information.

Let a normal game be given, in which an equilibrium is an optimal randomized security strategy for the defending player 1. If the defender goes ahead and mixes his strategies according to the equilibrium, then he should obtain a new pure and optimal strategy, expectedly outperforming all prior defenses (in being an optimized mix thereof). Somewhat unexpectedly, however, it turns out that this intuition is flawed as the counterexample showed, and that the choice of ordering under which the equilibrium is optimal can play a crucial role. In detail, if the ordering is the standard order of real numbers, the above intuition may fail. Attributing this failure to the way of purifying the strategy, however, is also flawed, since the phenomenon disappeared upon replacing the numeric order, and only the order, by a stochastic one.

The findings show that distribution-valued games' decisions can be more effective in practice. Their empirically assessed consequences meet the defender's satisfaction in terms of closeness to his aspiration level (*ideal point*) and the disappointment rate.

As it stands, the experimental study provides a new interpretation of mixed strategies in general, coming with the surprising property that optimality of an equilibrium can fail under certain, though natural, implementations of it. This reveals that the "optimality" of a defense is *not* the same as optimizing a security score since the means by which security is quantified and optimized play a much deeper role than intuitively expected. A resulting open research question derived from this observation concerns formal explanations of this effect, as well as conditions that characterize when this sub-optimality of equilibria can or cannot occur [2][1].

Part III

CYBERSECURITY MANAGEMENT

This part shows the application of the methodological approach presented in Part I to cybersecurity problems. In this part, the tweakable stochastic order is proposed, which can be adapted to individual risk attitudes of game players.

Publication references*:

- Ali **Alshawish** and Hermann de Meer. "Risk mitigation in electric power systems: Where to start?" *Energy Informatics* 2:34, 2019.
- Ali **Alshawish**, Korbinian Spielvogel, and Hermann de Meer. "A Model-based Time-to-Compromise Estimator to Assess the Security Posture of Vulnerable Networks." In *Proceedings of the International Conference on Networked Systems (NetSys)*, München, Germany, IEEE, 2019.
- Ali **Alshawish** and Hermann de Meer. "Prioritize when patching everthing is impossible!" In *Proceedings of the 44th IEEE Conference on Local Computer Networks (LCN)*, Osnabrück, Germany, IEEE, 2019.
- Ali **Alshawish** and Hermann de Meer. "Risk-based Decision-Support for Vulnerability Remediation in Electric Power Networks." In *Proceedings of the Tenth ACM International Conference on Future Energy Systems*, Phoenix AZ, USA, ACM, 2019.

* The research work (including some ideas and figures) from these papers, which is presented in Part III, was carried out and documented by the author of this thesis.

8 CYBERSECURITY GAMES

While Chapter 8 is based on the research work published in [3, 4, 6, 7]♮, its content is mainly adopted from [7]♮.

8.1 INTRODUCTION

CI systems (e. g., power systems, air traffic control, etc.) are becoming increasingly intelligent. In this regard, they benefit considerably from the IT networks, coupled with their underlying Operational Technology (OT) networks. While IT networks provide sufficient controllability and observability of critical assets and processes, they make them vulnerable to cyber threats and risks [7]♮.

On 12 May 2017, a very disruptive malware called WannaCry was observed. WannaCry is a ransomware instance, which infected about 250,000 computers in 150 countries as well as resulted in huge damage costs predicted about four billions of dollars [29]. The impact of WannaCry has been witnessed mainly in CI systems like healthcare organizations, transportation networks, telecommunications, among others. The destructive nature of this malware, the criticality of the infected systems, and the sheer scale of the attack are not the only interesting features of this malware but rather its attack vector. Interestingly, WannaCry has exploited a known Windows-specific vulnerability that was well-documented in the National Vulnerability Database (NVD) under CVE-2017-0144 [40]. The NVD entry was published on 16 March 2017, meaning that the vulnerability was discovered *at least* two months prior to the month of the event. On top of this, Microsoft released a vulnerability patch on 14 March 2017 towards fixing this vulnerability and providing protection against any potential attacks. That is, the infected systems would not have been subject to this attack, had these systems been updated during the two-month period before the attack. This raises the question, *why had these (critical) systems not been patched timely?*

To answer this question, one can review some security standards and guidelines as well as patch management plans such as *NIST Special Publication 800-40 Revision 2 – Creating a Patch and Vulnerability Management Program* [131], *NIST Special Publication 800-40 Revision 3 – Guide to Enterprise Patch Management Technologies* [202], and *BSI IT-Grundschutz-Kompendium* [23] . Based on the conducted review, the key reasons for this phenomenon are:

- **Strict patch validation process:** while standards encourage, if not oblige, organizations to perform maintenance and update of their assets in a timely manner, they impose very rigorous and time-consuming patch testing procedures before deployment.
- **Limited available security resources:** Lack of resources, including budget and time, is one of the key reasons for not timely upgrading such systems. This constraint imposes that available resources have to be intelligently allocated over the various security-related activities within an organization. This involves

a process of prioritizing risks and strategies, which ultimately vary from one system to another due to heterogeneity (i) in risk preferences of the involved decision-makers and (ii) in risk appetite and tolerance statements of organizations.

- **High reliability and availability requirements:** the increased reliance on such critical facilities requires high reliability (e. g., 99.999% availability; aka "five nines" uptime) of their networks. This corresponds to a very short downtime per year [95]. Hence, any maintenance and upgrade decisions have to be made very prudently. Moreover, some organizations lack the incentive to invest in more advanced and secure products. When thinking of (critical) control equipment, several operational organizations such as transportation organizations and power supply facilities are still using outdated software and unsupported operating systems such as Windows XP. The vulnerabilities of Meltdown and Spectre provide excellent evidence in this regard. Both vulnerabilities have been affecting a significantly wide spectrum of products that are still in use since 1995 [56].

In practice, such constraints would prevent organizations from fully resolving all of the vulnerabilities that their assets are at risk from. It is, therefore, very difficult – if not impossible – to have an operational system that is entirely vulnerability-free. Another complicating factor is the rapidly growing volume of released patches [7]♮. This can overburden security teams resulting in an imperfect patch management process. All these issues make the question (*where to start implementing remediation actions?*) pivotal in patch management processes.

A proper patch prioritization represents an efficient way of dealing with the aspects of security economics and risk management. It seeks to maximize the benefits of the available resources by focusing on the most critical issues first and hence minimize the inherent security risk in an effective manner [71, 74]. Such a process would certainly involve

(i) the use of some comparative judgments to define a ranking system, and

(ii) a decision-support technique to evaluate and compare the different options of a prioritization decision.

Given this fact, there is a need for assessing the security posture of vulnerable systems. Such assessments are of vital importance towards obtaining insights into existing security risks and hence supporting the decision-making process. Throughout this work, the TTC metric is pursued as a comparative security metric. TTC allows decision-makers to assess the security posture of such vulnerable systems and to intelligently compare different defense mechanisms (e.g., system hardening options) towards allocating available security resources in the most effective way. For the sake of simplicity, the model presented in this work is limited to only software (technical) vulnerabilities. Nevertheless, TTC can be leveraged to give an indication on a system's robustness against not only technical vulnerabilities but also social and organizational factors. The developed risk estimator considers several factors, including [7]♮:

(i) the inherent assessment uncertainty,

(ii) interdependencies between the network components,

(iii) different adversary skill levels, and

(iv) public vulnerability and exploit information.

As explained at several places in this thesis, the decision-makers involved in CI protection tend to enhance the system resilience against extreme events. Thus, they seek to avoid security decisions associated with likely severe risks. Practically, this risk attitude guides the decision-making process in such critical organizations and hence the sought-after prioritization as well. Therefore, this part of the thesis focuses on developing a decision-making process that employs the security management games presented in Section 4.1 to strategically prioritize vulnerability remediation actions towards minimizing the risk of compromise. Technically, actions are prioritized through successively playing a set of dependent zero-sum games. As shown in Section 4.1.2, the game-theoretical model considers the stochastic nature of risk assessments and the specific risk attitude of CI defenders carefully [7]♮.

8.2 STOCHASTIC TIME-TO-COMPROMISE MODEL

Commonly, TTC metric is used to compute single-point estimates such as Mean-Time-To-Compromise (MTTC) [117]. However, these estimates cannot robustly deliver an accurate risk prediction due to different uncertainties involved in real systems and underlying observational data [6]♮. Therefore, a generalized stochastic TTC model integrated with Monte Carlo simulation[1] techniques is presented in this work to account for the input data variability and inherent prediction uncertainty.

A TTC estimate denotes a prediction of the time needed for a potential adversary to exploit technical vulnerabilities of a system towards gaining unauthorized access to it. This corresponds to the time of a graph transition connecting a pair of nodes (`SOuRCe`, `DESTination`), given that the adversary controls the `SORC` node and seeks to compromise the `DEST` through exploiting its vulnerabilities. To estimate a Transition Time To Compromise (TTTC), a stochastic model is developed that takes into account a set of inputs summarized in Table 8.1. The inputs depend on existing statistical observations as well as outcomes of a security analysis of the system to be protected. The presented model is used to deliver comprehensive TTC estimates described using probability distributions instead of single-point estimates computed by the basic model described in [128].

Basically, the model rests on the following two probabilities:

- p_0: The probability that an adversary find "zero" fully functioning exploit (from his M available exploits) for the n vulnerabilities visible at `DEST`, given that there are totally N known vulnerabilities. Based on the definition of the hypergeometric distribution[2]:

$$p_0 = \frac{\binom{N-M}{n}}{\binom{N}{n}} \tag{8.1}$$

1 Monte Carlo simulation is a technique used to understand the impact of uncertainty and statistical behavior in prediction models. It depends on modeling input variables using probability distributions as well as performing an iterative empirical process to obtain the required predictions [136].

2 The hypergeometric distribution describes the probability of obtaining exactly m marked objects in n draws, without replacement, from a finite object population of size N that contains exactly M marked objects [66]: $P(m) = \frac{\binom{M}{m}\binom{N-M}{n-m}}{\binom{N}{n}}$

Table 8.1: A list of the TTC model inputs (adapted from [7]♮)

VAR.	DESCRIPTION AND INFORMATION SOURCE
N	The total number of disclosed vulnerabilities. Such information can be retrieved from major vulnerability databases such as the NVD and Rapid7 that catalogs about 141348 vulnerabilities [139, 160].
n_H	The number of known high-complex vulnerabilities (visible at DEST) that require a measurable amount of investments and efforts to be successfully exploited. One can use the Attack Complexity (AC) metric of the open standard CVSS to retrieve such details [41].
n_L	The number of known low-complex vulnerabilities (visible at DEST) exploitable without special conditions or circumstances [41].
n	The number of known vulnerabilities visible at DEST; $n = n_L + n_H$. The Attack Vector (AV) metric of the CVSS standard can be further used to identify the vulnerabilities' exploitation contexts; i. e., exploitable from (remote) network or adjacent/local access. This piece of information is used to identify which vulnerabilities are exploitable through inter-layer transitions or intra-layer transitions.
S	The adversary's experience and skill level function. S has a significant impact on the different time and probability computations of our model. For example, it is more certain that an expert adversary can employ existing exploits or even craft his own one with less time than the time needed by a beginner hacker. Based on an existing statistical study [117], S can equal to Expert=1.0, Intermediate= 0.55, Beginner= 0.3, or Novice=0.15.
E	The total number of existing exploits. Rapid7, a major exploit database, catalogs about 3859 readily available exploits [160].
M	The average number of readily available exploits that can be adapted or modified given the adversary skill level; $M = E \times S$ [117].
C	The average number of vulnerabilities for which an exploit can be found or crafted by an adversary given his S; $C = n \times S$ [117].
β_1*	The time needed for a successful compromise attempt using a readily available exploit code of known vulnerability. It is described by a random variable following the beta distribution with the mean of 1 day and a value range $[0 \ldots 5]$ days [129].
$\Gamma_{5.8}$*	The time needed to craft a working exploit code for a specific vulnerability. It is described by a random variable following the gamma distribution with the mean value of 5.8 days. 5.8 days has been derived based on the observed average time between a vulnerability announcement and the release of the first exploit [129].
Γ_{65}*	The time needed to find a new zero-day vulnerability. It is described, similar to $\Gamma_{5.8}$, by a random variable following the gamma distribution with the mean value of 65 days. 65 days is derived based on observations of the lifetime of zero-day vulnerabilities [130, 146].

* With more statistical and historical data, these values are expected to change to be more accurate.

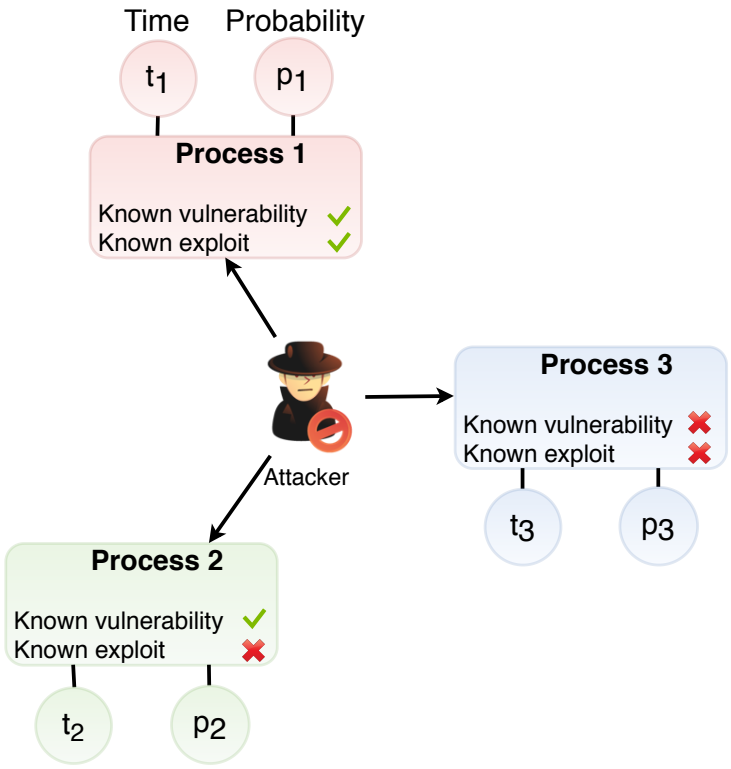

Figure 8.1: TTC model (adapted from [4]♮)

- $\hat{p}$: The probability that an adversary *fails* to craft any functioning exploit for the known vulnerabilities visible at DEST. $\hat{p}$ depends mainly on the adversary's skill level S ($0 \leqslant S \leqslant 1 \equiv$ Expert) and vulnerabilities visible at the target node (n_L, n_H) (see Table 8.1). More precisely, if DEST has no known vulnerability, then $\hat{p}$ should be 1. But, $\hat{p}$ should be very small if the adversary has in-depth knowledge (i. e., $S \approx 1$), and DEST has a known low-complex vulnerability; it can be approximated by $\hat{p} = 1 - S$. Under the assumption of independent vulnerabilities, $\hat{p}$ can be generalized as follows:

$$\hat{p} = (1 - S + \hat{l})^{n_L} \times (1 - S + \hat{h})^{n_H} \tag{8.2}$$

  where $\hat{l}$ and $\hat{h}$ are two control parameters[3] reflecting that an adversary's chance of failing is *higher* against high-complex vulnerabilities rather than low-complex ones.

In the TTC model (as illustrated in Figure 8.1), an adversary trying to compromise a node DEST can be in one of three random processes. For each process i, two quantities have to be assessed; namely

- p_i: the probability of being in process i, and
- t_i: the time needed for a successful compromise attempt given that the adversary is in the process i.

Concretely, these stochastic processes are:

3 Here, we use $\hat{l} = 0$, $\hat{h} = 0.10$, and the convention of $0^0 = 1$ in the $\hat{p}$ computations.

PROCESS 1: An adversary has identified one or more known vulnerabilities and has one or more exploits readily available. Therefore, the probability that the adversary is in Process 1 is the complement of the probability that an adversary has zero exploit readily available, which is p_0 as defined in Equation (8.1). This yields:

$$p_1 = 1 - p_0 = 1 - \frac{\binom{N-M}{n}}{\binom{N}{n}}$$

The time needed for an adversary in Process 1 can be described using the random variable β_1, as described in Table 8.1. Typically, the time and the adversary skill level vary inversely. Thus, the model modifies the time estimate in such a way that the time increases if the adversary skill level decreases. This yields:

$$t_1 = \beta_1 \times \frac{1}{S}$$

Remark. The number of readily available exploits that an adversary might have to compromise node DIST using its n known vulnerabilities might be modeled as a Poisson process with the average rate $\lambda = \frac{nM}{N}$ where the adversary has M exploits readily available for the totally N known vulnerabilities. That is, the probability that the adversary can find exactly m exploits from his M ones can be computed as follows:

$$P(X = m) = \frac{\lambda^m e^{-\lambda}}{m!}$$

Thus, the probability that an adversary find "zero" fully functioning exploit (from his M available exploits) for the n vulnerabilities visible at DEST can be computed as follows $p_0 = P(X = 0) = e^{-\frac{nM}{N}}$, while $p_1 = P(x > 0) = 1 - e^{-\frac{nM}{N}}$.

PROCESS 2: An adversary has identified one or more known vulnerabilities but could not find a functioning exploit readily available, and he tries to craft an own exploit. p_2 is defined as the product of the probability of having zero readily available exploit (i. e., p_0) and the probability of successfully developing at least one functioning exploit for at least one of the n visible vulnerabilities (i. e., $1 - \hat{p}$). This yields:

$$p_2 = p_0 \times (1 - \hat{p}) = \frac{\binom{N-M}{n}}{\binom{N}{n}} \times (1 - (1 - S + \hat{l})^{n_L} \times (1 - S + \hat{h})^{n_H})$$

Then, t_2 depends on (i) the time needed to craft a working exploit modeled as a random variable $\Gamma_{5.8}$ in Table 8.1, and (ii) the expected number of tries ET until the adversary can develop a fully working exploit code for one of the n vulnerabilities.

$$ET = S \times (1 + \sum_{k=2}^{n-C+1} [k \times \prod_{i=2}^{k} (\frac{n - C - i + 2}{n - i + 1})]) \tag{8.3}$$

This yields:

$$t_2 = \Gamma_{5.8} \times ET$$

Briefly, Equation (8.3) implies that the number of tries until developing one working exploit significantly depends on the adversary skill level; the higher the skill level, the less the number of tries. That is, as S increases, the expected number of vulnerabilities for which an exploit can be developed (C) increases, as well. Consequently, the number of useless vulnerabilities[4] due to lack of skills, defined as $(n - C)$, will decrease, and so do the number of tries ET. The detailed derivation of Equation (8.3) is shown in Appendix A.2.

PROCESS 3: An adversary does not have any working exploits, neither has he developed a functioning exploit for any known vulnerability at `DEST`. Therefore, he tries to discover an unknown (zero-day) vulnerability and then develop a working exploit therefor. For the sake of simplicity, a potential adversary can be in one of these processes. That is, the three identified processes are both "mutually exclusive" and "collectively exhaustive" and their probabilities can be added to yield a probability of 1. Thus, p_3 is equal to the product of the probability of having zero readily available exploit (p_0 defined in Equation (8.1)) and the probability of failing to develop any functioning exploit ($\hat{p}$ defined in Equation (8.2)):

$$p_3 = 1 - p_1 - p_2 = p_0 \times \hat{p} = \frac{\binom{N-M}{n}}{\binom{N}{n}} \times (1 - S + \hat{l})^{n_L} \times (1 - S + \hat{h})^{n_H}$$

In Process 3, t_3 involves three factors: (i) the time needed for discovering an unknown vulnerability, modeled as Γ_{65} in Table 8.1; (ii) the time needed to craft an own exploit $\Gamma_{5.8}$; and (iii) the skill level S. This yields:

$$t_3 = \frac{1}{S} \times (\Gamma_{65} + \Gamma_{5.8})$$

Ultimately, the transition time is the sum of the expected completion time of the three processes:

$$\text{TTTC} = t_1 \times p_1 + t_2 \times p_2 + t_3 \times p_3 \tag{8.4}$$

In light of this, a risk estimator is developed, integrating Equation (8.4) and its underlying processes with Monte Carlo simulation to assess the risk of compromise in critical networks (see Section 8.4.4 for further details) [7]♮.

8.3 CYBERSECURITY GAME

Broadly speaking, a cybersecurity game $G = \langle \{\mathcal{D}, \mathcal{A}\}, \{SP_{\mathcal{D}}, SP_{\mathcal{A}}\}, A \in \mathcal{F}^{n \times m}, \preceq \rangle$ is an instance of the zero-sum security management game introduced in Section 4.1.2. It is used to model the interaction between a defender $\mathcal{D}$ and an attacker $\mathcal{A}$. Here, the latter abstracts all external adversaries that seek to benefit from a network's technical vulnerabilities towards compromising a target component that is usually critical to the operation of the respective network. In contrast, $\mathcal{D}$ abstracts any decision-maker (e. g., chief security officer or patch management operation team) seeking to minimize the risk

4 The notion of useless vulnerabilities refers to the vulnerabilities that an adversary will not be able to use given his skill level S. The higher the number of useless vulnerabilities, the higher the expected number of tries until crafting a working exploit code.

of compromising the target. The finite action space $SP_{\mathcal{D}} = \{d_i\}$ includes the vulnerability remediation activities the defender is able to perform to defend the network in question towards minimizing the risk of compromise. In contrast, $SP_{\mathcal{A}} = \{a_i\}$ represents the potential ways the attacker can use to compromise the network. The utility function of G is modeled as a payoff matrix A, telling the estimated risk of compromise under each combination in $SP_{\mathcal{D}} \times SP_{\mathcal{A}}$. Here, the risk is quantified by means of the TTC security metric. The payoffs are described using more general risk descriptions $\mathcal{F}$ instead of $\mathbb{R}$, and the stochastic tail order ($\preceq$) is used to play the game. As explained in Section 4.1.2, this order is consistent with the aforementioned risk attitude. Hence, $\preceq$-based games have the appeal of minimizing the likelihood of extreme risks by doing optimization through shifting the risk mass towards low-risk levels rather than optimizing single statics such as the average values. This is achieved by choosing the security strategy $\delta^*_{\mathcal{D}} \in \Delta(SP_{\mathcal{D}})$ that puts more importance on $\mathcal{D}$'s actions that essentially remedy risks with high(er) likelihood for high(er) levels. The object $\delta^*_{\mathcal{D}}$, described in the form of a probability measure $\left(\delta^*_{\mathcal{D}} : SP_{\mathcal{D}} \to [0,1]\right)$, assigns probability $\delta^*_{\mathcal{D}}(d_i) \geqslant 0$ for each action $d_i \in SP_{\mathcal{D}}$ and satisfies $\sum_{d_i \in SP_{\mathcal{D}}} \delta^*_{\mathcal{D}}(d_i) = 1$.

Ultimately, one can interpret the obtained security strategy $\delta^*_{\mathcal{D}}$ (hereafter referred to as δ^*) as a belief function on the defense actions $SP_{\mathcal{D}}$. This belief function can be realized to the defender as advice on how to best defend the network of interest using the most effective remediation actions. Here, the most effective actions stand for those actions assigned with nonzero-probabilities by the belief function, i. e., $\delta^*(d_i) > 0$. In practice, $\mathcal{D}$ has no incentive to play actions assigned with zero-probabilities as they are dominated actions, and it is definitively better to play other actions given the equilibrium state defined by δ^* [7]♮.

8.4 PRIORITIZATION FRAMEWORK

Vulnerabilities are weaknesses that create possible points of security compromise for a target component that matters most for a network of interest [4]♮. In the cyber realm, remote attackers seek to exploit cyber vulnerabilities present in IT networks to obtain unauthorized access to interconnected OT networks, thereby causing significant damages. However, resolving all vulnerabilities at once could seem like an insuperable hurdle due to several technical and economic constraints. Therefore, a system defender needs to prudently assess priorities and make a decision on the importance of the possible remediation activities in order to implement them more effectively [3]♮. To this end, this work introduces a prioritization framework based on the methodological approach for security management presented in Section 4.2. The framework assists the defender in making risk-informed decisions on the action priorities. It addresses comprehensively the competitive nature of the decision environment, the specific risk attitude of the defender of CIs, and uncertainties inherent in risk assessments [7]♮.

8.4.1 *Context establishment*

As explained in Section 4.2.1, this step aims at understanding the system and the environment of interest. This can involve, just to name a few, (i) identification of different components and resources relevant to the examined system and the connections among them; (i) identification of possible exposures to risks using techniques such as

vulnerability assessment; and (iii) identification of a potential target component "T" that matters most to the system of interest. With regard to the latter point, Master Terminal Units (MTUs), Intelligent Electronic Devices (IEDs), data concentrator, and SCADA servers are of crucial importance for controlling and operating OT networks since they communicate and control critical machinery and processes [4]. The outcome of this step is (i) a topological map of the examined system, and (ii) a list of the known vulnerabilities of the system components with their CVSS-based characteristics such as the AV and AC metrics. These data are summarized in SQ table representing the "*status quo*" of the system before implementing any remediation action.

8.4.2 *Identification of strategies*

This step involves identification of possible actions that can be done by involved agents (i. e., defender and attacker) to accomplish their intended goals.

8.4.2.1 *Identification of potential attack strategies*

Attack (or compromise) strategies represent a set of entry points to the examined network and their corresponding (feasible) compromise paths. These paths can be used by a remote adversary to reach the identified target. Based on the topological map of the studied network (delivered by the step described in Section 8.4.1), one can model the possible attack strategies using asset-centric compromise graphs[5].

In a compromise graph, there are basically two node types based on the characteristics and the functionality of the corresponding physical component or subsystem: (i) `Network` nodes that are accessible from across the Internet or from a different layer (e. g., border devices, such as routers and firewalls, are always `Network` nodes as they can maintain connectivity between two layers); and (ii) `Local` nodes that are only accessible locally and from nodes located in the same network layer. The target node "T" can, therefore, be either a `Network` node or `Local` node based on its characteristics and connectivity pattern. Additionally, each compromise graph has one hypothetical root node (called "Launch") representing an adversarial remote node.

The transitions (or edges) of a compromise graph represent the possible compromise steps. They are classified into: (i) `Breach` edges (or inter-layer transitions; only possible if the transition's source and destination nodes belong to different layers and the destination is a `Network` node), and (ii) `Penetration` edges (or intra-layer transitions; only possible between two nodes of the same layer regardless whether they are `Network` or `Local` nodes). In this respect, it is worth mentioning that the involvement of experts with specialized domain knowledge and security skills can be of vital importance at this step to refine and simplify the final compromise graphs through discarding impractical and technically infeasible compromise paths. The output of this step describes the attacker's action set $SP_{\mathcal{A}}$.

5 Compromise graph is an asset-centric rather than a vulnerability-centric modeling technique. This aims at (i) avoiding the known "state explosion problem" due to the potentially large number of vulnerabilities in a system; and (ii) simplifying the model to the system's operators, who usually do not understand the language of technical vulnerabilities. In the asset-centric approach, nodes are the components of the examined network. Thus, if there are some components that approximately share the same profile (e. g., connectivity pattern, functions, patch level, etc.), they can be grouped into a single separate subsystem (one node in the graph). This facilitates an additional reduction of the graph complexity.

8.4.2.2 *Identification of potential defense strategies*

Defense strategies represent the different vulnerability remediation actions (patching, system hardening activities, etc.) or security investment plans the defender can implement in order to control and mitigate the risk of compromise. For the sake of simplicity, each set of changes and activities designed to fix and improve an individual node of the identified compromise graphs can be represented by one defense strategy. Since there are some vulnerabilities without any applicable patches or workarounds, each strategy d_i is characterized by its envisaged `Fix-rate(`d_i`)`. This metric is the ratio between the number of fixed vulnerabilities and the number of vulnerabilities identified in the respective node. The output of this step describes the defender's action space ($SP_{\mathcal{D}}$).

8.4.3 *Identification of goals*

This step aims at identifying the different goals relevant to the sought-after prioritization process. Throughout this work, the focus is merely laid on minimizing the risk of compromise quantified in terms of the presented TTC security metric. More specifically, the defense strategies identified in Section 8.4.2.2 are assessed in terms of their impact on risk reduction against all compromise strategies identified in Section 8.4.2.1.

8.4.4 *Effectiveness assessment*

Generally, this step aims at assessing the outcomes of all possible combinations of the (defender, attacker) actions, namely all $(d_i, a_j) \in SP_{\mathcal{D}} \times SP_{\mathcal{A}}$, in terms of the identified goals. Here, only one objective is to be optimized, which is the risk of compromise. The assessment process benefits from the stochastic TTC model described in Section 8.2. The model involves the use of a wide variety of observed and statistical data. That is, significant uncertainty and variability are associated with such data and can have a severe impact on the TTC estimation process. Single-point estimates fail to communicate comprehensive risk assessments to decision-makers. For instance, presenting assessments using mean values makes decision-makers indifferent between different uncertain options with equal values, even if one option might be riskier. To address this challenge, the assessment step incorporates an iterative TTC estimation process based on Monte Carlo simulation techniques, in which any input parameter that has inherent uncertainty is modeled using a proper probability distribution function. At each iteration, different values can be used for these parameters based on their distribution functions. In this way, the assessment outcomes will provide the decision-maker with a range of possible TTC estimates and the occurrence probabilities thereof. In addition to random sampling, each iteration of the risk assessment process of a scenario $(d_i, a_j) \in SP_{\mathcal{D}} \times SP_{\mathcal{A}}$, includes the following steps:

1) Identify the involved compromise graph based on the strategy a_j.
2) Retrieve values of some model inputs (such as n_H, n_L) from $\mathcal{SQ}_{d_i}$, which is a version of the state $\mathcal{SQ}$ locally modified according to `Fix-rate(`d_i`)`. For instance, suppose $\mathcal{SQ}$ states that nodes x and y have 5 and 3 high-complex vulnerabilities, respectively. If d_i fixes all vulnerabilities in node x, then the TTC model will use $\mathcal{SQ}_{d_i}$, in which $(x, y) \xrightarrow{n_H} (0, 3)$ to assess the risk of compromise.

3) Estimate a TTTC value of each transition in a_j through applying the model described in Section 8.2 and its Equation (8.4).

4) Estimate a TTC value of each identified path from node "Launch" to "T" in the graph a_j, denoted as Path Time To Compromise (PTTC). A PTTC value of a specific path z is simply the sum of the TTTC estimates of its constituting transitions $ct \in z$, i. e., $PTTC_z = \sum_{ct \in z} TTTC_{ct}$.

5) Record the obtained PTTC estimates for all identified compromise paths in a_j.

Subsequently, the outcomes of all iterations are merged using several techniques (e. g., frequency histogram, kernel density estimation, or the maximum entropy method) to generate the final TTC distribution function. It is worth mentioning that the assessment results of all scenarios $(d_i, a_j) \in SP_{\mathcal{D}} \times SP_{\mathcal{A}}$ will be used to construct the payoff matrices of security games that ultimately support the sought-after prioritization decisions.

8.4.5 *Prioritization process of the defense strategies*

This step aims at assisting the defender in arranging the possible defense strategies in the order of their risk mitigation effects. This involves an iterative process of playing security games, whose underlying model is sketched in Section 8.3. Each game supports the defender in choosing and ranking one action as dictated by its computed security strategy. As a result, this process yields a chain of security games, the length of which is equal to $(|SP_{\mathcal{D}}| - 1)$, where $|SP_{\mathcal{D}}|$ stands for the cardinality of the set $SP_{\mathcal{D}}$. This technique is called Iterated Prioritization of Risk Mitigation Actions (IPRMA), while the whole process is described in Algorithm 1. The first game in the chain G_1 is constructed using the complete action spaces $SP_{\mathcal{D}}$ and $SP_{\mathcal{A}}$ as well as their corresponding payoff matrix A_1, whose elements are assessed following the process explained in Section 8.4.4. The best action of G_1, denoted as d_1^*, will be chosen according to the probability distribution prescribed by the security strategy of G_1, i. e., $\delta_1^* \in \Delta(SP_{\mathcal{D}})$. Then, d_1^* is ranked top on the ordered action list, assigned with the highest priority to be implemented. Afterward, the system state $\mathcal{SQ}$ is *globally* updated according to the envisaged remediation effects of d_1^* (i.e., `Fix-rate(`d_1^*`)`). That is, $\mathcal{SQ}$ is modified as if d_1^* would have been implemented practically. Then, d_1^* will be removed from the possible action space $SP_{\mathcal{D}}$. The changes applied on $SP_{\mathcal{D}}$ and $\mathcal{SQ}$ result in a new and smaller game, the best action of which is assigned a lower priority than the previously removed action. This process is repeated, creating new and even smaller games until all security actions are ranked [7]‡.

8.5 SUMMARY

There is a rising need to quantify and measure security as a vital step towards a practical planning and decision-making process. Security metrics such as TTC can significantly help security officers assessing and prioritizing various security risks as well as mitigation strategies. Furthermore, such practice is essential to allocate scarce and expensive security resources in an effective manner. In this chapter, a stochastic TTC model for assessing the security posture of IT networks is described. That model can be employed to provide comprehensive TTC estimates, thereby offering valuable insights into existing security risks.

Algorithm 1 IPRMA process - chained games (adapted from [7]♮)

Require: SQ, $SP_{\mathcal{D}} \leftarrow \{d_1, \ldots, d_n\}$, $SP_{\mathcal{A}} \leftarrow \{a_1, \ldots, a_m\}$
Ensure: an ordered list of $SP_{\mathcal{D}}$ according to their risk mitigation impact

initialize $ol \leftarrow \{\}$ ▷ an empty ordered list
initialize $k \leftarrow 0$ ▷ the game index
while $length(SP_{\mathcal{D}}) > 1$ **do** ▷ $length(SP_{\mathcal{D}}) \equiv |SP_{\mathcal{D}}|$
    $k \leftarrow k + 1$
    $A_k \leftarrow assessRisk(SP_{\mathcal{D}}, SP_{\mathcal{A}}, SQ)$ ▷ assess the payoff matrix for all action combinations in $SP_{\mathcal{D}} \times SP_{\mathcal{A}}$
    $G_k \leftarrow constructGame(SP_{\mathcal{D}}, SP_{\mathcal{A}}, A_k)$
    $\delta_k^* \leftarrow lexNashEq(G_k)$ ▷ compute the security strategy of G_k
    $d_k^* \leftarrow bestAction(\delta_k^*)$ ▷ the best action(s) drawn acc. to the probability distribution prescribed by δ_k^*
    $ol.insert(d_k^*)$ ▷ add the best action into ol
    $SQ.update(d_k^*)$ ▷ update the (global) state SQ with the changes associated with d_k^*
    $SP_{\mathcal{D}} \leftarrow SP_{\mathcal{D}} \setminus \{d_k^*\}$ ▷ remove the best action from $SP_{\mathcal{D}}$
end while
$ol.insert(SP_{\mathcal{D}})$ ▷ insert the last (least important) action into ol
return ol ▷ return the ordered list of the defender actions

As shown in Part II, physical surveillance games focus on one facet of security management in CIs, namely security resources allocation. The focus in this chapter is laid on another extremely important facet, which is prioritization decisions. Therefore, the methodological approach for security management presented in Chapter 4 is adapted to successively prioritize possible vulnerability patch actions according to their risk remediation impact. The whole approach is illustrated by examining a case study in Chapter 9.

9 USE CASE

The content of Chapter 9 is based on the research work published in [7]♮.

9.1 INTRODUCTION

Given several technical and operational constraints, there is a growing need for developing a coherent patch management plan towards effectively reducing security risks posed by potential plans to compromise CIs such as electric power systems. In a recent study on the resilience of power systems, Bie et al. stress the vital importance of being able to mitigate extreme risks as a condition for having resilient electricity infrastructures [31]. In this respect, this chapter illustrates the application of the prioritization approach put forth in Chapter 8 to assist an involved defender in prioritizing possible remediation actions according to their mitigation effects of high-level risks. For illustrative purposes, this work considers a simplified network of an electricity provider, which controls the electricity provision process basically using SCADA systems [7]♮. The decision makers involved in the management operations of this system increasingly integrate IT devices into the OT space that had been designed with neither widespread connectivity nor adequate security in mind. On the one hand, this integration aims at leveraging all available resources for enhancing the grid efficiency and control. But on the other hand, it could pave the way for a broad spectrum of potential attackers, ranging from amateur (cyber) criminal to advanced terrorist and state-sponsored attackers, to take control of critical assets and operational resources.

9.2 CONTEXT ESTABLISHMENT

As a first step, it is necessary to analyze the IT network of the examined system. The analysis outcome is depicted in Figure 9.1. It illustrates the topological map of the examined electricity provider with the different technical subsystems and the connections among them. The electricity provider operates basically two different interconnected network layers. Layer (LA) includes the most networking components that are reflecting the business and the high-level control requirements. It is composed of the traditional office workstations and servers, as well as the control servers that are responsible for the high-level supervision and data acquisition of the devices located in the substation network. Based on their functions and connectivity characteristics, the devices in LA are grouped into three subsystems S1, S2, and S3, as depicted in Figure 9.1. Layer (LB) provides an abstract representation of an IEC-61850-based electric substation. This layer includes three subsystems S4, S5, and S6. Subsystem S4 includes the local substation workstations and Human Machine Interface (HMI) devices. Subsystem S5 comprises the substation management server for managing the substation asset integrity and reliability. Subsystem S6 represents the substation controller connected to the most critical process network and primary field devices. These devices include, just to name a few, transformers, circuit breakers, and capacitor banks. Controlling and protecting

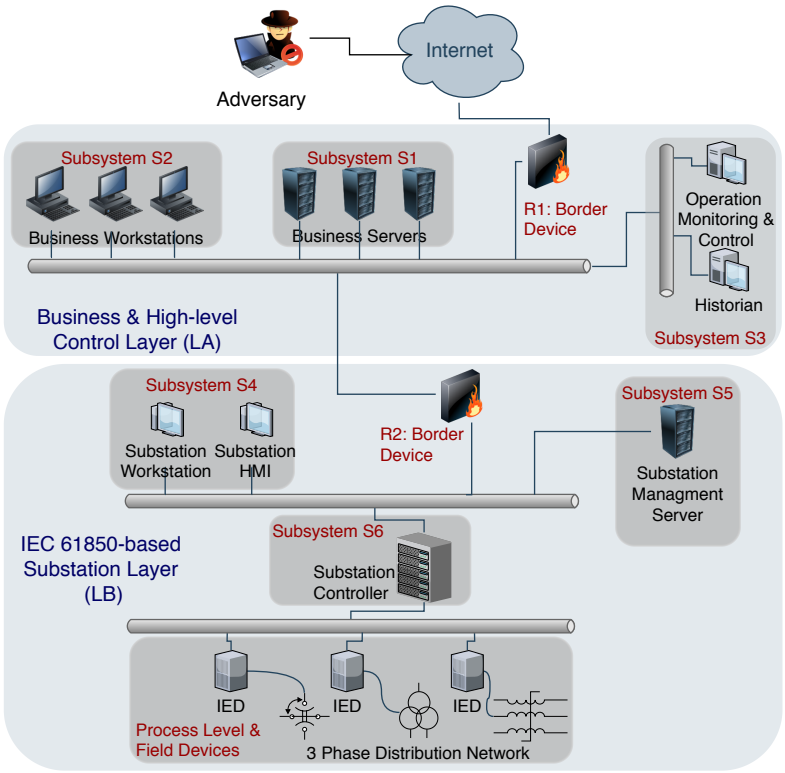

Figure 9.1: A topological map of the studied electricity provider network [7]♮

these critical devices involve the use of a set of programmable devices called IEDs. Additionally, the examined system utilizes two border devices R1 and R2 (with router and firewall functionality), to control the segregation between the whole system and the Internet as well as between the two identified layers. Concerning the accessibility type, S3 and S6 are `Local` components as they are not accessible from outside their respective layers. The other LB's devices are `Network` components but not accessible from across the Internet. In contrast, S1, S2, and R1 are `Network` components and Internet-accessible, marked as $\texttt{Network}^{+}$ nodes. S6 is identified as the target node "T" of this study based on its crucial role in controlling and operating the electric distribution network. More specifically, once a remote adversary $\mathcal{A}$ gains unauthorized access to S6 through a cyber intrusion path, $\mathcal{A}$ has control of important devices such as protective relays and circuit breakers. These devices are typically employed to protect critical and expensive assets such as transformers, generators, or transmission and distribution lines. Therefore, $\mathcal{A}$ can cause significant damage and a widespread power outage by manipulating the configuration settings of these devices. Exploiting cyber vulnerabilities of power grids can result in further consequences including, but not limited to,

- disruption of grid stability through controlling Volt-Amp Reactive (VAR) devices, thereby causing voltage and frequency fluctuations in the grid;
- loss of substation information essential to the reliable operation of power grids such as metering information and fault recordings; and

Table 9.1: The shared system state $\mathcal{SQ}$ [7]♮

Sub-system	Access. type	AV: Network		AV: Adj.+Local	
		n_H	n_L	n_H	n_L
R1	Network$^+$	1	5	0	0
S1	Network$^+$	2	5	3	7
S2	Network$^+$	3	4	5	5
S3	Local	1	3	2	6
R2	Network	2	5	0	0
S4	Network	5	5	5	5
S5	Network	3	6	4	4
S6 (T)	Local	0	3	3	6

- loss or interruption of communication and control channels and thus loss of engineering and maintenance access to IEDs and Remote Terminal Units (RTUs) [24].

As explained in Section 8.4.1, the process of vulnerability analysis gives additional insights into the number of vulnerabilities visible in the network, classified according to their CVSS-based characteristics; i. e., AV and AC metrics. These pieces of information form the shared system state $\mathcal{SQ}$, as summarized in Table 9.1.

9.3 IDENTIFICATION OF STRATEGIES

9.3.1 *Identification of potential attack strategies*

Based on the outcome of the above step described in Section 9.2, one can identify three entry points available for a remote adversary $\mathcal{A}$ attempting to compromise the identified target subsystem. These points are the three subsystems S1, S2, and R1, which are Internet-accessible. As previously explained in Section 8.4.2.1, each attack strategy can be modeled using a compromise graph describing the different feasible compromise paths from the respective entry point to the target. Figure 9.2 depicts three compromise graphs corresponding to the three possible attack strategies. The attack strategy a_1, for example, aims at exploiting the weaknesses of the border device (R1) to breach[1] Layer (LA) in the first place. After establishing an initial foothold in LA, $\mathcal{A}$ has two options: (i) spreading through LA to strengthen the gained foothold through penetrating an ordinary node S1, S2, or S3 and then breaching Layer LB (APT-like attack scenarios); or (ii) rushing forward towards the target through breaching a network node in Layer LB; i. e., R2, S4, or S5. At this stage, technical and domain knowledge from experts can be incorporated to refine the list of paths depending on their relevance and practical feasibility. Based on such knowledge, the back transitions, such as the one from S1 to R1 in the compromise graph a_1, are obviously meaningless. Likewise, the attack strategies a_2 and a_3 are established, exploiting the vulnerable network nodes S1 and

1 Breach stands for inter-layer transitions. Penetration stands for intra-layer transitions.

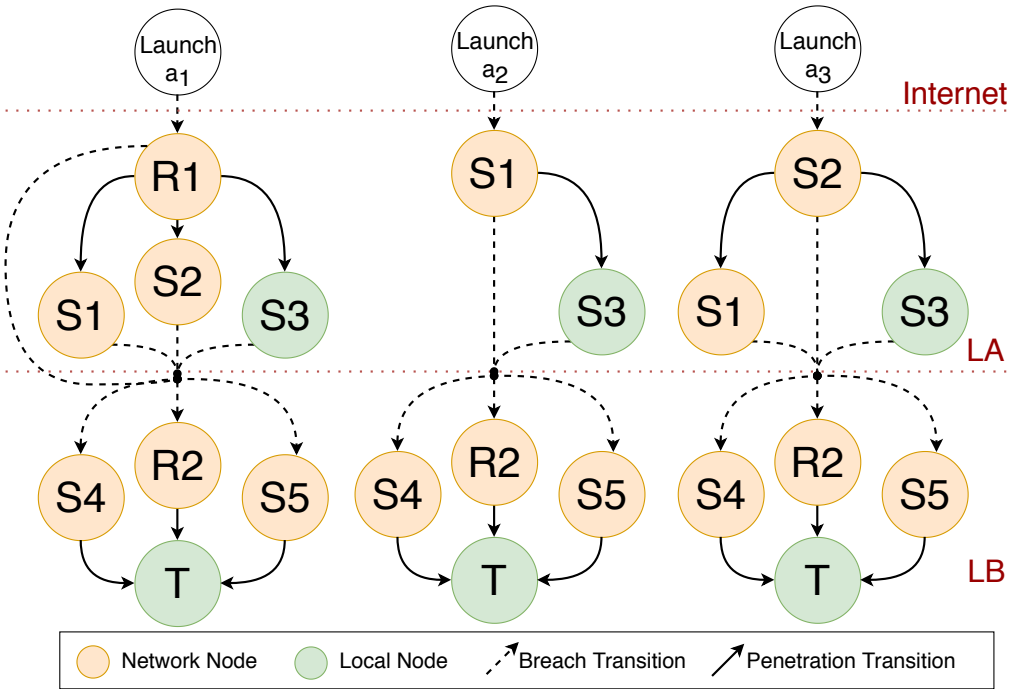

Figure 9.2: Three compromise graphs corresponding to the identified attack strategies a_1, a_2, and a_3 [7]‡

S2, respectively. Here, the breach transition from S1 to S2 is deemed as technically meaningless and can not offer potential adversaries with better chances to reach the target. Therefore, it is omitted from the compromise graph of a_2. It is worth mentioning that compromise graphs provide a powerful and compact representation of $\mathcal{A}$'s action space. Each graph can be easily updated upon the identification of new compromise steps/paths.

9.3.2 *Identification of possible defense strategies*

The defender $\mathcal{D}$ has identified 8 defensive actions corresponding to the patching solutions designed to fix the known vulnerabilities in the 8 nodes of the established compromise graphs. These strategies are $SP_{\mathcal{D}} = \{d_1 - R1,\ d_2 - S1,\ d_3 - S2,\ d_4 - S3,\ d_5 - R1,\ d_6 - S4,\ d_7 - S5,\ d_8 - T\}$, where the strategy $(d_1 - R1)$ stands for the defense strategy d_1 dedicated to fix the known vulnerabilities in the node R1. If there are some vulnerabilities without any applicable patches or workarounds, these vulnerabilities should not be removed from the shared state $\mathcal{SQ}$ when updating their respective nodes during the prioritization process. For the sake of simplicity, it is assumed here that each defense strategy is able to resolve all vulnerabilities visible at its respective node completely; $\texttt{Fix-rate}(d_i) = 1 \quad \forall d_i \in SP_{\mathcal{D}}$.

Table 9.2: Overview of the used risk levels; higher levels indicate a higher risk of compromise [7]♮

RISK LEVEL	TTC INTERVAL (IN DAYS)
10	0 - 14
9	15 - 28
8	29 - 45
7	46 - 90
6	91 - 150
5	151 - 230
4	231 - 300
3	301 - 360
2	361 - 540
1	> 540

9.4 EFFECTIVENESS ASSESSMENT

The risk of compromise is quantified using the TTC security metric. The Monte-Carlo-simulation-based assessment process has (1000) iterations[2] and utilizes the model described in Section 8.2. For each iteration, the input parameters accept different values according to their specified distribution functions. Regarding the adversary skill level parameter, each iteration chooses a random value based on the following probability mass function (Expert: 14%, Intermediate: 33%, Beginner: 34% and Novice: 19%), which is derived from the statistical findings of a previous research work on the classification of hackers by their observed behaviors [233]. The obtained TTC distributions can be further processed to generate corresponding risk probability distributions through categorizing the TTC assessments based on a set of risk categories that are predefined and approved by the system operator and other involved stakeholders (see Table 9.2). In Algorithm 1, the function assessRisk() realizes the risk assessment process described in this section to return the payoff matrices needed for the cyber security games.

9.5 PRIORITIZATION PROCESS OF DEFENSE STRATEGIES

Based on Algorithm 1, the prioritization process involves constructing a chain of 7 security games. Table 9.3 summarizes the input/output associated with each of those games. The chain begins with the game G_1, which is constructed using the whole action spaces $SP_{\mathcal{D}}$ and $SP_{\mathcal{A}}$, where $|SP_{\mathcal{D}}| = 8$ and $|SP_{\mathcal{A}}| = 3$. Using the shared state $\mathcal{SQ}$ described in Table 9.1, the function assessRisk() computes the payoff matrix A_1 of G_1. For the sake of clarity, Figure 9.3 shows the matrix A_1 used to compute the security

2 The number of iterations has been estimated by fixing a precision factor $\epsilon = 0.001$ and using the Kullback-Leibler divergence $D_{KL}(X_{k_a} \| X_{k_b})$ to measure the difference between two probability distributions representing two risk distributions of the same scenario estimated using a different number of iterations. A random test scenario is fixed and then different number of iterations $\{100, 200, \ldots, 10000\}$ are tested. The number of 1000 has been chosen since $D_{KL}(X_{1100} \| X_{1000}) \approx 0.000586 < \epsilon$.

strategy in G_1. The matrix has a shape of 8×3. Each matrix element (i, j) corresponds to the comprehensive TTC-based risk assessments of the respective action combination $(d_i, a_j) \in SP_{\mathcal{D}} \times SP_{\mathcal{A}}$. Figure 9.3 shows that the risk of compromise varies not only from one defense action to another (e. g., risk of level 10 and 9 is more probable under action d_4, as shown in the 4^{th} row in A_1, rather than action d_8 – regardless which compromise action is played) but also from one compromise action to another given a specific defense action (e. g., risk of level 10 and 9 is more probable under action d_2 if the attacker follows action a_1 or a_3 but not a_2). That is, even simple scenarios can be associated with a certain amount of complexity involved in answering important questions such as *where to start?* and *what to do next?*. Therefore, the present game-theoretical approach analyzes the situation as a whole towards supporting the defender when making prioritization-related decisions.

As Table 9.3 tells us, the security strategy of G_1 describes a pure equilibrium strategy, in which the action $(d_8 - T)$ is the most effective action in reducing the risk of compromise under the current state $\mathcal{SQ}$. Therefore, the defender assigns the highest priority to fix the vulnerabilities visible at the target node T (i. e., S6) immediately. Based on this result, the action $(d_8 - T)$ is placed at the top of the sought-after ranking and removed from $SP_{\mathcal{D}}$. Then, $\mathcal{SQ}$ is updated accordingly by removing all vulnerability in the target. This yields a new game G_2, which has the same attack action space but with a smaller defense action space $SP_{\mathcal{D}} \leftarrow SP_{\mathcal{D}} \setminus \{d_8\}$. The game chain proceeds forwards until all the defensive actions are ranked. It is worth mentioning that the function `bestAction()` uses the probability distribution dictated by the output strategy of each game to draw the corresponding best action. For example, `bestAction()` chooses the action $(d_4 - S_3)$ with the probability (0.375) and the action $(d_6 - S_4)$ with the probability (0.625) as dictated by the mixed equilibrium strategy δ_2^* of the game G_2. Table 9.3 shows only one prioritization option by pursuing the actions with the highest probabilities, i. e., $d_k^* \leftarrow \mathrm{argmax}_{d_i \in SP_{\mathcal{D}}} \delta_k^*(d_i)$. Afterward, the chain proceeds forwards until the last game G_7, which supports the decision on the prioritization of the last two actions. Ultimately, there are definitively at least two prioritization options if there is one game of the chain with a mixed equilibrium strategy. These options can be combined in a comprehensive prioritization tree, in which the nodes are the different defense actions connected by edges that have weights representing the action probabilities as assigned by the corresponding security strategies. Each tree has a hypothetical `root` node. The weight of each path l, starting from the `root` to any leaf node in the tree, can be computed as the product of the weights of its composing edges; i. e., $w(l) = \prod_{e_i \in l} w(e_i)$, where $w(e_i)$ stands for the weight of the edge e_i that is part of the path l. With regard to the studied use case, Figure 9.4 depicts the final prioritization tree. It includes three prioritization options: i) OptionA $= d_8 \rightarrow d_4 \rightarrow d_6 \rightarrow d_1 \rightarrow d_2 \rightarrow d_7 \rightarrow d_3 \rightarrow d_5$, ii) OptionB $= d_8 \rightarrow d_6 \rightarrow d_2 \rightarrow d_1 \rightarrow d_7 \rightarrow d_4 \rightarrow d_5 \rightarrow d_3$, and iii) OptionC $= d_8 \rightarrow d_6 \rightarrow d_2 \rightarrow d_7 \rightarrow d_4 \rightarrow d_3 \rightarrow d_1 \rightarrow d_5$ with the probabilistic weights of 0.375, 0.375, and 0.25, respectively.

Remark. The above results imply the assumption that the defender is limited to perform (or complete) only one defense strategy every step due some constraints in the sense of available time, human resources, costs, or even policy-related restrictions. That is, the remediation problem of the whole system is already quantized into a finite set of manageable remediation actions in the light of existing constraints. Nevertheless, if the defending team has the resources and is allowed by existing policy to

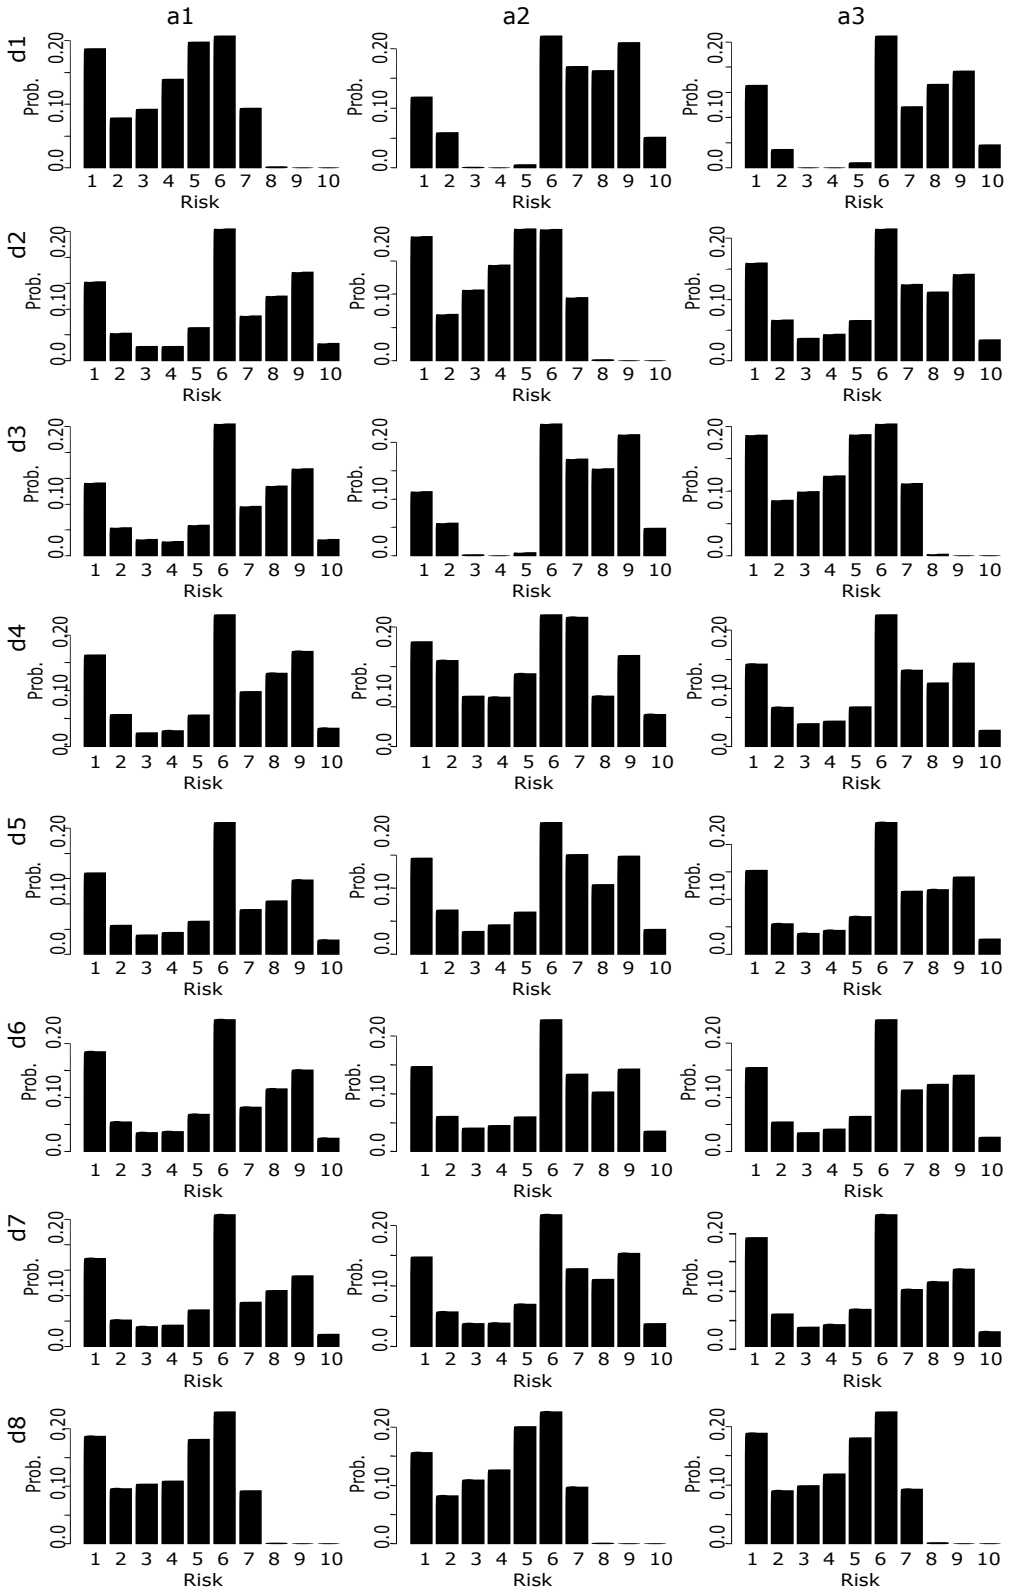

Figure 9.3: A_1: the 8×3 payoff matrix of the first game (G_1) in the chain [7]♮

Table 9.3: A chain of stochastic security games [7]♮

Game G	$SP_{\mathcal{A}}$	$SP_{\mathcal{D}}$	$\delta_k^* = \text{lexNashEq}(G)$	$\text{bestAction}(\delta_k^*)$
G_1	$\{a_1, a_2, a_3\}$	$\{d_1, d_2, d_3, d_4, d_5, d_6, d_7, d_8\}$	$(0,0,0,0,0,0,0,1)$	$d_8 - T$
G_2	$\{a_1, a_2, a_3\}$	$\{d_1, d_2, d_3, d_4, d_5, d_6, d_7\}$	$(0,0,0,0.375,0,0.625,0)$	$d_6 - S4$
G_3	$\{a_1, a_2, a_3\}$	$\{d_1, d_2, d_3, d_4, d_5, d_7\}$	$(0,1,0,0,0,0)$	$d_2 - S1$
G_4	$\{a_1, a_2, a_3\}$	$\{d_1, d_3, d_4, d_5, d_7\}$	$(0.6,0,0,0,0.4)$	$d_1 - R1$
G_5	$\{a_1, a_2, a_3\}$	$\{d_3, d_4, d_5, d_7\}$	$(0,0,0,1)$	$d_7 - S5$
G_6	$\{a_1, a_2, a_3\}$	$\{d_3, d_4, d_5\}$	$(1,0,0)$	$d_4 - S3$
G_7	$\{a_1, a_2, a_3\}$	$\{d_3, d_5\}$	$(0,1)$	$d_5 - R2$

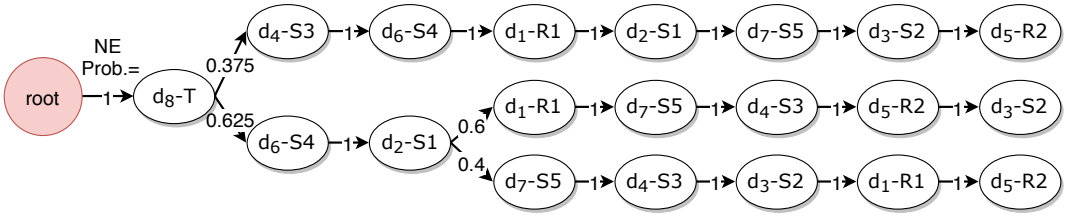

Figure 9.4: The decision-support tree for action prioritization [7]♮

perform multiple defense actions in parallel, there is a possibility to adapt the function $\text{bestAction}(\delta_k^*, n\text{Max})$ to return multiple actions if the equilibrium δ_k^* allows. Here, $n\text{Max}$ specifies the maximum number of actions the defender is able to perform in parallel; in this work $n\text{Max} = 1$.

9.6 EVALUATION OF OBTAINED PRIORITIZATION OPTIONS

This section analyzes the results of applying the proposed prioritization methodology. Moreover, it illustrates the performance of the delivered prioritization options. The fundamental goal of the developed framework is achieved by constructing the prioritization tree depicted in Figure 9.4. That tree supports the defender in making risk-informed decisions about the prioritization of the possible security actions. It represents a tremendous reduction of the decision space that the defender needs to explore. In the examined use case, the framework ends up with 3 prioritization options out of 40320 possible prioritization variations of the 8 identified defense actions[3].

As a risk-based prioritization approach, the defender is interested in investigating whether the three delivered decision options have comparatively equivalent risk mitigation effects. This analysis is achieved by utilizing the equilibrium payoffs obtained by the different games of the constructed chain. The equilibrium payoffs describe the expected risk distributions the defender can assure himself in the different games. To have a complete vision of the risk mitigation progress as the decision-support chain moves forward, two additional games G_0 and G_8 are constructed. The former delivers insights into the compromise risk distribution under the current network configuration before implementing any defense action, whereas the latter addresses the situation after all actions are performed. Broadly speaking, the three options exhibit a similar

3 n actions can be sequenced in $n!$ variations.

Table 9.4: Statistical quantities about the equilibrium risk distributions of all games [7]‡

Game	Q2 (95%)			Q2 (75%)			Q3 Pr(risk > 6)		
	Decision option			Decision option			Decision option		
	A	B	C	A	B	C	A	B	C
G_0	9	9	9	9	9	9	0.5875	0.5875	0.5875
G_1	7	7	7	6	6	6	0.0987	0.0987	0.0987
G_2	7	7	7	5	5	5	0.0561	0.0784	0.0784
G_3	6	6	6	5	4	4	0.0374	0.0017	0.0017
G_4	6	6	6	5	4	4	0.0348	0.0012	0.0012
G_5	5	5	5	3	3	3	0.0012	0.0011	0.0010
G_6	5	5	5	3	3	3	0.0009	0.0007	0.0008
G_7	5	5	5	2	2	3	0.0008	0.0006	0.0007
G_8	5	5	5	2	2	2	0.0	0.0	0.0

positive effect of reducing the compromise risk as the chain progresses. As shown in Figure 9.5a, Figure 9.5b, and Figure 9.5c, the three options squeeze the risk probability mass towards the lower risk levels, in much the same manner.

Unlike classical game models with scalar-valued payoffs, the outcomes of the constructed chain are more comprehensive, thereby enabling a detailed analysis of the remediation impact of the respective options. They allow for drawing conclusions that are of utmost interest to the defender of power systems. In the studied use case, the defender is interested in the performance of the three decision options with respect to

Q1) what are the average risk values expected by each game in the decision chain?;

Q2) what is the maximal risk level that occurs in 95% and 75% of the cases in each game?; and

Q3) what are the chances of suffering a compromise risk of the category "6" or above after each step in the chain?

The answer to the question Q1 is provided by the results depicted in Figure 9.6. They show that the three decision options approximately lead to similar expected risk values over the whole chain progress. The drastic risk reduction is obtained directly by the outcome of G_1, in which the average risk is reduced from (6.429) to (4.111). The answers to Q2 and Q3 are more crucial to the defender as they give insights into the impact of the three decisions on the occurrences of high-level risks. Table 9.4 presents detailed statistical quantities about the obtained equilibrium risk distributions. The results show that the probability of suffering from a risk at level 6 or higher is reduced from 58.75% to 9.87% when having applied the game G_1. Moreover, as can be seen from Table 9.4 as well, the maximal risk level in 95% cases is also reduced from 9 to 7 when having applied the game G_1. Based on the results shown in Figure 9.6 and Table 9.4, the three

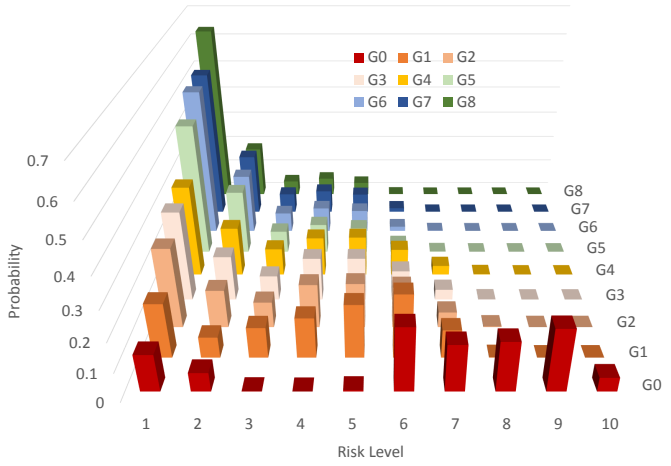

(a) OptionA

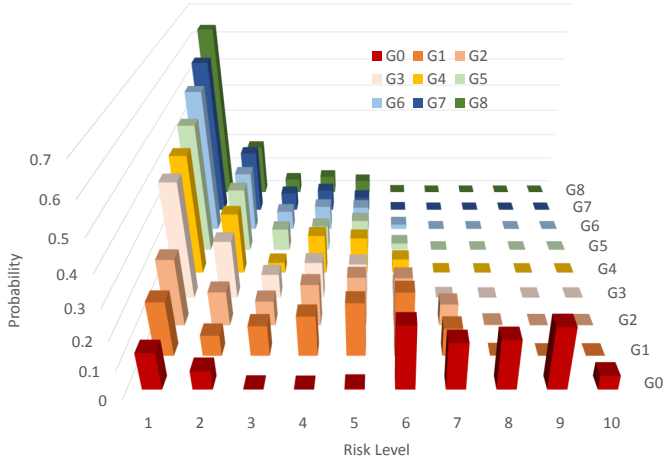

(b) OptionB

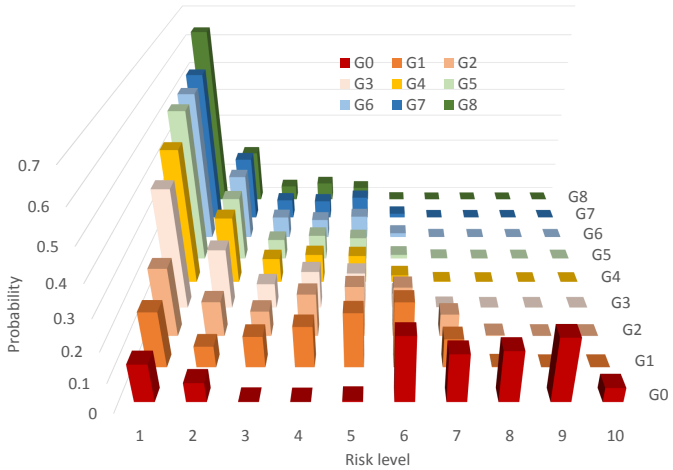

(c) OptionC

Figure 9.5: The comprehensive risk mitigation progress caused by the obtained prioritization options; i. e., OptionA, OptionB, and OptionC [7]♮

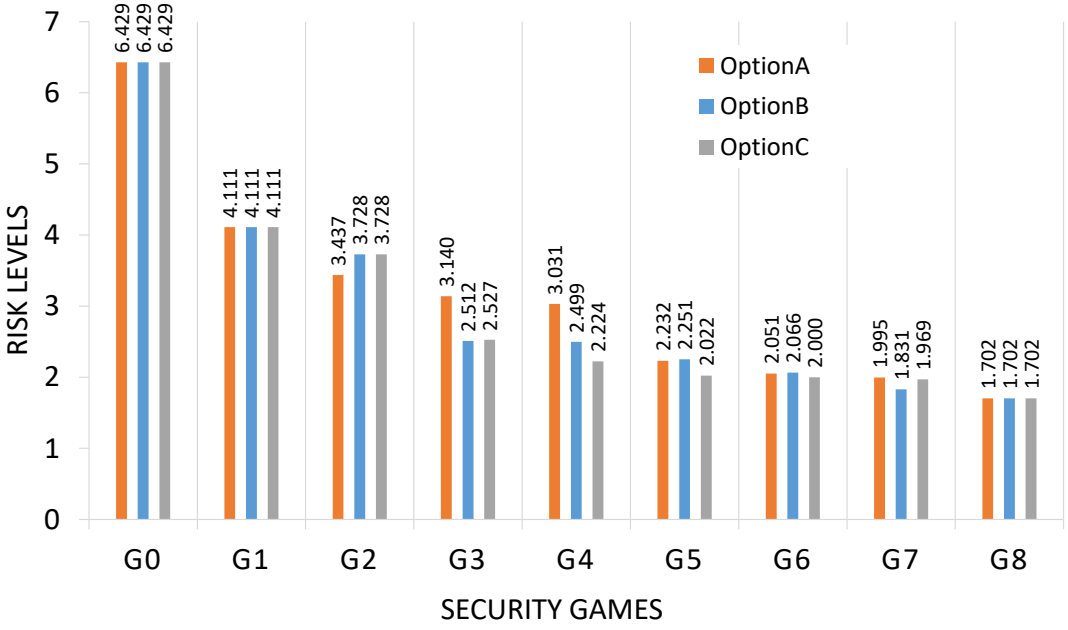

Figure 9.6: The average compromise risk values of the three obtained prioritization options (adapted from [7]↯)

options have almost similar remediation effects. More precisely, OptionA can result in a slightly better risk minimization after two steps (see G_2 effects). Nevertheless, OptionB and OptionC can compensate for this difference in the third step. That is, OptionB and OptionC can contribute slightly more beneficial effects if the decision constraints allow implementing three remediation actions in sequence.

9.7 SUMMARY

Due to their complexity and dynamic nature, electric power networks will always have a degree of vulnerability, making them attractive targets for remote adversaries with different intentions. An involved defender seeks to prioritize the possible remediation actions towards efficiently mitigating the risk of compromise stemming from exploiting vulnerabilities in such systems. As shown, even a small number of actions can create a large exploration space that demands a considerable effort for the defender. Unlike traditional IT defenders, who are commonly indifferent between decision options with equal expected utility (losses) even if one option might be riskier, defenders of electric power systems are more sensitive to extreme (risky) events due to the high criticality of such systems. Therefore, the cybersecurity games presented in this work can be employed to assist the defender in making risk-informed decisions on the action priorities. Given several constraints, the need for prioritization is evident in electric power systems. The presented prioritization approach enables the defender to quantize the remediation problem of the whole system into a finite set of manageable remediation actions. Even with scarce resources, the most critical actions will be performed first to help minimize the risk of compromise in an efficient manner [7]↯.

10 TWEAKABLE STOCHASTIC ORDER FOR CYBER INSURANCES

10.1 INTRODUCTION

While many game-theoretic models for cyber insurance try to set utility values so as to reflect a person's choices as accurately as possible, bounded rationality research has shown that many such attempts failed [229] (see Section 2.2.3 for further details on existing techniques for decision making under risk). This motivates paying more attention to the ordering relation itself upon which rational behavior is defined, which leads to the introduction of a *tweakable stochastic order*. This is a total ordering relation defined on random variables such as loss distributions known in actuarial science, which can be adapted to individual risk attitudes of players in an insurance game model. The idea and method are illustrated in this chapter using a straightforward bimatrix game model, in which the customer can decide to make a (false) claim, while the insurer is challenged with the decision of whether or not to audit the customer (and hence take additional costs and customer dissatisfaction into account). A uniform auditing policy applied to all customers obviously "approximates" all customers (honest and with a potential of fraud) by a single fixed customer model. Therefore, the presented tweakable stochastic order fills this room for improvement with an auditing policy tailored to the customer's risk attitude, so that the insurance can act more informed and accurate on the detection of fraud.

10.2 TWEAKABLE STOCHASTIC ORDER

Let X, Y be two random variables defined as follows:
$X = \{(x_0, p_0), (x_1, p_1), \ldots, (x_{n-1}, p_{n-1})\}$ and
$Y = \{(y_0, p_0), (y_1, p_1), \ldots, (y_{n-1}, p_{n-1})\}$ where

- $x_0 < \ldots < x_{n-1}$,
- $y_0 < \ldots < y_{n-1}$, and
- p_i is the occurrence probability of x_i or y_i.

The common support of X and Y can be defined as the interval $[s_0, s_1] = [\min\{x_0, y_0\}, \max\{x_{n-1}, y_{n-1}\}]$, though it is, for simplicity, assumed that both variables share the same support (i. e., $x_i = y_i$ for all i). This assumption is justified by the fact that most practical risk management uses categorical scales for impacts and likelihoods. The variables x_i, y_i may be directly associated with these categories, and the categorical scale is often fixed throughout the process as best practice standards recommend [93].

In practice, the random variables X and Y can represent the stochastic outcomes of assessment process of two security actions d_1 and d_2, respectively. They can also describe the assessment results of residual risks after applying d_1 and d_2. The question is, *which action is preferred to be performed?* In fact, answering this question depends first and foremost on the preferences of the involved decision-maker. These preferences

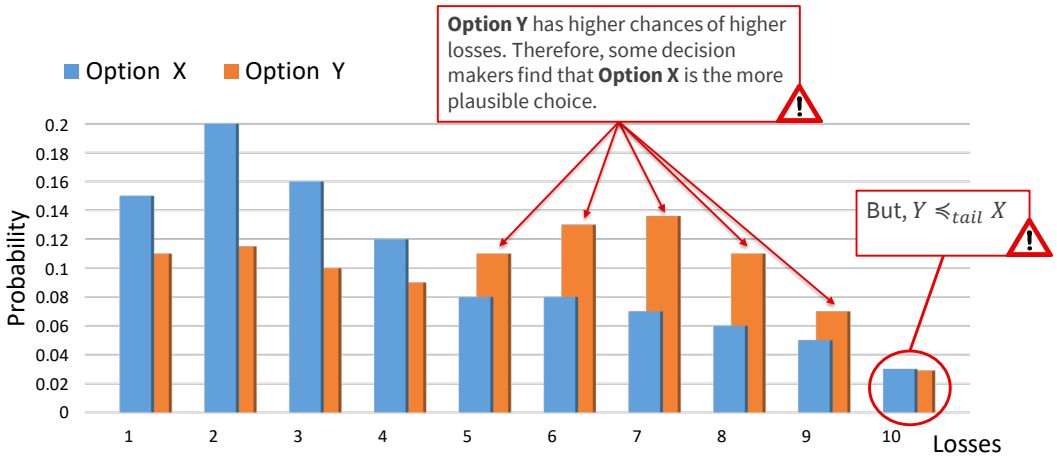

Figure 10.1: Drawbacks of the stochastic tail order (here, 10 is the highest loss category)

can be established based on several factors including the context and characterization of the investigated systems (risk-appetite and tolerance statements), the knowledge level of the decision-maker, and definitively the set of assumptions adopted by the decision-maker. For example, a decision maker involved in the protection process of CI systems seems to be more sensitive to extreme (risky) events due to the high criticality of such systems [7]$^{\natural}$. Therefore, he would prefer actions in which extreme events are less likely to happen. To model the aforementioned extreme risk preference, one can rely on the stochastic tail order $\preceq$ defined in [168], as shown above in this thesis. However, the drawback of such preference relation is that it is very pessimistic and very sensitive to assessment outcomes and observations (i. e., the decision is very sensitive to the occurrence probability (or frequency) of the most extreme events).

For the sake of illustration, consider the decision options shown in Figure 10.1. In this example, a very slight difference on the loss category 10 is enough to make a decision that Option Y is preferable over X. The reason is that the stochastic tail order gives excessive priority to the highest payoff (loss) dimension leading to potentially inefficient exploitation of other information encoded in the loss distribution. In other words, several CI decision makers can see that Option Y is riskier and X is the more plausible choice as the masses put on other high loss categories are less under Option X than under Y.

The second drawback of the tail order is that the $\preceq$-based decision analysis is static[1] and inconsistent (unable to cope) with the different attitudes of decision makers towards extreme risks. Moreover, it does not account for the changes in the decision caused by a learning process or changes in the context.

1 Static means that the comparison is the same for all scenarios independent of the decision-maker attitudes. The decision analysis process based on the expected value theory is also static as the calculation is the same for all scenarios. Being static means ignoring the fact that when decision makers are faced with (repeated) choices between X and Y, they choose differently and the same decision maker can choose X in some instances and Y in others due to domain and knowledge differences or changes caused by a learning process.

10.2.1 *Approach*

To overcome those limitations as well as to ensure the adaptivity to different attitudes towards risks, decision-makers should take into account not only the occurrence probabilities of extreme outcomes but also those of other high levels of risk. This can be done by looking at the risk contributions of possible outcomes belonging to upper tail areas of a loss distribution to the total expected risk value of the respective distribution. This could be similar to using higher moments to evaluate uncertain actions. However, using higher moments such as skewness and kurtosis are reliable only if the underlying sample data are normally distributed. When the assumption of normality is violated the results of the analysis can be at least misleading up to becoming completely wrong. Thus, these moments are unstable. The approach of tweakable stochastic orders targets prudent decision-makers who account for certain higher levels of risk rather than only extreme and perhaps less probable risks in order to manage available resources more reasonably. The approach should reward the action with mass function concentrated on lower risk and punish the one with mass function concentrated on higher risk categories based on preferences of the involved decision-maker. Basically, the goal is to develop a flexible comparison framework that allows decision-makers to accommodate the comparison according to their needs. The comparison uses the lexicographical ordering mechanism that has several theoretical properties enabling the development of an adjustable comparison framework. Therefore, each decision choice described using a random variable (loss distribution) will be evaluated as a sequence of values. Each value in the sequence represents the contribution of an upper-tail region to the total expected risk value, which is the last element in any evaluation sequence.

In other words, one can divide the probability density function of a random variable into $m \in \mathbb{N}^*$ overlapping partitions. Each partition represents an upper-tail region of the distribution, starting from a specific point to the end of the common support. Let us define a function $C : S \to \mathbb{R}^m$, which converts a distribution function into a sequence of risk values where S is the set of distribution functions with a common support $[s_0, s_1]$[2]. This function generates sequence values starting from the most right partition, which represents the most significant measure in the considered scenarios since higher damage is located to the right side of the considered distributions[3]. To evaluate a random variable X, the following steps can be applied:

i) Divide the support into m intervals so that $[s_0, s_1] = [a_0 = s_0, a_1, \ldots, a_{m-1}, a_m = s_1]$. That is, the points a_i for all $i \in \{0, \cdots, m-1\}$ represent the partitioning (cut) points of the respective support.

ii) Then, one can define the partitions $T_i = [a_i, a_m]\ \forall i \in \{0, \ldots, m-1\}$ and compute the value $c_i(X)$ in each interval T_i as given by

$$c_i(X) = \sum_{k=a_i}^{a_m} x_k \cdot p_k. \tag{10.1}$$

2 A prior work on stochastic orders [166] needed to assume $s_1 > s_0 \geqslant 1$; an assumption that one can abandon in this approach.

3 If the higher damage is located to the left side of a loss distribution as in economic risks, the order of the partitions should be reversed and the sequence values as well

This yields a sequence $C(X) = \{c_{m-1}(X), c_{m-2}(X), \cdots, c_0(X)\}$ of m values used to represent X in a decision problem. Each entry c_i is perceived as the contribution of the partition T_i to the overall expected risk.

iii) To make the setting adjustable to individual (i. e., subjective) risk attitudes, one needs to

a) fix the number m of partitions,

b) and more importantly, fix the points $a_1, a_2, \ldots, a_{m-1}$ at which the risk range is divided.

The question about how to pick m is easy to heuristically answer since one can just use the number of categories that the underlying risk management process proposes (typically, these are three to five categories; see [93] for examples).

The second question about how to determine the cut points a_i between s_0 and s_1 is more interesting, as it involves an understanding of the risk attitudes, as explained in Section 10.2.2. Following prospect theory, one could go ahead by rewriting Equation (10.1) into using $u(x_k) \cdot \pi(p_k)$ within the summation, where u can be a Bernoulli utility function to express the subjective valuation $u(x_i)$ of the objective outcome x_i, and with π being a weighting function to cancel out systematic effects of underrating high and overrating low risks (e. g., using a Prelec function [157] for π is one possible choice).

10.2.2 *Tailoring the ordering to subjective risk attitudes*

Psychological studies of risk attitudes (see [15, 89, 216] for examples) have discovered that risk loving, neutral or averse behavior is expressible by a utility weighting function that is either concave from below (for a risk avoider), linear (for risk neutrality), or convex (for risk seeker). More specifically, and referring to the aforementioned utility weighting function, the decision-making person will take action and receive some objective consequence (payoff or loss) measured by the choice $Z \in \{X, Y\}$, but subjectively valuates this value as $u(Z)$ rather than Z. Risk aversion then means that the expected subjective utility is less than the felt average, i. e., $E(u(Z)) < u(E(Z))$, a concavity condition on u. Likewise, risk seeking behavior rates the subjective utility higher than it actually is, which is a convexity condition on u being $E(u(Z)) > u(E(Z))$.

Given this information, one can "linearize" the individual risk ratings by defining the partitions to reflect *the regions that have the most influence on the subject's decision making*. To make this rigorous, let u_i be a function specifically associated with an *individual*, so the subscript i is added as a reminder here. Let us assume that u_i is

1. continuous
2. monotonously increasing
3. bounded, say, w.l.o.g., within $[0, 1]$.

One can equidistantly divide the range $[0, 1]$ into m intervals $q_0 = 0 < q_1 < \ldots < q_m = 1$, and define T_j for $j = 1, 2, \ldots, m-1$ as the q_j-th quantile of u_i, just like quantiles are defined for probability distribution functions. Note that this is similar to prior work of [46] on risk premium calculations; this work being mostly different from that in its

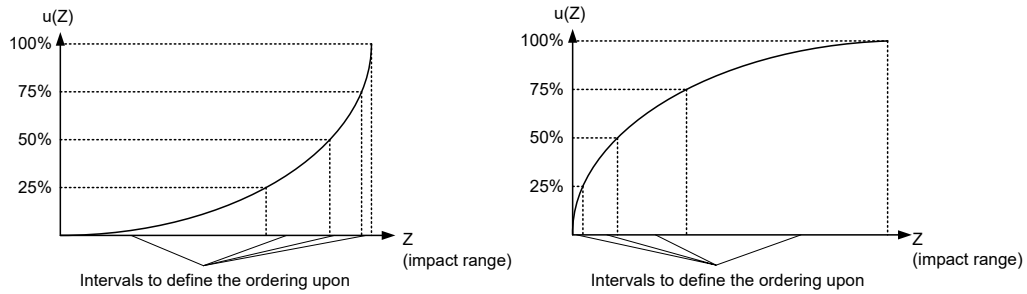

(a) Risk seeker's utility weighting (b) Risk avoiders's utility weighting

Figure 10.2: Definition of intervals based on subjective risk attitudes

focus on the decision making, and in its use of distributions themselves, rather than statistics derived from them.

Figure 10.2 shows examples on how risk aversion and risk loving behavior induces different such partitions (the case of risk neutrality is omitted, since it would be a humble straight line between the two images). It is worth noting that if the assessments (payoffs) are defined on a *loss scale*, the shape of the functions for an avoider vs. that for a risk seeker *interchange* from what is described previously and depicted in Figure 10.2. Namely, a risk seeker would have a convex utility weighting, expressed as the condition $\mathrm{E}(u(Z)) > u(\mathrm{E}(Z))$ for a *utility* Z. Losses are expressible as utilities with a negative sign, thus by multiplying the last inequality by -1 and using the linearity of the expectation operator, one can find the risk seeking behavior to require $(-1) \cdot \mathrm{E}(u(Z)) = \mathrm{E}(-u(Z)) < -u(\mathrm{E}(Z))$, in which one can think of $-u$ as the aforementioned loss. This, however, means that the risk seeker will have a concave shape of the utility weighting function on a loss scale, and likewise changed to a convex shape for the risk avoider. Consistently with what one intuitively expects, a risk averse person would probably apply a more fine-grained view in areas of high losses, while a risk seeking person may be less picky among outcomes of high loss magnitudes; reflected in the intervals becoming more fine-grained towards lower losses.

Remark. It is important to stress that the abscissa in both diagrams of Figure 10.2 reflects the actual loss scale *before* any categorization thereof. That is, the horizontal axis shows the real monetary or quantified loss on which categories for risk management are defined in the first place. The developed ordering does not directly work with these categories, and only uses their number as a guideline to choose the number of partitions, but makes the definition of categories depend on the subjective utility weighting function. This is important, because if the definition of categories were fixed, it may no longer match with the subjective understanding of a, say "low" or "high" risk, as these terms depend on a person's subjective attitude towards what *is* low or high (expressed in the shape of the function u_i).

10.2.3 *Defining the ordering*

Representing each distribution as a sequence of overlapping partitions provides a tractable solution to make the definition of preference relations more dynamic and tuneable. The overlapping partitioning provides the decision makers with more com-

plete, factual, and less-smoothed information about all viable decision options and their associated trade-offs. Moreover, it enables performing a comparison between two random variables using the lexicographical ordering $\leqslant_{lex}$, defined as:

$$\forall X, Y \in \mathbb{R}^n, C(X) \leqslant_{lex} C(Y) \iff \exists i \in \{0, \cdots, m-1\}, \\ (\forall j > i, c_j(X) = c_j(Y)) \wedge (c_i(X) < c_i(Y))$$

where $|C(X)| = |C(Y)| = m$.

The stochastic order $X \preceq Y$ between two random variables X, Y is then set equal to the lexicographic order of the representative vectors $C(X)$ and $C(Y)$; formally,

$$X \preceq Y \iff C(X) \leqslant_{lex} C(Y), \tag{10.2}$$

using the number m of partitions defined for the individual subject to choose between X and Y. Later in Section 10.3, an example of such a choice is given in the insurance domain.

It is important to remark that since the lexicographic ordering is defined upon an ordering within $\mathbb{R}$ (since $C(X), C(Y) \in \mathbb{R}^m$), it is automatically a *total* stochastic order (unlike many other stochastic orders; see [137]) without additional efforts.

10.3 CYBER RISK AND INSURANCE: A USE CASE

This section proposes a new view on game-theoretic models for cyber insurance, by incorporating subjective risk attitudes into the choice preference rules of players, rather than into the payoffs. As a simple showcase application, this work considers a use case of an insurance facing the decision of whether or not to audit a customer upon a claim.

10.3.1 *Game model*

Often, insurances collect profile information about their customers for matters of pricing and risk calculations, but they may also use this information to adapt their auditing policy to their customers individually. That is, an insurance may assign each customer i a risk attitude profile u_i, being precisely the utility weighting function discussed in prior sections, and based on this, determine the customer's likelihood of fraudulent claims. A simple way of computing a worst-case such estimate is offered by game theory: the situation is a two-person nonzero-sum game between the customer and the insurance, where the customer has the action set $AS_1 = \{\text{be honest, fraudulent claim}\} = \{\text{HO,FR}\}$, and the insurance has the action set $AS_2 = \{\text{pay, be sceptic and audit}\} = \{\text{PA,AU}\}$, both of which are abbreviated in the following.

It is straightforward to specify possible outcomes for both players in all four scenarios induced by this modeling, where the four cases are considered separately:

- honest claim and payment without audit (HO,PA): in that case, the insurer will pay L, while the customer has only a reduced cost of C, since most of the cost is taken by the insurance (if not all); setting $L = C = 0$ is also possible if both parties are supposed to have just followed their contract so that no loss or gains are made on either side.

- honest claim but skeptic insurer auditing (HO,AU): the insurer has audit costs A to check the eligibility, but learns that the customer has been honest. So, it has to pay the claim's amount L plus the costs of an audit, so the payoff comes to $L + A$. Actually, this may even offend the customer, which increases the potential loss for the insurer (as the customer may leave), so one may set the payoff for the customer to some larger loss $> C$ (not done in the following).
- fraudulent claim with the insurer paying (FR,PA): in that case, the customer gains L while the insurer loses L at the same time.
- fraudulent claim with the insurer auditing (FR,AU): the audit process may not be perfect, so the customer may have some probabilistic payoff that can be positive or negative, depending on whether the fraud was detected or not. Likewise, the insurer may have to pay either only the audit cost, or the audit cost plus the claim amount. For simplicity, let us assume that the audit is reliable; thus, the fraudulent claim will be discovered, and penalized with a fine F for the customer and with the insurer being left just with the audit costs A.

The resulting game is thus a straightforward 2×2 bimatrix game with payoff structure for both minimizing players:

	PA	AU
HO	$(+C, +L)$	$(+C, L + A)$
FR	$(-L, +L)$	$(+F, +A)$

(10.3)

with the model parameters all being $\geqslant 0$, and summarized as

C: residual costs (e. g., retained amount) after insurance payment
L: insurer's paid amount (insured lot)
A: audit costs
F: fines if a fraudulent claim is discovered

All four of these quantities are naturally random values, and the simple (naive) approach to solving the game is merely taking expectations over them to replace the random variables by real-valued quantities. This obvious method is exactly what the proposed stochastic order shall avoid, and in fact improve by letting the rationality of the players be implemented in a decision making process that respects the risk attitudes of both players.

The stochastic order, defined actually as a lexicographic order on the partitioning based on the individual risk attitude, then defines the choice rule for player 1 (i. e., the customer) to pick the best among the two strategies HO and FR. Likewise, the insurer can define its own choice rule based on internal risk management policies (following similar or other rationales as discussed here, but this is usually a confidential part of the insurer's business strategy).

The definition of equilibria for games over stochastic orders is doable along the same lines as for conventional games. For formal details, [167, 175] reconstructs the entire theory of games based on any total stochastic order.

10.3.2 *Practical use and meaning of the lexicographic Nash equilibrium*

An equilibrium in the game is, as for any other game, an optimal (possibly randomized) choice rule over an infinitude of repetitions of the game (10.3). It comes as a pair of distributions $((\mathrm{HO}, p^*), (\mathrm{FR}, 1-p^*))$ for the customer and $((\mathrm{PA}, q^*), (\mathrm{AU}, 1-q^*))$ for the insurer. The value p^* is thus the optimal probability of making honest claims, and the value $1-q^*$ is the optimal probability of auditing *that* particular customer. Since the ordering is individual for different customers, each of the insurance's clients may thus have its own value q^*. As a lexicographic Nash equilibrium, any deviation will indirectly cause losses for the deviating player in regards of a more important payoff dimension [175] (cf. Section 11.1).

10.3.3 *Equilibrium computation*

While the computation of equilibria is generally involved, the particularly simple structure of the game as being 2×2 admits a simple iterative online-learning algorithm, known as *fictitious play*. As explained in Chapter 4, this method just "simulates" gameplay between the two players, each of which keeps his own record about the opponent's choices made so far to adapt the next choice accordingly. Classical results of [134] assure that this learning process always converges for 2×2-games, as long as proper tie-breaking rules are imposed (see [135]), i.e., the game is non-degenerate, meaning that if the matrix game is given as

	PA	AU
HO	(a_{11}, b_{11})	(a_{12}, b_{12})
FR	(a_{21}, b_{21})	(a_{22}, b_{22})

the conditions

$$a_{11} - a_{21} - a_{12} + a_{22} \neq 0, \text{ and}$$
$$b_{11} - b_{21} - b_{12} + b_{22} \neq 0$$

are satisfied. The choice of parameters in the game (10.3) needs to respect these conditions if fictitious play should be used to compute an equilibrium[4].

Algorithm 2 illustrates the fictitious play process to solve the presented model of stochastic insurance games. The code is based on an implementation of fictitious play for stochastic orders of another kind [12]$^\natural$, and modifies this previous work from zero-sum arbitrary-shape to nonzero-sum 2×2 games. The fictitious play iteration can stop after a fixed number of iterations, or as soon as the empirical frequencies reach a steady state. With regard to the latter, the convergence of Algorithm 2 is based on the convergence analysis of the generalized version of fictitious play as shown in [164]. More precisely, let $\mathbf{x}_{(k)}$ denotes the empirical frequencies of strategy choices as recorded by Algorithm 2 in line 13. Fix any precision degree $\epsilon > 0$ and some vector-norm on $\mathbb{R}^2$ (e. g., $\|\cdot\|_\infty$), and terminate the algorithm as soon as $\frac{1}{k} \left\|\mathbf{x}_{(k+1)} - \mathbf{x}_{(k)}\right\| < \epsilon$. This work uses the precision value $\epsilon = 0.001$.

4 The tie-breaking rules (or diagonal property) of 2×2 games ensures that the games have fictitious play property (FPP); i. e., every fictitious process converges to an equilibrium.

Algorithm 2 Fictitious Play - 2×2 nonzero-sum insurance games

Require: Two non-degenerate 2×2 matrices of loss distributions $\mathbf{A}_{\text{customer}} = (a_{ij})$ and $\mathbf{A}_{\text{insurer}} = (b_{ij})$; two risk attitude functions – ins for the insurer and $cust$ for the customer.

Ensure: an approximation of an equilibrium pair $(\mathbf{x}, \mathbf{y})$

initialize $\mathbf{x} = (x_1, x_2) \leftarrow (0, 0), \mathbf{y} = (y_1, y_2) \leftarrow (0, 0)$
$r \leftarrow rand(rows)$ ▷ customer starts with a random (row) strategy
$c \leftarrow rand(cols)$ ▷ insurer starts with a random (column) strategy
$\mathbf{v} \leftarrow (b_{r1}, b_{r2})$ ▷ keep record of made choices
$x_r \leftarrow x_r + 1$
$\mathbf{u} \leftarrow (a_{1c}, a_{2c})$ ▷ keep record of made choices
$y_c \leftarrow y_c + 1$
$k \leftarrow 1$ ▷ iteration counter
while $not\ converged$ **do** ▷ exit the loop upon convergence
    $u^* \leftarrow$ the $\preceq$-$minimum(cust, \mathbf{u})$ ▷ the $\preceq$-$minimum$ is computed w.r.t. the tweakable order and the risk attitude of the customer
    $r \leftarrow$ the index of u^* in $\mathbf{u}$
    $\mathbf{v} \leftarrow \mathbf{v} + (b_{r1}, b_{r2})$
    $x_r \leftarrow x_r + 1$
    $v^* \leftarrow$ the $\preceq$-$minimum(ins, \mathbf{v})$ ▷ the $\preceq$-$minimum$ is computed w.r.t. the tweakable order and the risk attitude of the insurer
    $c \leftarrow$ the index of v^* in $\mathbf{v}$
    $\mathbf{u} \leftarrow \mathbf{u} + (a_{1c}, a_{2c})$
    $y_c \leftarrow y_c + 1$
    $k \leftarrow k + 1$
end while
Normalize $\mathbf{x}, \mathbf{y}$ to unit total sum ▷ turn $\mathbf{x}, \mathbf{y}$ into probability distributions
$\mathbf{ld}_{\text{customer}} = \mathbf{x}^T.\mathbf{A}_{\text{customer}}.\mathbf{y}$ ▷ the equilibrium loss distribution for the customer
$\mathbf{ld}_{\text{insurer}} = \mathbf{x}^T.\mathbf{A}_{\text{insurer}}.\mathbf{y}$ ▷ the equilibrium loss distribution for the insurer
return $p^* \leftarrow x_1$, $q^* \leftarrow y_1$, $\mathbf{ld}_{\text{customer}}$, and $\mathbf{ld}_{\text{insurer}}$

10.3.4 *Example*

For the sake of illustration, let us consider an application of empirical game theory with the tweakable ordering as defined above. Since the statistics are most conveniently handled with R, the exposition in the following will implicitly refer to this system [159].

Precisely, let us assume that there are empirical loss estimates available, being subjective expert opinions. Let us further presume a non-informative setting, i. e., the modeler cannot (or does not want to) specify any particular parametric setting, and instead, resorts to a nonparametric distribution by a mix of Gaussian densities to get an approximate loss distribution from the empirical data. The hypothetical modeler does so for all four scenarios of the game in Equation (10.3), giving rise to the payoff structures shown in Figure 10.3. Note that for the visual presentation, the matrices are separated into two distinct ones, constituting the payoffs for the players that are jointly (in pair notation) given in Equation (10.3).

In this use case, the payoffs are defined on a *loss scale*, and the choices of attitudes functions are $u(x) \propto x^2$ for the convex, and $u(x) \propto \sqrt{x}$ for the concave shape, with the proportionality factors set to reasonably span the function over the loss range on the abscissa, and within $[0, 1]$ on the vertical axis (so that one can use the internal functions of R to compute quantiles). Risk neutrality is expressed by the trivial straight line $u(x) = x$. The quantiles to define the partitioning are equidistant 20% steps along the vertical axis (Figure 10.2 used 25% steps).

The lexicographic Nash equilibrium in terms of the presented tweakable loss function is found from fictitious play (see Algorithm 2), by picking a random strategy for both players to start with. Then, they choose the individually $\preceq$-best reply among the so-far recorded choice frequencies $\mathbf{x}_{(k)} = \frac{1}{k} \cdot$ [number of rounds where HO has been played, number of rounds where FR has been played] and $\mathbf{y}_{(k)}$ defined alike. In round $k+1$, the next choice is then per player $\preceq$-minimal from the loss vector $\mathbf{x}_{(k)}^{\mathsf{T}} \cdot \mathbf{A}_{\text{insurer}}$, and $\mathbf{A}_{\text{customer}} \cdot \mathbf{y}_{(k)}$.

The game is analyzed using fictitious play for each of the nine possible combinations of both players being risk seekers, avoiders or neutral, and approximated the equilibrium for each of these combinations. Table 10.1 summarizes the results.

The results indicate that, at least for the studied hypothetical example, the risk attitude does have *some* impact: the likelihoods for the insurance to audit are largest for a risk avoiding attitude, and lowest for the risk seeking behavior (with risk neutrality locating the insurer in the middle). Similarly, if the customer is a risk avoider, it has higher chances to act honestly, while a risk seeking customer will have a larger probability for a fraud attempt. Interestingly, the numeric impact of the risk attitude on the probabilities is still not very strong, which can be attributed to a continuous dependence of the quantiles, and hence the ordering, on the shape of the utility weighting function u. That is, unless the graphs (like plotted in Figure 10.2) are significantly different in the sense of an extreme rise near the left or right end of the scale, the quantiles may come up "approximately equal" in the partitioning, thus yielding to roughly the same decisions. This, compared with a smooth loss distribution, can explain the experimental results on the differences to be there, but the risk attitude having not too much of an impact on the auditing policy or fraud incentive. The example was, for illustrative purposes, chosen with smooth loss distributions, but more "irregular" choices are of course possible, and

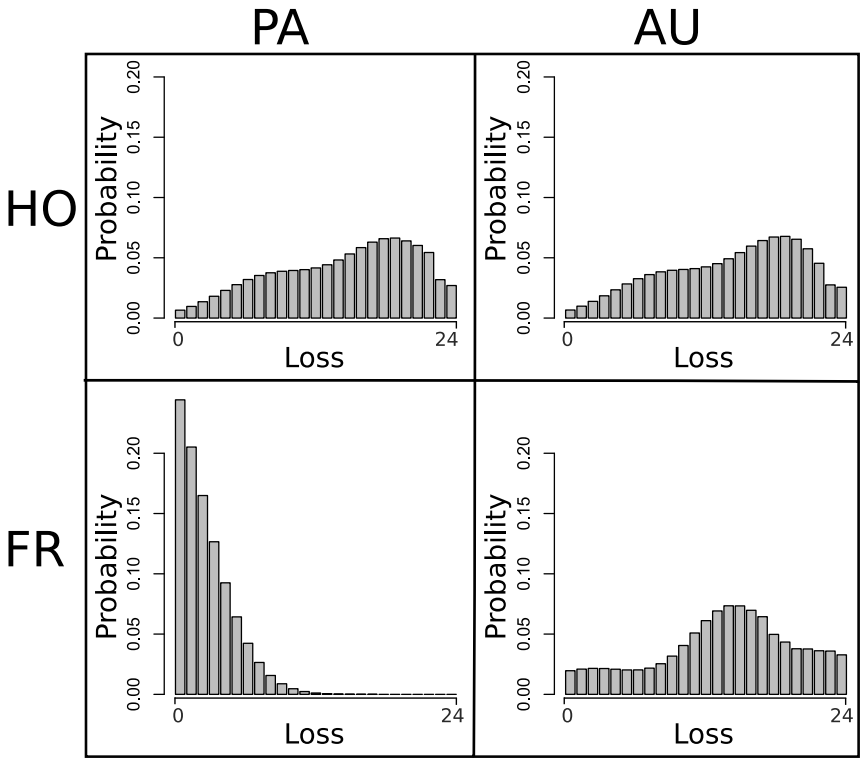

(a) Empirical loss estimates, payoff matrix $\mathbf{A}_{\text{customer}}$, for the customer

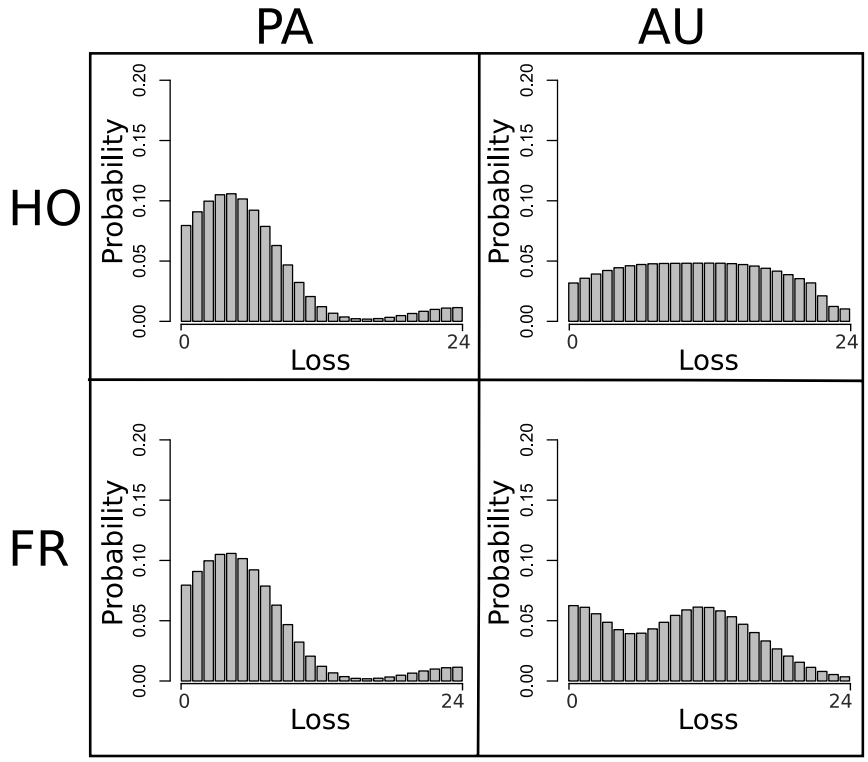

(b) Empirical loss estimates, payoff matrix $\mathbf{A}_{\text{insurer}}$, for the insurer

Figure 10.3: Empirical loss distributions for the insurance example

Table 10.1: Lexicographic Nash equilibria for the Example

PLAYER ATTITUDES (Customer-Insurer)	LEX. NASH EQUILIBRIUM			
	Customer		Insurer	
	Pr(HO)	Pr(FR)	Pr(PA)	Pr(AU)
Neutral-Neutral	0,582	0,418	0,056	0,944
Avoider-Neutral	0,586	0,414	0,210	0,790
Seeker-Neutral	0,000	1,000	0,000	1,000
Neutral-Avoider	0,985	0,015	0,001	0,999
Avoider-Avoider	0,985	0,015	0,193	0,807
Seeker-Avoider	0,000	1,000	0,000	1,000
Neutral-Seeker	0,280	0,720	0,061	0,939
Avoider-Seeker	0,282	0,718	0,210	0,790
Seeker-Seeker	0,000	1,000	0,000	1,000

yield larger differences in the decision making. Parametric models for losses, on the contrary, may be more likely to be smooth.

The author does not argue for or against a particular choice of loss distributions, nor can he, based on this example, anyhow claim that the utility weighting functions should or should not be quadratic or square root shaped (The choices were made to resemble a hyperbolic absolute risk aversion class, which is the most general class of utility functions based on the Arrow-Pratt coefficient[5]). Instead, the message is the *easy possibility* to include such attitudes in empirically justified decisions and game-theoretic models, at a reasonably cheap cost. A real benefit, however, seems only obtainable if the modeling of both, the loss distribution and the utility weighting function is made on a careful basis; a matter of research and methods beyond the scope of this current work. This work is limited to the conclusion that it is possible and "cheap" to account for risk attitudes, *but* the impact of such an account strongly depends on how "careful" the losses and utility weights are being modeled (see [68] for related empirical studies).

10.4 QUALITY ANALYSIS

As discussed above, the main idea of the presented model is to integrate the risk attitudes into the decision-making process. Nevertheless, the comprehensive nature of the outcomes of the presented distribution-valued insurance games can be further used to help the hypothetical modeler to judge on the quality of the attitude-based decisions in terms of several dimensions such as players' satisfaction. After having computed the equilibria of the different scenarios, the expected losses incurred from

5 The absolute risk aversion coefficient of a utility function $u(x)$ is defined as the ratio of the concavity and the slop of the utility function $A(x) = -\frac{u''(x)}{u'(x)}$ [156]. A utility function exhibits a hyperbolic absolute risk aversion if its Arrow-Pratt coefficient of absolute risk aversion is a hyperbolic function; that is, $A(x) = -\frac{u''(x)}{u'(x)} = \frac{1}{ax+b}$. The sign of $A(x)$, which equals that of $-u''(x)$, has a meaning close to the willingness to accept risks; a negative (positive) sign of $A(x)$ implies concavity (convexity) of $u(x)$.

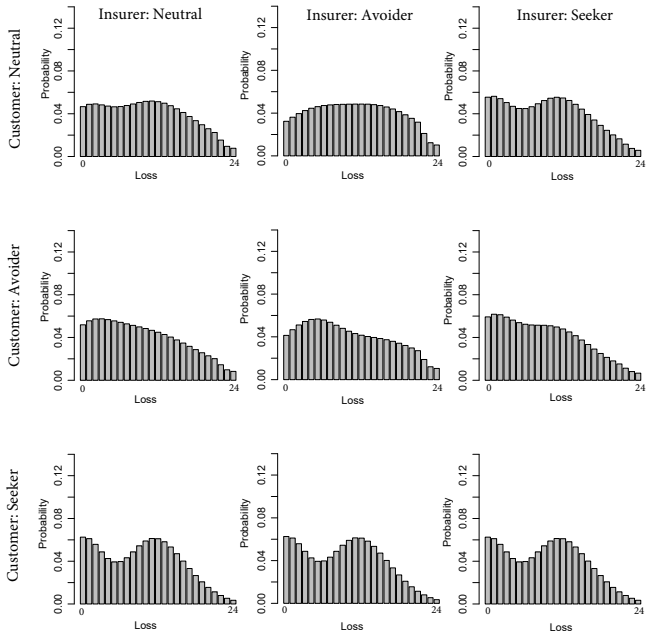

Figure 10.4: The nine equilibrium loss distributions for the insurer

implementing each equilibrium strategy are quantified using the loss distributions $\mathbf{ld}_{\text{customer}}$ and $\mathbf{ld}_{\text{insurer}}$ as described in line 21 and 22 of Algorithm 2. Those distributions allow drawing further conclusions and comparative measurements that can be of significant interest to the modeler or involved players. The analysis can involve one or more of the following points:

- Q1) what is the average loss value assured to each player following the computed equilibrium strategy in each specific scenario?
- Q2) what is the probability that the incurred loss exceeds the players' expectations?
- Q3) what is the maximum loss that could occur in a specific amount, say 95%, of the cases in each scenario?
- Q4) what is the probability of suffering the maximum possible loss (in the studied example, the maximum loss is 24)?
- Q5) what is the probability of suffering loss equal or higher than specific loss level, e. g., losses $\geqslant 20$?

The main goal of the aforementioned analysis dimensions is to summarize the impact of each attitude-based decision using a set of quantities of interest, which are ultimately employed to provide insights into how satisfying the different risk attitudes will be.

10.4.1 *The insurer*

As shown in Table 10.1, there are nine equilibria computed for the nine possible attitude scenarios (or combinations) of both players. Each equilibrium results in a loss distribution $\mathbf{ld}_{\text{insurer}}$ for the insurer. Figure 10.4 shows the nine equilibrium loss distributions, which represent the basis for answering the above questions using a set of summary statistics. More precisely, let $\mathbf{ld}$ be a random variable describing an equilibrium loss distribution for the insurer in one of the nine possible scenarios:

Table 10.2: Statistical quantities of the equilibrium loss distributions for the insurer

PLAYER ATTITUDES (Customer-Insurer)	STATISTICAL QUANTITIES				
	Q1	Q2	Q3	Q4	Q5
Neutral-Neutral	10.0522	0.4679	21	0.0077	0.0809
Avoider-Neutral	9.3484	0.4573	21	0.0083	0.0756
Seeker-Neutral	9.3175	0.5040	19	0.0035	0.0436
Neutral-Avoider	10.9921	0.5196	21	0.0102	0.1100
Avoider-Avoider	9.9761	0.4852	21	0.0104	0.0980
Seeker-Avoider	9.3175	0.5040	19	0.0035	0.0436
Neutral-Seeker	9.5455	0.4989	20	0.0058	0.0616
Avoider-Seeker	8.9387	0.4930	20	0.0067	0.0594
Seeker-Seeker	9.3175	0.5040	19	0.0035	0.0436

- The answer to Q1 is given by the mathematical expectation of **ld**, i. e., $\mathrm{E}(\mathbf{ld})$.
- The answer to Q2 is reached by computing the disappointment rate of **ld** given by $\mathrm{d}(\mathbf{ld}) = \Pr(\mathbf{ld} > \mathrm{E}(\mathbf{ld}))$.
- The answer to Q3 is computed using the 95-th quantile of **ld**.
- With regard to Q4 and Q5, the answers are given using the probabilities $\Pr(\mathbf{ld} = 24)$ and $\Pr(\mathbf{ld} \geqslant 20)$, respectively.

Given the studied example, the different statistical quantities are summarized in Table 10.2. Those quantities are further used to analyze the decision quality of each insurer attitude with respect to the three customer attitudes being risk neutral, avoider, and seeker. This can be achieved as follows: firstly, choose a specific customer attitude; secondly, for each analysis dimension $\{Q1, \ldots, Q5\}$, compare the responses of the three insurer attitudes. For example, if the customer has a neutral risk attitude, the best responses (here, best in the sense of lowest value) with regard to Q1, Q3, Q4, and Q5 are obtained by decisions associated with the risk-seeking insurer. In contrast, the lowest disappointment rate (i. e., Q2) is obtained by risk-neutral insurer. To simplify the comparison process, the results included in Table 10.2 are transformed into quality scores and illustrated using radar charts. The quality scores are computed as follows: for each analysis dimension Q_i, the best response is identified and denoted as $\min_{Q_i}$; then, each response x_{Q_i} with regard the respective dimension is assigned a quality score given by $\min_{Q_i}/x_{Q_i}$. The quality scores of the three insurer attitudes with respect to each customer attitude function are illustrated in Figure 10.5. The most interesting conclusion from this analysis is that for each customer risk attitude the seeker attitude of the insurer outperforms the other two attitudes in the sense of having the most satisfying impact. Based on such a result, the modeler might decide to adopt one specific attitude for the insurer regardless of the customer behavior.

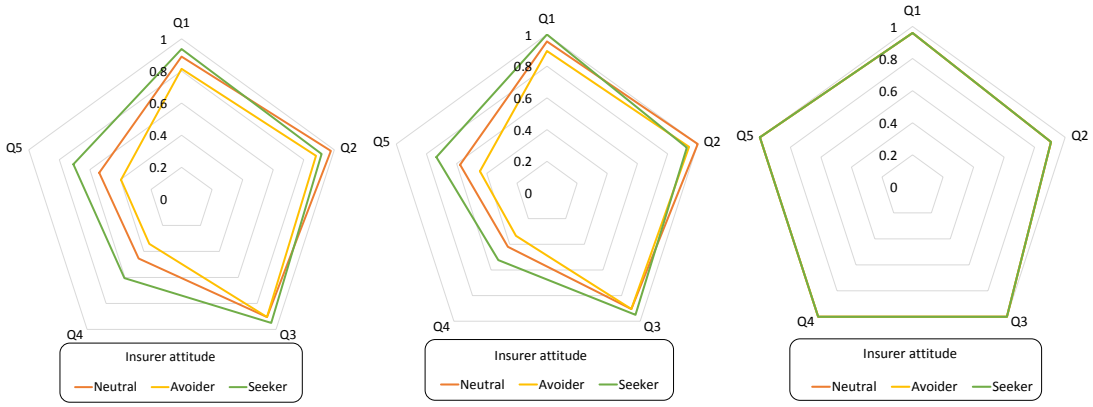

(a) Customer attitude: Neutral (b) Customer attitude: Avoider (c) Customer attitude: Seeker

Figure 10.5: The quality scores of the three insurer attitudes with respect to each customer attitude function

10.4.2 *The customer*

To judge on the quality of the customer risk attitudes, the same analysis procedure developed in Section 10.4.1 is applied on the produced equilibrium loss distributions for the customer $\mathbf{ld}_{\text{customer}}$ that are depicted in Figure 10.6. Those distributions are summarized using the aforementioned statistical quantities in Table 10.2, while the corresponding quality scores are computed and illustrated in Figure 10.7. In the same manner as the insurer case, the risk-averse attitude of the customer exhibits the most consistent behavior against the different insurer attitudes. As shown in Figure 10.7, the risk-averse customer delivers the best responses over all dimensions when the insurer attitude is seeker. When the insurer is neutral or avoider, a risk-seeking customer can deliver slightly better responses than a risk avoider only on the dimension Q5. That is, the customer can assure himself less chances of higher loss ranges $\geqslant 20$ if he adopted decisions associated with a risk-seeking attitude. Nevertheless, this benefit becomes less important in comparison with the dimension Q4 that gives insights into the probability of suffering the *maximum* possible loss that is evidently assured by the risk-averse customer. Finally, combining the results depicted in Figure 10.5 and Figure 10.7 leads to the conclusion that the best attitude scenario in the presented game is the one with the customer being risk avoider and the insurer being risk seeker. These results undoubtedly vary from one case study to another based on the available assessments.

10.5 SUMMARY

Even though the famous Debreu representation theorem [51] indicates that any continuous preference rule on $\mathbb{R}^n$ is expressible by a continuous real-valued function, not all choice rules may fall under this representability. The lexicographic ordering is one example of a discontinuous choice rule, yet it is easy and natural to use in risk management where categories of risks are important. Real numbers may, therefore, be insufficiently expressive to accurately describe a human's individual attitude towards risk. In some cases, it may be simpler to trade the complexity of a model over the reals for a more complex ordering on the payoffs in return of a simpler model to describe

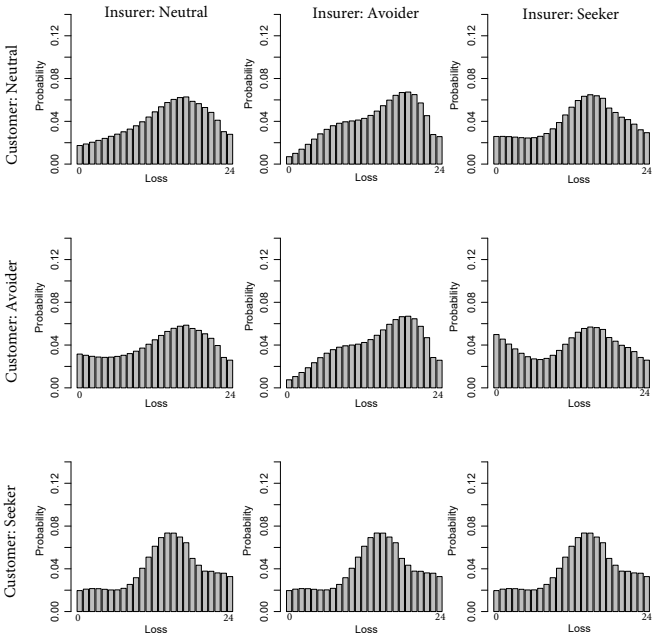

Figure 10.6: The nine equilibrium loss distributions for the customer

Table 10.3: Statistical quantities of the equilibrium loss distributions for the customer

PLAYER ATTITUDES (Customer-Insurer)	STATISTICAL QUANTITIES				
	Q1	Q2	Q3	Q4	Q5
Neutral-Neutral	13.7260	0.5588	23	0.0278	0.2005
Avoider-Neutral	13.0232	0.5246	23	0.0259	0.1907
Seeker-Neutral	13.6888	0.5544	23	0.0328	0.1803
Neutral-Avoider	14.1914	0.5338	23	0.0259	0.2207
Avoider-Avoider	14.1903	0.5342	23	0.0258	0.2233
Seeker-Avoider	13.6888	0.5544	23	0.0328	0.1803
Neutral-Seeker	13.3369	0.5385	23	0.0293	0.1842
Avoider-Seeker	12.1342	0.5318	23	0.0259	0.1657
Seeker-Seeker	13.6888	0.5544	23	0.0328	0.1803

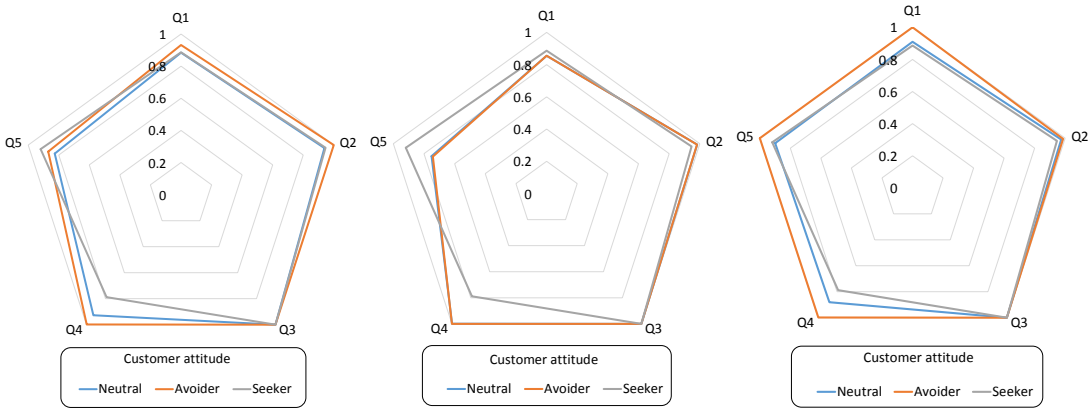

(a) Insurer attitude: Neutral (b) Insurer attitude: Avoider (c) Insurer attitude: Seeker

Figure 10.7: The quality scores of the three customer attitudes with respect to each insurer attitude function

the interaction between the players. For insurances, this work proposes the simplest game model to measure the likelihood of fraudulent vs. honest behavior and to get the best auditing strategy for the insurance. Such a model would have been widely inappropriate, because inaccurate, if it were defined over the reals, but can be made into a complex and flexible description of reality by resorting to a stochastic rather than a real ordering.

CONCLUSIONS, APPENDIX, AND BIBLIOGRAPHY

11 CONCLUSION AND OUTLOOK

This chapter recaps the contributions and results presented in this thesis. Furthermore, it discusses experiences and limitations distilled from investigating the new theory of distribution-valued games that can be taken into account in future approaches.

11.1 THE LEXICOGRAPHIC PARADOX

In this thesis, security strategies are perceived as the advice received from distribution-valued security games. They represent the best defense to be enforced towards minimizing the extreme risk (or equivalently minimizing the highest loss category) in the first place no matter what attack will be mounted, given a finite set of potential attacks. The core component of this kind of games is the preference order that enables judging on uncertain payoffs in a pairwise comparison. In this thesis, two variations have been investigated, namely the stochastic tail order and its generalization, which is the tweakable stochastic order. Both approaches establish basically a lexicographic order on payoff distributions. That is, let the payoff distributions X, Y be represented by (possibly infinite) vectors $\mathbf{x} = (x_1, x_2, x_3, \ldots)$ and $\mathbf{y} = (y_1, y_2, y_3, \ldots)$ respectively so that $X \preceq Y$ if and only if $\mathbf{x} \leqslant_{lex} y$. In light of this, the notion of "lexicographic games" is used here to denote games in which payoffs are vector-valued and each player seeks to optimize his own payoff in the sense of lexicographic ordering. As mentioned earlier in this thesis, the entire theory of games based on any total stochastic order has been put forth in [163, 164, 167] with formal proofs on the existence of Nash equilibria in such games, and hence security strategies thereof. The question is whether computed security strategies are lexicographically optimal as the theory promises. That is, if players play optimally, then there should not be any $\leqslant_{lex}$-better defense against anything that the adversary may do.

To compute security strategies, this thesis employs modified versions of fictitious play algorithm, which is an iterative learning scheme to compute a game equilibrium approximately. Following this algorithm, each player relies on information from previous iterations about choices made by his opponent to pick the action that satisfy his needs. In the presented zero-sum games, the defender picks in each iteration the action that lexicographically minimizes his losses while the attacker adapts to this but to maximize the defender's loss. Intuitively, the first coordinate, x_1, in each payoff vector $(x_1, x_2, x_3, \ldots)$ is the most important criteria in a lexicographic decision rule, followed by x_2, and so on. Roughly speaking, each coordinate x_i gets infinitesimal importance weight relative to its precedent x_{i-1} [81]. Thus, any decision-making process concerned with $\leqslant_{lex}$ is typically conducted in a sequence of steps starting from the most important coordinate; once a decision is made, all other less important coordinates will be completely neglected.

As a consequence, using lexicographic order to compare vector-valued payoffs of games is similar to run fictitious play on a stack of real-valued games. Each game has a real-valued payoff matrix corresponding to one coordinate of the respective payoff

structures. Hence, similar to the lexicographic choice mechanism, the best response in a lexicographic game will be determined by the real-valued game of the first coordinate (i. e., the highest risk category) regardless of other (less significant) games in the stack. Upon a tie (i. e., identical payoff values), the best response will be determined by the second game in the stack, and so on. More precisely, fictitious play, in the initial stage, is merely a game being played on the first coordinate of the payoff structures until it converges (i. e., the highest risk is minimized). By then, it breaks the tie[1] and moves on to the second coordinate, to continue the game from there, further adapting the strategies for both players. This, however, may invalidate the optimum found so far by the first coordinate game, thus causing fictitious play to immediately return to the first coordinate again to fix the issue due to the utmost importance of this coordinate. This means that the algorithm will end up with cycling between the first and the second coordinate until both have been optimized before moving further to other coordinates. This can practically mean that fictitious play may still take an infeasibly large number of iterations, unless the accuracy is set sufficiently coarse to accept ties within a certain proximity of the actual $\leqslant_{lex}$-optimum that might be impossible to reach in an efficient manner via online learning.

Therefore, optimizing less important coordinates can lead the defender to act suboptimally with respect to the most important one of the payoff structures. This violates the importance ordering of decision criteria (i. e., coordinates) imposed by the lexicographic principle. Consequently, the computation mechanism (here, fictitious play) responds immediately to fix this issue and keep the most important coordinate optimized. Therefore, fictitious play algorithm aims at satisfying the greatest possible number of games starting from the most important (dominant) game and going deeper in the respective stack. In the worst case, it converges to the optimum of the first game in the stack, which meets the concrete need to minimize the chance of extreme risk for security applications in CIs.

However, when players are at equilibrium in the most important game, the defender will be indifferent between the different defense options. As a result, his next decision will be made based on a less important coordinate of the payoff structures in compliance with the underlying lexicographic importance ordering. Theoretically, there can be another defense that improves the respective less coordinate given this specific circumstance (i. e., the attacker is fixed at his equilibrium strategy of the first game). Nevertheless, if the new defense breaks the optimum of the most important game, the defender will be prohibited from deviation by the lexicographic principle itself. Therefore, "lexicographic paradox" involves such situations, in which the computation mechanism seeks to strictly adhere to the lexicographic principle, leading to results that are not lexicographic-optimal. In light of this, the match between Nash equilibrium and lexicographic Nash equilibrium in lexicographic games cannot be always guaranteed. Further details on the definition of lexicographic Nash equilibrium can be obtained form a recent research work in [175].

Thus, to avoid this subtle issue, we need an exact procedure as is constructible by a sequence of linear programs. The aforementioned view of lexicographic optimization as a stack of real-valued games leads to this exact method of computing the optimum.

1 When the algorithm reaches a mixed strategy Nash equilibrium, the players must be indifferent between their actions realized by the mixed strategy. That is, those actions yield the identical expected payoff for the respective player (i. e., a tie occurs).

Briefly, the method starts with solving the most important real-valued game in the stack by converting it into a corresponding linear program. This will deliver the optimal assured payoff (aka the value) of the game ν_1. Afterward, the method moves on to solve the next game in the stack but with an additional constraint that the new solution should not worsen the achieved optimal value of the most important game (i. e., ν_1), and so forth. Hence, this method tries to refine the equilibrium in each step towards finding the most feasible one that optimizes the maximum number of games in the respective stack. The formal description of this method given in [164] in response to the lexicographic paradox.

11.2 CONTRIBUTIONS AND RESULTS

CI systems represent the main pillars for our modern society, as they provide essential services and fundamental networks upon which national economy and prosperity are significantly dependent. They are large-scale and complex systems that increasingly extend beyond their traditional borders to cope with the rapidly growing demand for their vital services. As a result, various systems, which were previously isolated from each other by clear and well-defined boundaries, might be spontaneously and seamlessly integrated into one system crossing the boundaries marked by their traditional individual perimeters [8]♮. The increasing connectivity and complexity of such systems make traditional best practice approaches insufficient to ensure security. The concept of extended perimeter is introduced to explain this phenomenon and how CIs are exposed to higher security risk than ever before. The identified extended perimeter components can significantly impair the security posture of existing CI systems and give potential attackers the opportunity to bypass security perimeter controls.

Moreover, risks are everywhere and involved in every serious or even trivial decision of our daily life. Thus, security decisions in CIs cannot be made seriously and meaningfully without a comprehensive assessment of their risk exposures. In this thesis, therefore, security management focuses essentially on how to configure and optimize security operations through changing the focus to (or aligning best-practice approaches with) security risk. This is aligned with the new definition of security as "*reducing the risk to critical infrastructure by physical means or defens(ive) cyber measures to intrusions, attacks, or the effects of natural or man-made disasters.*" [208].

In this thesis, security management is concerned with how to create a coherent security strategy that leaves organizations well-positioned in security games against potential adversaries. It defines a process of controlling and coordinating security practices such as conducting random patrols or regular spot-checks to prevent or deter potential intrusions as well as prioritization of vulnerability remediation actions. In such applications, there are several sources of uncertainty affecting security decisions, such as lack of reliable information about attackers' types and incentives as well as the randomness of consequences that actions have. Moreover, and due to their vital importance, CI systems should be designed and configured to cope with everyday situations; that is, decision-makers involved in protection operations of CIs aim to prepare those systems to withstand extreme conditions and worst-case scenarios. For that purpose, they prefer security strategies that prevent or minimize the occurrence of a catastrophic loss. Practically, this risk attitude guides decision-making processes in such critical systems and hence the security management process as well. Traditional

utility optimization techniques address situations in which decision-makers seek to optimize their (long-term) expected payoffs regardless of the severe consequences that the made decisions may imply. Therefore, security management in CIs adopts *the principle of extreme risk minimization*, not the traditional utility maximization principle. Towards fulfilling those requirements, this thesis casts security management problem into distribution-valued security games. This type of games accepts randomness as an inherent part of the payoffs and integrates the specific risk attitude into the decision-making process itself.

The core of security management involves a decision-making process, in which an involved decision-maker assesses possible defense choices towards finding optimal configurations and rules, thereby minimizing security risks. Towards streamlining security management operations, this thesis provides a methodological approach that integrates concepts and principles from risk management, decision theory, and game theory. The approach breaks down into smaller and manageable steps to support defenders of CIs to make risk-informed security decisions. Those steps are context establishment, identification of strategies, identification of goals, effectiveness assessment, identification of best response, and implementation of best response. The methodological approach presented in this thesis is compliant with the generic risk management methodologies like the one defined by ISO 31000 standard [99] and its other closely aligned standards such as ISO/IEC 27005 [98]. Furthermore, it can be easily integrated into another standard framework such as NIST 800-37 [58], too. Additionally, the presented methodology is suitable to develop a situational awareness process that aims to glue past, present, and future together to enhance the security posture of CI organizations.

To bridge the gap between defining a theoretical model and practically instantiating it, this thesis introduces the concept of physical surveillance games to address scenarios in which mobile agents perform random spot-checks within CI boundaries to improve flexibility and intrusion detection probabilities. Having dynamic and mobile surveillance strategies is highly important to maintain situational awareness even within a system complex so that potential intruders can still be detected. Physical surveillance games have several important real-life applications, such as physical border patrolling, scheduling random security checkpoints, public transit security, and fare enforcement planning, just to name a few. Those practices are particularly challenging because of limited security resources, uncertain consequences of surveillance actions, predictable monitoring patterns, a potential breach of employees' privacy and comfort, among others. The physical surveillance game model studied in this thesis takes all these factors into account and provides practitioners with security strategies that keep a balance between multiple goals.

To assess the consequences of identified surveillance strategies, a simulation model for physical intrusion problems in CIs is developed. The model allows establishing a faithful image of the studied physical environment, deployed personnel and their behavior, as well as potential attacks that may occur. Moreover, this thesis presents an entropy-based model to assess the impact of different inspection strategies on the preservation of location privacy. The model uses the technique of CTMC to quantify privacy as a function of time.

Using results obtained from physical surveillance games, a comparative analysis is conducted to achieve a better understanding of the differences between the new class of generalized games in which payoffs are probability distributions and its counterpart

of classical (real-valued) games. To achieve this, each randomized defense strategy is converted into a consistent security resource allocation (i. e., quasi-purified strategies). Afterward, the new defense is added to the defender's action set to validate whether it outperforms all previously identified defenses. One finding of this analysis is that this method works for stochastic orders, but fails for standard numeric orders, depending on how security is quantified [2]♮.

If we go for a (perhaps practically more reasonable) approach of including uncertainty in the simulation, then a stochastic order supports this resource sharing method. If, however, we average out the uncertainty to recover a perhaps more familiar numeric measure of security, then the equilibrium may be optimal, but the resource allocation derived from it may not be optimal. This counterintuitive phenomenon can be avoided by resorting to a more sophisticated ordering than the plain numerical order on the real numbers. The findings show that distribution-valued games' decisions can be more effective in practice [2]♮. Their empirically assessed consequences meet the defender's satisfaction in terms of closeness to his aspiration level (*ideal point*) and the disappointment rate.

Besides physical security, this thesis demonstrates and evaluates the application of the proposed security management approach to address cybersecurity problems. While IT networks offer sufficient controllability and observability of CI assets, they can expose those critical assets to cyber risks due to some cyber vulnerabilities that are not properly maintained. In CI systems, resolving all vulnerabilities at once might be infeasible because of several technical and economic factors that can significantly affect the patching and upgrading decisions of their components, including limited time and budget as well as legal constraints [7]♮. To figure out where to start, an involved defender has to prioritize the possible vulnerability remediation actions prudently. The key objective of prioritization is to efficiently reduce the inherent security risk to which the system in question is exposed [7]♮. Therefore, this thesis presents an integrated risk-based decision-support methodology for prioritizing possible remediation activities according to their risk mitigation impact. It leverages the TTC security metric to quantitatively assess the risk of compromise. The developed risk estimator considers several factors, including (i) the inherent assessment uncertainty, (ii) interdependencies between the network components, (iii) different adversary skill levels, and (iv) public vulnerability and exploit information.

Technically, the remediation actions are successively prioritized with the aid of a chain of cybersecurity games. The chain depends on a general game-theoretical model with distribution-valued payoffs to account for the process of decision-making under uncertainties. The game model benefits from a stochastic tail order to incorporate the risk attitude, imposed by the criticality of the investigated electric power systems, into the decision-making process. The power system case study shows that even a small number of remediation actions can create a large exploration space that demands a huge effort for the defender. The key goal of the developed framework is achieved by constructing the prioritization tree. It supports the defender in making risk-informed decisions about the prioritization of the possible security actions. The tree represents a tremendous reduction of the decision space that the defender needs to explore. In the examined system, the framework ends up with 3 prioritization options out of 40320 possible prioritization variations of the 8 identified defense actions [7]♮.

Finally, the thesis presents the approach of tweakable stochastic order, which seeks to extend and generalize the existing stochastic tail order by focusing on how to make an individual risk attitude part of a preference order. The approach employs existing risk attitude models (expressible by utility weighting functions) to linearize the individual risk ratings by defining partitions to reflect the regions that have the most influence on the subject's decision making. It enables performing a comparison between two random variables using lexicographic order. Representing actions with uncertain consequences as sequences of overlapping partitions provides a tractable solution to make the definition of preference relations more dynamic and tuneable. The overlapping partitioning provides decision makers with more complete, factual, and less-smoothed information about all viable decision options.

To illustrate the idea and method of tweakable stochastic order, this thesis considers a use case for an insurance facing the decision whether or not to audit a customer upon a claim (and hence take additional costs and customer dissatisfaction into account). A uniform auditing policy applied to all customers approximates all customers (honest and with a potential of fraud) by a single fixed customer model. To avoid this, the proposed tweakable stochastic order seeks to adjust the auditing policy according to the customer's risk attitude, so that the insurance can act more informed and accurate on the detection of fraud. To analyze the impact, this thesis casts the insurance problem into a two-player nonzero-sum game with vector-valued payoffs. The results indicate that, at least for the considered hypothetical example, the risk attitude does have some impact: the likelihoods for the insurance to audit are largest for a risk-avoiding attitude, and lowest for the risk-seeking behavior (with risk neutrality locating the insurer in the middle). Similarly, if the customer is a risk avoider, it has higher chances to act honestly, while a risk-seeking customer will have larger probability for a fraud attempt. This approach employs the lexicographic preference to construct a tweakable decision-making process. Hence, it provides a possible mitigation of the aforementioned lexicographic paradox by allowing the decision-makers to adapt the value on the most important coordinate of a lexicographic sequence according to their needs.

11.3 OUTLOOK ON FUTURE RESEARCH

It seems to be very interesting and promising to continue deepening on the investigation and improvement of the theory that combines games and stochastic orders to handle different challenges of real-life problems.

The study performed in Chapter 7 provides a new interpretation of mixed strategies in general, coming with the surprising property that optimality of an equilibrium can fail under certain, though natural, implementations of it. This reveals that optimality of a defense is not the same as optimizing a security score, since the means by which security is quantified and optimized play a much deeper role than intuitively expected. A resulting open research question derived from this observation concerns formal explanations of this effect, as well as conditions that characterize when this sub-optimality of equilibria can or cannot occur [2]♯.

Concerning the risk assessment approaches presented in this thesis, the TTC model can be extended to address the overall attack surface of organizations, including social and organizational factors. Social engineering attacks and changing policies might cause shorter paths to compromise target systems. Besides the compromise risk, decision

constraints such as limited time and budget can be integrated into the decisions-making process through defining proper action-response models. Moreover, the prioritization framework introduced in Section 8.4 exhibits a high degree of flexibility. Thus, it can support the defender to address multiple target components at the same time. This can be achieved by extending the attacker action space to include compromise graphs of different targets. Furthermore, the same framework can be exploited to obtain risk-based vulnerability prioritization through a proper adaptation of the defender attack space to address specific vulnerabilities [7]♮.

Tweakable stochastic order shows the ease of constructing stochastic orders in a form that can naturally embody risk attitudes of persons in a(ny) game-theoretic model. Lifting game theory from real-valued orderings to more complex orderings is a task that has been accomplished in the past literature and can be very simple in special cases like 2×2 insurance games that this thesis used. It will be an interesting aisle of unexplored potential to see what other game models are amendable to tweakable stochastic orders, and to study the degree to which bounded rationality can be captured by such extended game models.

Besides the lexicographic paradox observed in fictitious play algorithm, lexicographic order is generally criticized for preventing any tradeoff and balance between decision criteria (i. e., coordinates). Therefore, security management in less critical systems calls for a more relaxed approach. In light of this, one future research direction can include substituting the lexicographic choice rule used in the tweakable order approach by a multiobjective optimization problem, in which each element (i. e., criteria) of a generated sequence would correspond to a distinct decision objective with a predefined importance weight. In this regard, the common weighted sum approach can be used to keep balance between different criteria such as extreme and expected risk values according to their predefined weights.

A

APPENDIX

A.1 FICTITIOUS PLAY IN A TWO-PERSON ZERO-SUM GAME WITH DISTRIBUTION-VALUED PAYOFFS

Algorithm 3 shows a generalized version of the fictitious play for a minimizing defender (player 1) in a zero-sum game with distribution-valued payoffs. It represents an iterative method of computing the defender's security strategy $\mathbf{x}^*$ by letting each player starts from an initial guess for his optimal pair of (payoff, strategy) and updates (i. e., improves) his behavior according to the best of his so-far recorded knowledge about the opponent's choices.

Algorithm 3 Fictitious Play (with $\preceq$-minimizing defender), adapted from [164]

Require: an $(n \times m)$-matrix A of payoff distributions $A = (F_{ij})$
Ensure: an approximation $(\mathbf{x}, \mathbf{y})$ of an equilibrium pair $(\mathbf{x}^*, \mathbf{y}^*)$
1: initialize $\mathbf{x} \leftarrow \mathbf{o} \in \mathbb{R}^n$, and $\mathbf{y} \leftarrow \mathbf{o} \in \mathbb{R}^m$
2: $r \leftarrow$ the row index giving the $\preceq$-minimum over all column-maxima
3: $c \leftarrow$ the column index giving the $\preceq$-maximum over all row-minima
4: $\mathbf{u} \leftarrow (F_{1c}, \ldots, F_{nc})$
5: $y_c \leftarrow y_c + 1$ ▷ $\mathbf{y} = (y_1, \ldots, y_m)$
6: $\mathbf{v} \leftarrow \mathbf{o}$ ▷ initialize $\mathbf{v}$ with m function that are zero everywhere
7: $k \leftarrow 1$ ▷ iteration counter
8: **while** not *converged* **do** ▷ exit the loop upon convergence
9: $u^* \leftarrow$ the $\preceq$-minimum of $\mathbf{u}$ ▷ best response to player 2's actions
10: $r \leftarrow$ the index of u^* in $\mathbf{u}$ ▷ record the current choice of player 1
11: $\mathbf{v} \leftarrow \mathbf{v} + (F_{r1}, \ldots, F_{rm})$ ▷ update player 2' payoffs based on player 1's choice; pointwise addition of functions
12: $x_r \leftarrow x_r + 1$ ▷ update player 1's behavior; $\mathbf{x} = (x_1, \ldots, x_n)$
13: $v^* \leftarrow$ the $\preceq$-maximum of $\mathbf{v}$ ▷ best response to player 1's actions
14: $c \leftarrow$ the index of v^* in $\mathbf{v}$ ▷ record the current choice of player 2
15: $\mathbf{u} \leftarrow (F_{1c}, \ldots, F_{nc})$ ▷ update player 1' payoffs based on player 2's choice
16: $y_c \leftarrow y_c + 1$ ▷ update player 2's behavior
17: $k \leftarrow k + 1$
18: **end while**
19: Normalize $\mathbf{x}, \mathbf{y}$ to unit total sum ▷ turn $\mathbf{x}, \mathbf{y}$ into probability distributions
20: **return** $\mathbf{x}^* \leftarrow \mathbf{x}$, $\mathbf{y}^* \leftarrow \mathbf{y}$, and $F(\mathbf{x}^*, \mathbf{y}^*) \leftarrow \sum_{i,j} F_{ij} \cdot x_i \cdot y_j$ ▷ $\approx (\mathbf{x}^*)^T \mathbf{A} \mathbf{y}^*$

The fictitious play process can stop after a fixed number of iterations (e. g., 1000 iterations), or as soon as the empirical frequencies reach a steady state. With regard to the latter, let $\mathbf{x}_{(k)}$ denotes the empirical frequencies of strategy choices as recorded by Algorithm 3 in line 12. Fix any tolerance degree $\epsilon > 0$ and some vector-norm on $\mathbb{R}^2$ (e. g., $\|\cdot\|_\infty$), and terminate the algorithm as soon as $\frac{1}{k} \left\| \mathbf{x}_{(k+1)} - \mathbf{x}_{(k)} \right\| < \epsilon$ [164]. Upon termination of the algorithm, x will approximate the security strategy $\mathbf{x}^* \in \Delta(SP_{\mathcal{D}})$, and

y approximates an optimal adversarial mixed strategy $\mathbf{y}^* \in \Delta(SP_{\mathcal{A}})$. Throughout this thesis, the accuracy parameter ϵ is set to 0.001.

A.2 DERIVATION OF ET EQUATION

The basic derivation of Equation (8.3), which models the expected number of tries ET until the adversary can develop a fully working exploit code, is discussed in [117, 129]. As explained in Table 8.1, let

- n be the number of known vulnerabilities visible at node `DEST`, and
- C be the average number of vulnerabilities for which an exploit can be found or crafted by an adversary given his skill level S [7][♮].

Hence, one can define $n - C$ as the number of useless (unexploitable) vulnerabilities with the skill level S.

Let p_k refer to the probability that an adversary is successful in crafting a functioning exploit code to take advantage of a vulnerability randomly chosen from those remaining after $(k-1)$ failed tries and p'_k the probability of the corresponding complement event.

- For $k = 1$ (i. e., before excluding any vulnerability from previous tries): $p_1 = \frac{C}{n}$ and $p'_1 = \frac{n-C}{n}$.
- For $k = 2$ (here, the total number of vulnerabilities to choose from is $n-1$ after excluding the vulnerability chosen in the first try): $p_2 = \frac{C}{n-1}$ and $p'_2 = \frac{n-C-1}{n-1}$.
- This yields that: $p_k = \frac{c}{n-k+1}$ and $p'_k = \frac{n-C-k+1}{n-k+1}$ for $1 \leqslant k \leqslant n-C+1$. Here, one can remark that the probability that the adversary is not successful in exploiting a vulnerability chosen from those remaining after $n-C+1$ tries is zero because all useless vulnerabilities are tried in the previous $n-C$ tries; i. e., $p'_k = 0$.

Afterward, let Pr_k be the probability that an adversary with skill level S needs k tries to success. Under the assumption of independent events, one can find that $Pr_k =$ (the probability that an adversary is successful in crafting a functioning exploit to take advantage of a vulnerability randomly chosen from those remaining after $k-1$ failed tries)$\times$(the probability that the adversary is not successful in exploiting the vulnerabilities chosen in the first $k-1$ tries). This yields:

$$
\begin{aligned}
Pr_k &= p_k \times \prod_{i=1}^{k-1} p'_i \\
&= \frac{C}{n-k+1} \times \prod_{i=1}^{k-1} \frac{n-C-i+1}{n-i+1} \\
&= \frac{C}{n} \times \prod_{i=2}^{k} \frac{n-C-i+2}{n-i+1}; \quad 2 \leqslant k \leqslant n-C+1
\end{aligned}
$$

To compute the expected value of tries ET, one can apply:

$$
\begin{aligned}
ET &= \sum_{k=1}^{n-C+1} k \times Pr_k \\
&= \frac{C}{n} \times (1 + \sum_{k=2}^{n-C+1} [k \times \prod_{i=2}^{k} (\frac{n-C-i+2}{n-i+1})]) \\
&= S \times (1 + \sum_{k=2}^{n-C+1} [k \times \prod_{i=2}^{k} (\frac{n-C-i+2}{n-i+1})])
\end{aligned}
$$

BIBLIOGRAPHY

PUBLICATIONS BY THE AUTHOR

[1] Ali Alshawish, Mohamed Amine Abid, and Hermann de Meer. "Game-Theoretic Optimization for Physical Surveillance of Critical Infrastructures: A Case Study." In: *Game Theory for Security and Risk Management: From Theory to Practice*. Ed. by Stefan Rass and Stefan Schauer. Cham: Springer International Publishing, 2018. Chap. 15, pp. 353–389. ISBN: 9783319752679. DOI: 10.1007/978-3-319-75268-6_15.

[2] Ali Alshawish, Mohamed Amine Abid, and Hermann de Meer. "Quasi-purification of mixed game strategies: Sub-optimality of equilibria in security games." In: *Computers & Security* 87.101575 (2019). ISSN: 0167-4048. DOI: 10.1016/j.cose.2019.101575. URL: http://www.sciencedirect.com/science/article/pii/S0167404819300458.

[3] Ali Alshawish and Hermann de Meer. "Prioritize When Patching Everything is Impossible!" In: *2019 IEEE 44th Conference on Local Computer Networks (LCN)*. IEEE. Osnabrück, Germany: IEEE, 2019, pp. 125–128. DOI: 10.1109/LCN44214.2019.8990847.

[4] Ali Alshawish and Hermann de Meer. "Risk-based Decision-Support for Vulnerability Remediation in Electric Power Networks." In: *Proceedings of the Tenth ACM International Conference on Future Energy Systems*. ACM, 2019, pp. 378–380. DOI: 10.1145/3307772.3330157.

[5] Ali Alshawish, Stefan Rass, and Hermann de Meer. "A Game-theoretical Decision-making Framework for Physical Surveillance Games." In: *Workshop on Novel Approaches in Risk and Security Management for Critical Infrastructures*. Vienna, Austria: Austrian Institute of Technolog AIT, 2017.

[6] Ali Alshawish, Korbinian Spielvogel, and Hermann De Meer. "A Model-based Time-to-Compromise Estimator to Assess the Security Posture of Vulnerable Networks." In: *2019 International Conference on Networked Systems (NetSys'19)*. Garching b. München, Germany: IEEE, 2019. DOI: 10.1109/NetSys.2019.8854511.

[7] Ali Alshawish and Hermann de Meer. "Risk mitigation in electric power systems: Where to start?" In: *Energy Informatics* 2.34 (2019). DOI: 10.1186/s42162-019-0099-6.

[8] Ali Alshawish, Mohamed Amine Abid, Zhiyuan Sui, Hermann de Meer, Antonios Gouglidis, and Stefan Rass. *HyRiM Deliverable 4.3: How to Enhance Perimeter Security Using New Surveillance Technologies*. 2017. URL: https://hyrim.net/wp-content/uploads/2017/12/HyRiM-D4.3-How-to-Enhance-Perimeter-Security-using-new-Surveillance-Technologies.pdf.

[9] Ali Alshawish, Mohamed Amine Abid, Stefan Rass, and Hermann de Meer. "Playing a Multi-objective Spot-checking Game in Public Transportation Systems." In: *Proceedings of The 4th Workshop on Security in Highly Connected IT Systems (SHCIS'17)*. Neuchâtel, Switzerland: ACM, June 2017, pp. 31–36. DOI: 10.1145/3099012.3099019.

[10] Ali Alshawish, Mohamed Amine Abid, Hermann de Meer, Stefan Schauer, Sandra König, Antonios Gouglidis, and David Hutchison. "G-DPS: A Game-theoretical Decision-making Framework for Physical Surveillance Games." In: *Game Theory for Security and Risk Management: From Theory to Practice*. Ed. by Stefan Rass and Stefan Schauer. Cham: Springer International Publishing, 2018. Chap. 6, pp. 129–156. ISBN: 9783319752679. DOI: 10.1007/978-3-319-75268-6_6.

[11] Antonios Gouglidis, Benjamin Green, David Hutchison, Ali Alshawish, and Hermann de Meer. "Surveillance and security: protecting electricity utilities and other critical infrastructures." In: *Energy Informatics* 1.15 (2018). DOI: 10.1186/s42162-018-0019-1.

[12] Stefan Rass, Sandra König, and Ali Alshawish. *HyRiM: Multicriteria Risk Management using Zero-Sum Games with vector-valued payoffs that are probability distributions*. R package version 1.0.3. Austrian Institute of Technology (AIT), 2019. URL: https://hyrim.net/software/.

[13] Stefan Rass, Ali Alshawish, Mohamed Amine Abid, Stefan Schauer, Quanyan Zhu, and Hermann de Meer. "Physical Intrusion Games-Optimizing Surveillance by Simulation and Game Theory." In: *IEEE Access* 5.16904577 (2017), pp. 8394–8407. DOI: 10.1109/ACCESS.2017.2693425.

REFERENCES

[14] Daron Acemoglu, Azarakhsh Malekian, and Asuman Ozdaglar. *Network security and contagion*. Tech. rep. National Bureau of Economic Research, 2013.

[15] Binay Kumar Adhikari and Anup Agrawal. "Does local religiosity matter for bank risk-taking?" In: *Journal of Corporate Finance* 38 (2016), pp. 272–293. ISSN: 09291199. DOI: 10.1016/j.jcorpfin.2016.01.009.

[16] Peter Ahrendt. "The multivariate gaussian probability distribution." In: *Technical University of Denmark, Tech. Rep* (2005).

[17] Martin Aigner and Michael Fromme. "A game of cops and robbers." In: *Discrete Applied Mathematics* 8.1 (1984), pp. 1–12.

[18] Tansu Alpcan and Tamer Başar. *Network Security: A Decision and Game Theoretic Approach*. Cambridge University Press, 2010.

[19] Brian Alspach. "Searching and sweeping graphs: a brief survey." In: *Le matematiche* 59.1, 2 (2006), pp. 5–37.

[20] Eitan Altman, Thomas Boulogne, Rachid El-Azouzi, Tania Jiménez, and Laura Wynter. "A survey on networking games in telecommunications." In: *Computers & Operations Research* 33.2 (2006), pp. 286–311.

[21] Bo An, David Kempe, Christopher Kiekintveld, Eric Shieh, Satinder Singh, Milind Tambe, and Yevgeniy Vorobeychik. "Security games with limited surveillance." In: *Ann Arbor* 1001 (2012), p. 48109.

[22] BSI: Bundesamt für Sicherheit in der Informationstechnik. *BSI-Standard 100-1: Information Security Management Systems (ISMS)*. 2008. URL: https://www.bsi.bund.de/SharedDocs/Downloads/EN/BSI/Publications/BSIStandards/standard_100-1_e_pdf.pdf?__blob=publicationFile&v=1 (visited on 11/20/2019).

[23] BSI: Bundesamt für Sicherheit in der Informationstechnik. *IT-Grundschutz-Kompendium*. 2018.

[24] Kenneth Barnes and Briam Johnson. *National SCADA test bed substation automation evaluation report*. Tech. rep. Idaho National Laboratory (INL), 2009.

[25] Matthew P Barrett. *Framework for improving critical infrastructure cybersecurity, version 1.1*. NIST Cybersecurity Framework. National Institute Of Standards and Technology(NIST), 2018. DOI: 10.6028/NIST.CSWP.04162018.

[26] David E Bell. "Disappointment in decision making under uncertainty." In: *Operations research* 33.1 (1985), pp. 1–27.

[27] Rocco Bellanova and Michael Friedewald. "Deliverable 1.1: Smart Surveillance–State of the Art." In: *SAPIENT. FP7 Sapient Project, Brussels* (2011).

[28] Daniel Bernoulli. "Exposition of a New Theory on the Measurement of Risk." In: *Econometrica* 22.1 (1954), pp. 23–36. ISSN: 00129682, 14680262. URL: http://www.jstor.org/stable/1909829.

[29] Jonathan Berr. *WannaCry ransomware attack losses could reach $4 billion*. May 2017. URL: https://www.cbsnews.com/news/wannacry-ransomware-attacks-wannacry-virus-losses/ (visited on 09/19/2018).

[30] Sourabh Bhattacharya, Tamer Başar, and Maurizio Falcone. "Surveillance for Security as a Pursuit-Evasion Game." In: *Proceedings of The International Conference on Decision and Game Theory for Security*. Springer, 2014, pp. 370–379.

[31] Zhaohong Bie, Yanling Lin, Gengfeng Li, and Furong Li. "Battling the extreme: A study on the power system resilience." In: *Proceedings of the IEEE* 105.7 (2017), pp. 1253–1266.

[32] Pavlo R. Blavatskyy. "Stochastic expected utility theory." In: *Journal of Risk and Uncertainty* 34.3 (2007), pp. 259–286. ISSN: 08955646. DOI: 10.1007/s11166-007-9009-6.

[33] Rainer Bohme and Galina Schwartz. "Modeling Cyber-Insurance: Towards A Unifying Framework." In: *Proceedings of the Workshop on the Economics of Information Security (WEIS)*. Harvard, 2010.

[34] Zoltan Bojthe, Levente Meszaros, Benjamin Seregi, Rudolf Hornig, and Andras Varga. *INET Framework: An open-source OMNeT++ model suite for wired, wireless and mobile networks.* June 2016. URL: https://inet.omnetpp.org/ (visited on 12/10/2017).

[35] Gunter Bolch, Stefan Greiner, Hermann De Meer, and Kishor S Trivedi. *Queueing networks and Markov chains: modeling and performance evaluation with computer science applications*. John Wiley & Sons, 2006.

[36] Jean Bolot and Marc Lelarge. "Cyber Insurance as an Incentivefor Internet Security." In: *Managing information risk and the economics of security*. Springer, 2009, pp. 269–290.

[37] Richard B. Borie, Craig A. Tovey, and Sven Koenig. "Algorithms and Complexity Results for Pursuit-Evasion Problems." In: *IJCAI*. Vol. 9. 2009, pp. 59–66.

[38] Jack Brassil. "Technical challenges in location-aware video surveillance privacy." In: *Protecting Privacy in Video Surveillance*. Springer, 2009, pp. 91–113.

[39] Matthew Brown, Sandhya Saisubramanian, Pradeep Reddy Varakantham, and Milind Tambe. "STREETS: game-theoretic traffic patrolling with exploration and exploitation." In: (2014).

[40] *CVE-2017-0144 Detail*. Mar. 2017. URL: https://nvd.nist.gov/vuln/detail/CVE-2017-0144 (visited on 09/19/2018).

[41] CVSS. *The Common Vulnerability Scoring System (CVSS)*. 2015. URL: https://nvd.nist.gov/vuln-metrics/cvss (visited on 09/19/2018).

[42] Tracy Camp, Jeff Boleng, and Vanessa Davies. "A survey of mobility models for ad hoc network research." In: *Wireless communications and mobile computing* 2.5 (2002), pp. 483–502.

[43] Enrique Castillo. *Extreme value theory in engineering*. Elsevier, 2012.

[44] Thierry Chauveau. *Subjective risk and disappointment*. Documents de Travail du Centre d'Economie de la Sorbonne, revised version 2012.63. 2012.

[45] Thierry Chauveau and Nicolas Nalpas. *A Theory of Disappointment*. online: http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.318.9141, [retrieved: March 26, 2018]. 2005.

[46] Jean-Paul Chavas and Kwansoo Kim. "Aversion to Risk and Downside Risk in the Large and in the Small under Non-Expected Utility: A Quantile Approach." In: *Theoretical Economics Letters* 05.06 (2015), pp. 784–804. ISSN: 2162-2078, 2162-2086. DOI: 10.4236/tel.2015.56090. URL: http://www.scirp.org/journal/doi.aspx?DOI=10.4236/tel.2015.56090 (visited on 07/30/2019).

[47] Roger Clarke. "Information technology and dataveillance." In: *Communications of the ACM* 31.5 (1988), pp. 498–512.

[48] André van Cleeff and Roelf J Wieringa. "De-perimeterisation as a cycle: tearing down and rebuilding security perimeters." In: (2008).

[49] Germán Coloma. "The penalty-kick game under incomplete information." In: *University of CEMA Economics Serie Documentos de Trabajo* 487 (2012).

[50] DHS Risk Steering Committee (RSC). *DHS risk lexicon*. Tech. rep. US Deptartment of Homeland Security Washington, DC, 2010. URL: https://www.cisa.gov/dhs-risk-lexicon.

[51] Gerard Debreu. *Theory of value: an axiomatic analysis of economic equilibrium*. eng. 19. Dr. Monograph / Cowles Foundation for Research in Economics at Yale University 17. OCLC: 33936319. New Haven: Yale Univ. Press, 1987. ISBN: 978-0-300-01559-1.

[52] Francesco Maria Delle Fave, Albert Xin Jiang, Zhengyu Yin, Chao Zhang, Milind Tambe, Sarit Kraus, and John P. Sullivan. "Game-theoretic patrolling with dynamic execution uncertainty and a case study on a real transit system." In: *Journal of Artificial Intelligence Research* 50 (2014), pp. 321–367.

[53] Department of Homeland Security (DHS). *NIPP 2013: Partnering for critical infrastructure security and resilience*. National infrastructure protection plan. US Deptartment of Homeland Security Washington, DC, 2013. URL: https://www.cisa.gov/publication/nipp-2013-partnering-critical-infrastructure-security-and-resilience.

[54] American Heritage Dictionary. *Perimeter*. URL: https://www.ahdictionary.com/ (visited on 09/25/2019).

[55] Cambridge Dictionary. *Perimeter*. URL: https://dictionary.cambridge.org/ (visited on 09/25/2019).

[56] ENISA. *Meltdown and Spectre: Critical processor vulnerabilities*. Jan. 2018. URL: https://www.enisa.europa.eu/publications/info-notes/meltdown-and-spectre-critical-processor-vulnerabilities (visited on 09/26/2018).

[57] *External Events Excluding Earthquakes in the Design of Nuclear Power Plants*. Specific Safety Guides NS-G-1.5. Vienna: INTERNATIONAL ATOMIC ENERGY AGENCY, 2003. ISBN: 92-0-113603-X. URL: https://www.iaea.org/publications/6733/external-events-excluding-earthquakes-in-the-design-of-nuclear-power-plants.

[58] JOINT TASK FORCE. "Risk Management Framework for Information Systems and Organizations." In: *NIST Special Publication* 800 (2018), p. 37.

[59] Nicolas Falliere, Liam O Murchu, and Eric Chien. "W32. stuxnet dossier." In: *White paper, Symantec Corp., Security Response* 5.6 (2011), p. 29.

[60] Sadegh Farhang, Mohammad Hossein Manshaei, Milad Nasr Esfahani, and Quanyan Zhu. "A Dynamic Bayesian Security Game Framework for Strategic Defense Mechanism Design." In: *Decision and Game Theory for Security*. Ed. by Radha Poovendran and Walid Saad. Vol. 8840. Lecture Notes in Computer Science. Cham: Springer International Publishing, 2014, pp. 319–328. ISBN: 978-3-319-12600-5. DOI: 10.1007/978-3-319-12601-2-18.

[61] Federal Ministry of the Interior. *National Strategy for Critical Infrastructure Protection*. CIP Strategy. Federal Office for Civil Protection and Disaster Assistance (BBK) and Federal Office for Information Security (BSI), Germany, 2009. URL: https://www.kritis.bund.de/SharedDocs/Downloads/Kritis/EN/CIP-Strategy.pdf.

[62] Bogdan Filipič and Tea Tušar. "Visualization in multiobjective optimization." In: *Proceedings of the Genetic and Evolutionary Computation Conference Companion*. ACM. 2018, pp. 858–879.

[63] Melissa L Finucane, Paul Slovic, Chris K Mertz, James Flynn, and Theresa A Satterfield. "Gender, race, and perceived risk: The'white male'effect." In: *Health, risk & society* 2.2 (2000), pp. 159–172.

[64] James Flynn, Paul Slovic, and Chris K Mertz. "Gender, race, and perception of environmental health risks." In: *Risk analysis* 14.6 (1994), pp. 1101–1108.

[65] Fedor V Fomin and Dimitrios M Thilikos. "An annotated bibliography on guaranteed graph searching." In: *Theoretical computer science* 399.3 (2008), pp. 236–245.

[66] Catherine Forbes, Merran Evans, Nicholas Hastings, and Brian Peacock. *Statistical Distributions*. John Wiley and Sons Ltd, Nov. 2010. ISBN: 9780470390634.

[67] The Open Group Jericho Forum. *Business Rationale for De-Perimeterization*. White Paper. 2007. URL: https://publications.opengroup.org/w127.

[68] Irwin Friend. "The Demand for Risky Assets." en. In: *Financial Dec Making Under Uncertainty*. Elsevier, 1977, pp. 65–82. ISBN: 978-0-12-445850-5. DOI: 10.1016/B978-0-12-445850-5.50008-8. URL: https://linkinghub.elsevier.com/retrieve/pii/B9780124458505500088 (visited on 07/30/2019).

[69] Michael Gallagher. "Proportionality, disproportionality and electoral systems." In: *Electoral studies* 10.1 (1991), pp. 33–51.

[70] Alexander Gattig and Laurie Hendrickx. "Judgmental discounting and environmental risk perception: Dimensional similarities, domain differences, and implications for sustainability." In: *Journal of Social Issues* 63.1 (2007), pp. 21–39.

[71] Annarita Giani, Russell Bent, Mark Hinrichs, Miles McQueen, and Kameshwar Poolla. "Metrics for assessment of smart grid data integrity attacks." In: *2012 IEEE Power and Energy Society General Meeting*. IEEE. 2012, pp. 1–8.

[72] Gabriele Gianini, Marco Cremonini, Andrea Rainini, Guido Lena Cota, and Leopold Ghemmogne Fossi. "A game theoretic approach to vulnerability patching." In: *Information and Communication Technology Research (ICTRC), 2015 International Conference on*. IEEE. 2015, pp. 88–91.

[73] Austen D Givens and Nathan E Busch. "Realizing the promise of public-private partnerships in US critical infrastructure protection." In: *International Journal of Critical Infrastructure Protection* 6.1 (2013), pp. 39–50.

[74] G Gonzalez-Granadillo, Joaquín Garcia-Alfaro, Ender Alvarez, Mohammed El-Barbori, and Hervé Debar. "Selecting optimal countermeasures for attacks against critical systems using the attack volume model and the RORI index." In: *Computers & Electrical Engineering* 47 (2015), pp. 13–34.

[75] Srihari Govindan and Robert Wilson. "Refinements of Nash equilibrium." In: *Available at SSRN* (2005). DOI: https://dx.doi.org/10.2139/ssrn.772081.

[76] Vitor Graveto, Luís Rosa, Tiago Cruz, and Paulo Simões. "A stealth monitoring mechanism for cyber-physical systems." In: *International Journal of Critical Infrastructure Protection* 24 (2019), pp. 126–143.

[77] Terry Gray. *Network Security Credo*. University of Washington. 2002. URL: https://staff.washington.edu/gray/papers/credo.html (visited on 09/25/2019).

[78] Oleg Grodzevich and Oleksandr Romanko. "Normalization and other topics in multi-objective optimization." In: *The Fields-MITACS Industrial Problems Workshop*. 2006.

[79] Mordechai Guri, Gabi Kedma, Assaf Kachlon, and Yuval Elovici. "AirHopper: Bridging the air-gap between isolated networks and mobile phones using radio frequencies." In: *2014 9th International Conference on Malicious and Unwanted Software: The Americas (MALWARE)*. IEEE. 2014, pp. 58–67.

[80] Geňa Hahn and Gary MacGillivray. "A note on k-cop, l-robber games on graphs." In: *Discrete mathematics* 306.19 (2006), pp. 2492–2497.

[81] Joseph Y Halpern. "Lexicographic probability, conditional probability, and nonstandard probability." In: *Games and Economic Behavior* 68.1 (2010), pp. 155–179. DOI: 10.1016/j.geb.2009.03.013.

[82] Dong Hao, Yizhi Ren, and Kouichi Sakurai. "A game theory-based surveillance mechanism against suspicious insiders in MANETs." In: *Trusted Systems* 6802 (2010), pp. 237–252.

[83] John C. Harsanyi. "Games with randomly disturbed payoffs: A new rationale for mixed-strategy equilibrium points." In: *International Journal of Game Theory* 2.1 (1973), pp. 1–23.

[84] Hemantha S B Herath and Tejaswini C Herath. "Copula-based actuarial model for pricing cyber-insurance policies." en. In: 2.1 (2011), p. 14.

[85] Lonnie G Hibbard. *Communicating with the net generation*. Tech. rep. ARMY WAR COLL CARLISLE BARRACKS PA, 2011.

[86] Jeffrey Hightower and Gaetano Borriello. "Location systems for ubiquitous computing." In: *Computer* 34.8 (2001), pp. 57–66.

[87] David Hillson and Ruth Murray-Webster. "Managing risk attitude using emotional literacy." In: *PMI Global Congress EMEA Proceedings*. 2006.

[88] David Hillson and Ruth Murray-Webster. *Understanding and managing risk attitude*. Routledge, 2017.

[89] David Hillson and Ruth Murray-Webster. *Understanding and managing risk attitude*. Second edition. A Gower book. London: Routledge, 2017. ISBN: 978-0-566-08798-1.

[90] Erik Hollnagel and Yushi Fujita. "The Fukushima disaster–systemic failures as the lack of resilience." In: *Nuclear Engineering and Technology* 45.1 (2013), pp. 13–20.

[91] Bengt Hölmstrom. "Moral hazard and observability." In: *The Bell journal of economics* (1979), pp. 74–91.

[92] Bengt Holmstrom. "Moral hazard in teams." In: *The Bell Journal of Economics* (1982), pp. 324–340.

[93] Greg Hutchins. *ISO 31000: 2018 Enterprise Risk Management*. CERM Academy Series on Enterprise Risk Management. 2018. ISBN: 978-0-9654665-1-6.

[94] Esa Hyytiä and Jorma Virtamo. "Random waypoint model in n-dimensional space." In: *Operations Research Letters* 33.6 (2005), pp. 567–571.

[95] IEC61508. "IEC 61508:2010 functional safety of electrical/electronic/programmable electronic safety-related systems." In: *International electrotechnical commission* (2010).

[96] ISO/IEC JTC 1/SC 27. *Information technology — Security techniques — Guidelines for cybersecurity*. Standard ISO/IEC 27032:2012. International Organization for Standardization, 2012.

[97] ISO/IEC JTC 1/SC 27. *Information technology — Security techniques — Information security management systems — Overview and vocabulary*. Standard ISO/IEC 27000:2018. International Organization for Standardization, 2018.

[98] ISO/IEC JTC 1/SC 27. *Information technology — Security techniques — Information security risk management*. Standard ISO/IEC 27005. International Organization for Standardization, 2018.

[99] ISO/TC 262. *Risk management–Principles and guidelines*. Standard ISO 31000. International Organization for Standardization, 2009.

[100] ISO/TC 292 Security and resilience. *Security management systems for the supply chain — Best practices for implementing supply chain security, assessments and plans — Requirements and guidance*. Standard ISO 28001:2007. International Organization for Standardization, 2007.

[101] ISO/TMBG Technical Management Board - groups. *ISO Guide 73:2009: Risk management — Vocabulary*. Standard ISO Guide 73:2009. International Organization for Standardization, 2009.

[102] Yasunari Inamura et al. "Estimating continuous time transition matrices from discretely observed data." In: *Bank of Japan* (2006), pp. 06–07.

[103] Sushil Jajodia, Anup K Ghosh, VS Subrahmanian, Vipin Swarup, Cliff Wang, and X Sean Wang. *Moving Target Defense II: Application of Game Theory and Adversarial Modeling*. Vol. 100. Springer, 2012.

[104] Sushil Jajodia, Anup K. Ghosh, V. S. Subrahmanian, Vipin Swarup, Cliff Wang, and Xiaoyang Sean Wang, eds. *Moving Target Defense II – Application of Game Theory and Adversarial Modeling*. Vol. 100. Advances in Information Security. Springer, 2013. ISBN: 978-1-4614-5415-1. DOI: 10.1007/978-1-4614-5416-8. URL: http://dx.doi.org/10.1007/978-1-4614-5416-8.

[105] Albert Xin Jiang, Zhengyu Yin, Matthew P. Johnson, Milind Tambe, Christopher Kiekintveld, Kevin Leyton-Brown, and Tuomas Sandholm. "Towards Optimal Patrol Strategies for Fare Inspection in Transit Systems." In: *AAAI Spring Symposium: Game Theory for Security, Sustainability, and Health*. 2012.

[106] Daniel Kahneman and Amos Tversky. "Prospect Theory: An Analysis of Decision under Risk." In: *Econometrica* 47.2 (1979), pp. 263–292.

[107] Stanley Kaplan and B John Garrick. "On the quantitative definition of risk." In: *Risk analysis* 1.1 (1981), pp. 11–27.

[108] Athanasios Kehagias, Dieter Mitsche, and Paweł Prałat. "The role of visibility in pursuit/evasion games." In: *Robotics* 3.4 (2014), pp. 371–399.

[109] Jay Kesan, Ruperto Majuca, and William Yurcik. "Cyberinsurance as a market-based solution to the problem of cybersecurity: a case study." In: *Proc. WEIS*. 2005.

[110] Frank H Knight. *Risk, uncertainty and profit*. Courier Corporation, 1921.

[111] Sandra König, Stefan Rass, and Stefan Schauer. *HyRiM Deliverable 1.2: Report on Definition and Categorisation of Hybrid Risk Metrics*. 2016. URL: https://hyrim.net/wp-content/uploads/2017/12/HyRiM-D1.2-Report-on-Definition-and-Categorisation-of-Hybrid-Risk-Metrics.pdf.

[112] Yuichiro Kumamoto, Masatoshi Yamada, Michio Aoyama, Yasunori Hamajima, Hideki Kaeriyama, Hisao Nagai, Takeyasu Yamagata, Akihiko Murata, and Yukio Masumoto. "Radiocesium in North Pacific coastal and offshore areas of Japan within several months after the Fukushima accident." In: *Journal of environmental radioactivity* 198 (2019), pp. 79–88.

[113] David Lacey. *Managing the Human Factor in Information Security: How to win over staff and influence business managers*. John Wiley & Sons, 2009.

[114] Lifeng Lai, Siu-Wai Ho, and H Vincent Poor. "Privacy–Security Trade-Offs in Biometric Security Systems—Part I: Single Use Case." In: *IEEE Transactions on Information Forensics and Security* 1.6 (2011), pp. 122–139.

[115] Marc Lelarge and Jean Bolot. "A local mean field analysis of security investments in networks." In: *Proceedings of the 3rd international workshop on Economics of networked systems*. ACM. 2008, pp. 25–30.

[116] Christian Leuprecht, David B Skillicorn, and Victoria E Tait. "Beyond the Castle Model of cyber-risk and cyber-security." In: *Government Information Quarterly* 33.2 (2016), pp. 250–257.

[117] David John Leversage and Eric James Byres. "Estimating a System's Mean Time-to-Compromise." In: *IEEE Security & Privacy* 6 (Jan. 2008), pp. 52–60. ISSN: 1540-7993. DOI: 10.1109/MSP.2008.9. URL: doi.ieeecomputersociety.org/10.1109/MSP.2008.9.

[118] Ted G Lewis. *Critical infrastructure protection in homeland security: defending a networked nation*. John Wiley & Sons, 2019.

[119] Richard G Little. "Managing the Risk of Aging infrastructure." In: *IRGC, Public Sector Governance of Emerging Risks Council, Infrastructure Case* (2012).

[120] D. Lozovanu, D. Solomon, and A. Zelikovsky. "Multiobjective Games and Determining Pareto-Nash Equilibria." In: *Buletinul Academiei de Stiinte a Republicii Moldova Matematica* 3.49 (2005), pp. 115–122.

[121] ROGER B. MYERSON. *Game Theory: Analysis of Conflict*. Harvard University Press, 1991. ISBN: 9780674341166. URL: http://www.jstor.org/stable/j.ctvjsf522.

[122] Mark J Machina. ""' Expected Utility" Analysis without the Independence Axiom." In: *Econometrica: Journal of the Econometric Society* (1982), pp. 277–323.

[123] Enrique Machuca, Lawrence Mandow, and Lucie Galand. "An evaluation of best compromise search in graphs." In: *Conference of the Spanish Association for Artificial Intelligence*. Springer. 2013, pp. 1–11.

[124] Louai Maghrabi, Eckhard Pfluegel, Luluwah Al-Fagih, Roman Graf, Giuseppe Settanni, and Florian Skopik. "Improved software vulnerability patching techniques using CVSS and game theory." In: *Cyber Security And Protection Of Digital Services (Cyber Security), 2017 International Conference on*. IEEE. 2017, pp. 1–6.

[125] Mohammad Hossein Manshaei, Quanyan Zhu, Tansu Alpcan, Tamer Bacşar, and Jean-Pierre Hubaux. "Game Theory Meets Network Security and Privacy." In: *ACM Comput. Surv.* 45.3 (2013), 25:1–25:39. ISSN: 0360-0300. DOI: 10.1145/2480741.2480742. URL: http://doi.acm.org/10.1145/2480741.2480742.

[126] Robin McCusker. "E-commerce security: The birth of technology, the death of common sense?" In: *Journal of Financial Crime* 9.1 (2001), pp. 79–89.

[127] Richard D. McKelvey and Thomas R. Palfrey. "Quantal Response Equilibria for Normal Form Games." In: *Games and Economic Behavior* 10.1 (1995), pp. 6–38. DOI: 10.1006/game.1995.1023.

[128] Miles A McQueen, Wayne F Boyer, Mark A Flynn, and George A Beitel. "Quantitative cyber risk reduction estimation methodology for a small SCADA control system." In: *System Sciences, 2006. HICSS'06. Proceedings of the 39th Annual Hawaii International Conference on*. Vol. 9. IEEE. 2006, pp. 226–226.

[129] Miles A. McQueen, Wayne F. Boyer, Mark A. Flynn, and George A. Beitel. "Time-to-Compromise Model for Cyber Risk Reduction Estimation." In: *Quality of Protection*. Ed. by Dieter Gollmann, Fabio Massacci, and Artsiom Yautsiukhin. Boston, MA: Springer US, 2006, pp. 49–64. ISBN: 978-0-387-36584-8.

[130] Miles A McQueen, Trevor A McQueen, Wayne F Boyer, and May R Chaffin. "Empirical Estimates and Observations of oDay Vulnerabilities." In: *2009 42nd Hawaii International Conference on System Sciences*. Jan. 2009, pp. 1–12. DOI: 10.1109/HICSS.2009.186.

[131] Peter Mell, Tiffany Bergeron, and David Henning. *NIST Special Publication 800-40 - Creating a Patch and Vulnerability Management Program*. Nov. 2005.

[132] B. Miller. "Everything you need to know about biometric identification. Personal Identification News 1988 Biometric Industry Directory. Washington DC: Warfel & Miller." In: *Inc., Washington DC* (1988).

[133] Rich Miller. *Power Outage Affects Amazon Customers*. 2013. URL: https://www.datacenterknowledge.com/archives/2012/06/15/power-outage-affects-amazon-customers (visited on 09/25/2019).

[134] Koichi Miyasawa. *On the Convergence of the Learning Process in a 2x2 Non-zero-sum TwoPerson Game*. Tech. rep. Research Memorandum No. 33. Economic Research Program, Princeton University, 1961.

[135] Dov Monderer and Aner Sela. "A2 ×2Game without the Fictitious Play Property." In: *Games and Economic Behavior* 14.1 (1996), pp. 144–148. DOI: 10.1006/game.1996.0045.

[136] Christopher Z Mooney. *Monte carlo simulation*. 1997.

[137] Karl Mosler and Marco Scarsini. *Stochastic Orders and Applications*. Vol. 401. Berlin, Heidelberg: Springer Berlin Heidelberg, 1993. ISBN: 978-3-540-56956-5. DOI: 10.1007/978-3-642-49972-2.

[138] J Murray. "Securing critical infrastructure: Perimeter protection strategy at key national security sites including transport hubs, power facilities, prisons and correctional centres." In: *Senstar Corp., Ottawa, ON, Canada, White Paper* (2012).

[139] NVD. *National Vulnerability Database U.S (NVD)*. 2018. URL: https://nvd.nist.gov/ (visited on 09/19/2018).

[140] Rosemarie Nagel. "Unraveling in Guessing Games: An Experimental Study." In: *American Economic Review* 85.5 (1995), pp. 1313–1326. URL: https://EconPapers.repec.org/RePEc:aea:aecrev:v:85:y:1995:i:5:p:1313-26.

[141] John Nash. "Non-cooperative games." In: *Annals of mathematics* (1951), pp. 286–295.

[142] John von Neumann. "Zur theorie der gesellschaftsspiele." In: *Mathematische annalen* 100.1 (1928), pp. 295–320.

[143] Clive Norris, Mike McCahill, and David Wood. "The growth of CCTV: a global perspective on the international diffusion of video surveillance in publicly accessible space." In: *Surveillance & Society* 2.2/3 (2004).

[144] J. R. Norris. "Continuous-time Markov chains II." In: *Markov Chains* (1997), pp. 108–127. DOI: 10.1017/CBO9780511810633.005.

[145] Serguei Y Novak. *Extreme value methods with applications to finance*. Chapman & Hall/CRC Press, 2011. ISBN: 9781439835746.

[146] William Nzoukou, Lingyu Wang, Sushil Jajodia, and Anoop Singhal. "A Unified Framework for Measuring a Network's Mean Time-to-Compromise." In: *2013 IEEE 32nd International Symposium on Reliable Distributed Systems* (2013), pp. 215–224.

[147] Antti Oulasvirta, Aurora Pihlajamaa, Jukka Perkiö, Debarshi Ray, Taneli Vähäkangas, Tero Hasu, Niklas Vainio, and Petri Myllymäki. "Long-term effects of ubiquitous surveillance in the home." In: *Proceedings of the 2012 ACM Conference on Ubiquitous Computing*. ACM, 2012, pp. 41–50.

[148] Ranjan Pal and Pan Hui. "CyberInsurance for cybersecurity a topological take on modulating insurance premiums." In: *ACM SIGMETRICS Performance Evaluation Review* 40.3 (2012), pp. 86–88.

[149] Ravindra Pal, Leana Golubchik, Konstantinos Psounis, and Pan Hui. "Will cyber-insurance improve network security? A market analysis." In: *INFOCOM, 2014 Proceedings IEEE*. IEEE. 2014, pp. 235–243.

[150] Emmanouil Panaousis, Andrew Fielder, Pasquale Malacaria, Chris Hankin, and Fabrizio Smeraldi. "Cybersecurity games and investments: A decision support approach." In: *International Conference on Decision and Game Theory for Security*. Springer. 2014, pp. 266–286.

[151] Heming Pang, Linying Jiang, Liu Yang, and Kun Yue. "Research of android smart phone surveillance system." In: *Computer Design and Applications (ICCDA), 2010 International Conference on*. Vol. 2. IEEE, 2010, pp. V2–373.

[152] Torrence Douglas Parsons. "Pursuit-evasion in a graph." In: *Theory and Applications of Graphs*. Ed. by Yousef Alavi and Don R. Lick. Berlin, Heidelberg: Springer Berlin Heidelberg, 1976, pp. 426–441. ISBN: 978-3-540-35912-8.

[153] Sarah Perez. "Technology Populism: Risks & Rewards." In: *Read Write Web* 5.6 (2008), p. 2010. URL: http://www.readwriteweb.com/archives/technology_populism_risks_rewards.php (visited on 09/25/2019).

[154] Physorg. *IAEA warned Japan over nuclear quake risk: WikiLeaks*. 2011. URL: https://phys.org/news/2011-03-iaea-japan-nuclear-quake-wikileaks.html (visited on 12/04/2019).

[155] Stjepan Picek, Erik Hemberg, and Una-May O'Reilly. "If You Can't Measure It, You Can't Improve It: Moving Target Defense Metrics." In: *Proceedings of the 2017 Workshop on Moving Target Defense*. MTD '17. New York, NY, USA: ACM, 2017, pp. 115–118. ISBN: 978-1-4503-5176-8. DOI: 10.1145/3140549.3140558. URL: http://doi.acm.org/10.1145/3140549.3140558.

[156] John W. Pratt. "Risk Aversion in the Small and in the Large." In: *Econometrica* 32.1/2 (Jan. 1964), p. 122. ISSN: 00129682. DOI: 10.2307/1913738. URL: https://www.jstor.org/stable/1913738?origin=crossref (visited on 07/30/2019).

[157] Drazen Prelec et al. "The probability weighting function." In: *ECONOMETRICA-EVANSTON ILL-* 66 (1998), pp. 497–528.

[158] John Quiggin. *Generalized expected utility theory: The rank-dependent model*. Springer Science & Business Media, 2012.

[159] R Core Team. *R: A Language and Environment for Statistical Computing*. ISBN 3-900051-07-0. Vienna, Austria, 2018. URL: http://www.R-project.org.

[160] RAPID7. *The Rapid7 Vulnerability and Exploit Database*. 2018. URL: https://www.rapid7.com/db (visited on 02/03/2019).

[161] Stefan Rass. "On Game-Theoretic Network Security Provisioning." In: *Springer Journal of Network and Systems Management* 21.1 (2013), pp. 47–64. DOI: 10.1007/s10922-012-9229-1. URL: http://www.springerlink.com/openurl.asp?genre=article&id=doi:10.1007/s10922-012-9229-1.

[162] Stefan Rass. "On game-theoretic network security provisioning." In: *Journal of network and systems management* 21.1 (2013), pp. 47–64.

[163] Stefan Rass. "On Game-Theoretic Risk Management (Part One)-Towards a Theory of Games with Payoffs that are Probability-Distributions." In: *arXiv preprint arXiv:1506.07368* (2015).

[164] Stefan Rass. "On Game-Theoretic Risk Management (Part Two)-Algorithms to Compute Nash-Equilibria in Games with Distributions as Payoffs." In: *arXiv preprint arXiv:1511.08591* (2015).

[165] Stefan Rass. "Security Strategies and Multi-Criteria Decision Making." In: *Game Theory for Security and Risk Management*. Springer, 2018, pp. 47–74.

[166] Stefan Rass, Sandra Koenig, and Stefan Schauer. "Decisions with Uncertain Consequences—A Total Ordering on Loss-Distributions." In: *PloS one* 11.12 (2016), e0168583.

[167] Stefan Rass, Sandra König, and Stefan Schauer. "Uncertainty in Games: Using Probability-Distributions as Payoffs." In: *Proceedings of The International Conference on Decision and Game Theory for Security*. Springer, 2015, pp. 346–357.

[168] Stefan Rass, Sandra Konig, and Stefan Schauer. "Decisions with Uncertain Consequences-A Total Ordering on Loss-Distributions." In: *PLoS ONE* 11.12 (2016), e0168583. DOI: 10.1371/journal.pone.0168583.

[169] Stefan Rass, Sandra Konig, and Stefan Schauer. "Defending Against Advanced Persistent Threats Using Game-Theory." In: *PLoS ONE* 12.1 (2017), e0168675. DOI: 10.1371/journal.pone.0168675.

[170] Stefan Rass, Sandra König, and Stefan Schauer. "On the Cost of Game Playing: How to Control the Expenses in Mixed Strategies." In: *Decision and Game Theory for Security*. [S.l.]: Springer, 2017, pp. 494–505. ISBN: 978-3319687100.

[171] Stefan Rass and Benjamin Rainer. "Numerical computation of multi-goal security strategies." In: *International Conference on Decision and Game Theory for Security*. Springer. 2014, pp. 118–133.

[172] Stefan Rass and Stefan Schauer, eds. *Game Theory for Security and Risk Management: From Theory to Practice*. Static & Dynamic Game Theory: Foundations & Applications. Springer Birkhäuser, 2018. ISBN: 978-3-319-75267-9.

[173] Stefan Rass and Quanyan Zhu. "GADAPT: A Sequential Game-Theoretic Framework for Designing Defense-in-Depth Strategies Against Advanced Persistent Threats." In: *Decision and Game Theory for Security*. Ed. by Quanyan Zhu, Tansu Alpcan, Emmanouil Panaousis, Milind Tambe, and William Casey. Vol. 9996. Lecture Notes in Computer Science. Cham: Springer International Publishing, 2016, pp. 314–326. ISBN: 978-3-319-47412-0. DOI: 10.1007/978-3-319-47413-7-18.

[174] Stefan Rass, Jasmin Wachter, Stefan Schauer, and Sandra König. "Subjektive Risikobewertung-Über Datenerhebung und Opinion Pooling." In: *DACH Security* (2017), pp. 225–237.

[175] Stefan Rass, Sandra König, Stefan Schauer, Vincent Bürgin, Jeremias Epperlein, and Fabian Wirth. "On Game Theory Using Stochastic Tail Orders." In: *arXiv preprint arXiv:2108.00680* (2021).

[176] Lauren Regan. "Electronic communications surveillance." In: *Monthly Review* 66.3 (2014), pp. 32–42.

[177] Alfréd Rényi et al. "On measures of entropy and information." In: *Proceedings of the Fourth Berkeley Symposium on Mathematical Statistics and Probability, Volume 1: Contributions to the Theory of Statistics*. The Regents of the University of California. 1961.

[178] Marc Oliver Rieger and Mei Wang. "Cumulative prospect theory and the St. Petersburg paradox." In: *Economic Theory* 28.3 (2006), pp. 665–679.

[179] Risk. *Merriam-Webster.com*. 2019. URL: https://www.merriam-webster.com/dictionary/risk (visited on 09/25/2019).

[180] Risk. *BusinessDictionary.com*. WebFinance, Inc., 2020. URL: http://www.businessdictionary.com/definition/risk.html (visited on 03/04/2020).

[181] Risk. *The Cambridge Advanced Learner's Dictionary and Thesaurus*. Cambridge University Press, 2020. URL: https://dictionary.cambridge.org/dictionary/english/risk (visited on 03/04/2020).

[182] Christian P. Robert. *The Bayesian choice*. New York: Springer, 2001.

[183] Bernd Rohrmann. "Risk perception, risk attitude, risk communication, risk management: A conceptual appraisal." In: *15th Internaional Emergency Management Society (TIEMS) Annual Conference*. 2008.

[184] Ronald Ross. *Security and privacy controls for federal information systems and organizations*. Special Publication (NIST SP) 800-53 Rev. 4. National Institute Of Standards and Technology(NIST), 2007. DOI: 10.6028/NIST.SP.800-53r4.

[185] M Joan Saary. "Radar plots: a useful way for presenting multivariate health care data." In: *Journal of clinical epidemiology* 61.4 (2008), pp. 311–317.

[186] Bruce Schneier. "The psychology of security." In: *International Conference on Cryptology in Africa*. Springer. 2008, pp. 50–79.

[187] Aner Sela. "Fictitious play in one-against-all multi-player games." In: *Economic Theory* 14.3 (1999), pp. 635–651.

[188] R. Selten. "Reexamination of the Perfectness Concept for Equilibrium Points in Extensive Games." In: *Models of Strategic Rationality*. Ed. by Reinhard Selten. Vol. 2. Theory and Decision Library C, Game Theory, Mathematical Programming and Operations Research. Dordrecht: Springer, 1988, pp. 1–31. ISBN: 978-90-481-8446-0. DOI: 10.1007/978-94-015-7774-8-1.

[189] Andrei Serjantov and George Danezis. "Towards an information theoretic metric for anonymity." In: *International Workshop on Privacy Enhancing Technologies*. Springer. 2002, pp. 41–53.

[190] Roberto Setola, Eric Luiijf, and Marianthi Theocharidou. "Critical Infrastructures, Protection and Resilience." In: *Managing the Complexity of Critical Infrastructures: A Modelling and Simulation Approach*. Ed. by Roberto Setola, Vittorio Rosato, Elias Kyriakides, and Erich Rome. Cham: Springer International Publishing, 2016, pp. 1–18. ISBN: 978-3-319-51043-9. DOI: 10.1007/978-3-319-51043-9_1. URL: https://doi.org/10.1007/978-3-319-51043-9_1.

[191] Moshe Shaked and J. George Shanthikumar. *Stochastic Orders*. Springer, 2006.

[192] Alireza Shameli-Sendi, Rouzbeh Aghababaei-Barzegar, and Mohamed Cheriet. "Taxonomy of information security risk assessment (ISRA)." In: *Computers & security* 57 (2016), pp. 14–30.

[193] Claude Elwood Shannon. "A mathematical theory of communication." In: *ACM SIGMOBILE Mobile Computing and Communications Review* 5.1 (2001), pp. 3–55.

[194] Devendra Shelar and Saurabh Amin. "Security assessment of electricity distribution networks under DER node compromises." In: *IEEE Transactions on Control of Network Systems* 4.1 (2016), pp. 23–36.

[195] Heidi Shey, K Mak, S Balaouras, and B Luu. "Understand the state of data security and privacy: 2013 to 2014." In: *Forrester Research Inc* 1 (2013).

[196] Jin Shi, Yin Lu, and Li Xie. "Game Theory Based Optimization of Security Configuration." In: *International Conference on Computational Intelligence and Security* (2007), pp. 799–803. DOI: 10.1109/CIS.2007.25.

[197] Vivek K. Singh and Mohan S. Kankanhalli. "Adversary aware surveillance systems." In: *IEEE Transactions on Information Forensics and Security* 4.3 (2009), pp. 552–563.

[198] Paul Slovic and Elke U Weber. "Perception of risk posed by extreme events." In: *Regulation of Toxic Substances and Hazardous Waste (2nd edition)(Applegate, Gabba, Laitos, and Sachs, Editors), Foundation Press, Forthcoming* (2002).

[199] Geoffrey Smith. "On the foundations of quantitative information flow." In: *International Conference on Foundations of Software Science and Computational Structures*. Springer. 2009, pp. 288–302.

[200] Rebecca Smith. "Assault on California power station raises alarm on potential for terrorism." In: *Wall Street Journal* 5 (2014).

[201] Jin Ho Song and Tae Woon Kim. "Severe accident issues raised by the Fukushima accident and improvements suggested." In: *Nuclear Engineering and Technology* 46.2 (2014), pp. 207–216.

[202] Murugiah Souppaya and Karen Scarfone. *NIST Special Publication 800-40 - Guide to Enterprise Patch Management Technologies*. July 2013. DOI: http://dx.doi.org/10.6028/NIST.SP.800-40r3.

[203] Michael Stamatelatos. "Probabilistic risk assessment: what is it and why is it worth performing it." In: *NASA Office of Safety and Mission Assurance* 4.05 (2000), p. 00.

[204] National Institute of Standards and Technology (NIST). *Guidelines for smart grid cybersecurity*. Tech. rep. 2014.

[205] Chris Starmer. "Developments in Non-Expected Utility Theory: The Hunt for a Descriptive Theory of Choice under Risk." In: *Journal of Economic Literature* 38.2 (2000), pp. 332–382. ISSN: 00220515. URL: http://www.jstor.org/stable/2565292.

[206] Krzysztof Szajowski. "Multi-variate Quickest Detection of Significant Change Process." In: *GameSec*. Springer, 2011, pp. 56–66.

[207] Milind Tambe. *Security and game theory: Algorithms, deployed systems, lessons learned*. 1. publ. Cambridge u.a.: Cambridge Univ. Press, 2012. ISBN: 978-1-107-09642-4.

[208] The White House. *PPD-21 – Critical infrastructure security and resilience*. Presidential Policy Directive. White House Washington, DC, 2013. URL: https://obamawhitehouse.archives.gov/the-press-office/2013/02/12/presidential-policy-directive-critical-infrastructure-security-and-resil.

[209] Bob Toxen. "The NSA and Snowden: securing the all-seeing eye." In: *Commun. ACM* 57.5 (2014), pp. 44–51.

[210] Amos Tversky and Daniel Kahneman. "Advances in prospect theory: Cumulative representation of uncertainty." In: *Journal of Risk and uncertainty* 5.4 (1992), pp. 297–323.

[211] UP KRITIS. *Public-Private Partnership for Critical Infrastructure Protection*. Implementation Plan. Federal Office for Information Security (BSI), Germany, 2014. URL: https://www.kritis.bund.de/SharedDocs/Downloads/Kritis/EN/UP_KRITIS.pdf.

[212] U.S. Congress. *Uniting and Strengthening America by Providing Appropriate Tools Required to Intercept and Obstruct Terrorism (Act of 2001)*. USA Patriot Act. National Legislative Bodies – National Authorities, USA, 2001. URL: https://www.congress.gov/bill/107th-congress/house-bill/3162.

[213] André Van Cleeff and Roel J Wieringa. "Rethinking de-perimeterisation: Problem analysis and solutions." In: *Proceedings of the IADIS International Conference Information Systems 2009, 25-27 Feb 2009, Barcelona*. 2009, pp. 105–112.

[214] Pradeep Varakantham, Hoong Chuin Lau, and Zhi Yuan. "Scalable Randomized Patrolling for Securing Rapid Transit Networks." In: *IAAI*. 2013.

[215] Andras Varga. *OMNeT++ 5.1: Discrete Event Simulator*. Apr. 2017. URL: https://omnetpp.org/ (visited on 12/10/2017).

[216] Lev Virine and Michael Trümper. *ProjectThink: Why good managers make poor project choices*. 2013. ISBN: 978-1-4724-0403-9.

[217] John Von Neumann and Oskar Morgenstern. *Theory of games and economic behavior (commemorative edition)*. 2007.

[218] M Voorneveld. "Pareto-optimal security strategies as minimax strategies of a standard matrix game." In: *Journal of Optimization Theory and Applications* 102.1 (1999), pp. 203–210.

[219] Jasmin Wachter, Stefan Rass, Sandra König, and Stefan Schauer. "Disappointment-Aversion in Security Games." In: *Decision and Game Theory for Security*. Ed. by Linda Bushnell, Radha Poovendran, and Tamer Başar. Cham: Springer International Publishing, 2018, pp. 314–325. ISBN: 978-3-030-01554-1.

[220] Isabel Wagner and David Eckhoff. "Technical privacy metrics: a systematic survey." In: *ACM Computing Surveys (CSUR)* 51.3 (2018), pp. 1–38.

[221] Abraham Wald. "Statistical decision functions which minimize the maximum risk." In: *Annals of Mathematics* (1945), pp. 265–280.

[222] James L. Wayman. "National Biometric Test Center: Collected Works 1997-2000." In: *Biometric Consortium of the US Government interest group on biometric authentication) San Jose State University, CA* (2000).

[223] James Wayman, Anil Jain, Davide Maltoni, and Dario Maio. "An introduction to biometric authentication systems." In: *Biometric Systems* (2005), pp. 1–20.

[224] Elke U Weber, Ann-Renee Blais, and Nancy E Betz. "A domain-specific risk-attitude scale: Measuring risk perceptions and risk behaviors." In: *Journal of behavioral decision making* 15.4 (2002), pp. 263–290.

[225] Elke U Weber and Christopher K Hsee. "Models and mosaics: Investigating cross-cultural differences in risk perception and risk preference." In: *Psychonomic Bulletin & Review* 6.4 (1999), pp. 611–617.

[226] Elke U Weber and Christopher Hsee. "Cross-cultural differences in risk perception, but cross-cultural similarities in attitudes towards perceived risk." In: *Management science* 44.9 (1998), pp. 1205–1217.

[227] Elke U Weber and Richard A Milliman. "Perceived risk attitudes: Relating risk perception to risky choice." In: *Management science* 43.2 (1997), pp. 123–144.

[228] Elke Weber. "Decision and choice: Risk, empirical studies." In: *International Encyclopedia of the Social and Behavioral Sciences*. Oxford, UK: Elsevier, 2001, pp. 13347–13351.

[229] Gregory Wheeler. "Bounded Rationality." In: *The Stanford Encyclopedia of Philosophy*. Ed. by Edward N. Zalta. Fall 2019. Metaphysics Research Lab, Stanford University, 2019. URL: https://plato.stanford.edu/archives/fall2019/entries/bounded-rationality/.

[230] Alyson G. Wilson, Gregory D. Wilson, and David H. Olwell. "Statistical methods in counterterrorism." In: *Springer Science+ Business Media* 250 (2006), p. 281.

[231] Michael Wooldridge. "Does game theory work?" In: *IEEE Intelligent Systems* 27.6 (2012), pp. 76–80.

[232] Maochao Xu and Lei Hu. "Cybersecurity Insurance: Modeling and Pricing." en. In: (2017), p. 38. URL: https://www.soa.org/globalassets/assets/Files/Research/Projects/cybersecurity-insurance-report.pdf.

[233] Xiong Zhang, Alex Tsang, Wei T Yue, and Michael Chau. "The classification of hackers by knowledge exchange behaviors." In: *Information Systems Frontiers* 17.6 (2015), pp. 1239–1251.

[234] Yichi Zhang, Lingfeng Wang, Yingmeng Xiang, and Chee-Wooi Ten. "Power System Reliability Evaluation With SCADA Cybersecurity Considerations." In: *IEEE Transactions on Smart Grid* 6.4 (July 2015), pp. 1707–1721. ISSN: 1949-3053. DOI: 10.1109/TSG.2015.2396994.

[235] Quanyan Zhu and Tamer Başar. "Game-Theoretic Approach to Feedback-Driven Multi-stage Moving Target Defense." In: *4th International Conference on Decision and Game Theory for Security - Volume 8252*. GameSec 2013. New York, NY, USA: Springer-Verlag New York, Inc, 2013, pp. 246–263. ISBN: 978-3-319-02785-2. DOI: 10.1007//978-3-319-02786-9-15.

[236] Quanyan Zhu and Stefan Rass. "Game Theory Meets Network Security: A Tutorial." In: *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security*. CCS '18. Toronto, Canada: ACM, 2018, pp. 2163–2165. ISBN: 978-1-4503-5693-0. DOI: 10.1145/3243734.3264421. URL: http://doi.acm.org/10.1145/3243734.3264421.